USER GUIDE

STAR-CCM+® Version 9.06



© 2014 CD-adapco™

CD-adapco™ and any and all CD-adapco brand, product, service and feature names, logos and slogans are registered trademarks or trademarks of CD-adapco in the United States or other countries. All other brand, product, service and feature names or trademarks are the property of their respective owners.

User Guide Contents

Navigating Help 1

New Features for v9.06 2

New Tutorials for v9.06 3

Using STAR-CCM+ 4

What Is STAR-CCM+? 6

Using Online Help 11

Browsing the Online Help 12

Accessing Context Sensitive Help 14

Using the Navigation Frame 16

Using the Toolbar Frame 21

Changing the Default Browser 22

Understanding Typographic Conventions 24

What Are the Important Concepts? 25

What is the Client-Server Architecture? 26

What Is the Client Workspace? 27

What Are Simulations? 28

What Are 3D-CAD Models? 30

What Are Parts? 31

What Are Continua? 33

What Are Models? 37

What Are Regions? 40

What Are Boundaries? 44

What Are Interfaces? 50

What Are Solvers? 52

What Are Reports, Monitors and Plots? 54
What Are Scenes? 55

What Is The General Workflow? 56
Preparing the Geometry 58
Defining the Simulation Topology 59
Generating the Mesh 61
Defining the Physics 67
Preparing for Analysis 71
Running the Simulation 73
Analyzing the Results 74

Working with Mixed or Double Precision 77
Advantages and Limitations of a Double Precision Simulation 79

Launching STAR-CCM+ 80
Launching a Client on Windows 81
Launching a Client on Linux 84

Working with Simulations 86
Creating a New Simulation Interactively 87
Loading an Existing Simulation Interactively 90
Understanding Interactive Startup Options 94
Creating and Opening Simulations from the Command Line in Serial Mode 106
Starting Remote Servers from the Command Line 108
Restrictions on Names of Files and Paths in STAR-CCM+ 110
Using the Simulation Pop-Up Menu 113
Saving Simulations to File 114
Closing the Simulation 124
Closing the Client Workspace 125
Viewing and Editing the Log File 127
Using the Summary Report 130

Using the STAR-CCM+ Workspace 136
Using the Workspace 137
The Object Selection Dialog 184
Using Filters 204
Using the Title Bar 211
Using the Menus 214
Using the Toolbars 228

Using the Explorer Pane....242
Working with the Object Tree....246
Using the Table Editor....275
Using the Properties Window....284
Using the Graphics Window....295
Using the Output Window....368
Using Keyboard Shortcuts....384

Working With Clients and Servers....393
What Is Client-Server?....394
What Is a Client?....396
Running the Client in Batch....398
What Is a Server?....401
Disconnecting from the Server....405
Connecting to the Server....407
Connecting a Client from the Command Line....419
Configuring the Server to Shutdown When Idle....420
Closing the Server Automatically Upon Solver Completion....422
Working with Remote Servers....423
Starting a Server without a Client....426
Displaying Diagnostic Messages from the Server and Client....427
Setting Network Binding Order for Failed Visualization Connections....429
The Server Processes Manager....430

Working With Parallel Servers....431
What Is a Parallel Server?....432
Selecting the MPI Implementation....434
Using a Local Multi-Core Workstation....446
Using a Remote Multi-Core Workstation....449
Using a Cluster of Machines....452
Starting a Proxy Server on a Host....463
Troubleshooting....465
Balancing....467
Getting Parallel Help through the Command Line....468
Viewing the Partitioning....469
Using Batch Management Systems....470
Cluster Node Communication....478
Expert MPI Options....485

Running Distributed Parallel Simulations on Multiple Windows Machines....486

Command-Line Reference....488

Working with .ini Files....498

Setting Conditions and Values....500

What Are Conditions?....502
What Are Profiles?....503
What Are Values?....512
Setting Conditions....513
Setting Profile Methods....514
Setting Values Using the Composite Vector Profile Method....515
Setting Values Using a Constant....518
Setting Values Using a Composite Array....520
Setting Values Using a Constant Array....521
Setting Values Using a Field Function....522
Setting Values Using an Expression....524
Setting Values Using Radial Equilibrium....526
Setting Values Using a Normal Distribution....527
Setting Values Using a Normal Distribution Array....529
Setting Values Using a Table....531
Setting Values Using a Feature Curve....556
Setting Values Using a Uniform Distribution....558
Setting Values Using a Uniform Distribution Array....559
Setting Values Using User Code....560
Setting Values Using a Composite Symmetric Tensor....561
Setting Values Using a Composite Tensor....563
Setting Values Using a Principal Tensor....565
Setting Values Using an Axisymmetric Tensor....567
Setting Values Using an Isotropic Tensor....568
Setting Values Using a Polynomial....569
Copying Properties....577
Defining Dimensionality....578
Setting Multiple Corresponding Conditions Simultaneously....579
Setting Multiple Corresponding Values Simultaneously....581
Setting Reference Values....583

STAR-CCM+ Licensing....587

STAR-CCM+ Product Configurations....588

Power-on-Demand Licensing 598
STAR-CCM+ Add-on Modules 612
Power Tokens 617

Getting Support 621
Gathering Information on the Issue 624
Preparing a Case for Support 632

Suggesting Improvements of STAR-CCM+ 634

Troubleshooting Graphics 636
Resolving Hardware Issues 637
Working Without Graphics Hardware (Linux only) 638
A Scene Will Not Open 640
Improving the Rendering of Concave Faces 642
Z-Fighting Appears 644

Credits 646
IBM Platform MPI 647
lp_solve 648
sundials 649
SuperLU 651
LAPACK 652
MUMPS 653
fastprox 654
JFreeChart 655

Modeling Geometry 656

Using 3D-CAD 658
Introducing 3D-CAD Modeling 660
Sketches 697
Constraining a Sketch 759
Dimensioning Sketches 777
Creating Bodies From Sketches 797
Working with Bodies 839
Working with Faces and Edges 891
Working with Patterns 919
Working with Fillets and Chamfers 933
Working with Design Parameters 942

Working with Volume Extraction....952
Removing Unwanted Features....965
Updating the 3D-CAD Model....969
Importing Geometry into 3D-CAD....972
Exporting Geometry from 3D-CAD....983
Working with 3D-CAD Models....987
The Turbo Slicing Tool....995

Importing Geometry....1008
What Is a Surface Mesh?....1010
What Is Neutral Format Data?....1011
What Is Native CAD Data?....1012
What Is XML Data?....1013
Importing Surface Data into STAR-CCM+....1014
Importing Surfaces Using an XML File....1030
Importing Multiple Surface Data Files....1034
Reviewing the Surface Data Import Report....1036
Using Surface Data Import Representations....1037
Checking the Surface Validity....1039

Understanding Geometry Parts....1041
What Is a Geometry Part?....1043
Creating and Importing Parts....1060
Organizing Assemblies Using Composite Parts....1083
Using Part Contact Data to Create Interfaces....1090
Updating Part Interfaces....1117

Managing Geometry Parts....1120
What Are Tags?....1121
Creating User Tags....1122
Applying a Tag to a Part....1127
Removing a Tag from a Part....1139
Deleting a Tag....1143
What Is Metadata?....1144

Manipulating Parts....1149
Retessellating a Part....1150
Transforming a Part....1152
Duplicating a Part....1162
Combining Parts....1164
Performing Boolean Operations....1167

Splitting Non Contiguous Parts....1178
Splitting a Part by Surface Topology....1180
Replacing a Part....1184
Manipulating Part Surfaces....1186
Manipulating Part Curves....1209
Repairing Part Surfaces and Curves....1216
Exporting Parts....1219
Getting the Face Count of a Part....1225
Editing Imported Parts in 3D CAD....1227
Part Operations Reference Page....1232

Preparing Surfaces....1239
Launching The Surface Preparation Tool....1243
Diagnostics....1254
Repair Surface Mode....1286
Repair Features Mode....1394
Offset Faces Edges Mode....1405
Edit Connected Surfaces Mode....1420
Merge And Imprint Mode....1430
Intersect Boolean Faces Mode....1467
Entity Selection....1480
Display Controls....1527
Surface, Contact, and Curve Edit Controls....1543
Set Controls....1547
Undo And Redo....1555
Query Diagnostic Responses From Java Macros....1556
Keyboard Shortcuts....1571
Pop Up Menu Of The Surface Preparation Graphics Window....1575

Displaying the Geometry....1576

Defining the Simulation Topology....1578

Assigning Geometry Parts to Regions....1579
Assigning Parts to Regions....1581
Assigning Part Surfaces to Boundaries....1594
Assigning Part Curves to Feature Curves....1601

Transforming Regions....1605
Scaling Regions....1608

Translating Regions 1610
Rotating Regions 1613
Reflect Regions 1616

Region Manipulation 1619
Region Combination 1621
Region Splitting 1627
Region Creation 1637
Regions Manager Properties 1640
Regions Manager Pop Up Menu 1641

Manipulating Boundaries 1642
Combining Boundaries 1643
Fusing Boundaries 1646
Merging Boundaries 1651
Splitting Boundaries 1654
Creating New Boundaries 1682
Projecting Boundaries 1684
Selecting Boundaries 1690
Copying Boundary Properties 1691
Properties of the Boundaries Manager 1692
Pop-Up Menu of the Boundaries Manager 1693

Performing Surface Boolean Operations 1694
What are Surface Boolean Operations? 1696
Applying a Surface Boolean 1699

Working with Interfaces 1709
Working with Interfaces between Regions 1713
Working with Interfaces between Boundaries 1714
Boundary Interface Types Reference 1753

Meshing 1792

Meshing Overview 1794
Comparison of Meshing Strategies 1795
Mesh Construction 1798
Parallel Meshing 1802
Concurrent Meshing 1806
Base Size 1809

Parts-Based Meshing **1811**
Preparing Geometry for Meshing 1812
Surface Wrapping the Geometry 1825
Generating Surface and Volume Mesh 1846
Mesh Refinement in a Specified Volume 1854
Creating a Parts-Based Mesh with Baffles 1860
Mesh Operations 1864
Mesh Descriptions 1916
Mesh Contacts 1921
Custom Mesh Controls 1934
Parts Meshes Continua 1961

Region-Based Meshing **1963**
Creating a Region-Based Mesh 1964
Feature Curves 1990
Using Feature Curves to Repair a Surface 2005
Volumetric Controls 2015
Mesh Continua 2054

Surface Meshers **2061**
Surface Meshers Overview 2062
Surface Wrapper 2063
Surface Remesher 2090
Automatic Surface Repair 2104
Surface Mesher Controls and Values 2112

Volume Meshers **2126**
Volume Meshers Overview 2128
Polyhedral Mesher 2130
Tetrahedral Mesher 2140
Trimmed Mesher 2142
Prism Layer Mesher 2166
Advancing Layer Mesher 2203
Extruder Mesher 2224
Thin Mesher 2236
Generalized Cylinder Mesher 2254
Volume Mesher Controls and Values 2273

Mesh Refinement **2286**
Wake Refinement 2287
Field Function Mesh Refinement 2300

Directed Meshing 2313
Preparing CAD Parts for Directed Meshing 2316
Launching a Directed Mesh Operation 2318
Choosing Source and Target Surfaces 2321
Meshing a Stack of Parts 2326
Meshing a Single Part in a Closed Loop with a Baffle 2332
What Is Patch Meshing? 2335
Creating a Source Mesh Using Patches 2338
Creating a Source Mesh Using an Existing Mesh 2355
Deleting a Source Mesh 2365
Creating a Volume Distribution 2366
Deleting a Volume Distribution 2370
Generating the Directed Volume Mesh 2371

Checking the Volume Mesh 2373
What Is a Valid Mesh? 2375
What Is Mesh Quality? 2376
Running Mesh Diagnostics 2378
Checking The Mesh Validity 2385
Using The Connected Cells Field Function 2389
Checking The Skewness Angle 2390
Checking The Face Validity, Cell Quality, And Volume Change Cell Metrics 2393
Identifying And Visualizing Bad Cells 2396
Removing Invalid Cells 2413
Volume Meshing Field Functions Reference 2419

Working With Representations 2426
What Are Representations? 2427
Displaying Representations 2430
Using Proxy Representations 2434
Replacing a Surface Mesh 2437
Preparing a Surface Mesh 2440
Exporting a Surface Mesh 2441
Extracting a Surface Mesh 2446
Using Cell Sets 2449

Mesh Visualization 2491
Displaying the Mesh 2492
Changing the Surface Opacity 2495

Displaying the Volume Mesh on a Plane Section....2497
Comparing a Volume Mesh with the Geometry.... 2501

Two Dimensional and Axisymmetric Meshes.... 2504
Creating a Two-Dimensional Mesh.... 2505
Extracting The Two Dimensional Mesh....2511
Examples of Two-Dimensional Mesh Extraction.... 2516

Scaling the Entire Mesh.... 2519

Working with Baffles.... 2522
Generating Baffles with Zero Thickness....2523
Generating Baffles with Finite Thickness....2533

Multi-Region Meshing....2534
Multi-Region Surface Requirements.... 2535
Starting from a Single Region.... 2538
Starting From Multiple Regions.... 2543
CAD Data Export Guidelines for Assemblies....2545

Modeling Physics 2549

Setting Up Physics 2551

Using Physics Continua . . . 2552
Selecting Physics Models . . . 2557
Setting Types . . . 2571
Link Redirection . . . 2582

Using Overset Meshes 2662

Understanding the Overset Workflow . . . 2666
Preparing Overset Regions . . . 2666
Coupling Regions . . . 2668
Dynamic Overset Behavior . . . 2676
Overset Topology . . . 2681
Setting up Overset Mesh with Fluid Film . . . 2685
Generating Suitable Meshes . . . 2690
Visualizing and Analyzing Solution . . . 2692
Overset Mesh Best-Practice Guidelines . . . 2695
Overset Mesh Methodology . . . 2702
Overset Mesh Field Function Reference . . . 2705
Troubleshooting Overset Mesh . . . 2706

Modeling Space, Time, and Motion 2711

Modeling Space . . . 2712
Modeling Time . . . 2725
Modeling Motion . . . 2732

Modeling Materials 2802

What Is a Material Model? . . . 2802
Selecting a Material Model . . . 2804
Modeling a Single-Component Material . . . 2806
Modeling a Multi-Component Mixture . . . 2810
Modeling a Multiphase Mixture . . . 2814
Managing Single-Component Materials . . . 2816
Managing Multi-Component Materials . . . 2819
Managing Multi-Part Solids . . . 2829
Modeling Linear Isotropic Elastic . . . 2839
Setting Material Properties Methods . . . 2841
Using Material Databases . . . 2943

Modeling Flow and Energy 2971

General Navier-Stokes . . . 2973

Specific Flow Phenomena . . . 3113
Solving the Flow Adjoint . . . 3237
Boundary Conditions Formulation . . . 3270
Flow Reference . . . 3286
Values Reference . . . 3308
Fluid Flow and Energy Field Functions Reference . . . 3318
Fluid Flow Formulation . . . 3332
Fluid Flow Bibliography . . . 3398

Modeling Heat Transfer 3401

Heat Conduction . . . 3403
Convective Heat Transfer . . . 3404
Thermal Radiation . . . 3406
Verification and Validation of Results . . . 3408
Designing a Heat Transfer Optimization Study . . . 3408
Convective Heat Transfer Coefficients . . . 3411
Setting an Environment Boundary Condition for Heat Transfer . . . 3419
Modeling Fully Developed Energy . . . 3429
Modeling Boiling . . . 3437
Modeling Moisture Boiling in Solids . . . 3441
Modeling Thin Film . . . 3446
Modeling Heat Exchangers . . . 3463
Modeling Energy in Solids . . . 3502
Energy Reference . . . 3518
Heat Transfer Field Functions Reference . . . 3536
Energy Formulation . . . 3544

Modeling Species 3563

What Is the Coupled Species Model? . . . 3563
What Is the Segregated Species Model? . . . 3565
Working with Species . . . 3566
Segregated Species/Coupled Species Model Properties . . . 3568
Species Boundary and Initial Conditions . . . 3570
What Is the Segregated Species Solver? . . . 3578
Setting a Species Region Source . . . 3578
Species Transport Formulation . . . 3581
Species Transport Field Functions Reference . . . 3583

Modeling Turbulence and Transition 3584

Selecting a Turbulence Modeling Approach . . . 3586

Selecting a Transition Modeling Approach 3591
Using Reynolds-Averaged Navier-Stokes Turbulence Models 3593
Using Large Eddy Simulation. 3602
Using Detached Eddy Simulation 3613
Using the Synthetic Eddy Method for Initialization and Inflow 3616
Using Wall Treatment Models 3625
Using Spalart-Allmaras Turbulence 3642
Using K-Epsilon Turbulence 3666
Using K-Omega Turbulence 3745
Using Reynolds Stress Transport Turbulence 3794
Using Smagorinsky Subgrid Scale Turbulence 3827
Using Dynamic Smagorinsky Subgrid Scale Turbulence 3834
Using WALE Subgrid Scale Turbulence 3838
Using the Turbulence Suppression Model 3843
Using the Gamma ReTheta Transition Model 3853
Modeling Turbulent Heat Flux 3875
Shared Conditions and Values Reference 3884
Turbulence Field Functions Reference 3891

Modeling Radiation 3892

Selecting the Radiation Model 3894
Modeling Surface-to-Surface Radiation 3897
Modeling Participating Media Radiation 3927
Modeling Thermal Radiation 3951
Modeling Solar Radiation Loads 3967
Modeling View Factors 3979
Radiation Field Functions Reference 3995

Modeling Aeroacoustics 3999

What is Aeroacoustics 3999
How Do I Conduct an Aeroacoustics Analysis? 4002
Using Broadband Noise Source Models 4005
Using the Acoustic Wave Model 4038
Using the Direct Noise Simulation. 4064
Using the Ffowcs Williams-Hawkings Aeroacoustics Models 4064
Using the Lighthill Stress Tensor 4104
Aeroacoustics Bibliography 4109
Modeling Hydroacoustics 4112

Modeling Combustion and Other Reactions 4113
Using Combustion Models 4115
Using Eddy Break-Up (EBU) Models 4122
Using the Premixed Eddy Break-Up Model 4136
Using Eddy Contact Micromixing 4143
Using Surface Chemistry 4151
Using Homogeneous Reactor Models 4167
Using Turbulence-Chemistry Interactions 4230
Using the Coherent Flame Model (CFM) 4236
Using the Partially-Premixed Coherent Flame Model 4245
Using the Turbulent Flame Speed Closure (TFC) Model 4251
Using the Partially-Premixed Turbulent Flame Speed Closure (PTFC) Model 4263
Using PPDF Models 4270
Using the Progress Variable Model (PVM) 4322
Using the Thickened Flame Model 4342
Using the Nox Emission Model 4366
Using the Fuel Nox Model 4400
Using Coal Combustion 4418
Using the Soot Emission Models 4446
Using Lagrangian Particle Reaction 4472
Using the Polymerization Model 4497
Working with the Fuel Mass Fraction 4513
Defining Chemical Reactions 4515
Post-Processing Reaction Rates 4530
Importing Chemkin-Formatted Data. 4531
Using Ignitors 4540
Using Fluid Streams 4554
Using Laminar Flame Speed Methods 4559
Combustion Field Functions Reference. 4562

Modeling Multiphase Flow 4563
What Are Multiphase Flows? 4565
Defining Phase Interactions 4572
Using the Lagrangian Multiphase Model 4579
Using the Fluid Film Model 4883
Using the Discrete Element Method 4996
Using the Eulerian Multiphase Model 5090
Using the Eulerian Multiphase Mixture Model 5112

Using the Multiphase Segregated Flow Model 5130
Using the Dispersed Multiphase Model 5457
Using the VOF Multiphase Model 5477
Working with VOF Waves 5612

Modeling Dynamic Fluid Body Interaction 5636
What Is a Rigid Body? 5637
What Is a DFBI Rotation and Translation Motion? 5638
What Is a DFBI Embedded Rotation Motion? 5638
What Is a DFBI Morphing Motion? 5639
What Is a DFBI Superposed Rotation Motion?. 5640
What Is the Six DOF Solver? 5641
What Is the 6-DOF Motion Solver? 5643
Working with 6-DOF Bodies 5644
Working with 6-DOF Body Couplings 5683
DFBI Formulation 5695
DFBI Bibliography 5713

Modeling with Harmonic Balance 5714
Setting Up a Harmonic Balance Simulation 5716
What Is the Harmonic Balance Model? 5717
What Is the Harmonic Balance Flow and Energy Model? 5720
What Are the Harmonic Balance Turbulence Models? 5739
Defining a Wake 5741
What Is the Harmonic Balance Flutter Motion? 5755
What Is the Implicit Harmonic Balance Solver? 5766
Post-Processing Harmonic Balance Solution Data 5767
Filtered Mode Field Functions 5772
Harmonic Balance Formulation 5775
Harmonic Balance Field Functions Reference 5782

Modeling Solid Stress 5786
What Is the Finite Volume Solid Stress Model? 5787
Modeling Dynamic Problems 5804
Modeling Rayleigh Damping 5805
Modeling Large Displacements 5808
Setting Up the Finite Volume Solid Stress Model 5811
What Is the FVA Stress Solver? 5818
Solid Stress Formulation 5821
Shared Conditions and Values 5829

Solid Stress Field Functions Reference 5837

Modeling Electromagnetism 5840

Modeling Electric Currents 5841
Modeling Magnetic Fields 5878
Modeling Electrostatics 5925
What Is the Electric Potential Solver? 5930
Electromagnetism Formulation 5932
Electromagnetism Field Functions Reference 5949

Modeling Electrochemistry 5950

Modeling Li-Ion Battery Cells. 5951
Modeling Corrosion 5994
Modeling Ionic Species Flux 6008
Electrochemistry Nomenclature. 6021

Modeling Batteries 6024

What Is Battery Modeling? 6025
Selecting the Battery Model. 6036
What Is the Battery Solver? 6037
Working with Battery Design Studio. 6038
Working with a Battery Cell 6062
Setting Up a Battery Module 6083
Generating a Battery Module 6120
Generating a Battery Mesh 6128
Setting Up Connector Ohmic Heating 6131
Defining External Battery Casing 6137
Defining an Electrical Load. 6145
Creating a Circuit 6153
Exporting Electrochemistry Data 6179
Evaluating Electrode Structure Tortuosity 6181
Battery Modeling Formulation 6182
Battery Field Functions Reference 6191

Modeling Casting 6194

Casting Simulation with STAR-Cast 6194
About the Casting Module 6195
Selecting the Casting Model 6196
Working With the Casting Material Database 6198
Using Imported Casting Materials and Properties in STAR-CCM+. . . 6233
Using the Casting Shell Generator. 6238

Using the Macro Porosity Model to Detect Shrinkage-Related Defects . 6246
Defect Prediction Using Criteria Functions 6256
Artificial Compressibility 6272
Automated Time-Step Control 6274
Combined Thermal Contact Resistance 6279
Corrective Energy Source for Pure Thermal Simulations 6280
The Casting User Interface 6282
Casting Field Functions 6318
Casting Bibliography 6321

Modeling Region Sources **6322**
Marking Cells for Region Sources 6322
Setting Fluid Sources or Sinks. 6325
Prescribing Constant Temperature in Fluid Cells 6326

Remedying Cell Quality **6331**
Troubleshooting 6331
Cell Quality Remediation Field Functions Reference 6332

Controlling Domain Decomposition **6334**
The Partitioning Solver 6334

Solving Transport Equation **6340**
Finite Volume Discretization 6341
Gradient Computation. 6352
Implicit Iteration 6362
Algebraic Multigrid 6364
Transport Equation Nomenclature 6382
Transport Equation Bibliography 6384
Transport Equation Solvers Reference 6385

Using Tools 6389

Using Coordinate Systems **6391**
What Are Coordinate Systems? 6392
What Is the Laboratory Coordinate System? 6392
What Is a Local Coordinate System? 6393
Creating a Local Coordinate System 6402
Creating a Local Coordinate System Interactively 6404
Editing a Local Coordinate System Interactively. 6411
Transforming a Local Coordinate System. 6411
What Is a Block-Mapped Coordinate System?. 6415

Creating a Block-Mapped Coordinate System 6421
Displaying a Coordinate System in a Scene 6423
Copying Properties of Local Coordinate Systems 6425
Pop-Up Menu of the Coordinate Systems Manager 6425

Using Field Functions 6426
What Are Field Functions? 6427
What Are the Uses of Field Functions? 6438
How Do Field Functions Work in the GUI? 6438
Creating a User Field Function 6439
Copying a User Field Function 6447
Ignoring Boundary Values 6448
Interpolating Field Functions Using Inverse Distance Weighting . . . 6449
Field Function Programming Reference 6451
User Field Function Examples 6469
Context Menu of the Field Functions Manager 6472

Using Reference Frames 6473
What Is the Lab Reference Frame? 6475
What Is a Rotating Reference Frame? 6475
What Is a Rotating and Translating Reference Frame? 6478
What Are Relative Reference Frames? 6482
Working with Reference Frames 6485
Governing Equations in Moving Reference Frames 6491

Using Tables 6500
What Are Tables? 6501
Reading a Table Data File 6514
Creating and Filling a Table 6518
Exporting a Table 6521
Copying Part Data Between Simulations 6522
Displaying a Table 6525
Automatically Extracting and Exporting Tables 6526
Copying, Deleting, and Renaming Tables 6527
File Format Reference 6528
Pop-Up Menu of the Tables Manager 6531

Using Engineering Units 6532
What Is a Unit? 6532
Working with Units 6536
Using Units While Setting Values 6541

Interacting with CAD/CAE Products 6547

Importing into STAR-CCM+ 6549
Importing Geometry and Meshes . . . 6550
Importing STAR-CD Mesh and Solution Data . . . 6575
Importing CAE Models . . . 6579

Exporting from STAR-CCM+ 6608
Exporting Geometry . . . 6609
Exporting Mesh and Solution Data . . . 6610

Using File-Based Coupling 6618
Understanding File-Based Coupling . . . 6618
Mapping STAR-CCM+ Solution Data to a CAE Model . . . 6623
Exporting Mapped Data . . . 6633
Mapping and Exporting Transient Surface Data . . . 6638
Importing CAE Solution Data. . . . 6641
Mapping CAE Solution Data to a STAR-CCM+ Model . . . 6649
Using Data Mappers. . . . 6655
Exchanging Heat Transfer Coefficients . . . 6697
Working with Specific CAE Products . . . 6705
Coupling with RCAS . . . 6742

Co-Simulation with One-Dimensional Codes 6747
What are the 1D Coupling Models? . . . 6747
Activating a 1D Coupling Model . . . 6749
Using the 1D Coupling Solver . . . 6749

GT-POWER Coupling 6751
What Is GT-POWER Coupling? . . . 6752
What Is a GT-POWER Species? . . . 6756
What Is a GT-POWER Zone? . . . 6759
Preparing the GT-POWER Case . . . 6761
Setting Up the STAR-CCM+ Simulation . . . 6773
Coupling STAR-CCM+ to GT-POWER . . . 6787

OLGA Coupling 6801
What is OLGA Coupling? . . . 6802
What is an OLGA Species? . . . 6805
What is an OLGA Zone? . . . 6806
Setting Up the STAR-CCM+ Simulation . . . 6807
Coupling STAR-CCM+ to OLGA 5 . . . 6834

Coupling STAR-CCM+ to OLGA 7 6845
Area Mismatch Error in OLGA Coupling. 6856

WAVE Coupling **6857**
What is WAVE Coupling? 6857
What is a WAVE Species? 6859
What is a WAVE Zone? 6860
Preparing the WAVE Case 6861
Setting Up the STAR-CCM+ Simulation 6864
Coupling STAR-CCM+ to WAVE 6877

RELAP Coupling **6882**
What is RELAP Coupling? 6882
What is a RELAP Zone or Flow Zone? 6883
Preparing the RELAP Case 6886
Setting Up the STAR-CCM+ Simulation 6887
Coupling STAR-CCM+ to RELAP 6894

AMESim Coupling **6902**
Setting Up the AMESim System. 6903
Setting Up the STAR-CCM+ Simulation 6906

Reacting Channel Coupling **6920**
What is the Reacting Channel Coupling Model?. 6920
Reacting Channel Coupling Formulation 6937

STAR-CCM+ to STAR-CCM+ Co-Simulation **6940**
What is One-Way STAR-CCM+ to STAR-CCM+ Co-Simulation? . . . 6941
What is Two-Way STAR-CCM+ to STAR-CCM+ Co-Simulation? . . . 6942
STAR-CCM+ Co-Simulation Workflow 6943
Activating the STAR-CCM+ Co-Simulation Models 6944
Specifying Co-Simulation Conditions and Values 6947
Specifying Co-Simulation Exchange Conditions and Solver Settings . . 6969
Running the Co-Simulation. 6972

Partitioned Coupling Module **6974**
Using Partitioned Coupling with STAR-CCM+ to STAR-CCM+ Co-Simulation 6974

Co-Simulation API **6982**
Co-Simulation API Files 6983
Interacting with the STAR-CCM+ API Server 6985
Data Containers in Co-Simulation 6986

STAR-CCM+ Polyhedral Surface Mesh 6988
STAR-CCM+ Field Types for Exchange 6989
Loading the Co-Simulation API Library 6991
Checking the Co-Simulation API Version. 6991
Registering API Callback Functions 6991
Compiling Partner Code 6992
Running a Co-Simulation with Partner Code 6993

Co-Simulation with Abaqus 6994

Target Applications 6994
Preparing Your System for Co-Simulation 6996
Solving Fluid-Structure Interaction Problems 6998
Developing Standalone Models 6998
Preparing the Model in Abaqus 6999
Preparing the Model in STAR-CCM+ 7002
STAR-CCM+ Properties and Equivalent Abaqus Keywords. 7035
Running the Co-Simulation. 7036
Post-Processing the Co-Simulation Results 7037
Allocating Cores to Abaqus in Batch Queuing Systems 7038
Additional Considerations 7039

The STAR-CCM+ Component for Isight 7042

Best Practices 7045

Guidelines for Fluid Flow and Energy 7046
Mesh 7047
Thermophysical Properties 7049
Boundary Conditions 7050
Conjugate Heat Transfer (CHT) 7051
Cyclone Separators 7057
Freezing the Solvers 7058

Guidelines for Aeroacoustics Calculations 7059
Automobile Aerodynamics 7060
Sunroof Buffeting 7065
Radial Compressors 7067
Axial Fans 7071
Choosing DES, LES, or URANS for Aeroacoustics 7074
Mesh Issues 7075
Freestream Boundary Condition Issues 7076
Acoustics in Liquid Media 7077
Post-Processing the Acoustic Pressures and Pressure Perturbations 7078

Guidelines for Aerodynamics Calculations 7080
Incompressible External Aerodynamics: Steady State RANS Approach 7081
Transonic/Supersonic External Aerodynamics: Steady State RANS Approach 7085
Vehicle Simulations 7090

Modeling Fluid-Structure Interaction 7094
What Is Fluid-Structure Interaction? 7095
Modeling FSI Using Finite Volume Stress 7104

Modeling Atmospheric Boundary Layers 7110
What is the ABL? 7111
Setting ABL Parameters in STAR-CCM+ 7114
Recommended Values for ABL Modeling 7116
Bibliography 7118

Preparing CAD Models for Import 7119
Preparing CAD Models in CATIA 7120
Preparing CAD Models in SolidWorks 7121
Preparing CAD Models in NX 7122

Preparing CAD Models in Creo and Creo Elements/Pro....7123

Migrating from Other Codes....7125

Migrating from Fluent....7126

Migrating from STAR-CD....7128

Running....7130

Using Initial Conditions....7131

What Are Initial Conditions?....7132

Setting Initial Conditions....7133

Using Solvers....7140

Setting Up Stopping Criteria....7142

What Are Stopping Criteria?....7143

Using Automatically Generated Stopping Criteria....7144

Using Manually Generated Stopping Criteria....7154

Using Monitors as Stopping Criteria....7156

Using Multiple Stopping Criteria....7165

Pop-Up Menu of the Stopping Criteria Manager....7166

Obtaining a Solution Interactively....7167

Initializing the Solution....7168

Clearing the Solution....7171

Running the Simulation....7175

Stepping the Solution....7176

Watching Residuals....7178

Stopping the Solution....7183

Steering the Solution....7185

Saving Intermediate Solutions....7191

Scripting the Application....7198

What Is a Java Macro?....7199

Recording the Macro....7201

Saving the Macro....7206

Editing the Macro....7207

Playing the Macro....7219

Debugging a Macro Using an IDE....7225

Running Multiple Simulations Consecutively....7231

Playing Macros Within Macros....7233

Using the Java API 7234
Programming in Java 7236
Working with Default Macros 7237

Working with User Code 7243
What Is User Code? 7244
What Are User Libraries? 7246
What Are User Functions? 7248
Using User Functions in Parallel Simulations 7251
What is a Library Registration Function? 7252
Creating a New User Function 7253
Creating a New User Library 7254
Accessing a User Library 7257
User Function Interface Reference (C) 7261
Fortran User Interface Reference 7265
Variables Available to User Functions 7269
User Code Examples 7272
Using User Code on Particles and Parcels 7278
Pop-Up Menu of the User Code Manager 7291

Working with the Plugins Manager 7292
Deactivating an Installed Plugin 7296

Troubleshooting the Solution 7297
Getting Past a Poor Initial Solution 7299
Debugging the Simulation Setup 7303

Analyzing 7306

Accessing Solution Data through Derived Parts 7307
What Is a Derived Part? 7308
Working with Derived Parts 7309
Working with Implicit Sections 7338
Working with Scalar Value-Based Parts 7384
Working with Vector Value-Based Parts 7407
Working with Streamlines 7412
Working with (m', theta) Warps 7437
Working with Separation/Attachment Parts 7443
Calculating the Surface Area of a Part 7453
Working with Vortex Cores 7455

Working with Probes 7463
Working with Topological Parts 7495

Visualizing the Solution 7501
Working with Scenes 7502
Changing the Attributes of a Scene 7521
Displaying Geometry and Mesh 7566
Displaying Scalar Quantities 7578
Displaying Vector Quantities 7606
Displaying Streamlines 7640
Displaying Dynamic Fluid Body Interaction (DFBI) 7656
Animating Scenes 7666
Using Displayers 7683
Working with Volume Rendering 7732
Saving and Printing Scene Data 7765
Using Annotations 7772
Using Colormaps 7852
Using Transforms 7858
Working with Views 7875

Recording and Viewing Solutions 7894
Understanding the Simulation History File 7896
Recording Solution Data 7898
Retrieving Results from the Simulation History Files 7906
Using Solution Snapshots 7918
Animating Solution Views 7921

Reporting Results 7927
What Is a Report? 7928
Adding Reports 7930
Renaming Reports 7937
Running Reports 7938
Monitoring Reports 7942
Using Reports as Field Functions 7943
Using Coordinate Systems with Force and Moment Reports 7944
Using Bands with Force and Moment Reports 7946
Copying Properties of a Report 7947
Deleting Reports 7948
Report Types Reference 7949
Pop-Up Menu of the Reports Manager 8004

Reports Bibliography 8011

Monitoring the Solution 8012
What Is a Monitor? 8013
Customizing Monitors 8016
Using Report-Based Monitors 8023
Using Field Function Monitors 8025
Properties of the Monitors Manager 8037
Pop-Up Menu of the Monitors Manager 8038

Plotting Results 8039
Working with Monitor Plots 8043
Working with XY Plots 8051
Working with Histogram Plots 8063
Updating Plots 8071
Deleting Plots 8074
Renaming Plots 8075
Copying Plots 8076
Customizing Plots 8077
Adding and Removing Data in a Plot 8095
What Is a Tabular Data Set? 8096
What Is a Derived Data Set? 8098
Tabulating Data 8100
Exporting Data 8103
Printing and Saving Plot Images 8104
Plot File Formats Reference 8106
Pop-Up Menu of the Plots Manager 8108

Triggers and Update Events 8110
Setting the Trigger for an Update 8111
Trigger Types 8113
Trigger Properties 8114
User-Defined Update Events 8115

Using Data Set Functions 8130
Creating Data Set Functions 8132
What Are Point Time Histories? 8134
What Are Point Time Fourier Transforms? 8138
Using Point Time Fourier Transforms 8147
What Are Point-to-Point Time Fourier Transforms? 8149
What Are Point-to-Point Time Correlations? 8154

What Are Point Inverse Time Fourier Transforms? 8157
What Are Line Time Histories? 8162
What Are Line Time Fourier Transforms? 8165
Using Line Time Fourier Transforms 8169
What Are Line Spatial Fourier Transforms? 8170
Using Line Spatial Fourier Transforms 8172
What Are Surface Time Histories? 8173
What Are Surface Time Fourier Transforms? 8176
What Is Derived Data? 8180
Post-Processing Data Set Functions 8215
Pop-Up Menu of the Data Set Functions Manager 8238
Data Set Functions Formulation 8239

Wizard Guide 8247

Using the Turbo Wizard 8249
What is the Turbo Wizard? 8250
Turbomachinery Related Capabilities of STAR-CCM+ 8251
Setting Up the Turbo Wizard 8253
Saving and Loading Turbo Wizard Data 8296

Using the Fire and Smoke Wizard 8299
What is the Fire and Smoke Wizard? 8300
STAR-CCM+ Capabilities Related to Fire and Smoke 8301
Setting Up the Fire Properties 8302

Using the Thermal Comfort Wizard 8323
What is the Thermal Comfort Model? 8324
What is the Thermophysiological Occupant Model? 8327
What is the Thermal Comfort Wizard? 8345
What are TCM Passengers? 8346
What are TCM Boundaries? 8349
What is the Thermal Comfort Solver? 8351
STAR-CCM+ Capabilities Related to Thermal Comfort 8353
Setting Up a Thermal Comfort Case 8354
Performing a TCM Simulation 8381
Post Processing a TCM Simulation 8383
Working with TCM Input Files 8386
Working with TCM Output Files 8388

Bibliography 8395

Using the Wing Mesher 8396
What is the Wing Mesher? 8397
Setting Up the Wing Mesher 8399

Introduction to the SPEED Motor Tool Import Wizard 8414
Importing SPEED Data Into STAR-CCM+ 8415
Accessing the SPEED Motor Tool Import Wizard 8417
Using the SPEED Motor Tool Import Wizard 8419
Using SPEED Data in STAR-CCM+ 8447

Using The IDF Import Wizard 8453
Importing IDF Files 8454

Simulation Assistants 8458

The Simulation Assistant User Interface 8460
Tasks 8462
Collapsible Information 8465
Object Tree Focus 8466
Conditions 8468

Setting Up a Simulation Assistant Project in NetBeans 8474
Creating A Simulation Assistant Project In NetBeans 8475
Adding Java Packages, Java Classes, and XHTML Files to the Simulation Assistant Project 8477
Adding STAR-CCM+ Libraries to the Simulation Assistant Project 8483

Developing a Simulation Assistant 8488
Overview 8489
Files 8492
Tasks and Classes 8493
Star Actions 8495
Annotations 8496
Client Server Objects and Lookups 8497
Tasks 8499
Conditions, Triggers, and Predicates 8500
XHTML 8501
Supplied Classes 8502
Using the Simulation Assistant Wizard in NetBeans 8520

Navigating Help

New Features

Quick links to documentation for the latest features

How to use STAR-CCM+

Launch STAR-CCM+; understand the workflow; learn about the interface.

Working with Geometry

Learn how to create geometry in 3D-CAD, import geometry, manipulate geometry, and repair surfaces before meshing.

Simulation Topology

Transfer the geometry model to the CFD model using regions, boundaries and interfaces.

Generate Mesh

Understand mesh continua, mesh sizing, models for surface and volume meshing, and how to refine in chosen locations.

Modeling Physics

Explore all the physics models and options, accessed through physics continua.

Running the Simulation

Set initial conditions and solvers, apply stopping criteria, and run the simulation from the GUI or in batch.

Analyze the Solution

Create derived parts, reports and scenes; plot XY data; record solutions and play back later.

Using Tools

Learn how to create your own field functions, define local coordinate systems and reference frames, and work with tables and units.

Wizard Guide

STAR-CCM+ provides several wizards: Turbo, Fire and Smoke, Thermal Comfort, Wing Mesher, and SPEED Import.

Interacting with CAD/CAE

Integrate a wide range of CAD/CAE tools into your STAR-CCM+ simulation.

Working with CAD Clients

Embedded solutions to let you define your STAR-CCM+ simulation within your CAD interface.

Tutorials

Get step-by-step instructions on setting up many different cases with our models.

STAR-View+

A freely distributable and independent application for viewing 3D scenes from STAR-CCM+.

Getting Support

Access the Knowledge Base for extra help; prepare queries for CD-adapco Support Engineers.

New Features for v9.06

Top new features in the current version of STAR-CCM+® are:

- Optimate user guide and tutorials are now available within the STAR-CCM+ user guide
- Creation of blade neutral surface with the Turbo Slicing Tool
- Multi-component boiling model with Maxwell-Stefan diffusion
- Support for multiple coupling zones in STAR-CCM+ to STAR-CCM+ co-simulation and co-simulation with Abaqus
- Partitioned Coupling Module (PCM)
- Acoustic Wave Solver
- Morpher control point creation tool
- Offset mesher
- DEM compatibility with overset mesh
- Fluid film compatibility with overset mesh
- 2D meshing
- Tangential velocity for contact forces
- Hybrid MUSCL 3rd order discretization scheme
- Partial wrapping
- Plot preview for field functions
- Harmonic balance post-processing

A full list of new features for this release is provided in the Release Notes.

The next section provides a list of new tutorials for this release.

New Tutorials for v9.06

The following new tutorials are provided in this release:

- Body Force Propeller Method: Marine Self-Propulsion
- Acoustic Wave Modeling: Noise from a Cylinder
- Co-Simulation API: Spindle Valve
- Optimate: Design Sweeps Tutorial
- Optimate: Design of Experiments Tutorial
- Optimate: Pareto Optimization Tutorial
- Optimate: Use Existing Designs Tutorial
- Optimate: 2D Flap Optimization with Macro Tutorial

Using STAR-CCM+

This part of the documentation focuses on how to use the current version of STAR-CCM+, independent of the physics or models that may be encountered in a specific application.

It includes:

- A summary of the features of STAR-CCM+
- An explanation of the differences between the mixed precision and double precision versions of STAR-CCM+.
- An explanation of using the workspace—various aspects of the STAR-CCM+ graphical user interface (GUI)
- A description of the basics of working with simulations—specifically how to stop and start it and how to utilize the client/server architecture fully
- Information on getting support for STAR-CCM+
- An introduction to the innovation forum for STAR-CCM+, IdeaStorm

The other parts of the documentation cover the following:

- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Modeling physics—the concepts, theory, and implementation details of the models
- Using tools in STAR-CCM+ to facilitate your work
- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format

- A guide to using STAR-View+, a viewer for three-dimensional STAR-CCM+ scenes

What Is STAR-CCM+?

STAR-CCM+ provides the world's most comprehensive engineering physics simulation inside a single integrated package. Much more than just a CFD solver, STAR-CCM+ is an entire engineering process for solving problems involving flow (of fluids or solids), heat transfer, and stress. It provides a suite of integrated components that combine to produce a powerful package that can address a wide variety of modeling needs.

These components include:

- 3D-CAD modeler
- CAD embedding
- Surface preparation tools
- Automatic meshing technology
- Physics modeling
- Turbulence modeling
- Post-processing
- CAE Integration

STAR-CCM+ is based on object-oriented programming technology. It is designed to handle large models quickly and efficiently using a unique client–server architecture that seamlessly meshes and simultaneously solves and post-processes over multiple computing resources without requiring additional effort from the user. STAR-CCM+ recently became the first commercial CFD package to mesh and solve a problem with over one billion cells.

The object-oriented nature of the code can be seen in the user interface. An object tree is provided for each live simulation, containing object representations of all the data associated with the simulation. The objects presented on the simulation tree reside on the server, which can run as either a serial or a parallel process. A client interface connects to a server, and displays the simulation objects available on that server.

3D-CAD Modeler

3D-CAD is a feature-based parametric solid modeler within STAR-CCM + that allows geometry to be built from scratch. The geometry created with 3D-CAD is stored as 3D-CAD models, which can subsequently be converting to geometry parts for integration with the meshing and simulation process. A major feature of 3D-CAD is design parameters, which allow you to modify the 3D-CAD model from outside of 3D-CAD. These allow you to solve for a particular geometry, change the size of one or more components, and quickly rerun the case.

CAD Embedding

STAR-CCM+ simulations can be set up, run and post-processed from within popular CAD and PLM environments such as SolidWorks, CATIA V5, Pro/ENGINEER, and NX. STAR-CCM+'s unique approach gets you from CAD model to an accurate CFD solution more quickly or more reliably.

CFD results are linked directly to the CAD geometry (a process called associativity). After any modification in the CAD model the simulation results can be updated almost instantly by clicking the "update solution" button, allowing the rapid and thorough investigation of the design space.

Surface Preparation Tools

At the heart of STAR-CCM+ is an automated process that links a powerful surface wrapper to CD-adapco's unique meshing technology. The surface wrapper significantly reduces the number of hours spent on surface cleanup and, for problems that involve large assemblies of complex geometry parts, reduces the entire meshing process to hours instead of days.

The surface wrapper works by 'shrink-wrapping' a high-quality triangulated surface mesh onto any geometrical model, closing holes in the geometry and joining disconnected and overlapping surfaces, providing a single manifold surface that can be used to automatically generate a computational mesh without user intervention.

STAR-CCM+ also includes a comprehensive set of surface-repair tools that allow users to interactively enhance the quality of imported or wrapped surfaces, offering the choice of a completely automatic repair, user control, or a combination of both.

Automatic Meshing Technology

STAR-CCM+'s single integrated process provides the fastest, most automatic route from complex CAD to CFD mesh. Advanced automatic

meshing technology generates either polyhedral or predominantly hexahedral control volumes at the touch of a button, offering a combination of speed, control, and accuracy. For problems involving multiple frames of reference, fluid-structure interaction and conjugate heat transfer, STAR-CCM+ can automatically create conformal meshes across multiple physical domains.

An important part of mesh generation for accurate CFD simulation is the near-wall region, or extrusion-layer mesh. STAR-CCM+ automatically produces a high-quality extrusion layer mesh on all walls in the domain. In addition, you can control the position, size, growth-rate, and number of cell layers in the extrusion-layer mesh.

Physics Models

STAR-CCM+ delivers a powerful CFD tool in an easy-to-use environment. STAR-CCM+'s physics modeling capabilities include:

Solvers

- Segregated
- Coupled
- Finite volume solid stress

Time

- Steady state
- Implicit and explicit unsteady
- Harmonic balance

Turbulence

- RANS
- RSM
- LES/DDES
- Laminar-turbulent transition

Compressibility

- Ideal Gas
- Real Gas

Heat Transfer

- Conjugate Heat Transfer (CHT)
- Multiband and gray thermal surface to surface radiation
- Solar radiation

- Discrete ordinates radiation (DOM) including participating media

Multiphase

- Lagrangian particle tracking
- VOF (incompressible and compressible)
- Cavitation & boiling
- Eulerian multiphase
- De-icing & De-fogging
- Melting & solidification

Moving Mesh

- Dynamic Fluid Body Interaction (DFBI or 6DOF)
- Rigid body motion
- Mesh morphing
- Multiple reference frames (MRF)

Combustion & Chemical Reaction

- Reaction kinetics
- Eddy break up (EBU)
- Presumed probability density function (PPDF)
- Complex chemistry
- Ignition
- NOx modeling

Distributed Resistance (Porous Media)

- Anisotropic
- Orthotropic
- User defined
- Porous baffles

Turbulence Modeling

In addition to its provision for inviscid and laminar flow, STAR-CCM+ has a comprehensive range of turbulence models:

- k-epsilon (Standard, V2F, Realizable, Two-layer)
- k-omega (Standard, SST and BSL)
- Reynolds Stress (RSM – linear and quadratic)

- Spalart-Allmaras Turbulence models
- Boundary-layer transition
- Large Eddy Simulation (LES)
- Detached Eddy Simulation (DES, in the new Delayed Detached Eddy Simulation or DDES formulation)

Post-processing

STAR-CCM+ has a comprehensive suite of post-processing tools designed to enable you to obtain maximum value and understanding from your CFD simulation. This includes scalar and vector scenes, streamlines, scene animation, numerical reporting, data plotting, import, and export of table data, and spectral analysis of acoustical data.

CAE Integration

Several third-party analysis packages can be coupled with STAR-CCM+ to further extend the range of possible simulations you can do. Co-simulation is possible using Abaqus, GT-Power, WAVE and OLGA, and file-based coupling is possible for other tools such as Radtherm, NASTRAN and ANSYS.

Using Online Help

The online help is accessed through the Help menu for browsing, and through the `<F1>` key for context sensitive information.

The PDF file can be found within the `doc` directory of your STAR-CCM+ installation.

The STAR-CCM+ User Guide is provided in two formats:

Browsing the Online Help

To access the online Help for browsing, select **Help** > **Help** from the top menus, or press the `<F1>` key.



STAR-CCM+ will launch an HTML browser at the title page of the online Help. You can also jump directly to other parts of the online Help using the same menu:

- New Features
- Tutorial Guide
- Wizard Guide

The Help system is based entirely on JavaScript and displays help topics with the frameset view, which divides the browser window into three areas:

- A navigation frame on the left that displays the Contents tab (with a Table of Contents pane) and the Search tab
- A toolbar frame across the top that displays a set of buttons to navigate around the Help system
- A topic frame on the right to display the actual Help topics included in the Help system



This view gives the most control and is used to display Help requested from the Help menu. It will also display if you open `$STAR_HOME/doc/online/index.html` directly into your browser.

You can also read useful information about:

- The navigation frame
- The toolbar frame
- Changing the default browser
- Typographic conventions used in this documentation

Accessing Context Sensitive Help

You can jump directly to help topics that relate to specific objects in the simulation tree by pressing `<F1>` while the object is selected.

For example, you could press `<F1>` on the **Realizable K-Epsilon Two-Layer** model node within the **Models** manager of a physics continuum.



This will launch a browser showing the Help system open at the location of the documentation for the Realizable K-Epsilon Two-Layer model.


What Is the Realizable Two-Layer K-Epsilon Model? - Mozilla Firefox
File Edit View History Bookmarks Tools Help
Most Visited Customize Links Free Hotmail Windows Marketplace Development begins Windows Media Windows JIRA Style Guide
What Is the Realizable Two-Layer K-...
Contents
Search
Using STAR-CCM+
Modeling Geometry
Defining The Simulation Topology
Generating Mesh
Modeling Physics
Setting Up Physics
Modeling Space, Time and Motion
Modeling Materials
Modeling Flow and Energy
Modeling Species
Modeling Turbulence and Transition
Selecting a Turbulence Modeling
Selecting a Transition Modeling
Using Reynolds-Averaged Navier-
Using Large Eddy Simulation
Using Detached Eddy Simulation
Using Wall Treatment Models
Using Spalart-Allmaras Turbulence
Using K-Epsilon Turbulence
What Are the K-Epsilon Turbule
What Is a Low-Reynolds Nun
Modeling Physics > Modeling Turbulence and Transition > Using K-Epsilon Turbulence > W
Two-Layer K-Epsilon Model?
What Is the Realizable Two-Layer K-Epsilon Model?
The Realizable Two-Layer K-Epsilon model combines the Realizable K-Epsilon mo
models are identical, but model gains the added flexibility of an all y+ wall treat
Properties
Two-layer type
Selects a two-layer formulation.
Shear Driven (Wolfstein)
Selects the two-layer formulation of Wolfstein [62]. This is appropriate for flows that are not dominated by buoyancy forces. See Eqns. (448) to (451).
Buoyancy Driven (Xu)
Selects the two-layer formulation of Xu et al. [63]. This is appropriate for flows that are

Using the Navigation Frame

By default, the navigation frame displays two tabs:

- Contents
- Search

These tabs provide the core navigational features in the Help system.



In addition, you can make the best use of space by resizing.

Navigating Contents

The Contents tab displays the table of contents in the form of an expandable/collapsible tree view.

Closed book icons represent TOC entries that have subentries. Click a closed book to open it and see its contents. When you expand a book, the closed book icon is replaced by an open book icon. You can click the book again to close it. The page icons represent individual Help topics. Click a page icon or the tree text to display the contents of that page in the Contents pane. Clicking the text of an element with subentries will open the folder and display the selected page.



Searching the Help System

Using the *Search* tab, you can search the full text of the Help system.

Type the word or phrase to search for, and then press <`Enter`> or click **Go**. The search is case insensitive, and by default searches on whole words. So searching for “lin” would not find “Linux”. Use an asterisk to perform partial word searches, so for example, use “lin*” to find “Linux” as well as all other words beginning with “lin”.

The *Search* tab displays a list of all the topics in the Help system that contain the word or phrase you entered. If you search for multiple words, the search finds Help topics that contain all the words you entered. The topics found are ranked in order of relevance. The higher the ranking, the more likely the topic includes the word or phrase you searched for.



When you click one of the topics found by search, the page is loaded into the topic frame. If you are using Internet Explorer on Microsoft Windows, each occurrence of the search term(s) is/are highlighted as shown in the following screenshot. (The search highlighting feature does not work in other browser environments.)



Resizing the Navigation Frame

You can resize the navigation frame by dragging the vertical line that separates the navigation frame from the topic frame.

As you reduce the width of the navigation frame, the buttons will reduce in size until they start being hidden after they reach a minimum width, starting with the right-hand one.



By default, the width of the navigation frame is set to 300 pixels.

Using the Toolbar Frame

The toolbar frame displays a set of buttons to give you additional functionality, as shown below.



Show in Contents

Locates the currently displayed topic in the *Contents* tab. When you click `Show in Contents`, Help highlights the entry in the *Contents* tab that corresponds to the currently displayed topic. If the *Contents* tab is not currently the front-most tab, it will be brought to the front.

In the topic-only view, the **Show in Contents** button uses a different

icon. When you click the **Show in Contents** button in the topic-only view, the display switches to frameset view, and the navigation frame is displayed.

Previous

Displays the topic that precedes the currently displayed topic in book sequence. This is not the same as the browser's back button.

Next

Displays the topic that follows the currently displayed topic in book sequence. This is not the same as the browser's forward button.

Print

Prints the currently displayed topic.

Changing the Default Browser

You change the command that STAR-CCM+ uses to launch a browser.

This can be useful, for example, if the `<F1>` keyboard command fails to bring up the context-sensitive Help.

Select **Tools > Options**. This activates the *Options* dialog.



Open the **STAR-CCM+** node, then select the **Environment** node to access its properties, which include the *HTML Browser Command* property.



Select that property and enter the full path of the HTML browser into the text field.

Click **Close**. STAR-CCM+ will use that new path when it launches the online Help.

Understanding Typographic Conventions

This printed documentation online Help system uses the following conventions:

Convention	What it Means
BoldArial	Items you can select in the STAR-CCM+ interface, including buttons and menu items.
Italic	Components in the STAR-CCM+ interface such as windows, dialogs, and menus.
>	The greater-than sign indicates one of the following: • A menu choice. For example, “select `File > Open`” means “click the **File** menu, and then click **Open**”. • A node of a tree that has been opened. For example, "select the **Scenes** > **Part Scene** > **Part Displayer** node" means "expand the Scenes node, then expand the **Part Scene** node under that and finally click on the **Part Displayer** node under that".
`Code`	A fixed space font indicates text that you would type, such as on the command line or in macros.
`<Enter>`	Fixed space font between angled brackets indicates a single keystroke on the keyboard.

What Are the Important Concepts?

This section presents the key concepts you come across when using STAR-CCM+.

It begins with an introduction to the client-server architecture, which is the foundation on which the software is built. Further concepts are:

What is the Client-Server Architecture?

STAR-CCM+ uses a client-server architecture in which:

- The simulation objects are created and solved on the server.
- The workspace views these objects through the client.



The server can run in serial or parallel, with or without any client attached. You can launch a client workspace initially, and then use the new simulation option from within the workspace to launch a server. It is also possible to start a server process from the command line, and connect a client workspace to it later.

What Is the Client Workspace?

The client workspace is the graphical user interface that appears when STAR-CCM+ is launched interactively.

The key elements of this interface are shown in the clickable image below. You can jump to the documentation for different elements by clicking them.



What Are Simulations?

A simulation contains the solution state and data for a STAR-CCM+ analysis.

It exists on a server, and can only be accessed through a client connected to the server. When a client connects to a server, the object tree for the simulation on the server is shown in its own explorer pane in the client workspace. If the client is connected to multiple servers, the simulation object tree for each server is shown in its own explorer pane within the client workspace.

The root object within each simulation object tree is known as the simulation node. This node has its own properties and pop-up menu.



Simulation Properties

Name	The name of the simulation (read-only). To change the name of the simulation, use the save as action.
Collaboration	Controls whether other users may connect to this simulation's server.

	Activated	Allows other users to connect.
	Deactivated	Prevents other users from connecting.

What Are 3D-CAD Models?

STAR-CCM+ contains its own feature-based parametric solid modeler that allows you to create geometry for your analysis.

This integrated solid modeler is called the 3D-CAD module, and solid models generated from it are called 3D-CAD models. The screenshot below shows a single 3D-CAD model object within the simulation object tree.



Complete documentation on using the 3D-CAD solid modeler can be found in Using 3D-CAD.

What Are Parts?

There are three different types of part that can be used in a STAR-CCM+ simulation:

- Geometry parts

  These parts represent the model geometry, and are contained within the **Geometry** > **Parts** manager node of the object tree.



- Model parts, including regions, boundaries, and interfaces

  These represent the discretized portions of the geometry to be analyzed, on which physics models are applied. These parts are contained within the **Regions**, **Boundaries**, and **Interfaces** manager nodes.



- Derived parts

  These parts are created in a number of ways. They can be used in the definition of analysis reports and for visualizing solution data in scalar and vector scenes. All derived parts are contained within the **Derived Parts** manager node of the object tree.

Parts can be used as inputs to many other objects in STAR-CCM+, such as displayers, reports, tabular data sets, and x-y plots.

What Are Continua?

Continua are used to contain selections of physics or meshing models that are subsequently applied to one or more regions. Continua have no geometric definition associated with them.

The **Continua** manager node, which has its own pop-up menu, contains all of the continuum nodes in the object tree. Two types of continua may be present within this manager node: mesh continua and physics continua.

Mesh Continuum

A mesh continuum contains a selection of meshing models, such as the surface remesher, the prism layer mesher, and the polyhedral mesher. When a mesh continuum is applied to a region, the region is discretized according to the meshing models selected for the mesh continuum. An example of a mesh continuum with its selection of models is shown below.



See the section, Mesh Continua, for further information.

Physics Continuum

A physics continuum contains a selection of physics models, such as a chosen flow solver, material models, steady or transient time model, a turbulence model, and so on. Each physics continuum represents a

single substance that is present in all regions to which the physics continuum applies.



See the section, Using Physics Continua, for more information.

Managing Continua

Similar to other objects, a continuum may be renamed within STAR-CCM+. They can be deleted from an existing simulation as long as they are not in use. To prevent a continuum that is in use from being deleted, the **Delete** menu item for that continuum is grayed out. The contents of a continuum may also be copied from one continuum to another using the clipboard or by using drag-and-drop. The copy action causes all the models from the source continuum to be activated in the destination continuum, with the properties of each model also being copied. For physics continua, all reference values and initial conditions and values are copied from the source continuum to the destination continuum. For mesh continua, all reference values and volumetric

control settings are copied from the source continuum to the destination one.

The continuum copy feature can be used within the same simulation or between two different simulations, if both simulations have been opened in the same STAR-CCM+ session.

Further documentation is available on mesh continua, and on physics continua. Some documentation common to both types of continua is included in the following pages.

Storing Commonly Used Continua

Suppose that you repeatedly perform simulations using the same continua, that is, simulations involving the same models and materials. Rather than setting up the continuum each time a simulation is created on a new mesh, you can set up the continuum once only, save it, then reuse it multiple times.

To set up the continuum:

1. Launch STAR-CCM+ and start a new simulation.
2. Open the **Continuum** node, and select the required models using the model selection dialog.
3. Set the material properties and model coefficients.
4. Save the simulation.

To reuse this continuum:

1. Run STAR-CCM+ and open the existing simulation that contains only the saved continuum.
2. Import the required mesh(es).
3. Save the simulation with a new name.
4. Proceed with the simulation as normal.

Pop-Up Menu of the Continua Manager

These pop-up menu items are available when you right-click the **Continua** manager node.

New Continuum	Adds a continuum with a default name to the model and three manager subnodes (models, reference values, and initial conditions).

	Mesh Continuum	Creates a mesh continuum.
	Physics Continuum	Creates a physics continuum.

See also: common menu items.

What Are Models?

Model nodes in the simulation object tree represent simulation features that are active in the current case.

Model nodes are contained within the Models manager nodes of meshing and physics continua. Meshing models control the construction of surface and volume meshes, and physics models define the behavior of the specified material in response to the conditions imposed upon it. Examples of both types of model node are shown below.



Model nodes are added in response to selections made using the model selection dialogs that can be accessed using the pop-up menus of the

Models manager nodes. Further information on selecting and reviewing models is given in the following sections:

- Selecting meshing models
- Selecting physics models

Some model nodes have properties that can be set in the *Properties* window. These properties set the default behavior of the model in all regions to which the parent continuum is applied. It may be possible to override this behavior at region or boundary levels if appropriate controls are provided. An example of a physics model with properties is shown below.



Some model nodes have additional options on their right-click menus, in addition to common menu options.

Models Manager Pop-Up Menu

Select Meshing Models...	In the mesh continuum, executes the select models action to allow selection of the models to be activated.
Select Models...	In the physics continuum, executes the select models action to allow selection of the models to be activated.

See also: common menu items.

What Are Regions?

Regions are volume domains (or areas in a two-dimensional case) in space that are completely surrounded by boundaries.

They are not necessarily contiguous, and are discretized by a conformal mesh consisting of connected faces, cells and vertices.



Regions are created either when a volume mesh is imported, or when a surface is imported without using geometry parts. Each region is given a unique name during the import process, but it can be renamed. One region is joined to another using an interface, so that information can be passed between them. Boundaries are not shared between regions.

There may be multiple regions in a simulation, each with a unique name and properties, as well as a pop-up menu. Region nodes are created when you:

- Import a mesh (from native CAD, neutral format, surface mesh or volume mesh formats)
- Split a non-contiguous region

The node icon is based on the region type. You can also work with region nodes with drag-and-drop actions.

Region Properties

Index	Index reference for region (read only).
Mesh Continuum	Defines the mesh continuum associated with this region. Select from a drop down list of all the available continua in the simulation.

Physics Continuum	Defines the physics continuum associated with this region. Select from a drop down list of all the available continua in the simulation.	
Parts	Specifies the parts included in the region, using the parts selection dialog.	
Type	Defines the type of the region depending on the material model selected. This is primarily used to define porosity in a region associated with a gas or liquid physics continuum. A region associated with a solid continuum only allows one option. This property will be changed if the material model of the associated continuum is changed from gas or liquid to solid (and vice versa).	
	Fluid Region	The region is associated with a gas or liquid continuum. No porous modeling is applied.
	Solid Region	The region is associated with a solid continuum.
	Porous Region	The region is associated with a gas or liquid continuum, and porous media modeling assumptions apply.

Region Pop-Up Menu

Highlight	Highlights a region in a scene. This option is only available if the scene is open.
Edit...	Opens the table editor, which combines the functionality of the simulation tree and *Properties* window. It displays the selected region, all of its child nodes, and their properties in a single dialog.
Combine	Executes the combine regions action to combine the selected regions (two or more) into one region.

Boolean	Applies Boolean operations in the submenu to the selected regions.	
	Unite Regions	Performs a Boolean unite operation.
	Intersect Regions	Performs a Boolean intersect operation.
	Subtract...	Performs a Boolean subtract operation.
Split by Surface Topology	Executes the split by surface topology action to split closed, semi-manifold surface meshes that belong to a single region into their constituent regions.	
Split Non-Contiguous	Executes the split non-contiguous region action to split the selected regions into multiple contiguous regions if possible. If more than one region is selected, the action will operate on each region in sequence.	
Split by Function...	Executes the split by function action to split a region according to the rounded integer value of a selected field function.	
Remove Invalid Cells...	Executes the remove invalid cells action to move cells below a user defined threshold quality into another region so that they are not included as part of the analysis.	
Transform	Scales, translates or rotates the selected regions and all associated boundaries. These actions operate on one or multiple regions, and if more than one region is selected, the transform action will be applied to each region in sequence.	
	Scale...	Executes the scale region action.
	Translate...	Executes the translate regions action.
	Rotate...	Executes the rotate region action.
	Reflect...	Executes the reflect region action.

Reorder Mesh	Reorders the internal references to cells in the volume mesh region in order to optimize the solver performance in terms of memory usage.	
Replace Mesh...	Executes the replace mesh action to replace the mesh of the selected region with that of another.	
Solution	Applies operations in the submenu to the individual region.	
	Clear	Executes the clear region solution action.
Create Interface	Creates a heat exchanger interface between selected regions.	

See also: common menu items.

Using Drag-and-Drop with a Region Node

- Dragging a region node onto another region node will copy all the properties of the source region to the destination region, including region condition methods, but excluding the region name.
- Dragging a region onto the **Parts** node of a scene displayer will add all the boundaries to the scene.
- Dropping a continuum onto a region will associate the region with that continuum (the same as setting the Continuum property).

What Are Boundaries?

Boundaries are surfaces (or lines in a two-dimensional case) that completely surround and define a region.

Each boundary has a corresponding node in the simulation tree, and has its own properties and pop-up menu. The figure below shows a simple region where the lines surrounding the region represent the boundaries. In this example, the region could be surrounded by a single boundary or multiple boundaries. The choice depends on what conditions and values need be assigned.



The discretization of a boundary corresponds to the surface bounding the volume mesh for the region. Boundaries are never shared between regions; a boundary can belong to only one region. The figure below shows an example of three separate regions. The dotted lines represent the boundaries that separate each region. The boundaries at the interface of one region and its neighbor are coincident in space.



Boundaries are created by:

- Volume mesh import or surface mesh import;
- Interface creation; or
- Boundary splitting

Only boundaries that belong to a surface mesh can be deleted.

Each boundary is given a unique name during the mesh import process, but it can be renamed, unless it is of the interface type.

Boundary Properties

Index	Index reference for boundary (read only). This index is available through the `Boundary Index` field function.	
Interfaces	The interfaces that were created from this boundary (read-only). Blank if this boundary has not been in any interface creation.	
Part Surfaces	Specifies the surfaces included in the boundary, using the parts selection dialog.	
Type	The boundary type. The full list of possible boundary types is listed below, but the actual list you will see depends on the models you have selected. The list of types changes depending on whether the boundary is a regular one, or one created by an interface.	
	Axis	The axis of an axisymmetric region. The two icons are for a fluid and solid region respectively.
	Flow Split Outlet	A flow split outlet boundary.
	Free-Stream	A free-stream boundary.
	Mass Flow Inlet	A mass flow inlet.
	Pressure Outlet	A pressure outlet.
	Stagnation Inlet	A stagnation inlet.

	Symmetry Plane	A symmetry plane.
	Velocity Inlet	A velocity inlet.
	Wall	A wall.
	Displacement	A solid stress displacement boundary
	Traction	A solid stress traction boundary
	Fixed	A solid stress fixed boundary
	Free	A solid stress free boundary
	Pressure	A solid stress pressure boundary
	Normal Displacement	A solid stress normal displacement boundary
	Force	A solid stress force boundary
	The list of possible types, and the icons, for boundaries generated by interfaces are listed below.	
	Baffle	Generated from a baffle interface.
	Contact	Generated from a contact interface. The two icons are for the master slave boundaries respectively.
	Fan	Generated from a fan interface
	Fully Developed	Generated from a fully developed interface.
	Internal Interface	Generated from an internal interface.
	Porous Baffle	Generated from a porous interface.

	Mixing Plane	Generated from a mixing plane interface.
Parent Interface	The name of the interface that created the boundary for internal use (read-only). Only available on a boundary created by an interface.	

Boundary Pop-Up Menu

Highlight	Highlights a region in a scene. This option is only available if the scene is open.	
Edit...	Opens the table editor, which combines the functionality of the simulation tree and *Properties* window. It displays the selected region, all of its child nodes, and their properties in a single dialog.	
Create Interface	Creates an interface of a particular topology as defined in the submenu. The topology can be changed after creation.	
	In-Place	Creates an in-place interface.
	Periodic	Creates a periodic interface.
	Repeating	Creates a repeating interface.
	Indirect	Create an indirect interface.
Create Shell Region	Creates a shell region from the selected boundary.	
Combine	Executes the combine boundaries action to combines two or more boundaries into one. This is only active when more than one boundary is selected.	
Fuse...	Executes the fuse boundaries action to fuse the boundaries together. The selected boundaries from the tree are placed in the Selected list of the *Fuse Boundaries* dialog.	

Split Non-Contiguous	Executes the split non-contiguous boundary action to split non-contiguous boundaries (which may not be contiguous in space), into multiple boundaries each of which are contiguous. This will operate on all selected boundaries.
Split By Angle...	Executes the split boundary by angle action to allow you to split a boundary based on the angle between adjacent faces. All the selected boundaries from the simulation tree are placed in the Selected list of the dialog.
Split By Function...	Executes the split boundary by function action to split a boundary based on the value of a field function. All the selected boundaries from the tree are placed in the Selected list of the *Split Boundaries by Function* dialog.
Split by Patch...	Executes the split boundary function action to split a boundary based on the underlying patch definitions of the surface mesh. This function can only be used if more than one patch definition exists in the imported surface.
Split by Feature Curves...	Executes the split boundary function action to split a boundary based on the feature curves in the surface mesh region. Feature curves must exist for the region for this option to be active. All the selected boundaries from the simulation tree are placed in the Selected list of the dialog.
Split Interactively...	Executes the split boundary function action to split a boundary interactively in the current scene. A scene must be open for the option to be active.
Project to Plane...	Executes the project boundary to plane action allowing you project a boundary onto a user-defined plane. All the selected boundaries from the tree are placed in the Selected list of the *Split Boundaries by Function* dialog.
Convert to Interface(s)	Converts the boundary into an interface by creating an in-place interface between the

	original boundary and a copy in order to allow additional regions to be made. The boundary copy is automatically made by the process.	
Create New Region from Boundaries	Create a new region from boundaries that make up a closed volume. Boundaries are removed from the original region and moved into the new region definition.	
Solution	Applies operations in the submenu to the individual region.	
	Clear	Executes the clear region solution action.

See also: common menu items.

What Are Interfaces?

Interfaces provide a connection between boundaries during the simulation meshing setup and/or analysis process.

They are always composed of boundary pairs, either in the same region or in different regions. There can be many (or zero) interfaces in a simulation.

When using the surface remesher, tetrahedral, polyhedral or prism layer meshers, interfaces may be defined to ensure that a conformal mesh is obtained between regions or within the same region if a baffle plane is required.

In a volume mesh, interfaces permit mass, energy and other continuum quantities to pass from one region to another. In the illustration below, the dotted black line indicates the location of the interfaces that are created using the coincident boundaries between each region:



Interfaces have a topology and a type associated with them. In general, the topology refers to the physical arrangement of the interface boundaries. For example, in-place and repeating interfaces normally have coincident boundaries while a periodic interface has separated boundary pairs.

The interface type determines the modeling behavior of the interface. For example, depending on the topology selected, you can:

- Allow mass and energy to be exchanged between fluid and porous regions that have coincident boundary planes.
- Simulate baffles/porous baffles within a region.

- Provide a contact interface for heat transfer between fluid-solid and solid-solid regions.
- Specify a fully developed flow condition to simulate the pressure drop or mass-flow across a periodic (cyclic) across a region.

See Also

What Is a Generalized Cylinder?

What Are Solvers?

In STAR-CCM+, solvers control the solution and are activated once per iteration (or once per time-step for explicit unsteady simulations).

Typically, a model elects the solvers that are required. Different models can use the same solver and sometimes models need more than one solver.

Some models control other models and some only perform tasks when the solution is initialized. For this reason, not all models elect solvers. Models can also elect solvers that perform subsidiary tasks in addition to the solver that the model elected to control it

Why Are Solvers Different From Models?

The essential difference between a solver and a model is that the scope of a model is limited to the continuum in which is defined, while the solver can span continua.

This distinction between solvers and models facilitates the simultaneous solution of physical phenomena across regions that are associated with different continua, such as conjugate heat transfer. The figure below, which shows two continua, illustrates this distinction:



- one region that is called Fluid Continuum, representing a fluid
- one region that is called Solid Continuum, representing a solid

Suppose the following are defined:

- a flow model and a fluid energy model in a Fluid Continuum

- a solid energy model in a Solid Continuum

A solver would be activated in the Fluid Continuum to control the flow solution in that continuum and another solver would be activated in Solid Continuum to control the energy solution. In this situation, a single solver would control two different models (fluid energy and solid energy) in both continua, which allows the implicit solution of the conjugate heat transfer.

What Are Reports, Monitors and Plots?

Reports provide a computed summary of simulation or system data.

They are contained within the **Reports** manager node in the object tree, and have varying properties that depend on their type.

Monitors provide a mechanism for sampling summary information from the simulation during runtime. Some built-in monitors are provided for residuals that the various solvers compute. User-defined monitors can be created based on reports that are contained within the Reports manager node.

Plots provide line graphs containing information that is supplied from; monitors, extracted simulation data, tables, or derived data sets. These plots can be customized with different line and point properties, axes, scales, titles, and so on.

An example object tree showing all three of these object types is shown below. In this particular example, the Drag and Lift Coefficient reports have been used to create monitors, which in turn have been used within a single plot called Force Coefficients Plot.



What Are Scenes?

Scenes are used to visualize the geometry, mesh, and solution data that is contained within a simulation.

They are viewed through the *Graphics* window in general, and can also be written to file for a hard copy. Four general types of scene can be created:

- Geometry scenes to view outlines and surfaces of geometry parts and regions.
- Mesh scenes to view the surface and volume mesh grids on the surfaces of parts and cell sets.
- Scalar scenes to view contours of scalar quantities on part surfaces.
- Vector scenes to view arrow vectors generated from vector quantities on part surfaces.

Scenes are often used with derived parts to obtain solution quantities at particular locations within the simulation. Streamlines can also be created through the streamline derived part.

What Is The General Workflow?

A workflow is the sequence of operations you must work through in order to achieve a certain result.

In STAR-CCM+, there is a general sequence of operations that is helpful to follow. You will, through experience, develop a workflow that best suits your modeling environment.

The most general view of the STAR-CCM+ workflow is depicted in the illustration below. Elements on the illustration can be clicked to jump to further documentation for that step.


General Sequence of Operations
in a STAR-CCM+ Analysis
Prepare the Geometry
Construct the Simulation Topology
Generate the Mesh
Define the Physics
Prepare for Analysis
Run the Simulation
Analyze the Results

Preparing the Geometry

A STAR-CCM+ analysis requires a geometry to represent the actual object or scenario being modeled.

You can access this geometry by:

1. Creating it using the integrated 3D-CAD module in STAR-CCM+
2. Importing an existing CAD model in its native or neutral format
3. Creating a simple shape part directly (block, cone, cylinder, or sphere).
4. Importing an existing surface mesh for which the underlying CAD is no longer available

It is likely that the initial geometry requires some further manipulation before it can be used in a simulation. Some common operations include:

- Repairing the surface of the CAD model
- Imprinting adjacent parts
- Extracting the fluid volume
- Transforming parts

When you prepare your geometry for a STAR-CCM+ simulation, consider the physics models that you are going to be using in the simulation.

For example, if you are modeling fluid flow through a valve:

- assign an inlet boundary to the surface where fluid enters
- assign an outlet or pressure boundary to the surface where fluid leaves.

To assign boundaries to surfaces, there must be separately selectable surfaces for both boundaries. To split surfaces at the geometry stage, use the following tools:

- 3D-CAD models: set face names
- Geometry parts: split surfaces by patch, angle, part curves, or non-contiguous

Defining the Simulation Topology

Geometry is defined in terms of geometry parts, surfaces, and curves. The computational model to which physics can be applied is defined in terms of regions and boundaries. Defining the simulation topology is the process of mapping between the geometrical definition of the problem and the computational (or physical) definition.

Geometry parts can be assigned to regions, part surfaces can be assigned to boundaries, and part curves can be assigned to objects called feature curves. This process of mapping has to be complete before you can create any surface or volume meshes.

General Process for Mapping Geometry Parts to Regions

(1) Prepare Geometry for Analysis



(2) Choose Assign Parts Action



(3) Select Options on Dialog



(4) Review Newly Created Regions, Boundaries and Interfaces



Tip:
(a) Select all connected geometry parts;
(b) **Create a Region For Each Part**, **Create a Boundary for Each Part Surface**, and activate **Create Interfaces from Contacts.**

This will automatically create interfaces between the parts.

Legacy Workflow: Importing Directly to Regions

In this release of STAR-CCM+, you can import geometry directly as regions and boundaries without going through geometry parts. Regions can be transformed in the same way as geometry parts, and can be manipulated to create more regions, or to combine existing ones. Boundaries can be manipulated to give the correct geometrical definition that is required for the simulation.

Generating the Mesh

The first step in generating a mesh is to choose which meshing strategy to use and select the meshing models.

STAR-CCM+ provides two meshing strategies:

- Parts Based Meshing
- Region Based Meshing

There are advantages to using parts-based meshing, although currently not all meshing models are available using the parts-based meshing strategy. To help you decide which strategy to use, see Comparison of Meshing Strategies. The following graphic shows the parts-based meshing strategy:



Meshing Models: Surface and Volume

Meshing models are split into two groups: surface mesh models and volume mesh models. Work with the surface meshing models initially, and when you have a high-quality surface mesh, go on to choose volume meshing models.

Two surface meshing models are provided:

1. The surface remesher, which is always required for geometries that have been imported from CAD models or created in 3D-CAD.
2. The surface wrapper, which can be used to:
   - extract fluid volumes;
   - simplify complex assemblies; or to
   - provide a closed surface over poor quality CAD.

Note: In parts-based meshing, the surface wrapper is created as a separate mesh operation which allows you to refine the mesh independently to the other meshing models.

Several volume meshing models are provided. The main models are:

- Polyhedral mesher

  Suited to complex, multi-part geometries.
- Trimmer

  Uses a predominantly hexahedral mesh. It is efficient at filling large volumes and uses less memory per cell than the polyhedral mesh.
- Prism layer mesher

  Used with the previous two models to provide layers of prism cells next to the wall. The prism layer mesher is essential for turbulent flows.
- Thin mesher (Currently available in region-based meshing only)

  For thin geometries (such as solid plates).
- Generalized cylinder (Currently available in region-based meshing only)

  For cylindrical geometries such as internal pipe flows.
- Extruder (Currently available in region-based meshing only)

  Used to extend volume mesh as an offset from selected boundaries.

Meshing Model Settings and Sizes

For each meshing model you select, a corresponding model node is added within the Automated Mesh operation. Global settings in the *Properties* window are accessed by selecting these nodes.



- You can refine the mesh settings by expanding the **Automated Mesh > Default Controls** node.



Volumetric controls allow you to set custom mesh sizes within a volume of a specified shape or geometry part. See Mesh Refinement In A Specified Volume.

Meshing Rules of Thumb

- Assign parts to regions then set region and boundary types before running the meshers.
- Define a minimum of 4–5 cells across a thin channel.
- For flows where heat transfer or lift and drag is important:
  - Have at least 15 cell layers in the prism layer.
  - Check that final Y+ values are less than 3.
- Refine the mesh when high gradients are present:
  - Recirculation
  - Jet flows
  - High temperature differences

Always consider what engineering quantities you want to derive from the simulation, and construct your mesh to suit.

Visualizing the Mesh

To visualize the mesh:

- Right-click **Scenes** and choose **New Scene** > **Mesh**
- Ensure that the correct representation is active in the *Graphics* window. Change it by expanding the **Representations** node, choosing the one you want (for example, **Volume Mesh**), and dragging it onto the active scene.

Defining the Physics

STAR-CCM+ contains a wide range of physics models and methods for the simulation of single- and multi-phase fluid flow, heat transfer, turbulence, solid stress, dynamic fluid body interaction, aeroacoustics, and related phenomena. These physics models are all selected using a physics continuum.

Simulations can have multiple physics continua, although a mesh region can only be associated with one.

(1) Right-click the physics continuum **Models** node and choose **Select models...**



(2) Choose the relevant models in the *Physics Model Selection* dialog



Tip:
Auto-select option will choose most commonly used models from each group box on the left side. Clear it to access alternatives.

Material Properties

Material properties are stored on physics continua. A continuum represents a single material, although this continuum could be multi-

component or multi-phase. Properties can be constant, or set to vary according to one or more factors.

- Select a material model in the *Physics Model Selection* dialog.
- Right-click the model node and choose **Replace with...**to access the materials database.



Region Types

Regions can be set to fluid, porous or solid types using the *Type* property of the region node.

Boundary Conditions

Boundary conditions are applied with the physics models. The boundary type can be set on a boundary before the physics continuum has been populated with models, but actual values can only be set after the relevant models have been defined.

Initial Conditions

Initial conditions set the initial field values within the simulation--for example, velocity and pressure. The choice of values can have an impact on how quickly STAR-CCM+ arrives at the correct solution. These values are set using the nodes for each field under the **Initial Conditions**manager for each continuum.


Continua
Mesh 1
Physics 1
Models
Reference Values
Initial Conditions
Pressure
Static Temperature
Turbulence Specification
Turbulent Viscosity Ratio
Velocity
Constant
Regions

Preparing for Analysis

STAR-CCM+ uses an iterative procedure to reach the solution to the transport equations that satisfies the boundary conditions for a chosen scenario.

Knowing when this process has converged to an acceptable level is an essential factor in running a successful simulation. While a residual plot is opened by default during the simulation, monitoring engineering values also provides an additional check on the solution progress. These values can be obtained using reports.

Reports, Monitors, and Plots

Reports are used to extract computed numerical data from the simulation. Monitors use reports to record the reported value as the solution progresses. Monitored data can be plotted to show trends while the solution is running.

To create a report, monitor and plot:

- Right-click the **Reports** node; choose **New** followed by the type of report.
- Set the report properties (including the parts from which to extract the data).
- Right-click the report node; choose **Create Monitor and Plot from Report**.

The monitor and plot nodes appear within the relevant manager nodes as shown below.



Graphics Scenes

STAR-CCM+ also provides dynamic updating of graphics scenes showing scalar and vector data. These updates provide visual confirmation that the simulation has been set up correctly. Hard copies of these scenes can be written to file at selected frequencies in order to provide the input for animations.

Stopping Criteria

Formal stopping criteria can be defined to monitor progress of the simulation, and to halt it when all the criteria are met.

- Right-click the **Stopping Criteria** node, choose **Create from Monitor**, followed by the name of an existing monitor node.
- Select the new node and choose the *Criterion Option* in the *Properties* window. Set any further values as required.

Running the Simulation

Starting the simulation process is achieved by clicking **Run** on the **Solution** toolbar.

The solution is initialized and the solver is launched. For an interactive session, residuals are plotted in the client workspace (as shown below), and reported in the *Output* window. For batch sessions, residuals are echoed to the command console. The client-server architecture means that you can disconnect and reconnect to the running server at your convenience.



If the residuals increase from step to step, instead of decreasing, the solution is diverging. If left alone, the solution eventually ends with a floating point exception. To solve this problem, check the section Troubleshooting the Solution, for ideas on how to proceed. After changing the setup, clear the solution and run the simulation again.

Analyzing the Results

Results can be visualized using scenes. Scenes are composed using one or more displayers, where each displayer can be created independently of the others.

The following displayer types are available:

- Geometry
  - Used to display geometry
  - You can alter; color, opacity, and shading
- Scalar
  - Used to plot field functions on parts
  - Can be displayed as filled (with or without mesh), smooth filled and as contours
- Vector
  - Used to plot vector quantities
  - You can alter vector styles and distribution
- Streamline
  - Only used to display streamlines (which are created using derived parts)
  - You can plot streamlines as tubes, lines, or ribbons
  - You can animate streamlines using the toolbar



Creating New Scenes

- Right-click the **Scenes** node and choose **New Scene** followed by the type.
- To switch to a view showing only the objects for the active scene, click the **scene/plot** button above the simulation tree.
- Ensure that the volume mesh representation is active.
  - Expand the **Representations** node.
  - Drag the **Volume Mesh** representation onto the active scene.

To create an overlay effect in the scene, change the opacity of a displayer.

Derived Parts

If the model geometry alone does not provide sufficient surfaces on which to display simulation data, various alternative options are provided on the **Derived Parts** node. These options include:

- Plane sections
- Isosurfaces
- Threshold-based parts
- Streamlines

- Probes
- Vortex cores

To create a new derived part:

- Select a scene in the *Graphics* window
- Right-click the **Derived Parts** node and choose **New Part,** followed by the type of part required

  An in-place dialog appears above the simulation tree.
- Select options as required, then click **Create**.

Working with Mixed or Double Precision

STAR-CCM+ comes in two versions: a mixed precision version and a double precision version.

Refer to the release notes for a list of the platforms for which a double precision version is available.

To run simulations with double precision, install a double precision version of STAR-CCM+. The interface and workflow for both versions are identical.

This section describes the following information:

- What is precision?
- Which variables are double precision?
- Advantages and limitations of the double precision version

Precision

The term precision refers to the size of (and hence, number of digits of precision that are stored for) floating point data. A single precision float occupies 4 bytes of memory and has 7 digits of precision. A double precision float occupies 8 bytes of memory and has 15 digits of precision.

Double Precision Variables

In the mixed precision version of STAR-CCM+, the fields of double precision floating point data include:

- Coordinates (that is, vertices locations in space)
- Pressure
- Displacements (for example, solid-stress displacements).

In the double precision version of STAR-CCM+, all fields of floating point data are double precision.

Fields of other types (for example, integers and booleans) are the same in either version.

Advantages and Limitations of a Double Precision Simulation

The double precision version of STAR-CCM+ is generally intended for cases that involve physical phenomena of vastly differing scales.

Examples include wave propagation in water, three-dimensional microstructure battery modeling, electrodynamics simulations and multiphase flows.

The advantage of running these simulations in double precision is increased accuracy due to less round-off error. This increase in accuracy leads to robust convergences of problems that would otherwise not converge to an adequate result in single precision.

When running simulations with double precision:

- The memory requirements are 20% greater than with mixed precision.
- The time per iteration is larger by 15-20%.
- The size of `.sim` files increases by not more than 100% (depending on the proportion of floating point fields for the case).

Launching STAR-CCM+

This section describes how to launch STAR-CCM+ after a successful installation.

STAR-CCM+ operates using a client-server architecture. The client software provides an interactive Graphical User Interface (the workspace) that connects to a simulation running on one or more server processes. Two steps are required to launch an interactive session of STAR-CCM+ for a new simulation:

1. Launch the STAR-CCM+ client. There are two procedures for launching the STAR-CCM+ client:
   - one for Windows
   - another for Linux
2. Start a new simulation on one or more server processes.

Separate sections of the user guide provide detailed information on the client-server architecture that is used in STAR-CCM+, and also on running parallel servers for large simulations.

If you are using STAR-CCM+ for the first time, work through the Introductory Tutorial before going any further.

Launching a Client on Windows

Since the Windows installer sets up an entry in the **Start** menu, you can start the client from the **Start** menu assuming you did not change any of the defaults during the installation:



If everything is set up correctly, a splash screen is displayed:

The components being loaded are displayed at the bottom of the splash screen on the left.

The splash screen is displayed for up to 30 seconds, although it is often much less, depending on the capability and load on your machine. STAR-CCM+ then opens in graphical user interface (GUI) mode without loading an existing simulation or creating a simulation.



After launching the client you can go on to start a new simulation.

Launching a Client on Linux

Make sure that you know the full installation path to the STAR-CCM+ installation that you wish to use.

In the command example below, [INSTALL_DIR] shows the path with STAR-CCM+ version 9.02.003. To launch a new STAR-CCM+ client, enter the following at the command prompt:

```
% [INSTALL_DIR]/starccm+9.02.003/star/bin/starccm+ &
```

If everything is set up correctly, a splash screen is displayed:

The components being loaded are displayed at the bottom of the splash screen on the left.

The splash screen is displayed for up to 30 seconds, although it is often much less, depending on the capability and load on your machine. STAR-CCM+ then opens in graphical user interface (GUI) mode without loading an existing simulation or creating a simulation.



After launching the client you can go on to start a new simulation.

Working with Simulations

A STAR-CCM+ simulation exists on a server, and is accessed by connecting a client to the server on which the simulation exists.

New simulations (with their corresponding servers) can be created interactively from the client workspace, or can be created using the command line. Similarly, existing simulations can be loaded interactively or from the command line.

The server that is used to host new or existing simulations can be launched as a serial or parallel process.

Creating a New Simulation Interactively

With the client workspace open, create a simulation with any of the following methods:

- Selecting **File > New Simulation...** from the menus



- Clicking the new simulation button on the toolbar



- Pressing the `<Ctrl><N>` keyboard shortcut

Each of these triggers opens the *Create a New Simulation* dialog that allows you start a new simulation.



Further information on these options is provided in the section Understanding Interactive Startup Options. When you have made any necessary changes, click **OK** to launch the new simulation.

It is only necessary to activate one of the boxes in the *License* box if you have purchased one of the following license types:

- Lite

- Power Session
- Power-On-Demand

The progress of creating the simulation is displayed in the *Output* window at the bottom of the workspace. Once complete, a new simulation tree is displayed in the *Explorer* pane showing all the default objects of an empty simulation. The simulation is given a default name (`Star 1` for example), though it can be renamed.



Loading an Existing Simulation Interactively

There are a number of ways to open an existing simulation interactively:

- Using the Client Workspace
- Using the Recent Files List
- Double-clicking in Explorer (Windows only)

Using the Client Workspace

If the client is already running and you want to open a simulation, you have a number of options:

- Select the **File > Load Simulation...** menu option



- Click the Load Simulation toolbar icon



- Use the `<Ctrl><L>` keyboard shortcut.

Each of these actions opens a dialog for loading an existing simulation.



The **Browse** button on this dialog opens an operating system-specific *Open* dialog with the File Type filter set to `*.sim`. You could change the filter (to `*.*` for example) but opening a file without the `.sim` extension has no effect, even if the file is a simulation database. In the *Open* dialog, navigate to the simulation file you want to open, and click **Open** to select the file.

Once you select a file, pressing the **Query** button generates information about the file in the **File Info** text box. This information includes non-zero values that are saved in the file for Step, Time, and Iteration.



This information can be helpful in choosing server options that match the file configuration in order to avoid Partitioning.

The other options on the dialog are described in section, Understanding Interactive Startup Options. Only activate one of the boxes in the *License* box if you have purchased one of the following license types:

- Lite
- Power Session
- Power-On-Demand

Opening Using the Recent Files List

STAR-CCM+ keeps a record of simulation files opened in the past and makes these files available in a most recently used list. To reopen a recently used file click the **File** > **Recent Files** menu item to expose the list, and click the file of interest.



This action is always activated.

Until you exit STAR-CCM+, the client remembers the process option that you used for the prior simulation, whether it was new or pre-existing.

Opening by Double-Clicking in Explorer

On Windows, the installation program makes a file association with the `.sim` file extension. So, simply navigate to the directory where the simulation file is using your file explorer, then DoubleClick the file.



This action launches the splash screen before opening the client and the specified simulation file. When double-clicking the file in Windows, the last installed version is launched, opening in serial by default. If you want to launch a version other than the latest one installed, launch the workspace first, and then open the simulation file from there.

The simulation tree has the focus once the loading is complete.

Understanding Interactive Startup Options

STAR-CCM+ presents you with a dialog containing startup options when you create a simulation or load an existing simulation.

These options control how STAR-CCM+ launches the server. This section documents these options. Click the various elements of the dialog images below to jump to the relevant information.

Create a New Simulation Dialog



Load Simulation Dialog



Selecting the Process Options

The *Process Options* group box contains two options for how the STAR-CCM+ server is launched:

- Serial Mode
- Parallel Mode

The serial mode option is explained more fully in the sections below. Information on the parallel mode is summarized below, and explained fully in the section, Working with Parallel Servers.

Starting in Serial Mode

When a simulation is launched in serial mode, the server on which the simulation resides runs as a single process. This server process can be set to run on the local machine, or to run on a remote machine.

Starting on the Local Machine

If the **Remote Server** option is left deactivated when the *Process Options* option is set to Serial, the simulation is loaded into a serial server on the local machine. This means that both the client and the server are running on your local machine. This setting is the default setting.

Starting on a Remote Machine

In the *Process Options* group box of the dialog, activate the **Remote Server** checkbox. In the **Remote Host** text box, enter the hostname of the machine you wish the server to run on. If your account on the remote machine does not use the same user name as the machine on which the workspace was launched, also enter your remote user name in the **Remote User** text box. Any new simulation that is created or opened from file starts a server on the remote host. The **Remote Shell** text box lets you specify the system command that is used to start servers on remote hosts (the default entry is `rsh`).



When launching on a remote server, activate the Enable Collaboration or Use Remote File Browser options if necessary. **Enable Collaboration** must be activated if you are using a remote server with a different user name.

Starting in Parallel Mode

When a simulation is launched in parallel mode, the server runs over multiple cores or processes.

These processes can be set to run on the local machine (provided the hardware is adequate) or can be set to run on one or more remote host machines. In all cases, the interactive client workspace runs on the local machine.

Using Parallel on the Local Machine

By default serial mode is selected, so click the **Parallel on Local Host** radio button for parallel mode.



The number of processes can be entered in the **Compute Processes** box or incremented using the arrows. The equivalent command line which is used to start the server is displayed in the **Command** text box.

Accept the completed dialog by clicking the **OK** button. Now any simulation that you open, runs the server in parallel mode with the specified number of processes on the local host.

Using Parallel on Multiple Host Machines

By default serial mode is selected, so click the **Parallel on Named Hosts** radio button to specify the host machines for parallel mode. This option displays a small spreadsheet where the list of hosts can be entered.



Specify the machines by typing entries into the list. Click in the Host column of the first row of this list and type the first machine name. Press `<Tab>` or `<Enter>` to move to the next row.

In addition to the table format, this dialog contains buttons for modifying the list of hosts. Simply enter the name of the host in the **Add Host** text box and click the **Add** button. The **Delete** button deletes from the list a host whose name you select in the list. STAR-CCM+ stops using this host for parallel processing.

Verify all your hostnames before attempting to launch the simulation. If any of the hostnames is incorrect when you click the **OK** button, STAR-CCM+ generates an error message and no simulation is started.



Specify the number of processes in the Processes column. Add the other machine names in the same fashion until you are done.



Accept the completed dialog by clicking the **OK** button.

Now any simulation that you open, runs the server in parallel mode with processes on the machines specified. The parallel controller process is started on one of the hosts by default.

If you frequently use the same set of machines for your parallel jobs, you can create a machine file for your convenience.

Saving Startup Configurations

The Saved Configurations drop-down list allows you to type in a name and save the current server options by clicking the **Save** button.



You can also use it to select and restore a previously saved server configuration.



Using Alternative License Options

In addition to the normal STAR-CCM+ license, three alternative license options are provided:

1. If you have purchased a Power Session license.
2. If you have purchased a STAR-CCM+ Lite license.
3. If you have purchased a Power-On-Demand license.



If you have not purchased any of these license types, then you do not need to select anything in the *License* options group box. For more information of the different product configurations, refer to the STAR-CCM+ Product Configuration section.

Enabling Collaboration With Other Users

The **Enable Collaboration** checkbox, when activated, allows other users to connect to the server running that simulation.

If a remote server is used with a different user name, then this option must be selected. This option is available for both serial and parallel servers.

Browsing Files Remotely

The **Use Remote File Browser** checkbox, when activated, starts the file browser on a remote machine when you click the **Browse** button.

You would use this option if a remote server is specified and the remote machine does not share a file system with the local machine. However, sometimes it is faster to use the local file browser when the local and remote hosts do share a common file system (for example, NFS or Windows Shared Folders).The path to the simulation file must always be valid on the server host (controller host for parallel).

Using the Microsoft Job Scheduler Option

In Windows, the dialogs for loading and creating simulations include an option for submitting a job to a cluster using Microsoft Job Scheduler.



When this option is activated, a server command is submitted to the cluster using the Microsoft job scheduler. A `-ccs` argument is added to the *Command* property in the dialog to indicate this behavior. This functionality is available only on Windows HPC Server 2008 machines.

Creating and Opening Simulations from the Command Line in Serial Mode

Using the `starccm+` command on its own, without any other options, launches the STAR-CCM+ client workspace.

Options can be added to the `starccm+` command to achieve more than just launching the client workspace.

- Creating a simulation connected to an interactive client
- Creating a simulation without an interactive client
- Loading an existing simulation in an interactive client
- Loading an existing simulation without an interactive client

Creating a New Simulation Connected to an Interactive Client

A single command-line expression can be used to:

- Launch an interactive client on the local machine
- Start a serial server with a new simulation on the local machine
- Connect the interactive client to the serial server

The command-line expression for these tasks is:

```
% starccm+ -new
```

Creating a New Simulation Without an Interactive Client

To create a simulation on a serial server, without launching an interactive client, use the `-server` option:

```
% starccm+ -server
```

The output of this command includes a line that provides the port number of the newly launched server. An example is:

```
% Server::start -host allegro.ourdomain.com:47830
```

To connect a client to this server, you would use the command:

```
% starccm+ -host allegro.ourdomain.com:47830
```

This action would launch a client on the local machine and connect it to the server running on `allegro`.

Loading an Existing Simulation In An Interactive Client

To launch an interactive client on the local machine:

- start a serial server on the local machine
- load an existing simulation into this server
- connect the client to the server

  you can use the following command:

```
% starccm+ <file.sim>
```

where `<file.sim>` is the name of the simulation file you wish to load into the serial server.

Note: On Windows platforms due to a Windows Command Prompt limitation in character representation, ensure that the locale is set to the same language as the .sim filename that is being opened.

Loading an Existing Simulation Without an Interactive Client

To launch a serial server on the local machine and load an existing simulation into it, you can use the following command:

```
% starccm+ -server -load <file.sim>
```

The output from this command ends with a line containing the host and port numbers that are needed when connecting a client.

Starting Remote Servers from the Command Line

The functionality for starting a remote STAR-CCM+ server involves a simple command, around which you add options depending on what needs specified.

- Starting a server without an interactive client
- Starting a server with an interactive client

Starting a Server without a Client

The server can be started on a remote machine using a remote shell such as `rsh` or `ssh` in the same way as any other command. For example, using `ssh`:

```
ssh user@myremotemachine starccm+ -server
```

Make sure that remote shell access is permitted and configured.

To specify that a particular simulation file, in this example `file.sim`, be opened:

```
ssh user@myremotemachine starccm+ -server
<absolute_pathname>/file.sim
```

Make sure that `file.sim` is in a path accessible from `myremotemachine` (the sim file must be readable by the server), and that the appropriate version of STAR-CCM+ is being started on `myremotemachine`.

If you wish to use a version of STAR-CCM+ other than the one to which your `$PATH` environment variable points, provide the absolute pathname:

```
ssh user@myremotemachine <absolute_pathname>/starccm+ -
server <absolute_pathname>/file.sim
```

Using a Client When Starting a Server

To start a local client (in this example Windows) along with a remote server:

```
starccm+.exe -rsh ssh -on user@myremotemachine
<absolute_pathname>/file.sim
```

- The client starts locally, then attempts to start the server remotely. The *Output* window includes something like the following:

```
Starting remote server: ssh user@myremotemachine -n -
o StrictHostKeyChecking=yes
```

```
starccm+ -on user@myremotemachine -server
<absolute_pathname>/file.sim
```

- The `-rsh` command is instructing the STAR-CCM+ client in the above example to use `ssh`.
- The `-on` is only relevant for the following purposes:
  - In parallel—it is used to specify the hosts to run workers on.
  - In serial—it tells the client where to start a serial server.

You can also specify the path where STAR-CCM+ is installed:

```
"c:\Program Files\CD-adapco\STAR-CCM+ <version_number>
\star\bin\starccm+.exe" -rsh ssh -on
user@myremotemachine <absolute_pathname>/file.sim
```

Restrictions on Names of Files and Paths in STAR-CCM+

The characters in the following table have restrictions on their use in the names of files or paths in STAR-CCM+.

Character	Restriction
\	Usable as a path separator character in Windows; not usable as a file name character.
/	Usable as a path separator character in Linux (and generally usable as a path separator in the Windows 32-bit API functions); not usable as a file name character.
*	Not usable as a file name or path name character.
:	Usable as a drive letter separator in Windows; not usable as a file name character.
?	Not usable as a file name or path name character.
"	Not usable as a file name character.
!	Can cause problems if used in a file name in a shell, since this character recalls a previous shell command.
|	Not usable as a file name or path name character.
<	Not usable as a file name or path name character.
>	Not usable as a file name or path name character.

Also, on Windows, the strings listed below will not be recognized by STAR-CCM+ if you use them as entire file names. This is true for upper or lower case. For example:

- `AUX` is not recognized.

- `AUX.txt` is not recognized.

But:

- `AUXILIARY.txt` is recognized.
- `AUXILIARY` is recognized.

The strings are:

```
AUX
CLOCK$
COM1
COM2
COM3
COM4
COM5
COM6
COM7
COM8
COM9
CON
LPT1
LPT2
LPT3
LPT4
LPT5
LPT6
LPT7
LPT8
LPT9
NUL
PRN
```

Using the Simulation Pop-Up Menu

Save	Executes the save action to save the latest changes to disk.	
View Log	Executes the view log action and opens a dialog allowing you to view and edit the simulation log file.	
Summary Report	Makes a summary report of the simulation objects available through one of the following techniques.	
	Browser	Executes the view summary report action and displays the report in a web browser.
	To File...	Executes the save summary report action and generates a copy of the report in `.html` or `.xml` format.
Collapse Tree	Closes all open nodes in the object tree so that only the top-level manager nodes are visible.	

See also: common menu items.

Saving Simulations to File

A simulation file contains the complete set of simulation objects, state, and data for the problem.

It is the restart file for STAR-CCM+ and is stored with the `.sim` extension. Since it is a binary file, applications other than STAR-CCM+ cannot read it. The file can safely be moved, copied, and renamed using operating system commands as it is the only file that is associated with a simulation.

This section covers the topics that are relevant to saving simulations in a simulation file:

Saving Over an Existing Simulation File

Once you have changed a simulation, you can save it to a simulation file.

This action writes changes to the simulation file on the hard drive where the simulation resides. If the simulation being saved is new (has not been saved yet), then the save as action is executed. This action is only available when a simulation tree is active.

There are a number of ways to save a file:

- Selecting the **File** > **Save** menu option



- Clicking the **Save** toolbar icon



- Using the `<Ctrl><S>` keyboard shortcut

These methods are only available if you have changed the simulation since it was opened. If no changes were made, the menu option and the toolbar are grayed and the keyboard shortcut has no effect.

Saving a Copy of a Simulation

The **File > Save As** menu option allows you to save the simulation into a new simulation file.



Selecting this option opens an operating system standard dialog to locate and name files to save.



The default directory is the last directory that was accessed. If none have been accessed, STAR-CCM+ uses the directory where the current simulation file is located, if applicable. Otherwise, STAR-CCM+ points to your `home` directory on Linux, or a comparable default directory on Windows.

This file save dialog only lists `*.sim` files by default. If you enter a new FileName, you do not have to enter the `.sim` FileName extension since it is automatically appended for you. It is still automatically appended even if you have already added a file extension to the name. Even if you select `All Files` in the Files of Type drop-down list, it is still appended.



- Entering `fred` saves the file as `fred.sim`
- Entering `fred.dat` saves the file as `fred.dat.sim`
- Entering `fred.sim` saves the file as `fred.sim`

Unless you are saving a newly created simulation, the original simulation file remains in the original location, so that you are effectively making a copy of the original file.

Saving Multiple Simulations

Once you have changed one or more open simulation files, you can save changes to them with a single action.

To save multiple simulations:

- Selecting the **File** > **Save All** menu option



- Clicking the **Save All** toolbar button



These methods are only available if you have changed any simulation since it was opened. If no changes were made, the menu option and the toolbar button are grayed.

Renaming the Simulation

A simulation file can be renamed using the inline editor, except within the STAR-CCM+ client this technique works for simulation files only through slow clicking, and only within the *Open* dialog.



Do not, however, change the `.sim` extension as STAR-CCM+ has an association with files of that extension:

- The file open dialog only lists `*.sim` files by default.
- Opening a file that does not have a `.sim` extension does not open the file and display the simulation objects, even if the file is a simulation file.
- There is a Windows desktop association with the `.sim` extension that allows you to DoubleClick the file to open it in the STAR-CCM+ client.

Using Backup Files

STAR-CCM+ generates a set of files that can come in handy as backups of your work.

In your standard file-management software, along with your simulation files, sometimes there are more files in the same location with similar file names. Some file names only differ by a tilde character (~) added to each extension. For example, the files `intake.sim` and `intake.sim~` present in the same folder.

STAR-CCM+ automatically creates such a file once an existing simulation file is modified and saved; it is always the previous version

of that simulation. Therefore if you wish to undo some changes after saving them to a simulation file, you can do so by opening the copy and saving it over the original simulation file.

The STAR-CCM+ client does not always recognize this type of file, which prevents accidental opening and modification. However, it is easy to make the file accessible. Begin by selecting it in your file-management software.



Rename it in order to remove the tilde and give it a name different from the original simulation file.



With the tilde removed, the node of the file appears in the *Open* dialog, which is activated by the **File > Load Simulation...** menu item.



This means that the file is now readily accessible to STAR-CCM+.

If you do not have write-access privileges for a simulation file, STAR-CCM+ does not only fail to save your changes, but also fails to

create or modify the copy. That fact is reflected in an exception message like this one.



What Happens If a Disk Is Full?

If the disk becomes full while you save a simulation file, STAR-CCM+ notifies you that the file could not be saved due to insufficient disk space.

To save the simulation file, do either of the following:

- Save to a different volume.
- Remove data from the target volume and try to save again.

Creating a Template Simulation File

You can create a simulation file as a template for multiple cases. STAR-CCM+ lets you set up many elements of the simulation before generating a mesh, importing a mesh, or running the simulation.

These elements include:

- Setting up a mesh continuum
- Setting up a physics continuum
- Adding reports
- Adding scenes, monitor plots, and XY plots
- Creating regions and boundaries (when appropriate)

A template can be helpful when, for example, you want to use the same set of continuum models to compare meshes.

Regions and boundaries are appropriate for a template only when you are working with multiple cases that have the same region and boundary names. For example, the optimization of a pipe, in which the geometry stays the same while changes are made in the dimensions or the locations of bends.

Closing the Simulation

When you are done with a simulation, you can close it to free up the machine resources it has been using.

You can close the simulation with any of the following methods:

- Right-clicking the simulation window title and selecting **Close** from the pop-up menu;



- Clicking the close icon on the simulation window title



- Pressing the `<Ctrl><F4>` keyboard shortcut

Each of these options executes the close window action. If anything in the simulation has been changed since opening or creating it, you are prompted to save the simulation:



Closing the Client Workspace

When you are done with simulations, close the client workspace to regain resources that it was using:

- Select **File** > **Exit** from the menus



- Press the `<Alt><F4>` keyboard shortcut

The exit application action is executed, which warns you with the following dialog when there are simulations still open and require saving.



Select an entry in the list and click **Save** to save a particular file and remove it from the list. To discard some files, save the files that are needed one at a time, then click **Discard All** when the list has been reduced to the files that you can discard.

Viewing and Editing the Log File

This section defines what is meant by a log file, and describes how to view a log file and add comments to one.

What Is a Log File?

Whenever you save a simulation, an entry is added to a simulation log. The simulation log is a record of each save operation through the life of the simulation, which specifically records the following for every save operation:

- Date
- Time
- Version number
- Architecture that was used

The information is stored in the simulation file itself, and not in a separate file. The log could be useful if you are checking compatibility issues between architectures, or between greatly disparate version numbers, for example.

Viewing the Log

You can view a log file by selecting the **View Log** menu item from the pop-up menu on the root node of the simulation tree:



The view log action is executed, and displays the contents of the log in a separate window.



Adding Comments

You can also intersperse a log with comments of your own. If you want to indicate, for example, that you tried a particular CFL number, switched models, or imported a certain file type. Clearly, you do not get a record of the results at each of the log points (although you could add key results in the log yourself). However, these entries could be useful in the future when trying to remember the options that you have already tried. To add an entry to the log file, view the log and then click the **Add Entry...** button on the *View Log* dialog.



A simple text editor opens, allowing you to enter some text that is appended to the log.



Once you have accepted the editor dialog by clicking **OK**, you cannot edit or delete an added entry.

Using the Summary Report

This section defines what is meant by a summary report, and describes how to view it or save it to a file.

The following sections provide more detail on the parts of the report:

- System information
- Simulation property settings
- Solution details

What Is a Summary Report?

A summary report is a customized document that is generated by STAR-CCM+ in your web browser when you execute the view summary report action. It contains a comprehensive range of basic information about the simulation, allowing you to view this information all in one place rather than examine the properties of individual objects in the object tree. You can also save and/or print this report just like any other online document.

Viewing the Summary Report

You can view a summary report one of two ways:

- Right-clicking the root node of the simulation tree and selecting **Summary Report** > **Browser** from the pop-up menu



- Selecting **Summary Report** > **Browser** from the **File** menu



This action executes the view summary report action, and displays the report in your web browser.



If your browser fails to respond, this can be resolved by changing the browser path that is used by STAR-CCM+.

System Information

This section of the report displays the following:

- Date and time the report was generated
- File name and size
- Build architecture
- Software version



Simulation Property Settings

This section of the report lists all property settings for all objects of the simulation, such as continua, regions, and boundaries. The information is presented in a sequence that corresponds to the object tree in the STAR-CCM+ GUI.



Solution Details

This section of the report lists the following:

- The CPU and elapsed time while iterating
- The total number of iterations run

These figures reflect totals since the simulation was created, or since the last time a clear solution was performed.



Saving the Summary Report to a File

You can save the summary report directly to an `.html` or `.xml` file on your hard drive. There are two ways to save the summary report:

- Right-clicking the root node of the simulation tree and selecting **Summary Report > To File...** from the pop-up menu



- Selecting **Summary Report > To File...** from the **File** menu



This action opens a typical *Save* dialog in which you can navigate to the location for saving the file.

Using the STAR-CCM+ Workspace

This chapter shows you how to use the various components of STAR-CCM+, with a focus on the user interface.

Alternatively, if you prefer to run applications with line commands, see the Command Line Reference.

Note: The STAR-CCM+ workspace lets you display with unicode characters, change the scene highlight font, and change the workspace font size.

Specific topics addressed here are:

Using the Workspace

If you have used any GUI-based applications (such as Microsoft Word), you will have little difficulty using the STAR-CCM+ workspace.

It includes many widgets and operations that have become commonplace, such as tree views, toolbars, menu bars, tabs, docking windows, drag-and-drop and pop-up menus. However, if you are unfamiliar with such tools, the following sections provide information on:

Moving and Resizing the Workspace

The STAR-CCM+ workspace is typical of the kind of window that you use in your Windows or Linux GUI.

You can move the title bar or resize the sizing border or sizing grip as shown in the illustration below.



Sizing grips exist in both Linux and Windows interfaces, even though the ones in Linux are clearly visible.

More instructions are provided about resizing panes and working with windows.

See Also

Using the Title Bar
Resizing Panes
Working with the Workspace Windows

Resizing Panes

The STAR-CCM+ workspace consists of four panes that together occupy the entire area below the menus and toolbars.



These panes (or a subset of the panes) always occupy the entire work area, even if you only have one pane open. The panes do not float in the work area or overlap each other. They are separated by splitters that you can drag to reallocate space from one pane to another.



See Also

Moving and Resizing the Workspace
Dragging the Splitters

Dragging the Splitters

The narrow area between two panes is called a splitter bar.

Hover the mouse over one of these bars until it changes into a directional arrow.



Hold down the left mouse button and drag the mouse up or down (or left and right) to move the bar and change the size of the panes on each side of the splitter bar.

The mouse icon looks slightly different on Windows systems, but the functionality is the same.



See Also

Resizing Panes

Working with the Workspace Windows

Windows are containers for controls.

In the following screenshot, there are four windows:

- Servers (when open; closed by default)
- Properties
- Output
- A simulation window (*Star 1* in this example) when a simulation is open



There are two ways that windows can reside in panes:

- They can be the only occupant of a pane, like *Properties* and *Output* above.
- They can share a pane with other windows, like *Servers* and *Star 1* above. In this case, the contents of only one window are visible at one time and the title bars of the windows are arranged like tabs at the top of the pane.

You can do the following with these windows:

See Also

Moving and Resizing the Workspace

Maximizing and Restoring Windows

You can have a window in the *Explorer* pane occupy the entire workspace (apart from the space that is required by the menu bar and toolbar).

This action is called maximizing the window, and is shown in the following image.



Maximize a window using one of the following techniques:

- Double-click the title bar
- Press `<Shift>+<Esc>` when the relevant window has the focus
- Select the **Window > Maximize Window** from the menu system when the relevant window has the focus



You can restore a window to its original position by repeating any of the above steps, except for the menu item, which would be **Window > Restore Window** .

Moving Windows in a Pane

You can reposition windows within the *Explorer* pane with a drag-and-drop mouse movement.

To move a window:

1. Click the tab of the window that you wish to move. As you begin to drag the tab, the window is outlined in red.



2. Drag the window toward the desired position until you see the red outline jump to that position.



3. Release the mouse button. The window and its tab are in the new location.



Detaching Windows from a Pane

You can pull windows out of the *Explorer* pane with a drag-and-drop mouse movement.

To detach a window:

1. Click the tab of the window that you wish to detach. As you begin to drag the tab, the window is outlined in red.



2. Continue to drag the tab until your mouse pointer is over the splitter to the right of the pane. If you want the detached window alongside the rest of the pane, drag to the middle of the splitter. If you want it above, drag near the top.



3. Release the mouse button. The window appears next to the rest of the pane.



Reattaching Windows to a Pane

You can put windows back into the *Explorer* pane with a drag-and-drop mouse movement.

To reattach a window:

1. Click the tab of the window that you wish to move. As you begin to drag the tab, the window is outlined in red.



2. Continue to drag the tab until your mouse pointer is well within the active window of the destination pane.



3. Release the mouse button. The window is part of the pane.



Closing and Opening Windows

There are various ways to close windows.

It is possible to reopen them using the Window menu. If you are closing a simulation window, STAR-CCM+ will prompt you to save the simulation if you have modified it since opening it.

Using the Mouse

Click the cross at the right edge of the tab. Take care not to close a tab accidentally when you intend to bring it to the foreground.



Using the Keyboard

Press `<Ctrl><F4>` when the tab is visible and the window holding the tab has the focus.

Scrolling

As with typical windows, the ones in the STAR-CCM+ workspace can be navigated by dragging the scroll bar within the scroll box.



Using Auto Hide

The Auto Hide feature lets you minimize windows within the workspace.

To minimize a window, click the **Auto Hide** button on its title bar. This button is marked with an arrowhead and a rectangle.



The window is minimized in a new bar that appears in the direction where the arrowhead was pointing.



To view the window temporarily, hover with the mouse pointer over the button with the name of the window in the bar. The window reappears until you move the mouse pointer over an object other than the bar or the window.



If you click inside that window or its button on the bar, giving it the focus, it remains open until you select another window in the workspace, or click the button on the bar again.

To restore the window completely, click the button in its title bar with the diamond shape.



Alternatively, you can restore the window by right-clicking its button in the bar and selecting **Minimize Window**.



Window Pop-Up Menu

Access this menu by right-clicking the title bar of the window.

Close	Closes the window. For a simulation, this action terminates the STAR-CCM+ client connection to the simulation, along with the corresponding STAR-CCM+ server process (represented by a node in the *Servers* window).
Maximize	Maximizes the window within the STAR-CCM+ workspace; can be reversed with the **Restore** menu item.
Restore	Restores the extent of the window before the use of the **Maximize** command.
Auto Hide	Minimizes the window within the workspace.

Preserving the Window Layout

Window layouts are used to save and restore an arrangement of windows in the graphics display.



Window layouts are explained in the following definitions and processes:

What Are Window Layouts and Modes?

A window layout is an arrangement of windows in the graphics display, which consists of one or more modes. They can be used to save, restore, and edit a specified visual display.

A layout mode is an individual section of the graphics display, which consists of scenes and plots. Modes can be edited, to add, or remove scenes and plots from the graphics display. As an example, the following graphics display has four modes, with one of them highlighted.



Modes can also be undocked from the graphics display. You can save the position on the screen with the layout.

Saving a Window Layout

Saving a window layout makes the current graphics display available for restore. When saved, a window layout can be edited to meet visualization requirements.

Layouts can be:

- Saved as new
- Saved as the default layout
- Saved when the simulation is saved

Save Current Layout

To save the current window layout as a new layout, right-click the **Layouts** manager, and select **Save Current Layout**.



A layout is added to the **Layouts** manager.

Save Layout as Default

To save the current window layout as the default layout, right-click the **Layouts** manager, and select **Save Layout As Default**.



Alternatively, from the menu bar, select **Window > Save Layout As Default**.



This action overwrites the default layout under the **Layouts** manager.

Save Layout on Simulation Save

To save the layout as the default layout when the simulation is saved, from the menu bar, select **Tool > Options**. Then in the *Options* dialog, select the **Visualization** node, and activate the **Save layout on sim save** option.



Restoring a Window Layout

Restoring a window layout restores a saved layout in the graphics display.

To restore a layout, right-click a layout from within the **Layouts** manager, and select **Restore Layout**.



The saved window layout is shown in the Graphics window.

Alternatively, to restore only the default layout, from the menu bar, select **Window** > **Restore Default Layout**.



Editing a Window Layout

Once saved, window layouts can be edited to meet visualization requirements. Editing consists of adding, replacing, or removing scenes and plots from the layout modes.

Adding a Scene or Plot to Layout Modes

To add a scene or plot to a layout mode, select a scene or plot and drag the item to the **Mode** manager.



A node representing the scene or plot is then added under the **Modes** manager. The graphics display is updated when the scene is restored.

Replacing a Scene or Plot in Layout Modes

To replace a scene or plot in a layout mode, select a scene or plot and drag the item to the node that you want to replace.



A node representing the scene or plot then replaces the original node. The graphics display is updated when the scene is restored.

Removing a Scene or Plot from Layout Modes

To remove a scene or plot from a layout mode, right-click the node representing the scene that you want to remove, and select **Remove from Layout**.



The node representing the scene or plot is then removed from the **Modes** manager. The graphics display is updated when the scene is restored.

Executing Actions

Actions are those parts of the client that actually carry out a piece of work, such as scaling a mesh. Actions can be triggered through several workspace features, including menus, toolbars, pop-up menus, and keystrokes.

This section describes the following:

Using One of Multiple Methods to Trigger an Action

The same action can be triggered multiple ways.

For example, you can create a new simulation using one of the following methods:

- Selecting **File > New Simulation...** from the menus



- Clicking the **New Simulation** button on the toolbar



- Pressing the `<Ctrl><N>` keyboard shortcut

How STAR-CCM+ Executes Actions

Actions vary in their implementation depending on complexity. Some of them involve simple functions such as creating a mesh continuum.



In such a case the action is completed as soon as it is triggered.



Other actions may require more input; therefore they include dialogs that are modal--they must be closed before you interact with any other elements of the workspace, such as the object tree. One such action is **File > Import > Import Volume Mesh**.



In order for the action to work, you must select a file to import.



Still other actions, such as creating derived parts, use modeless dialogs.



A modeless dialog lets you interact with other parts of the workspace. This functionality is helpful when you need to check certain information before deciding what to specify in the dialog.



Common Actions Reference

This page contains links to explanations of the most common actions among the various features of STAR-CCM+.

Choose your page of interest from the list below.

- Close window
- Connect to server
- Create or Add
- Copy
- Cut
- Define dimensionality
- Define property with units
- Delete
- Edit
- Maximize window
- Open plot
- Open scene
- Paste
- Restore window
- Rename
- Save simulation

- Sort

Opening Files

This action activates the desktop-specific (for example, Windows XP versus Linux KDE) *Open* dialog.



Look In:

A drop-down list that corresponds to your directory hierarchy. From here, you can move rapidly to a higher directory level. The default directory is the last directory that you accessed, if it exists. If that directory does not exist anymore, STAR-CCM+ uses the directory where the current simulation file is located, if applicable. Otherwise, STAR-CCM+ points to your `home` directory on Linux, or a comparable default directory on Windows.

(Up One Level)

Points the dialog to the next higher directory level. This option is deactivated when the dialog points to the root directory.

(Home)

Points the dialog to your `home` directory (Linux).

(Create New Folder)

Lets you create a new folder directly in this dialog rather than going through your system's standard navigation interface.

(List)

Displays the objects in the current directory with their icons and names only, in alphabetical order.

(Details)

Displays the objects in the current directory with names and modification dates; for individual files, file size is also displayed. The labels at the top allow you to reposition the information columns.

File Name:

This text box shows the currently selected file name. You can type in the name of the file you wish to open. It also allows you to search for a particular file or file type using wildcard characters. For example, if you wish to locate all text files (.txt) in the current directory, use a wildcard ("*") as follows: enter `*.txt` and click `Open`. Only text files, if any, will be visible in the navigation pane along with the subdirectories.

Files of Type:

The default setting of this filter is for `*.sim` files. The `All Files` filter lets you look at other file types. Do not attempt to open files that do not have the `.sim` extension as STAR-CCM+ will not recognize the contents even if the file is a simulation database.

Working with Dialogs in STAR-CCM+

Dialogs are a medium for you and STAR-CCM+ to interact.

For information about the object selector, see The Object Selection Dialog.

This section introduces you to:

Components of Dialogs

These illustrations provide the names of various components in typical STAR-CCM+ dialogs.

They can be helpful for understanding instructions on the use of the code.



One useful component of dialogs is the inclusion of sizing borders and sizing grips, which are identical to the sizing borders and sizing grips used for resizing the workspace.


Create Section
Input Parts:
[Fluid]
Plane Parameters
insertion point
origin
normal
X
-0.5334
1
Y
1.268259
0
Z
0
0
group box
Plane Tool
mouse pointer
Snap To Part
Fluid
checkbox
Display
tick
Displayer
New Geometry Displayer
New Scalar Displayer
radio button
New Vector Displayer
Existing Displayer
Outline
Geometry 1
drop-down list
Create
Close



The Options Dialog

The *Options* dialog is used to control a number of global variables in STAR-CCM+.

To access this dialog, select **Tools > Options** in the menu bar.

The *Options* dialog appears.



The sections available are:

- **Environment** - This page contains the options that control the user environment.
- **Parallel**- This page contains the options that control the settings for parallel computation in STAR-CCM+.
- **Visualization** - This page contains the options that control how the model is displayed in the *Graphics* window, and how you interact with it.
- **Licensing** - This page contains the add-on licensing options. From here, you can reserve the desired add-on licenses.
- **SSH Tunnel** - This page contains the SSH tunnel options. From here, you can pre-configure the SSH options so that they are automatically applied each time that you want to use SSH tunneling with STAR-CCM+.

See Also

Changing the Appearance of the Scene Highlight Font

Message Dialogs

Message dialogs are the method by which STAR-CCM+ provides you with urgently needed information.



Defining Colors

The dialog for color control is used for a wide variety of objects within STAR-CCM+.

The most common method for activating this dialog is clicking the property customizer ( ) in the color property of an object. The following example controls the color of symbols in the plot of a data set.



The *Swatches* tab provides a set of colors to choose from. As you hover with the pointer over the color squares, the RGB (Red, Green, and Blue) values appear in a tooltip. The smaller swatch that is labeled Recent shows colors that you selected. In the *Preview* group box, you can see how the color that you choose would appear in various forms before you confirm the change.

The following tabs provide you with the tools to customize your color selections with the precision of a graphics tool. They can be useful to generate output for a professional-quality printed report.

- *HSV* (Hue, Saturation, and Value) -- works like mixtures of paint
- *HSL* (Hue, Saturation, and Lightness) -- resembles perceptual color combinations
- *RGB* (Red, Green, and Blue) -- the color model commonly used in video displays
- *CMYK* (Cyan, Magenta, Yellow, and Black) -- the color model commonly used in printing

The *Palette* tab lets you choose from the colors that are used to design the STAR-CCM+ GUI.



Colors are added to the custom color history palette, below the standard color palette, when you select a color that is not in the palette. This history contains up to 14 colors and is saved during the session.

To access the color selection of the *Palette* tab directly in the *Properties* window, click in the right column of the color property.



Identifying the Version of STAR-CCM+

To find out which version of STAR-CCM+ you are running, select **Help > About**.

This action opens an information dialog with details about the code version.



You can click the **View Log** button to display an error log file in your default web browser. The information in this log can be useful for support purposes.

Also, you can restore the default settings by clicking the **Reset User Settings** button.

Working with Unicode Characters

On systems that support unicode, named objects on the tree as well as scene annotations can include non-Western (UCS-16) character sets as shown below.



Changing the Appearance of the Scene Highlight Font

You can control the appearance of the font that is used to highlight objects in scene displays.

To change the font that appears in a scene display:

1. Select **Tools > Options**.
2. In the *Options* dialog, select the **STAR-CCM+ > Visualization** node.

3. To set the *Highlight Caption Font* property, do either of the following:
   - Enter the name of the font that you require.
   - Click the property customizer button ( ). In the dialog that appears, select the font that you require.

See Also

The Options Dialog

Changing the Workspace Font Size

In Linux or Windows you can adjust the size of the font that is used in the STAR-CCM+ client workspace, such as dialogs and the object tree, for that session.

A typical startup of the client activates the workspace in its standard font size.



Using the `-fontsize` option activates the workspace with the specified font size. The example in the following screenshot shows the workspace with a 16-point font size setting.



Linux

The line command for specifying the font size of the workspace is as follows:

```
% starccm+ -fontsize X
```

where `X` equals the point size of the fonts. For example, `-fontsize 16` would specify a 16-point font size.

Windows

The line command for specifying the font size of the workspace is as follows:

```
% starccm+.exe "-fontsize X"
```

where `X` equals the point size of the fonts. For example, `"-fontsize 16"` would specify a 16-point font size.

If you typically work with Windows shortcuts, you can create a copy of your STAR-CCM+ shortcut and modify the contents of the *Target* text box in the shortcut properties with the `-fontsize` option:



(This example may differ from the syntax for your shortcut, depending on your installation of Windows.)

Resetting Factory Defaults

There may be occasions when changes to the STAR-CCM+ workspace are no longer needed, such as when a project ends or when the client computer passes to a different user. There are multiple methods for restoring the default settings.

In the **Tools** menu, select **Reset User Settings...**



The *Reset User Settings* dialog appears. From here, you can choose to revert changes to log files, the appearance of windows, and other features.



Select what to reset and click **OK**. You can also activate this dialog by clicking a button with the same name in the About STAR-CCM+ dialog.

Alternatively, the full range of items can be reset by using the `-reset` line command as described in the Command-Line Reference.

Finally, this can also be done by deleting the user settings directory in either Linux or Windows. You may need to take this step in case you attempt to launch STAR-CCM+ with an unsupported version of Java that is installed on your computer. This attempt can corrupt user settings files, causing the STAR-CCM+ workspace to appear incomplete.

Linux

Exit the program, delete the hidden folder `.star-*-*-*` (for example, `.star-2.06.008` for STAR-CCM+ Version 2.06.008) on your `home` directory, and then restart the program.

To find the complete path to this directory, select **Help > About...** and look at the User Settings Directory: line in the *About STAR-CCM+* dialog.

Windows

Delete the folder `STAR-CCM+ *-*-*` (for example, `STAR-CCM+ 2.06.008` for STAR-CCM+ Version 2.06.008) in `[Profiles]\Local Settings \Application Data\CD-adapco`.

The `[Profiles]` folder may be:

- `C:\Winnt\Profiles\[login name]`; or
- `C:\Documents and Settings\[login name]`

depending upon your installation of Windows.

To find the complete path to this directory, select **Help > About...** and look at the User Settings Directory: line in the *About STAR-CCM+* dialog.

The Object Selection Dialog

Certain options in STAR-CCM+ allow selection of objects from an object selection dialog.



Using this dialog includes the following operations:

See Also

Using the Part Selector

Activating the Object Selection Dialog

For the specific method to open the object selection dialog, see the instructions of the particular STAR-CCM+ feature that you are using.

In general, to open the dialog:

- Click the right half of a *Parts* property (or a similar property by another name), or its property customizer button ( ).



- Click a specific button in a dialog, such as **Select** for *Input Parts*.



Accessing Objects in the Dialog

Use the Tree view or the List view. In either view, you can use the `<Ctrl><Click>` technique to highlight objects without selecting them.

The Tree View

By default, the objects in the dialog are arranged in a hierarchical node tree similar to what you see in a simulation window.

- To open the tree view of an object and all its children at once, select **Expand** from the pop-up menu of the object.



- To open objects selectively, use the handle ( ).



- To identify where in the node hierarchy an object is, hover over it with the mouse pointer in the view panel. The path field, a read-only text box near the bottom of the dialog, shows the location.



The text that is directly above the tree window displays the portion of the path that all objects in the tree share in common. In the following example, only regions are displayed, so the path **Regions** appears above the tree.



This common portion of the path is not duplicated in the path field beneath the tree. Only **Default_Solid** appears in the path field, and not **Regions** > **Default_Solid**.

The List View

If you know the name of the object that you want to select, switch to the **List View** by clicking its button.

- 

- With the objects in an alphabetical list, scroll until you find the desired object. If the list is long, filter the displayed objects.



- To return to the tree view, click the **Tree View** button.



Highlighting Objects

You can highlight one or more objects in the view panel without selecting them.

To highlight one object or a non-contiguous group of objects:

1. Click in the row of each object while holding down the `<Ctrl>` key.



To highlight a contiguous group of objects:

2. Click the first object in the group while holding down the `<Shift>` key or the `<Ctrl>` key.
3. Click the last object in the group while holding down the `<Shift>` key.



To remove the highlighting, use one of the following techniques:

4.
- Click anywhere in the view panel outside the highlighted group of objects.

  If you click an object in this situation, the selection state of the object does not change.
- Right-click the highlighted group and select **Remove Highlight** from the pop-up menu.



5. To select the highlighted items, click any one of them.



Selecting Objects

You can select objects in an object selector dialog using a variety of techniques, depending on the hierarchy of the object that you want to select.

- To select an object, click anywhere in the row of that object.



- To deselect the object, click in the row of the object again.
- To select or deselect all the children of manager objects that have no solution or mesh data (such as a **Boundaries** manager), click anywhere in the row of the manager object.



Note: A different procedure exists when parent objects, such as regions, are available to select.Selection transitions that occur in manager or parent objects can only affect visible children. Children that are currently filtered out of the view are not affected.

A counter at the bottom left of the object selector dialog indicates the number of objects selected. It only appears when your mouse pointer is not hovering over an object or highlighting it.



- To select all objects in the dialog, click the **Column Selection** button in the upper left portion of the dialog. This global selection also applies to all objects that are contained within objects.



- To deselect all selections, click the **Column Selection** button a second time.
- To apply your selections, click **OK**.
- To close the dialog without applying the selections, click **Cancel**.

Note: When you choose multiple nodes before accessing the object selector dialog on a common property of the nodes, the dialog displays the intersection of any objects that are available or already selected. For example, if you choose multiple table nodes and open the object selector dialog to set the *Parts* property, STAR-CCM+ displays the intersection of any objects that were already selected for the individual nodes.

Selecting Certain Parent and Child Objects (Tree View Only)

Parent objects that can contain solution or mesh data (such as a region) are available for selection, and there are four different states of selection possible.

To choose one of the selection states, right-click the parent object. The current state of the object is unavailable (grayed out).



Alternatively, you can click the row of the parent object. Each successive click of that row executes one of the following actions, in a cyclical sequence that is outlined below.

1. Click once to select the parent only (**Turbo Blade Row 1** in this example), leaving the selection status of the children unchanged.



2. Click a second time to select all children in addition to the parent.



3. Click a third time to deselect the parent (leaving only the children selected).



4. Click a fourth time to deselect all children as well as the parent.



5. On the fifth click, the sequence repeats to select the parent only.

Filtering Displayed Objects

To view only the objects that are selected or unselected, use the Filter menu. To filter specific objects by path, enter text into the **Filter by Path** text box near the top of the dialog. To see results without the need to open parent objects, use the List view.

You can use either plain text, regular expressions, or query filters.

Filtering Views

To view only the objects that are selected or unselected, use the menu in the upper left part of the dialog.

To view results without opening parent objects, use the List view. By default, the menu is set to **Show All**.



Note: In the Tree view, if an item with children matches the filter, the item and all related children are displayed regardless of whether they themselves match the criteria. In the List view, only items that match the criteria are shown.

- To view only the objects that are selected, use the **Selected** option.



- To view only the unselected objects, use the **Unselected** option.



Filtering Using Plain Text

To filter the paths of the objects of interest, enter text that is a part of the paths of the objects as shown in the path field.

Plain text filtering is not case-sensitive, but regular expressions are case-sensitive by default.

The following example is a search for boundary objects using the text `Boundaries`.



Filtering Using Regular Expressions

These brief instructions presume prior familiarity with the syntax of regular expressions. For more details about using this technique with the object selection dialog, see the section Regular Expressions.

- To enter regular expressions, click the **Use Regular Expressions** button to the right of the **Filter by Path** text box.

- To find objects with paths that include the same text strings, use the `.*` metacharacter. The example below shows items that have the word `Clearance`, followed by the `.*` portion (which means any number of characters, possibly no characters), followed by a `1`.



- To find objects using an OR relationship, use the | metacharacter. The example below contains the results for interface boundaries.



- To switch back to searching using plain text, click the **Use Regular Expressions** button.

Filtering Using Filter Objects

If you created filters in your simulation, the object selection dialog detects them and shows an additional button in the top left corner.



- To make these filters available in the dialog, click (**Show/Hide Filters**). A new panel with a list of available filters appears in the left portion of the dialog.



- To apply one of the available filters, activate the checkbox next to its name.

  The filters are applied in addition to any filters that you have already applied to the view.

- To extend the view of the filter name, drag the divider between the filter and view panels.



You can also select a group of objects that are listed in the search results of a filter. See Selecting Nodes through a Filter.

Regular Expressions

In Regular Expressions mode, any text that is entered in the text box is interpreted as a regular expression.

Regular expressions are case-sensitive by default. Objects are filtered on whether their paths, or paths of their parents, contain the regular expression. Regular expressions behave quite differently from simple text fragments, in that they follow a particular set of rules and conventions that must be learned. For example, the period, '.', in a regular expression represents any character, rather than just a period. In the context of a regular expression, the period can be called a metacharacter, as it has a programmatic function. For more information on using regular expressions, consult a suitable textbook. The full syntax permissible within the object selection dialog is detailed in the help page of the Java Pattern class. Some examples are given in the table below.

Example	Description
`mir.*rt`	Search for any object whose name contains text that starts with `mir` and ends with `rt`, and has any number of characters in the middle. The '*' metacharacter indicates zero or more times.
`bar|log`	The '|' metacharacter indicates an OR relationship. This search finds any object whose name contains either `bar` or `log`.

Example	Description
`[0-9]+`	The `[0-9]` represents a single digit in the range 0 to 9, and the '+' metacharacter indicates one or more times. This search finds any object that contains one or more digits in its name.
`[0-9]{2}`	The `{2}` after the digit range indicates two times. This search finds any object that contains at least two consecutive digits in its name. Any number can be used within the `{}` brackets.

Single-Selection Mode

The object selection dialog appears in a special single-selection mode for STAR-CCM+ features that are not designed to have multiple objects selected at one time.



Examples of such features include:

- Defining a battery module configuration
- Selecting parts for an external battery casing
- Magnetic fields -- controls for:
  - The excitation coil reference cross section
  - The local direction field
- Defining topology -- controls for:
  - Assigning parts to regions
  - Assigning part surfaces to boundaries
  - Assigning part curves to feature curves
- Surface repair -- controls for:
  - Editing part surfaces
  - Editing part curves

Single-selection mode works the same as the typical multiple-selection mode, with the following exceptions:

- You can select only one object in the panel at a time.
- Clicking a selection button toggles between selection and deselection of an object.
- If you select an object when another object in the panel is already selected, that other object becomes deselected automatically.
- The selection button will only be active for objects that are available for selection. Therefore the button is not active for manager objects, which cannot be selected.
- The header of the selection column is not a button but an indicator. Its icon appears colored when any object is selected, and clear when there is no selection.

Alternate Formats

Other formats of the object selection dialog, commonly used in earlier versions of STAR-CCM+, appear as follows:





In the inline format immediately above, the buttons between the two lists function as follows:

- Transfers the part that has the focus from the Select From list to the Selected list.

- Transfers the part that has the focus from the Selected list to the Select From list.

- Transfers all parts in the Select From list to the Selected list.

- Transfers all parts in the Selected list to the Select From list.

You can select multiple parts at the same time by holding down the `<Ctrl>` key while selecting them.

Using Filters

A filter uses queries to narrow groups of STAR-CCM+ objects. A query consists of one or more conditional statements, also called "predicates", by which STAR-CCM+ searches among the objects in the simulation.

Each filter is stored in the **Tools** > **Filters** node of the STAR-CCM+ object tree.



Each predicate is a logical statement for searching, and consists of three components: class, operator, and criteria. An example of a predicate is:

```
Name contains temperature
```

where the class is `Name`, the operator is `contains`, and the criteria are `temperature`.

Classes

The available classes of a STAR-CCM+ predicate are:

- Type
- Tag
- Metadata
- Name
- Object
- Filter

Note: The Tag and Metadata classes are for geometry parts only. The Filter class is for filters only.

Operators

The operator of a predicate specifies the relationship between the class and the criteria. The set of available operators varies with the class. For details, see Filter Dialog Reference.

Criteria

Possible criteria include text strings, such as the beginning of a name, as well as the following STAR-CCM+ objects:

- Geometry parts
- Part surfaces, part curves, and feature curves
- Regions, boundaries, and interfaces
- Derived parts, either all of them or individual types
- Reports, monitors, and plots
- Filters

Compound Predicates

Compound predicates are groups of predicates that appear in the query filter dialog as AND and OR:

- AND requires search results to comply with all predicates within it.
- OR requires search results to comply with one or more of the predicates within it.

Creating a Filter

To create a filter:

1. Right-click the **Tools > Filters** node and select **New**.
2. In the *Filter* dialog, select the first predicate beneath AND and complete its associated properties.
3. To finish specifying the predicate, click anywhere in the view outside the components of the predicate, or press `<Enter>`.
4. Use the actions available on the dialog to create further predicates. See Filter Dialog Reference for more information.

In the following example, the filter contains results of this search:

```
(Type is Boundary) AND ((Name contains "Memory Chip")
OR (Name contains "Capacitor"))
```



5. Click **OK**.
6. If necessary, rename the filter to an appropriate mnemonic.
7. To apply the filter in an object selector, see Filtering Using Filter Objects

Working with Existing Filters

To view the query definition of an existing filter:

1. Select the node of the filter.
2. Hover over the right half of the *Query* property with the pointer. The *Query* definition appears in a tooltip.
3. To modify an existing filter, do one of the following:
   - Select the filter node and click the right half of the *Query* property.
   - Right-click the filter node and select **Edit**.

The filter dialog appears with the query parameters of the existing filter.

Selecting Nodes through a Filter

You can use a filter to access a group of nodes in the object tree directly.

To select the nodes in the tree:

1. Right-click the desired filter and choose **Select Results**.

   An object selection dialog appears with all the search results of the filter. This feature gives you the opportunity to omit some of the nodes from your selection in the object tree.
2. Click **OK**.

   The nodes that were included among the search results appear selected in the object tree.

Filter Dialog Reference

The filter dialog contains editors for creating predicates, working with compound predicates, and tools for modifying predicates.

Editors of Predicate Components

Three editors exist for creating or modifying a predicate:

- The class editor lets you select a class from the drop-down list.



- The operator editor becomes available after you select a class.



The set of available operators varies with the class, as indicated in the following table:

	Class					
Operator	Type	Tag	Metadata	Name	Object	Filter
is	X				X	
is not	X				X	
contains		X	X	X		
does not contain		X	X	X		
matches				X		X
does not match				X		X
starts with				X		
ends with				X		
does not start with				X		
does not end with				X		

Note: Operators that exclude items from a search, such as `is not` and `does not match`, are the equivalent of the search term NOT.

- The criteria editor lets you specify what to search for. The criteria editor varies in its form, depending on the class that you select:
  - The Type class uses a drop-down list.
  - The Tag class uses the object selector.
  - The Metadata class uses two text components.
  - The Name class uses one text component. You can search with only part of the name if necessary.
  - The Object class uses the object selector.
  - The Filter class uses the object selector.

Compound Predicate Editors

In a new query, the compound predicate editor appears near the top of the view as an AND with a drop-down list in a new query.

- You can change **AND** to **OR** using a drop-down list. To activate the drop-down list, double-click **AND** or click (**Edit**).
- To add the OR compound predicate, indent the predicate. The OR compound predicate appears.

Query Editing Tools

The following table lists the tools that are available, depending on the situation, at the top of the filter dialog. Many of these tools are also available in a context menu when you right-click in the view.

(**Edit**), also `<Space>`	Accesses a predicate for modification.
(**Add**), also `<Insert>`	Creates a predicate.
(**Delete**), also `<Delete>`	Deletes a predicate.
(**Indent**)	Indents a predicate. Use this to create a subdomain such as OR.
(**Outdent**)	Outdents a predicate.
(**Undo**), also `<Ctrl><Z>`	Reverses the latest change.

(**Redo**), also `<Ctrl><Y>` or `<Ctrl><Shift><Z>`	Repeats a change that was subjected to the Undo command.
(**Cut**), also `<Ctrl><X>`	Removes the selected object and stores it on the clipboard.
(**Copy**), also `<Ctrl><C>`	Stores a copy of the selected object on the clipboard.
(**Paste**), also `<Ctrl><V>`	Places the contents of the clipboard into the active area.
(**Display query results as a tree**)	Lets you view search results in a hierarchical node tree. The tree view resembles what you see in a simulation window.
(**Display query results as a list**)	Lets you view search results in a list (default setting).

Using the Title Bar

The title bar provides information, access, and window control for the workspace.

The title bar:

- identifies the application name and the simulation that has the focus, if at least one simulation is open (Application name);
- serves as a control point for moving the window when it is in a restored state (by dragging with the mouse);
- serves as an access point for commands that apply to the whole workspace (Application icon); and
- has buttons for changing the state (Window buttons).



Also, shortcuts exist for working with the title bar.

See Also

Moving and Resizing the Workspace

The Application Icon

The characteristics of the application icon are common to both Windows and Linux.

Clicking the application icon displays the following menu.

Windows	Linux (KDE)

This menu is also displayed if you right-click anywhere in the title bar apart from on any buttons or icons.

All of these commands are available elsewhere:

- Restore - from the restore button on the right when the application is maximized;
- Move - by dragging the title bar when the application is in the restored state;
- Size - by dragging one of the sides when in the restored state;
- Minimize - from the minimize button on the right of the bar;
- Maximize - from the maximize button on the right of the bar when the application is in the restored state; and
- Close - from the **File > Exit** menu item or the Close button at the right of the bar. This is the default action for this button, and will be activated if you double-click the icon.

The Window Buttons

The title bar includes a standard set of window buttons.

These characteristics are common to both Windows and Linux.

minimize the workspace

maximize the workspace (only shown when the workspace is in the restored state)

restore the workspace (only shown when the workspace is in the maximized state)

close the workspace (client)

Shortcuts

By clicking the mouse on the title bar, you can carry out window operations quickly and easily.

Double-clicking the application icon closes the application.

Right-clicking the title bar (not on an icon or button) displays the application menu.

(Windows) Double-clicking the title bar (not on an icon or button) toggles state between maximized and restored.

(Linux KDE) Double-clicking the title bar (not on an icon or button) toggles state between shade (title bar visible only) and restored.

Using the Menus

The menu bar is typical of GUI-based applications. It can be navigated with the mouse or with the keyboard.

The STAR-CCM+ menus provide a function-oriented view of some capabilities that are not bound to objects. They also duplicate a few of the double-click functions where it was deemed to improve user efficiency. The menus do not provide all of the available functionality. Rather, most functionality is accessed by right-clicking, double-clicking, and dragging objects around the workspace.

Using the Mouse

You can use the mouse for navigating the menus in a variety of ways.

The menu bar is activated by clicking a menu button with the mouse.

Once it has been activated, you can navigate around the menu system simply by moving the mouse. Moving horizontally in the menu bar activates different menus.



Moving vertically highlights a menu item.



If a menu item holds a submenu (indicated by the ▸ icon), the submenu is opened as the item is highlighted. The focus remains on the menu system even if you move the mouse away from the menus.



Click a highlighted menu item to execute that action.

Using the Keyboard

Menus are available with the keyboard as well as the mouse.

The keyboard can be used to access the menu bar when you hold down the `<Alt>` key and simultaneously press the underscored character of the menu button of interest. For example, press `<Alt><T>` to pull down the **Tools** menu.



Once the menu system has the focus, use the cursor keys to change menu buttons (`<Left>`, `<Right>`),



or change menu items (`<Up>`, `<Down>`).



Pressing `<Right>` on a menu item that has a submenu opens the submenu.



Press <Enter> on a selected menu item to execute that action.

Many menu items have access keys and quick keys that are associated with them. An access key is shown by the underscored letter in the menu item name, for example the "D" of**Diagnostics**.



With the **Mesh** menu open, pressing <d> on the keyboard has the same effect as navigating to that menu item with the cursor keys and then pressing <Enter>. In this case, the *Mesh Diagnostics* dialog opens.



A quick key (if one is assigned) is displayed to the right of the menu item, such as this one for **Step** in the **Solution** menu.



The quick key can be used to activate that menu item directly, even if the menu system does not have the focus.

Menu Reference

This section provides a reference to all the menu items that appear in the supplied menu system.

The menu bar is always located immediately under the title bar and cannot be moved or hidden. The terms that are used with the menu bar are shown below. The following sections describe each of the menu

items in the workspace and provide links to reference information where necessary.



If a menu item is grayed, it means that function is not available at the moment. Typically the function is not available because the workspace is not in a suitable mode, or the relevant object is not selected.

If a menu item has "..." appended to the name, it indicates that this function opens a dialog.

File

The File menu items apply to an entire server, simulation, or other file.

New Simulation...

Executes the new simulation action. This menu item is always activated.

Load Simulation...

Executes the load simulation action. This menu item is always activated.

Connect to Server...

Executes the connect to server action. This menu item is always activated.

Recent Files

Executes the recent files action. This menu item is always activated.

Save

Executes the save simulation action. This menu item is deactivated when there is no simulation open, when there are no changes to the simulation whose window is active, or when the active simulation is running. It becomes activated when you change the simulation, including opening a scene in the Graphics window.

Save As...

Executes the save as action. This option is deactivated when the simulation is running.

Save All

Executes the save all action. This menu item is activated when changes are made to at least one simulation (including the opening of a scene in the *Graphics* window).

Auto Save...

Executes the auto save action. This menu item is activated when a simulation window is active.

Macro

This menu item contains the following items in a submenu:

- Play Macro... -- Executes the play macro action. This menu item is activated when a simulation window is active.
- Start Recording... -- Executes the start recording action. This menu item is only activated when a simulation window is active.
- Pause Recording -- executes the pause recording action. This menu item is activated when a macro is recording.
- Stop Recording -- Executes the stop recording action. This menu item is activated when a macro is recording.
- Play Recent -- Executes the play recent macro action. This menu item is activated with submenu items once one macro has been run in any simulation.

Import

This menu item contains the following items in a submenu:

- Import Surface Mesh... -- Executes the import surface mesh action. This menu item is activated when a simulation window is active.

- Recent Surface Files -- Provides a list of the most recently imported surface files in a submenu. This menu item is activated with submenu items once one surface file has been imported into any simulation.
- Import Volume Mesh... -- Executes the import volume mesh action. This menu item is activated when a simulation window is active.
- Recent Volume Mesh Files -- Provides a list of the most recently imported volume mesh files in a submenu. This menu item is activated with submenu items once one volume mesh file has been imported into any simulation.
- Import CAE Model... -- Executes the import CAE model action. This menu item is activated when a simulation window is active.
- Recent CAE Files -- Provides a list of the most recently imported CAE files in a submenu. This menu item is activated with submenu items once one CAE file has been imported into any simulation.

Export...

Executes the export solution action, which opens a typical *Save to File* dialog. This menu item is activated when a simulation window is active.

Auto Export...

Executes the auto export action. This menu item is activated when a simulation window is active.

Page Setup...

Executes the page setup action. This menu item is always available.

Print...

Executes the print action. This menu option is deactivated if there is no display active in the *Graphics* window.

Disconnect

Executes the disconnect action. This menu item is activated when a simulation window is active.

Exit

Executes the exit application action. This menu item is always activated.

Edit

The following menu items let you edit objects using the clipboard of the operating system, and to delete items from the simulation.

Cut

Executes the cut action. This menu item is deactivated when the selected object cannot be cut.

Copy

Executes the copy action. This menu item is deactivated when there is no selection of an object that can be copied.

Paste

Executes the paste action. This menu item is deactivated when the clipboard is empty, or when there is no selection of an object into which the material on the clipboard can be pasted.

Delete

Executes the delete action. This menu item is deactivated when the selected object cannot be deleted, such as manager nodes of a simulation tree.

Mesh

The Mesh menu allows you to carry out a wide variety of tasks with elements of the mesh, making it unnecessary to adjust the original mesh file.

Generate Surface Mesh

Executes the generate surface mesh action. This menu item is activated when a simulation window is active, if the simulation contains a mesh that is generated in STAR-CCM+, and that no mesh generation process is running.

Generate Volume Mesh

Executes the generate volume mesh action. This menu item is activated when a simulation window is active, if the simulation contains a mesh that is generated in STAR-CCM+, and that no mesh generation process is running.

Clear Generated Meshes

Executes the clear generated meshes action. This menu item is activated when a simulation window is active, if the simulation contains a mesh that is generated in STAR-CCM+, and that no mesh generation process is running.

Scale Mesh...

Executes the scale mesh action. This menu item is always activated when a simulation window is active.

Diagnostics...

Executes the mesh diagnostics action. This menu item is always activated when a simulation window is active.

Convert to 2D...

Executes the convert mesh to 2D action. This menu item is always activated when a simulation window is active.

Solution

The Solution menu provides control over solution data and over running the simulation.

Clear Solution...

Executes the clear solution action. This menu item is always activated when a simulation window is active.

Initialize Solution

Executes the initialize solution action. This menu item is always activated when a simulation window is active.

Run

Executes the run action. This menu item is always activated when a simulation window is active.

Step

Executes the step action. This menu item is always activated when a simulation window is active.

Step Control

Executes the step control action. This menu item is always activated when a simulation window is active.

Tools

The Tools menu contains items for customizing your client workspace.

Reset User Settings

Opens a *Reset User Settings* dialog where you can select items to restore to their default settings.



Tick the items to reset and click **OK**. You are prompted to close STAR-CCM+ to complete the reset.

Add-ons

Launches the Add-on Licensing dialog. This dialog manages the use of more licensing features.

Options

Opens a dialog with a host of advanced settings, which are covered under their respective topics. This menu item is always activated.

Plugins

Launches the Plugins dialog. This dialog manages the use of NetBeans modules in STAR-CCM+, which are plugins that permit the STAR-CCM+ GUI to be extended dynamically.

Window

The Window menu lets you control the appearance and position of windows.

Only three windows are accessible through the **Window** menu: *Properties*, *Servers*, and *Output*. These three windows always appear as the system defaults; others, such as a named simulation window or the *Graphics* window, appear as a result of user actions.

Close Window

Executes the close window action. This menu item is always activated unless all windows have been closed.

Maximize Window

Executes the maximize window action. This menu item is activated when a window in the *Explorer* pane has the focus and is hidden when one of the windows has been maximized.

Restore Window

Executes the restore window action. This menu item is activated when any maximized window has the focus and is hidden if no windows are maximized.

Properties

Executes the view Properties window action. This menu item is always activated.

Servers

Executes the view servers window action. This menu item is always activated.

Output

Executes the view output window action. This menu item is always activated.

Toolbars

Provides a submenu for controlling the display and customization of toolbars.

- Macro -- An option that displays (activated) or hides (deactivated) the macro toolbar. This menu item is always available.
- Solution -- An option that displays (activated) or hides (deactivated) the solution toolbar. This menu item is always available.
- Edit -- An option that displays (activated) or hides (deactivated) the edit toolbar. This menu item is always available.
- Memory -- An option that displays (activated) or hides (deactivated) the memory toolbar. This menu item is always available.
- System -- An option that displays (activated) or hides (deactivated) the system toolbar. This menu item is always available.
- View -- An option that displays (activated) or hides (deactivated) the view toolbar. This menu item is always available.
- Small Toolbar Icons -- An option that displays small (activated) or large (deactivated) versions of the images on the toolbar buttons. This menu item is always available.
- Standard -- A radio button menu item that is only visible when a second toolbar configuration has been saved.
- Customize -- Executes the customize toolbar action. This menu item is always available.
- Save Configuration -- Executes the save toolbar configuration action. This menu item is always available.

Help

The Help menu provides access to the online documentation, the Java API resources, and product information.

Help...

Executes the help action. For Help pages on specific topics, select the item and press <`F1`>. This menu item is always activated.

Java API...

Accesses STAR-CCM+ Java API resources through a web browser. This menu item is always activated.

Steve Portal

Accesses the CD-adapco support site, the Steve Portal, through a web browser. This menu item is always activated.

STAR-CCM+ IdeaStorm

Accesses the CD-adapco innovation forum, STAR-CCM+ IdeaStorm, through a web browser. This menu item is always activated.

About

Executes the about action. This menu item is always activated.

Using the Toolbars

Like the menu bars, the toolbars are typical of GUI-based applications, providing easy access to the most common actions.

You can control the appearance of toolbars with the following methods:

- Moving a toolbar
- Hiding a toolbar
- Displaying a toolbar
- Customizing a toolbar
- Creating a toolbar

This section also includes a toolbar reference.

Moving a Toolbar

For easy access to frequently used toolbars, move them with the mouse.

The toolbars always remain near the top of the workspace, just below the menu bar. They cannot be floated freely, or docked on the sides or bottom. However, they can be moved relative to each other.

To move a toolbar, position the mouse over the dimpled end (left side) of the bar and then drag it to a new position.



Dragging downwards creates a row for the toolbar, dragging left and right repositions it relative to the other bars.

Hiding a Toolbar

It is possible to hide a toolbar using the menu or the right-mouse click.

Using the Menu

Select `Window > Toolbars` to display the following submenu.



All the visible toolbars have the checkbox ticked. Select the toolbar name to hide and click the left mouse button to clear the checkbox and hide that toolbar. The cleared checkbox appears only if you bring up the submenu a second time.

Using the Right-Mouse Click

Right-click over any of the gray space of the toolbar (not directly over a button) to display a pop-up menu identical to the **Tools > Toolbars** submenu.

Displaying a Toolbar

It is possible to display a toolbar using the menu or the right-mouse click.

Using the Menu

Select `Window > Toolbars` to display the following submenu.



A hidden toolbar has the checkbox cleared. Select the toolbar name to show and click the left mouse button to set the checkbox and show that toolbar. The ticked checkbox appears only if you bring up the submenu a second time.

Using the Right-Mouse Click

Right-click over any of the gray space of the toolbar (not directly over a button) to display a pop-up menu identical to the **Window > Toolbars** submenu.

Customizing a Toolbar

You can add, delete, or relocate buttons to toolbars.

In the **Window** menu, select **Toolbars > Customize...**An explorer tree view of the current toolbar opens with definitions, as shown below.



This feature uses a drag-and-drop operation between the dialog and the toolbars themselves in the workspace. In the example below, the **Delete** button is being added to the Edit toolbar, into the position that is shown:



This method allows you to make any of the following changes to toolbars:

- Adding new buttons by dragging them from the dialog to the toolbar
- Deleting existing buttons by dragging them from the toolbar to the dialog
- Relocating existing buttons (by dragging them from the dialog to the new location) either within a toolbar or between toolbars

Creating a Toolbar

You can add a toolbar to your workspace.

Click the **New Toolbar** button at the bottom of the *Customize Toolbars* dialog.



This button opens a *New Toolbar* dialog allowing you to name the new toolbar.



After you click **OK**, customize your new toolbar before you close the *Customize Toolbars* dialog.

Once a new toolbar has been saved, it appears along with other toolbars on the **Tools** > **Toolbars** menu.



It is also possible for a NetBeans plugin to create a toolbar when it is installed. To learn how to create your own modules, contact CD-adapco support.

Toolbar Reference

Toolbar buttons are duplicates of menu item functions, and are provided for easy mouse access to common functions.

All the buttons are grouped onto several toolbars. In the figure below, the toolbars have been arranged so that they lie under each other. Click a toolbar to find out what each button does.



The following subsections describe each of the toolbars and the buttons on them:

- Animation
- Edit
- Macro
- Mesh generation
- Plot
- Selection
- Solution
- System
- Visualization

Animation

The Animation toolbar buttons resemble the controls of a common video player.

(Play/Pause) -- this toggles between playing the scene animation and pausing it.

(Stop) -- Stops the animation in the scene.

(Single Step Forward) -- Moves the paused animation forward one frame.

(Single Step Backward) -- Moves the paused animation backward one frame.

(Write Movie) -- Records the animation onto a video file.

Edit

The Edit toolbar lets you copy and paste certain objects.

(Copy) -- Stores a copy of the selected object on the clipboard.

(Paste) -- Places the contents of the clipboard into the active area.

Macro

The Macro toolbar lets you record and run macros.

(Play Macro) -- Launches an existing macro.

(Start Recording) -- Begins recording a new macro.

(Pause Recording) -- Interrupts the recording of the macro until you click the button again.

(Stop Recording) -- Finishes recording the macro.

Mesh Generation

The Mesh Generation toolbar provides quick access to important meshing operations.

(Import Surface) -- Imports surface data for generating a mesh.

(Launch Surface Repair) -- Opens the *Surface Preparation* dialog for regions or parts.

(Clear Generated Meshes) -- Deletes meshes generated in STAR-CCM+.

(Initialize Meshing) -- Provides the starting surface for the meshing models.

(Generate Surface Mesh) -- Begins the generation of a surface mesh.

(Generate Volume Mesh) -- Begins the generation of a volume mesh.

Plot

The Plot toolbar lets you open plot displays and navigate within them.

(Create/Open Plots)	Opens a drop-down menu that allows you to add a plot or open an existing one.	
	XY Plot	Creates a XY plot and opens a display for it.
	Monitor Plot	Creates a monitor plot and opens a display for it.
	Histogram Plot	Creates a histogram plot and opens a display for it.
	Open	Contains a submenu with names of existing plots.
	Open All Plots	Displays all existing plots in the *Graphics* window.

(Fit) -- Returns the data portion of the plot display to the default zoom and pan position.

(Pan Selected Plot) -- A toggle button that, when selected, pans (horizontally or vertically) the data portion of the plot as the mouse is dragged with the left mouse button depressed. If the **Zoom** button was selected, then it is deselected as you press this button.

(Zoom Selected Plot) -- A toggle button that, when selected, zooms the data portion of the plot approximately one order of magnitude each time the left mouse button is clicked. The click point is centered on the axes, so you could use this point to determine more accurately the x and y values of a particular point. A drag operation begins a rubber-band

zoom that is effected when the button is released. If the **Pan** button was selected, then it is deselected as you press this button.

Selection

The Selection toolbar lets you refine your selection of objects in the scene display.

(Rubberband Select) -- Selects a specific area in the display by clicking two points that define the opposing corners of the area. Press `<Esc>` to cancel.

(Zone Select) -- Lets you draw a zone by clicking and dragging to create a series of connected points. Objects within that zone are selected. Press `<Esc>` to cancel.

(Selection Filter) -- Provides options for selecting volumes, surfaces, or curves with either the rubberband or zone tool.

Solution

The Solution toolbar gives you quick access to basic solution controls.

(Initialize Solution) -- Initializes the solution.

(Step) -- Adds a user-specified number of iterations to a solution.

(Run) -- Begins running the simulation.

(Stop) -- Stops the running of the simulation.

System

The System toolbar provides controls for starting a simulation and saving it to file.

(New Simulation) -- Creates a simulation.

(Load Simulation) -- Opens an existing simulation.

(Save) -- Saves the simulation.

(Save All) -- Saves all open simulations at once.

Visualization

The Visualization toolbar lets you open scene displays, work with them interactively, and store view settings.

(Create/Open Scenes)	Opens a drop-down menu that allows you to add a scene or open an existing one.	
	Geometry	Creates a geometry scene and opens a display for it.
	Mesh	Creates a mesh scene and opens a display for it.
	Scalar	Creates a scalar scene and opens a display for it.
	Vector	Creates a vector scene and opens a display for it.
	Empty	Creates an empty scene and opens a display for it.
	Open	Contains a submenu with names of existing scenes.
	Open All Scenes	Displays all existing scenes in the *Graphics* window.

(Reset View) -- Restores the view to a zoom factor such that the model takes up most of the display. The view angle is unchanged. Use this function if the model has somehow zoomed or panned to such a position that you can no longer see it on the display.

(Rotate) -- Rotates the view around a focal point that your most recent mouse click in the display sets.

(Rubberband Zoom) -- Allows you to zoom to a specific area in the display by clicking two points that define the opposing corners of the zoomed area. Press `<Esc>` to cancel.

(View)	Opens a drop-down menu that allows you to set, save or restore a view.	
	Store Current View	Records properties about the projection mode, zoom factor, pan position and view angle and adds a node to the **Views** node of the simulation tree. It is named View N, where N is the first available number that is not used by existing views with the same prefix. So if you have a view, View 1, then this operation creates a view, View 2.
	Restore View	Contains a submenu of saved views, including a thumbnail of the view. Select one of the submenu items to restore that view.
	Projection Mode	Contains a submenu that toggles between perspective and parallel mode for the current view.
	Look Down	Contains a submenu of standard view angles using the axes as reference vectors.
	Standard Views	Contains a submenu of standard views using the aspects of the object in the scene.
	View Coordinate System	Contains a submenu of coordinate systems in the simulation that a view can use.

(Previous View) -- Reverses the latest change to the view.

(Next View) -- Repeats a change that was subjected to the `Previous View` command.

(Make Scene Transparent) -- Makes all parts transparent so that you can see all of them from any angle, and overrides the opacity settings of all the displayers in the scene. Restores the displayer settings when you click the button again.

(Create Planar Section) -- Allows you to create a plane section derived part, which is normal to the view, by clicking two points in the display.

(Measure Distance) -- Lets you click two points in the display to measure the distance. After your second click, the information on point coordinates and distance appears in the *Output* window. If you want to cancel this process after clicking the button, simply press `<Esc>`.

(Show All Meshes) -- Adds mesh lines to the display. Clicking this button once adds the mesh lines, and overrides the settings of all the displayers in the scene. Clicking it a second time reverts to the displayer settings.

(Hide All Meshes) -- Removes mesh lines from the display. Clicking this button once removes the mesh lines, and overrides the settings of all the displayers in the scene. Clicking it a second time reverts to the displayer settings.

Using the Explorer Pane

The *Explorer* pane is typically docked on the left of the workspace.

This pane can hold either of the following kinds of window:

- Zero or one Servers window (this window indicates servers on the network and is closed by default)
- Zero or more named simulation windows (each running a simulation; *intake* in the example below)



Each window type holds a tree view and, sometimes, a status bar of some kind. There are general rules for navigating the trees (such as opening and closing nodes) that apply to all the trees, and also general rules for using the windows.

The tabs near the top of the simulation window allow you to toggle among views of simulation objects, scene or plot objects, or editing panels when applicable.

Servers Window

This window shows the Servers object tree.

The nodes in this tree represent STAR-CCM+ server processes running on your local network.

The *Servers* window is closed by default. To open it, select **Window > Servers** from the top menus.



By default, only processes that you start are displayed. STAR-CCM+ server processes that others start can be displayed by entering different user names in the *User* property of the **Servers** node. This feature can be useful if your user name on a remote machine differs from your local user name.

Simulation Windows

This window shows the named simulation tree, which contains nodes that represent the various objects of the open simulation.

Multiple windows of this kind can be present at the same time--one for each open simulation.



Near the top of a simulation window are one or more tabs that allow you to use this window for specific tasks, where applicable:

- Viewing the objects of a scene or plot, when its display is open and active



- Editing an object, particularly creating a derived part



The *simulation* tab is always present in a simulation window.

Working with the Object Tree

The STAR-CCM+ user interface is highly object-oriented, with objects representing various aspects of a simulation, and with most objects having properties that you can set as required for your particular problem.

The objects are presented as nodes on a tree view, and properties are presented in a spreadsheet-like view tightly linked to the selected node in the tree.

This section includes explanations of what the object tree is, how to navigate around and use the tree view, and how to select nodes in the tree view. This section also covers how to create, copy, paste, delete, rename, sort, and group objects in the simulation tree, as well as how to copy properties among them. Information about object dependencies is included.

What Is the Object Tree?

This section explains the use of objects in STAR-CCM+.

The object-oriented nature of the GUI reflects itself in several ways, some of which make it feel a little different than some other CFD applications. For this reason, it is important to understand the underlying principle. “Object-oriented” involves regarding all aspects of STAR-CCM+ in terms of objects that exist on the server, are reflected on the client, and are presented to you in multiple ways. This principle applies to high-level concepts like regions as well as small concepts like material properties.

For example, a data plot can be viewed as a hierarchy of component objects on a tree, or it can be represented pictorially as a graph. Similarly, a derived part (such as a surface cut) exists on the tree as an object in its own right, but can also appear in one or more three-dimensional scenes. Yet the same underlying object is being viewed in each situation.

This strong object-oriented nature of the GUI also reflects itself in the content of the menu bar. In many other applications, you can learn about all of the features by browsing through the menus, with the unwritten understanding that all features are presented in the menus. In addition, you have access to shortcuts like toolbars and pop-up menus.

In STAR-CCM+, however, the core of the application is in the objects, and the objects are represented on a tree. So typically you select an

object on the tree and then work with it by right-clicking the object and selecting from a pop-up menu. Much of the functionality is not available in the menus. To "browse" the features of STAR-CCM+, right-click, double-click, and drag on objects in the workspace, in addition to browsing the menus.

Some functionality, such as global settings for the application as a whole, appear only on the menus. Other functionality is duplicated on the objects and on the menus. Duplication is available where it has the potential to increase efficiency.

Another aspect of the object-oriented nature of the GUI is that "objects have properties". To support this aspect, the GUI uses a *Properties* window where all of the properties for a selected object can be viewed and manipulated. This feature significantly reduces the number of large and potentially complex dialogs that are required. In some cases, to increase efficiency in the use of the application, properties of multiple objects get combined in a larger dialog.

Object Dependencies

When an object in STAR-CCM+ requires data from another, a dependency relationship is formed between the objects.

Dependency relationships affect which objects you can delete and which you cannot, as well as the consequences of modifying an object. Examples include:

- A monitor plot showing data from a monitor
- A scalar scene displaying data from a field function
- A report using data from a coordinate system

The types of STAR-CCM+ objects with dependency relationships include:

- Annotations
- Coordinate systems
- Field functions
- Monitors
- Plots
- Reports

Checking an Object for Dependencies

To check whether an object has dependencies:

1. Right-click the node of that object.
2. If the pop-up menu includes a **Dependencies** item, select it.



Before you delete or modify an object, you can assess the potential impact of your change by looking at the dependent objects. To examine an object under the *Dependencies* tab:

1. Right-click its node and choose **Select in Simulation**.



STAR-CCM+ locates the selected object in the object tree.

2. When you are finished working with the *Dependencies* tab, click **Close**.

Deleting Objects with Dependencies

To delete an object under the *Dependencies* tab:

- Examine the node hierarchy. Each parent node is referenced for data by its immediate sub-node.



- To delete all the objects, click **Delete All**. Click **Yes** in the confirmation dialog.
- To delete the object at the bottom of the hierarchy (which has no dependencies), right-click its node and select **Delete**. Click **Yes** in the confirmation dialog.

  Deleting any node in the hierarchy other than the one at the bottom also deletes all of its dependencies.
- To delete an object and all of its dependencies, right-click its node and select **Delete**. The confirmation dialog will ask whether you want to delete all of those objects. If so, click **Yes**.

Using Tree Views

This section gives guidance on how to work with and navigate around object trees.

Tree views appear in several places in STAR-CCM+.



Specific topics include:

- Understanding the icons
- Opening and closing parent nodes
- Searching the tree
- Actions and default actions
- Using drag-and-drop
- Navigating during drag-and-drop
- Collapsing the object tree

Understanding the Icons

Icons represent objects and their hierarchical relationships.

There are two types of graphical widgets that are used in the tree views:

- A handle indicates whether the node has children, and serves as a mouse-click target to open and close the parent node. The handle can be:
  - Closed (the sub-nodes are not visible)
  - Open (the immediate children are visible)
- An icon provides a visual cue as to what object the node represents. Icons fall into two groups:
  - A folder ( ) indicates that the node is a "manager" or "collection" of several nodes of the same type. For example, the **Scenes** node is a manager for several individual scene nodes.
  - An object other than a folder has a pictorial icon to represent the type of object. Under boundaries for example, each boundary of a different type has a different icon, and as the boundary type is changed in the properties, the icon changes.

Opening and Closing Parent Nodes

Whether you prefer using the mouse or the keyboard, you can use either device to open and close parent nodes.

Using the Mouse

Click the handle. If it is closed, the node opens. If it is open, the node closes.

It is also possible to collapse the object tree.

Using the Keyboard

With the focus on the node of interest, press the `<Left>` arrow to close the node (assuming it is open) and the `<Right>` arrow to open the node (assuming it is closed). Adding `<Ctrl>` to each keystroke closes and opens all nodes and sub-nodes within the selected node respectively.

If the node is already closed, or has no children, when you press the `<Left>` arrow, the selection jumps to the parent node of that node. If the node has no children, or is already open when you press the `<Right>` arrow, the selection goes to the next node down (child or sibling).

Searching the Tree

You can search the tree using keystrokes.

If the focus is on the tree, that is, on one of its nodes, pressing an alphanumeric key (lower case only) has the following effects:

- Moving the focus to a node beginning with the typed character (if one exists)
- Opening a search edit box with that character



The preceding screenshot was taken after typing an `<fl>` when the root of the tree was selected. Continue typing to jump to the first node that matches the character set typed so far.



If no match is found, then no node is selected.

The search facility is not case-sensitive:

- Typing a capital letter does not open the search box.
- Once the search box is open, case is ignored. So if you have one scene that is called `Fred` and another called `fred`, the search always highlights `Fred`.

To deactivate the search box, do either of the following:

- Press `<Esc>`.
- Press `<Enter>` to execute the default action on the selected node.

Actions and Default Actions

Most tree nodes have actions that are associated with them.

If a tree node has any associated actions, right-clicking the node opens a pop-up menu. If multiple nodes are selected and the actions do not apply to all selected nodes, some menu items are grayed.

STAR-CCM+ remembers your actions with a node during a simulation session. For example, when you double-click a node that you have not accessed previously during the session, the node expands and a *Customizer* dialog appears for that node. However, if you close the node back and double-click again, the *Customizer* dialog appears but the node remains closed, as you left it.

One of the actions is a default action--it is executed once you press `<Enter>` on the focused node or double-click the node. Unfortunately, there is no way to tell which of the actions is the default action just by looking at the pop-up menu.

Using Drag-and-Drop

Drag-and-drop is a productivity feature that can often save several key or mouse strokes.

To use drag-and-drop:

1. Place the mouse over a node.
2. Press and hold the left mouse button.
3. Move the mouse to another node or another window, and navigate if necessary.
4. Release the mouse button.

During step 3, the mouse cursor changes to indicate whether a drop would do one of the following when you release the mouse:

- Be invalid (a no-entry icon)



- Perform a copy operation (a plus icon)



Navigating during Drag-and-Drop

If the target node is not visible when you start the drag operation, you can navigate while you have the left mouse button depressed during the drag.

Scroll Up or Down

Move the mouse cursor near the top or bottom of the tree window. If necessary, move slightly outside the border to start the tree scrolling. The tree continues scrolling until the top or bottom, or until you move the mouse back into the window.

Open Nodes with Children

Hold the mouse over the node for a second or so, and the node opens. The node must be highlighted for this action to work.

Collapsing the Object Tree

After opening many different nodes in the object tree and navigating among them for a while, closing everything can clarify your overall view of the simulation.

You can use a single menu action to shrink the tree back to its default size. Right-click the simulation node at the root of the object tree and select **Collapse Tree** from the pop-up menu.



Selecting Nodes

Nodes on the object tree can be in various states of selection depending on how you click them.

The selection state of nodes affects what actions can be performed on them. You can select single nodes, or a contiguous block of multiple nodes or a non-contiguous block of multiple nodes.

Selection Modes

Access to a node or its properties depends on the selection state of the node.

A node can be:

- Selected (can view properties and perform common actions such as delete) or unselected
- Focused (can navigate with the keyboard) or unfocused

Combining these options gives four possibilities with visual cues for each state:

- Unselected and unfocused Scalar Scene 1 - no background or border
- Unselected and focused Scalar Scene 1 - no background, darker border
- Selected and unfocused Scalar Scene 1 - lighter background, no border
- Selected and focused Scalar Scene 1 - lighter background, darker border

Selecting Individual Nodes

Whether you prefer using the mouse or the keyboard, you can use either device to select individual nodes.

Using the Mouse

Simply click a node.

Using the Keyboard

To move the selection up or down the tree, press the `<Up>` or `<Down>` key. A node must already have the focus for this action to work.

Selecting Multiple Nodes

It is possible to select more than one node at a time.

The selection can be contiguous or non-contiguous. In these modes, the *Properties* window displays properties that are common to all of the selected nodes, and if there are no common properties, then the *Properties* window is empty. Right-clicking one of the selected nodes displays a pop-up menu of actions that are common to all nodes.

To change properties of more than one object at a time, use multiple selections.

Selecting a Contiguous Block of Nodes

Whether you prefer using the mouse or the keyboard, you can use either device to select a contiguous block of nodes.



Using the Mouse

Click the first node of the block, press and hold down the `<Shift>` key, and click the last node of the block.

Using the Keyboard

Select the first node of the block using the `<Up>` or `<Down>` arrow keys, hold down the `<Shift>` key, then press the `<Down>` arrow until you reach the last node of the block.

In the preceding screenshot, three nodes are selected, and the last one is also focused.

Selecting a Non-Contiguous Block of Nodes

Whether you prefer using the mouse or the keyboard, you can use either device to select a non-contiguous block of nodes.



Using the Mouse

Click the first node of the block, press and hold down the `<Ctrl>` key, and click the other nodes that you want to include in the block.

Using the Keyboard

Select the first node of the block using the `<Up>` or `<Down>` arrow keys, hold down the `<Ctrl>` key, and press the `<Down>` arrow until you reach the next node of the block. Press the `<Space>` bar. Continue navigating

up or down and pressing `<Space>` on required nodes until you are done.

In the preceding screenshot, two nodes are selected, and the last one is also focused.

Creating (Adding) Objects

Objects are created in STAR-CCM+ in various ways, some through a deliberate action by you, others as the automatic result of a separate step.

Details of how objects are created are included in the explanation of each particular object.

Examples of objects that are created through a deliberate action include:

- Simulations
- Generated meshes
- Reports
- Monitor plots and XY plots
- Scenes
- Derived parts

Examples of objects that are created automatically include:

- Continua -- during a mesh import. A continuum can also be created manually.
- Regions and boundaries (from an imported mesh)
- Models
- Residual plots
- All the various manager nodes that folder icons ( ) represent in the object tree

Copying Objects

To give one node the same characteristics (such as properties or child nodes) as another of the same type, copy one node over the other using the clipboard or drag-and-drop.

Copying objects does not always transfer dependent objects automatically. In particular, when you copy an object from one

simulation to another, copy all of its dependent objects first. For example, to copy a scene displaying derived parts into another simulation, you would have to copy the derived parts beforehand so that the scene could access them properly.

The specifics about what gets copied (some properties, all properties, or child objects) are detailed in the user guide explanation of the specific object.

Copying Objects Using the Clipboard

You can copy STAR-CCM+ objects using standard copy and paste operations.

Select the node from which you want to copy properties and/or child objects, and press `<Ctrl><C>` on the keyboard, select the **Edit > Copy** menu item or click the (**Copy**) toolbar button.

With some nodes, you can also right-click on the node and select **Copy** from the pop-up menu.



This action places the object properties on the clipboard ready for pasting onto another node.

Now select the node to which you want to copy the properties and press `<Ctrl><V>`, or click the (**Paste**) toolbar button. With some nodes, you can also right-click on the node and select **Paste** from the pop-up menu.



The node where you applied the **Paste** command resembles the node from which you copied properties.



The data that are placed on the clipboard by this operation cannot be viewed. If you attempt to paste it into a text editor, for example, you get unintelligible material that is usable only by STAR-CCM+.

Using Drag-and-Drop to Copy Objects

You can copy STAR-CCM+ objects by dragging and dropping them on particular target objects.

Click the node that has been set up and, while still holding down the left mouse button, drag the mouse until the pointer hovers over the target node.



To finish copying all the node properties from the first node to the second, release the mouse button.

		
Before copy		After copy

Pasting Objects

The paste action places the contents of the clipboard into the selected node.

This action is unavailable when the clipboard does not contain material that can be added in this manner.

Depending on the context, the paste action could do one of the following:

- Make a duplicate of the object
- Copy the properties of the clipboard object to the current object
- Make an association between the clipboard object and the current object

Select the node where you want to paste properties, and press `<Ctrl><V>` on the keyboard, select the **Edit > Paste** menu item or click the (**Paste**) toolbar button.

Deleting Objects

Many of the objects on the tree that are added through importing a mesh or setting up the simulation can be deleted.

To delete an object, select **Edit > Delete**, use the pop-up menu, or use the keyboard.

Use this option with care, since it removes the object permanently; the deleted objects are not moved to the Trash or Recycle Bin of your computer.

Some objects cannot be deleted.

Using the Pop-Up Menu to Delete Objects

The delete action appears in the pop-up menu of many objects.

Select the object in the simulation tree that you want to delete (a region in the example below) then right-click the selection and select **Delete** from the pop-up menu. (Deleting multiple objects is shown under deleting objects using the keyboard.)



If this menu item is grayed, then the object currently cannot be deleted for some reason.

If the delete action worked, a confirmation dialog is displayed showing the name of the selected object and asking if you really want to delete it.



If you click **Yes** on the confirmation dialog, the object is deleted and the simulation tree is updated accordingly. No messages are printed in the *Output* window.



Using the Keyboard to Delete Objects

The `<Delete>` key can be used to delete a variety of objects.

Select the objects in the simulation tree that you want to delete (two regions in the example below) then press the `<Delete>` key. (Deleting a single object is shown under deleting objects using the pop-up menu.)



If nothing happens when you click the <Delete> key, then one or more of the selected objects currently cannot be deleted for some reason.

If the delete action worked, then a confirmation dialog is displayed asking if you really want to delete the number of items you selected. The names of these objects are not displayed.



If you click **Yes** on the confirmation dialog, the objects are deleted and the simulation tree is updated accordingly. No messages are printed in the *Output* window.



Objects That Cannot Be Deleted

The implications of deleting a specific object, and when it cannot be deleted, are included in the user guide for the specific objects.

In general, objects that cannot be deleted include:

- All the objects that are automatically created on the tree of an empty simulation
- Boundaries that are used as interfaces
- Objects that other objects use, such as when a monitor plot receives data from a monitor, or a continuum is assigned to a region

When an object uses another object, there is a dependency. When you attempt to delete an object that has a dependency, STAR-CCM+ opens a warning dialog:



- To find out what the dependency is, click **Details**. A *Dependencies* tab appears in the *Explorer* pane.



The above example shows the results of an attempt to delete a report that has dependencies. A monitor requires the data from that report, and a plot requires the data from the monitor. Due to these dependencies the report could not be deleted.

Renaming Objects

Most objects on the tree can be renamed.

Typically manager ( ) objects cannot be renamed. Bear in mind that most objects within one simulation must have unique names. The

uniqueness rules are case-sensitive, so `Wall-1` is different from `wall-1`. If you do inadvertently use a duplicate name for an object, STAR-CCM+ appends a space and a digit to what you typed so that the result is a name that is still unique.

One major exception to this rule, however, is with region and boundary names. You can have a region named `Element` and a boundary named `Element` in the same simulation. Also, you can give the same name to multiple boundaries, provided each of those boundaries is within a different region.

To change the name of an object, you have three options:

- Selecting the node and pressing `<F2>`
- Slowly clicking the label of the node
- Right-clicking the node and selecting **Rename...**

Any of these options activates the inline editor in the label of the node.

To apply your change, press `<Enter>`

To keep the existing name, even if you have changed the name in the inline editor, press `<Esc>` or click outside the node label.

When you rename a node, it is not sorted with its siblings. To alphabetize the nodes, use the sort action.

Using the Inline Editor to Rename Objects

You can rename objects directly in the tree by using an inline editor that opens a small input box at the node you are trying to change.

You can access the inline editor in one of two ways.

Pressing <F2>

Select the object of interest so that it is in focus, then press the `<F2>` key. This procedure allows you to carry out an inline rename of the object.

Clicking Twice Slowly

Clicking the object name twice, slowly enough that you are not double-clicking the object, puts the object into inline rename mode.

In both cases, the edit box is opened on the selected node, and the entire text is selected.



When the node label is in this state, typing any character erases the selection and starts typing a new name. Alternatively, use the `<Left>` or `<Right>` cursor keys to remove the selection and then edit the name as normal.

Press `<Enter>` once you have modified the name or press `<Esc>` to cancel the edit. The normal rules of uniqueness are applied when you press `<Enter>`.

Using the Pop-Up Menu to Rename Objects

The rename action appears in the pop-up menu of many objects.

Right-click the object (a scene in the example below) and select **Rename...** from the pop-up menu.



Enter the new name of the object.

Press <Enter>.

If this object appears in a scene or a plot, then the new name appears in both places.



Sorting Objects

You can sort child nodes of a selected object alphabetically.

This feature is typically available to manager nodes and would be used if any of the sub-nodes were renamed. Nodes are not sorted automatically after a name change, and this action will force a sort. You can sort by pressing <F5> or by right-clicking a manager node and selecting **Sort**.

Linux (Unix) and Windows systems sort differently:

- Linux systems sort taking case into account with upper case being "smaller" than lower case (Wall, cyclic).
- Windows systems sort by ignoring case (cyclic, Wall).

Grouping Objects

When an object tree contains many nodes of the same type, such as interfaces, it helps to be able to organize them into groups.

The following types of objects can be handled this way by creating subfolder nodes in their manager nodes:

- Regions
- Boundaries
- Interfaces
- Derived parts
- Reports

- Monitors
- Plots
- Scenes

Placement of objects in a group has no effect on how they function in the simulation.

To use the feature right-click, for example, the **Reports** node and select **New Group**.



A **New Group** node appears in the object tree.



You can then drag-and-drop nodes into this group node. It is possible to rename the group node to represent its contents.

If you have a set of items so large that you want subcategories, it is also possible to create groups within groups.



To reverse the process, right-click the manager node (**Reports** in this example) and select **UnGroup**. This action moves the objects (the report nodes in this example) back to their original location, and automatically deletes any empty group nodes.

The **Group By** pop-up menu item is an alternative way to use this feature. It takes all the properties of the objects in that manager node and makes them available as group categories.



Group nodes that are created through this automated, property-based approach can also have subfolder nodes that are created within them.

Groupings are only possible based on one property at a time. Pre-existing group nodes, whether they were created manually or automatically, are deleted as the objects are moved into the new groups.

Copying Properties

Copying properties is one result of copying objects.

If you have two (or more) nodes that you want to have the same properties, then you can copy the properties from one node to any others using the clipboard functionality:

- Keyboard commands
- Toolbar buttons
- Pop-up menus (in many cases)

You can also use drag-and-drop.

These techniques can be useful even if the two nodes only have part of their properties matching. Set up one node, copy it to the similar node, then adjust the unique properties of the second node.

In most cases, you can copy properties between similar nodes that are in different simulation files.

The effects of these actions (copying some or all properties) vary with the type of object. The specifics about exactly which properties get copied are detailed in the user guide explanation of the specific object.

Common Pop-Up Menu Items

Many objects on the simulation tree have items in the pop-up menu that are common. This section lists all those items, although any single object does not necessarily use all of the menu items.

Edit	Allows customizing of the selected object and its children using a separate modeless stay-on-top dialog.
Copy	Copies the object and selected properties onto the clipboard.
Delete	Removes the selected objects.
Highlight	Highlights an object in a scene. This option is only available if the scene is open.
Paste	Replaces the properties of the current object by pasting the ones from the clipboard.
Refresh	Updates the selected object or display so it reflects the latest changes to the simulation.
Rename	Allows you to change the name of the currently selected object.
Sort	Sorts child nodes in alphabetical order. This item is typically available on manager nodes.

Using the Table Editor

The table editor combines the functionality of the simulation tree and *Properties* window.

It displays the selected node, all of its child nodes, and their properties in a single dialog.



The left-hand column includes nodes and their properties underneath. The right-hand column contains the settings of those properties.

When a row in the table has the focus, its cells have a colored background. The cell in the row that has the focus has a colored background and border.



Note: There is no function available to right-click on a node within the edit dialog. To access pop-up menu options for a node, exit the edit dialog and right-click the appropriate node in the simulation tree.

This section provides instructions for the following:

- Activating the table editor
- Working with multiple objects at a glance
- Adjusting the table editor
- Navigating with the keyboard

Combining the Object Tree and Properties Window

The table editor combines much of the functionality of the object tree and *Properties* window in a single dialog.

- The description area near the bottom left of the table editor summarizes the currently selected object or property, just like the properties description area does.
- You can access the online Help for objects within the table editor just as you would from the *Properties* window. When a node has the focus, pressing the question mark button in the lower right corner of the table has the same effect as the `<F1>` key.
- As with the object tree, *Properties* window and toolbars, hovering with the mouse activates a brief description of the item of interest.

Activating the Table Editor

The table editor can be activated from several places in the workspace.

- Right-clicking on a node in the simulation tree and selecting **Edit...**



- Right-clicking on an object in the *Graphics* window while in interact mode, and selecting one of the edit options from the pop-up menu



These techniques can also be used to activate the multi-edit mode.

Activating the Table Editor in Multi-Edit Mode

The multi-edit mode of the table editor is a useful way to work with multiple objects at once, provided all of those objects are of the same general type, such as boundaries.

To activate this mode, select all the objects you wish to view in the table editor using the `<Shift><Click>` method, or `<Ctrl><Click>` for non-contiguous items. Right-click on these nodes and select **Edit...**



You can also use the `<Ctrl><Click>` selection technique in the *Graphics* window while in interact mode.



Working with Multiple Objects at a Glance

The table editor lets you apply settings to multiple separate objects easily since you can view them all at once, provided their parent object has the focus when you activate this feature.

You can expand the node hierarchy in the table editor by clicking the **Expand/Contract Tree** toolbar button in the top of the dialog. In

addition, if you want to enter different values for a property that is common to multiple nodes, click the **Expand/Contract Values** toolbar button.



This will give you access to the properties of the subnodes. These buttons toggle the rows and columns open and closed.

For example, if you wish to apply the same setting to multiple boundaries, you can do so by referring to a previous entry elsewhere in the table.



Note: There is no function available to right-click on a node within the Multi-edit dialog. To access pop-up menu options for a node, exit the Multi-edit dialog and right-click the appropriate node in the simulation tree.

Adjusting the Table Editor

You can adjust the table editor for easier access to the objects of interest.

The handles on the nodes work just like the ones in the object tree, allowing you to show or hide subnodes or properties.



The column headers, **Nodes** and **Values**, are buttons that let you sort the visible objects in the respective column by name, toggling between ascending and descending order.



Navigating with the Keyboard

If you prefer using the keyboard, you can use keystrokes to access various parts of the table editor.

For navigating among the cells in the body of the table, the following keys are supported:

- `<Right>` - shifts the focus from one column to the one right of it. When the focus is in the tree column, pressing this arrow opens the node of the row (assuming it is closed).
- `<Ctrl>+<Right>` - When the focus is in the tree column, pressing this combination opens all nodes and subnodes within the selected node of the row.

- <Left> - shifts the focus from one column to the one left of it. When the focus is in the tree column, pressing this arrow closes the node of the row (assuming it is open).
- <Ctrl>+<Left> - When the focus is in the tree column, pressing this combination closes all nodes and subnodes within the selected node of the row.
- <Tab> - works like the <Right> arrow.
- <Shift><Tab> - works like the <Left> arrow.
- <Up> - shifts the focus from one row to the one above it.
- <Down> - shifts the focus from one row to the one below it.
- <Home> - shifts the focus to the left cell of the row.
- <End> - shifts the focus to the right cell of the row.
- <Space bar> - goes to edit mode in a property cell.
- <Esc> - goes to view mode in a property cell, assuming that cell is in edit mode; it does not change the value if you have changed it.

  This key closes the dialog if pressed when there is no property cell in edit mode.

To access a drop-down list, navigate with the arrows until the cell of that list has the focus, then press the <Space> bar to display the list. Use the <Up> or <Down> arrow keys to reach the desired item in the list, and then press <Enter>. To exit the list without entering any change, press <Esc>.

To move to the top row of the table body, press <Ctrl>+<Home>. To move to the bottom row, press <Ctrl>+<End>.

To select a contiguous block of nodes, select the first node of the block using the <Up> or <Down> arrow keys, hold down the <Shift> key, then press the <Down> arrow until you reach the last node of the block.

To select a non-contiguous block of nodes, select the first node of the block using the <Up> or <Down> arrow keys, hold down the <Ctrl> key, then press the <Down> arrow until you reach the next node of the block and press the <Space> bar. Continue navigating up or down and pressing <Space> on required nodes until you are done.

Using the Properties Window

The *Properties* window is a single location where all of the properties for a selected object can be viewed and manipulated.



It is designed for easy navigation with the mouse or the keyboard, and includes a description area at the bottom of the window.

To open this window, select **Window** > **Properties**.



The dialog that appears under the *edit* tab of the *Explorer* pane, typically used for derived parts and surface repair, causes the *Properties* window to disappear from the workspace. It reappears when the *edit* tab is closed.

A set of editing techniques is common to virtually all STAR-CCM+ objects.

Additional methods of adjusting properties include:

- Using the part selector
- Types of inline editors
- Using a property customizer
- Using a slider

See Also

Preparing Surfaces
Using the In-Place Dialog

Navigating Around the Properties Window

Various techniques exist for finding items in the *Properties* window.

The search field is a powerful tool that brings a specific property into focus immediately. Whether you prefer using the mouse or the keyboard, you can use either device to work with the *Properties* window. Navigating this window becomes easier once you are familiar with the property modes.

See Also

Accessing the Properties Window with the Keyboard

Using the Search Field

This section explains how to use the search field in the *Properties* window.

When the *Properties* window has the focus, typing a keystroke causes a search field to appear.



If you type one or more letters that are in the name of one of the properties, that property gets the focus.

Using the Mouse

This section describes how you can use the mouse to edit properties.

Click the left-hand side of a property to select it, click the right-hand side or the ... (property customizer), if one is there, to edit it.

The handle in a category bar works just like that of a parent node in the object tree:

- closed (the properties are not visible); or
- open (the properties are visible).

Property Modes

Properties are displayed in several modes which depend on functionality.

A line in the property list has the following modes:

- View only, unfocused - name and value show on white background.
- View only, focused - name and value show on blue background.
- Edit (has the focus by default) - name shows on blue background, value is in edit mode.

Common Editing Techniques

This section explains how you can edit property values.

Put the property into edit mode by:

- Clicking the property value (in the right-hand column);
- Using the keyboard to navigate to the property and press the `<Space>` bar; or
- Clicking the ellipsis button in the right-hand column (if one is there) to activate the property customizer.

Change the value by typing, making a selection from a drop-down list or making the changes within the property customizer (if one is activated).

Accept the new value by pressing `<Enter>` in the inline editor or clicking **OK** on the *Customizer* dialog. The property is then put into view mode and retains the focus. If the value that you entered was invalid for some reason, a warning box is displayed and the value reverts to its previous value.

Cancel your editing by pressing `<Esc>` or clicking **Cancel** on the *Customizer* dialog.

Using the Part Selector

This section explains how to use the part selector dialog.

A special kind of property, usually labeled *Parts*, activates the object selection dialog. This dialog lets you select parts of the model for a derived part.



To select an object, click the selection button to the left of it in the selection column, then click **OK**.

See also: Selecting Objects

See Also

The Object Selection Dialog

Types of Inline Editors

There are four kinds of inline editors.

Text Box

This inline editor accommodates all the keys you would expect to work while typing and editing text. Simply type the new value and press `<Enter>`. For some objects, you can also enter an expression, such as `10*sin(3.14*$Time)`, directly.

Checkbox

Click with the mouse to tick or clear the box, toggling between the two settings of the property.

Drop-Down List

Select with the mouse or use the `<Up>` and `<Down>` cursor keys followed by `<Enter>`.

You can use keystrokes to search drop-down lists.

Mini-Dialog

Make a selection using the mouse, there is no keyboard navigation.

See Also

Accessing the Properties Window with the Keyboard
Setting Conditions and Values

Using a Property Customizer

Use the property customizer for greater control over a property's value, typically with control over units.

The most common method for activating this dialog is clicking the ... (property customizer) in an object's value property.



For some objects, you can also enter an expression directly into the value field.



Alternatively, if you click the ... (property customizer) in that field, it activates a larger text box (helpful for longer entries).



See Also

Setting Values Using an Expression

Using a Slider

For any property with a finite range of values, a slider is available for adjustments.

To use this feature:

1. Click the property value. This action makes the slider available.



2. Drag the slider button along the range until you reach the value you want.



3. If the relevant property is displayed in the current scene, the scene automatically updates to the new value.
4. To close the tool, click anywhere else in the STAR-CCM+ window.

   The new value is automatically applied.

   Alternatively, you can enter a value by typing it.

Showing and Hiding the Description Area

At the bottom of the *Properties* window is a description area that summarizes the currently selected object or property by displaying the name of that object or property along with a brief description.

The example below shows the description area when an object on the simulation tree has the focus.



To hide the description area:

1. Right-click anywhere within the *Properties* window, except on the title bar or within a property value that is in edit mode.
2. In the pop-up menu that appears, select **Show Description Area** so that the checkbox of that menu item is cleared.

To restore the description area:

3. Repeat the step so that the checkbox is ticked.



Properties Window Reference

This section describes the structure of the properties window.

The *Properties* window contains one or more collapsible lists, or categories, each of which has a category bar at the top. This window also has its own pop-up menu.

Using Properties Categories

This section explains how to use the different categories of properties in the *Properties* window.

Properties are grouped in different categories for ease of use, depending on the tree that is visible in the *Explorer* pane, with the most commonly accessed properties in the top category, which is visible by default.

- A simulation tree always has a *Properties* category, and may optionally have an *Expert* category depending on which node is selected.

  You are typically not required to visit the *Expert* category, as this holds properties that usually do not need to be modified.
- The Servers tree always has a *Properties* category, and may optionally have an *Expert* category depending on which node is selected.
- A handle controls each category inside its category bar. The handle can be:
  - closed (the properties are not visible); or
  - open (the properties are visible).

Each category in the *Properties* window contains a list of properties for the selected node in the *Explorer* pane arranged in two columns. The property name is on the left and the property value is on the right.



Pop-Up Menu

Pop-up menu items give you more control over the information in the *Properties* window.

Sort by Category	Sorts the properties by type. This feature acts as a radio button.
Sort by Name	Sorts the properties in ascending alphabetical order. This feature acts as a radio button.
Show Description Area	Controlled by a checkbox: Shows the description area when ticked, and hides it when cleared.

Using the Graphics Window

The Graphics window can contain Graphics windows and/or plot windows -- powerful tools that let you view developments in the simulation in real time, as well as plot data.



The following sections provide explanations of:

- Continuous updating
- Working with displays
- Splitting the window
- Viewing two simulations side-by-side
- Using scenes to interact with the object tree
- Working with the Graphics window
- Selecting parts by type in the display
- Changing a three-dimensional visualization display
- Changing a two-dimensional visualization display
- Interacting with the plot display

Continuous Updating

The scenes in STAR-CCM+ always display an up-to-date view of the server objects.

When the solution changes (after an iteration for example), all objects that are viewed in the GUI are updated to reflect the new values. This feature is valuable for monitoring the progress of a solution and being able to make corrections to the settings and conditions to allow you to get to a faster solution.

The downside of this feature is that continuous updating can be expensive and can slow down the server in some situations. The solution would be to close any open scene displays, or at least scenes which are no longer essential.

Working with Displays

Visualization or plot displays are a way to make the best use of the STAR-CCM+ GUI.

This section explains the following:

- Refreshing the display
- Rearranging displays
- Bringing a display tab to the foreground
- Controlling the page setup of a display for printing
- Printing a display
- Using pop-up menus

See Also

Plot Tab Reference
Working with Graphics Windows

Refreshing the Display

Typically the display of a scene or XY plot updates automatically to reflect the latest changes to the simulation. However, in case it fails to do so, you can force an update.

To force a scene to update:

1. Right-click the scene node, or in the background of its display, and select **Refresh**.



To force an XY plot to update:

2. Right-click the XY plot and select **Refresh**.



If you select this menu item when the display of this scene or plot is not open or active in the *Graphics* window, the open scene or open plot action is triggered.

See Also

Interacting with the Scene
Plotting Results
What Is an XY Plot?
What Is a Scene?

Rearranging Displays

You can reposition displays inside the *Graphics* window with a drag-and-drop mouse movement.

To rearrange displays:

1. Click the tab of the display that you wish to move. As you begin to drag the tab, the display has a red outline.



2. Drag toward the desired position until you see the red outline jump to that position.



3. Release the mouse button. The display and its tab are in the new location.



Bringing a Display Tab to the Foreground

Click the tab with the left mouse button.

There are no keyboard alternatives for navigating through the tabs.

Controlling the Page Setup of a Display

Follow these instructions to adjust the print settings of a display.

1. Select **File > Page Setup...**

The operating system-specific page setup dialog appears.



2. Here, you can specify detailed aspects of the print job, such as margins and page orientation.

Printing a Display

Follow these instructions to print a display.

1. Select **File > Print...** or press `<Ctrl><P>` on the keyboard.

The operating system-specific standard print dialog appears.



2. Click **Print** (or **OK**) to print the contents of the active display in the *Graphics* window.

   The contents of this dialog are not documented further. Refer to your operating system documentation for details of using this dialog.

   This action is only available if a simulation is visible and a display (scene or plot) is visible in the *Graphics* window (it does not have to have the focus).

Using Pop-Up Menus

The pop-up menu that appears on each display tab gives you an extra measure of control over these windows.

To activate the pop-up menu:

1. Right-click on the display tab.



The menu items refer to these displays as "documents" because they can be printed, saved to a graphic, and updated in real time.

Splitting the Window

STAR-CCM+ lets you display two or more scenes simultaneously, as well as revert to displaying a single scene.

You may need to compare subtle patterns in two different displays, say, scalars and vectors. For this reason, it would be helpful to look at multiple scenes at the same time, rather than constantly working with tabs.

Viewing Two Scenes Side-by-Side

Scenes can share space in the *Graphics* window.

Using the left mouse button, drag the tab of one of the displays in the direction where you want it to appear. In the illustrations below, the *Geometry Scene 1* display tab is dragged to the right so that the display occupies the right half of the *Graphics* window.

A red outline appears around the scene display as you begin to move it.



As you continue to move the mouse in that direction, the outline changes to show how the *Graphics* window will be split between the scenes.



Let go of the mouse button. The scenes share space in the *Graphics* window.



You can use the same technique to divide the *Graphics* window further among more scenes.

Showing More Scenes Simultaneously

You can arrange more scenes in a column alongside a main scene.

To do this, drag the tab of the additional scene downward over the open scene with which it should be stacked.



As before, keep dragging until the outline changes to indicate a split.



This action creates the stack.



Reverting to Displaying One Scene

The drag-and-drop technique that is used for splitting scenes can also be used to reverse the splits.

To reverse a split, drag one tab over another. In the illustrations below, the scene in the lower left corner of the *Graphics* window is being dragged over the one in the upper left corner.



As a result, the scene that you drag over the other now takes the entire left half of the *Graphics* window. Only two of the three scene displays are active, while the third is open but inactive.



You can keep following this procedure until only one scene display is active at a time.

Viewing Scenes from Two Simulations Side-by-Side

View two scene displays from separate simulations side-by-side for easy comparison.

At times, you may need to view scenes from two or more simulations simultaneously to compare data. This can be done with the window-splitting capability in STAR-CCM+.

1. With one or more scene displays open in one simulation, open a second simulation in the same STAR-CCM+ client. This results in two open simulation windows in the *Explorer* pane that you can switch between.



2. Open a scene display in the second simulation.



3. Split the *Graphics* window.



The two scene displays from the separate simulations are side-by-side for easy comparison.



When more than one simulation is open, the tab includes the name of the corresponding simulation. This feature lets you know at a glance which simulation a display represents.

Using Scenes to Interact with the Object Tree

The sets of instructions below show how you can use the various connections between the *Graphics* window and the object tree.

- Finding a scene in the tree
- Highlighting objects, including multiple objects and objects obscured from view
- Highlighting objects in a scene by using their tree nodes
- Highlighting objects in a scene without labels
- Dragging parts onto a scene

Finding a Scene in the Tree

This section explains how to determine which node in the object tree corresponds to a display window.

1. Click the background of a display window.



The label of the node corresponding to that window is highlighted.



When more than one simulation is open, the tab includes the name of the corresponding simulation. This feature lets you know at a glance which simulation a display represents.

Highlighting Objects

When you click an object in either the display window or on the object tree, that object is highlighted in both locations simultaneously.

1. Click a part inside a window.



The part becomes labeled and highlighted.



The label of the boundary node corresponding to that part is also highlighted in the simulation tree.



Note: This default behavior can be stopped by deactivating the *Highlight Parts in Scene* option in the **Visualization** node of the *Options* dialog. Deactivating this option can improve the speed of response in the *Graphics* window when working with models that contain many surfaces. It is also possible to highlight objects in a scene without the labels.

2. To access this property, select **Tools** > **Options** from the main menu in the STAR-CCM+ workspace.

   The *Options* dialog appears.
3. Expand the **STAR-CCM+** node in the **Options** node.
4. Select the **Visualization** node and deactivate the *Highlight Parts in Scene* option to stop highlighting selections in scenes.



5. Click **Close**.

To highlight parts, part contacts, part surface contacts, and objects in custom mesh controls in the *Graphics* window when the *Highlight Parts in Scene* function is deactivated:

6. Right-click the node you wish to highlight.
7. Select **Highlight**.

See Also

Highlighting Objects Without Labels

Highlighting an Object Using the Tree Node

You can highlight an object in a scene by right-clicking on its node in the object tree.

The same works for a region or part node. This functionality also works for multiple objects. Additionally, it is possible to highlight objects in a scene without the labels.

1. In the object tree, right-click the node of a boundary and select **Highlight**.



The object is highlighted in the display.



2. Use the (**Reset**) button in the toolbar to view the entire shape.

See Also

Synchronizing Multiple Objects
Highlighting Objects Without Labels

Synchronizing Multiple Objects

It is possible to work visually with multiple objects by selecting them in a scene.

It is possible to highlight objects in a scene without the labels.

1. To select multiple objects in a scene, hold down the `<Ctrl>` key while clicking them.

The labels of the corresponding nodes in the object tree are also highlighted.



Depending on what type of objects you select, right-clicking on one of the selected objects activates certain pop-up menu options. For example:

2. Right-clicking on one of two selected boundaries in a scene allows you to combine those boundaries.



See Also

Highlighting an Object Using the Tree Node
Combining Boundaries
Highlighting Objects Without Labels

Highlighting Objects Without Labels

It is possible to highlight a boundary in a scene without the label of the boundary appearing.

This facility is useful if your simulation contains many boundaries, when the display of the labels can take a long time to load.

To change the setting:

1. Select **Tools > Options**.
   The *Options* dialog appears.



2. Expand the **STAR-CCM+** node in the **Options** node.
3. Select the **Visualization** node and deactivate the *Include Captions with Highlights* option to omit labels from highlighted objects in the display.
4. Click **Close**.

See Also

Highlighting Objects
Highlighting an Object Using the Tree Node
Synchronizing Multiple Objects

Accessing Objects Obscured from View

There are several ways to get access to an object in a scene when it is obscured by other objects.

Hiding Objects

1. The most common technique is to hide any objects that are in the way.



Rotating Objects

2. Rotating the object slightly can give you access to distant points. Move the mouse until you begin to see the object that you want to access.



Zooming Toward Objects

3. Click that part of the object and hold down the mouse button. This action establishes the focal point at that distant location; you can either rotate the object around that point, or zoom toward it.



4. You can also use nodes in the object tree.

See Also

Using Tree Nodes to Access Objects Obscured from View
Zooming
Rotating
Hiding a Part in a Scene

Dragging Parts onto a Scene

You can use the drag-and-drop technique to add parts to an active scene display.

The nodes in the object tree function as tangible, viewable objects.

1. Drag the boundary node into the *Graphics* window.



A pop-up menu appears letting you choose which part displayer receives the part.



In the display window, the part that you added has the properties that are specified in the part displayer that you selected from the pop-up menu.



See Also

Selecting Parts by Type in the Display

Working with Graphics Windows

A Graphics window contains a tab and a display (the white space that holds the image and other widgets).

The tab is named the same as the node in the object tree. If multiple simulations are open at once, then the tab includes the name of the corresponding simulation. You can open a Graphics window by double-clicking the corresponding scene node in the tree. If the window for a particular scene has been opened, but is not active, you can double-click the relevant simulation tree node or click its tab in the *Graphics* window.



The visualization display is the space inside the Graphics window where the model is shown with various display objects such as the axes, scalar and/or vector bars, and textual legend.



In addition to the selection capabilities, the following operations let you get the most out of the Graphics window:

- Renaming parts in the scene

- Using the mouse
- Activating menus by right-clicking
- Accessing hidden objects using tree nodes
- Selecting scalars and vectors using scalar and vector bars
- Using the distance measuring tool
- Changing the view
- Moving the image around the x, y or z-axis
- Obtaining scalar data with the cell probe

See Also

Working with Displays

Renaming Parts in the Scene

This section explains how to change the name of a part in the scene.

1. To change the name of a part in the scene, click once on it in the visualization display.
   This highlights that part.
2. Right-click on it to activate the pop-up menu and select **Rename...**



3. Enter the new name of the part in the inline editor.
4. Press `<Enter>`.

   The new name appears in both the simulation tree and the label of the part (when you select it).



5. Alternatively, right-click the boundary node in the object tree and select **Rename...**



Using the Mouse in the Visualization Display

Each mouse button, sometimes combined with a keystroke, affects the view.

In the following text:

- LMB = left mouse button.
- MMB = middle mouse button.
- RMB = right mouse button.

The default configuration for mouse interaction in STAR-CCM+ is shown below. The rotate, roll, and zoom functions operate around the point where the mouse was first clicked. So in zoom, for example, the point on the model where the mouse was clicked remains at the same point on the display as the rest of the model shrinks or expands around that point.

<LMB>	Rotates the view. Moving the mouse horizontally rotates the view around the vertical axis. Moving the mouse vertically rotates the view around the horizontal axis. Moving the mouse at an angle, combines the two rotations.
<MMB>	Zooms the view in or out. Moving the mouse up, makes the view larger, moving the mouse down,

	makes the view smaller. Moving the mouse horizontally has no effect.
<RMB>	Pans the model across the screen in the direction of the mouse movement.
<Ctrl><LMB>	Rolls the view, that is, rotates it around the axis perpendicular to the screen. Moving the mouse down rotates clockwise, and moving it up rotates the view counter-clockwise. Moving the mouse in the left or right direction has no effect.

Within STAR-CCM+, it is possible to change the configuration of the mouse in order to suit your preferred working environment. You can change the mouse configuration in the *Options* dialog (**Tools > Options**) under the **Visualization** node by selecting an alternative *Interactor Style*. The default option, **STAR Default**, is described above, and the rest of the styles are described below.

Pro-E / NX

This interactor style matches the mouse configuration in the Pro-Engineer and NX CAD packages. The rotate, roll, and zoom functions operate around the point where the mouse was first clicked.

<MMB>	Rotates the view. Moving the mouse horizontally rotates the view around the vertical axis. Moving the mouse vertically rotates the view around the horizontal axis. Moving the mouse at some angle, combines the two rotations.
<Ctrl><MMB>	Zooms the view in or out. Moving the mouse up, makes the view larger, moving the mouse down, makes the view smaller. Moving the mouse horizontally has no effect.

<Shift><MMB>	Pans the model across the screen in the direction of the mouse movement.
<Ctrl><LMB>	Rolls the view, that is, rotates it around the axis perpendicular to the screen. Moving the mouse down rotates clockwise, and moving it up rotates the view counter-clockwise. Moving the mouse in the left or right directions has no effect.

SolidWorks

This interactor style matches the mouse configuration in the SolidWorks CAD package. The rotate, roll, and zoom functions operate around the point where the mouse was first clicked.

<MMB>	Rotates the view. Moving the mouse horizontally rotates the view around the vertical axis. Moving the mouse vertically rotates the view around the horizontal axis. Moving the mouse at some angle, combines the two rotations.
<Shift><MMB>	Zooms the view in or out. Moving the mouse up, makes the view larger, moving the mouse down, makes the view smaller. Moving the mouse horizontally has no effect.
<Ctrl><MMB>	Pans the model across the screen in the direction of the mouse movement.
<Ctrl><LMB>	Rolls the view, that is, rotates it around the axis perpendicular to the screen. Moving the mouse down rotates clockwise, and moving it up rotates the view counter-clockwise. Moving the mouse in the left or right direction has no effect.

Joystick

All the operations are executed simply by clicking the mouse, no mouse movement is necessary. The function of each mouse button is the same as for **STAR Default**, described above.

The rotate, roll, and zoom functions operate around the center of the model. So in zoom, for example, if the center of the model is at the right of the display, it will appear that the model is moving towards a point out to your right, rather than coming towards you, as you zoom in.

Trackball

All the operations are executed by dragging the mouse with the relevant mouse button held down. The function of each mouse button is the same as for **STAR Default**, described above.

The rotate, roll, and zoom functions operate around the center of the model. So in zoom, for example, if the center of the model is at the right of the display, it will appear that the model is moving towards a point out to your right, rather than coming towards you, as you zoom in.

3Dconnexion (Windows Only)

Currently STAR-CCM+ supports 3Dconnexion hardware only on Windows. Only basic view navigation through the controller cap/handle is supported. Its function keys and other special keys are not supported.

To use this hardware with STAR-CCM+:

1. From the main menu, select **Tools > Options**.

   The *Options* dialog appears.

2. In the **Visualization** node, activate **Enable 3Dconnexion devices**.

   When this option is active, STAR-CCM+ periodically polls to see if a 3Dconnexion device has been added. If it has, manipulating the device changes the view in the currently active scene.

3. Adjust the sensitivity of the device in *3Dconnexion sensitivity*.

   This value ranges from 0.0 to 1.0 and controls how much the view changes.

See Also

Changing a Three-Dimensional Visualization Display
Changing a Two-Dimensional Visualization Display

Right-Clicking

If you right-click inside the display, you get a pop-up menu identical to that of the node associated with that display.



A similar functionality exists for boundaries in the object tree and related objects in the scene display. The pop-up menu for a part that is selected in the visualization display contains all of the boundary-related options that exist in the corresponding node's pop-up menu.





Using Tree Nodes to Access Objects Obscured from View

This section explains how to access an object behind other objects in the visualization display.

To access an object behind other objects in the visualization display

1. Click the tree node of that object.
   The object is highlighted in the display.



2. Click the (**Reset**) button in the toolbar.
   The view focuses directly to that object, allowing you to click it for rotating, zooming, or other actions.



See Also

Accessing Objects Obscured from View
Zooming
Rotating

Using the Scalar and Vector Bars

The scalar and vector bars appear in scalar and vector scenes, respectively. Each type of bar maps values to colors.

The bar contains value indicators and a title reflecting the scalar or vector type. The Color Bar node of the corresponding displayer controls the position, colors, and other visual characteristics of the bar.

	
Scalar bar	Vector bar

When a scalar bar is used in a display of scalars on streamlines, it may be described as a “streamline bar”.

The displayer of the parent scene controls the types and ranges of scalar quantities and vector quantities. You can control scalar and vector types from pop-up menus that are activated by right-clicking the bars themselves.

The following operations can be carried out on the scalar or vector bars:

- Repositioning by dragging the bar -- change the orientation between vertical and horizontal by dragging it near a horizontal or vertical border.
- Resizing by dragging one of the edges of the bar -- the text resizes accordingly.
- Right-clicking to activate a drop-down list for the following:
  - Selecting scalar functions



- Selecting vector functions



- Selecting the individual components of certain functions, such as positions and centroids, in local coordinate systems



You can use keystrokes to search drop-down lists.

See Also

Accessing the Properties Window with the Keyboard
Displaying Vector Quantities
Displaying Scalar Quantities
Using the Color Bar

Using the Distance Measuring Tool

You can measure distances in the scene. Depending where in the scene you click, these distances can be among vertices, cells and/or the background of the scene.

1. To begin using this feature, click the (**Measure Distance**) toolbar button.
2. If you want to cancel this process after clicking the button, simply press `<Esc>`.

   You can rotate, pan, and make other standard movements with the scene while measuring. You can also choose one of the standard or stored views using the **View** menus in the Visualization toolbar. The measurement is not final until you complete the second click and release the left mouse button.
3. To complete the measurement, click again at the end of the distance.



4. Click the starting point in the scene for your measurement, and then drag the mouse in the direction you want to measure.

   A line (the "length bar") with the distance appears in the scene.



If you click near a vertex at the beginning or the end, the point that you create snaps to that vertex.



The actual distance for snapping depends on the zoom level.

If you move the mouse anywhere in the scene other than inside a cell or vertex, the distance appears in red.



Detailed information on your measurement also appears in the *Output* window.



In addition to the distance itself, this information includes:

- The coordinates of the end points.
- The type of each end point -- if the point is on a vertex (or mesh node) (`Node Point`), in a cell (`Cell Point`), or off the model (`Scene Point`).
- Distances in the x, y, and z planes.

5. The length bar remains in the display, and the toolbar button for this feature is deactivated, until you do any of the following:
   - Click anywhere in the scene.
   - Click elsewhere in the GUI.
   - Press `<Esc>`.

See Also

Restoring a View
Setting a Predefined View
Changing a Three-Dimensional Visualization Display

Changing the View

You can change the view in the scene display using keystrokes.

- `<F>` - brings the view to the front of the object.



Pressing this key again flips to the back.



Additional strokes of this key toggle between front and back.

- <S> - brings the view to the side of the object.



Pressing this key again flips to the other side.



Additional strokes of this key toggle from one side to another.

- Similarly, the <C> key flips the picture horizontally along a vertical axis in the center of the view.
- <T> - brings the view to the top of the object.



Pressing this key again flips to the bottom.



Additional strokes of this key toggle between top and bottom.

Moving the Image Around the X_ Y_ or Z-Axis

You can use the `<X>`, `<Y>`, and `<Z>` keys to move the image around the corresponding axis.

For example, `<Z>` rotates the object in a negative direction around the z-axis, while `<Shift><Z>` rotates the object in the opposite direction.



Each keystroke turns the image 10 degrees.

The shifts are in relation to the focal point of the display.



See Also

Changing a Three-Dimensional Visualization Display

Obtaining Scalar Data with the Cell Probe

You can obtain scalar data for a particular cell in an open scalar scene.

1. To use the cell probe in a scalar scene, click in the display of that scene.
2. To obtain the scalar value for a cell, hover with the pointer over the location of that cell in the scene and press the period (“.”) key.



A report appears in the *Output* window that includes the scalar value and position.



The reported position is based on the coordinate system, which is Laboratory by default. To use a different coordinate system, see Setting the Coordinate System of the View.

If you point outside the mesh (in the background of the scene display), the *Output* window simply shows the message: `No cell picked`.

Note: In other types of scenes, such as geometry, this action only results in the name of the boundary and the position.

Selecting Parts by Type in the Display

STAR-CCM+ provides a toolbar that lets you select objects in a scene with precision.

The selection tools come in two varieties: rubberband and zone. You can control each one by a filter that lets you select parts by curve, surface, or volume.

Add any volumes, surfaces, or curves that you wish to select using these tools to the scene displayers. For example, when importing CAD files, part curves are not included in the geometry displayer by default, so are not available for selection.

See Also

Dragging Parts onto a Scene

Selecting in a Rubberband

This technique lets you define a specific area in the display by clicking two points that define the opposing corners. Objects within that area are selected. You can use the selection filter to control whether this selection includes volumes, surfaces and/or curves.

1. Click the (**Rubberband Select**) button to activate this tool.
   After you click the first point in the display, the tool draws a rectangle and finishes it once you click the second point.



2. Press <Esc> to cancel the drawing process.
   The resulting selection is highlighted in the display. In the example that illustrated here, the setting is for curves in a 3D-CAD view scene.



See Also

Filtering the Selections

Selecting by Zone

This technique lets you draw a zone by clicking and dragging to create a series of connected points. Objects within that zone are selected.

You can use the selection filter to control whether this selection includes volumes, surfaces and/or curves.

1. Click the (**Zone Select**) button to activate this tool. Each click in the display adds a point that is connected by a segment to the previous point.



2. Right-click the last point to finish the zone.
3. Press <Esc> to cancel the drawing process.

   The resulting selection is highlighted in the display. In the example that illustrated here, the setting is for surfaces and curves in a 3D-CAD view scene.



See Also

Filtering the Selections

Filtering the Selections

If you have a wide variety of parts in the object being simulated, you can filter the selection based on curve, surface, or volume.

The result will be that, as you carry out mouse actions in the display, only the type of part you specify is selected.

Click the (**Selection Filter**) button. It activates a menu.



By ticking the checkbox next to an option, you can specify the type of part selected. Since selecting volumes always includes the selection of curves and surfaces, the options in the menu change their selection automatically depending on which checkbox you tick, providing the following set of possibilities:

- Volumes
- Surfaces
- Curves
- Surfaces and curves

The meaning of these terms for the specific part type varies depending on what type of display you are accessing:

- For 3D-CAD views:
  - "Volumes" are 3D-CAD bodies.
  - "Surfaces" are the faces of the 3D-CAD bodies.
  - "Curves" are the edges of the 3D-CAD faces.
- For geometry views displaying geometry parts:
  - "Volumes" are geometry parts.
  - "Surfaces" are part surfaces.
  - "Curves" are part curves.
- For views displaying regions and boundaries:
  - "Volumes" are regions.

- "Surfaces" are boundaries.
- "Curves" are feature curves.

The filter also controls how the left-click selects objects. For example, if it is set to select surfaces only, a click on a part selects only the surface.



However, if selection of volumes is specified, then a click in the exact same place selects the volume.



See Also

Selecting in a Rubberband
Selecting by Zone

Changing a Three-Dimensional Visualization Display

This section explains various methods of changing the three-dimensional visualization display.

The following operations let you manipulate a three-dimensional visualization display:

- Shifting the image by panning
- Moving the part in the visualization display around a specific axis by rotating and rolling

  (You can also rotate the image around the x, y, or z-axis using the keyboard.)
- Getting a close-up view by zooming
- Aligning the view to a part normal

It is also possible to reset the scene to its original view.

These techniques are based on the default interactor style. Within STAR-CCM+, it is possible to change the mouse configuration to suit your preferred working environment. Available styles include Pro-Engineer/NX and SolidWorks. It is also possible to specify joystick, trackball, and 3Dconnexion interactor styles.

See Also

Using the Distance Measuring Tool
Using the Mouse in the Visualization Display
Moving the Image Around the X_ Y_ or Z-Axis
Changing a Two-Dimensional Visualization Display

Panning

This section explains how to pan in a scene.

To pan the image in the visualization display, hold down the right mouse button and drag the mouse in the direction where you want to move the object.



The tiny colored square that appears while you click and drag represents the point at which you began the movement. This is helpful for knowing where you started the movement. However, the point from which you start panning has no effect on the movement; the movement is governed exclusively by how you reposition the mouse.

The image appears to change its view angle as you pan. This is because the scene's view is in the `Perspective` projection mode, which depicts three-dimensional space. In the `Parallel` projection mode, the view angle stays the same.



You can achieve the same effect with the middle mouse button if you hold down the `<Shift>` key at the same time.

Rotating

This section explains how to rotate scenes.

STAR-CCM+ uses a point on a part to rotate (left image below). If you select a point somewhere other than on a part (right image below), STAR-CCM+ uses the depth to that last point; the depth is normal to the view plane.



The scene can be rotated around various axes, depending on the direction you move the mouse.

Rotating around a Vertical Axis

Choose a point on or near the object in the visualization display, and click and drag (holding down the left mouse button) in a horizontal direction from that point.



This action rotates the object around the vertical axis. The small colored square represents the point at which you began the click-and-drag motion. The colored number near it indicates degrees of rotation.

Rotating around a Horizontal Axis

Now try clicking and dragging in a vertical direction.



The axis is horizontal.

Rotating Around an Arbitrary Axis

Click and drag at an angle.



The two directions of the axes are combined.

This technique rotates the object around the axis perpendicular to the direction of your drag. You set the location of the axis when you chose a point to begin clicking and dragging.

Rolling

This section explains how to roll scenes.

1. To roll the scene, hold down the `<Ctrl>` key, and click and drag the mouse toward you (downward).



2. Using the `<Ctrl>` key rotates the object along an axis perpendicular to your screen. The downward movement of your mouse creates a clockwise rolling motion. The movement stops when your mouse pointer goes outside the edges of the visualization display. To roll the object further, simply release the mouse button and move the mouse back to its earlier position.
3. Now click and drag in the opposite direction.



This results in a counter-clockwise roll.

If you try clicking and dragging left or right in this mode, there is no effect.

Zooming

This section explains how to zoom in a scene.

There are different ways to zoom with the mouse:

Choosing a focal point

Choose a point inside the display, then click and drag downward using the middle mouse button.



This brings your view closer to that point.

Moving the mouse straight up has the opposite effect.



Rubber banding

To zoom to a specific area in the display, click the (Rubberband Zoom) button, and then click two points that define the opposing corners of the zoomed area.



Using the mouse wheel

Turning the mouse wheel toward you brings your view closer to the visualization display.



The opposite wheel movement takes your view farther away.

The mouse pointer must be within the visualization display for this operation to work. Its particular position within the display determines the direction of this type of zooming.

See Also

Accessing Objects Obscured from View
Using Tree Nodes to Access Objects Obscured from View

Aligning to a Normal

This section explains how to align the view to the surface normal of a part.

1. In the scene display, right-click on the part and select **Align View to Surface Normal**.



The view aligns accordingly.



2. Similar functionality exists in the pop-up menu of the scene's part node: the **Align View to Part Normal** menu item.



Resetting

This section explains how to reset the view of a scene.

1. To restore the original view of the scene, click the (**Reset**) button in the toolbar.

   The <R> key behaves the same as the (**Reset View**) toolbar button.

Changing a Two-Dimensional Visualization Display

This section explains various methods of changing the two-dimensional visualization display.

A two-dimensional mesh appears in the visualization display as perfectly flat; therefore the following display operations differ somewhat from the display options for three-dimensional displays:

- Panning
- Rotating
- Zooming

Resetting the scene to its original view can be done the same as for a three-dimensional display.

These techniques are based on the default interactor style. Within STAR-CCM+, it is possible to change the mouse configuration to suit your preferred working environment. Available styles include Pro-Engineer/NX and SolidWorks. It is also possible to specify joystick, trackball, and 3Dconnexion interactor styles.

See Also

Changing a Three-Dimensional Visualization Display
Using the Mouse in the Visualization Display

Panning

This section explains how to pan in a scene.

Panning a two-dimensional image is similar to doing so with a three-dimensional one: holding down the left mouse button and dragging the mouse in the direction where you want to move the mesh.



However, there are two important differences:

- The right mouse button is also available for panning, because no rotation is possible.
- There is no change in perspective as you pan, since there is no third dimension in the view.

Rotating

A two-dimensional mesh cannot be rotated, since it has no third dimension to display. However, it is possible to roll it.

To roll a two-dimensional mesh, hold down the `<Ctrl>` key, and click and drag the mouse toward you (downward).



See Also

Accessing Objects Obscured from View
Using Tree Nodes to Access Objects Obscured from View

Zooming

This section explains how to zoom in a scene.

The zooming functionality matches that for a three-dimensional mesh.



Interacting with the Plot Display

This section explains how to interact with the plot display.

The typical plot display contains a wealth of information. Therefore, it can be difficult to read all of it from the default appearance. That is why the plot display in STAR-CCM+ includes panning and zooming tools to navigate within the plot, as well as a feature for resetting the original plot view. In addition, you can navigate the object tree to plot objects. You can also change the plot characteristics by right-clicking.

Panning and Zooming

This section explains how to pan and zoom in the plot display.

By clicking the (**Pan**) button in the toolbar, you allow the mouse to shift the plot within the confines of the axes. When you release the mouse, the values of the axes shift in response to your mouse movement.



However, the pan feature is especially useful when you zoom in on (magnify) a portion of the plot to look at smaller changes in values, and then pan to examine adjacent portions.

1. Click the (**Zoom**) button in the toolbar.

   This action activates the plot zooming feature, which operates on the rubber-banding principle.
2. Drag the mouse through the area of the plot that you want to magnify. This will show a quadrangle expanding like a rubber band.



That quadrangle bounds the zoomed area.



3. Click the **Pan** button again. As you move the plot with the mouse, you get a close-up view of other parts of the plot.



Resetting the Original Plot View

4. To restore the original view of the plot, click the (**Fit**) button in the toolbar.

Navigating the Object Tree

This section explains how to interact with the plot display using the object tree.

1. To select a variable type in a plot, simply click its label in the plot display.



The tree navigates to the corresponding node.



For XY plots, this action takes you to the **X Type** or **Y Type** nodes. For other plots, it navigates to the corresponding node under the **Axes** node.

2. Clicking any other part of the plot display navigates to the plot node itself.

Right-Clicking

If you right-click inside the plot display, you get a pop-up menu almost identical to that of the node associated with that display.



Plot Tab Reference

A plot tab consists of a plot display. The tab provides a representation of a plot node in the simulation tree and is named the same as the node.

You can create a plot tab by double-clicking the corresponding node in the tree. If the tab for a particular plot has been created, but is not active, you can double-click the relevant simulation tree node or click its tab in the *Graphics* window.



See Also

Working with Displays

Plot Display

The key elements of a plot display are shown below.



You can modify these elements (series, X and Y axes, legend, symbols and line styles) by customizing their common properties.

See Also

Customizing Plots

Using the Output Window

The *Output* window gives a precise text description of all developments within the simulation.

This output can be saved to a file, and handled like a text document in other ways.

This window is open in the STAR-CCM+ workspace by default. However, to open the *Output* window, select **Window** > **Output**.



The illustration below shows the output during the loading of an existing simulation.

```
Output - intake
Starting local server: /home/dfg/install/starccm+2.03.096/star/bin/starcc

STAR-CCM+ Version 2.03.096 (linux-x86-2.2.5/gnu4.1)
License version 14-dec-2004
Checking license file: /usr/local/STAR-CD/v3150/license/license.dat
1 copies of prostar checked out from /usr/local/STAR-CD/v3150/license/lic
Feature prostar expires in 96 days

Server::start zx6r:56026
Loading simulation database: /home/dfg/install/starccm+2.03.096/doc/tutor
Loading module: CoupledFlowModel
Loading module: KeTurbModel
Loading module: MaterialModel
Loading module: MotionModel
Database saved by:
  STAR-CCM+ Version 2.03.072 (linux-x86_64-2.2.5/gnu4.1)
  Tag date = 2006-08-21 13:35:18 GMT
Loading into:
  STAR-CCM+ Version 2.03.096 (linux-x86-2.2.5/gnu4.1) Serial
  Tag date = 2006-09-22 23:16:59 GMT

Number of partitions stored in database: 1
Simulation database load completed.
```

The example below is of running a simulation 1000 iterations.

```
Output - intake
    995   1.276528e-02 7.746027e+00 4.657106e-01 4.458370e-01 3.843451e+0
    996   1.275689e-02 7.750090e+00 4.712672e-01 4.377354e-01 3.840665e+0
    997   1.276532e-02 7.746091e+00 4.656880e-01 4.458260e-01 3.843463e+0
    998   1.275709e-02 7.750113e+00 4.712879e-01 4.377322e-01 3.840719e+0
    999   1.276538e-02 7.746105e+00 4.657168e-01 4.458107e-01 3.843474e+0
   1000   1.275703e-02 7.750091e+00 4.712908e-01 4.377354e-01 3.840699e+0
Stopping criterion Maximum Steps satisfied.
```

The following is an example of importing a mesh.

```
Output - intake
Boundary type[4] id: 5, name: "Wall-5", label: "Wall-5", type: "wall"
Boundary type[5] id: 6, name: "Wall-6", label: "Wall-6", type: "wall"
Reading cell Ids...
Reading faces...
Reading boundaries...
Distributing vertices...

Processing mesh
Finished reading /home/dfg/install/starccm+2.03.096/doc/tutorials/intake/
Reordering Reordering cells in region Fluid_1... done
Reordering faces in region Fluid_1... done
Reordering cells in region Fluid_1 2... done
```

The illustration below shows the *Output* window with separate tabs, each with output for an open simulation.

```
Output
intake   Star 7
   Mesh is topologically valid and has no negative volume cells.
-> FACE VALIDITY STATISTICS:
Minimum Face Validity: 9.842612e-01
Maximum Face Validity: 1.000000e+00
        Face Validity <  0.50          0    0.000%
0.50 <= Face Validity <  0.60          0    0.000%
0.60 <= Face Validity <  0.70          0    0.000%
0.70 <= Face Validity <  0.80          0    0.000%
0.80 <= Face Validity <  0.90          0    0.000%
0.90 <= Face Validity <  0.95          0    0.000%
```

The *Output* window displays errors in red text, as shown in the following screenshot.

```
Output - intake
Loading module: KeTurbModel
Simulation database saved by:
  STAR-CCM+ 6.05.072 (linux-x86_64-2.3.4/gnu4.5) Mon Aug 15 20:32:34 UTC 2
Loading into:
  STAR-CCM+ 8.02.004 (win64/intel12.1-r8) Thu Dec 20 17:25:06 UTC 2012 Ser
Simulation database load completed.
Started default macro:
C:\Users\dfg\AppData\Local\CD-adapco\STAR-CCM+ 8.02.004\var\journal\star47
Loading/configuring connectivity (old|new partitions: 1|1)
  Fluid : 58800 cells, 171560 faces
Configuring finished
Calculating grid flux in region Volume Mesh:Fluid
Error: field function 'User Field Function 1 (UserFieldFunction_1)' compil
Unable to compile expression "AreaMagnitude1"
 in function User Field Function 1 (UserFieldFunction_1)
```

The following example is of a warning, with text colored in black rather than the usual gray.

```
Output - intake
Simulation database saved by:
  STAR-CCM+ 6.05.072 (linux-x86_64-2.3.4/gnu4.5) Mon Aug 15 20:32:34 UTC 2011
Loading into:
  STAR-CCM+ 8.02.004 (win64/intel12.1-r8) Thu Dec 20 17:25:06 UTC 2012 Serial
Simulation database load completed.
Started default macro:
C:\Users\dfg\AppData\Local\CD-adapco\STAR-CCM+ 8.02.004\var\journal\star39747
Loading/configuring connectivity (old|new partitions: 1|1)
  Fluid : 58800 cells, 171560 faces
Configuring finished
Continuum or phase "Physics 1" is missing required models: [K-Epsilon High y+
Command: StepSimulation
   CompletedCommand: StepSimulation
   error: Server Error
```

Saving the Output to a File

You can save the contents of the *Output* window to an external file.

The automated saving approach retains the entire session's data in a `.log` file. Alternatively, you can right-click in the *Output* window to save the contents as a standard text (`.txt`) file in a location of your choice, but such a file would only retain the current contents of the *Output* window.

Automatically Retaining the Output of the Entire Session

You can save the content of the *Output* window to a file.

In the *Options* dialog, select the **STAR-CCM+ > Environment** node. Make sure that the *Log output window to file* property has its checkbox ticked.

This file will be named `output****.log`, with "`****`" representing a series of numeric characters. It can be found in the following locations:

- On Linux/Unix: `~/.star-X.XX.XXX/var/log/`, where ~ is your `home` directory and `X.XX.XXX` is the version of STAR-CCM+, for example 9.04.006
- The Windows path may look like the following:
  - On Vista: `C:\Users\username\AppData\Local\CD-adapco\STAR-CCM+ X.XX.XXX\var\log\`
  - On XP: `C:\Documents and Settings\username\Local Settings\Application Data\CD-adapco\STAR-CCM+ X.XX.XXX\var\log\`

Once you access this file, be sure to save a copy of it in a location where you keep data related to the simulation project.

Saving the Current Output as a Text File

This technique only captures what still appears in the *Output* window.

Earlier portions of the output may be lost unless you have automated saving enabled.

To save the contents of the *Output* window in a standard text file:

1. Right-click in the window and select **Save As...**



A *Save Output As* dialog will appear allowing you to navigate to where you want to save the output.



2. Click **Save** to create the text file.

Working with Output Text

You can work with the contents of the *Output* window as you would with a typical text document.

Right-clicking in the window reveals a set of options similar to those that exist in common word processors.



These options are described in the following sections.

Copying Output Text to the Clipbord

To copy material from the window to the clipboard, highlight that material:

1. Right-click in the window and select **Copy** from the pop-up menu.



This menu item is disabled (grayed out) when no text in the *Output* window is highlighted.

Finding a Term in the Output Text

To find a term in the text:

1. Right-click in the window and select **Find...** from the pop-up menu.



A separate dialog appears in which you can enter a search term.



2. Click **Find** to locate that term.

The dialog also contains a drop-down list of all your search entries from the current STAR-CCM+ session.



Finding the Next Occurrence of a Term in the Output Text

Sometimes a term appears many times in a text.

To repeat your search for a term efficiently:

1. Right-click in the window and select **Find Next** from the pop-up menu.



Finding the Previous Occurrence of a Term in the Output Text

You can search for a term backwards.

To search for a term backwards:

1. Right-click in the window and select **Find Previous** from the pop-up menu.



Filtering the Output Text

To view only the lines with a certain text string in them:

1. Right-click in the window and select **Filter...** from the pop-up menu.



In the *Filter* dialog, enter the string that you want and click **Filter**.



2. Tick the **Regular Expression** checkbox to include variations of the string.

   The window shows only those lines of text.



The dialog also contains a drop-down list of all your entries from the current STAR-CCM+ session.



3. To restore the full content of the window, right-click and select **Filter...** again.

Wrapping the Output Text

By default the text lines appear to run off the side of the *Output* window.

To make the text "wrap", or break lines according to the window width:

1. Right-click in the window and select **Wrap text** from the pop-up menu.



2. To restore the default setting, select the same menu item again.

Increasing Font Size in the Output Text

To increase the font size incrementally:

1. Right-click in the window and select **Larger Font** from the pop-up menu.



It is also possible to customize fonts.

Decreasing Font Size in the Output Text

To decrease the font size incrementally:

1. Right-click in the window and select **Smaller Font** from the pop-up menu.



It is also possible to customize fonts.

Customizing Fonts

It is possible to customize the appearance of fonts in the *Output* window.

1. Right-click in the window and select **Choose Font...** from the pop-up menu.



2. In the *Font Chooser* dialog, you can specify the font family, font style, and font size. Alternatively, you can revert any customized settings by clicking **Default Font**.
3. Click **OK** to apply your changes.



Clearing the Output Text

To remove the contents of the window for the current simulation:

1. Right-click in the window and select **Clear** from the pop-up menu.

This action results in the loss of these contents if you did not save them as a file previously.



If multiple simulations are open, the *Output* window contains multiple tabs with output for each simulation, and this action in one tab will not affect the contents of any other tab.

Closing the Output Text

To close the active text output:

1. Right-click in the window and select **Close** from the pop-up menu.

   This action closes the entire *Output* window unless there are multiple tabs in the window at the time. Additionally, this action results in the loss of the contents if you did not save them as a file previously.



2. To reopen the window, select **Window** > **Output**.

Using Keyboard Shortcuts

STAR-CCM+ accommodates keystrokes for a host of activities.

You can use the keyboard to navigate the menu bar, work with the table editor, access elements of the *Properties* window, manipulate the scene display.

Keyboard shortcuts for surface preparation are also provided in the Preparing Surfaces section.

Using the Keyboard to Navigate the Menu Bar

The keyboard can be used to access the menu bar when you hold down the `<Alt>` key and simultaneously press the underscored character of the menu button of interest. You can also use the `<Right>` and `<Left>` arrow keys to move from one menu to the next.

File

Action	Keyboard Shortcut
Create a new simulation.	`<Ctrl>+<N>` or `<Alt>+<F>+<N>`
Load an existing simulation.	`<Ctrl>+<L>` or `<Alt>+<F>+<L>`
Select a recently opened simulation file.	`<Alt>+<F>+<R>` then select with `<Up>` or `<Down>` arrow key and press `<Enter>`.
Save changes to the current simulation.	`<Alt>+<F>+<S>`
Save a copy of the current simulation file with a new name and/or in a new location.	`<Ctrl>+<S>`
Play a macro.	`<Alt>+<F>+<C>` then `<Alt>+<P>`
Replay a recently played macro.	`<Alt>+<F>+<C>+<P>+<T>` then select with `<Up>` or `<Down>` arrow key and press `<Enter>`.
Start recording a macro.	`<Alt>+<F>+<C>+<R>`
Stop recording a macro.	`<Alt>+<F>+<C>+<S>`
Save the default macro.	`<Alt>+<F>+<C>+<D>`
Export simulation data.	`<Alt>+<F>+<E>`

Action	Keyboard Shortcut
Print contents of display.	<Alt>+<F>+<P>
Disconnect from the server.	<Alt>+<F>+<D>

Edit

Action	Keyboard Shortcut
Copy the selected object.	<Alt>+<E>+<Y>
Paste the contents of the clipboard into the selected location in the object tree.	<Ctrl>+<V> or <Alt>+<E>+<P>
Delete the selected object.	<Alt>+<E>+<D>

Mesh

Action	Keyboard Shortcut
Scale the mesh.	<Alt>+<M>+<S>
Run mesh diagnostics.	<Alt>+<M>+<D>
Extract a two-dimensional mesh.	<Alt>+<M>+<2>

Solution

Action	Keyboard Shortcut
Clear the solution data of the simulation.	<Alt>+<S>+<C>
Run the simulation.	<Alt>+<S>+<R>
Step the solution.	<Alt>+<S>+<S>
Set step size.	<Alt>+<S>+<T>
Stop the solution.	<Alt>+<S>+<P>

Tools

Action	Keyboard Shortcut
Reset the workspace to default settings.	<Alt>+<T>+<R>
Show or hide a toolbar. This action toggles between showing and hiding the selected toolbar.	<Alt>+<T>+<T> then select with <Up> or <Down> arrow key and press <Enter>.

Action	Keyboard Shortcut
Switch to or from small toolbar icons.	<Alt>+<T>+<T>+<S>
Reset toolbars to default configuration.	<Alt>+<T>+<T>+<T>
Open a dialog to customize toolbars.	<Alt>+<T>+<T>+<C>
Open the *Options* dialog.	<Alt>+<T>+<O>
Access the plugins.	<Alt>+<T>+<G>

Window

Action	Keyboard Shortcut
Open the *Output* window, or make it active if it is already open.	<Ctrl>+<4> or <Alt>+<W>+<O> +<O>
Restore the default layout of the workspace.	<Alt>+<W>+<R>
Save the existing layout as the default.	<Alt>+<W>+<S>
Close the selected window (*Properties* or windows in the *Explorer* pane only).	<Ctrl>+<F4> or <Alt>+<W>+<C>
Maximize (or restore when maximized) the selected window (*Properties* or windows in the *Explorer* pane only).	<Shift>+<Esc> or <Alt>+<W> +<Z>
Undock the selected window.	<Alt>+<Shift>+<D>

Help

Action	Keyboard Shortcut
Access the online Help.	<F1>
Get information about the current version of STAR-CCM+.	<Alt>+<H>+<A>

Navigating the Table Editor with the Keyboard

For navigating among the cells in the body of the table, the following keys are supported:

Keyboard Shortcuts

Action	Keyboard Shortcut
Shift the focus from one column to the one right of it. When the focus is in the tree column, pressing this arrow opens the node of the row (assuming it is closed).	`<Right>` arrow
When the focus is in the tree column, pressing this combination opens all nodes and subnodes within the selected node of the row.	`<Ctrl>+<Right>` arrow
Shift the focus from one column to the one left of it. When the focus is in the tree column, pressing this arrow closes the node of the row (assuming it is open).	`<Left>` arrow
When the focus is in the tree column, pressing this combination closes all nodes and subnodes within the selected node of the row.	`<Ctrl>+<Left>` arrow
Shift the focus from one column to the one right of it.	`<Tab>`
Shift the focus from one column to the one left of it.	`<Shift>+<Tab>`
Shift the focus from one row to the one above it.	`<Up>` arrow
Shift the focus from one row to the one below it.	`<Down>` arrow
Shift the focus to the left cell of the row.	`<Home>`
Shift the focus to the right cell of the row.	`<End>`
Go to edit mode in a property cell.	`<Space bar>`
Go to view mode in a property cell, assuming that cell is in edit mode. This does not change the value if you have changed it. This key closes the dialog if pressed	`<Esc>`

Action	Keyboard Shortcut
when there is no property cell in edit mode.	
Display drop-down list -- when the cell of that list has the focus.	<Space bar>
Move from a drop-down list item to the one above it.	<Up> arrow
Move from a drop-down list item to the one below it.	<Down> arrow
Select a drop-down list item.	<Enter>
Exit a drop-down list without entering any change.	<Esc>
Move to the top row of the table body.	<Ctrl>+<Home>
Move to the bottom row of the table body.	<Ctrl>+<End>

To select a contiguous block of nodes, select the first node of the block using the <Up> or <Down> arrow keys, hold down the <Shift> key, then press the <Down> arrow until you reach the last node of the block.

To select a non-contiguous block of nodes, select the first node of the block using the <Up> or <Down> arrow keys, hold down the <Ctrl> key, then press the <Down> arrow until you reach the next node of the block and press the <Space> bar. Continue navigating up or down and pressing <Space> on required nodes until you are done.

Accessing the Properties Window with the Keyboard

Once the *Properties* window has the focus, the following keys are supported:

Navigating

Action	Keyboard Shortcut
Close an open category.	<Left> arrow
Open a closed category.	<Right> arrow
Move the focus to the previous property in the list, stopping at the category bar of the top list.	<Up> arrow

Action	Keyboard Shortcut
Move the focus to the next property in the list, stopping at the last item of an open category or the category bar of a closed category.	<Down> arrow
Move to the next property, rewrapping from the bottom category to the top one.	<Tab>
Move to the next property, rewrapping from the top category to the bottom one.	<Shift>+<Tab>
Put the property into edit mode if it is not read-only (with the text selected if the inline editor is a text box).	<Space bar>

In edit mode, the keys available depend on the type of inline editor.

Text Box

Action	Keyboard Shortcut
Accept the new value (subject to error checking) and go to view mode. If the error checking fails, the value reverts to the previous value.	<Enter>
Revert the property to the previous value, going to view mode, provided you have not typed <Enter>.	<Esc>

Drop-Down List

Action	Keyboard Shortcut
Open the list.	<Space bar>
Change the selection to the item above.	<Up> arrow
Change the selection to the item below.	<Down> arrow

Action	Keyboard Shortcut
Accept the list selection, reverting to view mode.	`<Enter>`
Go to view mode, not changing the value if you have changed it.	`<Esc>`

When you have the drop-down list open, pressing any letter will always be interpreted as the first letter of the required item. Repeated presses of that key will continue to cycle through the set of items that begin with that letter.

For example, suppose you open a drop-down list with scalar functions, and you want to select `Temperature`. When you press `<T>`, you may highlight `Tangential Velocity` (depending on the simulation). If you try pressing `<T>` followed by `<E>`, you may highlight `Tangential Velocity` and then `Effective Conductivity`. Instead, keep pressing `<T>` until `Temperature` is highlighted. Then press `<Enter>` to complete the selection.

Mini-Dialog

There is no keyboard support for mini-dialogs.

See Also

Navigating Around the Properties Window
Types of Inline Editors
Using the Scalar and Vector Bars

Using Keyboard Shortcuts in the Scene Display

Keyboard shortcuts allow you to work with the scene display as follows:

Changing the View

Action	Keyboard Shortcut
Bring the view to the front of the object. Flip it to the back with a second press. Additional strokes of this key toggle between front and back.	`<F>`
Bring the view to the side of the object. Flip it to the other side with	`<S>`

Action	Keyboard Shortcut
a second press. Additional strokes of this key toggle between front and back.	
Flip the picture horizontally along a vertical axis in the center of the view.	<C>
Bring the view to the top of the object. Flip it to the bottom with a second press. Additional strokes of this key toggle between front and back.	<T>

Moving the Image on an Axis

Action	Keyboard Shortcut
Move the image in a negative direction around the x-axis. Each keystroke turns the image 10 degrees. The shifts are in relation to the focal point of the display.	<X>
Move the image in a positive direction around the x-axis.	<Shift>+<X>
Move the image in a negative direction around the y-axis.	<Y>
Move the image in a positive direction around the y-axis.	<Shift>+<Y>
Move the image in a negative direction around the z-axis.	<Z>
Move the image in a positive direction around the z-axis.	<Shift>+<Z>

Obtaining Scalar Data for a Cell

Action	Keyboard Shortcut
Determine the scalar value for a cell when hovering with the pointer over the location (scalar scenes only).	<.>

Resetting the View

Action	Keyboard Shortcut
Restore the original view of the scene.	`<R>`

Setting the Focal Point

Action	Keyboard Shortcut
Set the focal point of the camera to the position where you hover with the mouse.	`<H>`

Working With Clients and Servers

This section describes the concept of the client-server architecture used in STAR-CCM+.

This architecture lets you start a simulation running, disconnect the client from the running server, and then connect again at a later time. This is useful for saving resources on your machine by shutting down the client as you proceed with other work, while the server continues to solve your problem. An alternate technique for doing this involves moving the client from interactive to batch mode. You can also load a simulation without starting the client, and display diagnostic messages from the server and client.

Detailed information is also provided about remote servers, running the client in batch, and the server processes manager.

You can also configure STAR-CCM+ to shutdown the server when it does not detect any activity for a set amount of time.

What Is Client-Server?

STAR-CCM+ uses a client-server architecture in which:

- The simulation objects are created and solved on the server.
- The workspace views these objects through the client.



Multiple clients can connect to a server at the same time, allowing you to check the progress of a batch job, and also to permit collaboration.

The client and the server are two distinct processes, explained in the following sections.

See Also

What Is a Client?
What Is a Server?
What Are Simulations?
What Is the Client Workspace?
Connecting from Multiple Clients
Running the Client in Batch

Client and Server on Different Machines

The client and the server processes can run on the same machine, or they can run on remote machines, and even different architectures. For example, the client can be run on your Windows workstation while the server is run on a Linux cluster. Ideally, the client process should be run

on your local machine, which should have hardware graphics acceleration available, while the server should run on a machine or machines with the computing resources suitable for your simulation.

This architecture also allows you to break and make a connection between the client and the server processes, without disrupting the server calculations. So you could start a job from a Windows client, disconnect, and then reconnect the next morning from a Linux client to check on progress.

Typically the client will run in GUI mode, where you have full richness of a modern windows-based user interface, but even if you run STAR-CCM+ in batch mode, a client process will still be started. In this situation the client loads the batch script (Java file), has it interpreted by the JRE (Java Runtime Engine) and passes the interpreted commands to the server.

One server process node, labeled in the form user@machine path, will be displayed for each serial server process or parallel master process.

Note that if usernames are not the same when the client is on one machine and the server is on another, you should either:

- enable the collaboration option
- create a new user account on one of the machines to match the other machine

See Also

Working with Remote Servers
What Is a Parallel Server?
Connecting to Another User's Server

What Is a Client?

The client is the part of the client-server architecture through which you launch a simulation and work with it, typically through the STAR-CCM+ workspace.

You can see client processes running on your computer (using `ps` in Linux or the Task Manager in Windows) when a simulation is open:

- In Linux every thread is listed, and the STAR-CCM+ process names vary depending on the particular operating system. The client processes typically include `java` in their listings.



- In Windows, depending on which type of client you are running, the batch process is named `starccm+.exe`, while the GUI-based process has the name `starccmw.exe`.



Tasks performed by the client include starting the simulation, inputting commands and manipulating views for analysis. However, it is possible to load a simulation without starting the client.

See Also

What Is Client-Server?
What Is the Client Workspace?
Starting a Server without a Client

Running the Client in Batch

Some computer resources do not allow you to use the strong interactive capabilities of STAR-CCM+, or maybe you need to run a long job that does not need continuous monitoring in the client workspace.

In these cases you can use a batch mode of operation that executes a series of commands from a macro file without user interaction. An alternative, however, is the ability to move the client from interactive to batch mode.

Running in batch mode essentially means running the client in batch mode, passing to it the name of a macro script that contains all the commands necessary to load a mesh, set properties, run a number of iterations and render images of scenes and/or plots. In this mode, the non-interactive client compiles and runs the script, passing on all the information and instructions to the server.

The instructions for recording, editing and running a macro file are found under scripting the application. Similar macro recording techniques can be used to set up a large job on a small machine that would not normally handle the large job. These batch jobs can also be submitted to batch management systems using additional command line parameters.

A batch job can also be accessed while in progress, as multiple clients can connect to the same server. By connecting through an additional client, you can perform interactive tasks such as opening scenes and adjusting solver parameters. You can use the auto save feature to save automatically at the end of a run.

It is also possible to generate a summary report of a batch simulation.

See Also

Closing the Server Automatically Upon Solver Completion
Using Batch Management Systems
Connecting from Multiple Clients
Generating a Batch Summary Report

Setting up a Large Job on a Small Machine

It is possible to prepare the batch script for a large problem that would not necessarily fit onto your computer. This would typically happen if your machine did not have enough memory for a large mesh.

In this case then set up a smaller, similar mesh and record the macro using the smaller mesh. Then edit the script and change the name of the mesh file that is imported to represent the original problem. For example, the code in question may look something like:

```
getSimulation().getImportManager().importFile(resolvePat
h("smallMesh.ccm"));
```

Simply change the file name in quotes in addition to any other changes you were needing to make on that file.

Generating a Batch Summary Report

A summary report of a batch simulation can be obtained by providing the following command line option:

```
-batch-report
```

The report in `.html` format provides details of the system the simulation was run on, as well as system information about the simulation itself including elapsed times. It will also contain a copy of the STAR-CCM+ console output written during the simulation.

When you enter the line command, the report will be generated in your default browser.


Mozilla Firefox
File Edit View History Bookmarks Tools Help
Batch Report:
Session Summary
Date Aug 19, 2009 2:17:34 PM
Simulation /home/ncr /sims/suprBump@00166.sim
File size 0.68 MB
Number of Partitions 1
Number of Restored Partitions 3
Software Summary
Version BuildArch: linux-x86_64-2.3.4
BuildEnv: gnu4.3
ReleaseDate: Fri Jul 31 22:29:48 UTC 2009
ReleaseNumber: 4.05.061
Hardware Summary
Number of Hosts 1
Host Name dougie.cdnany.com
Chip Intel(R) Xeon(R) CPU X5570 @ 2.93GHz
Number of CPUs 2
Number of Cores per CPU 4
Clock Speed (MHz) 2926
Cache (KBytes) 8192
Ram (MBytes) 12003.0
Operating System Linuxversion2.6.18-128.2.1.el5(mockbuild@hs20 (gccversion4.1.220080704(RedHat4.1.2-44))#1S glibcversion(glibc 2.5)
glibc Version glibc 2.5
Process ID's 5719
Server Summary
Done

What Is a Server?

The server is the part of the client-server architecture that implements your commands in the simulation, such as importing data or running the solver.

It is represented by a node that has its own properties and pop-up menu.



The *Servers* window displays a list of processes on the local network. These processes are also listed in your computer's system. The *Servers* window is closed by default. To open it, select **Window > Servers** from the top menus.

What Is a Local Network?

For purposes of using STAR-CCM+, a local network is all machines accessible without crossing a router, subject to any firewalls blocking the server locator port (8406).

See Also

What Is Client-Server?

Server Process Properties
Server Process Pop-Up Menu
Which Server Processes Are Displayed in the Servers Window?
Viewing Server Processes in Your Computer's System

Which Server Processes Are Displayed in the Servers Window?

The *Servers* window is populated by sending a broadcast to the local network.

Servers that respond are listed in the window. Normally each of these servers needs a matching STAR-CCM+ version number and user name.



Viewing Server Processes in Your Computer's System

You can see server processes running on your computer (using `ps` in Linux or the Task Manager in Windows):

- In Linux every thread is listed, and the STAR-CCM+ process names vary depending on the particular operating system. As a general rule, the server threads are listed as `star-ccm+`.



- In Windows, the server process is named `star-ccm+.exe`.



Server Process Properties

The server properties are listed in the table below:

Host	The name of the computer where the STAR-CCM+ server process is running (read-only).
Port	The network port on which the STAR-CCM+ server process is listening for clients (read-only).
Compute Processes	Indicates the number of processes running for the selected STAR-CCM+ server (read-only).
File	The path name of the simulation file (read-only). This value may be blank if the server process is for a new, unsaved simulation.
Clients	The user name and computer IP address of each connected client computer. This read-only value may be blank if no client is currently connected to the server.
User	The name of the user who started the STAR-CCM+ server process (read-only).

Server Process Pop-Up Menu

The pop-up menu items are listed in the table below:

Connect	Executes the connect to server action to connect this client to the selected server process. More than one client can be connected to a given server process.
Properties	Displays the properties of the object in a separate stay-on-top modeless dialog.

Disconnecting from the Server

Let's assume that you have started a simulation locally on your machine.

1. First of all, open the *Servers* window by selecting **Window > Servers** from the top menus and see your simulation listed in the window.



2. Reselect the simulation window in the *Explorer* pane, and select the **File > Disconnect** menu item.



This will disconnect the client from the server, but both will still be running, although you will no longer see the simulation window in the client. Also, if all clients are disconnected from a running server, no graphics scenes will be saved to files; you will see a warning dialog to this effect when you select this menu item.



To close the client and leave the server running, select the **File > Exit** menu item after you disconnect, but you could leave the client running (to open another simulation, for example). An alternative, however, is moving the client from interactive to batch mode.

To shut down the server as well as the client, simply use the **File > Exit** menu item while the server is still connected to the client.

Connecting to the Server

You can reconnect your client to your own server after having disconnected or started without a client. It is also possible to connect to another user's server if necessary. The results of this action depend on the trigger:

- **File > Connect to Server...** - this opens the connect to server dialog to allow manual specification of the host machine that it is running on and the port number that it is listening on.
- Right-click a server in the *Servers* window and select **Connect** - this attempts to connect directly to a given process based on the properties of the currently selected object.

In addition, multiple clients can connect to the server. Clients can be launched from the command line.

Reconnecting to Your Own Server

When you disconnect or load a simulation without the client, then the server process will be running but you will not be able to see what the server is doing or be able to change any properties. To connect to a running simulation, launch the workspace and then open the *Servers* window by selecting **Window > Servers** from the top menus; the window will appear in the *Explorer* pane. This lists all servers (simulations) that are running on the local network.



Selecting the node in the tree will display the server process properties in the *Properties* window (the bottom left of the workspace if it has not been customized).



Connect by right-clicking the server process node in the Servers tree and selecting **Connect** from the pop-up menu.



The simulation will opened in a new simulation tree in the *Explorer* pane. You will now be able to work with the simulation as though you were sitting at your own machine and had never disconnected.

If you select **Connect** when you are already connected to that server process, the window of the simulation tree will simply be brought to the foreground.

Connecting to Another User's Server

The principles shown below can be used for any serial or parallel simulation.

By default, only the user who starts a server is allowed to connect to it. To allow other users to connect to your server, you must add the option `-collab` to the startup syntax in either of the standard methods:

- Including it in the command line:

```
% starccm+ -server -collab [portnumber] file.sim
```

- Ticking the **Enable Collaboration** checkbox, which is near the bottom of the workspace dialog for creating a new simulation or loading an existing one:



The option `-collab` will appear in the dialog's command line automatically.

1. Launch the STAR-CCM+ client if it is not already running. Then open the *Servers* window by selecting **Window > Servers** from the top menus; the window will appear in the *Explorer* pane. Assuming there were no simulations running on the machine hosting this client, the servers tree would be empty.



2. Select the *User* property and enter the user name into the inline editor.



3. As you hit `<Enter>` the client will search the network for server processes running under that user name.



4. Right-click the server process node in the **Servers** tree and select **Connect**.



The simulation will be opened in a new simulation tree in the *Explorer* pane.

Connecting from Multiple Clients

You can check the progress of a batch job easily, since more than one client can connect to a server at the same time. When this is the case, each client has a connection to the same server and that single server has a single connection to the file. When anything changes on the server, that change is reflected on all clients. Some actions, such as selecting entities during mesh repair, only take place on the client and will not be seen on the other clients.

Suppose two clients, Client 1 (Linux) and Client 2 (Windows), are connected to a simulation with a velocity of `5 m/s`. Client 1 changes the velocity to `10` without saving.



Client 2 sees this change if that node is open.



This change would be immediately visible to Client 2 regardless of when Client 2 connected to the server.

Also, if either client saves, the new velocity value of `10` is stored in the simulation file.

A similar real-time connection works with scene displays and XY plots. In this example, Client 1 and Client 2 have the same scene open.



Client 1	Client 2

Client 1 rotates the scene, and Client 2 sees the scene rotate.

	
Client 1	Client 2

Similar behavior will be seen for XY plots.

See Also

Running the Macro in Batch Mode

The Connect to Server Dialog

The *Connect to Server* dialog lets you specify the host and port number of a server.

You can also connect using an ssh tunnel.



The basic connection features of this dialog are the following:

- *Host* —The host name of the machine where the server process is running.
- *Port* —The port number that the server process is listening on. When a STAR-CCM+ server process is started, it reports the port number that it is listening on. In the following example, the server is started on the host cindy and is listening on the port 47828:

```
Server::start -host cindy:47828
```

- **Scan Host** button —This button checks for STAR-CCM+ server processes on the specified host, and indicates how many it finds.

The dialog includes a scroll-down list of past client-server connections, allowing you to access a host multiple times while entering its information only once.



If an unknown host or port number has been entered into the dialog the following message is displayed.



If a server process cannot be found listening on the specified port the following error is displayed.



Clicking the **Show Details** button provides more information.



Connecting with an SSH Tunnel

You can connect to a running server through an ssh tunnel. This is necessary when the client machine cannot see the server host, but you have ssh permissions to another machine which can see the server host.

The most common example of this is when the server is running on the nodes of a cluster, you want to run the client on your local workstation, and your workstation can see the head-node of the cluster but not the nodes used to run jobs. Therefore you must connect using an ssh tunnel through the head-node.



The ssh connection features of this dialog are as follows:

- *Connect through ssh tunnel* — Tick this checkbox to enable the ssh tunnel capability.
- *SSH Tunnel Host* — This is the host through which you wish to tunnel. In a cluster situation this would be the cluster head-node.
- *SSH Tunnel Host Username* — Your user name on the ssh host
- *SSH Options* — Default settings are provided but this field lets you edit them.

The preconditions for use of this feature are that you must be able make an ssh connection to the ssh host without password authentication or any other interactive input. (Typically you must have made such a connection to the ssh host from your client workstation to ensure the host is included in the ssh `known_hosts` file.)

If you do not connect properly STAR-CCM+ generates an error message. The example shown below involves an incorrect user name.


Information
ssh exited with error : 255
Permission denied (publickey,gssapi-with-mic,password).
Ensure that you are able to ssh to the tunnel host without password authentification or any other interactive input.
ssh command : ssh -o PasswordAuthentication=no -o StrictHostKeyChecking=yes -N -L 57615:bobby:47827 -l wronguser filmore
OK

Connecting a Client from the Command Line

When STAR-CCM+ starts a new server, it prints the hostname and port number in the *Output* window in a form that matches the syntax required by the STAR-CCM+ launch command.



Using the output shown above, the command to launch a client that connects to this server would be:

```
starccm+ -host scorpion:47828
```

The same output will be provided when a server is launched remotely from the command line; it is then convenient to copy and paste the relevant portion in order to launch a local client.

Configuring the Server to Shutdown When Idle

You can configure STAR-CCM+ to release a license automatically (by shutting down the server) when it detects no activity for a set amount of time.

You control the automatic server shutdown through a configuration file, `sitesettings.props`, that was added to your STAR-CCM+ installation directory during the installation process. By default, the STAR-CCM+ server does not shutdown automatically. If no configuration file is found, the default behavior stands and no timeout occurs.

When the automatic server shutdown feature is active, the server issues a message in the output window upon startup. If you set the timeout value to two hours, the message reads:

```
Server idle time limit is 2 hours.
```

STAR-CCM+ checks for the following activity on the server every five minutes:

- reading/saving files
- visualization requests
- editing properties
- meshing
- solving

If STAR-CCM+ detects no activity, and the timeout time is approaching, the server issues a warning message in the client output window. For any timeout value less than two hours, the idle warning is issued at 90% of the timeout value. For timeout values over two hours, the idle warning is issued one hour before the timeout.

If no activity is detected after the warning message, the server shuts down at the set timeout value, releasing the license. The client, however, remains open. A copy of the simulation file, named `[filename]@idle.sim`, is saved in your working directory.

Note: If you set the timeout value to 5 minutes or less, the actual server shutdown takes place 10 minutes after the warning message appears. This behavior is designed to give you enough time to react to the warning message before the server shuts down.

If multiple clients are connected to the same server, the warning message is sent to all clients, and only one client needs to respond to reset the timeout. Performing any action that invokes the server resets the timeout counter.

Automatic Server Shutdown in Batch Mode

If no client is connected to the server, the batch job stops by itself, closes the server, and automatically releases the license. If, however, one or more clients are connected to a running batch job, the server does not close down, and the license remains checked out. To avoid this issue, set up the automatic server shutdown feature. See Setting Up the Server to Shutdown Automatically.

Setting Up the Server to Shutdown Automatically

In order to set up the automatic server shutdown, you need to be able to access the STAR-CCM+ installation directory. If STAR-CCM+ is installed on a network location, contact your system administrator and ask them to perform the changes to the configuration file.

To set up the automatic server shutdown:

1. Navigate to your STAR-CCM+ installation directory:
   - For Windows: [starccmp_installation]\config
   - For Linux: [starccmp_installation]/star/config
2. Open the `sitesettings.props` file in a text editor of your choice.

An example `sitesettings.props` file is shown below:

```
# STAR-CCM+ Site Settings
#Mon Jul  2 15:45:59 EDT 2012
ServerIdleTimeout=0
```

3. Enter the timeout value (in hours) next to the `ServerIdleTimeout=` argument. If you leave this field as `0`, the automatic server shutdown feature is disabled.
4. Save the configuration file.

Closing the Server Automatically Upon Solver Completion

Often you may use the interactive client to monitor the mesh generation, the early stages of the solution, and to ensure that everything is set up correctly with no sign of problems. At this point, you may decide to leave the case running overnight, or over the weekend. You can instruct STAR-CCM+ to close the server when the solver is finished. This feature is particularly useful if you are using Power-On-Demand (PoD) licensing, as you are being charged for the duration of time the server is active, even though you may not be actively using it.

To instruct STAR-CCM+ to automatically shut down the server when the solver is finished:

1. Click the **X** button next to the simulation in the *Explorer* pane.

The *Close Options* dialog appears.



The following options are available:

2. **Wait and Close** - this option leaves the server running, with the interactive client open and connected, until all convergence criteria are met. When the convergence criteria are met, the server will save the simulation file and close automatically. The *Output* window in the client will remain visible. This option is ideally suited to running an analysis using PoD licensing.
3. **Disconnect** - this option disconnects the interactive client from the server. The solver continues running and the server will remain active indefinitely until you manually close it. You can reconnect to the server and close it.

Note that the above dialog and options are only available when you close the simulation while the solver is running.

Working with Remote Servers

The client-server model employed in STAR-CCM+ allows you to start the workspace on a different machine to the server.

This is particularly useful if you wish to set up, and monitor a simulation from your workstation but let another machine compute the solution.

Separate sections cover the specifics of starting the server remotely in:

- Serial with the workspace
- Serial with the command line
- Parallel with the workspace
- Parallel with the command line

This section describes the requirements for starting the server remotely, along with the following issues:

- Running remote servers using Windows
- Dealing with a protocol failure
- The general use of alternative remote shell programs to launch remote servers

Clients may be launched from the command line that connect to specified remote servers.

What Are the Requirements for Running Remotely?

Starting remote servers requires you to have remote shell access to the remote server machine without password prompting and the `starccm+` command must be available on each of the remote machines.

This can easily be checked using

```
rsh "host" which starccm+
```

If a permission problem is encountered, you may have to add entries into their `$HOME/.rhosts` file.

If remote shell permission problems persist after ensuring correct entries in the `$HOME/.rhosts` file you should ensure that your user name has fewer than twelve characters.

If problems continue assistance should be sought from the network administrator.

Running Remote Servers Using Windows

Windows systems do not provide server-side support for remote shell access, so starting remote servers where the remote machine is running Windows is not supported. However, starting remote servers from a client running on a Windows machine is supported if an appropriate command (such as rsh or ssh, for example) is available for STAR-CCM+ to use. Before STAR-CCM+ is launched, you must ensure that the remote shell program on Windows will connect to the server without requiring further action by you. Specifically, the connection must be such that the remote shell connection will not prompt for additional inputs.

Dealing With a Protocol Failure

On some Unix based systems there is a limit to the number of open remote shell sessions. If this limit is exceeded, the system may report a protocol failure when starting a remote server or a parallel server. This problem is usually the result of defunct rsh processes and you should first ensure all such processes are killed.

If the problem persists the administrator should ensure the limit on the number of remote processes is not set too low. On Linux systems the file `/etc/inetd.conf` can be modified to allow more rsh processes per minute. For example,

```
shell stream tcp nowait root /etc/tcpd2 in.rshd
```

to

```
shell stream tcp nowait.200 root /etc/tcpd2 in.rshd
```

Using Alternative Remote Shell Programs

By default STAR-CCM+ uses `rsh` as the remote shell program used to launch remote servers. Alternative remote shell programs may be used with STAR-CCM+ by specifying the remote shell on the command line using:

```
-rsh <remote shell command>
```

The remote shell specified must conform to the syntax:

```
“remote shell command” <HOST> -l <USER>
```

where `HOST` is the remote machine to which the remote shell will connect and the username is specified using the `-l` syntax.

Before STAR-CCM+ is launched, you must ensure that the remote shell program will connect to the server without requiring further action by you. Specifically, the user name and password on both systems must be the same so that the remote shell connection will not prompt you for additional inputs.

The ssh Program

For example, `ssh` can be used with STAR-CCM+ by specifying it as:

```
-rsh /usr/bin/ssh
```

You must test the authorization in order to initialize the relationship between the host and the remote host before running STAR-CCM+:

```
ssh remote-host ls -lsa
```

Contact your system administrator if the automatic authorization is not working.

Starting a Server without a Client

You can open an existing simulation in a server without explicitly starting the client.

In this scenario, the server will start and the simulation file will be opened and the server will simply wait for a client to connect to it. This may be useful if you want to start a remote server from a Windows machine for example.

To open a simulation on a server without launching the client, issue the following command:

```
% starccm+ -server -port [portnumber] file.sim
```

The server will use the port specified via the `-port` option (default `47827`) or the first open port after the specified port, until it has tried the number of ports specified via `-portrange` (default 100). If it cannot find an open port in that range (default `47827` - `47927`), the server will print a message and exit.

Displaying Diagnostic Messages from the Server and Client

To troubleshoot client-server connections, the server locator allows you to gather diagnostic information from the client and server in one place.

To make the server locator available, select **Tools** > **Options**. This activates the *Options* dialog.



Open the **STAR-CCM+** node, then select the **Environment** node to access its properties, which include the *Server locator diagnostics* expert property.



Activate the checkbox of this property to have messages generated by the client and server. The client and any connected server will print messages related to the server locator in the *Output* window.

For complete information about client-server activity, use this feature by following these steps:

- Start a server with a line command:

```
% starccm+ -server -cout file.sim
```

- Launch the workspace and connect to the server.
- Activate the server locator option. After a few seconds diagnostic information from the client and server will be generated.

Setting Network Binding Order for Failed Visualization Connections

When sockets between client and server are opened Java and C++ use slightly different rules for picking up the network interface, so on some machine with multiple interfaces the Visualization socket fails to open. This problem can be overcome by ensuring that your default LAN interface is at the top of the network adapter binding order on your machine. To do this on Windows:

1. Open the Network Connections folder. To do this: click **Start** on the desktop, select **Run...**, enter `ncpa.cpl` in the *Run* dialog and click **OK**.
2. Press the `<ALT>` key to make the menu show up on the *Network Connections* folder. Select **Advanced > Advanced Settings...** Enter the Administrator password if prompted.
3. Click the *Adapters and Bindings* tab.
4. Use the arrow buttons to move your Local Area Connection to the top of the Connections list.
5. Click **OK** to close the dialog.

This should resolve issues with the Visualization connection error.

The Server Processes Manager

All of the server process nodes are under a **Servers** manager node that has its own properties.



The *Servers* window is closed by default. To open it, select **Window** > **Servers** from the top menus.

Server Process Manager Properties

Refresh Time	This is the upper limit on time (in milliseconds) that may elapse before STAR-CCM+ refreshes the list of processes. The list also gets updated in response to events such as the opening or closing of a simulation. However, the window must be in the foreground to be updated.
User	Lets you change the user name to view STAR-CCM+ server processes launched by another user. This may be useful if your user name on a remote computer differs from your user name on the local computer, or if you want to see processes being run by others. If left blank, the user name of the current user is used. Multiple user names may be separated via commas, and all user processes can be displayed by entering `*`.

Expert Properties

Time to Live	Essentially, this is the number of routers that the broadcast message will be allowed to jump when searching for servers on the network. However, the accuracy of this number depends on how routers on the internal network are configured.

Working With Parallel Servers

A parallel computation is one where the work is computed by more than one process.

In distributing the parallel processes across multiple cores, the time that is taken to complete the work can be reduced. It is highly desirable to compute CFD solutions in parallel because of the computationally intensive work that is required to acquire converged solutions. STAR-CCM+ can be run in parallel mode providing a scalable solution to solving problems faster and to attempting larger problems.

The core of this functionality is based upon a data parallelism paradigm in which the mesh is distributed across a number of processes. Each process applies numerical methods identical to those used in a serial solution. Communication between the processes ensures consistency with a serial solution.

This section describes the following:

What Is a Parallel Server?

A parallel server is a STAR-CCM+ server that is set to run on multiple cores or a cluster of machines when it is first launched.

Launching a server in parallel is a transparent process in STAR-CCM+. You can launch STAR-CCM+ in parallel either through the command line or through the client GUI.

The parallel server is composed of N processes. Each process computes an equal share of the computational workload. The client only ever communicates with the first process in the parallel server which is referred to as the master process. This master process participates equally in the computational workload.



Processes communicate with each other by passing messages. These messages conform to a programming standard known as the Message Passing Interface (MPI). STAR-CCM+ is designed to work with different implementations of MPI, such as IBM Platform MPI, Microsoft MPI, and MPICH2. One or more of these implementations have been installed on your system along with STAR-CCM+. You may also have access to vendor-specific implementations, depending on your system. A command line option, `-mpi`, allows you to select a particular implementation when you launch the parallel server.

What Are the Requirements to Run in Parallel?

In order to run in parallel, STAR-CCM+ must be installed on each of the machines you wish to run on. On Windows machines, make sure that MPI is installed on each machine, and that the STAR-CCM+ installation path is the same on all the machines. On Unix machines, make sure that you have `rsh` access to each machine.

Currently, STAR-CCM+ only supports homogeneous parallel servers so, for example, you cannot have some processes running on a Windows machine, and other processes running on Linux machines. The client process does not need to be on the same architecture as the server processes.

Also make sure that you have the correct number and type of licenses that are required for a parallel server.

Hyper-Threading

Generally, for best performance, it is recommended that you turn off any features that artificially increase the number of cores on a processor, such as hyper-threading. In situations where hyper-threading is turned on to benefit other applications, you should generally avoid loading more processes than there are physical cores.

In addition to turning off hyper-threading, it is recommended that you set your BIOS settings to favor performance rather than power saving. Almost all heavily parallelized applications suffer performance problems if the CPU frequencies are spun up and down. In most parallel applications, and STAR-CCM+ in particular, when one core is operating at a reduced frequency, all other cores are running at lower frequencies as well.

See Also

What Is Client-Server?
Selecting the MPI Implementation

Selecting the MPI Implementation

Regardless of whether you use a local multi-core workstation, remote machines, or a remote cluster to run your parallel server, all approaches require an implementation of the Message Passing Interface (MPI).

STAR-CCM+ uses a default implementation of MPI depending on the platform you are working with:

- For Windows and Linux, the default MPI is IBM Platform MPI (supplied with STAR-CCM+).
- For Windows HPC Server 2008, the default MPI is Microsoft MPI (supplied with the Operating System).

You can select an alternative implementation of MPI using the `-mpi` argument on the command line. For example:

```
-mpi mpich2
```

selects the MPICH2 distribution on Windows. Command line options for supported distributions are given in the tables below.

Windows

MPI	Supplied with STAR-CCM+	Command line option
IBM Platform MPI (default)	Yes	`platform`
MPICH2	Yes	`mpich2`
Intel MPI	Yes	`intel`
Microsoft MPI (HPC Server 2008)	No	`msmpi`

Linux

MPI	Supplied with STAR-CCM+	Command line option
IBM Platform MPI (default)	Yes	`platform`
Open MPI 1.4.3	Yes	`openmpi`
Intel MPI	Yes	`intel`
SGI MPI	No	`sgi`

MPI	Supplied with STAR-CCM+	Command line option
Cray MPI	No	`crayxt`

Some notes on these distributions are as follows:

- IBM Platform MPI is a commercial product and is considered one of the best MPI distributions on the market. It supports many different Interconnects such as Infiniband (using different drivers OFED, QLogic, PSM) and Myrinet. IBM Platform MPI uses the fastest network channel that it detects without any use specification.
- Open MPI is an actively developed and widely used open source MPI distribution. Open MPI supports a range of interconnects. However, due to limited access to the full range, STAR-CCM+ provides a mechanism to use locally installed Open MPI installations that are built to support specific hardware. The version of the local Open MPI must be 1.4.3 or later.
- Intel MPI is a mature MPI that has become more competitive in recent years. Resident memory usage with Intel MPI is 40-50% less compared to other MPIs.
- SGI MPI is an MPI distribution that is optimized for SGI architectures.
- Cray MPI is required when running on Cray clusters.

Using Open MPI

Open MPI is an actively developed and widely used open source MPI distribution.

The Open MPI libraries are specified at runtime using one of the following commands (note that these commands perform the same action):

```
-mpidriver openmpi
```

or

```
-mpi openmpi
```

The Open MPI distribution that is packaged with STAR-CCM+ has support for shared memory, TCP, and OpenIB interfaces. If you would like to take advantage of Open MPI features that are not supported in the packaged distribution, you can use a locally installed Open MPI

distribution that supports your specific hardware. The version of the local Open MPI must be 1.4.3 or later.

Set the environment variable `OPENMPI_DIR` to the local installation of Open MPI. An example is shown below:

```
setenv OPENMPI_DIR /opt/openmpi
```

Open MPI command line options can be passed into STAR-CCM+ using any of the following commands:

```
-mpiflags <mpirun options>
```

or

```
-mpidriver openmpi:<mpirun options>
```

or

```
-mpi openmpi:<mpirun options>
```

Open MPI Runtime Dependencies

The Open MPI packaged with STAR-CCM+ has a runtime dependency on the rdmacm communication manager. If this library is not installed on your system, warning messages appear in the console window at runtime.

This runtime library is usually installed by your network fabric. If this is not installed, your system administrator can configure your system using the `yum` command:

```
yum install librdmacm.x86_64
```

Using Cray MPI

STAR-CCM+ supports the Cray Extreme Scaling Mode through Cray's optimized MPI libraries on Cray's XT/XE machines.

This MPI takes full advantage of Cray's optimizations.

The following prerequisites must be satisfied in order to run STAR-CCM+ on Cray XT/XE machines:

- STAR-CCM+ must be installed on the Lustre file system.
- The Data Virtualization Service (DVS) must be running on the cluster so that the compute nodes can access both the STAR-CCM+ installation and a fully functional Linux IOS environment. (Without

DVS running the compute node's Cray Linux Environment only provides a subset of Linux functions which do not satisfy the STAR-CCM+ runtime requirements).

The native Cray MPI libraries are specified at runtime with the command line options:

```
-mpidriver crayxt
```

or

```
-mpi crayxt
```

By default STAR-CCM+ attempts to locate the native MPI libraries in the default install locations, these locations are:

- Cray XE6 (Gemini) - `/opt/cray/mpt/default/xt/gemini/mpich2-gnu`
- Cray X4 or X5 (SeaStar) - `/opt/cray/mpt/default/xt/seastar/mpich2-gnu`

You can specify the location of the native MPI libraries using an environment variable, for example:

```
setenv CRAYXT_HOME /opt/cray/mpt/default/xt/seastar/
mpich2-gnu
```

The `aprun` command is similar to the `mpirun` command used in alternative MPI distributions, and as such, more `aprun` arguments can be passed into STAR-CCM+ using the `-mppflags` command-line argument:

```
-mppflags "-N 12 -S 3"
```

Jobs are submitted to the Cray XT/XE machine using the PBS batch scheduling system.

An example of a PBS script file `Jstar` for launching a STAR-CCM+ job is provided below:

```
------------------Beginning of
Jstar---------------------
#!/bin/csh

#

#PBS -N starccm+

#PBS -l walltime=00:20:00

#PBS -j eo
```

```
#PBS -o $PBS_O_WORKDIR

#PBS -V



cd $PBS_O_WORKDIR



# setup STAR-CCM+ directory

setenv STARCCM_DIR /lus/scratch/STAR-CCM+/starccm
+7.02.004/star/bin



# add STARCCM_DIR to PATH

setenv PATH $STARCCM_DIR:$PATH



setenv PBS_NNODES   `qstat -f $PBS_JOBID | grep
Resource_List.mppwidt | awk '{print $3}'`



# MPI library for XE6 (GEMINI)

setenv CRAYXT_HOME /opt/cray/mpt/default/xt/gemini/
mpich2-gnu



# MPI library for X4 or X5 (SeaStar)

# setenv CRAYXT_HOME #/opt/cray/mpt/default/xt/seastar/
mpich2-gnu



# Run starccm+

starccm+ -power -collab -np $PBS_NNODES -time -arch
linux-x86_64-2.3.4 -mpidriver crayxt -mppflags “-S 3” -
batch my.java -rsh ssh mycase.sim



-----------------------End of
Jstar-----------------------
```

Where:

- `-collab` is an option necessary to avoid user name checking.
- `my.java` is an optional java file.
- `mycase.sim` is the STAR-CCM+ simulation file.

The command setenv PBS_NNODES `qstat -f $PBS_JOBID | grep Resource_List.mppwidt | awk '{print $3}'` queries the number of processes and stores this number in the `PBS_NNODES` variable. This variable is then used in the command line option: `-np <nprocesses>`.

To launch the STAR-CCM+ job, enter the following in the command line:

```
qsub -l mppwidth=128 Jstar
```

Selecting the Fabric

The *fabric* describes the interface between the MPI layer and the network hardware. For instance, TCP is the fabric interface to Ethernet. Infiniband hardware can be interfaced using different software layers, each described using a different fabric selection.

Each MPI implementation defaults to the most suitable fabric. In most cases, you do not need to specify the fabric. You need to specify the fabric only if the MPI distribution fails to automatically pick the correct fabric or if you wish to use a non-default fabric. The fabrics available for selection depend upon their availability on the system, and the MPI implementation being used. For example, if the system has Myrinet MX hardware installed, it is unlikely that the QLogic PSM interface is available. In this case, the default MPI should detect the Myrinet MX hardware and automatically select the correct fabric at runtime.

On Cray systems, Cray MPI must be used and you have no control over fabric selection. On these systems, the locally installed MPI distribution is tuned with the hardware to provide optimal performance.

Fabric selection is not supported when using SGI MPI at this time. SGI MPI is tuned for optimal performance on SGI architectures.

The fabric can be specified using the command line option:

```
-fabric <fabric>
```

If you enter the fabric option in uppercase (e.g. `-fabric TCP`), this specifies no fall-back. If you specify the fabric option in lowercase (e.g. `-fabric tcp`), STAR-CCM+ falls back to an alternative fabric if the specified fabric is not available, and if an alternative fabric is available.

Fabric Option	Fabric
`-fabric TCP`	Ethernet
`-fabric IBV`	Infiniband OpenFabrics Alliance Verbs API (IBV)
`-fabric UDAPL`	Infiniband User Direct Access Programming Library (UDAPL)
`-fabric PSM`	Infiniband QLogic Performance Scaled Messaging (PSM)
`-fabric GM`	Myrinet GM
`-fabric MX`	Myrinet MX

The `-fabric <fabric>` option is equivalent to the `-mpp "-<fabric>"` option used in earlier scripts for HPC clusters. Thus `-fabric TCP` is equivalent to `-mpp "-TCP"`.

The`-fabric <fabric>` option is equivalent to the IBM Platform MPI options passed using the `-mpp "-<fabric>"` option. Thus `-fabric TCP` is equivalent to `-mppflags "-TCP"` in IBM Platform MPI.

To show the fabric selected by MPI, where applicable, use:

```
-fabricverbose
```

SLURM

Use the `-slurm` command line option to spawn new processes under the control of SLURM using the `srun` command.

This option is equivalent to these options in earlier scripts:

- PC MPI: `-mppflags "-srun"`
- Intel MPI: `-mppflags "--bootstrap=srun"`

CPU Binding

CPU binding constrains a simulation to run only on a particular set of CPUs or cores, and not to migrate to others during the life of the simulation.

In the STAR-CCM+ command line, you use the `-cpubind` option to specify the default CPU binding. CPU binding is turned on by default as binding processes can improve the performance of STAR-CCM+.

This option applies to both serial and parallel cases on Windows and Linux. In serial, the option pins the serial server process to the core that is assigned to the process by the scheduler. In parallel, the server process on each host is bound to the cores alternating between available sockets on multi-socket chips.

For parallel simulations, there are two binding policies that you can specify using the following argument:

```
-cpubind <policy>
```

Where `<policy>` is:

- `bandwidth` - default option that allocates processes to cores alternating between sockets
- `latency` - allocates processes to all cores on the same socket before allocating to additional sockets

If you do not specify a policy in the argument, the default policy is used.

This option is equivalent to these options in earlier scripts:

- IBM Platform MPI: `-mppflags "-aff=automatic:bandwdith"`
- Intel MPI: `-mppflags "-binding pin=1;cell=core;map=scatter"`
- Open MPI: `-mppflags "-bind-to-core -bysocket"`

To request a report of the CPU binding in parallel, type:

```
-cpubind v
```

or

```
-cpubind <policy>, v
```

When STAR-CCM+ identifies another STAR-CCM+ server running on any of the hosts that are being used, it disables any requested CPU binding for all processes on all hosts that are being used. This avoids processes being bound to identical cores, which can result in significant performance degradation.

You can manually deactivate CPU binding by typing the following argument in the command line while launching STAR-CCM+:

```
-cpubind off
```

Configuring MPICH2 for Network Installations on Windows

If you install STAR-CCM+ onto a shared network drive, ensure that the MPICH2 services that are installed on the local workstations have permission to access this shared drive. If the MPICH2 service cannot access the shared drive, it may not start automatically and must be configured.

To configure the MPICH2 service:

1. Access the *Services* dialog from the local workstations where MPICH2 is installed. For instruction on how to access the *Services* dialog, click one of the links below:

   a) Click **Start** and enter `Services` in the search field.

   b) Click **Services** in the list of search results.

   The *Services* dialog appears.
2. In the *Services* dialog, locate the **MPICH2 Process Manager, Argonne National Lab** service, right-click on it and select **Properties**.

The *MPICH2 Process Manager, Argonne National Lab Properties* dialog appears.

3. Click the *Log On* tab.
4. Click the **This account:** radio button under the *Log on as: property*.
5. Enter the details of an account that have administrator privileges with access to the network share.
6. Click **OK**.

You can now start the MPICH2 service.

Registering Your Credentials for Distributed Simulations on Windows

Your credentials must be registered before running distributed simulations on multiple nodes. This section describes how to:

- register your credentials on non-HPC Server versions of Windows
- register your credentials on Windows HPC Server 2008

Registering Your Credentials on Non-HPC Server versions of Windows

To register your credentials on non-HPC Server versions of Windows:

For MPICH2:

1. Start a **Command Prompt** from the **Start** menu.
2. Enter the command:

```
"<path>\mpiexec" -register
```

   Where `<path>` is the full path to the `mpiexec.exe` executable file in your STAR-CCM+ installation directory. The following is an example of a typical path for a 64-bit installation using STAR-CCM+ version 7.04.006:

```
C:\Program Files\CD-adapco\STAR-CCM+7.04.006\mpi
\mpich2_ch3\win64\bin
```

3. Enter your user name when requested.
4. Enter your password when requested.

For IBM Platform MPI:

5. Start a **Platform-MPI Command Prompt** using the action that is provided within the Platform-MPI group on the **Start** menu.
6. Enter the command:

```
mpidiag -s 127.0.0.1 -cache -at
```

7. Enter your password when requested.

Registering Your Credentials on Windows HPC Server 2008

To register your credentials on a Windows HPC Server 2008 machine:

1. Launch STAR-CCM+ with the desired number of cores. See Launching STAR-CCM+.

If your credentials are not already registered, the *Connect to [host_name]* dialog appears.

2. Enter your user name and password.
3. To store these credentials, activate *Remember my password*.
4. Click **OK**.

Some additional points must be noted when using IBM Platform MPI on Windows Server 2008 machines for distributed simulations on multiple nodes.

5. STAR-CCM+ must be launched from within a PowerShell command window. The current working directory must be defined as a UNC path that includes the name of the head node on which you are launching STAR-CCM+. An example of a UNC path is:

```
\\hpcmachine\WORKAREA\fluidsims\case1
```

The working directory must be shared and accessible from all compute nodes.

If you used the `-NOADMIN` command line option to install STAR-CCM+ without having administration privileges, you may receive the following error message when you attempt to register your credentials:

```
connect() to IP 127.0.0.1/port 8636, failed:
GetLastError(10061)Socket Connect failed: 10061
```

Your credentials are still registered even though this message appears.

Using Alternative Versions of IBM Platform MPI

On Linux platforms, STAR-CCM+ can be configured to use an alternative version of IBM Platform MPI to the one that is included in the STAR-CCM+ installer. This facility is useful if your cluster requires features that are not supported by the version that is provided in the installer. This facility also allows older versions of STAR-CCM+ to use more recent versions of IBM Platform MPI.

Provide the location of the alternative installation of IBM Platform MPI using the environment variable, `PCMPI_DIR`.

Only use this feature when the bundled version of IBM Platform MPI prohibits the use of the machine.

To create the environment variable, enter the following in the command prompt:

- For BASH/KSH shells:

```
% export PCMPI_DIR=<path to IBM Platform MPI
installation root directory>
```

- For BASH CSH shells:

```
% setenv PCMPI_DIR <path to IBM Platform MPI
installation root directory>
```

Using a Local Multi-Core Workstation

If a workstation is a multi-core machine it is convenient to exploit these resources with STAR-CCM+, either from the workspace or from the command line. Typically, you can experience a significant decrease in the time that is required to reach a solution, although limits on the parallel efficiency can be experienced due to the sharing of system resources under this architecture.

See Also

Launching STAR-CCM+

Starting from the Workspace

To start a simulation on a cluster using the workspace:

1. Launch the workspace.
   - If you are starting a new simulation, select **File > New Simulation...** to activate the *Create a New Simulation* dialog.
   - If you are opening an existing simulation, select **File > Load Simulation...** to activate the *Load Simulation* dialog.
2. By default serial mode is selected, so click the **Parallel on Local Host** radio button to change to parallel mode.



3. Enter the number of processes in the **Compute Processes** box. You can also increment this value using the arrows. The equivalent command line which is used to start the server is displayed in the **Command** text box.
4. Click **OK** to accept the settings. Now, any simulation that you open runs the server in parallel mode with the specified number of processes on the local host.

Starting from the Command Line

If your machine has 2 cores, two processes can be started using the `-np` option. The full command line would be:

```
starccm+ -np 2
```

In this simple case, you do not need to specify the `-on` option or the `-machinefile` option. STAR-CCM+ presumes that you are referring to the local host when you omit these options.

Using a Remote Multi-Core Workstation

To start a simulation on remote machines using the workspace:

1. Open the workspace.
   - If you are starting a new simulation, select **File > New Simulation...** to activate the *Create a New Simulation* dialog.
   - If you are opening an existing simulation, select **File > Load Simulation...** to activate the *Load Simulation* dialog.
2. By default serial mode is selected, so click the **Parallel on Named Hosts** radio button to specify the host machines for parallel mode. This option displays a small spreadsheet where you can enter the list of hosts.



3. Specify the machines by typing entries into the list. Click in the Host column of the first row of this list and type the first machine name. Press `<Tab>` or `<Enter>` to move to the next row.

In addition to the table format, this dialog contains buttons for modifying the list of hosts. Simply enter the name of the host in the **Add Host** text box and click the **Add** button. The **Delete** button deletes from

the list a host whose name you select in the list. STAR-CCM+ stops using this host for parallel processing.

Verify all your hostnames before attempting to launch the simulation. If any of the hostnames is incorrect when you click the **OK** button, STAR-CCM+ generates an error message and no simulation is started.



4. Specify the number of processes in the Processes column. Add the other machine names in the same fashion until you are done.



5. Accept the completed dialog by clicking the **OK** button. Now, any simulation that you open runs the server in parallel mode and distributes the processes on the machines specified. The parallel controller process is started on one of the hosts by default.

See Also

Launching STAR-CCM+

Using a Cluster of Machines

A cluster of machines can also be utilized in parallel mode, using the workspace or the command line. It is also possible to start a cluster of machines manually. This facility provides a much more scalable solution if you wish to decrease significantly the time that is required to reach a solution while also increasing the total memory available for a simulation. In addition, you can edit hostnames of machines based on your network requirements. With the Power Session license feature, you can run one job on an unlimited number of machines. Finally, for machines that support it, a parallel I/O option is available to improve the save and restore performance.

Information is provided separately about how STAR-CCM+ works with specific network devices. You can start a proxy server on the cluster to expose running simulations to workstations outside the cluster's private network.

See Also

How Many Cluster Nodes Should I Use for a Simulation?
Starting from the Command Line
Starting a Cluster of Machines Manually
Using the Power Session License Feature
Editing Hostnames of Machines
Using Parallel I/O
Cluster Node Communication
Starting a Proxy Server on a Host

How Many Cluster Nodes Should I Use for a Simulation?

It is important to scale the simulation size with the number of cluster nodes being used. The optimal scaling is determined from a ratio of the time to compute and the time that is taken to exchange data between cluster nodes.

If the simulation is distributed between too many cluster nodes, the time that is taken to perform the necessary data exchange dominates the total time that is taken to compute each simulation iteration. Consequently, no benefit is achieved from using more cluster nodes.

The time that it takes to exchange data between cluster nodes is determined by the size of the simulation, the number of cluster nodes

that are used, and the speed at which the network hardware can exchange data.

CFD calculations require pulling large amounts of data through each core continuously. Since the typical calculations performed on the data are relatively simple, the key limitation on overall performance comes from the ability to access the data. Multiple cores have to share memory bandwidth, which may not be able to keep up with the data flow requirements.

Therefore, the optimal number of computational processes is dependent on the specification of the nodes and the interconnects used. It is recommended that when using 100/10BaseT Ethernet cards the minimum number of mesh cells per process is 100,000, whereas on higher-performance network hardware this limit can be decreased.

Using the Workspace (with a Machine File)

This section covers specifying machines in the workspace with a machine file. Other approaches in the workspace are covered in a separate section.

It is useful to specify the machines using an MPI file when multiple machines are being used. The format of the file is simply a list of machines on which you wish to place processes. The name of the machine can be concatenated with a colon and an integer to represent the number of processes you wish to place on the machine. An example of such a file could be:

```
node1
node2:3
node3:2
```

This file dictates that processes are allocated to machines `node1`, `node2`, and `node3` with 3 processes allocated to `node2`, and 2 processes allocated to `node3`.

If the machine file specifies more resources than are requested using the `-np` option the machine file is parsed allocating processes to the machines until the requested number have been allocated. In the above example, if you requested 3 processes using `-np 3` on the command line the processes would be placed on `node1` and `node2`.

If the machine file specifies fewer resources than are requested using the `-np` option the allocation of processes is continued from the start of the machine file once the end is encountered. In the above example, if you requested 8 processes the first would be allocated to `node1`, the second,

third, and fourth allocated to `node2`, and the fifth and sixth allocated to `node3`. Then the seventh would be allocated to `node1`, and the eight allocated to `node2`.

To use the machine file option from the command line, specify the file using the machine file option `-machinefile <file name>`. A typical command line looks like:

```
% starccm+ -machinefile mymachines -np 4 ~/airfoil.sim
```

The machine file can also be specified from the workspace. In the *Create a New Simulation* dialog or the *Load Simulation* dialog click the **Parallel Specified by Machine File**radio button. Then do the following:

- Enter the number of required processes in the **Compute Processes** spin box.
- Either enter the name of the machine file in the **Machine File** text box or click the ellipsis button to open the file browser.



The **Command** text box displays the command line that is used for launching the server. You can edit this command line directly to override the command line parameters that are generated by the GUI elements.

Starting from the Command Line

Assume that you have a cluster of four machines that are called `node1`, `node2`, `node3`, and `node4`, and you want to run `airfoil.sim` in your home directory using one process per machine. If you are starting this job from another machine that shares a common view of the filesystem where your simulation file is located (for example, your home directory on your Linux machine is shared through NFS with the cluster machines), enter the command:

```
% starccm+ -on node1,node2,node3,node4 ~/airfoil.sim
```

This action starts the STAR-CCM+ client on your local machine, which in turn starts the STAR-CCM+ server on the cluster machines with one parallel process per machine. The named simulation file is read in to the server.

There is no mechanism for specifying a different user name in parallel mode, and you cannot start the master on a remote machine using the `-machinefile` option.

STAR-CCM+ assumes that the install path is the same on all the cluster machines as on the machine on which the server is started. If NFS is being used, the install path must be the same.

Starting a Cluster of Machines Manually

If the filesystem associated with the cluster is not visible from your local machine, if you are starting a parallel job from a Windows machine for example, you can start the remote server manually. To start the remote server manually, remotely log in to the cluster machine (through telnet or rlogin), and then connect to the server from the client. (To facilitate this process, read about how STAR-CCM+ works with specific network devices).

Telnet to the master by typing the command:

```
telnet node1
```

into the *Run* dialog (use **Start > Run...** to get to that dialog) or in a command shell window and login to that machine as you normally would.

From your telnet window, issue the command:

```
% starccm+ -server -on node1,node2,node3,node4 ~/
airfoil.sim
```

The `-server` option instructs STAR-CCM+ to start the server rather than the client process. The `-on` option prescribes the hosts where processes are started. When starting a server, the `-on` argument must specify the local host. The workspace is not started.

Back on your Windows machine or workstation open the workspace of STAR-CCM+ and connect to the server.

Editing Hostnames of Machines

A command line option permits you to edit the hostnames of the machines that are prescribed to the server. This facility is primarily for the use with batch systems.

Some clusters are set up so that each cluster node (machine) has two network cards. The first network card can be a default ethernet card while the second one can be a higher-performance network card. In these circumstances, the hostnames that are associated with the second card are not the same as those allocated to the first card. Typically the hostnames are appended by a string, but it depends on how the machine is configured. The batch systems do not see second cards, so when the batch system allocates nodes on which jobs can run, they refer to the default cards.

However, when you run STAR-CCM+ you may want to take advantage of the higher-performance card. To do this, you would need to translate the hostname that is provided by the batch system to the hostname you actually want to use.

Suppose that you have a cluster of machines with names:

```
green-node1
green-node2
green-node3
...
green-nodeN
```

The batch system can allocate four nodes:

```
green-nodeN1
green-nodeN2
green-nodeN3
green-nodeN4
```

However, the application wants to use the higher performance network cards which could be named:

```
green-nodeN1-eth1
green-nodeN2-eth1
```

```
green-nodeN3-eth1
green-nodeN4-eth1
```

To use the higher performance network cards, enter the following in the command line:

```
-hostregexpr /<pattern>/<substitution>/
```

where the "regular expression" `pattern` is replaced by the regular expression `substitution`. Regular expressions are a general way of pattern matching strings.

Also, `/<pattern>/<substitution>/` is passed directly into the command:

```
echo $machinename | sed -s 's/<pattern>/
<substitution>/g'
```

Examples of how you could use this technique are:

- Appending a string to the end of every hostname:

  ```
  -hostregexpr /\$/<string>/
  ```

  So to append `-eth1` you would enter:

  ```
  -hostregexpr /\$/-eth1/
  ```

  The dollar sign represents the end of line (or end of hostname) and replaces it with the string. At the same time, it is necessary to escape the dollar sign to prevent the command-line shell from interpreting it. Hence, the syntax is `/\$/-eth1/` instead of `/$/-eth1/`.

- Adding a string to the start of every hostname:

  ```
  -hostregexpr /^/eth1-/
  ```

  This option gives entries such as `eth1-green-nodeN1`. (The ^ is the start of a line or hostname in a regular expression. There is no need to escape this character as it has no special meaning to the shell.)

- Substituting a common string with another:

  ```
  -hostregexpr /green/green-eth1/
  ```

  This option results in `green-eth1-nodeN1`.

Using the Power Session License Feature

The Power Session license feature lets you run one job on an unlimited number of cores. A list of the different license features is provided in the

STAR-CCM+ Licensing Guide section. When this type of license is used, other licenses such as `ccmpsuite` remain inactive.

If a job is started using the Power Session, only one copy of the `ccmppower` license is checked out. Contact CD-adapco to obtain this license.

Starting from the Workspace

You can activate the Power Session feature using a checkbox in the *Create a New Simulation* and *Load Simulation* dialogs.



When you tick this checkbox, the option `-power` appears in the command line of the dialog automatically.



Starting from the Command Line

The command line includes the option `-power` as follows:

```
starccm+ -power -np 2
```

If the license is not available when you attempt to activate it, STAR-CCM+ generates a message like this one:



Using Parallel I/O

For systems that support it, simulation files can be saved and restored using parallel input/output (I/O).

With serial I/O (the default behavior), a single STAR-CCM+ server process performs all the data input and output on behalf of all server processes, scattering and gathering data as needed. This data flow creates a serial I/O bottleneck that can limit save and restore performance. With parallel I/O activated, multiple STAR-CCM+ server processes perform the I/O operations concurrently using the parallel file system. This option mitigates the serial I/O bottleneck, and can potentially reduce the time it takes to save and restore data.

To activate parallel I/O, you require a parallel file system. Using the feature on a locally mounted file system or network file system (NFS) usually provides no performance benefit, and can actually reduce performance. NFS is not a parallel file system.

The format of simulation (`.sim`) files that are saved with parallel I/O activated is the same as the format of simulation files that are saved using default serial I/O. A simulation file that is saved using parallel I/O can be read using serial I/O. Similarly, a simulation file that is saved using serial I/O can be read using parallel I/O.

Backward Compatibility

You can restore all old simulation files using parallel I/O. To obtain the full performance benefit of parallel I/O on restore, the simulation file must have been saved using STAR-CCM+ version 6.06 or higher. Although the simulation file format for serial and parallel I/O are the same, the format was revised for version 6.06 to improve its application to parallel I/O.

Limitations

Current limitations are:

- Linux only. Currently, parallel I/O is only available on Linux platforms.
- Simulation files only. Currently, the capability is only available for saving and restoring `.sim` files.

What are the Requirements for Parallel I/O?

The main requirement for this feature is a parallel file system. A command line option is required to activate the feature with STAR-CCM+. Further requirements are:

- Common visible path on all compute nodes
- Sufficient hardware resources
- Sufficiently large simulations

Parallel File System

A parallel file system is required to get benefits from this feature. Examples of parallel file systems include PanFS from Panasas, Lustre, and the General Parallel File System (GPFS) from IBM. Of parallel file systems, the current capabilities have been tested and tuned for PanFS. The feature should work on other parallel file systems that are supported by the I/O framework of the Message Passing Interface (MPI) standard.

Common Visible Path

A common visible file path on all compute nodes is required for this feature to work correctly. The volume on the parallel file system (where the simulation file is accessed) must:

1. be mounted on all compute nodes where the associated STAR-CCM+ server processes are being run; and,
2. have a common mount point. For example, if a volume is mounted as `/foo` on some nodes and `/fu` on other nodes, the simulation file cannot be saved as `/foo/bar.sim` with parallel I/O since `/foo` is not visible on some of the nodes.

STAR-CCM+ attempts to revert to serial I/O in cases where these requirements are not met, but this is not guaranteed to be successful.

Sufficient Hardware Resources

Even with a parallel file system, sufficient hardware resources must be available to obtain significant performance benefits. For example, with Panasas storage systems, multiple shelves (or blade racks) are required

to obtain significant performance improvements relative to serial I/O (especially on a fast network).

Sufficiently large simulations

The benefits of parallel I/O can only be realized for sufficiently large simulation files (typically containing more than 1 million cells). This is because sufficiently large storage blocks are required to effectively stripe across the available file system servers.

Activating Parallel I/O

You can activate parallel I/O for simulation file saves and restores globally from the command line using the `-pio` option. For example:

```
$ starccm+ -pio -np 8 foo.sim
```

When this flag is included, the given simulation file is restored using parallel I/O, and all subsequent simulation files are saved using parallel I/O.

Starting a Proxy Server on a Host

You can specify a host on which to start a proxy server using the `-proxy <host>` option.

You can use the `-proxy` option with both serial and parallel simulations, and all methods of specifying the hosts where the simulation runs, including the `-batchsystem` command-line argument. To expose simulations that are running on clusters to workstations outside the cluster's private network, use the `-proxy` option.

A typical command to enter in the command line is:

```
starccm+ -batch -np 64 -machinefile hostfile -rsh ssh -
proxy cluster-head mycase.sim
```

where:

- `-batch` is the argument that is used to run the case in batch mode
- `-np 64` specifies the number of cores to use (64 in this example)
- `-machinefile hostfile` is the configuration file for the host you want to spawn the proxy server on
- `-rsh ssh` is the shell that you want to use to spawn a proxy server on the specified host. By default, the secure shell (`rsh`) is used. You can specify to use the remote shell (`ssh`) by adding -rsh ssh to the command line (as shown in the example).
- `-proxy cluster-head` is the head node of the cluster you want to connect to.
- `mycase.sim` is the name of the simulation file you want to connect to

The proxy server appears in the *Servers* tab in STAR-CCM+. You can connect a client to this server. The proxy server continues running until the connection with the parent server is lost.

On startup, the *Output* window shows:

```
Server::start -host cluster-node-01:47827
```

This message identifies which host the server is running on and the port number on which the server is listening for client connections. In the example above, the host is `cluster-node-01` and the port number is `47827` (which is the default STAR-CCM+ port number).

Additionally, the output also displays:

```
ProxyServer:start -host cluster-head:47827
```

This message identifies the host and port number of the proxy server. In the example above, the cluster's head node is `cluster-head` and the port number is `47827`. The port numbers of the server and proxy server do not have to be the same.

You can now connect to either the server or proxy server using the respective host and port combinations. If connecting to the server from outside the cluster, it is likely that you cannot see the cluster nodes (such as `cluster-node-01`) but are able to see the cluster's head node (`cluster-head`). In this case, connect through the proxy server.

Connecting to the proxy server grants you the same access to the simulation as connecting directly to the server.

Starting a Proxy Server from a Windows Machine

The proxy server requires that rsh/ssh is able to start a remote process from the server's host to the proxy host. Windows machines do not support rsh/ssh by default—you can obtain an additional package (Windows Services for Unix) that adds this support from Microsoft: http://www.microsoft.com/download/en/details.aspx?id=274. Alternatively, you can use a separate program such as PuTTY.

Troubleshooting

This section contains some troubleshooting tips.

Specifying parallel hosts yields invalid entries

- Make sure that the specified hostnames are spelled and configured correctly.

Command line option fails

- Ensure that the command line shown in the GUI conforms to the options described in Command-Line Reference, p.459.

A Windows cluster does not work

- Install the most recent release version of STAR-CCM+ available using the same installation path on all machines, and ensuring that the MPI component is chosen and set up correctly.
- Verify on each machine that you can start serial and local parallel servers.
- Start the multi-host parallel server, trying one host at a time. Starting remote servers is not supported under Windows, so when starting a multi-host parallel server the master process must be run on the local host.
- If it does not work, turn off the firewalls on the machines and try again. If you need help, start the server from the command line with the `-v` option to get debug output, as follows:

  ```
  "c:\Program Files\CD-adapco\STAR-CCM+
  <version_number>\star\bin\starccm+.exe" -server -on
  host1,host2 -v
  ```

  where `host1` and `host2` are the two machines that you want to run on, and one of them is the current local machine.

Memory limits are set too low

Linux workstations are often configured with low limits on the amount of allowable shared memory. This behavior restricts how much memory is pinned by the libraries that MPI uses. These libraries can print warning messages even when using a single host and they usually indicate that the limits are set too low—even if the library is not being used. Set workstations to have high or preferably unlimited limits.

The error messages are:

From IBM Platform MPI:

```
libibverbs: Warning: RLIMIT_MEMLOCK is 32768 bytes
```

From Open MPI:

```
The OpenIB BTL failed to initialize while trying to
allocate some locked memory. This typically can
indicate that the memlock limits are set too low. For
most HPC installations, the memlock limits should be
set to "unlimited".

The failure occurred here:

Host: compute_node.example.com
OMPI source: btl_opebib.c:114
Function: ibv_create_cq()
Device: Out of memory
Memlock limit: 32767
```

Instructions on how to increase your system limits are provided in the IBM Platform MPI User Guide and Release Notes under the IBM Platform MPI directory in the installation. The Open MPI Frequently Asked Questions also provide a guide. (See answer to How can a system administrator (or user) change locked memory limits?)

Balancing

Care must be taken to balance a parallel configuration to achieve maximum performance. Generally, allocate each process to a single core where no other jobs are being executed. Any additional load on the core to which a process is allocated results in an imbalance in the parallel configuration and paging, resulting to a degradation in performance. In general, the slowest core in the machine limits the parallel simulation: so if one core has a greater load than the remaining cores, all cores are limited by the excessively loaded core.

In the case of a distributed parallel simulation, performance is limited by the slowest machine in the cluster. If one node has a greater load than the remaining nodes, all nodes are all limited by the excessively loaded node.

Getting Parallel Help through the Command Line

You can get information on the parallel options through the command line.

To get detailed help about parallel in the console, enter the following in the command line:

```
starccm+ -phelp
```

This command can also be used with specific parallel options by adding the syntax `:<parallel option>`. For example, to get specific help on the `-np` option, enter:

```
starccm+ -phelp:-np
```

Viewing the Partitioning

The distribution of the simulation data can be visualized in the same way as any other scalar quantity: adding a scalar scene or adding a scalar displayer to an existing scene, and then changing the scalar quantity selection (the field function) to `Partition`. The cells are colored according to the partition that they have been allocated to by STAR-CCM+.



Using Batch Management Systems

Batch systems manage the allocation of computational jobs to a single large computer or a cluster of smaller computers.

You would typically submit a job to the batch system requesting computational resources through the attributes of that particular batch system. The batch system would then manage the job, allocating it specific machines when they become available based on the requested resources. In this way, the resources can be delegated to computational jobs in a fair and efficient manner that is based on rules that are configured by your system administrator.

STAR-CCM+ can be run in batch mode under the control of a batch system. If you wish to run STAR-CCM+ in parallel under the control of a batch system, certain resource attributes must be obtained from the batch system and used to construct the correct command-line arguments for STAR-CCM+. A common way of doing this is to examine attributes of the batch system in the batch system's submission script and pass these attributes to the STAR-CCM+ command line. For example, if you requested two machines from the batch system when you submitted the job, the batch system should provide a way of finding out which machines have been allocated. These machine names can then be used to construct a machine file for the parallel execution of STAR-CCM+.

Automatic transfer from the batch system to STAR-CCM+ of a limited number of attributes is supported for the following batch systems:

Working in Unix-Based Batch Systems

STAR-CCM+ helps you interact with certain batch systems by simplifying the submission process. You can either write a submission script manually or permit STAR-CCM+ to submit itself automatically to the system.

Using Manual Submission

STAR-CCM+ simplifies manual submission by providing the `-batchsystem` argument, which automatically extracts batch system properties and translates them into STAR-CCM+ command-line arguments. This method requires you to construct a small submission script manually, and is the most popular submission technique today. This technique is most easily explained through an example of a simple

submission script, which is in the section on using the PBS Batch System.

Due to the multitude of batch systems and possible configurations for them, the details of the particular submission script that you may wish to construct are beyond the scope of this section. Refer to the documentation of your own specific batch system and, if necessary, consult with your system administrator to identify the attributes to forward to the STAR-CCM+ command line.

If you want to use a different parallel configuration, perhaps one more complex, write a custom script. The standard transfer of information from the batch system to STAR-CCM+ will not be adequate. This method involves bypassing the `-batchsystem` option, extracting the batch system properties manually, and forwarding them to your STAR-CCM+ script.

An example of a custom script is included in the section on using the PBS Batch System.

Using Automatic Submission

STAR-CCM+ can submit itself automatically to a selection of popular batch systems, so you do not need to write a submission script. In this setup, each process is placed on its own batch system node by default.

Note: You cannot use automatic submission if you have a custom parallel configuration.

The general syntax is as follows:

```
starccm+ <options> -batchsystem <batchsystem>:"-submit
<batchoptions submitoptions>" <sim file>
```

where the batch system is: `pbs`, `lsf`, `sge`, or `ll`.

Among the submit options are those that can be forwarded to the `submit_starccm+` script, and those that can be forwarded to the batch system submission command. Options that are forwarded to the `submit_starccm+` script include:

- `-logfile <file>` Specifies the log file to receive the output.
- `-submitfile <file>` Specifies the filename of the submission script generated.
- `-writescript` Writes the submission script and exits (that is, it does not actually submit the job).

  This option is useful when you want to generate a template submission script, rather than launch an actual submission.

- `-tmpdir <dir>` Specifies the temporary directory.
- `-v` Increases verbosity.
- `-cleanup` Specifies running the `kill_procfile` script after completing the job.

An example of automatic submission is included in the section on using the PBS Batch System.

Using PBS

When you specify PBS manually on the command line, STAR-CCM+ will automatically distribute a process to each PBS node allocated to the PBS job.

An alternative is automatic submission.

To specify PBS manually, create a simple script file (for example, `pbsbatch`) that contains the full command line that you would use if you were not using a batch management system, but include the `-batchsystem` option:

```
#!/bin/sh
cd $PBS_O_WORKDIR
/install_path/star/bin/starccm+ -batchsystem pbs -batch
\
 <macro file>
```

Next, submit the job to the batch management system, specifying the number of resource nodes to use:

```
% qsub pbsbatch -q <queue name> -l nodes=<# processes>
```

To see how the job is progressing, you can use the following command:

```
% qstat -a
```

The `-batchsystem` argument only supports the case in which one process is distributed to each of the PBS nodes that are allocated for the job. If processes are to be distributed to the PBS nodes in a different manner, you cannot use the `-batchsystem` argument and must write a custom submission script. An example of such a script is provided below.

```
#!/bin/sh
# This script requests 4 OpenPBS nodes each with
# two cores. It then allocates two starccm+
# processes to each OpenPBS node.
# Request 4 nodes each with two cores
```

```
#PBS -l nodes=4:ppn=2
##################################################
# Additional PBS directives could be placed
# here. For example the queue could be specified
# or the maximum cpu or wall time specified.
##################################################
# specify 8 processes
NCPU=8
echo
-------------------------------------------------------
echo 'This job is allocated on '${NCPU}' cpu(s)'
echo 'Job is running on node(s): '
cat $PBS_NODEFILE
echo
-------------------------------------------------------
cd $PBS_O_WORKDIR
starccm+ -batch <macro file> -np $NCPU -machinefile
$PBS_NODEFILE
exit
```

Automatic Submission to PBS

For PBS, an example automatic submission command is displayed below:

```
starccm+ -rsh ssh -mpidriver platform:"-UDAPL -prot" -
batch profile.java -batchsystem pbs:"-submit -logfile
logfile -cleanup -lwalltime=1:00:00 -
lselect=14:ncpus=8:mpiprocs=8" /path_to_simulation_file/
setup.sim
```

An example of the submission script that is automatically generated by STAR-CCM+ is provided below.

```
#!/bin/sh
# This file has been automatically generated by the
submit_starccm+ script.
# It is used to submit a STAR-CCM+ command to the pbs
batch management
# system and should be removed automatically once the
job has finished.
# The existence of this script implies the job is still
being executed.
# run exit_hook on exit
trap 'eval $exit_hook' 0
trap 'eval $exit_hook; echo "`basename $0`
Interrupted"; exit_hook=; exit' 2
trap 'eval $exit_hook; echo "`basename $0` Terminated";
exit_hook=; exit' 15
trap 'eval $exit_hook; echo "`basename $0` Aborted";
exit_hook=; exit' 6
exit_hook="rm -f /home/cdadapco/noel/data/f1/47million/
pids.19762; $exit_hook"
exit_hook="kill_procfile /home/cdadapco/noel/data/
```

```
f1/47million/pids.19762; $exit_hook"
### PBS ###
### PBS Options added automatically
### Name of job
#PBS -N setup
## Set the PBS shell to sh
#PBS -S /bin/sh
## Request wall clock time of 24 hours
#PBS -lwalltime=24:00:00
### Join stdout and stderr
#PBS -j eo
### Set output file
#PBS -o /home/cdadapco/noel/data/f1/47million/setup.
19762
### Print some useful PBS output
echo
"-------------------------------------------------------"
>> logfile.19762
echo "Job is running on node" >> logfile.19762
cat $PBS_NODEFILE >> logfile.19762
echo
"-------------------------------------------------------"
 >> logfile.19762
echo "PBS: qsub is running on $PBS_O_HOST" >> logfile.
19762
echo "PBS: originating queue is $PBS_O_QUEUE" >>
logfile.19762
echo "PBS: executing queue is $PBS_QUEUE" >> logfile.
19762
echo "PBS: working directory is $PBS_O_WORKDIR" >>
logfile.19762
echo "PBS: execution mode is $PBS_ENVIRONMENT" >>
logfile.19762
echo "PBS: job identifier is $PBS_JOBID" >> logfile.
19762
echo "PBS: job name is $PBS_JOBNAME" >> logfile.19762
echo "PBS: node file is $PBS_NODEFILE" >> logfile.19762
echo "PBS: current home directory is $PBS_O_HOME" >>
logfile.19762
echo "PBS: PATH = $PBS_O_PATH" >> logfile.19762
echo
"-------------------------------------------------------"
 >> logfile.19762
### Set working directory
WORK_DIR="/home/cdadapco/noel/data/f1/47million"
cd $WORK_DIR
### Simulation File
SIM_FILE="setup.sim"
### Macro
MACRO="profile.java"
### Run Command
starccm+ -rsh ssh -mpidriver platform -mppflags "-UDAPL
-prot" -batch $MACRO -batchsystem pbs $SIM_FILE -
batchsystemoptions " -nosubmit" -procfile /home/
cdadapco/noel/data/f1/47million/pids.19762 >> logfile.
```

```
19762
### Sleep for 2 minutes. This helps ensure the STAR-CCM
+ server
### has shut down before any prologue scripts kill
remaining processes
sleep 120
```

The submission command that automatically gets executed looks like this:

```
qsub -lwalltime=1:00:00 -
lselect=14:ncpus=8:mpiprocs=8 /tmp/pbs-cdadapco20273/
pbs-submit.20273
```

Using LSF

By specifying LSF on the command line, STAR-CCM+ will automatically distribute a process to each LSF node allocated to the job.

An alternative is automatic submission.

Submit the job to the batch management system, specifying the number of resource nodes to use (\ indicates a line continuation):

```
% bsub -q <LSF queue> -n <number of processes> \
/install_path/star/bin/starccm+ -batchsystem lsf \
-batch <macro file> /path_to_simulation_file/simfile.sim
```

To see how the job is progressing:

```
% bjobs
```

Using Sun Grid Engine

By specifying Sun Grid Engine on the command line, STAR-CCM+ will automatically distribute a process to each Sun Grid Engine node allocated to the job.

An alternative is automatic submission.

Submit the job to the batch management system, specifying the number of resource nodes to use:

```
% qsub -pe <SGE Parallel Environment> <number of SGE
nodes> <Sun Grid Engine Script>
```

where the minimal Sun Grid Engine script would be

```
#!/bin/sh
/install_path/star/bin/starccm+ -batchsystem sge -batch
<macro file> /path_to_simulation_file/simfile.sim
```

To see how the job is progressing:

```
% qstat
```

Using Load Leveler

By specifying Load Leveler on the command line, STAR-CCM+ will automatically distribute a process to each Load Leveler node allocated to the job.

An alternative is automatic submission.

Submit the job to the batch management system, specifying the number of resource nodes to use:

```
% llsubmit <Load Leveler Script>
```

where the minimal Load Leveler script would be:

```
#!/bin/sh
# @ job_type = parallel
# @ node = <number of processes>
/install_path/star/bin/starccm+ -batchsystem ll -batch
<macro file> /path_to_simulation_file/simfile.sim
```

To see how the job is progressing:

```
% llq
```

Using Windows HPC Server 2008

STAR-CCM+ batch jobs can be submitted to the Windows HPC Server 2008 job scheduler interactively or using the -msjob option. Parallel jobs that are submitted to the job scheduler use MS-MPI by default. Before you can submit a batch job to the Windows job scheduler, register your credentials. See Registering Your Credentials on Windows HPC Server 2008.

To submit the job:

```
[INSTALL PATH]\star\bin\starccm+.exe -msjob:[job
options] [STAR-CCM+ options] -batch [PATH TO SIMULATION
FILE]\simfile.sim
```

The working directory, standard output, and standard error are set to the directory that the simulation file is in.

If the job is successfully submitted, the job scheduler prints the job ID:

```
Job has been submitted. ID: 230.
```

Cluster Node Communication

This section explains the following:

What is a Network Device?

A cluster of machines is connected using a network which provides the means for communication between the cluster nodes. The hardware that is used for the network determines the maximum speed that data can be exchanged between cluster nodes.

When running STAR-CCM+ in parallel, data is exchanged between the different cluster nodes several times during each simulation iteration, and so the performance of the network affects the performance of STAR-CCM+. Typically, clusters are not accessible from outside the head node. The nodes often have both ethernet and another flavor of network card. The network configuration prohibits direct access to the cluster nodes.

Support for each different network option is provided by using different MPI libraries. The MPI library provides the software to accomplish data transfer between different cluster nodes. Different MPI libraries interface with, and take optimal advantage of, the underlying network hardware, and the dependent communication protocol. By default STAR-CCM+ uses the included IBM Platform MPI or SGI MPI if it is available on the hosts. SGI MPI takes advantage of the SGI virtual shared memory architecture while IBM Platform MPI automatically probes the fastest network device available, and reverts to TCP/IP if none are found.

Shared Memory Limits Too Low

Linux workstations are often configured with low limits on the amount of allowable shared memory. This restricts how much memory can be pinned by the libraries that MPI uses. These libraries can print warning messages even when only using a single host and they usually indicate the limits are set too low—even if the library isn't being used. Set workstations to have high or preferably unlimited limits.

The error messages are:

- From Platform MPI:

```
libibverbs: Warning: RLIMIT_MEMLOCK is 32768 bytes
```

- From Open MPI:

```
The OpenIB BTL failed to initialize while trying to
allocate some locked memory.
This typically can indicate that the memlock limits
are set too low. For most HPC installations, the
memlock limits should be set to "unlimited".

The failure occurred here:

Host: compute_node.example.com
OMPI source: btl_opebib.c:114
Function: ibv_create_cq()
Device: Out of memory
Memlock limit: 32767
```

The Open MPI Frequently Asked Questions describe how the limits are changed (see answer to How can a system administrator (or user) change locked memory limits?)

What Different Network Devices Does STAR-CCM+ Support?

Enter a short description of your concept here.

STAR-CCM+ can take advantage of different network devices that are installed between cluster nodes. A list of supported devices includes:

- Shared Memory
- Ethernet
- Myrinet GM and Myrinet MX
- Infiniband:
  - OpenFabrics Enterprise Distribution using Verbs API Infiniband Interface (VAPI)
  - SilverStorm using User-Level Direct Access Programming Infiniband Interface (uDAPL)
  - QLogic Infiniband Interface
- SGI Shared Memory

Different network hardware can be selected with the command line by specifying the respective MPI driver and options:

```
-mpidriver <MPI driver><:driver options>
```

The `<driver options>` portion of the command line permits you to incorporate options that you want to add, depending on your use of the MPI application and the machines on which you are running

STAR-CCM+. However, these options must only be added by a support person or someone with expert knowledge of the MPI application.

Using Shared Memory

STAR-CCM+ always uses shared memory between processes running on the same host, regardless of the network hardware that is chosen to pass information between processes on different hosts.

Using Ethernet

If IBM Platform MPI cannot detect a faster network card, it reverts to using Ethernet. Ethernet can be explicitly chosen using:

```
-mpidriver platform:-tcp
```

In some cluster configurations, each node can be equipped with more than one network device. For example, a simple 100/10BaseT Ethernet device might be accompanied by a higher-speed Gigabit Ethernet device. Since both devices use the same underlying communication protocol (TCP/IP), the same MPI library can be used, but when two network devices are present using the same communication protocol, it is often necessary to tell STAR-CCM+ explicitly which one to use. This is normally achieved by specifying the correct hostnames that are associated with each network device.

Using Myrinet

Myrinet network devices provide support by two software interfaces: the GM Interface and the more recent MX (Myrinet Express) Interface. STAR-CCM+ with Myrinet can be faster than Ethernet, especially when large numbers of processes are involved.

You need a Myrinet card (in addition to the Ethernet network card) in each machine you want to use, as well as the appropriate drivers.

Using the Myrinet GM Interface

If the device supports the GM Interface, you can choose it by using:

```
-mpidriver myrinet
```

By default when you choose Myrinet GM using the above command, or let IBM Platform MPI choose Myrinet GM automatically using:

```
-mpidriver platform
```

IBM Platform MPI will try and locate the Myrinet GM lib `libgm.so`. It will first check whether the file is in the default system path, then it will look in the default installation path `/opt/gm/lib/`.

If the file is not located in the default installation path or in the system path, you may specify the location by including the environment variable in your `.cshrc` file:

```
setenv MPI_ICLIB_GM <installation path>/libgm.so
```

Using the Myrinet MX Interface

If the device supports the MX Interface, it may be chosen using:

```
-mpidriver myrinet-mx
```

By default when you choose Myrinet MX using the above command, or let IBM Platform MPI choose Myrinet MX automatically using:

```
-mpidriver platform
```

IBM Platform MPI will try and locate the Myrinet MX lib `libmyriexpress.so`. It will first check whether the file is in the default system path, then it will look in the default installation path `/opt/mx/lib/`.

If the file is not located in the default installation path or in the system path, you may specify the location by including the environment variable in your `.cshrc` file:

```
setenv MPI_ICLIB_MX <installation path>/
libmyriexpress.so
```

The Myrinet MX driver has not been tested.

Using Infiniband

Infiniband is a specification of a high-performance network device. Different Infiniband network solutions exist, and supported devices include:

- Open Fabrics Enterprise Distribution (OFED) Infiniband

  ```
  -mpidriver infiniband:-ibv
  ```

- SilverStorm via the User Level Direct Access Programming Interface (uDAPL)

  ```
  -mpidriver infiniband:-udapl
  ```

- QLogic Infiniband Interface

  ```
  -mpidriver infiniband:-psm
  ```

However, rather than specify the Infiniband interface, you can simply use the generic command to activate it:

```
-mpidriver infiniband
```

STAR-CCM+ translates this action as `-mpidriver platform` and thus lets IBM Platform MPI decide which Infiniband interface is most suitable. If no Infiniband interface exists, an alternative transport layer is chosen by IBM Platform MPI, but no warning is given to you.

Using SGI Shared Memory

SGI machines are often configured such that a distributed memory system appears as a shared memory system. Such SGI machines are often equipped with an MPI library that can take advantage of this specific memory architecture. STAR-CCM+ can take advantage of this memory architecture on single-host, multi-core machines where the SGI MPI library has been installed. STAR-CCM+ probes for SGI MPI and uses it automatically, provided you have not used the `-machinefile` or `-on` options. You do not need to enter a separate command.

SGI MPI must be installed separately on your machine; it is not distributed with STAR-CCM+.

If SGI MPI is not installed under the default directory (`/usr`), the environment variable `SGI_MPI_HOME` must be set such that `libmpi.so` and `libmpi++.so` exist in `$SGI_MPI_HOME/lib/`. Also, the executable `mpirun` must exist in `$SGI_MPI_HOME/bin/`. The `starccm+` script can show how this works. As only single-host machines are supported, prescribe the number of processes using the `-np` rather than the `-on` or `-machinefile` options.

SGI MPI uses a concept of cpusets to control which processes are run on which CPUs. To use SGI MPI, either specify a particular cpuset, otherwise STAR-CCM+ runs on the default cpuset. For more information about cpusets, refer to the SGI Technical Documentation.

Processes are placed on the CPUs within the cpuset in a round-robin fashion unless the `-oncpu` option is invoked to specify which CPUs relative to the cpuset are used. The SGI command `dplace` must be available to use `-oncpu`. The `starccm+` script checks for this option.

On some SGI systems, you can prescribe the CPUs where the STAR-CCM+ processes are run using the following command line option:

```
-oncpu <CPU1,CPU2,...,CPUN>
```

The series that is shown in brackets can either be a comma-separated list or a range that is given by a hyphen. The exact format is the same as the

SGI `dplace -c` command. This option is helpful when you wish to run on one shared-memory machine while restricting which CPUs of that machine STAR-CCM+ uses. For example:

```
starccm+ -np 10 -oncpu 0-12
```

To achieve good performance, make sure that the cpuset has enough CPUs allocated to it to permit each process to be placed on their own CPUs.

Polling is off by default. This driver is able to support polling to enhance performance by adding the option:

```
-mpidriver sgi:-polling
```

To deactivate it again, use the following line command:

```
-mpidriver sgi
```

Note: Do not over-subscribe a node with processes as it can result in performance degradation.

When starting SGI MPI, you may encounter the following warning:

```
"MPI: This executable was not linked with MPT version
1.6.1 or later. Memory mapping features such as single
copy transfers, MPI one-sided communication and shmem
functions are not available unless you relink your
executable with the current MPT library."
```

It is a known problem with SGI and can be ignored.

On Which Platforms Are Different Network Devices Used?

The table below represents the support for the different interconnects on different release platforms.

Platform	Interconnect
linux-x86	ethernet (IBM Platform MPI) Myrinet-GM (IBM Platform MPI)* Myrinet-MX (IBM Platform MPI)** Infiniband (IBM Platform MPI)**
linux-x86_64	ethernet (IBM Platform MPI) ethernet (SGI MPI) Myrinet-GM (IBM Platform MPI)* Myrinet-GM (SGI MPI) Myrinet-MX (IBM Platform MPI)**

Platform	Interconnect
	Myrinet-MX (SGI MPI) Infiniband (IBM Platform MPI)* Infiniband (SGI MPI)
win64	ethernet (MPICH2) ethernet (MS-MPI) ethernet (IBM Platform MPI) Infiniband (MS-MPI)

* Tested but not quality-assured support.

** No testing done at this time.

Expert MPI Options

IBM Platform MPI is the default MPI used on Linux and Windows. Typically, you should not need to work with IBM Platform MPI expert options.

However, if you wish to do so, refer to the IBM Platform MPI User's Guide (`pcmpi.XX.XX.XX.ug.pdf`) and Release Notes (`pcmpi.XX.XX.XX.rn.pdf`) that are located in the STAR-CCM+ installation directory:

- Linux: `[INSTALL DIR]/mpi/platform/X.X.X.X/[ARCHITECTURE]/doc/`
- Windows: `[INSTALL_DIR]\mpi\platform\X.X.X.X\[ARCHITECTURE]\help\doc\`

Running Distributed Parallel Simulations on Multiple Windows Machines

On Windows platforms you can run distributed parallel simulations using cores from multiple machines, in conjunction with IBM Platform MPI. To do this, use Windows PowerShell to launch STAR-CCM+ from within a shared directory, which is accessible to all machines that you want to use. You must have read/write access to the shared directory.

To run distributed parallel simulations on multiple machines:

1. Make sure that your credentials are registered on all machines that you want to use. See Registering Your Credentials for Distributed Simulations on Windows.
2. Click **<Start>** and type `powershell` in the search field.
3. Right-click on **Windows PowerShell** from the list of search results, and select **Run as Administrator**. You must have administrator privileges.

The *Administrator: Windows PowerShell* command prompt appears.

4. Change to the shared working directory:

```
> cd \\<MACHINE_NAME>\<SHARED_DIR>
```

5. Launch STAR-CCM+:
   - If STAR-CCM+ has the same installation path on each machine, enter:

```
> starccm+
```

   - If STAR-CCM+ has a different installation path on each machine, use the UNC path to the STAR-CCM+ executable:

```
> \\<MACHINE_NAME>\<STAR-CCM+_INSTALL_DIR>\starccm
+
```

   The STAR-CCM+ GUI appears.
6. Create or load a simulation file as you normally would. See Creating a New Simulation Interactively and Loading an Existing Simulation Interactively.
7. In the *Create a New Simulation* or *Load Simulation* dialog, select **Parallel on Named Hosts** from the *Process Options*. See Using Parallel on Multiple Host Machines.
8. Enter the host names and the corresponding number of cores to use from each host.



9. If you are opening an existing file, select the file that you want to load.
10. Click **OK**.

STAR-CCM+ creates or opens the file, and uses the allocated cores from each machine. The relevant information is displayed in the *Output* window.

Command-Line Reference

The `starccm+` command can start either the STAR-CCM+ client:

```
starccm+ [options] [simulation]
```

or server process:

```
starccm+ -server [options] [simulation]
```

Without any arguments, the STAR-CCM+ GUI client is started. No server process is started, as that is only required when a simulation is created or loaded.

If a `simulation` file is named without any options, both the STAR-CCM+ client and a server are started, and the named file is loaded into the server.

If the `-server` argument is specified, only the STAR-CCM+ server process is started. Any server option can also be specified on the command-line to start the client; the client passes those options on when it starts the server process.

General Options

Argument	Description
`-info`	Prints information about the simulation file: the version of STAR-CCM+ with which that file was saved (in this example 9.06.004), number of partitions, file size, and number of iterations in the file: `STAR-CCM+ 9.06.004 (linux-x86_64-2.5/gnu4.6) Serial 680.34 KB  Iter 66`
`-ini <file>`	Activates an `.ini` file that contains general and client options.
`-new`	Launches the STAR-CCM+ workspace and starts a new simulation without accessing a simulation file.

Argument	Description
`-loc`	For the client, this option turns off broadcasting of queries to locate servers on the network; it also prevents the server from responding to any locator queries. The broadcast from the client is a UDP multicast, along with TCP pings, to known hosts. This multicast search is started by default when you launch the STAR-CCM+ client without a simulation file. In Windows this option can cause network access warnings from your personal firewall.
`-v, -vv, -vvv`	Activates Verbose mode, prints environment changes and sub-commands.
`-version`	Prints the STAR-CCM+ version information (in this example 9.06.004): `9.06.004`

License Options

Argument	Description
`-lite`	Requests a Lite Session license as opposed to a full license. Only a reduced functionality set is available.
`-power`	Starts STAR-CCM+ with the Power Session license feature.
`-nosuite`	Prevents main licenses from being used to enable cores for parallel computation. Only parallel licenses are used to enable cores for parallel computation.
`-powerpre`	Uses 10 `DOEtoken` licenses to launch a STAR-CCM+ pre- and post-processing session.
`-tokensonly`	Specifies to only use `DOEtoken` licenses for optimization studies. This option has no effect when launching STAR-CCM+ in a non-Optimate/ Optimate+ environment.
`-notokens`	Specifies to not use `DOEtoken` licenses for optimization studies.
`-doe-prefer-hpcdomains`	Changes the order in which the `DOEtoken` license feature is checked out. Using this option checks out the `hpcdomains` license feature before the `DOEtoken` license feature for optimization studies.

Argument	Description
`-podkey <value>`	Specifies a Power on Demand license key.
`-doeuuid <value>`	Specifies a DOE UUID (universal unique identifier).
`-licpath <path:...>`	Specifies a license path that overrides the default license path.
`-reserve <lic1[,lic2,...]>`	Allows you to specify add-on licenses that can be automatically checked out.
`-noreserve`	Prevents add-on licenses from being checked out unless a model operation requires them.
`-noretry`	Prevents add-on licenses from being retried when a significant operation is carried out.
`-norelease`	Prevents add-on licenses from being released after importing geometry.

Client Options

Argument	Description
`-batch [command]`	Starts STAR-CCM+ in batch mode (without a GUI). The following additional commands are available: • `mesh` — runs the meshing pipeline • `run` — runs the solver until the stopping criteria are satisfied • `step` — steps the solution • `<macro>.java` — runs the specified Java macro file If no command is specified, the default action is to run the solver. In each case, the simulation is auto-saved through the auto save feature (unless that property is turned off). You can specify multiple commands within the `-batch` argument. In the following example, three macros are run in order before running the mesh pipeline and the solver: `-batch macro1.java,macro2.java,macro3.java,mesh,run`

Argument	Description
-batch-report	Generates a summary report of a batch simulation. The report appears in .html format. This report includes details about the computer system where the simulation was run, as well as details about the simulation itself.
-classpath <classpath>	Overrides the *User macro classpath* setting from the *Options* dialog. This option only affects the macro compilation in the current session of STAR-CCM+.
-clientldlibpath <path>:<path>:<path>	This option permits you to add directory paths explicitly to the STAR-CCM+ client library search path (LD_LIBRARY_PATH) for the client runtime environment.
-config <id>	Accesses a saved configuration of server options. (These options are also listed in the Saved Configurations list of the dialog that loads a simulation or creates a new one.)
-save-config <id>	Saves a configuration of server options and adds it to the Saved Configurations list of the dialog that loads a simulation or creates a new one.
-fontsize <size>	Adjusts the size of the font that is used in the STAR-CCM+ client workspace, such as dialogs and the object tree, for that session. Add the desired font size with this command. For example, a command of starccm+ -fontsize 16 sets the font size to 16 points.
-host <host> -port <port>	Specifies the host and port of a running STAR-CCM+ server process to which the client connects. If the -host and -port options are used, a simulation file cannot be specified.
-locale <locale>	Overrides the default locale <en:US, ja, ru:RU>.
-macro <file>	Runs the macro file upon client startup. This option is similar to -batch in that a macro file is read, but in this case a GUI is opened. This option requires specification of a simulation file. Since it is impossible to record a macro that includes the opening of a simulation file, there is never a conflict.

Argument	Description
-nbuserdir <dir>	Lets you specify a location, other than $HOME, in which the .star-*-*-* directory is written. Enter this command after client startup. This option changes the location of key STAR-CCM+ information, including the user settings directory and client log messages. Consider using this option when you do not have write access to your $HOME directory, or when you want to share modules.
-noexit	Prevents the server process from shutting down after batch file execution has completed. By default, the client and server processes both shut down after the macro completes.
-on <host>	Specifies the hostmachine on which to start the server process. By default, when a new simulation is created or an existing simulation file is opened, the STAR-CCM+ client starts a server process on the local machine. If the -on option is specified however, the client attempts to start the server process on the named host machine using rsh. Additionally, this option can be used directly as a parallel option to specify the hosts that the parallel server uses.
-reset	Restores the default client layout (including window positions) and settings (options and mounted directories).

-time	Prints server-elapsed and CPU times after a batch run. Specifically, it prints a summary of computational resource use of the server on close-down.
	In serial:
	`Saving: /home/ncr/sims/suprBump@00077.sim`
	`Server Elapsed Time :  5.362  (s)`
	`Server CPU Time  :  0.830  (s)`
	`Closing connection to server and exiting`
	In parallel:
	`Auto save: /home/ncr/sims/suprBump@00077.sim`
	`Saving: /home/ncr/sims/suprBump@00077.sim`
	`|Global Rank|Server Elapsed Time (s)|Server CPU Time (s)|MPI Overhead (s)|MPI Overhead %Elapsed Time (s)|`

Argument	Description

	`|  0|`
	`7.120|  1.250|`
	`0.461|  6.476|`
	`|  1|`
	`7.120|  5.150|`
	`2.371|  33.302|`
	`|  2|`
	`7.120|  5.120|`
	`2.425|  34.056|`
	`|  3|`
	`7.120|  5.160|`
	`2.451|  34.421|`
	`Maximum worker communication time on rank : 3`
	`Minimum worker communication time on rank : 0`
	`Average worker communication time : 1.927`
	`Standard deviation worker communication times : 0.847`
	`Closing connection to server and exiting`

`-userdir <name>`	On Linux, adds a user settings directory to your `home` directory with `name` optionally added to the folder name.
`-nbuserdir <name>`	On Linux, specifies the NetBeans directory.
`-clientcore`	On Linux, reserves one core for running the client in a batch system. Use this flag if you perform a lot of post-processing while running in parallel.
`-clientldlibpath <path:...>`	On Linux, adds directory paths explicitly to the STAR-CCM+ client library search path, `LD_LIBRARY_PATH`.
`-clientldpreload <file:...>`	On Linux, sets the client `LD_PRELOAD` environment variable.

Server Options

Argument	Description
`-arch <architecture>`	This option allows you to override the automatic determination of the architecture of the machine. The specified architecture must exist within the

Argument	Description
	installation and must be compatible with the machine architecture. If these conditions are not met, STAR-CCM+ exits, or runtime errors occur. For example, this argument can work when a 32-bit Linux installation is executed on a 64-bit Opteron-based architecture: `-arch linux-x86-2.2.5`
`-batchsystem <system>`	This option can be used to aid the startup of parallel STAR-CCM+ under the control of a supported Batch Management System (for usage see Using Batch Management Systems). Currently the supported systems are: • PBS (`<system>=pbs`) • LSF (`<system>=lsf`) • Sun Grid Engine (`<system>=sge`) • Load Leveler (`<system>=ll`)
`-collab`	In the startup of a server, allows others to connect to it. (By default only the user who starts a STAR-CCM+ server is allowed to connect to it.)
`-cout`	Messages are directed to the console rather than to the STAR-CCM+ client.
`-ldlibpath <path>:<path>:<path>`	This option permits you to add directory paths explicitly to the STAR-CCM+ server library search path (`LD_LIBRARY_PATH`) for the server runtime environment. The `-ldlibpath` option is required when STAR-CCM+ relies upon third-party software that is installed in custom locations. In addition, consider using it for user code in the following instances: • Linking against a math library like IMSL or NAG. (Getting access to the functions in them is one reason to write user code.) • Most compilers that can build a `.so` file dynamically link it against their own runtime libraries. Add paths to these libraries, unless they happen to be the system libraries (as they are with GNU compilers on Linux).

Argument	Description
`-load`	Loads the specified `simulation` file as soon as the server starts. Without this option, the simulation database is loaded on demand when the client connects to the server.
`-msjob[:<opt>]`	Submits a cluster job to the Windows HPC Server job scheduler. This functionality is available only on Windows HPC Server 2008 machines.
`-port <port>`	Specifies server port for STAR-CCM+ to open.
`-portrange <range>`	Specifies the number of ports the server can attempt to bind to.
`-server`	Starts the STAR-CCM+ server with an automatically generated port number.
`-rsh <rsh command>`	Specifies the remote shell command with which to start remote servers and parallel servers. The default remote shell command is rsh.

Visualization

Argument	Description
`-hardwarebatch`	Uses hardware acceleration for rendering graphics scenes when running in batch. This option can provide faster performance when post-processing large numbers of scenes. Hardware drivers must be installed properly and be in use on the machine where the client process of STAR-CCM+ is running. For Linux/Unix machines, run an X server on the machine where the client process of STAR-CCM+ is running. This X server could be your interactive X session, or another X session that you have permission to run.
`-mesa`	Uses the alternative mesa graphics libraries in the following situations: • There is a problem with the graphics driver • You use software rendering for: ◦ Running in batch ◦ Displaying a client remotely with an X-server

Parallel Options

Argument	Description
-machinefile <file>	MPI machine resource file name. Only used with parallel operation. The processes are allocated to the machines in the resource list starting with the first entry. If the number of requested processes (defined by -np) exceeds the number of machines that are listed, the remaining processes continue to be allocated starting with the first entry in the resource list. Cannot be used with -on.
-mpi <driver> -mpidriver <driver>	Specifies the MPI driver to use <mpich2, ms, platform>. If omitted, STAR-CCM+ uses a default driver. To specify the HP MPI distribution enter: -mpi hp2 In addition, this command can help if you do not know the specific inter-connect hardware available on a machine. When activated, HP MPI chooses the fastest transport layer that it detects on the machine.
-mpiflags <options>	MPI driver expert options.
-mpitest [:<options>]	Performs MPI tests and exits.
-np <np>	Number of processes to use. Only used for parallel operation. If the -np option is omitted and -machinefile or -on is specified, the number of processes is assumed to be the number that the resource list prescribes. If neither -on or -machinefile is specified with this option, then -on localhost:<np> is assumed.
-np <np[,npcosim]>	Number of processors to use for co-simulation, where np is the number of processes for STAR-CCM+ and npcosim is the number of processes for the other tool.
-on <host[:np][,host[:np]]>	List of hosts to run parallel processes on. If :np is specified, then np processes are put on that machine. There must be no spaces in the list of hosts. The processes are allocated to the machines in the resource list starting with the first entry adding np processes before moving on to the next machine in the list. If the number of requested processes (defined by -np) exceeds the number of machines that are listed, the remaining processes continue to

Argument	Description
	be allocated starting with the first entry in the resource list. Cannot be used with `-machinefile`.
`-pio`	Activates parallel input/output on systems that support it. When this flag is included, the given simulation file is restored using parallel I/O, and all subsequent simulation files are saved using parallel I/O.
`-cpubind <policy>`	Specifies what binding policy `<off`, `bandwidth`, `latency>` to use when binding CPUs. You can request a report to display the binding layout by entering: `-cpubind v`

See Also

Working with .ini Files
Using the Power Session License Feature
Using Power Tokens for Pre- and Post-Processing Sessions
Controlling the Check Out Order of DOE Parallel Licensing
Power-on-Demand Licensing
STAR-CCM+ Add-on Modules
Generating a Batch Summary Report
Using Parallel I/O
CPU Binding

Working with .ini Files

You can use an `.ini` file as a shortcut for general and client command line options when starting local STAR-CCM+ simulations.

To use this feature, add the `-ini` option to your line command:

```
% starccm+ -ini myinifile.ini <simfile>
```

The file contains simple listed commands, with any line beginning with `#` considered a comment. For instance, if you usually prefer to use `ssh` rather than `rsh` on your system, the `.ini` file could look like the following:

```
# set default rsh command
-rsh ssh
```

A more complex `.ini` file could be constructed as follows:

```
# set default rsh command
-rsh ssh
# set batchsystem and always allocate the controller
process to a cpu/core shared with a worker process
-batchsystem pbs:
# MPI options
#   - HP MPI 2.3.0
#   - SilverStorm Infiniband Driver using UDAPL layer
#   - Bind processes to cores
-mpidriver hp2:"-prot -UDAPL -cpu_bind=MAP_CPU:
0,1,2,3,4,5,6,7"
```

Arguments that are passed on the command line take precedence over the ones that are listed in the `.ini` file. For example, suppose that you issue the following command line:

```
% starccm+ -ini myinifile.ini <simfile> -rsh rsh
```

The `-rsh` argument in the command line takes precedence over the one in the `.ini` file. Currently there is no way to merge these two arguments.

Your system administrator can create a global `.ini` file that applies these options to all servers.

See Also

Creating a Global .ini File

Creating a Global .ini File

System administrators can create a global `.ini` file that contains command line options that affect all users. This ability is particularly useful if there are system-wide options that all users should be using.

For a list of command line options, refer to Command Line Reference.

Create the global `.ini` file as described in Working with .ini Files and place it in the following directory of your STAR-CCM+ installation:

```
[INSTALL_DIR]/config/star/config/ini
```

Name the `.ini` file according to the architecture of the system where it is installed, for example:

```
[INSTALL_DIR]/config/star/config/ini/starccm+-<arch>.ini
```

Where `<arch>` is the architecture of your system.

Naming the Linux Global .ini File

On Linux, you can find the architecture by using `maparch`, which is located in `[INSTALL_DIR]/config/star/bin/maparch`. An example command to run `maparch` is:

```
% ./star/bin/maparch
linux-x86_64-2.5
```

The `.ini` file is then named:

```
[INSTALL_DIR]/config/star/config/ini/starccm+-linux-
x86_64-2.5.ini
```

Naming the Windows Global .ini File

On Windows, the `.ini` file is named:

```
[INSTALL_DIR]/config/star/config/ini/starccm+-
windows.ini
```

See Also

Command-Line Reference
Activating Parallel I/O

Setting Conditions and Values

Along with types, conditions and values are used to provide to the continuum models complete information for dealing with boundaries, regions, and interfaces.

Types are covered in detail in a separate section. In addition to explanations of conditions, profiles and values, this section describes how to set conditions and values for the regions, boundaries, and interfaces. Setting conditions generally includes selecting a method for specifying the value, and then values can be specified using:

- The composite vector method
- Constants
- Composite arrays
- Constant arrays
- Field functions
- Expressions
- Radial equilibrium (for pressure profiles)
- Normally distributed values
- Normally distributed arrays
- Tables
- Feature Curves
- Uniformly distributed values
- Uniformly distributed arrays
- User code
- The composite symmetric tensor method
- The composite tensor method
- The principal tensor method

- The axisymmetric tensor method
- The isotropic tensor method
- The polynomial method

It is possible to do the following:

- Copying properties from one value or condition node to another
- Defining dimensionality of a value or unit
- Setting several conditions or several values of the same type at the same time.

See Also

Types of Inline Editors

What Are Conditions?

Although the concept of a type is fairly unambiguous, models need additional information to deal with the type. Conditions provide this information.

For example, with a boundary of type wall, the conditions tell you whether this wall is to be a no-slip wall, a slip wall, or a moving wall. The conditions also tell you whether you want to apply a specified temperature (Dirichlet) thermal boundary condition or a specified heat flux (Neumann) thermal boundary condition.

What Are Profiles?

There are many conditions (boundary, initial, region, interface) that are entered as profiles, instead of as simple numbers since different values can be set at different points on, for example, the boundary.

The five basic types of profiles are:

- Scalar
- Pressure
- Scalar Array
- Vector
- Tensor

If you want to use the same value at all points on the boundary, then you can enter constant as the profile. You also have the flexibility of setting conditions that are dependent on other values (using field functions), or of reading in a table of values from a file.

Using Scalar Profiles

Scalar profiles are used to specify scalar quantities such as Mach number, temperature, and mass flow.

In some cases, the node exists in the **Values** node independent of the corresponding **Conditions**. However, in other cases, the node is added to the **Values** node according to the method selected in the **Conditions** node. For example, mass flow can be entered as mass flux or as mass flow rate.

Properties

Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	
Method	Selects the method to use for entering the scalar value.	
	Constant	Specifies the scalar value as a single, constant value. A **Constant** node is

		added as a child to this node.
	Field Function	Defines the scalar value using a field function (typically user defined). A **Function** node is added as a child to this node.
	Normally distributed	Defines the scalar value randomly using a normal distribution. A **Normally distributed** node is added as a child to this node.
	Table (iteration)	Defines the scalar value as a function of iteration number. A **Table (iteration)** node is added as a child to this node.
	Table (r)	Defines the scalar value as a function of radius. A **Table (r)** node is added as a child to this node.
	Table (time)	Defines the scalar value as a function of physical time. A **Table (time)** node is added as a child to this node.
	Table (x,y,z)	Defines the scalar value as a function of position. A **Table (x,y,z)** node is added as a child to this node.
	Table (xyz, time)	Defines the scalar value as a function of position and time. A **Table (xyz, time)** node

		is added as a child to this node.
	Tabular	Defines the value using tabular data. A **Tabular** node is added as a child to this node.
	Uniformly distributed	Defines the scalar value randomly using a uniform distribution. A **Uniformly distributed** node is added as a child to this node.
	User Code	Defines the scalar value using a user subroutine. A **User Code** node is added as a child to this node.

Using Pressure Profiles

Pressure profiles are similar to scalar profiles except that they have an additional option in the *Method* property—`Radial Equilibrium`.

There is no corresponding node in the **Conditions** node.

Properties

Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	
Method	Selects the method to use for entering the pressure.	
	Constant	Specifies the pressure as a single, constant number. A **Constant** node is added as a child to this node.
	Field Function	Defines the pressure using a field function (typically user

		defined). A **Function** node is added as a child to this node.
	Radial Equilibrium	Defines the pressure using a radial equilibrium condition.
	Table (iteration)	Defines the pressure as a function of iteration number. A **Table (iteration)** node is added as a child to this node.
	Table (r)	Defines the pressure as a function of radius. A **Table (r)** node is added as a child to this node.
	Table (time)	Defines the pressure as a function of physical time. A **Table (time)** node is added as a child to this node.
	Table (x,y,z)	Defines the pressure as a function of position. A **Table (x,y,z)** node is added as a child to this node.
	Table (xyz, time)	Defines the pressure as a function of position and time. A **Table (xyz, time)** node is added as a child to this node.
	User Code	Defines the pressure using a user subroutine. A **User Code** node is added as a child to this node.

Using Scalar Array Profiles

A scalar array is used to specify multiple values of mass-fraction or volume-fraction where each value corresponds to one component of the mixture.

Properties

Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	
Method	Selects the method to use for entering the scalar value.	
	Composite	Defines the method as a composite array profile, which serves as a placeholder for separate scalar nodes.
	Constant	Defines the method as a single constant array.

Using Vector Profiles

A vector profile is used to specify a vector quantity.

Typical vector quantities are momentum source, flow direction, and velocity. This profile's *Coordinate System* property lets you switch from the default laboratory coordinate system to one of multiple local coordinate systems, depending on how you want to set the conditions in your simulation.

Properties

Coordinate System	The coordinate system to use.	
	Cartesian	Specifies the Cartesian coordinate system.
	Cylindrical	Specifies the cylindrical coordinate system.

	Laboratory	Specifies the laboratory coordinate system.
	Spherical	Specifies the spherical coordinate system.
Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	
Method	Selects the method to use for entering the vector.	
	Composite	Defines the vector as three (for example, X, Y, Z) components (or two (for example, X, Y) for two-dimensional problems).
	Constant	Specifies the vector as a single three-part (or two-part for two-dimensional), comma-separated, number. A **Constant** node is added as a child to this node.
	Feature Curve	Specifies the curve parts, representation, and orientation from which the local direction is calculated internally. A **Feature Curve** node is added as a child to this node.
	Field Function	Defines the vector using a field function (typically user defined). A **Field Function** node is added as a child to this node.

	Normally distributed	Specifies the vector randomly using a normal distribution. A **Normally distributed** node is added as a child to this node.
	Table (iteration)	Defines the vector as a function of iteration number. A **Table (iteration)** node is added as a child to this node.
	Table (r)	Defines the vector as a function of radius. A **Table (r)** node is added as a child to this node.
	Table (time)	Defines the vector as a function of physical time. A **Table (time)** node is added as a child to this node.
	Table (x,y,z)	Defines the vector as a function of position. A **Table (x,y,z)** node is added as a child to this node.
	Table (xyz, time)	Defines the vector as a function of position and time. A **Table (xyz, time)** node is added as a child to this node.
	Uniformly distributed	Specifies the vector randomly using a uniform distribution. A **Uniformly distributed** node is added as a child to this node.
	User Code	Defines the vector using a user

		subroutine. A **User Code** node is added as a child to this node.

Using Tensor Profiles

A tensor profile is used to specify tensor quantities such as porous resistances.

Properties

Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	
Method	Selects the method to use for entering the tensor.	
	Axisymmetric Tensor	Defines the method as an axisymmetric tensor profile, which serves as a placeholder for separate scalar nodes.
	Composite Symmetric Tensor	Defines the method as a composite symmetric tensor profile, which serves as a placeholder for separate scalar nodes.
	Composite Tensor	Defines the method as a composite tensor profile, which serves as a placeholder for separate scalar nodes.
	Isotropic Tensor	Defines the method as an isotropic tensor profile, which serves as a placeholder for separate scalar nodes.
	Principal Tensor	Defines the method as a principal tensor profile, which serves

		as a placeholder for separate scalar and vector nodes.

What Are Values?

The types and conditions inform models how to deal with a boundary, a region, or an interface, but they do not specify actual numerical input. Values provide these numbers.

Some examples of how certain conditions require values follow:

- The region condition specifies that a volumetric heat source must be applied to a region with an energy model, but the value specifies the magnitude of the volumetric heat source.
- The boundary condition specifies that a wall is to have a specified heat flux, but the value specifies the magnitude of the heat flux.
- The interface condition specifies that a thermal contact resistance (say due to a layer of paint) is associated with a contact interface between a solid region and a fluid region, but the value specifies the value of the contact resistance.

Setting Conditions

To set the conditions for a region, boundary or interface, open the node corresponding to the specific region, boundary, or interface.

Then open the **Conditions** node and select the required condition. In the *Properties* window, select one of the available options from the drop-down list corresponding to the *Method* property.



Setting Profile Methods

To set a value for a region, boundary or interface, first set the condition. Once the condition is set, open the **Values** node and select the node of the required profile.

By default the simplest method is assumed, a constant method for most profile types. To change this method, select one of the items (`Table (iteration)` in the following screenshot) in the drop-down list corresponding to the *Method* property in the *Properties* window.



Setting Values Using the Composite Vector Profile Method

The composite vector profile method differs from the others, as its node serves as a placeholder for the scalar profile quantity nodes. This feature lets you specify profile methods for X, Y, Z components separately.

To do this, open the **Composite** node and work with the scalar profile quantities as you did with the profile node.



Using the Scalar Profile Quantity

Three of these nodes appear as X, Y, Z components, or two (X, Y) for two-dimensional problems, when the composite vector profile method is selected.



Properties

Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	
Method	Selects the method to use for entering the scalar value.	
	Constant	Specifies the scalar value as a single, constant value A **Constant** node is added as a child to this node.
	Field Function	Defines the scalar value using a field function (typically user defined). A

		Function node is added as a child to this node.
	Table (iteration)	Defines the scalar value as a function of iteration number. A **Table (iteration)** node is added as a child to this node.
	Table (r)	Defines the scalar value as a function of radius. A **Table (r)** node is added as a child to this node.
	Table (time)	Defines the scalar value as a function of physical time. A **Table (time)** node is added as a child to this node.
	Table (x,y,z)	Defines the scalar value as a function of position. A **Table (x,y,z)** node is added as a child to this node.
	User Code	Defines the scalar value using a user subroutine. A **User Code** node is added as a child to this node.

Setting Values Using a Constant

After specifying the method of a profile:

1. Open the node of that profile.
2. Select the **Constant** node
3. In the *Properties* window, enter the appropriate values (scalar, pressure, or vector).



Techniques also exist for using units while setting values. Many constant values, particularly ones with units, can be entered as expressions.

Properties of Scalar and Pressure Constants

Value	Specifies the value with a text entry. You can also enter an expression, such as `10 * sin(3.14 * $Time)`, directly.

Properties of a Vector Constant

Value	Specifies the x, y, and z values. These values are separated by commas in a text entry. You can also enter an expression, such as `10 * sin(3.14 * $Time)`, directly. The property customizer opens a dialog which contains separate fields for these values.

Setting Values Using a Composite Array

The composite array profile method node serves as a placeholder for separate nodes, each of which represents a selected component in a mixture.

The component nodes function just like the scalar profile quantity nodes in a typical composite vector profile method. This feature lets you specify profile methods for each component separately.

To do this, open the **Composite** node and work with the scalar profile quantities as you did with the profile node.



Setting Values Using a Constant Array

The constant array node differs from a typical scalar constant node in that its properties can contain multiple values. The number and sequence of these values correspond to the number and sequence of the components of the array.

After specifying the method of a profile,

1. Open the node of that profile.
2. Select the **Constant** node.
3. In the *Properties* window, enter the appropriate values.



Properties of a Constant Array

Value	Specifies a series of comma-separated values with a text entry.

Setting Values Using a Field Function

After specifying the method of a profile:

1. Open the node of that profile.
2. Select the **Field Function** node.
3. Select the appropriate field function from the object selector dialog for a scalar or vector function.



Properties of a Scalar Field Function Method

Scalar Function	Activates an object selector dialog from which you can choose a field function.

Properties of a Vector Field Function Method

Vector Function	Activates an object selector dialog from which you can choose a field function.

Setting Values Using an Expression

In addition to constants, you can enter expressions that use field function syntax, such as `10 * sin(3.14 * $Time)`, directly. This is an easy alternative to creating a macro for calculations of non-constant physical values.

Examples of applications include:

- Valve openings or closings
- Tank sloshing
- A vehicle starting at rest and accelerating to full speed
- Prescribed rigid body motion
- A varying gravity vector
- A solar radiation vector

Simply type the function definition into the property value as you would a constant. The example in the following screenshot specifies oscillation with a period of 2 seconds.



This technique can also be used for vector entries.



Alternatively, you can create a field function that uses the expression, and then copy the name of the field function into the property.



The name comes from the *FunctionName* property of the field function. When you enter it into a constant value property, prefix it with `$`. In the preceding example, the function name of the field function is `SolarFlux`.

At the same time, other objects in the simulation can reference these field functions.

The field functions that can be used to define the expressions cannot be a function of space; they have to evaluate to a single number as the expressions are constants in space. Variables such as `$Time`, `$TimeStep` and `$Iteration`, as well as the field functions associated with reports, are valid. Variables such as `$Temperature` or `$Pressure` have different values at every different point in space and cannot be used in expressions.

Using Units

You can include units in these expressions, for example:

- `$Time - <1 min>`
- `<180 deg>`
- `<65 mph>`

See Also

Using a Property Customizer

Setting Values Using Radial Equilibrium

This node represents the location of the amount of pressure at a particular part of the model.

After specifying the method of a profile, open the node of that profile, select the **Radial Equilibrium** node and enter the appropriate value in the *Hub Pressure* property.



Properties

Hub Pressure	Specifies the pressure.

Setting Values Using a Normal Distribution

Use this method to set the value of a scalar randomly using a normal distribution.

After specifying the method of a profile, open the node of that profile, select the **Normally distributed** node and enter the appropriate values in the *Properties* window.



Properties of a Normally Distributed Method

Mean	Specifies the mean value of the scalar. Clicking the ... property customizer button activates a property customizer that enables you to specify both the value and the appropriate units.
Deviation	Specifies the deviation of the scalar. Clicking the ... property customizer button activates a property customizer that enables you to specify the value.
Minimum	Specifies the maximum value of the scalar. Clicking the ... property customizer button activates a property customizer that enables

	you to specify both the value and the appropriate units.
Maximum	Specifies the minimum value of the scalar. Clicking the ... property customizer button activates a property customizer that enables you to specify both the value and the appropriate units.

Setting Values Using a Normal Distribution Array

Use this method to set the value of a vector randomly using a normal distribution.

After specifying the method of a profile, open the node of that profile, select the **Normally distributed** node and enter the appropriate values in the *Properties* window.



Properties of a Normally Distributed Method

Mean	Specifies the mean values of the x, y, and z components of the vector. These values are separated by commas in a text entry. Clicking ... (**Property Customizer**) activates a property customizer that enables you to specify both the values and the appropriate units.
Deviation	Specifies the deviations of the x, y, and z components of the vector. These values are separated by commas in a text entry. Clicking ... (**Property Customizer**) activates a property customizer that enables you to specify the values.
Minimum	Specifies the maximum values of the x, y, and z components of the vector. These values are

	separated by commas in a text entry. Clicking (**Property Customizer**) activates a property customizer that enables you to specify both the values and the appropriate units.
Maximum	Specifies the minimum values of the x, y, and z components of the vector. These values are separated by commas in a text entry. Clicking (**Property Customizer**) activates a property customizer that enables you to specify both the values and the appropriate units.

Setting Values Using a Table

There are several common types of tabular methods, each represented by nodes.



These types are:

- Table (iteration)
- Table (particle size CDF)
- Table (r)
- Table (time)
- Table (x,y,z)
- Table (xyz, time)
- Tabular

The first three have similar properties. The properties of the Table (x,y,z) and the Table (xyz, time) are similar in that they are structured to

accommodate the three coordinates, though the Table (xyz, time) also takes simulation time into account.

Additionally, there are two specific table types, Table (er) and Table (particle size CDF), which use the same properties as the first three types that are listed previously.



Although a table data file is treated simply as a table when it is read into STAR-CCM+, each of these methods looks for a particular type of column in the table that accompany the data values.

After specifying the method of a profile as a table, open the profile node (**Mach Number** in the following screenshot) if necessary, select the tabular node. Set the appropriate properties from the drop-down lists corresponding to *Table: Data* and *Table*. You must have read a table into STAR-CCM+ to do this; otherwise these drop-down lists are empty and the *Table: Data* and *Table* properties show as `Null`. The other property (*Table (iteration)* in the following screenshot), is read-only since it simply indicates the title of the column that the tabular method expects.



Separate table value nodes exist for setting vector profiles.

Sections below describe examples of the following:

- Creating a table for a profile
- Extracting and exporting a table
- Using data with sharp discontinuities

Setting Vector Profiles Using a Table

You can also use a single table method to specify a vector profile.



Separate *Data* properties exist for the x-, y-, and z-components. As is the case with the table methods for scalar and pressure profiles, properties are customized for iteration, radius, and time, and for (x,y,z) and (xyz, time).

What Is Table (iteration)?

An (iteration) table has a column of 'iteration' values that specify the iteration of each point (row) in the table.

The points must be ordered by increasing values of iteration. Multiple methods are available for interpolating between data points in the table, as described in the properties section. A profile in iteration could also be specified analytically with a user-defined field function that uses the $Iteration variable.

The following is a simple example of how an (iteration) table might look with actual data:

```
iteration,Value
5,0.5
40,1.0
100,1.5
```

What Is Table (r)?

An (r) table has a column of 'R' values that specify the radius, in default units, of each point (row) in the table.

The points must be ordered by increasing values of radius. A B-spline is used on the data to calculate face and cell values, requiring extra measures for data with sharp discontinuities.

Radius is computed relative to the frame of reference of each region, which is specified in the region values. The axis from which the radius is computed is specified using the Axis node of the region.

If a Rotation or Rotation and Translation reference frame is specified under the **Motion Specification** node, the axis that is used to define that reference frame is used to calculate the radial coordinate for the region.

For example, to initialize velocity using an (r) table:

- Import the table.

  The tabular data can look as follows:



Tabular Data

	R	Velocity Magnitude
0	0	62.95
1	4.9968E-4	62.54
2	9.9936E-4	61.36
3	0.0015	59.21
4	0.002	56.73
5	0.0025	53.34
6	0.003	48.8
7	0.0035	41.99
8	0.0036	0
9	0.004	3.45
10	0.0045	11.46
11	0.005	15.18
12	0.0055	15.45
13	0.006	15.15
14	0.0065	15.97
15	0.007	15.56

Export... Close

Specify the axis for the region that is used to define the radial coordinate:

1. Select the **Regions** > **[Region]** > **Physics Values** > **Motion Specification** node and make sure that *Reference Frame* is set to **Lab Reference Frame**.

   The **Physics Values** > **Axis** node for the region appears only when the *Reference Frame* is set to **Lab Reference Frame**.

   In this example, a local cylindrical coordinate system is chosen to define the axis for the region. The *Direction* `[0.0, 0.0, 1.0]` in the **Laboratory-** > **Cylindrical 1** system is rooted at *Position* `[0.0, 0.0, 0.0]`, as shown in the following screenshot.



2. Select the **Regions** > **[Region]** > **Physics Values** > **Axis** node and set the *Coordinate System* property.



In this example, the initial jet velocity component is set using a composite specification. A composite specification allows you to specify each component of the velocity using a different method. To set a component using an imported (r) table:

1. Select the **Continua** > **[Physics Continuum]** > **Initial Conditions** > **Velocity** node and set Method to **Composite**.
2. Select the **Velocity** > **Composite** > **[Component]** node and set *Method* to **Table (r)**.
3. Select the **[Component]** > **Table (r)** node and set the *Table* and *Table:Data* properties.



Similarly, to set the boundary condition:

1. Select the **Regions** > **[Region]** > **Boundaries** > **[Boundary]** > **Physics Values** > **Velocity Magnitude** node and set *Method* to **Table (r)**.
2. Select the **Velocity Magnitude** > **Table (r)** node and set the *Table* and *Table:Data* properties.

The initialized velocity for the flow field in this example looks as follows.



What Is Table (time)?

A (time) table has a column of 'time' values that specify the physical time, in seconds, of each point (row) in the table.

The points must be ordered by increasing values of time. Multiple methods are available for interpolating between data points in the table, as described in the properties section. This type of table is only applicable to unsteady calculations where "time" is defined.

Note: A profile in time can also be specified analytically with a user-defined field function that uses the `$Time` variable.

What Is Table (x,y,z)?

An (x,y,z) table has a column of 'X' values, 'Y' values and, in three-dimensional cases, 'Z' values that specify the position, in default units, of each point (row) in the table:

```
X Y Z VAR1 ... VARn
```

where, for each column header, `VARn` is the name of a particular variable.

The points are considered unordered in the file, and no interpolation is done to calculate face and cell values. Each face or cell is assigned the value from the closest point in the table.

What Is Table (xyz, time)?

Similar to the (x,y,z) table, an (xyz, time) table has a column of 'X' values, 'Y' values and, in three-dimensional cases, 'Z' values that specify the position, in default units, of each point (row) in the table. The points are considered unordered in the file.

In addition, this method causes simulation time to be considered when STAR-CCM+ reads the table. The data columns would appear as follows:

```
X Y Z VAR1[t=T1] VAR1[t=T2] VAR1[t=T3] VAR2[t=TT1]
VAR2[t=TT2]
```

where, for each column header, `VARn` is the name of a particular variable and `x` in `[t=x]` is the time stamp for the column. The file that is shown above contains two time series, `VAR1` and `VAR2`. `VAR1` has data points for all X,Y,Z at times `T1`, `T2` and `T3`, and `VAR2` has data points for times `TT1` and `TT2`. (The values of `Tn` and `TTn` for each `n` can be different.)

The following is a simple example of how an (xyz, time) table might look with actual data:

```
X Y Z Temperature(K)[t=1s] Temperature(K)[t=2s]
Temperature(K)[t=3s]
```

```
1 2 3 300 400 500
```

In the STAR-CCM+ GUI, the time stamps are removed and you are able to select a variable name for the data property in the method. As the simulation iterates, the profile is filled with values from the table based

on the time, using the interpolation method that is set for the *Interpolation* property.

What Is Table (er)?

An (er), or equivalence ratio, table has a column of 'er' values and 'lfs' (laminar flame speed) values. This method is available with the User Defined Laminar Flame Speed option of the Coherent Flame and Thickened Flame models.

This example shows the contents of a file with equivalence ratio data:

```
er      lfs
0.6     0.146586604
0.7     0.231662295
0.8     0.329079651
0.9     0.42048291
1       0.483519189
1.1     0.500554535
1.2     0.466632794
1.3     0.391806544
1.4     0.296351954
1.5     0.201946692
1.6     0.12399419
1.7     0.068601963
1.8     0.03420356
1.9     0.015368442
```

Example of Creating a Table for a Profile

This section illustrates how you would go about creating a table for a profile.

Suppose that you must specify a velocity inlet as a function of Z. If you have a few points of data, you can create a table and use it to specify the inlet.

First, create, and save your table file in a typical text editor, specifying the .csv format as the file extension.



Even if you only want calculations in relation to one direction, you need columns that correspond to the two others (or the other in a two-dimensional case). Hence “empty” columns must contain entries of zero.

Next, read the table file into STAR-CCM+ using the **Tools** > **Tables** node. In this example, the table file is named `profile.csv`.



Open the **Regions** > **[individual region node]** > **Boundaries** > **[individual boundary node]** > **Physics Values** nodes, and select the profile node (**Velocity Magnitude** in the following screenshot).



Set the *Method* property to a tabular method (`Table (x,y,z)` in the following screenshot).



As a result of this action, the method node inside the profile node changes to a tabular method node. Its label matches the selection in the drop-down list, in this case **Table (x,y,z)**. Select that node and set the *Table* property to the name of your table file. In this example, the name is `profile`.



Set the *Table:Data* property to the column of the scalar value, in this case `Velocity`.



When you use `Table (x,y,z)` as the tabular method, once you set the *Table* and *Table:Data* properties, the other properties for the X, Y, and Z coordinates are read-only, since they map automatically to the columns in your table file with those headings.

This completes the example of setting a profile with tables.

Example of Extracting and Exporting a Table

This section illustrates how you would extract a table from one simulation to use as a boundary condition in another.

In this example, a channel flow simulation with periodic boundaries is used to create a fully developed channel flow profile. Values from an x-y plane cut across the fully developed channel are extracted into a table, exported and then used at the inlet of another simulation. This obviates the need to create a long inlet section in the second simulation to obtain the required fully developed conditions.

Create an xyz internal table. Make sure that the node of the new table is selected.



In the properties, set the *Parts* property by clicking the right half to activate the in-place object selector. In this example, `plane section 1` is used.



Once you select the part, set the *Scalars* property. While the property allows text entries, it is much more convenient to use the property customizer ( ), which activates a scalars dialog.



This dialog works the same as the one you used for the part selection. In this example, `Velocity: X-Component`, `Turbulent Kinetic Energy`, and `Turbulent Dissipation Rate` are selected.

Extract the table data, then export the data to a `.csv` file.



In this example, the table is named `inlet_rke.csv`.

Open the simulation in which you want to use the profile. Then read the table file that you created in the other simulation.



In this example, the profiles that receive the data of this file are:

- Turbulent Dissipation Rate
- Turbulent Kinetic Energy
- Velocity: X-Component

They correspond to the scalars selected for the xyz table that was created earlier.

Open the **Regions > [individual region node] > Boundaries > [individual boundary node] > Physics Values** nodes, and select the profile node (**Turbulent Dissipation Rate** in the following screenshot).



Set the *Method* property to a tabular method (`Table (x,y,z)` in the following screenshot).



As a result of this action, the method node inside the profile node changes to a tabular method node. Its label matches the selection in the drop-down list, in this case **Table (x,y,z)**. Select that node and set the

Table property to the name of your table file. In this example, the name is `inlet_rke`.



Set the *Table:Data* property to the column of the scalar value, in this case `Turbulent Dissipation Rate`.



When you use `Table (x,y,z)` as the tabular method, once you set the *Table* and *Table:Data* properties, the other properties for the X, Y, and Z coordinates are read-only, since they map automatically to the columns in your table file with those headings.

In this example, in the same **Physics Values** node, these steps are repeated for **Turbulent Kinetic Energy**, its node depicted below.



In the same **Physics Values** node, a profile node for velocity is selected.



In this example, the velocity specification boundary condition has been set to `Components`, since only the X-component is being used in the table file. Hence, for this **Velocity** node, the *Method* property must be set to `Composite`, to use the composite vector profile method.

The next step, therefore, is to open the **Velocity** and **Composite** nodes. To use this approach, select the node of the components for which you want to plot data.



In this example, the X-component is selected, while the Y-component is a constant. (There is no Z-component because this case is two-dimensional.)

Set the *Method* property (`Table (x,y,z)` in the following screenshot).



As before, set the *Table* property to the name of your table file. Then set the *Table:Data* property to the column of the scalar value, in this example it is `Velocity X-Component`.

This completes the example of using a table file to transfer boundary conditions from one simulation to another.

Example of Using Data with Sharp Discontinuities

In earlier versions of STAR-CCM+, it was not possible to choose the Step interpolation option for tables based on iteration and time; all interpolation was based on a B-spline approach.

The B-spline approach is still used for tables based on a radial profile, and so some care must be taken to deal with any sharp discontinuities across the radial profile.

- The essential approach is to add additional data points before and after the sharp discontinuity so that STAR-CCM+ can more accurately interpolate data in that interval.
- It is also possible to generate data with sharp discontinuities by using a user-defined field function, and so avoiding the use of a table altogether.

```
($Time < t) ? 0 : 1e-4
```

Table Properties Iteration Radius and Time

Table: Data	Identifies the selected column in the table file as the data column.	
Table: [heading]	Looks for the column in the table with the heading `Iteration`, `R`, or `time` (read-only).	
Table	Selects a table from among the ones that are represented by nodes in the **Tables** node of the **Tools** node.	
Interpolation	(Iteration and Time table types). Selects the method that is used to interpolate between data points in the table.	
	Spline	Uses B-spline interpolation between table data values. This method is set automatically for Radial profiles.
	Step	Uses the table data value corresponding to the highest iteration or time that is less than or equal to the current iteration or time.
	Linear	Specifies a linear interpolation between table data values.

Properties of Table (x,y,z) and Table (xyz, time)

Table: X-Coordinate	Looks for the column in the table with the heading `X` (read-only).
Table: Y-Coordinate	Looks for the column in the table with the heading `Y` (read-only).
Table: Z-Coordinate	Looks for the column in the table with the heading `Z` (read-only).
Coordinate System	Specifies the coordinate system (laboratory by default). The drop-down list of this property includes any local coordinate systems that exist in the simulation.
Table: Data	Identifies the selected column in the table file as the data column.
Table	Selects a table from among the ones that are represented by nodes in the **Tables** node of the **Tools** node.

Expert Properties (Table (xyz, time) only)

Interpolation	Selects the method that is used to interpolate between data points in the table.	
	Spline	Uses B-spline interpolation between table data values.
	Step	Uses the table data value corresponding to the highest time that is less than or equal to the current time.
	Linear	Specifies a linear interpolation between table data values.

Vector Table Properties Iteration Radius and Time

Table: [heading]	Looks for the column in the table with the heading `Iteration`, `R`, or `time` (read-only).	
Table	Selects a table from among the ones that are represented by nodes in the **Tables** node of the **Tools** node.	
Table: X-Data	Identifies the selected column in the table file as the data column for the x-component.	
Table: Y-Data	Identifies the selected column in the table file as the data column for the y-component.	
Table: Z-Data	Identifies the selected column in the table file as the data column for the z-component.	
Interpolation	(Iteration and Time table types). Selects the method that is used to interpolate between data points in the table.	
	Spline	Uses B-spline interpolation between table data values. This method is set automatically for radial profiles.
	Step	Uses the table data value corresponding to the highest iteration or time that is less than or equal to the current iteration or time.
	Linear	Specifies a linear interpolation between table data values.

Properties of Vector Table (x,y,z) and Vector Table (xyz, time)

Table: X-Coordinate	Looks for the column in the table with the heading `x` (read-only).	

Table: Y-Coordinate	Looks for the column in the table with the heading `Y` (read-only).
Table: Z-Coordinate	Looks for the column in the table with the heading `Z` (read-only).
Coordinate System	Specifies the coordinate system (laboratory by default). The drop-down list of this property includes any local coordinate systems that exist in the simulation.
Table: X-Data	Identifies the selected column in the table file as the data column for the x-component.
Table: Y-Data	Identifies the selected column in the table file as the data column for the y-component.
Table: Z-Data	Identifies the selected column in the table file as the data column for the z-component.
Table	Selects a table from among the ones that are represented by nodes in the **Tables** node of the **Tools** node.

Expert Properties (Vector Table (xyz, time) only)

Interpolation	Selects the method that is used to interpolate between data points in the table.	
	Spline	Uses B-spline interpolation between table data values.
	Step	Uses the table data value corresponding to the highest time that is less than or equal to the current time.
	Linear	Specifies a linear interpolation between table data values.

Setting Values Using a Feature Curve

Use this method to specify a field of direction vectors using line curves or surfaces.

After specifying the method of a profile, open the node of that profile, select the **Feature Curve** node and enter the appropriate values in the *Properties* window.



This method is similar to the **interpolateDirection** field function.

Properties of a Feature Curve Method

Curve Parts	A surface or a line curve. For a surface, the direction field is given using the normals of that part. For a line curve, the direction field is given using the line segment directions.
Representation	The representation that is used to extract the part.

Direction	Specifies the orientation of the direction vectors.	
	FORWARD	Set the orientations of the vectors based on the direction of the normals of the input surface or the direction of the line segment of the input line.
	BACKWARD	Reverses the direction vectors.

Setting Values Using a Uniform Distribution

Use this method to set the value of a scalar randomly using a uniform distribution.

After specifying the method of a profile, open the node of that profile, select the **Uniformly distributed** node and enter the appropriate values in the *Properties* window.



Properties of a Uniformly Distributed Method

Minimum	Specifies the maximum value of the scalar. Clicking the ... property customizer button activates a property customizer that enables you to specify both the value and the appropriate units.
Maximum	Specifies the minimum values of the scalar. Clicking the ... property customizer button activates a property customizer that enables you to specify both the value and the appropriate units.

Setting Values Using a Uniform Distribution Array

Use this method to set the value of a vector randomly using a uniform distribution.

After specifying the method of a profile, open the node of that profile, select the **Uniformly distributed** node and enter the appropriate values in the *Properties* window.



Properties of a Uniformly Distributed Method

Minimum	Specifies the maximum values of the x, y, and z components of the vector. These values are separated by commas in a text entry. Clicking the ... property customizer button activates a property customizer that enables you to specify both the values and the appropriate units.
Maximum	Specifies the minimum values of the x, y, and z components of the vector. These values are separated by commas in a text entry. Clicking the ... property customizer button activates a property customizer that enables you to specify both the values and the appropriate units.

Setting Values Using User Code

After specifying the method of a profile, open the node of that profile, select the **User Code** node and select the appropriate user code function from the drop-down list in the properties.



Properties of User Code

Function	Provides a drop-down list from which you can choose a user function.

Setting Values Using a Composite Symmetric Tensor

The composite symmetric tensor profile method node serves as a placeholder for separate nodes.

Each node represents a component of a symmetric tensor: three in two-dimensional problems (XX, XY &YX and YY) or six in three-dimensional problems (XX, XY & YX, XZ & ZX, YY, YZ & ZY, and ZZ). The component nodes function just like the scalar profile quantity nodes in a typical composite vector profile method. This feature lets you specify profile methods for each component separately.

To do this, open the **Composite Symmetric Tensor** node and work with the scalar profile quantities as you did with the profile node.



For example, to specify the tensor

$$\begin{bmatrix} 10 & 15 \\ 15 & 20 \end{bmatrix}$$

set the **XX Component** to a *Constant* value of `10`, the **XY & YX Component** to `15`, and the **YY Component** to `20`.

Properties of a Composite Symmetric Tensor

Coordinate System	The coordinate system in which the tensor components are defined.

Setting Values Using a Composite Tensor

The composite tensor profile method node serves as a placeholder for separate nodes.

Each node represents a component of a tensor: four in two-dimensional problems (XX, XY, YX, and YY) or nine in three-dimensional problems (XX, XY, XZ, YX, YY, YZ, ZX, ZY, and ZZ). The component nodes function just like the scalar profile quantity nodes in a typical composite vector profile method. This feature lets you specify profile methods for each component separately.

To do this, open the **Composite Tensor** node and work with the scalar profile quantities as you did with the profile node.



For example, to specify the tensor

$$\begin{bmatrix} 10 & 15 \\ 25 & 20 \end{bmatrix}$$

set the **XX Component** to a *Constant* value of `10`, the **XY Component** to `15`, the **YX Component** to `25`, and the **YY Component** to `20`.

Properties of a Composite Tensor

Coordinate System	The coordinate system in which the tensor components are defined.

Setting Values Using a Principal Tensor

In the principal tensor profile method, a tensor is specified using its diagonal (principal) components with respect to its principal axes. The node serves as a placeholder for separate component and axis nodes.

The component nodes represent the diagonal components of a tensor: two in two-dimensional problems (XX and YY) or three in three-dimensional problems (XX, YY and ZZ). These nodes function just like the scalar profile quantity nodes in a typical composite vector profile method.

The axis nodes represent the directions of the principal axes of the tensor. Only two principal axes need to be specified (XX and YY), which must not be parallel to each other. The third principal axis is deduced internally as being orthogonal to the other two. The magnitude of the axis vectors is ignored; only their directions are important. If the entered axes are parallel to each other or are zero length (the degenerate condition), the principal axes fall back to the axes of the standard coordinate system. The axis nodes function just like vector profile quantity nodes, including the ability to specify a coordinate system for each axis vector. The coordinate system may often be the same for both axes but need not be.

The design allows you to specify profile methods for each component and axis separately.

To do this, open the **Principal Tensor** node and work with the scalar and vector profile quantities as you did with the profile node.



For example, to specify the tensor

$$\begin{bmatrix} 20 & 0 & 0 \\ 0 & 30 & 0 \\ 0 & 0 & 30 \end{bmatrix}$$

set the **XX Component** to a *Constant* value of `20`, and the **YY Component** to `30`. Then specify the **XX Axis** as `[1, 0, 0]` and the **YY Axis** to `[0, 1, 0]`.

Setting Values Using an Axisymmetric Tensor

In the axisymmetric tensor profile method, a tensor is specified using its components along a user-specified axis and those normal to the axis. The node serves as a placeholder for separate component and axis nodes.

The Axis node represents the axisymmetric direction. There are two component nodes: **Axial component** which is along the user-specified Axis, and **Cross-Stream Component** which is along the direction normal to the Axis. These nodes function just like the scalar profile quantity nodes in a typical composite vector profile method.



For example, to specify the tensor

$$\begin{bmatrix} 20 & 0 & 0 \\ 0 & 30 & 0 \\ 0 & 0 & 30 \end{bmatrix}$$

specify the **Axis** as a **Constant** value of `[1, 0, 0]`. Then set the **Axial Component** to `20`, and the **Cross-Stream Component** to `30`.

Setting Values Using an Isotropic Tensor

In the isotropic tensor profile method, a tensor is specified using a single component which remains invariant in all the spatial directions. The node serves as a placeholder for the isotropic component.

The isotropic component represents the single invariant component of the tensor. This functions just like the scalar profile quantity nodes in a typical composite vector profile method.



For example, to specify the tensor

$$\begin{bmatrix} 20 & 0 & 0 \\ 0 & 20 & 0 \\ 0 & 0 & 20 \end{bmatrix}$$

specify the **Isotropic Component** to a *Constant* value of `20`.

Setting Values Using a Polynomial

STAR-CCM+ allows a number of quantities to be specified using a polynomial.

Examples are density and C_p as polynomial functions of temperature. The properties for a polynomial are set in the *Properties* window, as in this example for a C_p polynomial.



A polynomial $y(x)$ is specified in piece-wise fashion where each interval i spans the range x_i to x_{i+1}, has m_i number of terms, and is of the form:

$$y(x) = \sum_{k=0}^{m_i} a_{i,k} \cdot x^{b_{i,k}};\ x_i \le x \le x_{i+1} \tag{1}$$

The polynomial has n intervals with interval ranges: $x_0, x_1, x_2, \ldots, x_{n+1}$

The number of terms in each interval are: $m_0, m_1, m_2, \ldots, m_n$

The exponents are:

$$\begin{matrix} b_{0,0}, & b_{0,1}, & b_{0,2}, \ldots & b_{0,m_0} \\ b_{1,0}, & b_{1,1}, & b_{1,2}, \ldots & b_{1,m_1} \\ \vdots & \vdots & \vdots & \vdots \\ b_{3,0}, & b_{3,1}, & b_{3,2}, \ldots & b_{3,m_3} \end{matrix}$$

The coefficients for the polynomial are:

$$\begin{matrix} a_{0,0}, & a_{0,1}, & a_{0,2}, \ldots & a_{0,m_0} \\ a_{1,0}, & a_{1,1}, & a_{1,2}, \ldots & a_{1,m_1} \\ \vdots & \vdots & \vdots & \vdots \\ a_{3,0}, & a_{3,1}, & a_{3,2}, \ldots & a_{3,m_3} \end{matrix}$$

You must specify a range of the polynomial sufficiently large such that the polynomial is not evaluated outside the range during the analysis. If the polynomial is ever evaluated outside the range, the property is assigned a value of the polynomial at the respective range limit.

A section below provides an example of working with this feature.

Polynomial Properties

In the *Properties* window, a single property that is called *Polynomial* includes a property customizer ( ).

Click that button to activate a customizer dialog with properties that correspond to the components of Eqn. (1).



These properties are grouped in a separate dialog so that when you click **OK**, STAR-CCM+ can check for consistency among the properties.

The plot in this dialog gets updated instantly with your changes. The *Input Units* property specifies units for the x-axis, while the *Output Units* property does so for the y-axis. Unlike constant values, the polynomial coefficients and polynomial ranges (and consequently x- and y-axes on the polynomial plot) do not rescale when you change the units. These unit selections serve as an indicator of the units that STAR-CCM+ uses when calculating quantities with this polynomial.

Properties that accommodate a series of entries have their own customizer buttons. When one of these buttons is clicked, it opens a text box with more space for entering values.



The property *Use Real Exponents* controls how polynomials are calculated. When the property is deactivated (the default setting), the exponents cannot be edited; they are determined by the number of coefficients in each range. When the property is activated, you can enter any values for exponents. However, entering exponent values manually can increase computing time considerably, because when coefficients are arbitrary, fast polynomial evaluation cannot be used.

For the *Coefficients* property, the entries are separated by commas and the coefficients are in the following order:

$$y(x) = a_0 + (a_1 \cdot x) + (a_2 \cdot x^2) + \ldots + (a_n \cdot x^n)$$

The following section provides an example of working with these properties using specific heat as a polynomial in temperature.

Example of Using a Polynomial

This section illustrates how you would use a STAR-CCM+ polynomial: the Polynomial in T.

In the *Properties* window, a single property that is called *Polynomial* includes a property customizer ( ). Click that button to activate a customizer dialog.



The plot in this dialog gets updated instantly with your changes.

Suppose that you have the following, defined with two intervals:

$$C_p = 1 + 2T,\ 100 \le T < 500$$

$$C_p = 3 + 4T + 5T^2,\ 500 \le T < 1000$$

First specify the *Number of Intervals*, in this case `2`.



Enter the *Interval Ranges*, `100,500,1000` in the illustration below.



You must specify a temperature range of the polynomial sufficiently large such that the polynomial is not evaluated outside the range during the analysis. If the polynomial is ever evaluated outside the range, the property is assigned a value of the polynomial at the respective range limit.

Specify the number of coefficients for each interval. In this example, the first interval has two coefficients and the second interval has three, so enter `2,3` in the *Number of Coefficients* property.



For the *Coefficients* property, specify all for each interval in succession. In this example, the entries are both coefficients for the first interval, `1,2`, followed by all three for the second, `3,4,5`.



Enter the exponents using the same technique as for the coefficients, in this example `0,1,0,1,2`.



Work with the units as needed. You can select units separately for each axis of the graph: *Input Units* for the x-axis, and *Output Units* for the y-axis.

Click **OK**.

This completes the example of using polynomials.

Copying Properties

The conditions and values for a region, boundary, or interface can be copied from a region, boundary, or interface of the identical type using the clipboard or by using drag-and-drop.

The copy action causes all the conditions and values in the destination region, boundary, or interface to be identical to the source.

This copy feature can be used between regions, boundaries, and interfaces of the same simulation or between two different simulations, provided that both the simulations have been opened in the same STAR-CCM+ session.

Defining Dimensionality

For various objects in the simulation tree, clicking the property customizer ( ) of the *Dimensions* property opens a *Dimensions* dialog.

It allows you to set the dimensionality of a value or unit by setting the exponent of each of the base dimensions.



Setting Multiple Corresponding Conditions Simultaneously

Corresponding conditions for different regions, boundaries, and interfaces can be set at the same time.

To do this:

1. Open the nodes for the corresponding regions, boundaries, or interfaces.
2. Open the **Conditions** nodes.
3. Select the required conditions using the mouse to do multiple selection.
4. In the *Properties* window, select one of the available options from the drop-down list corresponding to the *Method* property.

   This option is applied to all the selected conditions.


intake
Simulation
Fluid_1
Boundaries
Free_Stream-4
Free_Stream-6
Physics Conditions
Flow Direction Specification
Free Stream Option
Turbulence Specification
Physics Values
Symmetry-7
Symmetry-8
Symmetry-9
Wall-1
Wall-2
Wall-3
Wall-5
Physics Conditions
Physics Values
Fluid_1 2
Boundaries
Free_Stream-31
Free_Stream-32
Physics Conditions
Flow Direction Specification
Free Stream Option
Turbulence Specification
Physics Values
Multiple Objects - Properties
Properties
Method
Intensity + Length Scale

Setting Multiple Corresponding Values Simultaneously

Corresponding conditions for different regions, boundaries, and interfaces can be set at the same time.

This is illustrated by setting a **Constant**, but the principle would hold true for field functions or tables.

1. Open the nodes for the corresponding regions, boundaries, or interfaces.
2. Open the **Values** nodes in each of the corresponding regions, boundaries, or interfaces.
3. Select the required values using the mouse to do multiple selection.
4. In the *Properties* window, select the **Constant** node and enter the appropriate values.

   This method is applied to all the selected values.


intake
Simulation
Free_Stream-6
Physics Conditions
Physics Values
Flow Direction
Mach Number
Constant
Pressure
Static Temperature
Turbulence Intensity
Turbulent Length Scale
Symmetry-7
Symmetry-8
Symmetry-9
Wall-1
Wall-2
Wall-3
Wall-5
Physics Conditions
Physics Values
Fluid_1 2
Boundaries
Free_Stream-31
Free_Stream-32
Physics Conditions
Physics Values
Flow Direction
Mach Number
Constant
Multiple Objects - Properties
Properties
Value
0.21

Setting Reference Values

Reference values comprise numerical values that are unique to the continuum and may be used by one or more models within the continuum.

The **Reference Values** manager node, which is subject to common menu items, contains the reference value nodes for the physics or mesh continuum.



Each model activated within the continuum can place one or more reference values in the continuum. In most situations, the reference values are critical to the simulation and must be appropriately set.

To view the reference values that are required for the continuum, open the **Reference Values** node. Selecting specific reference values allows the properties to be changed.



The table below lists the available reference values and the models that placed them in the continuum.

Reference Values	Registering Models
Pressure	Flow Models
Temperature	Boussinesq Model
Gravity	Gravity Model
Density	Gravity Model
Altitude	Gravity Model

Setting Reference Values by a Field Function

If you wish to use the same reference value in more than one location in a simulation, you can set the value using a user-defined field function, and access this where required. This would include property specifiers, analysis objects, and other user-defined field functions. For example, to do this with reference pressure:

1. Create a user-defined field function called `Reference Pressure`, and assign it a single value in the *Definition* property.



2. Enter the expression `$ReferencePressure` in the property value of the **Reference Pressure** reference value node.



As a result, you have a single source for the continuum reference value and the scalar function.



STAR-CCM+ Licensing

STAR-CCM+ has multiple product configurations. Each product configuration may require a different type of license, depending on what add-on modules you want to use.

There are two methods that are used to license CD-adapco products:

- standard, yearly leases (floating licenses)
- cloud-based, pay-per-hour licensing (Power-On-Demand)

Note: You can configure STAR-CCM+ to release a license automatically when it detects no activity for a set amount of time.

The following STAR-CCM+ product configurations are available:

- STAR-CCM+
- STAR-CCM+ Lite
- STAR-CCM+ Power Session
- STAR-CCM+ Power-On-Demand
- STAR-CCM+ Add-on modules

STAR-CCM+ Product Configurations

Purchasing a standard license allows you to use CD-adapco's software for a year. Licenses can be node-locked or floating, stored locally or on a network server.

The following license features are used with the various STAR-CCM+ product configurations:

Product Configuration	License Feature	Description
STAR-CCM+	ccmpsuite	The STAR-CCM+ product license contains all the standard features in STAR-CCM+. The default `ccmpsuite` feature license is used for this configuration.
STAR-CCM+ Lite	ccmplite	STAR-CCM+ Lite is a version of STAR-CCM+ that provides a limited set of features. STAR-CCM+ Lite configuration uses the `ccmplite` feature license. Models that are not available in STAR-CCM+ Lite are grayed out in the selection dialog that is used to set up the physics for a case. You can use STAR-CCM+ Lite to post-process the results of a simulation that were solved using STAR-CCM+. It is also possible for STAR-CCM+ Lite to run the simulation for further iterations providing all models are available in

Product Configuration	License Feature	Description
		STAR-CCM+ Lite. STAR-CCM+ Lite can run Java macros that were generated by STAR-CCM+, but if the macro activates any models that are not permitted under the STAR-CCM+ Lite license, the simulation fails to run.
STAR-CCM+ Power Session	ccmppower	STAR-CCM+ Power provides all the features of STAR-CCM+ with the added benefit of running one job on an unlimited number of cores. When this type of license is used, other licenses such as `ccmpsuite` remain inactive. If you start a job using the Power Session, only one copy of the `ccmppower` license is checked out.
STAR-CCM+ Power-On-Demand	ccmppower	The Power-On-Demand (PoD) session is a licensing model that enables you to use a STAR-CCM+ Power Session from a CD-adapco public license server rather than a license server on your own network. For more information, refer to the Power-On-Demand section.

A comparison table listing the STAR-CCM+ add-ons is provided below. A short description of what the configuration includes, along with any pre-requisites is also provided for each configuration.

STAR-CCM+ Add-on	Description	Pre-requisites	Feature LicenseName
STAR-CAT5	Enables the STAR-CAT5 client	STAR-CCM+ or STAR-CCM+ Lite and CATIA V5 with ME2 license	star-cat5
STAR-Works	Enables the STAR-Works client	STAR-CCM+ or STAR-CCM+ Lite and SolidWorks license	star-works
STAR-Pro/E	Enables the STAR-Pro/E client	STAR-CCM+ or STAR-CCM+ Lite and Pro/E license	star-proe
STAR-Creo	Enables the STAR-Creo client	STAR-CCM+ or STAR-CCM+ Lite and Creo or Pro/E license	star-proe
STAR-NX	Enables the STAR-NX client	STAR-CCM+ or STAR-CCM+ Lite and NX license	star-nx
STAR-Inventor	Enables the STAR-Inventor client	STAR-CCM+ or STAR-CCM+ Lite and Inventor license	star-inventor
STAR-SpaceClaim	Enables the STAR-SpaceClaim client	STAR-CCM+ or STAR-CCM+ Lite and SpaceClaim license	star-spaceclaim
CAD Exchange	Adds native CAD reading capability to STAR-CCM+	STAR-CCM+ or STAR-CCM+ Lite license	cadexchange

JTOpen Reader	Adds a JTOpen reader	STAR-CCM+ or STAR-CCM+ Lite license	jtopen
Battery Simulation Module	Enables Li-Ion Battery modeling in STAR-CCM+	STAR-CCM+ license	batterysim
Vibro-Acoustics Analysis	Enables wave number decomposition	STAR-CCM+ license	aerovibroacoustics
Optimate	Enables Optimate wizard and HEEDS post	STAR-CCM+, STAR-CCM+ Lite, or STAR-CCM+ Power license	doesuite
Optimate+	Enables Optimate wizard, HEEDS post, and SHERPA	STAR-CCM+, STAR-CCM+ Lite, or STAR-CCM+ Power license	doesuiteplus
Power Tokens	Run time tokens for Optimate and Optimate+	Optimate or Optimate+ license	doetokens
STAR-Cast	Enables casting modeling	STAR-CCM+ or STAR-CCM+ Power license	star-cast

Capabilities of STAR-CCM+ Lite

The relative capabilities of STAR-CCM+ and STAR-CCM+ Lite are given in the table below.

This list is not exhaustive and is intended to give a broad summary of the features.

Modeling Capabilities		**STAR-CCM+ Lite**	**STAR-CCM+**
3D-CAD	Parametric geometry modeling	✓	✓
	Native CAD import (requires	✓	✓

Modeling Capabilities		**STAR-CCM+ Lite**	**STAR-CCM+**
	an appropriate add-on license)		
Space	2D/ Axisymmetric	✓	✓
	Cyclic/periodic symmetry	✓	✓
	3D	✓	✓
Time	Steady state	✓	✓
	Transient	✓	✓
Motion	Stationary	✓	✓
	Moving Reference Frame	✓	✓
	Mixing Planes	-	✓
	Rigid Body Motion	-	✓
	Morphing	-	✓
	Superimposed Motions	-	✓
Overset Meshing			✓

Modeling Capabilities		**STAR-CCM+ Lite**	**STAR-CCM+**
Frequency Domain	Multi-Row Harmonic Balance	-	✓
Material	Single component	✓	✓
	Multi-component gas or liquid mixture	-	✓
	Multiphase mixture	-	✓
	Lagrangian multiphase	-	✓
	Thin film	-	✓
Solution Algorithm, Differencing Schemes and Solver Options	Segregated flow and fluid energy	✓	✓
	Coupled flow and energy	-	✓
Flow Regimes	Subsonic/ Transonic/ Supersonic	✓	✓
	Hypersonic	-	✓
	Inviscid	✓	✓
	Laminar	✓	✓

Modeling Capabilities		**STAR-CCM+ Lite**	**STAR-CCM+**
	Turbulent (RANS)	✓ (K-Epsilon Models)	✓
	Turbulent Transition	-	✓
	LES/DES	-	✓
Heat Transfer and Conjugate Heat Transfer	Conduction	✓	✓
	Convection	✓	✓
	Radiation	-	✓
Porous Media		✓	✓
Combustion and Chemical Reactions	Solid (e.g. coal)	-	✓
	Liquid	-	✓
	Gaseous	-	✓
	NOx and Soot models	-	✓

Modeling Capabilities		**STAR-CCM+ Lite**	**STAR-CCM+**
	Ignition models (both SI and CI)	-	✓
Multiphysics	Aeroacoustics	-	✓
	Dynamic Fluid Body Interaction (6-DOF motion)	-	✓
	Solid Stress	✓	✓
	Fluid Structure Interaction	-	✓
	Melting and Solidification	-	✓
Boundary Conditions	Inlet	✓	✓
	Outlet	✓	✓
	Pressure	✓	✓
	Stagnation conditions	✓	✓
	Non-reflective pressure and stagnation conditions	-	✓
	Wall	✓	✓

Modeling Capabilities		**STAR-CCM+ Lite**	**STAR-CCM+**
	Symmetry	✓	✓
	Cyclic	✓	✓
	Periodic Cyclic	✓	✓
	Transmissive free stream	✓	✓
	Riemann	✓	✓
	User-defined options	✓	✓
General	Cell types	✓	✓
	Fluid injection	-	✓
	Material property database	✓	✓
	Mixture properties	-	✓
	Temperature-viscosity relationships	✓	✓
	Specific heat and enthalpy heat relations	✓	✓
	Passive scalars		

Modeling Capabilities		STAR-CCM+ Lite	STAR-CCM+
		✓	✓
	Multi-stream	✓	✓

Power-on-Demand Licensing

The Power-On-Demand (PoD) session is a licensing model that allows you to use a STAR-CCM+ Power Session from a CD-adapco public license server rather than a license server on your own network.

You can purchase Real Time credit to use with the CD-adapco public license server. When you purchase Real Time credit, this credit is added to your account where it is available for you to use at any time. Real Time credit has a 12-month validity period. A unique key is used for authenticating the session before it starts. STAR-CCM+ requires this key when you launch a new session with PoD licensing. You can create multiple keys to manage and track your time credit within your company.

STAR-CCM+ requires some ports to be open to facilitate communication between your machine and the PoD license server. Also define the `CDLMD_LICENSE_FILE` environment variable.

In the *Create a New Simulation* or *Load Simulation* dialog, the correct format for the license path (the **Server** text box) is `Port@Host` or `Port@Host1,Port@Host2,Port@Host3`. `Port` is an integer. `Host` is an IP address or hostname. The Power-On-Demand Key must be a Base64-encoded string.

Verify all your license information before entering it. If the **Power-On-Demand** checkbox is activated and if either of these text boxes is empty or incorrectly formatted, an error message appears. In addition to the error message, the invalid text box is highlighted with a red border.



If neither text box has a valid entry, error notifications appear for both.



Additionally, a list of potential error messages with a brief explanation of what they mean is provided in the section on Power-On-Demand server errors.

Purchasing Real Time Credit

Real Time is defined as the difference between the time at which a STAR-CCM+ Power Session task ends and the time when it starts. It accounts for the summation of I/O times, CPU times, multiple concurrent sessions, communication channel delays, and all other types of wait.

Real Time credit is purchased from your CD-adapco sales representative.

An example of how Real Time works is outlined below.

- You have 500 hours of time credit in your account. You could run:
  - 1 Power Session job for 500 hours
  - 5 Power Sessions jobs for 100 hours each

- Or 500 Power Sessions jobs for 1 hour each

Important: when you start a STAR-CCM+ server using PoD licensing, two things happen:

- STAR-CCM+ reserves a number of hours from the available credit. This action is known as a deposit.
- Real Time credit counting starts.

Viewing Your Power-On-Demand Account Online

The details of your PoD account can be checked at any time by logging in to your Steve Portal account.

The remaining amount of PoD credit is displayed on the left of the page.

To get a more detailed view of your PoD account, click the **Licensing** link at the top of the page. If you do not have a Steve Portal account, contact your local support office.



The *Account Activity* tab is selected by default. Here, you can check your Real Time statement.



The available balance is shown in green on the graph, the used time is shown in red, while the deposits appear in yellow.

In the *Account Usage* tab, you can create and view reports to monitor the Real Time license usage in your company.

Managing License Keys

You can manage existing keys or generate new ones for each STAR-CCM+ session you wish to run.

In the *Manage Keys* tab, you can manage existing keys or generate new ones for each STAR-CCM+ session you wish to run.



To generate more keys:

1. Click (**Create License Key**) at the top of the screen.

   The *Security Key Generation* dialog appears.



Here, you can select an alias for the key and the dates you want it to be valid for. IP address restrictions that you want to place on the key, and the number of hours you want to allocate to it, are also specified in this dialog.

2. Click **Add key** to create the key and add it the list.
3. To delete a key, click (**Delete the key on this row**) next to the key you want to delete.

Which Ports Are Used in a Power-On-Demand Session?

Ports 1999 and 2099 must be open to allow communication traffic between your machine and the PoD license server.

On Linux, you can use the Nmap utility to verify that these ports are open. On Windows, you can use the Port Query tool that is provided in the FLEXlm directory of your STAR-CCM+ installation directory.

Verifying Ports on Linux

On Linux, you can use the Nmap utility to check that ports `1999` and `2099` are open.

To check your ports:

1. Enter the following in the command prompt:

```
% nmap flex.cd-adapco.com -PN -p 1999,2099
```

   where:

   - `[PATH]` is the full path to the nmap.exe executable file.
   - `flex.cd-adapco` is the host that you want to check.
   - `-PN` is the command for Nmap to ping the host.
   - `-p 1999,2099` are the ports that you wish to check.

   A report showing the results is displayed.

```
puriri3/jcarline% nmap flex.cd-adapco.com -PN -p 1999,2099

Starting Nmap 5.21 ( http://nmap.org ) at 2012-04-04 10:37 BST
Nmap scan report for flex.cd-adapco.com (12.39.225.146)
Host is up (0.21s latency).
PORT     STATE SERVICE
1999/tcp open  tcp-id-port
2099/tcp open  unknown

Nmap done: 1 IP address (1 host up) scanned in 0.47 seconds
puriri3/jcarline%
```

   In this example, both ports are `open`. If your report shows that the ports are anything but `open`, contact your system administrator and request that they open these ports for your machine. For more information on what the various *STATE* results mean, visit the Nmap website.

Verifying Ports on Windows

On Windows, you can use Port Query or Nmap to check that ports `1999` and `2099` are open.

Using Port Query

Port Query is included with your STAR-CCM+ installation. The default location is:

```
C:\Program Files\CD-adapco\FLEXlm\11_9_1_0\bin
\portqueryui.exe
```

To check your ports:

2. Launch `portqueryui.exe`.

   The following message may appear, warning you that no configuration file could be found.



3. Click **OK** to ignore this message.

   The *Port Query* dialog appears.
4. In the *Enter destination IP or FQDN to query:* entry box, enter `flex.cd-adapco.com`
5. Select the **Manually input query ports:** radio button.
6. In the *Ports to query:* entry box, enter `1999,2099`
7. Ensure *Protocol:* is set to **TCP**.



8. Click **Query**.

   The output displays the query results.



In this example, port 1999 and 2099 are open. If your ports are reported as anything but `LISTENING`, contact your system administrator and request that they open these ports for your machine.

Using Nmap on Windows

Nmap is a free utility that you can download from the Nmap download page. To check that the ports are open:

9. Download the stable command-line zip file.
10. Extract the contents of the file to a location of your choice.
11. Open a command prompt window.
12. Enter the following command:

```
[PATH]\nmap.exe flex.cd-adapco.com -PN -p 1999,2099
```

where:

- `[PATH]` is the full path to the `nmap.exe` executable file.
- `flex.cd-adapco` is the host that you want to check.
- `-PN` is the command for Nmap to ping the host.
- `-p 1999,2099` are the ports that you wish to check.

A report showing the results is displayed.

```
C:\Users\jcarline>C:\nmap-5.51\nmap.exe flex.cd-adapco.com -PN -p 1999,2099

Starting Nmap 5.51 ( http://nmap.org ) at 2012-04-04 13:29 GMT Daylight Time
Nmap scan report for flex.cd-adapco.com (12.39.225.146)
Host is up (0.085s latency).
PORT     STATE SERVICE
1999/tcp open  tcp-id-port
2099/tcp open  h2250-annex-g

Nmap done: 1 IP address (1 host up) scanned in 0.32 seconds

C:\Users\jcarline>_
```

In this example, both ports are `open`. If your report shows that the ports are anything but `open`, contact your system administrator and request that they open these ports for your machine. For more information on what the various *STATE* results mean, visit the Nmap website.

Creating the CDLMD_LICENSE_FILE Environment Variable

The `CDLMD_LICENSE_FILE` environment variable must be present on your machine before you can use PoD licensing.

On Linux, use the `setenv` command to define this variable. On Windows, you can manually add this variable to the machine.

Setting the `CDLMD_LICENSE_FILE` Variable on Linux

Use the `setenv` command to define the `CDLMD_LICENSE_FILE` variable.

1. Enter the following text in the command line:

```
% setenv CDLMD_LICENSE_FILE=1999@flex.cd-adapco.com
```

Setting the `CDLMD_LICENSE_FILE` Variable on Windows

Add the `CDLMD_LICENSE_FILE` variable using the *Environment Variables* dialog:

2. For Windows XP:
   a) Click **Start**, locate, and right-click on **My Computer** in the **Start** menu and select **Properties**.

      The *System Properties* dialog appears.

   b) Click the *Advanced* tab.

   c) Click the **Environment Variables** button.
3. For Windows Vista and Windows 7:
   a) Click **Start** and type `env` into the search field near the bottom of the **Start** menu.

b) Select **Edit environment variables for your account** from the list of results.

The *Environment Variables* dialog appears.

4. In the *System variables* group box, click **New...**.

   The *New System Variable* dialog appears.

5. In the *Variable name* entry box, enter `CDLMD_LICENSE_FILE`.
6. In the *Variable value* entry box, enter `1999@flex.cd-adapco.com`



7. Click **OK** to close the *New System Variable* dialog.
8. Click **OK** to close the *Environment Variables* dialog.

Launching STAR-CCM+ with the Power-On-Demand Option

Launch STAR-CCM+ with the Power-On-Demand option.

To enable PoD licensing:

1. Tick the PoD checkbox in the *Create a New Simulation* dialog.

   STAR-CCM+ automatically enables the **Power Session** option, and presents entry boxes for the Server and PoD Key. By default, the STAR-CCM+ public license server is entered in the **Server** entry box: `1999@flex.cd-adapco.com`.

2. Enter your PoD key in the **PoD Key** entry box.

   You can copy and paste this line from the *License Key* column of the PoD license management webpage.



Power-On-Demand Deposits

When you start a PoD session, a deposit is automatically taken out of your account.

The deposit is calculated as the number of hours from when your job began until the next midnight in Coordinated Universal Time (UTC). If your job runs beyond the next midnight, a further deposit of 24 hours is taken. This procedure is repeated until the job completes. When you end the simulation, any hours unused in the final deposit are returned to your overall credit account.

Note: If you run out of credits while your job is running, this job continues to run until completion. Purchase more credit to cover the hours used.

Power-on-Demand Server Errors Reference Page

A list of PoD server error messages—with an explanation of what they mean—is provided below. These error messages are displayed in the *Output* window of STAR-CCM+ when the error occurs.

`FAILED_TOKEN_NOT_RECOGNIZED -20001`

This error shows that the key that is passed via `LM_PROJECT` or `-podkey` does not match any keys that are stored in the PoD database. Make sure that the key you entered is valid. A key is a binary string and is converted to text using base64 encoding, so it is typically 24 or 28 characters long. Valid characters are: any of the 52 upper and lower case letters, the numbers 0–9, and the `+` and `/` characters.

`FAILED_TOKEN_EXPIRED -20002`

This error means that the keys that you created in your Steve Portal account have expired. To resolve this issue, create new keys.

`FAILED_TOKEN_NO_CREDIT -20003`

In your Steve Portal account, you can define a cap when creating PoD keys. This error means that the capped usage of the key has been exceeded and the key is no longer valid. To resolve this issue, create a new key.

`FAILED_TOKEN_IP_ADDRESS_NOT_IN_RANGE -20004`

This error appears when your IP address does not match the allowed range that you set up in the keys that are stored in your Steve Portal account. Your organization probably uses NAT (network address translation), so a machine can have one IP address on internal networks but another one when it accesses external networks. As the PoD server is external, make sure that you use the right IP range if you choose this option when creating a key. Alternatively, create a new key and do not enable the IP range restriction option.

`FAILED_NO_CREDIT_AT_ALL -20005`

This error means that there is no credit left in your PoD account. Credit is valid for a 12 month period, so this error can mean that your unused credit has expired. To resolve this issue, purchase more credit.

`FAILED_NO_PROJECT -30001`

This error means that you have not set `LM_PROJECT` either via the environment variable or via the `-podkey` command. To resolve this issue, create an environment variable named `LM_PROJECT` and enter your PoD key as the variable value. Alternatively, use the `-podkey <yourkey>` command when launching STAR-CCM+.

`FAILED_NO_CONNECTION -30002`

This error means that the PoD server cannot connect to the CD-adapco database. This error indicates a problem with the PoD server.

`FAILED_REFUSED_CONNECTION -30003`

This error means that the connection from the PoD server to the CD-adapco database has been refused. This error indicates a problem with the PoD server.

`FAILED_EXPECTED_FAILURE -40001`

This error only appears when the CD-adapco server undergoes routine maintenance.

`FAILED_UNCHARGED_CHECKOUT -50001`

This error appears for STAR-CD PoD and occurs when the `hpcdomains` license checks out before the `starsuite` license. This error can mean that you are checking out `starsuite` from a local server and the `hpcdomains` from the PoD server. To resolve this issue, look at the license output and check the order of the checkouts; the `starsuite` license must check out before the `hpcdomains` license.

STAR-CCM+ Add-on Modules

STAR-CCM+ automatically checks out any license that it needs for add-on modules.

A list of the add-on modules, along with the feature licenses that they use, is provided below:

- Battery Simulation Module (which uses the `batterysim` feature license)
- STAR-Cast (which uses the `star-cast` feature license)
- Aero-Vibro Acoustics Analysis (which uses the `aero-vibro` feature license). The Spatial Fourier Transform is licensed under this feature.
- JTOpen Reader (which uses the `JTOpen` feature license)
- CAD Exchange Reader (which uses the `cadexchange` feature license)

For example if you load a simulation file that contains items from the battery module, STAR-CCM+ checks out an available `batterysim` feature license.

Add-on licenses are checked out when:

- an add-on feature is used in a simulation
- the license is requested from the command line
- it is specified in the Add-on Licensing dialog
- it is reserved in the options dialog

Add-on licenses are checked in when:

- a simulation is closed
- a requiring feature is removed
- it is deactivated from the *Add-on Licensing* dialog

Note: By default, the JTOpen Reader and CAD Exchange Reader licenses are automatically released after importing a CAD geometry into STAR-CCM+. The licensing options for the STAR-CCM+ add-on modules can be manually controlled through the add-on licensing manager.

You can set up your simulation in your native CAD package environment using one of the STAR-CCM+ CAD client add-ons.

Working with the Add-on Licensing Manager

Licenses for the STAR-CCM+ add-on modules can be managed through the *Add-on Licensing* dialog.

To access this dialog:

1. Select **Tools** > **Add-ons** from the menu bar.



2. Alternatively, right-click on the root node in the STAR-CCM+ object tree and select **Add-ons**.



The *Add-on Licensing* dialog appears.



3. You can manually check out licenses by activating the checkbox next to the respective add-on, and clicking **OK**.

   A message then appears in the *Output* window informing you that the license was checked out successfully. Checked out licenses can be checked back in by deactivating the checkbox next to the respective add-on, and clicking **OK**. A message appears in the *Output* window informing you of this action.

The *Retry licenses on requiring operation* option controls whether STAR-CCM+ tries to check out the required license (if it is not already checked out). If this option is cleared, STAR-CCM+ does not try and check out any of the add-on licenses.

The *Release licenses after geometry import* option controls whether the JTOpen Reader and CAD Exchange licenses are released after importing a CAD geometry into STAR-CCM+. If this option is cleared, these two licenses will remain checked out until you manually release them or the simulation file is closed.

Reserving a License

Add-on licenses can be configured so that they are automatically checked out when a STAR-CCM+ session is started.

This option is useful when a certain add-on module is regularly used and can ensure that a license is always available when you need it.

Note: A reserved license is released when the simulation file is closed.

To reserve an add-on license:

1. Select **Tools** > **Options** from the menu bar.



A dialog displaying the various options for STAR-CCM+ appears. This dialog can also be accessed when no STAR-CCM+ session is active.

2. Expand the STAR-CCM+ node and select **Licensing**.



The pane in the right-hand side of the dialog displays the reserve options for each add-on.

3. Ticking these checkboxes and clicking **Close** will place a reserve on the respective add-on modules, so that the corresponding licenses will be automatically checked out next time a STAR-CCM+ session is started.
4. To release a license, manually check it in by clearing the respective checkbox in the Add-on Licensing dialog.

CAD Client Add-ons

The CAD clients plug into your native CAD work package and allow you to set up a simulation within its workspace.

You can also export your CAD geometries from your CAD package directly into STAR-CCM+ by clicking the **Transfer Geometry** button.

The following CAD client add-ons are available.

- STAR-Works
- STAR-CAT5
- STAR-Inventor
- STAR-Pro/E
- STAR-Creo
- STAR-NX
- STAR-SpaceClaim

Power Tokens

If you have an Optimate or Optimate+ license, you can use power tokens (`DOEtoken` feature licenses) with certain features of STAR-CCM+.

This section contains the following topics:

- Using Power Tokens for Pre- and Post-Processing Sessions
- Using Power Tokens as a Parallel License Feature for HPC Runs
- Controlling the Check Out Order of DOE Parallel Licensing

Using Power Tokens for Pre- and Post-Processing Sessions

If you have an Optimate or Optimate+ license, with power tokens, you can use 10 power tokens to perform pre- and post-processing in STAR-CCM+.

You use power tokens instead of a STAR-CCM+ license. You can run the pre- and post-processing session interactively or in batch mode. During the pre- and post-processing session, the STAR-CCM+ solver is deactivated.

To use power tokens with STAR-CCM+:

1. Launch STAR-CCM+ and load or create a simulation with the `-powerpre` argument.

STAR-CCM+ checks out 10 copies of DOEtoken.

Using Power Tokens as a Parallel License Feature for HPC Runs

If you have an Optimate or Optimate+ license, with power tokens, you can use power tokens as parallel license features.

Prior to STAR-CCM+ v9.02, the license check-out procedure for a HPC run was as follows:

1. Try to satisfy the parallel requirement with `hpcdomains`.
2. If insufficient licenses, try to satisfy the remaining requirement with `ccmplite` (if using STAR-CCM+ lite).
3. If insufficient licenses, and the runtime option `-nosuite` is not specified, try to satisfy the remaining requirement with `ccmpsuite`.

From STAR-CCM+ v9.02 onwards, you can use power tokens to fulfill requirements for a parallel license.

The new procedure is as follows:

4. Try to satisfy the parallel requirement with `hpcdomains`.
5. If insufficient licenses, try to satisfy the remaining requirement with `DOEtoken`.
6. If insufficient licenses, try to satisfy the remaining requirement with `ccmplite` (if using STAR-CCM+ lite).
7. If insufficient licenses, and the runtime option `-nosuite` is not specified, try to satisfy the remaining requirement with `ccmpsuite`.

Controlling the Check Out Order of DOE Parallel Licensing

When running Optimate cases, you can control the check out order of DOE parallel licenses.

The standard process for checking out DOE parallel licenses is as follows:

1. Try to satisfy the parallel requirement with `DOEtoken`.
2. If insufficient licenses, try to satisfy the remaining requirement with `hpcdomains`.
3. If insufficient licenses, and the runtime option `-nosuite` is not specified, try to satisfy the remaining requirement with `ccmpsuite`.

You can switch the order that the `DOEtoken` and `hpcdomains` license features are checked out using the `-doe-prefer-hpcdomains` runtime option. The license features are then checked out in the following order:

4. Try to satisfy the parallel requirement with `hpcdomains`.
5. If insufficient licenses, try to satisfy the remaining requirement with `DOEtoken`.
6. If insufficient licenses, and the runtime option `-nosuite` is not specified, try to satisfy the remaining requirement with `ccmpsuite`.

Example Scenario

You have the following feature licenses available:

- 20 `ccmpsuite`
- 20 `DOEtoken`
- 20 `hpcdomains`
- 1 `DOEsuite`

You want to run 3 jobs simultaneously, with 10 processors for each job. Job 1 is run first, then Job 2, followed by Job 3. Each job requires the following licenses:

- 1 suite license
- 1 core license
- 9 parallel licenses

For each DOE study, a single `DOEsuite` feature license is checked out. All STAR-CCM+ jobs that are spawned by a DOE study will share the `DOEsuite` license. Each job also requires a license for the core process. The default behavior is to use `DOEtoken` licenses to fulfill this requirement. If there are insufficient `DOEtoken` licenses for the core process, a `ccmpsuite` license is checked out instead. Similarly, `DOEtoken` licenses are used for each parallel process. If there are no `DOEtoken` licenses available for the parallel processes, `hpcdomains` are used instead.



Using the `-doe-prefer-hpcdomains` runtime option makes `hpcdomains` licenses the preferred license for fulfilling the parallel license requirement. If there are insufficient `hpcdomains` licenses available for the parallel processes, `DOEtoken` licenses are used to fulfill

the remaining requirement. `DOEtoken` is still used to satisfy the license requirements for the core process.



Getting Support

When you need assistance with using STAR-CCM+, you can first check the CD-adapco support site, The Steve Portal.

The Steve Portal lets you search information and open a support case. Access The Steve Portal at:

http://steve.cd-adapco.com

Alternatively, select **Help** > **Steve Portal**.



Choose from among the following features on this site:

- A Knowledge Base of Frequently Asked Questions, Videos, Best Practices, and other types of articles that cover a wide range of topics. You can search thoroughly for materials that relate to your problem. Perhaps your question has already been answered.
- A searchable library of documentation and training materials.
- The Macro Hut forum for sharing macros with the community.

- IdeaStorm to suggest new features, and vote on others that are submitted by the community, for STAR-CCM+.
- Access the status of your Power-on-Demand license and the next training sessions, as well as other technical resources.
- Tools for opening and tracking support cases.

  The Steve Portal can guide you through the process of opening a support case. In addition, it can suggest possible solutions that are based on the details of your case, and help you to track the progress of the case through to resolution.

The screenshot below shows the Home page of The Steve Portal. To switch between different areas of the site, use the tabs on the top-level navigation menu.



If you cannot find the answer to your question by searching The Steve Portal resources, you can open a support case for your support representative by clicking **OPEN SUPPORT CASE**. Alternatively, if Chat is available, select **Chat with Support**.

To help your CD-adapco representative, please prepare appropriately for reporting the issue. Further instructions on this subject are:

Gathering Information on the Issue

Accurate information is vital for dealing with your issue effectively. Your affected simulation file, if you can provide it, would be a key source of this information. See Preparing a Case for Support. Discuss the details of this file transfer with your CD-adapco support representative.

In addition, there are two ways to obtain information, depending on the issue:

- Copying data from the workspace
- Getting an external data file
- Retrieving graphics driver information

Copying Data from the Workspace

Error messages in STAR-CCM+ commonly appear in an exception dialog, which contains extensive information.

To access this information:

1. Click the **Show Details** button.



   The details appear in the dialog.
2. Simply highlight the information with your cursor and, when you have reached the end of the text, copy it to the clipboard and paste it into an e-mail message.



Click the **Hide Details** button to restore the brief dialog if you prefer.

3. In some cases, the error message also appears in the *Output* window, where you can also highlight and copy information.



4. Alternatively, you can use the pop-up menu of the *Output* window to save the full content as a text file.



If the resulting text file contains substantial information that does not pertain to the error message, you may want to edit it.

Getting an External Data File

Various errors cause STAR-CCM+ to generate separate reports. You should send these reports to CD-adapco support along with your query. The details they contain can help resolve your issue much sooner.

The most common external source of information on STAR-CCM+ errors is the `messages.log` file.



This file can be found in the following locations:

- On Linux: `~/.star-X.XX.XXX/var/log/messages.log`, where ~ is your `home` directory and `X.XX.XXX` is the version of STAR-CCM+, for example 9.04.006
- The Windows path may look like the following:

- On Vista: `C:\Users\username\AppData\Local\CD-adapco\STAR-CCM+ X.XX.XXX\var\log\`
- On XP: `C:\Documents and Settings\username\Local Settings\Application Data\CD-adapco\STAR-CCM+ X.XX.XXX\var\log\`

The file also contains prior error messages, so you may want to delete the earlier material for brevity.

Some errors are known as Java HotSpot errors because they are related to the Java functionality. As a result, a special log file may be generated called `hs_err_pidXXXX.log`, where `XXXX` is the process ID.



5. Look for the report in your main directory:
   - Your `home` directory on Linux
   - `C:\Users\user\Documents\` on Windows

Unlike `messages.log`, this type of report is generated one time due to a single error. Therefore all of its information pertains to that event.

Retrieving Graphics Driver Information

Some issues may originate with your graphics driver. You may want to get detailed information about that driver and send it to CD-adapco support. At a minimum you should provide the driver version, card model and operating system.

The technique for obtaining this information depends on your operating system:

6. On Linux, use the line command:

```
% glxinfo
```

This results in detailed output, which you can either paste in the body of your e-mail message or save as a separate text file.



7. On Windows, some drivers might have a special all-in-one interface. The example shown in the following screenshot is accessible on Windows Vista by right-clicking the desktop and selecting **NVIDIA Control Panel**.



8. Alternatively on Windows Vista, right-click the **Desktop** icon in the object tree of Windows Explorer and select **Properties**.
9. In the separate dialog that appears, click **Display Settings** in the list of options under the *Personalize appearance and sounds* heading.
10. In the separate *Display Settings* dialog, click the **Advanced Settings...** button.

    This activates another dialog with tabs that provide detailed information on the graphics hardware and drivers.



11. On Windows XP, right-click the desktop and select **Properties**.
12. In the separate *Display Properties* dialog that appears, select the *Settings* tab.
13. Click the **Advanced** button.

    This activates another dialog with tabs that provide detailed information on the graphics hardware and drivers.


(Multiple Monitors) and NVIDIA Quadro FX Go700 Pr...
Color Management
Quadro FX Go700
General
Adapter
Monitor
Troubleshoot
Adapter Type
NVIDIA Quadro FX Go700
Properties
Adapter Information
Chip Type: Quadro FX Go700
DAC Type: Integrated RAMDAC
Memory Size: 128 MB
Adapter String: Quadro FX Go700
Bios Information: Version 4.31.20.52.C2
List All Modes...
OK
Cancel
Apply

Preparing a Case for Support

This section explains how to get your case information to CD-adapco support.

To create a case that is routed directly to your support representative, when in The Steve Portal, click **OPEN SUPPORT CASE**, and fill the details as shown in the following image.



Specify the product and the exact version, as well as the computer type and operating system if they are relevant. The subject can include the error message that you obtained. Write the description thoroughly.

If you have any files to attach, such as screenshots of scenes, or error message text files, you can do so once the case has been created.


Add Attachment(s)
BROWSE
NOTICE:
This file upload service uses a third party PUBLIC cloud system. Data submitted will be stored and transmitted on this cloud. Export controlled data may not be uploaded using this site. By checking the box below, you certify that uploading data to this cloud system is not in violation of any applicable export control regulations including those of the US, EU, and/or any other jurisdiction,as applicable, and that the data is not found on the US Munitions or Commerce Control Lists, the UK or other EU member state Military Lists, the EU Dual Use List or any other list of export controlled items, as applicable. If you need to transfer export controlled data, please contact your DSE for alternative file transfer options and provide the export control classification information.
I confirm
SAVE
CANCEL

Suggesting Improvements of STAR-CCM+

If you want to suggest an enhancement or join the discussion about improving STAR-CCM+, CD-adapco provides a special forum, STAR-CCM+ IdeaStorm.

As a STAR-CCM+ user, you can participate in this forum in the following ways:

- Adding ideas
- Searching for ideas
- Commenting and voting on ideas
- Tracking ideas by subscribing to them

Access STAR-CCM+ IdeaStorm at:

http://cd-adapco.brightidea.com

Alternatively, select **Help** > **STAR-CCM+ IdeaStorm**.



STAR-CCM+ IdeaStorm is not intended for the following input:

- Discussions on support issues or best practices -- visit The Steve Portal.
- Reporting possible bugs -- contact your Dedicated Support Engineer (DSE).

The screenshot below shows the Home page of STAR-CCM+ IdeaStorm.

Detailed information on participating in STAR-CCM+ IdeaStorm is available in the IdeaStorm User Guide.

To access it:

1. Enter STAR-CCM+ IdeaStorm.
2. Click **FAQ** on the top navigation menu.
3. Under the *Extra Links* heading, click **IdeaStorm User Guide**.

Troubleshooting Graphics

This section provides techniques for dealing with the following visualization problems:

- Resolving hardware issues
- Working without hardware acceleration
- A scene does not open
- Concave faces display badly
- Z-fighting appears

Resolving Hardware Issues

If you experience problems in displaying graphics within the STAR-CCM+ *Graphics* window, try the following solutions:

- Install the latest driver for your graphics hardware. The process for obtaining and installing the latest driver depends on your graphics hardware. Refer to the manufacturers instructions for details.
- On Windows, decrease the hardware acceleration slider in the *Troubleshooting* tab of the advanced display properties dialog. If the graphics display becomes slow as a result, contact your CD-adapco representative.
- If problems persist after installing the latest drivers, or after changing hardware acceleration on Windows, then you can bypass the graphics hardware altogether by using the `-mesa` command-line option when starting STAR-CCM+. Report the problems to your CD-adapco representative.

Some specific issues reported to CD-adapco support in the past are:

- With some older NVIDIA drivers, memory leaks can occur when using OpenGL display lists to render scene graphics. OpenGL display lists are used when the Immediate Mode option is turned off in a scene displayer. To fix the memory leaks, try upgrading to the latest driver, or alternatively turn on *Immediate Mode* for scene displayers. (See section, Scalar Displayer Properties).
- On Linux systems, the 169.XX series of NVIDIA drivers are unstable and are not supported. Upgrade to the latest series of drivers.
- On Windows laptops using certain NVIDIA graphics chips (such as the Quadro FX770M), STAR-CCM+ can crash, hang, or show unexpected artifacts in the graphics display. If you experience problems:
  - upgrade to the latest graphics driver for your graphics card
  - choose the **Dassault Systemes CATIA - compatible** option for the global 3D setting in the NVIDIA control panel

Working Without Graphics Hardware (Linux only)

When hardware-accelerated drivers are not available on your machine, or when you must run STAR-CCM+ remotely through an X-server (such as VNC), you can activate Mesa graphics libraries.

The Mesa graphics libraries provide software rendering without requiring hardware acceleration. Graphics performance is much slower when using Mesa than when using graphics hardware with up-to-date hardware-accelerated drivers. To explicitly force the use of software rendering with Mesa, launch STAR-CCM+ from a command prompt with the `-mesa` option appended.

To increase responsiveness:

1. Select **Tools > Options**. In the *Options* dialog that appears, open the **Options** and **STAR-CCM+** nodes, and select the **Visualization** node.



2. Increase the *Refresh Rate* property.

   This adjustment demotes surfaces to a simpler representation (and hence drawn faster) when interaction occurs in the scene.

Limitations of Software Rendering with Mesa

Software rendering is known to be much slower than hardware rendering. When rendering with advanced visualization options, such as Line Integral Convolution (LIC) or volume rendering, the time to

render a scene can be 100 times longer than with hardware rendering. Consequently, you are strongly advised to run STAR-CCM+ with hardware rendering whenever graphics hardware is present.

Note: You can enable hardware rendering in batch (assuming that you have the appropriate graphics hardware) with the `-hardwarebatch` option.

A Scene Will Not Open

If you are unable to open a scene, you can obtain information for troubleshooting.

1. Right-click on the Scenes node and select **Test Graphics**.



2. Otherwise, right-click on the node of the particular scene and select **Report**.



Either command produces a report in the *Output* window containing technical details on the graphics hardware.



If you wish, simply send the information to support by selecting the text in the *Output* window, copying it and pasting it into an e-mail message.

Improving the Rendering of Concave Faces

Some objects with highly non-planar polygons show artifacts. This problem is because STAR-CCM+ attempts to triangulate these kinds of surfaces--a process that requires extra memory.



To activate triangulation:

1. Select **Tools > Options**. In the *Options* dialog that appears, select the **Options > STAR-CCM+ > Visualization** node.
2. Activate the checkbox of the *Triangulate Part Surfaces* property.



Use this option sparingly, since increased memory usage slows graphics rendering.

The triangulated surfaces appear in the scene display.



Z-Fighting Appears

The model shows rear polygons “bleeding” from behind the front ones.



This problem is known as z-fighting. The “fighting” occurs in three-dimensional rendering when multiple polygons are in the same plane or nearly so, and their pixels displace one another as you move the image in the scene. The competing portions of these polygons are called artifacts.

In general the cause of z-fighting is insufficient precision in the z-buffer, the element of the computer hardware/software that is supposed to manage the depth of the image (in the z-direction). Yet there are multiple aspects to this problem that require different solutions:

- Hardware/software requirements

  If your depth buffer size is less than 32 bits, then an upgrade of your graphics hardware is necessary to manage the depth imaging properly.

- Multiple surfaces in the same plane (creating a shimmering effect)

  Review the distances among the various parts of your model. If they are close enough, they could be treated as being in the same plane. Also, make sure that if a part is in multiple displayers in the same scene, then only one of those displayers (at most) has surface mode activated.

- Excessive viewing distance

In rare cases the problem can be due to the model being shifted too far from the origin. Typically distances of millions of units are required to cause this effect, although these distances can be shorter if you are using obsolete graphics hardware. To eliminate excessive distance, shift the model back to the origin.

Credits

STAR-CCM+ makes use of several third party software components to provide certain features within its code.

Some of these components are as follows:

- HOOPS 3D Exchange v7.2 from Tech Soft 3D.
- Parasolid Kernel Modeler v26.1.217 by Siemens PLM Software, Inc.
- DCM v63 by Siemens PLM Software, Inc.

IBM Platform MPI

STAR-CCM+ supports IBM Platform MPI from Platform Computing Inc. IBM Platform MPI is the default message passing interface for STAR-CCM+ and is included with the STAR-CCM+ installer.

lp_solve

STAR-CCM+ references the `lp_solve` library within its code. This library is governed by the GNU Lesser General Public Licence. Full citation data and disclaimer notice is given as follows:

```
lpsolve citation data
----------------------

Description : Open source (Mixed-Integer) Linear
Programming system
Language : Multi-platform, pure ANSI C / POSIX
source code, Lex/Yacc based parsing
Official name : lp_solve (alternatively lpsolve)
Release data : Version 5.1.0.0 dated 1 May 2004
Co-developers : Michel Berkelaar, Kjell Eikland,
Peter Notebaert
Licence terms : GNU LGPL (Lesser General Public
Licence)
Citation policy : General references as per LGPL.
Module specific references as specified therein
```

The `lp_solve` library is distributed "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

sundials

Copyright ® 2002, The Regents of the University of California.
Produced at the Lawrence Livermore National Laboratory.
All rights reserved.

This file is part of SUNDIALS.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the disclaimer below.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the disclaimer (as noted below) in the documentation and/or other materials provided with the distribution.
3. Neither the name of the UC/LLNL nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OF THE UNIVERSITY OF CALIFORNIA, THE U.S. DEPARTMENT OF ENERGY OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Additional BSD Notice

1. This notice is required to be provided under our contract with the U.S. Department of Energy (DOE). This work was produced at the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-ENG-48 with the DOE.

2. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately-owned rights.
3. Also, reference herein to any specific commercial products, process, or services by trade name, trademark, manufacturer or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

SuperLU

Copyright ® 2003, The Regents of the University of California.
Produced at the Lawrence Berkely National Laboratory (subject to receipt of any required approvals from U.S. Dept. of Energy).
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Lawrence Berkeley National Laboratory, U.S. Dept. of Energy nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

LAPACK

```
LAPACK++ 1.1a Linear Algebra Package 1.1a
University of Tennessee, Knoxvilee, TN.
Oak Ridge National Laboratory, Oak Ridge, TN.
Authors: J. J. Dongarra, E. Greaser, R. Pozo, D.
Walker

® 1992-1996 All Rights Reserved
```

NOTICE

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted provided that the above copyright notice appear in all copies and that both the copyright notice and this permission notice appear in supporting documentation.

Neither the Institutions (University of Tennessee, and Oak Ridge National Laboratory) nor the Authors make any representations about the suitability of this software for any purpose. This software is provided ``as is'' without express or implied warranty.

LAPACK++ was funded in part by the U.S. Department of Energy, the National Science Foundation and the State of Tennessee.

MUMPS

This version of MUMPS is provided to you free of charge. It is public domain, based on public domain software developed during the Esprit IV European project PARASOL (1996-1999). Since this first public domain version in 1999, research and developments have been supported by the following institutions: CERFACS, CNRS, ENS Lyon, INPT(ENSEEIHT)-IRIT, INRIA, and University of Bordeaux.

The MUMPS team at the moment of releasing this version includes Patrick Amestoy, Maurice Bremond, Alfredo Buttari, Abdou Guermouche, Guillaume Joslin, Jean-Yves L'Excellent, Francois-Henry Rouet, Bora Ucar and Clement Weisbecker.

We are also grateful to Emmanuel Agullo, Caroline Bousquet, Indranil Chowdhury, Philippe Combes, Christophe Daniel, Iain Duff, Vincent Espirat, Aurelia Fevre, Jacko Koster, Stephane Pralet, Chiara Puglisi, Gregoire Richard, Tzvetomila Slavova, Miroslav Tuma and Christophe Voemel who have been contributing to this project.

Up-to-date copies of the MUMPS package can be obtained from the Web pages: http://mumps.enseeiht.fr/ or http://graal.ens-lyon.fr/MUMPS

THIS MATERIAL IS PROVIDED AS IS, WITH ABSOLUTELY NO WARRANTY EXPRESSED OR IMPLIED. ANY USE IS AT YOUR OWN RISK.

User documentation of any code that uses this software can include this complete notice. You can acknowledge (using references [1] and [2]) the contribution of this package in any scientific publication dependent upon the use of the package. You shall use reasonable endeavours to notify the authors of the package of this publication.

[1] P. R. Amestoy, I. S. Duff, J. Koster and J.-Y. L'Excellent, A fully asynchronous multifrontal solver using distributed dynamic scheduling, SIAM Journal of Matrix Analysis and Applications, Vol 23, No 1, pp 15-41 (2001).

[2] P. R. Amestoy and A. Guermouche and J.-Y. L'Excellent and S. Pralet, Hybrid scheduling for the parallel solution of linear systems. Parallel Computing Vol 32 (2), pp 136-156 (2006).

fastprox


```
Copyright (C) 2011 Paul Mineiro
All rights reserved.
```

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Paul Mineiro nor the names of other contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Contact: Paul Mineiro <paul@mineiro.com>

JFreeChart

This software incorporates JFreeChart, (C)opyright 2000-2009 by Object Renery Limited and Contributors.

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the

```
Free Software Foundation, Inc.,
59 Temple Place,
Suite 330,
Boston, MA
02111-1307 USA
```

Modeling Geometry

This part of the User Guide describes how to use 3D-CAD, the parametric solid modeler available within STAR-CCM+. It also describes the features and tools available for further preparation of the geometry model and its surfaces.

For many simulations in STAR-CCM+, the first step involves preparing the geometry of the case under investigation. This geometry could come from a model that was prepared previously in a CAD package, or could be created directly using the 3D-CAD module within STAR-CCM+.

3D-CAD

This fully integrated CAD tool allows you to build geometry from scratch within the STAR-CCM+ environment, as well as import external CAD models and modify them before running a simulation. 3D-CAD is tailored to suit the needs of a CFD analysis and allows you to change the geometry quickly and easily, before rerunning the simulation.

This fully integrated CAD tool allows you to build geometry from scratch within the STAR-CCM+ environment, as well as import external CAD models and modify them before running a simulation. 3D-CAD is tailored to suit the needs of a CFD analysis and allows you to change the geometry quickly and easily, before rerunning the simulation.

The chapters in this section covering 3D-CAD are as follows:

- Creating geometry with 3D-CAD
- Importing geometry into 3D-CAD
- Exporting geometry from 3D-CAD
- Using 3D-CAD models in a simulation

Preparation of the Geometry Model

The concepts that are associated with manipulating and preparing geometry have been distinguished from the ones that relate to regions and boundaries through the introduction of geometry parts. Geometry parts are contained within the Geometry Parts node of the object tree. They are fully described in the section Using Geometry Parts.

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- How to generate or import meshes and work with them

- Modeling physics -- the concepts, theory, and implementation details of the models
- Using tools in STAR-CCM+ to facilitate your work
- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format
- A guide to using STAR-View+ , a viewer for three-dimensional STAR-CCM+ scenes

Using 3D-CAD

The first part of this chapter is an introduction to 3D-CAD modeling that covers the fundamental concepts behind 3D-CAD. If you are new to solid modeling or 3D-CAD, read this section first.



Subsequent chapters describe how to create geometry using the tools available in 3D-CAD. These tools are:

- Working with sketches
- Converting sketches to bodies

- Working with bodies
- Working with fillets and chamfers
- Working with design parameters
- Working with fluid volume extraction

To create geometry that is based on an imported CAD file, refer to the following section:

- Importing geometry into 3D-CAD

Introducing 3D-CAD Modeling

3D-CAD is a feature-based parametric modeler within STAR-CCM+ that allows you to build geometries from scratch. The geometries that you create with 3D-CAD are stored as 3D-CAD models, which you can then convert to geometry parts for integration with the meshing and simulation process.

The first step in building a geometry is to create a 3D-CAD model. This action activates 3D-CAD and you can start to create your geometry. Models are created by carrying out a series of operations such as sketching, extruding, and cutting. Each operation is a feature of the model, and the chronological overview of these features is displayed in the feature tree.

In general, the process of building a model in 3D-CAD begins with the creation of a sketch. You can convert this sketch to a body by extruding or revolving. You can add or remove material from this body by creating more sketches and performing extrusions, revolutions, sweeps, lofts, and cuts. Body operations allow you to modify the geometry further. In addition to solid modeling, 3D-CAD also supports surface modeling.

You can modify geometry that you created in 3D-CAD at any time by reactivating 3D-CAD and modifying the model. When the model is updated the changes are applied throughout the simulation, requiring you only to remesh and rerun the case.

A major feature of 3D-CAD is design parameters, which allow you to modify the 3D-CAD model from outside of 3D-CAD. These design parameters allow you to solve for a particular geometry, change the size of one or more components, and rerun the case quickly. If a solution has already been computed, it is mapped across to the modified geometry, decreasing the overall run time.

You can check the validity of geometries that you create as well as geometries that you import in 3D-CAD.

See Also

What Is a 3D-CAD Model
Creating Geometry Parts from 3D-CAD
Working with Features and the Feature Tree
Sketches
Surface Modeling in 3D-CAD
What Is a Design Parameter?

Understanding Geometry Parts

What Is a 3D-CAD Model

A 3D-CAD model is a solid model that is created in STAR-CCM+ using the feature-based parametric solid modeler, 3D-CAD.

You create models within the familiar environment of STAR-CCM+. These models are stored under the **Geometry** > **3D-CAD Models** node in the simulation tree.



Creating a 3D-CAD Model

To create a 3D-CAD model:

- Right-click on the **Geometry** > **3D-CAD Models** node and select **New**.



This action creates a 3D-CAD model node under the **Geometry > 3D-CAD Models** node and activates 3D-CAD. When 3D-CAD is open, the 3D-CAD model tree is shown in the explorer pane and a *3D-CAD View* scene is displayed in the *Graphics* window. You can begin adding features to the model by creating a sketch or sketch plane, or by importing geometry from an external CAD file. To finish working on the model, close 3D-CAD.

Exiting 3D-CAD

To exit 3D-CAD and return to the main simulation tree:

- Click the **Close 3D-CAD** button.



The *3D-CAD View* scene automatically closes, and view settings in the scene are discarded. You can save these settings before exiting 3D-CAD by storing the current view. You can then restore the view in another scene later on. To reactivate 3D-CAD, you can edit an existing 3D-CAD model or create a model.

Editing an Existing 3D-CAD Model

You can edit existing 3D-CAD models and resume from where you left off. To edit a 3D-CAD model:

- Right-clicking on the model node under the **Geometry > 3D-CAD Models** manager node and select **Edit...**



3D-CAD launches and opens the *3D-CAD View* scene. You can resume working on your 3D-CAD model. When you have finished editing the model, you can close 3D-CAD and return to the simulation tree. If you already used the model to create geometry parts, update these parts for the changes to take effect.

Note: Further steps are required if you used the geometry to generate a mesh or run a simulation. See Modifying a 3D-CAD Model.

See Also

Sketches
Sketch Planes
Creating Geometry Parts from 3D-CAD

Visualizing a 3D-CAD Model

3D-CAD models are visualized in the *3D-CAD View* scene. This scene opens automatically when you activate 3D-CAD by creating or editing a model.



When you close 3D-CAD, the *3D-CAD View* scene also closes. You can restore the scene by reactivating 3D-CAD as described above. Closing 3D-CAD also discards the view settings—when a new *3D-CAD View* scene is opened, the default view is shown. You can save the view settings before exiting 3D-CAD by storing the current view. You can restore these view settings in a scene later on.

To visualize a 3D-CAD model outside of 3D-CAD:

1. Convert the model to geometry parts. See Creating Geometry Parts from 3D-CAD.
2. Display the parts in a geometry scene. To create a geometry scene:
   - Click the (**Create/Open Scenes**) button in the **Vis** toolbar and select **Geometry**.
   - Alternatively, right-click on the **Scenes** node and select **New Scene** > **Geometry**.

See Also

What Is a 3D-CAD Model

Checking Geometry Validity

STAR-CCM+ can check the validity of geometries within the 3D-CAD module.

Checking for valid geometries is important before setting up a simulation—particularly when using imported geometries. Invalid geometries can cause problems during the meshing process. You can check the validity of an entire model, specified bodies, and specified faces in 3D-CAD. You can check the validity of sheet bodies as well as solid bodies.

Use one of the following actions to run a validity check:

- In the feature tree, right-click a body that you want to check and select **Check Validity**.
- In the feature tree, right-click the model node and select **Check Validity**.
- In the 3D-CAD View scene, right-click on a face and select **Check Validity** from the *Body* group box.
- In the 3D-CAD View scene, right-click on a face and select **Check Validity** from the *Face* group box.

  A message with the results of the validity check is printed in the *Output* window. For example, if a body is not valid, the following message is displayed:

  ```
  Body Body 2 is not valid.
  ```

Checking Validity Upon Import

You can automatically check the validity of a geometry when you import it into 3D-CAD. See Importing Geometry into 3D-CAD.

3D-CAD Feature Previews

For most 3D-CAD features, STAR-CCM+ displays a preview of the expected result in the *3D-CAD View* scene before you complete the action to create the feature. If you have a complex model, and find that the preview function takes too long, you can deactivate the function in the feature panel.



To deactivate the preview display:

1. Deactivate the *Show Preview* option in the relevant feature panel.



Working with the 3D-CAD View Scene

You can interact with the 3D-CAD model geometry in the *3D-CAD View* scene.

Right-clicking on a face or edge in the *3D-CAD View* scene brings up a menu with options separated into three categories: `Face`, `Body`, and `Feature`. This menu also informs you of the name of the selection, what body it belongs to, and what feature it belongs to. Options that are specific to either a face, feature, or body are listed in the relevant category.

In the following example, the right-click menu informs you that the highlighted selection is a **Face** that belongs to **Fillet 3** and is found on **Body 2**.



Right-clicking on features in the *3D-CAD View* scene is a quick and easy way of identifying them without needing to search the 3D-CAD feature tree. You can get more information about each option that is available in the right-click menu can be found by hovering the mouse pointer over the desired menu item and pressing the `<F1>` key.

You can select either a face or an edge in the *3D-CAD View* scene. To facilitate edge selection on a geometry:

- Set the **Selection Filter** in the Selection Toolbar to **Select Curves**. See Selecting Parts by Type in the Display.

This action makes right-clicking on a specific edge of a geometry model easier, especially on larger, and more complex models.

Working with Features and the Feature Tree

Creating geometry in 3D-CAD involves working with features and the feature tree. This section describes how to roll the feature tree back, forward, and to the end.

This section also explains how to suppress and unsuppress features. Manipulating the feature tree using these tools changes the state of the features and allows you to control where new features are inserted in the feature tree. Different feature state icons appear next to the feature nodes when their state changes.

Note: Right-clicking on a feature node presents a pop-up menu relevant to that particular feature. Some of the actions are common to most of the features, while others are more specific.

What are Features?

STAR-CCM+ uses a feature-based modeling approach in its tools for building 3D-CAD models. In this approach, you construct the 3D-CAD model using a series of operations, or features, each of which represents a particular action.

For example, a sketch feature is added to the model when you construct points, lines and curves on a sketch plane. An extrude feature is added when you select an existing sketch feature and extrude it in a particular direction. The entire set of features that is used to construct a model is stored in the feature tree in the order that you created them. For a full list of the features available in 3D-CAD, see Feature Tree Node Reference.

What is the Feature Tree?

The feature tree of a 3D-CAD model is an ordered sequence of all the actions (or features) that are required to build the final model. The **Features** manager node acts as a parent to all feature nodes within the tree. A full list of these nodes is given in the feature tree node reference section.

The feature tree is dynamic—that is, editing an existing feature in the tree impacts the outcome of subsequent features. The result from one feature acts as the input for the next feature, so that the resulting model

is the cumulative effect of applying all the features to the original canonical sketch planes.

The dynamic nature of the feature tree means that you manipulate it to achieve different forms of the final model. If you roll back the feature tree to a particular point, then only those features that appear before the roll point are executed in the CAD engine. When you add a feature to a rolled-back tree, it is inserted at the roll point, rather than at the end of the tree. Rolling the tree forward (either to a specific feature or to the end of the tree) results in all intermediate features being reactivated within the CAD engine. These items are discussed in more detail in the relevant sections.

You can also suppress features to deactivate them. Suppressed features are not removed from the tree, and can be unsuppressed to reinstate their contribution.

Initial State

In its initial state, the feature tree contains nodes for each of the three canonical sketch planes. Right-clicking on any of these three nodes provides actions to create a sketch on the sketch plane, or to create a further sketch plane that is based on the initial canonical sketch plane.



See Also

Rolling Back
Understanding Feature States
Understanding Where New Features Are Inserted
Suppressing Features

Common Feature Actions Reference

The following actions are common to nodes across 3D-CAD:

- Edit
- Delete
- Rename

Note: You cannot edit, delete, or rename the three default sketch planes.

Editing a Feature

By editing a feature, the underlying properties of the feature can be altered, such as angle of revolution, sketch dimension, or selected bodies. To edit a feature:

1. Right-click on the feature node in the feature tree and select **Edit**.



This action automatically rolls back to the selected feature and opens the corresponding panel.

2. Perform the changes that you want.
3. Click **OK** to confirm the changes.

Additional information is available for editing:

- Sketches
- Lofts
- Sweeps
- Extrusions
- Revolves

Deleting a Feature

You can permanently remove a feature from your 3D-CAD model by deleting it from the feature tree.

To delete a feature:

- Right-click on the feature node and select **Delete**.



The feature is removed from the feature tree and the 3D-CAD model. If the selected feature is a parent feature— that is, a feature upon which other features depend, it cannot be deleted until its child features have been deleted.

Caution: This action is irreversible.

Additional information is available for deleting design parameters. See Deleting a Design Parameter.

Renaming a Feature

You can rename most feature nodes in the 3D-CAD feature tree. This capability is useful for complex models containing many features.

To rename a feature:

- Use the inline editor



- Alternatively, use the right-click menu.

Similarly, you can rename body and design parameter nodes.

See Also

Renaming Objects
What Is a Design Parameter?
Overview of Bodies in 3D CAD

Rolling Back

When you roll back to a particular feature, you temporarily deactivate that feature and all subsequent features.

While in the rolled back state, you can insert new features in the feature tree. These features are placed directly above the rolled-back features, and will remain there after you roll forward. The 3D-CAD model is only constructed from active features. Rolled-back features are overlaid with the roll back symbol .

To roll back to a particular point:

- Right-click the feature node that you want to roll back to and select **Roll Back**.



The selected node, and those following to the end of the feature tree, are deactivated.



Rolling Forward

When the 3D-CAD feature tree contains rolled back features, you can roll forward to one of those features. Rolling forward to a particular feature reactivates that feature and all features that come before it. The CAD engine includes these reactivated features in the construction of the 3D-CAD model.

To roll forward to a feature:

- Right-click the feature node that you want to roll forward to and select **Roll Forward**.



This node, and all rolled-back nodes that come before it, are reactivated in the feature tree.

Rolling to the End

When the 3D-CAD feature tree contains rolled-back features, rolling to the end reactivates all features in the feature tree. The CAD engine includes all active features in the construction of the 3D-CAD model.

Note: Features that were suppressed before applying the roll back will remain suppressed after rolling to the end.

To roll to the end:

- Right-click a rolled-back node in the feature tree and choose **Roll To End**.



All rolled back features are reactivated.

Understanding Where New Features Are Inserted

The point at which new features are added depends on whether you have rolled back the feature tree. If you rolled back the feature tree, the new features are added after the last active feature. If you did not roll back the feature tree, new features are added at the end.

In the following example, the feature tree is rolled back to the **TransformSketchPlane** feature.



If you create a new feature at this point (for example, a sketch), it is inserted before the first rolled-back feature, **TransformSketchPlane**.



Suppressing Features

Suppressing a feature deactivates it and any features that depend on it.

Suppressed features are indicated by the symbol. You can suppress multiple features at once by selecting multiple feature nodes in the feature tree using the <`Ctrl`> and <`Shift`> keys.

To suppress a feature:

- Right-click the feature that you want to suppress and choose **Suppress**.



The selected feature is suppressed. The child features of the sketch plane (**Sketch 2** and **Extrude Cut 1**) are also suppressed.



Note: Suppressing a feature can invalidate subsequent features in the feature tree, and cause warning dialogs to appear if the CAD engine fails to reconstruct the model.

Unsuppressing Features

You can unsuppress a previously suppressed feature to reinstate its operation in the construction of the model.

To unsuppress a feature:

- Right-click the suppressed features in the feature tree and select **Unsuppress**.



The selected features are unsuppressed and the suppressed icon overlays are removed. If the suppressed features have dependents (which would also have been suppressed), it is possible to unsuppress the dependent features first. This action also unsuppresses the parent feature. The precise behavior depends on the nature of the feature tree. You must unsuppress child features if you first unsuppressed the parent feature. You can unsuppress multiple features at once by selecting them in the feature tree using the <`Ctrl`> and <`Shift`> keys.

Understanding Feature States

Each feature node has a unique icon to indicate what the feature does.

If the feature is in its normal state, and it is correctly constructed by the CAD engine, the icon appears as normal. If the feature is not in its normal state, an additional symbol appears on the feature icon to indicate the alternate state. The possible states are shown in the table below:

Status	Description
Rolled Back	The feature is further down the tree than the roll back point and so is not executed in the CAD engine.
Suppressed	The feature, or its parent feature, is suppressed and is not executed in the CAD engine.
Invalid	The CAD engine has failed to update or execute the feature during model construction. Editing an invalid feature can allow you to work out the cause of the problem, and correct it accordingly.
Out-of-Date	Indicates that a feature is out-of-date and requires updating.

See Also

Feature Tree Node Reference
Rolling Back
Suppressing Features

Creating Basic Geometric Shapes in 3D-CAD

In electronics cooling applications, most electronic components are represented by basic geometric shapes such as blocks and cylinders. 3D-CAD allows you to create these simple shapes in a single operation without having to create a sketch first.



Note: In the electronics cooling industry, the circuit board and its various components are defined in the XY plane. 3D-CAD uses this convention when you create simple shapes.

To create a basic shape:

1. Right-click the **Features** node and select:
   - **Solid Primitive > Block** to create a block.
   - **Solid Primitive > Cylinder** to create a cylinder.
2. In the subsequent dialog, set the shape parameters.
   - For blocks, see Create Block Panel.
   - For cylinders, see Create Cylinder Panel.
3. Expose any properties that you want as design parameters by clicking **(Expose Parameter)** next to the property.

   The *3D-CAD view* scene updates to show a preview of the body.



4. Click **OK**.

Create Block Panel

Use the following properties to define a simple block shape in 3D-CAD. The block extends symmetrically from the center point.

Length (X Distance)	Specifies the total size of the block in the X direction.	
Width (Y Distance)	Specifies the total size of the block in the Y direction.	
Height (Z Distance)	Specifies the total size of the block in the Z direction.	
Center	Specifies the center point of the base of the shape.	
	X	X coordinate.
	Y	Y coordinate.
	Z	Z coordinate.
Coordinate System	Specifies the coordinate system in which you define the center point.	

Create Cylinder Panel

Use the following properties to define a simple block shape in 3D-CAD.

Radius	Specifies the radius of the cylinder.	
Height	Specifies the amount that the cylinder extends outwards from the center point.	
Center	Specifies the center point of the base of the shape.	
	X	X coordinate.
	Y	Y coordinate.
	Z	Z coordinate.
Coordinate System	Specifies the coordinate system in which you define the center point.	

See Also

What Is a Design Parameter?

Feature Tree Node Reference

The following features can be used to create a 3D-CAD model:

	XY	The XY canonical sketch plane. It cannot be edited or deleted.
	YZ	The YZ canonical sketch plane. It cannot be edited or deleted.
	ZX	The ZX canonical sketch plane. It cannot be edited or deleted.
	Chamfer	Added when a chamfer is applied to edges of a body.
	Circular Pattern	Added when a circular pattern is created from a body.
	Circular Pattern Cut	Added when a circular pattern cut is created from a body.

	Create Sheet Bodies	Added when a sheet body is created from a solid body, or another sheet body.
	Delete Body	Added when a body is deleted.
	Delete Faces	Added when a face is deleted.
	Duplicate Body	Added when a body is duplicated.
	Extract External	Added when the external volume of space around a model is extracted and used to create a new solid body.
	Extract Internal	Added when the internal volume of space within a model is extracted and used to create a new solid body.
	Extrude	Added when a sketch is extruded.
	Extrude Cut	Added when an extrude cut feature is used to remove material from the model by sweeping a sketch in a specific direction.
	Face Sketch	Added when a sketch is created by extracting the edges of a face.
	Filled Surface	Added when a filled surface is created.
	Fillet	Added when a fillet is applied to the edges of a body.
	Import Cad	Added when geometry is imported into 3D-CAD from an external CAD file.
	Imprint	Added when two or more bodies are imprinted onto each other.

	Intersect Bodies	Added when two or more bodies are intersected to find the common volume that is shared by the bodies.
	Linear Pattern	Added when a linear pattern is created from a body.
	Linear Pattern Cut	Added when a linear pattern cut is created from a body.
	Loft	Added when two or more sketches are lofted to create a body.
	Loft Cut	Added when a lofted cut, created from two or more sketches, is used to remove material from the model.
	Mirror Body	Added when one or more bodies is mirrored.
	Move Body	Added when one or more bodies is translated.
	Revolve	Added when a sketch is revolved.
	Revolve Cut	Added when a revolve cut feature is used to remove material from the model by revolving a sketch around an axis.
	Rotate Body	Added when one or more bodies is rotated.
	Scale Body	Added when one or more bodies is scaled.
	Sew Sheet Bodies	Added when one or more sheet bodies are sewn together.
	Sketch	Added when a sketch is created.

	Subtract Bodies	Added when one or more bodies is used to subtract material from a target body.
	Sweep	Added when a sketch is swept along a profile to create a body.
	Sweep Cut	Added when a swept cut, which is created by sweeping a sketch along a profile, is used to remove material from the model.
	Unite Bodies	Added when two or more bodies are united.

See Also

Sketches
Creating Bodies From Sketches
Working with Patterns
Working with Volume Extraction
Working with Fillets and Chamfers
Working with Bodies

Working with the 3D-CAD Toolbar

You can create features in 3D-CAD using multiple methods. The default method is to right-click on objects in the 3D-CAD feature tree and in the *3D-CAD View* scene. An alternate method is to use the 3D-CAD toolbar.

This section contains the following topics:

- Activating the 3D-CAD Toolbar
- Using the 3D-CAD Toolbar

Activating the 3D-CAD Toolbar

By default, the 3D-CAD toolbar is deactivated. To activate the 3D-CAD toolbar:

- Select **Window** > **Toolbars** > **3D-CAD**.



The 3D-CAD toolbar is activated and appears next to the Selection toolbar. See Selecting Parts by Type in the Display.

Using the 3D-CAD Toolbar

There are two ways to use the 3D-CAD toolbar:

- Pre-selecting
- Post-selecting

The first method involves selecting an item before using the toolbar. For example, selecting a sketch profile before using the extrude function. The second method does not have any item active when using the toolbar, instead the desired item is selected after using the toolbar, when the feature panel is open.

The 3D-CAD toolbar is shown below.

The operations that are provided in the toolbar include:

- Create Sketch
- Create Transform Sketch Plane
- Create Extrude
- Create Extrude Cut
- Create Revolve
- Create Revolve Cut
- Create Sweep
- Create Loft
- Create Fillet
- Create Chamfer

Pre-Selecting

Pre-selecting means that you select the face, edge, or sketch to use with a feature before you click the relevant button in the 3D-CAD toolbar.

Using this method takes you directly to the relevant feature panel.

For example, if you select a sketch in the feature tree and then select the **Create Extrude** button in the 3D-CAD toolbar, you are taken to the *Extrude* panel.



Post-Selecting

Post-selecting means that you use the 3D-CAD toolbar with nothing pre-selected.

If no selection is active, the *Select Sketch Plane or Face* panel appears. Select a suitable face, edge, or sketch to continue to the relevant feature panel.

This panel differs slightly depending on what feature you are creating.

Creating a Sketch and Transform Sketch Plane

If you click **Create Sketch** on the 3D-CAD toolbar, and no selections are active in 3D-CAD, the *Select Sketch Plane or Face* panel appears.



1. Select the desired sketch plane from the drop-down list or a face from an existing body.
2. Click **Create Sketch** to continue to the *Sketch* panel.

The same procedure applies when you create a transform sketch plane. See Creating a Transform Sketch Plane.

Creating Any Other CAD Features

When you create any other CAD using the 3D-CAD toolbar, the normal feature panel appears.

1. Select the sketch or edge that you want to use from the *3D-CAD View* scene.
2. Set up the feature as described in the relevant feature section.

Surface Modeling in 3D-CAD

Surface modeling is the representation of three-dimensional geometry using zero-thickness surfaces.

In 3D-CAD, zero-thickness surfaces are called sheet bodies. 3D-CAD provides several functions that allow you to create and modify sheet bodies. You can also convert a sheet body to a solid body, or create a sheet body from a solid body.

Two important functions of surface modelling in 3D-CAD are that it allows you to repair imported CAD geometry, and also create zero-thickness baffles.

In 3D- CAD, sheet bodies are created in the same way as solid bodies—there are no separate features or toolbars. When it is possible to create a sheet body or a solid body, the *Body Type* option appears in the relevant feature panel, allowing you to choose what type of body you want to create. Currently, you cannot trim or remove areas from sheet bodies. The following tools are available for sheet bodies:

- Extrude
- Revolve
- Loft
- Sweep
- Chamfers
- Fillets
- Boolean operations
- Transform operations

3D-CAD enables you to perform the following useful functions with your existing geometries:

- Filling holes with sheet bodies
- Sewing sheet bodies together to form a solid body
- Creating a sheet body from a solid body

See Also

Sheet Bodies in 3D-CAD
Solid Bodies in 3D CAD

Sketches

Sketches are two-dimensional profiles that are created on a planar surface. They are composed from edges and points.

Many of the operations for creating solid bodes take sketches as their inputs. For example, the extrude operation requires a sketch as an input, along with the direction vector in which to extrude the sketch.

Every sketch is created on a sketch plane, which defines the position and orientation of the sketch in the 3D-CAD model.

You construct a sketch from sketch entities: lines, circles, arcs, splines, and more. Sketch entities are defined using the local coordinate system on the sketch plane.



Face Sketches

A Face Sketch is a sketch that you create by extracting the edges of a face on a solid body. You cannot modify these converted entities. The link between these sketch entities and the face is retained so that changes you make to the body face are propagated to the Face Sketch. The sketch entities that are linked to the face cannot be edited through

splitting or deleting operations, but you can convert them to construction entities.

See Also

Splitting and Trimming Sketch Entities
Deleting and Duplicating Entities
Working with Construction Entities

Creating a Sketch

Creating a sketch is often the first step in creating a solid body. Sketches are associated with planar surfaces, which are known as sketch planes.

To create a sketch:

1. Identify the planar surface on which to create the sketch. This planar surface can be one of the default sketch planes, a previously created transform sketch plane, or the surface of an existing body.

   See Sketch Planes.
2. Activate the sketch tool on the chosen surface. The method that you use depends on the type of surface:

   - For a default sketch plane, or a transform sketch plane, right-click on the corresponding feature node in the feature tree and choose **Create Sketch**.
   - For the face of an existing body, right-click on the face in the *Graphics* window and choose **Create Sketch**.

   The *Sketch* panel appears above the 3D-CAD object tree; a sketch grid with X and Y axes appears within the *Graphics* window.
3. In the *Sketch* panel, the *Display Options* box provides controls to orient and size the grid in the *Graphics* window. Configure the grid for the sketch using these controls.

   An example is to click **View Normal to Sketch Plane** so that you look directly down the normal axis of the sketch plane.

   See Sketch Display Options.
4. In the Sketch panel, the *Create Sketch Entities* box provides several controls for creating sketch entities. Use these controls as necessary to create the sketch.

   An example is to click **Create circle**, move the mouse pointer over the grid, and click twice to define the circle. The first click sets the center; the second click set the radius.

See Sketch Entity Reference for details of sketch entities.

5. Add constraints that fix the position of sketch entities relative to each other, and to specific locations.

   - Right-click on sketch entities in the *Graphics* window and choose the desired constraint from the context menu.

   Some constraints require that you select multiple sketch entities before they appear on the context menu.

   See Constraining a Sketch.

6. Set dimensions using constant sizes or parametric expressions.

   - Right-click on sketch entities in the *Graphics* window and choose the desired dimension from the context menu.

   See Dimensioning Sketches.

7. When the sketch is complete, click **OK** in the *Sketch* panel.

A **Sketch** node is added to the feature tree.

Editing a Sketch

You can modify existing sketches by editing them. Changes that are made to a sketch are propagated to dependent features further down the feature tree.

To edit a sketch:

1. Right-click the corresponding sketch node and select **Edit**.



This action automatically rolls back the 3D-CAD model to the selected sketch and activates Sketch Mode.

2. Modify your sketch using the tools available in the *Sketch* panel. You can create, edit, and delete sketch entities.
3. Click **Apply** to confirm the changes and exit the sketch.

Note: If the changes cause any feature in the model to fail to rebuild properly, the feature state of affected features changes. To resolve the issue, update the model, or edit the affected features.

See Also

Rolling Back
Understanding Feature States
Updating the 3D-CAD Model

Duplicating Sketches

You can make an identical copy of a sketch. To duplicate a sketch:

1. Right-click a **Sketch** node and select **Duplicate**.



An identical sketch feature is added to the feature tree. When you carry out this operation on a FaceSketch, an ordinary sketch feature is created.

Note: Duplicating a sketch also retains any design parameters that are associated with that sketch. Editing the design parameter affects both the original sketch and the duplicate sketch.

See Also

Sketches
Working with Design Parameters

Creating a Sketch from Existing Edges

A face sketch is a sketch that comes from the edges of a planar face on existing bodies.

To create a face sketch:

1. Right-click on the target face in the *3D-CAD View* scene and select **Create Sketch From Face Edges**.

STAR-CCM+ activates sketch mode, and adds the edges that define the face to the sketch.



2. If you intend to use the face edges for reference only, convert them to construction entities.
   - Multi-select the edges on the sketch, right-click, and choose **Set As Construction**.
3. Add other sketch entities as required.
4. When complete, click **OK** on the *Sketch* panel.

A **FaceSketch** node is added to the feature tree.

Sketch Planes

A sketch plane is a flat surface on which you can draw a sketch.

The sketch plane definition comes from the coordinates of its origin and two direction vectors. The two direction vectors provide the local x and y axes on the plane.

There are three types of sketch plane:

Default Sketch Planes

Default sketch planes are labeled **XY**, **YZ**, and **XZ**. These three sketch planes are the starting point for every 3D-CAD model. The default sketch planes are predefined—you cannot create, edit, or delete them. Typically, the first step when creating a 3D-CAD model is to create a sketch on one of the default sketch planes.

Transform Sketch Planes

Transform sketch planes, indicated by **TransformSketchPlane** nodes, are sketch planes defined relative to an existing sketch plane or body face. Transformation parameters define the relation of the sketch plane to the reference sketch plane or face. These parameters are a translation vector and three angles of rotation.

See Creating a Transform Sketch Plane.

User Sketch Planes

User sketch planes, indicated by a **Sketch Plane** node, are custom sketch planes that you create. You can use any available local coordinate system to define this type of sketch plane.

See Creating a User Sketch Plane.

A 3D-CAD model can contain any number of transform sketch planes and user sketch planes. You can also export the coordinate system of a sketch plane in 3D-CAD for use in STAR-CCM+.

See Also

Sketches
Creating a Sketch
Creating a Transform Sketch Plane
Exporting the Coordinate System from a Sketch Plane

Creating a Transform Sketch Plane

To create a transform sketch plane:

1. Select the base sketch plane using one of the following methods:
   - Right-click on one of the three default sketch plane nodes in the feature tree and select **Create Transform Sketch Plane**.
   - Right-click on an existing face of a geometry and select **Create Transform Sketch Plane**. A transform sketch plane that you

create on a face with zero translation or rotation lies directly on that face.

- Click **Create Transform Sketch Plane** in the 3D-CAD toolbar. If you do not pre-select a sketch plane or face when using the toolbar, select the desired sketch plane or face from the Select Sketch Plane or Face panel.

The *Transform Sketch Plane* panel appears.



2. Enter the transform parameters. See Transform Sketch Plane Panel.
3. If required, create design parameters.
4. Click **OK**.

A transform sketch plane is created and a **TransformSketchPlane** node is added to the 3D-CAD feature tree.

See Also

Creating a New Design Parameter

Transform Sketch Plane Panel

You can create a transform sketch plane from one of the default sketch planes or from a face on a geometry. Define the transform sketch plane using the following parameters:

Translation Vector	Specifies the direction vector for the translation of the sketch plane.

Rotation Angle	Specifies the angle of rotation for the sketch plane. Note that if you define a Translation Vector and a Rotation Angle, the sketch plane is translated according to the translation vector before the angle of rotation is applied.

Creating a User Sketch Plane

To create a user sketch plane:

1. Right-click **Features** and select **Create Sketch Plane**.



The *User Sketch Plane* panel appears.



2. Enter the sketch plane parameters. See User Sketch Plane Panel.
3. If required, create design parameters.
4. Click **OK**.

A user sketch plane is created and a **Sketch Plane** node is added to the 3D-CAD feature tree.

See Also

Creating a New Design Parameter

User Sketch Plane Panel

You can create a custom sketch plane using any available local coordinate system in STAR-CCM+. Define the sketch plane using the following parameters:

Coordinate System	Specifies the local coordinate system to use when defining the sketch plane. The coordinate system can be Cartesian, Cylindrical, or Spherical.
Origin	Specifies coordinates of the origin of the sketch plane.
X Direction	Specifies the direction vector for the X direction of the sketch plane.
Y Direction	Specifies the direction vector for the Y direction of the sketch plane.
Renormalize	Makes sure that the X and Y directions of the sketch plane are normal to each other. Clicking this button can result to changes to some of the specified values.

Changing the Sketch Plane of a Sketch

You can move sketches onto different sketch planes or faces to modify the position of its child features. For example, this facility permits you to relocate the position of an extruded body by associating its starting sketch with a different sketch plane.

To change the sketch plane on which a sketch is located:

1. Right-click on the relevant **Sketch** node in the feature tree and select **Edit Sketch Plane**.

The *Edit Sketch Plane* panel appears.



2. Select one of the following options:
   - To relocate the sketch onto a particular sketch plane, click the **Select Sketch Plane** radio button, then choose a sketch plane from the Sketch Planes list. You can only select sketch planes that are higher up in the feature tree than the selected sketch. Sketch planes that are below the selected sketch are not shown in the list. If necessary, you can insert alternative sketch planes further up in the feature tree.
   - To relocate the sketch onto a face, click the **Select Face** radio button then click the desired face in the *3D-CAD View* scene.
3. Click **Apply**.
4. Click **Close**.

The sketch and any dependent features are relocated according to the settings that you applied in the panel.

See Also

Sketch Planes
Understanding Where New Features Are Inserted

Exporting the Coordinate System from a Sketch Plane

You can export the coordinate system of a sketch plane in 3D-CAD for use in STAR-CCM+ so that it becomes available as a local coordinate system under **Tools > Coordinate Systems**.

To export the coordinate system:

1. Right-click on a sketch plane in the 3D-CAD object tree and select **Export Coordinate System**.



A new local Cartesian coordinate system is created and added to the **Local Coordinate Systems** folder of the **Laboratory** coordinate system. This local coordinate system is linked to the coordinate system of the original sketch plane in 3D-CAD.



In the *Properties* window of the sketch plane, the *Exported Coordinate System* property confirms the successful export of the coordinate system.



Caution: If you change the settings of the coordinate system in **Tools > Local Coordinate Systems > 3D-CAD [model_name] > 1**, these changes will be lost the next time that you update the 3D-CAD model. If you modify the sketch plane in 3D-CAD, however, the exported coordinate system is automatically updated with the changes.

You cannot create Cylindrical or Spherical coordinate systems directly from 3D-CAD. Instead, create these coordinate systems starting from the exported local coordinate system in the **Tools > Coordinate Systems** objects tree.

Sketch Entity Reference

Sketches are comprised of sketch entities. The following table describes the sketch entities that are available in 3D-CAD.

Sketch Entity	Description and Procedure
Point	A point is a single vertex. To create a point, click the location where you want to place it.

Sketch Entity	Description and Procedure
	
Line	A line is a straight edge that is constructed from a start point and an end point. To create a line, you require two mouse clicks: the first defines the start point, the second defines the end point. To create multiple lines, continue to define points after defining the end point of the first line. Click <`Esc`> to exit the line tool 
Rectangle	A rectangle is constrained to be aligned with the X and Y-axis directions on the sketch plane. To create a rectangle, you require two mouse clicks: one to define the location of the first corner and the second to define the opposite corner. Rectangles are automatically

Sketch Entity	Description and Procedure
	constrained so that one set of opposite lines is horizontal, and the other is vertical. 
Three-Point Rectangle 	A three-point rectangle can be oriented in any direction. To create this type of rectangle, you require three clicks: the first two clicks set the length and orientation of the rectangle; the third click sets the width. The rectangle is automatically constrained with perpendicular and parallel constraints. 

Sketch Entity	Description and Procedure
Circle	To define a circle, you require two mouse clicks: one to set the center of the circle, and the other to set the radius of the circle. 
Center-Point Circular Arc	A circular arc is a segment of a circle. To create a center-point circular arc, you require three mouse clicks: one to define the center of the arc, and two further clicks to set the start point and the end point. The arc is drawn anti-clockwise from the start point to the end point. STAR-CCM+ provides a function that converts and arc into its complement. This function is useful when the original arc is drawn in the opposite direction to the one that you require. See Converting an Arc Into its Complement.

Sketch Entity	Description and Procedure
	
Three-Point Circular Arc	A circular arc is a segment of a circle. To create a three-point circular arc, you require three mouse clicks: the first two clicks set the start and end points of the arc, and the third click sets the radius of the arc. STAR-CCM+ provides a function that converts and arc into its complement. This function is useful when the original arc is drawn in the opposite direction to the one that you require. See Converting an Arc Into its Complement. 
Ellipse	An ellipse is a closed loop in the shape of an oval or flattened circle. To create an ellipse, you require three

Sketch Entity	Description and Procedure
	mouse clicks: one to set the center of the ellipse, followed by two further clicks to define the minor and major radii. You can create an elliptical arc by splitting and trimming an ellipse. It is possible to apply minor and major radius dimensions by right-clicking on the sketch and selecting Apply Minor Radius or Apply Major Radius. 
Spline	A spline is a smooth curve that is composed from a series of points (called spline points). To create a spline, you require a minimum of three mouse clicks, which set the position of each spline point. You can add or remove points from the spline, as well as edit the position of each spline point individually. See Splines.

Sketch Entity	Description and Procedure
	

See Also

Constraint Types
Dimension Types

Sketch Display Options

This section describes how you can control the view in the *3D-CAD View* scene while you are working on a sketch.

Show Grid	Controls whether the grid on the sketch plane is visible or not. This option is stored in each sketch. When you edit or resume a sketch, these settings are maintained.
Snap to Grid	Controls whether the mouse pointer snaps to intersections in the grid lines. This option is activated by default. When the grid is turned off, the mouse pointer still snaps to it. This option is stored in each sketch. When you edit or resume a sketch, these settings are maintained.

Show/Hide Relations	Controls whether the dimension markers and constraint glyphs are displayed in the sketch. See Showing/Hiding Relations.
View Normal to Sketch Plane	Aligns the view so that it is normal to the sketch plane. The sketch plane appears automatically when in sketch mode. It displays the local axes of the sketch plane and the sketch plane origin.
Set Sketch Grid Spacing	Controls the distance between gridlines on the sketch plane. Enter the new value (with units) in the *Grid Spacing* dialog and click **OK**. Model dimensions in 3D-CAD must not exceed 10,000 m. Setting large values for grid spacing (>100 m) can result in errors when creating sketch entities.

Splines

A spline is a smooth curve that is composed from a series of points (called spline points).

In 3D-CAD, splines must be composed of at least three points.

Each point on a spline has additional properties that define the angle and curvature of the spline passing through the point. You can modify these properties using splines handles that appear when you select a spline.

A list of topics that cover more advanced features of splines is provided below:

Placing Splines on Existing Points

When you place a spline point on an existing point in the sketch, 3D-CAD assumes that it is the end of the spline, and exits the spline creation tool.



If you want to carry on adding spline points after passing through an existing point, press and hold the <`Ctrl`> key before clicking the existing point. This action allows you to create splines that pass through existing points in the sketch.



The spline point is placed on top of the existing point. To connect the two points, apply a coincidence constraint.

Note: Do not split splines using the **Split for Trim** operation. Doing so modifies the spline curvature.

See Also

Constraint Types
Splitting and Trimming Sketch Entities

Adding and Removing Points on a Spline

You can edit splines by selecting individual spline points and modifying them using the same technique as for points. You can also add and remove points from a spline.

To add a point to a spline:

- Click the (**Create point**) button, and click or near the required location on the spline.



You can modify this new point in the same manner as any other point on the spline.

To remove a spline point:

- Right-click the point and select **Delete**, or select the point and press the <`Delete`> key.



The spline redraws based on the remaining spline points.

See Also

Sketch Entity Reference

Controlling Spline Curvature

STAR-CCM+ provides more controls for the curvature of a spline passing through a point. These controls are:

- Angle: the angle of the spline, measured anti-clockwise from the x axis of the sketch plane.
- Magnitudes: parameters that control the curvature of a spline on either side of a point. The radius of the curve increases as the magnitude increases.

To modify spline curvature interactively:

1. Click the spline whose curvature you wish to modify.

Spline handles appear for each point on the spline. The curvature of the spline at a point is always tangential to the spline handle.



2. Click and drag the arrowheads of each spline handle as follows:
   - To change the amount of curvature in the direction of the arrow, extend or reduce the length of the arrow. Press `<Ctrl>` while moving the arrow to keep the current angle.
   - To change the angle of the spline, rotate the arrow around the spline point. Press `<Shift>` while rotating the arrow to keep the current length.
3. To complete the operation, click the sketch away from the spline.

You can also modify spline curvature using the *Spline Properties* panel that appears when you select a spline.

Note: Adding or deleting points to a spline with a modified curvature reverts the spline back to its original curvature. It is recommended that you add or remove points on a spline before editing its curvature.

Sketch Entity Operations

There are two general methods that you can use to edit a sketch entity:

- Editing sketch entity properties in the relevant properties panel.

  See Editing Sketch Entity Properties.
- Moving sketch entities interactively in the *Graphics* window.

  See Editing Sketch Entities Interactively.

Further operations are available, as described in subsequent topics within this section.

Editing Sketch Entity Properties

To have more control over sketch entities, you can edit the sketch entity settings through the *Properties* panel.

To edit a sketch entity:

1. In the *3D-CAD view* scene, click the sketch entity that you want to modify.
2. In the *Properties* panel, enter the values that you want to use for the sketch entity. The properties differ for each sketch entity



The sketch entity updates with the new settings. All coordinates are defined in the local coordinate system on the sketch plane.

Editing Sketch Entities Interactively

You can modify sketch entities in the 3D-CAD view by interacting with the lines and points that make up the entity.

To edit an entity:

- In the *3D-CAD View* scene, click and drag an entity. Clicking different parts of an entity affects it in different ways:
  - Dragging a point, stretches the entities that are connected to the point.

- Dragging a line, moves the entire entity and maintains its size and orientation. Neighboring entities stretch to maintain the connection to the moved entity.
- Dragging the edge of a circle, arc, or ellipse, modifies its radius.
- Dragging the center point of a circle, arc, or ellipse, moves the entire entity, maintaining the radius.

Working with Construction Entities

Construction entities are used within a sketch to aid the sketching process. 3D-CAD ignores construction entities unless the sketch is used to create a revolved feature.

When working with a sketch, the cursor snaps to a construction entity when it gets near to it. They are created by converting an existing sketch entity into a construction entity, and are only visible when in sketch mode. Construction entities are displayed as dashed lines within the sketch.

To create a construction entity:

1. Create a normal sketch entity.
2. Right-click on the sketch entity in the *3D-CAD View* scene and select **Set As Construction**.



Note: You can select multiple sketch entities by holding down the <`Ctrl`> key.

The solid entity changes from a solid line to a dashed line, to show that it is now a construction entity.

Using Construction Entities for Revolved Features

When you use a sketch to create either a revolved or a revolved cut feature, it must contain at least one construction line to define the axis about which the sketch revolves.


3D-CAD View 1
STAR-CCM+
X
Y
Z
X
Y


3D-CAD View 1
STAR-CCM+
Z
X
Y

See Also

Sketches
Sketch Entity Reference
Revolved Features

Splitting and Trimming Sketch Entities

You can trim overlapping sketch entities by splitting them and deleting the unwanted segments. When using the **Split for Trim** operation, the entity that you select first is split where it intersects the sketch entities that you select next. The sketch entities that are used for splitting can either fully overlap the target, or have a point lying on the target.

Note: When a circle is split, the result is two circular arcs, therefore it must be intersected at a minimum of two points by the splitting entities. The same must be done when splitting an ellipse.

To split a sketch entity:

1. Select the target sketch entity in the *3D-CAD View* scene (the circle that is highlighted in the screenshot below).
2. Hold down the <`Ctrl`> key and select the sketch entities that you want to use to split the target (the two perpendicular lines that are highlighted below).
3. Right-click on any of the selected entities and select **Split For Trim**.



The target sketch entity splits at the points of intersection. The entities that intersected the target sketch entity remain unchanged. You can use the delete operation to trim any unwanted entities, such as the arc highlighted below. To complete the sketch below, you could use **Split for Trim** to split each of the perpendicular lines, using the circular arc as the splitting entity in each operation.



It is good practice to check the connections between points after using the **Split for Trim** operation. To check whether the points are connected following a split operation:

4. Click the intersection point and drag it to one side.
5. If the two sketch entities come apart when the point is moved, apply a coincidence constraint between the two points to form a closed sketch. You can drag the point back onto the existing point on the sketch plane; a coincidence constraint is automatically applied.

Dimensions that are attached to a sketch entity are lost when the **Split for Trim** action is carried out.

See Also

Deleting and Duplicating Entities
Sketch Entity Reference
Constraint Types
Dimensioning Sketches

Creating An Elliptical Arc

An elliptical arc cannot be created using the center-point or three-point arc tools. Create elliptical arcs by splitting an ellipse.

To construct an elliptical arc:

1. Create an ellipse.
2. Create a sketch entity that intersects the ellipse at two points.



3. Use the sketch entity to split the ellipse by the Split for Trim operation. The result, after deleting the unwanted entities remaining, is an elliptical arc.



You can apply minor and major radius dimensions to elliptical arcs by right-clicking on the sketch and selecting Apply Minor Radius or Apply Major Radius. If you applied the minor or major radius dimension before splitting the ellipse, the dimension is retained, although the dimension annotation (and constraint) is removed from the scene. Reapply the dimension to see the annotation and add the constraint.

See Also

Sketch Entity Reference
Splitting and Trimming Sketch Entities
Dimension Types

Converting An Arc Into Its Complement

The complement of an arc is the arc that you get when you swap the start and end points of the original arc. This feature is useful when the original arc is drawn opposite to the arc that you require.



To convert an arc into its complement:

1. Right-click an arc in the *3D-CAD View* scene and select **Convert Into Complement**.

Offsetting Sketch Entities

You can uniformly displace sketch entities within the sketch plane. This operation is called an offset.



You can control the offset direction and the offset distance in each operation. This operation works with all sketch entities: lines, circles, ellipses, arcs, and splines. The offset operation is compatible with both open and closed sketches, and projected entities from outside the active sketch. The original sketch entities can be converted to construction lines or left as they are.

Note: The offset operation does not work with faces.

To offset a sketch entity:

1. Select the desired entity (or entities) that you wish to offset.
2. Right-click on one of the selections and select **Offset Sketch Primitive**.



The *Offset Sketch* dialog appears.



3. In *Offset Distance*, set the size of the offset. This value controls the distance between the base sketch entity and the offset sketch entity.
4. In *Offset Options*, select the type of offset that you want to perform. Three options are available. These options control the direction of the offset. In the examples that are shown below, the base entities are shown as highlighted and have been converted to construction lines:
   - **Outward**—in a closed sketch, the offset is on the outside of the enclosed sketch area.



- **Inward**—in a closed sketch, the offset is on the inside of the enclosed sketch area.



- **Bi-directional**—this causes the offset to take place on both sides of the selected sketch entities.



5. If you want to convert the base sketch to a construction line during the operation, activate *Make Base Construction*.
6. Click **OK**.

The offset operation is performed. You can undo/redo the operation using the appropriate buttons in the *Sketch* in-place dialog.

Offsetting Open Sketches

In an open sketch, the offset direction is set according to the direction of the sketch (that is, the order in which you created the vertices). For example, consider the case that is shown below:



Creating an offset with the *Offset Options* property set to **Outward** results in the offset being placed on the right-hand side of the sketch entity. If the *Offset Options* property is set to **Inward**, the offset appears on the left side of the original, as you look along the direction of the line.

See Also

Sketch Entity Reference
Projecting Entities to the Active Sketch

Deleting and Duplicating Entities

To duplicate a sketch entity:

- Right-click on the desired entity and choose **Duplicate sketch primitive**.



The sketch entity is duplicated. The duplicated sketch entity occupies the same location as the original. Move and edit it as required. You can duplicate multiple sketch entities at once by holding the <`Ctrl`> key while selecting them.

Deleting Sketch Entities

To delete a sketch entity:

- Right-click on the desired entity and choose **Delete**.
- Alternatively, select the sketch entity and press the <`Delete`> key.

You can delete multiple sketch entities at once by holding the `<Ctrl>` key while selecting them.

Moving Sketch Entities

To move a sketch entity:

1. Click and drag the sketch entity to its new position on the sketch plane. You can move several sketch entities at the same time by first holding down the <`Ctrl`> key and selecting the sketch entities that you want to move. Then, keeping the <`Ctrl`> key down, click on one of the selected entities, and drag it to a new position.



All of the selected sketch entities are moved according to the translation vector defined by the drag-and-drop operation.



Creating Sketch Fillets

The **Sketch Fillet** option rounds off the corner that is formed between two lines on a sketch.

You can only create a fillet between two lines at a time: these lines can be two straight lines, a straight line and a curve, or two curved lines. The two lines can be disjointed, intersecting or have a common end point.

To create a sketch fillet:

1. Select the two sketch lines that you want to fillet.
2. Right-click on one of these lines and select **Sketch Fillet**.



The *Dimensions* dialog appears.



3. Enter the radius of the fillet in the *Fillet Radius* textbox. You can define the fillet radius as a numerical value, or as an expression using design parameters.
4. If you set up a design parameter and want to expose it, activate the *Expose Parameter?*.
5. Click **OK**.

The fillet is applied to the selected entities.

Two sets of tangent and coincidence constraints are added to the sketch when the fillet is created. A radius dimension is also created. You can edit the radius dimension to control the fillet radius.

Note: The fillet radius must be smaller than the length of the shortest sketch line.

You can undo/redo the sketch fillet feature using the appropriate buttons in the *Sketch* in-place dialog.

Intersecting Lines

You can create fillets between intersecting lines. STAR-CCM+ retains the longest portions of the selected lines, and creates a fillet at the point where they cross. This operation is shown in the following diagrams:

There are two fillet locations (the two intersections), STAR-CCM+ chooses the lower intersection as this leaves the longer line segments. The remaining parts of the intersecting lines are discarded:

The sketch entities remain selected after the operation so that you can create a fillet again without the need to reselect the entities. In the example, creating a second fillet would connect the lines at the top:

Disjointed Lines

You can create fillets between disjointed lines. STAR-CCM+ extends each line to a point where it meets with the other line, and creates the fillet there. This operation is shown in the following diagrams:

3D-CAD View 1
STAR-CCM+
Y
Z
X
Y
X

Non-Intersecting Lines

You can create fillets between two non-intersecting sketch features. STAR-CCM+ places a fillet at the point that retains the two longest segments of the features.



In the following example, there are two possible places where a fillet could be placed (based on the specified fillet radius). The longest segments are retained when the fillet is placed on the top right, as shown highlighted in red:

3D-CAD therefore chooses this location, and the remaining segments are discarded:

See Also

What Is a Design Parameter?
Constraint Types
Dimension Types

Creating Sketch Chamfers

A sketch chamfer is a feature that makes a straight cut across the corner that is formed between two lines on a sketch.

You can only create chamfers between two lines at a time; these lines can be disjointed, intersecting or have a common end point. When the lines intersect each other, the chamfer is created between the two longest segments of the lines.

To create a sketch chamfer:

1. Select the two sketch lines that you want to chamfer.
2. Right-click on one of these lines and select **Sketch Chamfer**.



The *Sketch Chamfer* dialog appears.



3. Enter values for *Distance 1* and *Distance 2*. These values define the size of the chamfer and are the distances that the chamfer will be from the point where the two lines meet.

Note: You can define the chamfer using a numerical value, or by an expression using design parameters. Activate *Expose Parameter ?* for each distance property to expose it as a design parameter.

4. Click **OK**.

The chamfer is applied to the selected entities.

Two sets of coincidence constraints are added to the sketch when the chamfer is created. Two distance dimensions are also created; you can edit these dimensions to control the chamfer size.

Note: The dimensions controlling the chamfer must be smaller than the length of the shortest original sketch line. Once created, the chamfer dimensions can be edited. The sketch chamfer feature can be undone/redone using the appropriate buttons in the *Sketch* in-place dialog.

See Also

What Is a Design Parameter?
Constraint Types
Dimension Types

Projecting Entities to the Active Sketch

You can project sketch entities from other sketches onto the active sketch.

Additionally, you can project edges and faces from other bodies to the active sketch.



To project entities:

1. Select the desired edges or faces that you wish to project.
2. Right-click on one of the selections and select **Project to sketch**.



The selected entities (or in the case of a face, the edges of the face) are projected onto the active sketch. Projected entities are displayed in orange for easier identification.



Note: The projected entities are associative, meaning that if you modify the original entity, the projection is automatically updated to match the modification.

You can use projected entities with other sketch entities in the current sketch. You can undo or redo the project feature operation using the appropriate buttons in the *Sketch* in-place dialog.

You cannot move or modify projected entities. You can, however, create a duplicate of the projected entities and modify those.

Note: The duplicated entities are not associated with the original source.

See Also

Sketch Entity Reference

Showing/Hiding Relations

You can turn off relations, such as constraints and dimensions when the sketch scene becomes too cluttered, as shown in the following screenshot.



You can visualize dimensions and constraints independently of each other.

To control the visualization of the relations:

1. Click the (**Show/hide relations**) button on the *Sketch* panel.

   The *Show/hide relations* dialog appears.



2. Here, you can control what is displayed in the sketch scene:
   - Controlling the display of dimensions
   - Controlling the display of constraints
3. Click **OK** to close the dialog and apply the changes.

Controlling the Display of Dimensions

To control the display of the dimensions in the sketch scene, select one of the options from the drop-down menu:

- **Show dimensions**—shows the dimensions in the sketch scene, including expressions and design parameter names.



- **Show dimension values only**—shows only the dimension values in the sketch scene. Expressions and design parameters are converted to dimension values. This option reduces the amount of text that appears in the sketch scene.



- **Hide dimensions**—hides all dimensions in the sketch scene.



Controlling the Display of Constraints

To control the display of the constraints, activate or deactivate the *Show constraints* checkbox:

- *Activated* – the constraints are displayed in the sketch scene.



- *Deactivated* – all constraints are hidden in the sketch scene.



See Also

Constraining a Sketch
Dimensioning Sketches

Constraining a Sketch

Constraints are used to impose geometric limitations on sketch entities as well as to control their position on the sketch plane.

You can apply a constraint to a single sketch entity to define how it relates to the sketch plane, or to a pair of sketch entities to define the relationship between them. It is possible to constrain the active sketch using sketch entities from an existing sketch or geometry. The constraints that are shown in the right-click menu differ depending on the selection of sketch entities.

To apply a constraint:

1. In the *3D-CAD view* scene, select the sketch entities that you want to constrain.
2. Right-click one of the selections and select the type of constraint you want to apply. Only constraints that are relevant to your selection appear in the right-click menu. See Constraint Types.

When you apply a constraint to a pair of sketch entities, 3D-CAD can move one of the sketch entities in order to satisfy the constraint. You cannot influence which sketch entity moves—3D-CAD determines which entity to move internally. If, however, you want to ensure that a particular sketch entity does not move when the constraint is applied, apply a fixation constraint to one of the sketch entities.

To use the sketch plane origin when constraining a sketch, create a point at `[0, 0]` and fix it in position. You can then use this point to constrain sketch entities in relation to the origin.

If the new constraint cannot be satisfied, due to any existing constraints or dimensions on the sketch, an error message is shown and the constraint is not applied. To resolve this issue, delete any existing constraints or dimensions that are counteracting the action of the new constraint.

3D-CAD considers lines to be infinitely long and circular arcs to be complete circles when constraints are applied. Therefore, a point that is constrained to coincide with a line does not always appear to lie on the line. Similarly, a point that coincides with a circular arc does not appear to lie on the arc. In each case, the point is coincident with some point on the line segment that defines the infinite line, or the arc that defines the full circle.

See Also

Sketches
Sketch Planes
Sketch Entity Reference
Constraint Types
Removing Constraints
Dimensioning Sketches

Constraint Types

A list of available coincidence constraints is provided in the following table.

Constraint	Description
Coincidence	A coincidence constraint forces a point on a sketch entity to coincide with another sketch entity on the sketch plane. The second sketch entity can be a line, circle, circular arc, or another point. You can apply a coincidence constraint between two circular arcs, causing the arcs to share center point and radius. 
Collinear	A collinear constraint forces two sketch entities to lie on the same infinite line. You can only apply collinear constraints to lines. Upon applying the constraint, one line maintains its position while the second line moves to meet the requirements of the constraint.

Constraint	Description
	
Concentric	A concentric constraint forces the selected sketch entities to share center points. You can apply concentric constraints to circles, circular arcs, and ellipses. You can also apply concentric constraints between these sketch entities and a point, which forces the center point of the circle/arc/ellipse to coincide with the point. 

Constraint	Description
Equal Length	An equal length constraint forces two line sketch entities to be the same length, and can only be applied to lines. You can use equal length constraints with a length dimension, which fixes the length of both lines. 
Equal Radius	An equal radius constraint sets the radius of two circular sketch entities to be equal. You can apply this constraint to circle and circular arc sketch entities. To specify which circle governs the size of the radius, apply a radius dimension to that particular circle before applying the equal radius constraint.

Constraint	Description
	
Fixation	A fixation constraint anchors the position of a single sketch entity on the sketch plane. You can apply fixation constraints to points, circles, circular arcs, ellipses, and spline points. When this constraint is used with a circular sketch entity, it is automatically applied to the center point. To fix the size of the circular entity, apply a radius dimension. To fix the position of a line sketch entity, apply a fixation constraint at the start and end points of the line.

Constraint	Description
	
Horizontal	A horizontal constraint aligns a line sketch entity with the local X axis on the sketch plane. In most cases, the start point of the line is moved to satisfy the constraint, but this behavior depends on which constraints and dimensions have already been applied to the sketch. The length of the line can change when you apply the constraint (unless you apply a length dimension). 

Constraint	Description
Vertical	A vertical constraint aligns a line sketch entity with the local Y axis on the sketch plane. In most cases, the start point of the line moves to satisfy the constraint, but this behavior depends on which constraints and dimensions are already applied to the sketch. The length of the line can change when you apply the constraint (unless you apply a length dimension). 
Midpoint	A midpoint constraint forces a point sketch entity to lie on the midpoint of a line. You can apply midpoint constraints between a point and a line sketch entity.

Constraint	Description
	
Parallel	A parallel constraint forces two line sketch entities to be parallel. You cannot apply a parallel constraint to two lines are at right angles to each other. 
Perpendicular	A perpendicular constraint forces two line sketch entities to be aligned at right angles to each other. You cannot apply a perpendicular constraint to two parallel lines.

Constraint	Description
	
Tangent	A tangent constraint forces an edge (of a circular, elliptical, or line sketch entity) to be tangential to a selected circle, circular arc, or ellipse. You can apply this constraint to circle, circular arc, ellipse, or line sketch entities. 

Constraint	Description
Pierce	When two sketch planes intersect, you can apply a pierce constraint between a point sketch entity (from the active sketch) and an edge sketch entity (from the reference sketch). This type of constraint forces the point sketch entity to pierce the edge sketch entity. This feature is useful when creating sweeps and guided lofts, as it ensures coincidence between the guide sketch and profile sketch.  In the above example, the two sketch entities are on different planes that are perpendicular to each other. This method would also work for any two sketches that have intersecting planes. If two locations exist where the point can pierce the reference sketch, the edge closest to the point is pierced.

See Also

Dimension Types
Sketch Entity Reference
Swept Features
Using Guide Wires in Loft Operations

Applying Constraints that Reference Other Sketches or Existing Geometries

You can select an edge from another sketch or geometry when creating a constraint in the active sketch.

The selected edge is projected onto the active sketch and is used to constrain the target entity in the active sketch. If the plane of the active sketch intersects the plane of the reference sketch, you can use the pierce constraint. See Pierce Constraint. This constraint is only available when selecting a point from the active sketch and an edge from the reference sketch, and ensures that the two sketch entities are coincident.

An example demonstrating how the projection works is shown below.



This feature gives you the ability to align sketch entities with one or more existing sketches. The same is also true for reference geometries, where the selected edge on the reference geometry is projected onto the active sketch. You can only use edges from a reference sketch/geometry to constrain an active sketch. You cannot use vertices from the reference sketch/geometry. When you modify a reference entity, the projection of that entity on the constrained sketch is also modified. This action can change the constrained sketch. If you delete a reference entity, the constraint that is based on that projection is also deleted. Using a similar method, you can pierce another sketch on a different plane, ensuring that the active sketch and the reference sketch are coincident.

In the following example, the gray geometry is used to constrain the active sketch (shown in blue).



To constrain the active sketch using a reference geometry:

1. Select the edge that you want to use as a reference from the reference geometry. This edge highlights in purple.

2. While holding down the `<Ctrl>` key, right-click the desired entity on the active sketch and select the required constraint from the right-click menu. In this case, the line that is on a diagonal must be parallel to the selected edge of the reference geometry, so **Apply Parallel Constraint**is selected.



The diagonal edge from the reference geometry is projected to the active sketch and the constraint operation is carried out.

Note: The projection does not appear on the active sketch. This operation is internal.

A symbol denoting the type of constraint is placed on the constrained sketch entity.

The two diagonal lines are now parallel to each other.

Reviewing Applied Constraints

When constraints have been used in a sketch, glyphs are displayed next to the constrained sketch entities. By selecting a constraint glyph in the *3D-CAD View* scene, you can see which entities the constraint applies to, and also delete the constraint. Constraints can be hidden from the sketch scene.

To review constraints that are applied to a particular sketch entity:

1. Select the sketch entity in the *3D-CAD View* scene. A list applied constraints are displayed in the *Properties* box in the *Sketch* panel.
2. Select a constraint from this list. All sketch entities that the constraint applies to become highlighted in the *3D-CAD View* scene.



You can use this list to remove a constraint from the sketch.

See Also

Removing Constraints
Showing/Hiding Relations

Removing Constraints

You can remove constraints from a sketch in two ways:

- By selecting the constraint glyph in the *3D-CAD View* scene.

- By using the constraints list in the *Sketch* panel.

Removing Constraints From the 3D-CAD View Scene

To remove a constraint directly in the *3D-CAD View* scene:

- Right-click on the constraint glyph in the *3D-CAD View* scene, and select **Delete**.



- Alternatively, select the constraint glyph and press the `<Delete>` key. Only select one of the glyphs where a constraint is applied to a pair of sketch entities.

Removing Constraints From the List of Constraints

To remove a constraint using the list of applied constraints in the *Sketch* panel:

1. Select the sketch entity in the *3D-CAD View* scene, from which you wish to remove a constraint. In the *Sketch* panel, a list of constraints that are applied to the selected sketch entity is displayed.
2. Right-click on the constraint you wish to remove and select **Delete**.



See Also

Constraining a Sketch
Reviewing Applied Constraints

Dimensioning Sketches

Dimensions are used to control sizes, distances, and angles in a sketch.

When you apply a dimension to one or more sketch entities, it fixes the size, distance, or angle to a specific value. The value can be a numerical quantity, or specified using design parameters, and can only be changed by editing the dimension.

Caution: There is a numerical limit to the geometrical precision of 3D-CAD models. Do not apply dimensions less that 1 micron (1.0E-6 m).

To dimension a sketch:

1. In the 3D-CAD view scene, right-click on the desired sketch entities and select the type of dimension that you want to apply. See Dimension Types.
2. Enter the dimension values in the *Dimension* dialog. You can specify the units by typing them directly in the textbox.



3. To expose the dimension as a design parameter, activate *Expose parameter?*
4. Click **OK**.

If necessary, you can hide dimensions in the sketch to give you a clearer view of the sketch geometry. Alternatively, you can reposition the dimension.



See Also

Sketches
Working with Design Parameters
Showing/Hiding Relations

Dimension Types

A list of available dimensions, with a short description, is provided in the following table. For information on how to apply dimensions, see Dimensioning Sketches.

Dimension	Description
Angle	An angle dimension specifies the angle between two line sketch entities. The angle between the two lines must be greater than zero and less than `360` degrees. Additionally, the angle must not be close to zero or `180` degrees. The tolerance for an angle dimension is `0.001` radians, which is approximately `0.057` degrees.

Dimension	Description
	 Selecting the lines in the opposite order allows you to apply the opposite angle dimension. For example, 25 degrees rather than 335 degrees.
Distance	A distance dimension sets the straight-line distance between two selected sketch entities. You can apply this dimension between a point and a line, or between two points. 

Dimension	Description
Horizontal Distance	A horizontal dimension sets the distance in the X-direction between two points. 
Vertical Distance	A vertical dimension sets the distance in the Y-direction between two points. 
Length	A length dimension sets the length of a line sketch entity.

Dimension	Description
	
Radius	A radius dimension sets the radial size of a circle or circular arc sketch entity. 
Minor Radius	The minor radius dimension sets the small diameter of an ellipse sketch entity.

Dimension	Description
	
Major Radius	The major radius dimension sets large diameter of an ellipse sketch entity. 

See Also

Sketch Entity Reference

Reviewing Applied Dimensions

You can review sketch dimensions by displaying them in the *3D-CAD View* scene.

If you exposed dimensions as a design parameter, an expression, such as `Length = 0.01 m`, is shown next to the dimension arrow. If you defined a dimension by an expression, such as `2*$Length`, the expression is displayed next to the dimension arrow. You can select dimensions that are displayed in the *3D-CAD View* scene and edit or delete them as necessary.

To view a list of dimensions that are applied to a particular sketch entity:

1. Select a sketch entity in the *3D-CAD View* scene.

   The list appears in the *Properties* box in the *Sketch* panel.



2. Select a dimension from this list.

   The relevant sketch entities become highlighted in the *3D-CAD view* scene. You can also use this list to edit or delete a dimension, as necessary.

See Also

Showing/Hiding Relations
What Is a Design Parameter?

Editing the Value of Dimensions

You can edit the value of a dimension in two ways:

- By selecting the dimension in the *3D-CAD View* scene.
- From the Dimensions list in the *Sketch* panel.

Either of these methods opens the *Dimension* dialog, where you can enter a new numerical value or expression in the textbox. Editing the dimension also allows you to expose the dimension as a design parameter, or conceal a previously exposed design parameter.

Editing Dimensions Using the 3D-CAD View Scene

To edit a dimension directly in the *3D-CAD View* scene:

1. Right-click on the dimension marker in the sketch and select **Edit Dimension**.



The *Dimension* dialog appears.

1. Enter the new value.
2. Click **OK**.

Editing Dimensions Using the Sketch Panel

To edit a dimension using the list of applied dimensions in the *Sketch* panel:

1. Select the sketch entity to which you applied the dimension. In the *Sketch* panel, a list of dimensions that are applied to the selected sketch entity is displayed in the *Properties* box.
2. Right-click on the dimension that you want to edit and select **Edit Dimension**.



The *Dimension* dialog appears.

1. Enter the new value.
2. Click **OK**.

See Also

What Is a Design Parameter?

Repositioning Dimensions

You can reposition sketch dimensions in the sketch scene.

This feature is useful for moving dimensions that are blocking the view of the sketch entities and sketch constraints, or for improving the presentation of the sketch. The following screenshot shows a cluttered sketch.



Rearranging the dimensions makes it easier for you to work with the sketch.



You can move all dimension types in 3D-CAD.

To move a dimension:

1. Click any part of the dimension in the `3D-CAD View` window. The selected dimension, and the sketch entities that correspond to it, are highlighted in the sketch scene.



2. Drag the dimension to its new position. The dimension label snaps to your mouse pointer and follows it until you release the mouse button.



While moving the mouse pointer, the dimension label also snaps to each grid intersection.


3D-CAD View 1
STAR-CCM+
0.12 m
Y
Y
X
Z
X


3D-CAD View 1
STAR-CCM+
0.12 m
Y
Y
X
Z
X

If the dimension label is moved beyond the extent of the sketch entity, a projection appears.



Angle dimensions are projected in the same way as radius dimensions.


3D-CAD View 1
STAR-CCM+
45 deg
Y
Z
X
Y
X

Removing Dimensions

You can remove dimensions from a sketch in two ways:

- By selecting the dimension in the *3D-CAD View* scene.
- From the Dimensions list in the *Sketch* panel.

Note: Deleting a dimension that is exposed as a design parameter does not delete the associated design parameter.

Removing Dimensions Using the 3D-CAD View Scene

To remove a dimension directly in the *3D-CAD View* scene:

- Right-click on the dimension marker in the sketch and select **Delete**.



- Alternatively, select the dimension and press the `<Delete>` key.

Removing Dimensions Using the Sketch Panel

To remove a dimension using the list of applied dimensions in the *Sketch* panel:

1. Select the sketch entity in the *3D-CAD View* scene from which you want to remove a dimension. In the *Sketch* panel, a list of dimensions that are applied to the selected sketch entity are displayed in the *Properties* box.
2. Right-click on the dimension that you want to remove and select **Delete**.



See Also

What Is a Design Parameter?

Creating Bodies From Sketches

In 3D-CAD, you create bodies by extruding, revolving, lofting, or sweeping sketches. Once a body is created, you add or remove material from it.

Use the following features to create bodies from sketches:

- Extrusion Features—An extrusion feature projects a chosen sketch in a specified direction to create a body or a cut.
- Revolved Features—A revolve feature revolves a sketch around a chosen axis to create a body or cut. The axis of revolution is defined by a construction line in the same sketch.
- Lofted Features—A loft fills the space between two sketches to form a solid or sheet body. At least two sketches containing the loft profiles are required. You can also use guide wires to guide the loft from one profile to the next. A loft cut uses the same principle as a loft to remove material from a solid body. You require at least two sketches containing the profiles of the loft. You can also use guide wires to guide the loft.
- Swept Features—A sweep requires two sketches, one to define the sweep profile, and the other to define the sweep path. The profile is swept along the path. A sweep cut requires two sketches, one to define the sweep profile, and the other to define the sweep path.

You can also create a sheet body from a closed sketch. See Creating Sheet Bodies from Sketches.

See Also

Sketches
Overview of Bodies in 3D CAD

Extrusion Features

An extrusion projects a chosen sketch in a specified direction to create a body or a cut.

If the sketch is a closed profile, the extrusion forms a solid body. If the sketch is an open profile, the extrusion forms a sheet body. Open profiles do not form cuts; they only form new sheet bodies. In an extruded cut, the solid body that is formed by projecting the closed sketch is subtracted from any other solid bodies in its path.





When creating an extruded body or cut, you can set a draft angle that expands or contracts the profile along the extrusion distance.





Requirements

For solid bodies:

- The sketch must contain only closed loops, with no free edges.
- Sketch entities must not overlap or lie on top of each other.
- For sketch entities containing two closed profiles, one profile must be entirely enclosed by the other. You cannot define more than two closed profiles for an extrusion.

For sheet bodies:

- Open sketches are supported.

- Sketch entities can overlap or lie on top of each other.

Construction lines are ignored in the extrusion process.

See Also

Sketches
Working with Construction Entities

Creating an Extrusion Feature

Use the same procedure to form an extrusion body or cut.

To create an extrusion feature:

1. Either:
   - Right-click the sketch feature that you wish to extrude and choose **Create Extrude** or **Create Extrude Cut**.



   - Click (**Create Extrude**) or (**Create Extrude Cut**) in the 3D-CAD toolbar and choose the source sketch for the extrusion.

The *Extrude* panel appears, with the identity of the chosen sketch shown in the bottom of the panel.



2. Enter the extrusion parameters. See Extrude Panel.
3. Expose any properties that you want as design parameters by clicking the (**Expose Parameter)** button next to the property.

   The *3D-CAD view* scene updates to show a preview of the body.



4. Click **OK**.

See Also

What Is a Design Parameter?

Extrude and Extrude Cut Panels

Use the following properties to define the behavior of extrusion features.

Body Type	Defines the type of body that results from the operation.	
	Solid	The resulting body is a closed solid.
	Sheet	The resulting body is an open sheet body.
Method	Defines how to set the end-point of the extrusion.	
	Blind	Extrudes the sketch to the distance that you set in the **Distance** property.
	Through All	Extends the sketch through all bodies, cutting everything in its path.

Up To Face	Extends the sketch to a selected face or its underlying surface. Click the desired face to select it. The face that you select is displayed in the *Selected Face* property. 
Up to Body	Extends the sketch up to a selected body. Click the desired body to select it. The body that you select is displayed in the *Selected Body* property. 
Beyond Face	Extends the sketch beyond the selected face by the amount that you set in the **Offset Distance** property.

		
Direction	Specifies the direction of the extrusion.	
	Normal	Extrudes the sketch in the normal direction.
	Reverse	Extrudes the sketch opposite to the normal direction.
	Specified	Extrudes the sketch in a custom direction, whose vector components are specified in the *Direction Vector* property.
Direction Vector	Specifies the custom extrusion direction. This property is activated when you set the *Direction* property to **Specified**.	
	X	X component.
	Y	Y component.
	Z	Z component.
Coordinate System	Specifies the coordinate system that is used in the extrusion operation.	
Extrusion Options	Specifies whether the sketch extends in a single direction or both directions.	
	One Way	Extrudes the sketch in one direction (specified in the *Direction* parameter).
	Two Way Symmetric	Extrudes the sketch the same distance in both directions.
	Two Way Asymmetric	Extrudes the sketch in both directions. The extrusion on either side of the sketch can be a different size. This

		option activates the **Asymmetric Distance** text box.
Distance	Specifies the length of the extrusion.	
Asymmetric Distance	Specifies the asymmetric extrusion distance. This option is activated when you set the *Extrusion Options* property to **Two Way Asymmetric**.	
Offset Distance	Specifies the offset distance of the extrusion from the selected surface. The offset distance is measured from the surface that you select.	
Draft	Specifies whether the extrusion has a constant cross-sectional profile along its length, or whether it tapers inwards or expands outwards.	
	None	Extrusion has a constant cross-section along the extrusion.
	Inward	Extrusion tapers inwards. When applying an inward draft to a concave profile, the profile must not contract to the point where its edges overlap and intersect. Similarly, if the draft splits the original profile into two profiles, the operation fails.
	Outward	Extrusion expands outwards.
Angle	Specifies the draft angle (in degrees or radians). This option is activated when you set the *Draft* property to **Inward** or **Outward**. Enter values in the range of 0 through 90 degrees, or 0 through p/2 radians	
Body Interaction	Specifies how the solid body that is created by the extrusion interacts with other bodies adjacent to it.	
	None	The new body is added to the **Bodies** manager on its own. No change is made if the body is adjacent to, or overlaps, surrounding bodies.
	Merge	The new body is united with any adjacent or overlapping bodies; otherwise it is added on its own to the **Bodies** manager.
	Imprint	The edges of the extruded body are imprinted on any adjacent or overlapping bodies.

Draft Angle Limitation

When applying a draft angle to an inward curving profile, make sure that the projected surface does not split along the extruded distance.

The source profile that is shown in the following example splits when a draft angle is applied along the extrusion.

	Valid: the extrusion stops before the profile is split.
	Invalid: the profile splits due to the combination of draft angle and extrusion length.
	Valid: there is no profile.

Revolved Features

A revolved feature is an object whose cross-sectional profile is constant around a specified axis.

It is created by sweeping a two-dimensional profile, defined in a sketch, around a specified axis by a given angle.



A revolved cut is defined in the same way; the volume that is defined by the revolved profile is subtracted from any bodies through which it cuts.



Requirements

For solid bodies:

- The sketch must contain only closed loops, with no free edges.
- Sketch entities must not overlap or cross each other.
- For sketch entities containing two closed profiles, one profile must be entirely enclosed by the other. You cannot define more than two closed profiles for a revolve feature.

For sheet bodies:

- Open sketches are supported.
- Sketch entities can overlap or lie on top of each other.

The sketch must contain at least one construction line, which defines the axis about which the sketch revolves.

See Also

Sketches

Working with Construction Entities

Creating a Revolved Feature

The Revolve and Revolve Cut features share the same procedure. To create a revolved feature:

1. Either:

   - Right-click the sketch feature that you wish to revolve and choose **Create Revolve** or **Create Revolve Cut**.



   - Click (**Create Revolve**) or (**Create Revolve Cut**) in the 3D-CAD toolbar and choose the source sketch for the revolve feature.

   The *Revolve* panel appears.



2. Enter the revolve parameters. See Revolve Panel.
3. Expose any properties that you want to use as design parameters by clicking the (**Expose Parameter)** button next to the property.

   A preview of the revolution is shown in the *3D-CAD View* scene.



4. Click **OK** to perform the revolve.

See Also

What Is a Design Parameter?

Revolve and Revolve Cut Panels

Use the following properties to define the behavior of the revolve or revolve cut feature.

Body Type	Defines the type of body that results from the operation.	
	Solid	The resulting body is a closed solid.
	Sheet	The resulting body is an open sheet body.
Direction	Specifies the direction of the revolution.	
	Normal	Revolves the sketch in the normal direction.
	Reverse	Revolves the sketch opposite to the normal direction.
Angle	Specifies the angle through which the sketch is revolved.	

Axis	Specifies the axis of about which the sketch revolves.		
	Line n	Specifies a construction line as the axis or rotation. If the selected sketch contains multiple construction lines, these are listed here.	
	Specified	Allows you to manually specify the axis of rotation direction and position.	
		Axis Direction	Specifies the direction of the axis of revolution. Enter the [X, Y, Z] components of the direction vector.
		Axis Position	Specifies the position of the axis of revolution. Enter the [X, Y, Z] coordinates for the axis.
		Coordinate System	Defines the coordinate system for the axis of revolution. If you have specified an alternative coordinate system within the simulation, you can define the axis with reference to this coordinate system instead of the default, **Laboratory**.
Body Interaction	Specifies how the solid body that is created by the revolve feature interacts with other bodies adjacent to it.		
	None	The new body is added to the **Bodies** manager on its own. No change is made if the body is adjacent to, or overlaps, surrounding bodies.	
	Merge	The new body is united with any adjacent or overlapping bodies; otherwise it is added on its own to the **Bodies** manager.	
	Imprint	The edges of the revolved body are imprinted on any adjacent or overlapping bodies.	

Lofted Features

A lofted feature creates a body by fitting surfaces between a series of profile sketches.

The cross-section of this body varies along the length of the feature according to the shape of each profile sketch. The start and end sketches define the faces that cap either end of the solid body. You require a minimum of two sketches to create a single loft.



A lofted cut is defined in the same way; the volume that is defined by the lofted profile is subtracted from any bodies through which it cuts. You can only use lofted cuts on solid bodies.



When 3D-CAD creates the lofted feature, it matches each vertex on one sketch with a vertex on the next sketch, and fits curves through the pairs for the fitted surfaces. You can manually control which vertices make up these pairs using a feature called Vertex Matching. This technique is useful when you are creating a loft between several sketches, between sketches that have complex geometry, or to define an alternative loft path.

You can further control the shape of the loft by using guide wires. This feature enables you to create complex shapes with twists along the length of the solid.

Requirements

The following criteria must be satisfied for the loft to complete successfully:

- Each profile sketch must contain a single, closed loop to define the shape of the loft.
- The number of points that lie on the closed loop must be the same for each sketch, unless, for two consecutive sketches, one loop has

no points, such as a circle or ellipse. For lofting, splines are considered to have no points. In this case, a loop that defines a single, closed spline is treated in the same way as a circle or ellipse, regardless of the number of points on the spline.

Some examples of the closed loops that can be used to created lofts are shown in the following images:

- A circle and an ellipse (two closed loops that contain no points):



- A circle and a polygon (a closed loop with no points and a closed loop with any number of points):



- Two polygons (two closed loops that have the same number of points):



When designing a lofted feature, observe the following points:

- Sketches can lie on sketch planes that are oriented in any direction, but no two consecutive sketches can lie on the same plane.
- You cannot create a lofted feature that passes through itself at any point.
- The shape and position of sketches must avoid any radical changes in the cross-sectional profile of the body.
- Ensure that sketches do not result in any self-intersections in the solid body.
- Construction lines have no impact on the loft.

See Also

Using Guide Wires in Loft Operations
Overview of Bodies in 3D CAD

Creating a Lofted Feature

The Loft and Loft Cut features share the same procedure. To create a lofted feature:

1. Press `<Ctrl>` key and select the sketch profiles in the order in which the loft passes through them.
2. Either:
   - Right-click on one of the profiles and select **Create Loft** or **Create Loft Cut**.



   - Click (**Create Loft**) in the 3D-CAD toolbar.

   The *Loft* panel appears.



3. If the loft requires guide wires to correctly define the loft, activate the *Use Guide Wires* option. Follow the procedure in Using Guide Wires in Loft Operations.
4. If vertices on successive profiles must be matched manually to get the correct orientation, activate the *Match Vertices Manually* option. Follow the procedure in Matching Vertices Manually.
5. Set the remaining loft parameters as you require. See Loft and Loft Cut Panels.
6. Expose any properties that you want to use as design parameters by clicking the (**Expose Parameter)** button next to the property.

   A preview of the lofted feature appears in the *3D-CAD View* scene.



7. Click **OK** to create the loft.

See Also

What Is a Design Parameter?

Loft and Loft Cut Panels

Use the following properties to control the behavior of the loft feature or loft cut.

Body Type	Defines the type of body that results from the operation.	
	Solid	The resulting body is a closed solid.
	Sheet	The resulting body is an open sheet body.
Use Guide Wires	Provides the option to constrain the loft using guide wires. See Using Guide Wires in Loft Operations.	

Create Body Normal to Start Sketch	Specifies whether the loft body is perpendicular to the start sketch.	
Create Body Normal to End Sketch	Specifies whether the loft body is perpendicular to the end sketch.	
Match Vertices Manually	Provides the option to manually match vertices across successive sketch profiles. See Matching Vertices Manually.	
Body Interaction	Specifies how the solid body that is created by the loft interacts with other bodies adjacent to it.	
	None	The new body is added to the **Bodies** manager on its own. No change is made if the body is adjacent to, or overlaps, surrounding bodies.
	Merge	The new body is united with any adjacent or overlapping bodies; otherwise it is added on its own to the **Bodies** manager.
	Imprint	The edges of the lofted body are imprinted on any adjacent or overlapping bodies.

Matching Vertices Manually

By default, 3D-CAD predicts the pairs of matching vertices between successive profiles in lofts. Alternatively, you can use the Vertex Matching feature to manually define the matching pairs.

You can also use vertex matching to define an alternative loft path, which enables you to twist the body between the profile sketches. During the process of matching vertices, the path between sketch profiles is determined by matching a vertex on one sketch profile with a vertex on the next profile.

Note: To use vertex matching, the loft profiles must have the same number of vertices.

To carry out vertex matching:

1. In the *Loft* panel, activate the **Match vertices manually** option.

The *Vertex Matching* box expands and displays a list of sketches that are available in the loft, along with the vertices that belong to each sketch.

2. For each sketch, select a vertex from the corresponding list in the *Vertex Matching* box.

   You can find out the name of a particular point by editing the corresponding sketch before creating the loft, and right-clicking on the point in the *3D-CAD View* scene; the point name is displayed at the top of the right-click menu.



The points that you select from each sketch define a loft path. This path is displayed in the *3D-CAD View* scene. You can check that you

have selected the correct vertices by rotating and zooming in on the model. A preview of the loft is displayed in the *3D-CAD View* scene.



3. Click **OK** to create the loft.

The final loft is shown below.



Using Guide Wires in Loft Operations

You can create more complex lofts by using guide wires.

The guide wires constrain the loft along its length and give you greater control on the shape of the loft. You provide guide wires as open sketch profiles along the length of the loft.

Requirements

The following rules apply when using guide wires:

- All loft profiles must have the same number of vertices. For example, a square profile cannot be linked to an elliptical profile using guide wires. However, a square profile can be connected to another profile that contains 4 vertices, as both profiles contain the same number of points.
- The end points of the guide wires must be coincident with the vertices on the start and end sketch profiles for the loft. Guide wires must also pass through the relevant vertices of intermediate profile

sketches. Use the pierce constraint to force the points on one sketch to coincide with the edge of another sketch. See Pierce Constraint.

- Only open sketches can be used as guide wires. Closed sketches are automatically assumed to be loft profiles.
- Guide wires must be formed from single, continuous, sketch entities.
- A single sketch can contain more than one guide wire.

Example

The following example demonstrates how the guide wires constrain the loft. 3D-CAD automatically calculates the shape of the loft based on the shape of the profiles and the distance that the profiles are apart.



By adding guide wires to the loft, you can control its final shape.



Swept Features

A swept feature creates a body by sweeping a profile along a specified path.

The profile and path are defined in separate sketches.



A swept cut is defined in the same way as a sweep, however the feature determines where material is removed from a body. You can only use swept cuts on solid bodies.



The order in which you select the sketches affects how the feature is created. The sketch that you select first defines the sweep profile. The second sketch that you select defines the sweep path. Renaming the sketches according to their purpose makes it easier to choose them correctly.

Requirements

The following criteria must be satisfied for the sweep to complete successfully:

- The profile for sweeping must contain only closed loops.
- For sketch entities containing two closed profiles, one profile must be entirely enclosed by the other. You cannot define more than two closed profiles for a sweep feature.
- Sketch entities must not intersect.
- The path can be either open or closed, but must be defined by a single loop.

- The path may be located anywhere in relation to the profile, however the position of the path affects the creation of the body. For predictable results, it is recommended that the starting point of the path intersects the profile itself.
- The path must not allow the created feature to pass through itself at any point.
- The path must avoid curves with radical changes in direction.

The sketches may contain any number of construction entities.

See Also

Sketches
Working with Construction Entities

Creating a Swept Feature

The Sweep and Sweep Cut features share the same procedure. To create a swept feature:

1. Select the sketch that contains the sweep profile from the **Features** manager.
2. Hold `<Ctrl>` and select the sketch that contains the sweep path.
3. Either:
    - Right-click on either sketch in the feature tree and select **Create Sweep**.



- Click (**Create Sweep**) in the 3D-CAD toolbar.

The *Sweep* panel appears.



4. Enter the sweep parameters. See Sweep Panel.

5. Expose any properties that you want as design parameters by clicking the (**Expose Parameter)** button next the desired feature.
6. Click **OK**.

Sweep and Sweep Cut Panel

Use the following properties to control the behavior of the sweep or sweep cut feature.

Body Type	Defines the type of body that results from the operation.	
	Solid	The resulting body is a closed solid.
	Sheet	The resulting body is an open sheet body.
Twist Angle	Defines the angle that the profile is rotated through as it sweeps along the path. The sweep path defines the axis of twist, which allows the twist to be offset from the center-line of the body. 	
Scale Factor	Defines the enlargement or shrinkage of the profile as the feature sweeps along the path. The position of the sweep path in relation to the profile affects the outcome of the scaling by defining the projection of the body. A scale factor of 1.0 maintains the original scale of the profile through the sweep.	

		
Body Interaction	Specifies how the solid body that is created by the sweep feature interacts with other bodies adjacent to it.	
	None	The new body is added to the **Bodies** manager on its own. No change is made if the body is adjacent to, or overlaps, surrounding bodies.
	Merge	The new body is united with any adjacent or overlapping bodies; otherwise it is added on its own to the **Bodies** manager.
	Imprint	The edges of the extruded body are imprinted on any adjacent or overlapping bodies.

Imprinting Sketches onto Solid Bodies

You can split a surface of a solid body by imprinting a sketch onto it.

You can then use the split surface as a boundary condition. This operation does not apply to sheet bodies.

Requirements

To successfully imprint a sketch onto a surface, the following criteria must be satisfied:

- Sketches to be used in the imprint feature must lie on the same plane as the surface they are imprinting.
- The sketch must only contain closed loops or profiles. Multiple profiles in one sketch are permitted.

- You can only imprint one sketch on a surface at time.
- If the sketch lies partially on a surface, this part of the sketch is used to split the surface. In the following example, the sketch that is used to imprint the surface contains two closed profiles: a square that lies completely within confines of the surface, and a circle that partially lies on the surface.



Creating The Imprint

In the following example, the sketch of a circle is imprinted onto the top surface of a solid block. The sketch lies on the plane of the top surface.



To create a sketch imprint:

1. Right-click on the desired sketch in the 3D-CAD object tree and select **Sketch Imprint**.



The *Imprint Sketch* panel appears and the selected sketch is highlighted in the *3D-CAD View* scene.



2. Click **OK** to imprint the sketch.

The imprinted sketch splits the original face of the block. The new face that is created by the imprint feature is shown highlighted below.



Creating Sheet Bodies from Sketches

You can create sheet bodies from closed sketches.

The sketch can contain multiple closed loops as shown below. This feature is useful for creating baffles.



To create a sheet body from a sketch:

1. Right-click the desired sketch in the 3D-CAD feature tree and select **Create Sheets from Sketch**.



The *Convert Sketch to Sheet* panel appears. The selected sketch is listed in the panel.

A preview of the operation appears in the *3D-CAD View* scene.



2. Click **OK**.

Working with Bodies

Bodies represent the individual objects that make up the 3D-CAD model.

There are two types of bodies in 3D-CAD: Sheet Bodies and Solid Bodies. All bodies are stored under the **Bodies** manager node for the model, and can be operated on by executing the commands that are shown below. The faces of a body can also be operated on, and are accessed by selecting them in the *3D-CAD View* scene.

You can perform the following actions with a 3D-CAD body:

- Rename
- Delete
- Show/Hide
- Duplicate
- Creating Sheet Bodies from a Solid Body
- Sewing Sheet Bodies Together
- Shelling a Body
- Thickening a Body
- Transform (including Move, Scale, Rotate, and Mirror)
- Boolean (including Unite, Intersect, Subtract, Imprint, and Slice)
- External volume extraction

Some features require you to select multiple bodies: a source and a target. The **Select From Scene** tool in 3D-CAD is designed for this purpose. See The Select From Scene Tool.

You can import individual, or groups of bodies into the main simulation. You can also export bodies from 3D-CAD as a neutral CAD format. The following options are available from the main simulation tree as well as 3D-CAD:

- Exporting Geometry from 3D-CAD
- Creating Geometry Parts from 3D-CAD

See Also

Faces and Edges in 3D-CAD
Working with Faces and Edges

Overview of Bodies in 3D CAD

In 3D-CAD, a body is a single object, formed from a series of operations and features.

You can have two types of bodies: Solid or Sheet (zero-thickness). Bodies that you create and import in 3D-CAD are stored under the **Bodies** manager node in the 3D-CAD model tree.



When you begin modeling in 3D-CAD, you create bodies from sketches, or by importing geometry.

See Also

Sketches
Importing Geometry into 3D-CAD

Solid Bodies in 3D CAD

Solid bodies are closed bodies that have a finite volume.

Solid bodies appear in a gray color in the *3D-CAD View* scene, and identified by a icon in the **Bodies** manager node.



The following actions are available for solid bodies:

Sheet Bodies in 3D-CAD

Sheet bodies are composed from surfaces that have zero-thickness.

Sheet bodies appear in a brown color in the *3D-CAD View* scene, and are identified by a icon in the **Bodies** manager node.



The green edges on a sheet body correspond to free edges, and indicate the outer perimeter of the sheet body. See Free Edge. The following actions are available for sheet bodies:

Renaming a Body

The names that are given to bodies in 3D-CAD are preserved when you use them to create geometry parts.

To rename a body:

1. In the **Bodies** manager node, right-click on the relevant body node and select **Rename**.



2. In the inline editor, enter the new name for the body.
3. To confirm the change, press `<Enter>`.

See Also

Creating and Updating 3D-CAD Parts

Deleting a Body

When you delete a body, it is removed from the 3D-CAD model. In 3D-CAD, there are two methods that you can use to delete an object:

To delete a body:

1. In the **Bodies** manager node, right-click on a body node and select **Delete**. To select multiple bodies, hold down the `<Ctrl>` key



The *Delete Bodies* panel appears. The current number of bodies that are selected is shown.



2. In the *3D-CAD View* scene, hold down the <`Ctrl`> key and select any additional bodies that you want to delete. To select concealed bodies, hide faces on the obstructive parts. See Hiding Faces.

3. Click **OK**.

The selected bodies are deleted from the model, and a single **DeleteBody** node is added to the feature tree.



- Alternatively, use the <Delete> key. See Using the Keyboard to Delete Objects.

Hiding Bodies

The hide feature allows you to hide bodies temporarily from the *3D-CAD View* scene.

You can also hide bodies while performing a CAD operation. This function is useful for gaining a clear view of other CAD components. You can restore hidden bodies again by using the show feature. If you

want to hide a particular face on a body, use the hide face action. See Hiding Faces.

Hiding a Body

To hide a body:

- Right-click the body from within the **Bodies** manager node and select **Hide**. You can select multiple bodies by holding down the <`Ctrl`> key.



- Alternatively, while performing a CAD operation in *Edit* mode, right-click the body in the *3D-CAD View* scene and select **Hide Body**.

In the feature tree, you can identify hidden bodies by the grayed out icon .

Restoring a Body

To restore a body:

- Right-click the body from within the **Bodies** manager node and select **Show**. You can select multiple bodies by holding the <`Ctrl`> key.



- Alternatively, while performing a CAD operation in *Edit* mode, right-click a blank area in the *3D-CAD View* scene and select **Restore Hidden Bodies**.

When you restore a hidden body in the *3D-CAD View* scene, the body icon returns to normal .

Duplicating a Body

When you duplicate a body, an identical copy of the body is created and added to the **Bodies** manager node. The new body is given a default

name, for example `Body 5`, and can be operated on. It is not necessary to duplicate a body before transforming it. In this case, use the copy option in the desired transform panel. See Creating a Copy During the Transform Operation. It is possible to duplicate several bodies at once, as described below.

To duplicate a body:

1. In the **Bodies** manager node, right-click on a body node and select **Duplicate**. You can select multiple bodies by holding down the `<Ctrl>` key.



The *Duplicate Bodies* panel appears. The current number of bodies that are selected is shown in the *Duplicate Bodies* panel.



2. Click **OK**.

The selected body is duplicated and a body node is added to the object tree for each of the selected bodies. A single **DuplicateBody** feature node is also added to the feature tree.



See Also

Transforming a Body

Creating Sheet Bodies from Solid Bodies

You can copy the faces on a solid body to create sheet bodies. You can copy the entire solid body or specific faces. A separate sheet body is created for each face on the original solid body. The original solid body remains unchanged.

To create sheet bodies from solid bodies:

1. Select the bodies or faces that you want to convert, using one of the following methods:
   - Select one or more body nodes in the **3D-CAD Bodies** node.
   - Select one or more faces in the *3D-CAD View* scene.
2. Right-click one of the selected items and choose **Create Sheet Bodies** from the context menu.

   In the 3D-CAD View scene, the context menu provides the Create Sheet Bodies action in both the Faces and the Bodies sections. If you choose the option from the Faces section, the selected faces are converted to sheet bodies. If you choose the option from the Bodies section, all faces on all selected bodies are converted to sheet bodies.

   The *Create Sheet Bodies* dialog appears.



3. Click **OK** to complete the feature.

A sheet body is created from each face in the solid model and placed in the **Bodies** manager node. These sheet bodies appear in a different color in the 3D-CAD view scene. A **CreateSheetBodies** feature node is also added to the feature tree.



Create Sheet Bodies Panel

Use the Create Sheet Bodies panel to define what STAR-CCM+ includes in the feature.

Select Bodies	Filters your selections to only select entire bodies.
Select Faces	Filters your selections to only select faces.

Sewing Sheet Bodies Together

When the edge of a sheet body coincides with the edge of an adjacent sheet body, STAR-CCM+ can sew the edges together to form a single sheet body.

If the sheet bodies form a closed geometry, you have the option to convert the sheet bodies into a solid body.

Only bodies whose edges are within the specified tolerance can be sewn together. The maximum tolerance for sewing is 1.0E-5 m. To fill gaps larger than this maximum tolerance, use the Fill Surface feature instead. See Filling Holes.

To sew sheet bodies:

1. Select the sheet bodies to sew, using one of the following methods:
   - Select one or more sheet bodies in the *3D-CAD View* scene.
   - Select one or more sheet body nodes in the 3D-CAD **Bodies** node.
2. Right-click on a selection and choose **Sew Sheet Bodies**.

   The *Sew Sheet Bodies* panel appears.



3. Define the sew parameters. See Sew Sheet Bodies Panel.
4. Click **OK**.

See Also

Sheet Bodies in 3D-CAD

Sew Sheet Bodies Panel

The Sew Sheet Bodies action merges sheet bodies whose edges coincide. For a set of sheet bodies that form a closed volume, you can choose to convert them to a solid body.

Sewing Tolerance	Defines the maximum distance that can exist between the edges of adjacent sheet bodies. If the sheet bodies have larger gaps between them, increase the tolerance accordingly.
Attempt To Form Solid	When activated, the Sew Sheet Bodies action attempts to form a solid body from the volume enclosed by the selected sheet bodies.

Shelling a Body

The Shell tool converts a solid body into a thin-walled, hollow, solid body.

You can choose to remove faces and expose the hollow volume of the solid. The shell tool can extrude the original walls inwards or outwards according to your choice.



To shell a solid body:

1. Select the body to shell using one of the following methods:
   - In the 3D-CAD feature tree, right-click the body and select **Shell**.
   - In the graphics window, right-click the body to shell and select **Shell**.

   The *Shell* panel appears.



2. Define the shell parameters. See Shell Panel.

3. If you want to remove any faces, use the **Select From Scene** tool. See The Select From Scene Tool.

   A preview of the operation is displayed in the graphics window.



4. Click **OK**.

The solid body is converted to a thin-walled structure and a**Shell** feature is added to the feature tree.

Shell Panel

The Shell action hollows out a solid body. You specify the wall thickness of the shell feature.

Select a Body	Specifies the solid body to shell. You can shell only one solid body per operation.
Thickness	Defines the thickness of the solid walls.
Shell Outward	When activated, the Shell action uses the existing surface of the solid as the internal surface of the shell. Otherwise, the tool creates a new internal surface that is offset inwards from the existing surface of the solid.
Select Faces to Remove	Removes the selected faces from the solid body and exposes the hollow volume.

Thickening a Body

You can thicken a sheet body by a specified amount.

This operation converts the sheet body into a solid body. You can thicken multiple bodies in a single operation.



To thicken a sheet body:

1. Right-click a sheet body in the *3D-CAD View* scene and select **Thicken**.

   The *Thicken Sheet Body* panel appears.



2. Set the parameters for the thicken operation. See Thicken Sheet Body Panel.

   A preview of the operation is displayed in the graphics window.



3. Click **OK**.

The sheet body is thickened and converted to a solid body. A **ThickenSheetBodies** feature is also added to the feature tree.

Thicken Sheet Body Panel

The Thicken action adds thickness to a sheet body, converting it into a solid body.

Direction	Specifies the direction to thicken the sheet body.	
	Forward	Thickens the sheet body along the face normal. 
	Backward	Thickens the sheet body against the face normal.

		
	Symmetric	Thickens the sheet body by an equal amount on both sides. The total thickness of the body is the specified thickness. 
Thickness	Specifies the amount to thicken the sheet body.	

Transforming a Body

You can transform bodies by performing the following operations. It is possible to create a copy of the original body during a transform operation.

- Translate
- Scale
- Rotate
- Mirror

A general workflow to follow when transforming a body is provided below. In the following example, a translation transformation is applied.

1. In the **Bodies** manager node, right-click on the desired body and select **Transform** > **[Operation Name]**. To select multiple bodies, hold down the <`Ctrl`> key.



The relevant panel appears for the transformation. The number of selected bodies is displayed here.



2. Define the properties for the transformation. See the relevant sections:
   - Translate Bodies
   - Scale Bodies
   - Rotate Bodies
   - Mirror Bodies
3. Expose any properties as design parameters by clicking the (**Expose Parameter**) button.
4. Choose whether to apply the transformation to the original body, or to a copy of the original. See Creating a Copy During the Transform Operation.

   A preview of the transformation is shown in the *3D-CAD View* scene.



5. Click **OK**.

The transformation is applied to the selected bodies. Each operation creates a feature node in the 3D-CAD feature tree.



Translating a Body

Translating a body moves it from one position to another.

You can translate multiple bodies simultaneously. The *Translate Bodies* panel options are listed below.

Translation Vector	Defines the direction of translation. Enter the [X, Y, Z] components of the translation vector
Coordinate System	Defines the coordinate system for the transformation. If you have specified an alternative coordinate system within the simulation, you can define the transformation with reference to this coordinate system instead of the default, **Laboratory**.

Action	Specifies whether to carry out the transformation on the original body or a copy of the original body. See Creating a Copy During the Transform Operation.

Scaling a Body

Scaling a body proportionally increases or decreases its size with respect to the selected coordinate system.

You can scale multiple bodies at once. When a body is scaled, the scale factor is applied from the origin of the chosen coordinate system, as shown in the following diagram.



The origin of the coordinate system may not always be the same as the origin of the sketch plane. If you are using a coordinate system other than the default, you can check the coordinates of the origin by selecting the relevant node under **Tools > Coordinate Systems** in the main simulation tree. The origin is displayed in the *Properties* window.

Scale Body Panel

The *Scale Bodies* panel options are listed below.

Scale all dimensions uniformly	Specifies whether the selected bodies are scaled uniformly in all directions. If you want to scale the body uniformly in all directions, activate this option. The value that you enter in the first *Scale Factor* text box applies to the remaining text boxes (which are grayed out)

Scale Factor	Defines the amount that the bodies are scaled. Enter the scale factor for each direction: **X**, **Y**, and **Z.**
Coordinate System	Defines the coordinate system for the transformation. If you have specified an alternative coordinate system within the simulation, you can define the transformation with reference to this coordinate system instead of the default, **Laboratory**.
Action	Specifies whether to carry out the transformation on the original body or a copy of the original body. See Creating a Copy During the Transform Operation.

Rotating a Body

Rotating a body turns it around a specified axis.

You can rotate several bodies around the same axis simultaneously. The *Rotate Bodies* panel options are listed below.

Rotate Axis	Defines the axis about which the selected bodies rotate.	
	Axis Direction	Defines the direction vector of the axis of rotation. Enter the `[X, Y, Z]` components for direction vector.
	Axis Position	Defines the origin of the axis of rotation. Enter the `[X, Y, Z]` coordinates of a point that lies on the axis of rotation.
Coordinate System	Defines the coordinate system for the transformation. If you have specified an alternative coordinate system within the simulation, you can define the transformation with reference to this coordinate system instead of the default, **Laboratory**.	
Angle	Specifies the angle of rotation.	
Action	Specifies whether to carry out the transformation on the original body or a copy of the original body. See Creating a Copy During the Transform Operation.	

Mirroring a Body

Mirroring a body reflects it about a specified plane.

You can mirror several bodies simultaneously.

Mirror Plane	Defines the plane about which the selected bodies are mirrored.	
	Normal	Defines the direction vector that is normal to the mirror plane. Enter the `[X, Y, Z]` components of this direction vector.
	Position	Defines a point that lies on the mirror plane. Enter the `[X, Y, Z]` coordinates of this point.
Coordinate System	Defines the coordinate system for the transformation. If you have specified an alternative coordinate system within the simulation, you can define the transformation with reference to this coordinate system instead of the default, **Laboratory**.	
Action	Specifies whether to carry out the transformation on the original body or a copy of the original body. See Creating a Copy During the Transform Operation.	

Creating a Copy During the Transform Operation

The **Copy** option eliminates the need for you to duplicate a body before transforming it. Each transform feature panel contains an **Action** drop-down list that allows you to make a copy of the original body and transform the copy, maintaining the original body.



By default the original body is transformed. An example that demonstrates how the **Copy** option works is shown below.



A copy of the body is being translated `0.5 m` in the X direction.



The result of this operation is shown below. The original body has not moved; the transform operation translated the copy of the original body.



See Also

Duplicating a Body

Performing Boolean Operations

You can perform the following Boolean operations in 3D-CAD:

- Unite
- Intersect
- Subtract
- Imprint
- Slice

A general workflow to follow is provided below. In the following example, two bodies are united.

1. Hold down the <`Ctrl`> key and select the relevant bodies from the **Bodies** manager node.

2. Right-click on one of the selected bodies and select **Boolean** > **[Operation Name]**.



The relevant panel appears for the Boolean operation. The number of selected bodies is displayed here.



3. Define the properties for the boolean operation.
4. Click **OK**.

A feature node is added to the feature tree for each Boolean operation.



Uniting Bodies

The Unite action merges two or more contacting bodies into a single body.

The *Unite Bodies* panel options are listed below.

Merge Type	Defines the type of merge parameter to use in the operation.	
	Precise	Uses the default internal model precision of `1.0E-8 m` as the tolerance. This tolerance setting requires that faces of separate bodies be within 1.0E-8 m of each other. This option takes a minimum amount of computational time and works well for all coincident faces and surfaces.
	Tolerant Edges	Reduces the model precision by using a higher tolerance value that you provide on all the edges in the model. This option causes any further geometry operations to use the lower precision model.

	Tolerant Faces	Finds the coincident faces within the specified tolerance. This option is well-suited for skewed faces without sacrificing the precision on the rest of the model.
	Tolerant Faces and Edges	Provides a combination of the **Tolerant Edges** and the **Tolerant Faces** options. Only use this option if the other options fail.
Tolerance	Defines the tolerance for the operation. This property only becomes active when you choose a tolerant merge type.	

Intersecting Bodies

Bodies that are intersecting share a common volume.

An Intersect operation combines these bodies and leaves only this common volume.

If the bodies do not all share a common volume, the operation deletes all selected bodies. If bodies do not intersect, use the transform operation to move the bodies until they share a common volume. Insert the transformation feature before the **IntersectBodies** feature.

See Also

Transforming a Body
Understanding Where New Features Are Inserted

Subtracting Bodies

In a subtract operation, the volume of one or more bodies is removed from another body.

The body from which material is removed is known as the target, and the bodies that are used to remove the material are known as tools. You can use any number of tool bodies to subtract material from the target body. You can choose whether to delete or whether to keep the tool bodies in this operation.

Each tool body must share some of its volume with the target. If a tool body does not overlap the target, it is deleted from the model during the operation and the target body is unchanged—a **SubtractBodies** feature node is added to the feature tree. If the bodies do not overlap, use the transform operation to move them so that they do overlap.

You can specify the designation of target and tool bodies in the *Subtract Bodies* panel.

Target Body	Specifies the body from which material is removed. The common volume of all other selected bodies is removed from the target body.
Keep Tool Bodies	When activated, retains the tool body after the operation. By default, this option is deactivated and the tool body is consumed in the operation.
Imprint	Imprints the tool body onto the target body after performing the operation. This option is only available when you activate the **Keep Tool Bodies** option.

See Also

Transforming a Body

Imprinting Bodies

The imprint operation is carried out between two or more bodies that are overlapping or have coincident faces.

This operation pastes the lines and curves from each body onto the faces where the two bodies meet or intersect. When you perform an imprint operation, the affected faces are not deleted or merged, but the face is split where the new lines and curves are added to each face.

This operation allows you to manipulate a face for further sketching operations, or to automate the creation of contact data before importing the model into a simulation. If the faces are coincident, imprinting also generates part contact data. STAR-CCM+ uses part contact data to create interfaces between regions.

Requirements

There are two types of imprinting in 3D-CAD: precise imprinting, and tolerant imprinting.

For precise imprinting, the following criteria must be satisfied:

- The bodies must overlap by less than `1.0E-8 m`, or have coincident faces.
- A minimum of two bodies are required for imprinting.

If the selected bodies do not meet the criteria that are mentioned above, the surface of each body remains unchanged after the operation, but an **Imprint** feature node is still added to the feature tree. If bodies are not

in contact, use the transform operation to move the bodies so that they overlap or make contact between faces.

Imprint Bodies Panel

The *Imprint Bodies* panel options are listed below.

Imprint Option	Specifies whether to imprint all the selected bodies together or whether to imprint two groups of bodies together.	
	Imprint Within First Group of Bodies	Imprints all selected bodies together.
	Imprint Between First Group and Second Group of Bodies	Imprints the bodies in the first group with the bodies in the second group. You can have bodies that are common to both groups. When you select this option, two **Select From Scene** tools are activated, one for each group. See The Select From Scene Tool.
Imprint Type	Defines the tolerance method to use in the operation.	
	Precise	Uses the default internal model precision of `1.0E-8 m` as the tolerance. This tolerance setting requires that faces of separate bodies be within 1.0E-8 m of each other. This option takes a minimum amount of computational time and works well for all coincident faces and surfaces.
	Tolerant Edges	Reduces the model precision by using a higher tolerance value that you provide on all the edges in the model. This option causes any further geometry operations to use the lower precision model.
	Tolerant Faces	Finds the coincident faces within the specified tolerance. This option is well-suited for skewed faces without sacrificing the precision on the rest of the model.
	Tolerant Faces and Edges	Provides a combination of the **Tolerant Edges** and the **Tolerant Faces** options. Only use this option if the other options fail.
Tolerance	Specifies the tolerance value for the operation. This property only becomes active when you choose a tolerant imprint type.	

See Also

Transforming a Body
Using Part Contact Data to Create Interfaces

Imprinting Groups of Bodies

In 3D-CAD, you can define and imprint two groups of bodies together rather than imprinting all bodies together. Imprinting all bodies together is computationally expensive, especially if you have hundreds of bodies, and in some cases is unnecessary. For example, if you want to imprint multiple bodies onto a single body, such as in electronics cooling, it is more efficient to perform only the imprint that is required.



An example workflow is provided below:

1. In the *Imprint Bodies* panel, set *Imprint Option* to **Imprint Between First Group and Second Group of Bodies**.

   Two **Select From Scene** tools are activated.

2. To define the first group of bodies, click the top **Select From Scene** tool and select the printed circuit board.

The selection is highlighted in the *3D-CAD View* scene.

3. To define the second group of bodies, click the bottom **Select From Scene** tool and select the desired components.

   There are two recommended methods to select a group of bodies:

   - Using the 3D-CAD feature tree:

     1. Click the **3D-CAD** tab to switch to the 3D-CAD feature tree.
     2. Expand the **Bodies** node and multi-select the bodies that you want to include in the group.



     3. Click the **Edit** tab to return to the *Imprint Bodies* panel.

   - Using the Rubberband Select tool:

     1. In the 3D-CAD View scene, hide the printed circuit board. See Hiding Bodies.
     2. Set the Selection Filter to only select Volumes. See Filtering the Selections.
     3. Use the Rubberband Select tool and drag a box around the bodies that you want to include in the group. See Selecting in a Rubberband.



The selections are highlighted in the *3D-CAD View* scene.

4. Set the imprint options. See Imprint Bodies Panel.
5. Click **OK** to perform the imprint.

Slicing a Body

Use the slicing operation to split an existing body into two or more bodies.

You can slice a body using a sketch plane, a face of another body, or a sheet body. You can slice a body using a single slicing tool, or multiple slicing tools. When using a single tool for the slice operation, you can choose whether to retain both of the resulting bodies, or only one on a particular side of the slice. When using multiple slicing tools, all resulting bodies are retained. The *Slice Bodies* panel options are listed below.

Slice by	Defines the method to use for slicing the selected body.	
	Sketch	Slices the body using an existing sketch plane.

	Plane	
	Face	Slices the body using any face from another body, or a sheet body. 
Extend Face	Extends the selected planar or cylindrical face to slice through the body. This is useful for when the slicing face either does not intersect the target body, or only partially intersects the target body.	

		
Select Side	Specifies which side of the slice is kept after performing the slice operation. This option is greyed out if you select multiple faces or sheet bodies.	
	Both Sides	Retains both sides of the slice. The body is split into two separate bodies. 
	Front Side	Retains the front side of the slice.

		
	Back Side	Retains the back side of the slice. 
Sketch Planes	Specifies the sketch plane to use in the slice operation. All sketch planes, including transform sketch planes, are listed here. This option is only available when slicing by sketch plane.	
Tool Type	Specifies whether the body is sliced using a face and/or a sheet body.	
	Face	Slices using the selected faces. Select the faces to slice the body using the

		Select From Scene tool. See The Select From Scene Tool. 
	Sheet Body	Slices the body using the select sheet bodies. Select the sheet bodies from the scene using the **Select From Scene** tool. See The Select From Scene Tool. 

Slicing a Body Using Multiple Faces or Sheet Bodies

You can slice a body with multiple faces or sheet bodies. This functionality is useful for creating a section from a periodic geometry, such as is commonly found in turbomachinery.



An example workflow is provided below:

1. Create the sheet bodies to act as the slicing tools.



2. In the 3D-CAD feature tree, right-click the body you want to slice and select **Boolean** > **Slice**.

   The *Slice Bodies* panel appears.
3. Set *Slice by* to **Face**. This selection allows you to use faces and sheet bodies as the slicing tool.
4. Set *Tool Type* to **Sheet Body**. If you want to select specific faces from a body, set this option to **Face**. See Tool Type.
5. Use the **Select From Scene** tool to select the slicing tools (the sheet body faces). You can select any number of slicing tools in the same operation. See The Select From Scene Tool.
6. Perform the slice.

The section is created. Hide the other bodies in the 3D-CAD tree to view the section.



The Select From Scene Tool

Some features require you to define multiple bodies, such as a source and a target. For example, when extending a surface up to a body, you select a source (the free edge on the sheet body) and the target (the body that you want to extend up to). The **Select From Scene** tool in 3D-CAD is designed to help you define multiple bodies for these types of operations.

Using the Select From Scene Tool

An example workflow is provided below:

1. In the relevant feature panel, select the option that requires you to select multiple bodies (such as a source and target).



The lower **Select From Scene** button becomes available.



2. Click the lower **Select From Scene** button to activate the selection mechanism.



3. In the *3D-CAD View* scene, select the desired body, face, or sheet body for the operation.



The description next to the **Select From Scene** tool updates to reflect the selection.



Note: You can switch between the source and target selection mechanisms by clicking on the relevant **Select From Scene** button. The selections are stored in memory. Active selections are highlighted in the *3D-CAD View* scene.

4. When done, click **OK**.

Working with Faces and Edges

Edges bound faces. Faces bound bodies. These faces and edges define the surfaces and curves in a geometry part, and the boundaries within a region.

You cannot access faces and edges in the 3D-CAD object tree. You can only select faces and edges in the *3D-CAD View* window. You can create new features from faces and edges, such as sketches and sketch planes, chamfers and fillets, as well as internal volume extractions. You can also label a face or edge with a specific name. A named face becomes its own surface when you convert the model to a geometry part.

This section covers the following topics:

See Also

Overview of Bodies in 3D CAD

Faces and Edges in 3D-CAD

Faces

A face is a bounding surface of a body. Faces are created and updated automatically when you generate or modify a body. You can operate on faces through the *3D-CAD View* window.



Note: Double-clicking a face selects other faces that are connected and tangent to it.

You can give names to faces in 3D-CAD (these then appear in the **Named Faces** folder within the body node). These surfaces then become boundaries when you assign the part to a region. When you convert a 3D-CAD model to geometry parts, the faces become surfaces.

Edges

An edge is the line joining two adjacent faces of a body. An edge is created and updated automatically when you generate or modify a body. You can operate on edges through the *3D-CAD View* window. When you converted a 3D-CAD model to geometry parts, the edges become curves.



Note: Double-clicking an edge selects other edges that are connected and tangent to it.

Free Edge

A free edge is an edge that lies on only one face. Free edges are indicated by a green color in the *3D-CAD View* scene. If a free edge appears between two sheet bodies, this indicates that the sheet bodies are not connected. They can be sewn together if their edges coincide. If free edges form a closed loop, you can created a filled surface. You cannot use free edges with a Fillet or Chamfer operation.



Note: Double-clicking a free edge selects all neighboring free edges that form a closed loop with that edge.

See Also

Working with Bodies
Creating Geometry Parts from 3D-CAD
Sheet Bodies in 3D-CAD
Working with Fillets and Chamfers
Setting Face and Edge Names

Hiding Faces

You can gain access to concealed parts by hiding obstructive parts.

You can use this technique to rename faces that were previously obscured, or to create features from them. You can also hide faces on bodies when you are performing body operations, allowing you to select bodies that are positioned within other bodies. If you want to hide an entire body from the scene, use the hide body action. See Hiding Bodies.

Hiding a Face

To hide a face:

- Right-click the face in the *3D-CAD View* scene and select **Hide**. To select multiple faces, hold down the `<Ctrl>` key.



- Alternatively, while performing a CAD operation in *Edit* mode, right-click the face in the *3D-CAD View* scene and select **Hide Face**. You can only hide one face at a time using this method.

The selected faces are hidden in the scene, exposing the geometry behind.



Restoring a Face

To restore hidden faces in the *3D-CAD View* scene:

- Right-click in the white space around the model and select **Restore Hidden Faces**. It is not possible to restore individual faces; this action restores all hidden faces in the model.



Deleting Faces

The Delete Faces action removes faces that you select from the 3D-CAD model.

You can delete faces from solid bodies or sheet bodies. A **DeleteFaces** feature is also added to the 3D-CAD feature tree. If you want to temporarily hide a face, and not delete it, use the Hide operation instead. See Hiding Faces.

After deleting a face from a solid body, the solid body is converted to a sheet body. If the deletion action results in a disconnected topology, multiple bodies are created. If the deletion action results in a topology with face groups that are connected only by vertices, then the feature creates separate bodies for each such face group. If you select all the faces on a body, the entire body is deleted. The edges that bound the deleted surface become free edges. You can also delete faces from sheet bodies.



To delete a face from a body:

1. In the *3D-CAD View* scene, right-click on the face that you want to delete and select **Delete Faces**.



The *Delete Faces* panel appears. The total number of selected faces is shown in the panel.

2. Click **OK**.

See Also

Sheet Bodies in 3D-CAD

Extending Surfaces

The Extend Surface feature allows you to extend an existing face or free edge. You can only extend faces and free edges on sheet bodies.

To extend a surface:

1. Right-click a surface or free edge and select **Extend Surface**.



The *Extend Surface* dialog appears.



2. Enter the parameters for the extended surface. See Extend Surface Panel.

3. If extending to a surface or body, use the **Select From Scene** tool. See The Select From Scene Tool.

   A preview feature is displayed in the *3D-CAD View* scene.



4. Click **OK**.

Extend Surface Panel

The extend surface operation extends existing surfaces on sheet bodies.

Selection Method	Specifies which part of the sheet body is used for this operation.	
	Faces or Free Edges	Faces: all free edges on the selected faces are extended. 

		Free Edges: the selected free edges are extended. 
	Sheet Body	All free edges of the sheet body are extended. 
Extend Type	Specifies whether the extension follows the curvature of the sheet body from which the free edge belongs or not.	
	Linear	The extended surface is planar and tangent to the existing sheet body.

		
	Same Surface	The extended surface maintains the curvature of the sheet body. 
Extend Shape	Specifies whether the extension follows the free edges of the face or whether it extends the surface normal to the edge.	
	Follow Side	Extension maintains shape of source surface.

		
	Normal to Base	Extension is normal to the selected edges. 
Extension Method	Species how to define the extension.	
	Distance	Extends blindly by the specified amount.

		
	Up to Face	Extends up to a selected face. Select a face using the **Select From Scene** button. See The Select From Scene Tool. 
	Up to Body	Extends up to a selected body. Select a body using the **Select From Scene** button. See The Select From Scene Tool.

		
Distance to Extend	Specifies how far to extend the surface.	

Filling Holes

The Fill Surface feature allows you to create a surface from free edges on the perimeter of a hole in a sheet body.

The resulting surface is automatically sewn into the sheet body containing the hole. If the final sheet body forms a closed volume, STAR-CCM+ converts it to a solid body.

Free edges must:

- belong to a single sheet body. Edges on solid bodies are not supported.
- form a single closed loop. Filling between two concentric circles, for example, is not supported.

You can control the tangency of the filled surface. By default the filled surface is planar.



Alternatively, you can choose for the filled surface to be tangent to adjacent faces.



Creating a Filled Surface

To create a filled surface:

1. In the *3D-CAD View* scene, select all the free edges surrounding the hole.

   These edges must form a closed loop. You can double-click a free edge to select all other free edges in the same loop.
2. Right-click one of the edges and select **Fill Surface**.

   The *Fill Surface* panel appears. The number of selected edges is shown in the panel.



3. Enter the parameters for the filled surface. See Fill Surface Panel.

A preview of the filled surface is displayed in the *3D-CAD View* scene.



4. Click **OK** to complete the feature.

A **Fill Surface** node is added to the 3D-CAD feature tree.

Fill Surface Panel

The fill surface operation creates a surface within free edges that form a closed loop.

The *Fill Surface* panel options are listed below.

Set Tangency at All Edges	Specifies whether to make the resulting filled surface tangential to all adjacent faces.	
Tolerance Value	Controls the overall accuracy of the result. Larger tolerance values can result in simpler geometries.	
Selected Edges	This list contains all the selected free edges. Clicking on each item in the list highlights the relevant edge in the *3D-CAD View* scene.	
	Edge Name	Lists the names of the selected free edges.

	Continuity	When activated for an edge, STAR-CCM+ maintains tangency between the fill surface and the face adjacent to the edge.

Bridge Surfaces

The Bridge Surface action creates a sheet body between two sets of edges.

You can use this feature to repair geometries that have disjointed sheet bodies.



Prerequisites

The following conditions must be satisfied to create a bridge surface:

- Each set of edges must originate from a single body.
- Edges on each set must be connected such that they form a manifold topology.
- Both sets of edges must be either open or closed:

- For open profiles, the number of edges in each set can be different.
- For closed profiles, the number of edges in each set must be the same, except where one profile only has a single edge.

Creating a Bridge Surface

To create a bridged surface:

1. In the 3D-CAD View scene, select all the edges (from both bodies) that you want to use to create the bridge surface.
2. Right-click on a selection and select **Bridge Surface**.

   The *Bridge Surface* panel appears.
3. Modify the selection of source and target edges using the **Select From Scene** tool. See The Select From Scene Tool.
4. Set any additional options in the panel. See Bridge Surface Panel.

   A preview of the operation is displayed in the *3D-CAD View* scene.



5. Click **OK**.

The bridge surface is created and a **Bridge Surface** feature is added to the feature tree.

Bridge Surface Panel

The Bridge Surface action creates a surface between two sets of edges.

Sew	Specifies whether to sew sheet bodies together. If the resulting geometry is closed, a solid geometry is created.
Tolerance Value	Controls the tolerance for sewing sheet bodies together.

Setting Face and Edge Names

When you convert a 3D-CAD model to geometry parts, the faces of the model become part surfaces. By naming a face on the 3D-CAD model, it is defined as a separate surface within the part definition, which makes it easier to specify boundaries later on. All renamed faces and edges are listed in the 3D-CAD feature tree, where you can edit and delete them.

To set the name of a face:

1. Select the face in the *3D-CAD View* window. To select multiple faces, hold down the **<Ctrl>** key.
2. Right-click the selection and select **Rename**.



3. In the inline editor, enter the new name for the face.
4. Press `<Enter>`.

Assigning the same name to multiple faces adds them to the same surface in the part definition. You can give faces the same name in one operation (as described above), or give faces the same name individually.

It is recommended that you rename faces after you complete the 3D-CAD model. Changes that you make to the model while working on it can affect the name settings. In general, a face name is retained when you modify a face, providing that some of the original face still exists. If the face itself is split across multiple bodies, only one segment of the face retains the name, and the other segments of the face have no name.

Setting Edge Names

You can name edges in the same way as faces. To set the name of an edge:

- Select the edges in the *3D-CAD View* window, right-click on one of the selections and select **Rename**.



See Also

Creating Geometry Parts from 3D-CAD

Visualizing Named Faces and Edges

Named faces are grouped under the **Bodies** > **Body** > **Named Faces** node in the 3D-CAD object tree. Named faces in imported CAD geometries (for supported formats only) are maintained and added to the **Named Faces** node at the end of the import process.

In the example below, the named faces of **Main Pipe** are shown.



Clicking the nodes highlights the relevant faces in the *3D-CAD View* scene.



Multiple faces that share a name are stored under one node in the object tree and appear in the *3D-CAD View* scene when you select the node.

Editing Face Names From the 3D-CAD Feature Tree

To rename a face that is listed:

1. In the 3D-CAD feature tree, right-clicking the desired face name node and select **Rename...**



1. Enter the new name in the inline editor.
2. Press `<Enter>`.

Face names that you edit this way are not sorted in alphabetical order. To sort the face names by alphabetical order:

- Right-click on the **Named Faces** manager node and select **Refresh**.

Removing Names From Faces

To remove a name from a face:

- In the 3D-CAD feature tree, right-click on a face name node and select **Delete**.
- Alternatively, edit the face name and enter a blank name in the inline editor.

Visualizing Named Edges

Named edges are treated in the same way as named faces. All named edges appear under the **Bodies** > **Body Name** > **Named Edges** node.



Working with Patterns

The linear pattern and circular pattern features are used to reproduce a single seed (original) body multiple times, within a 3D-CAD model.

The arrangement of the bodies follows a regular pattern, which is defined in the pattern dialog.

Working with Linear Patterns

The linear pattern tool is used to reproduce a series of identical bodies, from a seed body, which follow a regular formation along straight lines.

You can only create a pattern feature from a single body, with the properties defined in the linear pattern dialog. A linear pattern cut is created in a similar way, however the seed body and pattern determine where material is removed.

You define the pattern by one or two linear directions extending from the original body. The number of bodies and the distance between them determine the extent of the pattern.



This section covers:

Creating a Linear Pattern

When creating a linear pattern, you require a seed body. To create a linear pattern:

1. In the **Bodies** manager node, right-click on the seed **Body** and select **Linear Pattern**.



The *Body Linear Pattern* panel appears.



2. Enter the details of the linear pattern. See Body Linear Pattern Panel.
3. Expose any options as design parameters by clicking the corresponding (**Expose Parameter**) button.

   A preview of the pattern is displayed in the *3D-CAD view* scene.



4. Click **OK**.

The pattern is created. Additional bodies are added to the **Bodies** node of the 3D-CAD object tree. Each new body has the name of the seed body plus its position index. For example, if you use a body (named `Tree`) to create a 2 x 2 matrix pattern, the new bodies are called `Tree (1 x 2)`, `Tree (2 x 1)`, and `Tree (2 x 2)`. The generalized naming convention is `Tree (i x j)`, where `i` is the row number and `j` is the column number in the pattern. If you rename the seed body and click the **Update** button, all child bodies adopt the new name plus the position index, as previously described.

You can select each new body independently from the rest of the pattern. Therefore you can apply body operations, such as transform, unite, intersect, subtract, and imprint, to individual bodies.

See Also

Transforming a Body

Body Linear Pattern Panel

A list of parameters that are found in the *Body Linear Pattern* panel is provided in the following table.

Direction (n)	Specifies the parameters for one of the two directions in the linear pattern.	
	Direction (n)	Specifies the direction that pattern extends towards. Enter the [X, Y, Z] components of the direction vector using the global coordinate system.
	Distance	Specifies the gap between each body instance, in meters. To use different units, enter the value with the new units.
	Number of Instances	Specifies how many times the body is reproduced in each direction. The value that you enter must be a whole number. If you enter a decimal value, it is rounded down to a whole number.
Coordinate System	Specifies the coordinate system that is used in the operation. Custom coordinate systems that you specified in the simulation are listed here.	
Skip	Allows you to ignore specified bodies in the feature creation. To skip a body, select the body position from the drop-down menu and activate this option. To skip multiple bodies, select another body from the drop-down menu and activate the option for each body.	
Body Interaction	Specifies how the solid bodies interact with other existing bodies. The bodies must coincide or overlap for these options to work.	
	None	Creates new bodies from the linear pattern only. This option does not affect any existing bodies.
	Merge All Bodies	Merges all linear pattern bodies into the existing body. No new bodies are created from the pattern.
	Imprint All Bodies	Creates new bodies from the pattern. The edges and vertices are imprinted on the affected faces of each body.
	Merge With Selected Bodies	Creates new bodies from the pattern and merges them with existing bodies that you select. Use the *Select From Scene* tool to select the bodies that you want to merge with the new pattern bodies. See The Select From Scene Tool.

	Imprint With Selected Bodies	Creates new bodies from the pattern and imprints them onto existing bodies that you select. The edges and vertices are imprinted on the affected faces of each body. Use the *Select From Scene* tool to select the existing bodies that you want to imprint with the new pattern bodies. See The Select From Scene Tool.

Creating a Linear Pattern Cut

Linear patter cuts behave in the same way as linear patterns, except that the seed body defines a pattern that removes material from existing bodies.

You require at least two bodies to use a linear pattern cut.



To create a linear pattern cut, see Creating a Linear Pattern. An example of a linear pattern cut is shown below.



Working with Circular Patterns

The circular pattern tool is used to reproduce a series of identical bodies, from a seed body, which follow a regular formation around a circle.

The feature can only be created from a single body, with the properties defined in the circular pattern dialog. A circular pattern cut is created in a similar way, however the seed body and pattern determine where material is removed from the model.

The pattern is created around a circle, which is defined by the position and direction of a central axis. The extent of the pattern is determined by the number of bodies and the angular increment between them, or the full circumference of the circle.



This section covers:

Creating a Circular Pattern

When creating a circular pattern, you require a seed body. To create a circular pattern:

1. In the **Bodies** manager node, right-click on the seed **Body** and select **Circular Pattern**.



The *Body Circular Pattern* panel appears.



2. Enter the circular pattern parameters. See Body Circular Pattern Panel.
3. Expose any options as design parameters by clicking the corresponding (**Expose Parameter**) button.

   A preview of the operation is displayed in the 3D-CAD View scene.



4. Click **OK**.

The pattern is created. Additional bodies are added to the **Bodies** node of the 3D-CAD object tree. Each new body has the name of the seed body plus the instance number. For example, if you use a body (named `Tree`) to create a circular pattern with 4 instances, the new bodies are called `Tree (n)`, where `n` is the instance number of the body. If you rename the seed body and click the **Update** button, all child bodies adopt the new name plus the instance number, as previously described.

You can select each new body independently from the rest of the pattern. Therefore you can apply body operations, such as transform, unite, intersect, subtract, and imprint, to individual bodies.

See Also

Transforming a Body

Body Circular Pattern Panel

A list of parameters that are found in the *Body Circular Pattern* panel is provided in the following table.

Revolve Axis	Determines the axis about which the circular pattern is created.	
	Axis Direction	Specifies the direction of the axis of revolution. Enter the `[X, Y, Z]` components of the direction vector using the global coordinate system.
	Axis Position	Specifies the origin of the axis of revolution. Enter the `[X, Y, Z]` coordinates of the origin.
	Coordinate System	Specifies the coordinate system that is used in the operation. Custom coordinate systems that you specified in the simulation are listed here.
Number of Instances	Specifies how many the seed body is reproduced around the circle. The value that you enter must be a whole number. If you enter a decimal value, it is rounded down to a whole number.	
Evenly Distributed	Activate this option to space the bodies in the pattern apart equally.	
Angular Increment	Specifies the angle difference between each body in the feature, relative to the axis of revolution and the seed body.	
Skip	Allows you to ignore specified bodies in the feature creation. To skip a body, select the body position from the drop-down menu and activate this option. To skip multiple bodies, select another body from the drop-down menu and activate the option for each body.	
Body Interaction	Specifies how the solid bodies interact with other existing bodies. The bodies must coincide or overlap for these options to work.	
	None	Creates new bodies from the circular pattern only. This does not affect any existing bodies.
	Merge All Bodies	Merges all circular pattern bodies into the existing body. No new bodies are created from the pattern.
	Imprint All Bodies	Creates new bodies from the pattern. The edges and vertices are imprinted onto the affected faces of each body.

	Merge With Selected Bodies	Creates new bodies from the pattern and merges them with existing bodies that you select. Use the *Select From Scene* tool to select the bodies that you want to merge with the new pattern bodies. See The Select From Scene Tool.
	Imprint With Selected Bodies	Creates new bodies from the pattern and imprints them onto existing bodies that you select. The edges and vertices are imprinted on the affected faces of each body. Use the *Select From Scene* tool to select the existing bodies that you want to imprint with the new pattern bodies. See The Select From Scene Tool.

Creating a Circular Pattern Cut

Circular pattern cuts behave in the same way as circular patterns, except that the seed body defines a pattern that removes material from existing bodies.

You require at least two bodies to use a circular pattern cut.



To create a circular pattern cut, see Creating a Circular Pattern. An example of a circular pattern cut is shown below.



Working with Fillets and Chamfers

Fillets and chamfers are used to blend two faces together, either by creating a smooth curve between them, or by cutting off the corner to leave a flat surface.

These features are applied to edges and faces, which can only be selected through the *3D-CAD View* window. These features are described in the following sections:

What Is a Fillet?

A fillet is a feature that rounds-off the edge, or corner, of a body by replacing it with a smooth curve of constant radius.

Fillets can be applied to interior and exterior corners.



Applying a fillet successfully depends on what edges or faces you select and the size of the fillet in relation to the model. If you are having difficulty in applying a fillet, try decreasing the size of the radius or selecting fewer edges/faces. In some cases, a fillet that is applied to a particular edge can extend automatically to include adjacent edges.

Applying a Fillet Feature

To apply a fillet:

1. Either:

   - Right-click on an edge or face in the *3D-CAD View* window and select **Create Fillet**.



   - Click (**Create Fillet**) in the 3D-CAD toolbar.

   The *Fillet* panel appears.



2. Select any additional faces and edges for the fillet. To make the edge or face selection process easier, use the Selection Filter. You can select multiple edges or faces at once by holding down the `<Ctrl>` key.
3. In the **Radius** textbox, enter the radius of the fillet. If you selected a face, the fillet is applied to all edges from the face.
4. If necessary, expose the radius as a design parameter by clicking the (**Expose Parameter**) button.

   The number of edges and faces that are selected is shown next to the Number of Selected Edges and Number of Selected Faces fields respectively. A preview of the fillet is shown in the *3D-CAD View* scene. The fillet is shown in yellow, while the edges and faces that are used to create the fillet are shown in purple.



5. Click **OK** to apply the fillet to the selected edges.

A **Fillet** node is added to the feature tree.



To change which edges or faces the feature is applied to, or to update the size of the fillet, right-click on the **Fillet** node and select **Edit...**

See Also

What Is a Design Parameter?

What Is a Chamfer?

A chamfer is a feature that makes a straight cut across the corner of a body, resulting in a flat surface.

The size of the chamfer is defined as two values which are the distances from the selected edge into each of the faces. Chamfers can be applied to interior and exterior corners. When applied to interior corners, chamfers add material to the body.



When you select an edge, the chamfer tool:

1. Offsets the faces that are adjacent to the selected edge with the specified distances, *Distance 1* and *Distance 2*.
2. Finds the point where the two offset faces intersect.

3. Projects this point back to the original faces, following the path that is normal to the original faces.
4. Creates the chamfer through these projected points.

An illustration of how the chamfer tool works is shown below.



You can also apply a chamfer to the edges of a selected face. The chamfer tool creates the chamfer for each edge following the same principles as described above.

Applying a chamfer successfully depends on which edges or faces you select, and the size of the chamfer in relation to the model. If you are having difficulty in applying a chamfer, try decreasing the dimensions of the chamfer or selecting fewer edges/faces. In order to solve the model geometry, a chamfer that is applied to a particular entity can extend automatically to include adjacent edges.

Applying a Chamfer Feature

To apply a chamfer:

1. Either:

- Right-click on an edge or face in the *3D-CAD View* scene and select **Create Chamfer**.



- Click (**Create Chamfer**) in the 3D-CAD toolbar.

The *Chamfer* panel appears.



2. Select any additional edges and faces in the *3D-CAD View* window. To make the selection of faces or edges more easily, use the Selection Filter. You can select multiple edges/faces at once by holding down the <`Ctrl`> key.
3. In the **Distance** textboxes, enter the chamfer size. These values determine the distance that the chamfer extends from the selected edge, or edges. It is not possible to determine which face each distance is applied to.
4. If necessary, you can expose each distance as design parameter by clicking the (**Expose Parameter**) button.

   The number of edges and faces that have been selected is shown next to the Number of Selected Edges and Number of Selected Faces fields respectively. A preview of the chamfer feature is displayed in the *3D-CAD View* scene, and allows you to visualize the chamfer feature and switch the two distances. The chamfer is shown in yellow, while the edges and faces that are used to create the chamfer are shown in purple.



5. Click **OK** to apply the chamfer to the selected edges.

A **Chamfer** node is added to the feature tree.



To modify the dimensions of the chamfer, right-click on the **Chamfer** node and select **Edit...**

See Also

What Is a Design Parameter?

Working with Design Parameters

Design parameters are used in 3D-CAD to specify characteristic dimensions of a model.

They can be based on a number of different features including a dimension in a sketch, the depth of an extrusion or the angle of a revolved feature. They are stored under the **Design Parameters** manager node for the model, which can also be accessed from the main simulation tree. Since design parameters can be edited from outside 3D-CAD mode, they can be used to control the size of various features in the model.

This section covers the following:

What Is a Design Parameter?

A design parameter is a reference variable that you can use to define model dimensions and feature sizes.

Each design parameter has a name and a value. When you enter this name into a dimension dialog, for example, the dimension is set to the corresponding value of the design parameter. Design parameters are versatile, and you can use them to form larger expressions. Another benefit of using design parameters is that they allow you to change feature values from outside of 3D-CAD.



There are three types of design parameter:

- Angle
- Length
- Generic

You can explicitly create design parameter, or you can expose an existing dimension (Angle and Length design parameters only). Angle design parameters are defined in either degrees (deg) or radians (radian). Length design parameters are defined in one of the standard units of length.

Creating a New Design Parameter

To manually create a design parameter:

- Right-click on the **Design Parameters** node and select **New** followed by one of the available options:



  - **Angle**
  - **Length**
  - **Generic**

Note: Selecting either **Angle** or **Length** allows you to choose which units define the design parameter. Selecting **Generic** lets you enter a dimensionless quantity.

A new design parameter node, of the selected type, is added to the object tree under the **Design Parameters** manager node. Design

parameters that you create in this way have the default value of `0.0`. To set the new value:

1. Select the design parameter node.
2. In the *Properties* window, enter the parameter value in the *Value* property.

Note: You can set the units of an **Angle** or **Length** design parameter by typing them after the quantity, for example `10 cm`. Alternatively, click the (**Custom Editor**) button next to the *Value* property, and select an option from the Units drop-down menu.

Design parameters that are created using this method are referred to as User-defined. You can also create design parameters by exposing a parameter. These types of parameter are referred to as Feature-exported.

Exposing a Parameter

You can expose an existing dimension as a design parameter and rename this parameter.

A corresponding node is placed in the **Design Parameters** manager node in the 3D-CAD feature tree. This node allows you to use the dimension as a variable when defining other model attributes. You can also edit the value of the dimension by changing the value of the associated design parameter. Use this technique if you want to control a model dimension without editing the corresponding feature directly, or if you want to define other variables in terms of the selected dimension.

There are two methods of exposing a design parameter:

- Exposing a dimension that defines a feature
- Exposing a dimension that defines a sketch entity

A design parameter that is created using this method is referred to as Feature-exported. You can also create design parameters explicitly using the **Design Parameters** manager node. These types of design parameters are referred to as User-defined.

When you delete a feature that contains an exposed dimension, the associated design parameter is not deleted. Instead, it is reclassified as User-defined.

Exposing a Parameter from a Feature

To expose a feature dimension as a design parameter:

1. In the corresponding feature dialog, click the (**Expose Parameter**) button.



The *Design Parameter* dialog appears.



2. Activate **Expose Parameter** to confirm that you want to expose this dimension.
3. In the *Name* field, enter the name for the design parameter.
4. Click **OK**.

A design parameter node is added to the 3D-CAD model tree under the **Design Parameters** manager node.



Note: If you cancel the creation of the feature, the design parameter node still exists in the model tree. To reverse this process, conceal the parameter.

Examples of feature dimensions that you can expose as design parameters are:

- The **Distance** of an extrusion or extruded cut feature
- The **Angle** of a revolved or revolved cut feature
- The **Radius** of a fillet feature
- The **Distances** of a chamfer feature

Exposing a Parameter from a Sketch Entity

You can expose a dimension that defines the size of a sketch entity as a design parameter. To expose this dimension:

1. Applying a dimension to a sketch entity, or edit an existing dimension.

   The *Dimension* dialog appears.



2. Activate **Expose Parameter**.
3. In the **Name** field, enter the name of the design parameter.
4. Click **OK** to expose the parameter.

A design parameter node is added to the 3D-CAD model tree under the **Design Parameters** manager node.



Note: If you cancel the creation of the feature, the design parameter node still exists in the model tree. To reverse this process, conceal the parameter.

See Also

Creating Bodies From Sketches
Sketch Entity Reference

Concealing a Design Parameter

To remove design parameters from the 3D-CAD feature tree, conceal them. To conceal a design parameter:

1. Access the *Dimension* dialog or the *Parameter Name* dialog. See Exposing a Parameter.
2. Deactivate **Expose Parameter**.



This action hides the design parameter and removes the design parameter node from the 3D-CAD feature tree. Any model dimensions that you specified using this design parameter are no longer satisfied, and errors may be reported for the affected features. To resolve the related errors, check the value of existing design parameters and dimensions.

You can also delete design parameters.

Using Expressions in a Design Parameter

You can use design parameters to specify the value of model dimensions by typing $ followed by the parameter name into the dimension textbox. This format is the simplest form of expression.



Note: Design parameters are case-sensitive.

You can also use design parameters as variables in longer expressions by using operators `* / - +` and parentheses `( )`. Spaces are acceptable between terms and operators in the expression. If the name of the design parameter, however, contains any spaces, enclose it in `{ }`, for example `${Length 2}`.



3D-CAD assumes that design parameters are in SI units. Dimensions that are defined by expressions are displayed in SI units only. When conflicting types of design parameter are used in the same expression, for example Angle and Length, 3D-CAD interprets the values in the SI units that correspond to the type of dimension.

Changing the Value of a Design Parameter

Changing the value of a design parameter modifies the size of any dependent features. This ability allows you to control the size of model attributes without having to edit specific features. By using this

technique, you can also modify the 3D-CAD model from outside of 3D-CAD.

To change the value of a design parameter:

1. Select the relevant node under the **Design Parameters** manager node.
2. In the *Properties* window, enter the new value for the parameter in the *Value* field. Modified features require an update for the changes to take effect. These features are indicated as an altered feature state.
3. Right-click the 3D-CAD model node and select **Update**.



The affected features are updated with the new parameter value.

See Also

Modifying a 3D-CAD Model
Understanding Feature States

Deleting a Design Parameter

There are two ways that you can delete a design parameter:

- Select the relevant node under the **Design Parameters** manager node and press the <`Delete`> key.
- Alternatively, right-click the relevant node under the **Design Parameters** manager node and select **Delete**.



You can also remove an exposed parameter from the feature tree by concealing it.

If you used the deleted design parameter to define any dimensions in the model, 3D-CAD is unable to solve the model geometry for the affected features. This situation is indicated in the feature tree by a change in the feature states, and the 3D-CAD model node displays a change of state. To remedy this situation, update any dimensions that are defined using the deleted design parameter. Alternatively, create another design parameter and specify the same parameter name as the deleted design parameter.

See Also

Understanding Feature States

Working with Volume Extraction

3D-CAD can automatically extract the volume of space surrounding a model or volume of space inside a model, and use it to form a new, separate solid body.

This process is often used for creating geometry that represents a fluid domain, which you can then import into a simulation to analyze internal or external flows. This chapter describes the methods that are used to extract the internal or external volume of a model:

What Is Volume Extraction?

When you extract a volume, 3D-CAD finds a continuous surface within, or around, a solid model and uses it to create a new body.

The body is extended up to arbitrary limits, which are determined by whether the extraction is internal or external.

In internal volume extraction, the process can be thought of as capping all inflow and outflow openings in the model, and filling the internal void with a new solid body. If there are any other holes in the geometry, material can leak out, so the geometry must be closed, apart from any inlets and outlets. You identify the location of these holes during the feature creation process, and 3D-CAD caps automatically them. The model must meet certain criteria for the volume extraction process to be successful.

In external volume extraction, a bounding body is placed around the model, and the volume between the outer surface of the model geometry. The extent of the bounding body is used to create a new solid body. In this process, any holes, hollows, and crevices in the model geometry are filled with the new solid body, in the same way that they are filled by a surrounding volume of gas or liquid. The bounding body can be created around a number of separate bodies, can intersect bodies, and can be any shape. This process is useful for creating a fluid domain for the analysis of external flows.

Internal Volume Extraction

In internal volume extraction, 3D-CAD caps holes in the model geometry, and fills the void inside the model with a new solid body.

The holes are often entry and exit points for a fluid, such as the start and end of a pipe. You identify these openings by selecting faces and/or edges. When you select a face, 3D-CAD analyzes it for the presence of any holes. Selecting edges specifies the exact location of a hole. If a solid body can be created by capping some or all of the holes that are found, the extraction process is successful. The extraction process fails if material can leak out of the model through any other holes.

The internal volume can span multiple bodies. During the operation, all bodies are united. When you create a solid body using internal volume extraction, an imprint operation is carried out automatically so that the edges of the new body are imprinted on existing bodies, and vice versa. The internal volume extraction tool recognizes any bodies that intersect or are located within the internal volume, and molds the new solid body around them.

Requirements

The 3D-CAD model must meet the following criteria in order to carry out internal volume extraction:

- The edge loop of each hole must lie on a single face, or cross a maximum of two connected faces. If the loop crosses three or more faces, 3D-CAD does not recognize the hole.
- The edges of each hole must lie of the same plane.
- The faces or edges that define the holes in the model are selected manually, and there can be any number of holes in the geometry.
- If the internal volume spans multiple bodies, these bodies must be able to unite successfully.

Extracting the Internal Volume

To carry out internal volume extraction:

1. Rotate the model in the *3D-CAD View* scene so that a face containing a hole is visible.
2. Right-click on the face, or the edge of the hole, and select **Extract Internal Volume**.



The *Extract Internal Volume* panel appears.



Here, you can see how many edges, faces, and bodies are currently selected.

3. In the *3D-CAD View* scene, hold down the <`Ctrl`> key and continue to select faces and edges to indicate the position of each hole. You can select any combination of edges and faces. To specify the location of a hole that crosses two connected faces, select the edges.

4. If the internal volume spans multiple bodies, use the **Select From Scene** tool to select all bodies that are located between the outermost bodies of the volume.

   3D-CAD attempts to fill the internal space with a new solid body, placing caps over any holes that are identified during selection. This action results in the creation of a new body, which is previewed in yellow in the *3D-CAD view* scene.



5. Click **OK** to proceed with the volume extraction.

An **Extract Internal** feature node is added to the feature tree, and the new body is added to the **Bodies** manager node. To visualize the new internal body, hide the original bodies so that you are only left with the new body.



To edit the edges and faces that define the feature:

- Right-click on the **Extract Internal** node and select **Edit**.

Troubleshooting

If you experience difficulty in extracting the internal volume of a model, try the following steps:

- Check that you selected all holes in the model geometry, as described above.
- Modify the geometry so that holes are on planar faces.
- Alter the model so that each hole is on a single face.
- If the internal volume spans multiple bodies, make sure that all bodies are selected and that it is possible to unite all the bodies together.
- If the geometry contains multiple holes on the same face and you cannot extract the desired hole, select the edge of the hole that you want to extract, rather than the face.

External Volume Extraction

External aerodynamics simulations require a fluid domain surrounding the object of interest. You can create this fluid domain as a solid body in 3D-CAD and extract the required external volume, leaving a cavity of the geometry that you want to analyze.

When the model itself contains a fully enclosed internal void, the external volume extraction feature also creates a solid body within this void.

Extracting the External Volume

There are two steps to extracting the external volume. The first step is to create a separate body around the model which defines the extents of the fluid domain. This body can be any size and shape, and can intersect the model geometry. Typically you define this body by creating a sketch, then extruding or revolving this sketch. Other methods can also be used such as duplicating an existing body, or importing a body.

To create the fluid domain body in 3D-CAD:

1. Create a sketch. See Creating a Sketch.
2. Right-click on the **Sketch** feature node and select **Create Extrude** or **Create Revolve**, depending on the type of fluid domain you want to create.
3. In the feature creation panel for the extruded or revolved feature, make sure that **Body Interaction** is set to **None**. You do not want the new body to merge with the existing bodies.



In the example below, the fluid domain surrounding a truck and trailer is extracted. The size of the fluid domain is purposely small for this example. As the case is symmetric, only half of the truck (and fluid domain) are required to model the case accurately .



This new fluid domain body is added to the feature tree. It is recommended you rename bodies to keep track of them.

To extract the external volume:

1. Right-click on the fluid domain body under the **Bodies** manager node and select **Extract External Volume**.



The *Extract External Volume* panel appears.



2. Set the extraction properties. See Extract External Volume Panel.

Note: You can modify the external body selection using the **Select From Scene** tool. See The Select From Scene Tool.

3. Click **OK** to perform the volume extraction operation.

3D-CAD creates a duplicate of the fluid domain body, unites the other bodies (the truck and trailer), and subtracts these bodies from the duplicate of the fluid domain. The duplicate body ends up with a cavity where the truck and trailer are positioned.

A new **Extract External** feature node is added to the feature tree. If the body that was used in the extraction process is no longer required, you can delete it.



In the following *3D-CAD View* scene, the original fluid domain, truck, and trailer are hidden. Scene transparency is also activated.



See Also

Duplicating a Body
Importing Geometry into 3D-CAD

Extract External Volume Panel

Use the following properties to define the external volume extraction.

Select External Body	Click to select the external body with the **Select From Scene** tool. See The Select From Scene Tool.	
Merge Type	Defines the tolerance method to use in the operation.	
	Precise	Uses the default internal model precision of `1.0E-8 m` as the tolerance. This tolerance setting requires that faces of separate bodies be within 1.0E-8 m of each other. This option takes a minimum amount of computational time and works well for all coincident faces and surfaces.
	Tolerant Edges	Reduces the model precision by using a higher tolerance value that you provide on all the edges in the model. This option causes any further geometry operations to use the lower precision model.
	Tolerant Faces	Finds the coincident faces within the specified tolerance. This option is well-suited for skewed faces without sacrificing the precision on the rest of the model.
	Tolerant Faces and Edges	Provides a combination of the **Tolerant Edges** and the **Tolerant Faces** options. Only use this option if the other options fail.
Tolerance	Specifies the tolerance value for the operation. This property only becomes active when you choose a tolerant merge type.	
Allow non-manifold condition	Temporarily allows non-manifold conditions while performing volume extraction.	
Extraction Bodies	Specifies the extraction method when extracting the external volume.	
	All Inner Bodies	Extracts all inner bodies from the outer body.
	Selected Inner Bodies	Extracts only the selected bodies from the outer body. Select the inner bodies using the

		appropriate **Select From Scene** tool. See The Select From Scene Tool.
Select Inner Bodies	Click to select inner bodies to extract from the external body with the **Select From Scene** tool. See The Select From Scene Tool.	

Removing Unwanted Features

The defeature function in 3D-CAD helps you to remove unwanted features from a geometry.

Unwanted features can include holes, protrusions, chamfers, and fillets. The defeature function can also merge divided surfaces into a single surface.



To defeature a geometry:

1. Right-click anywhere on a body in the *3D-CAD view* scene and select **Defeature**.



The *Defeature* panel appears.

2. Select the faces of the features that you want to remove on the geometry. To select multiple features, hold down the <`Ctrl`> key.



As you make selections, STAR-CCM+ provides a yellow preview of the defeatured geometry.

3. Click **OK** to perform the defeature.

When the defeature operation completes, STAR-CCM+ adds a **Defeature** node to the feature tree.



The defeatured geometry is shown below.


3D-CAD View 1
STAR-CCM+
Z
Y
X

Updating the 3D-CAD Model

When you edit features in the 3D-CAD model, update the model for the changes to take effect.

For example, increasing an extrusion distance by changing the value of a design parameter, requires you to update the model. Updating the model causes 3D-CAD to rebuild the model based on active features. Features that you have rolled back or suppressed are not affected.

When a model needs updating, the **3D-CAD Model** node and affected feature nodes display the **Out-of-Date** feature state icon. To update the model:

- Click the **Update 3D-CAD** button.



- Alternatively, right-click on the 3D-CAD Model node and select **Update**.

You can use the same process to update a model that is showing the **Out-of-Date** icon under the **Geometry > 3D-CAD Models** node in the main simulation tree. This action is necessary if you modified the model by changing the value of a design parameter outside of 3D-CAD.



If you created geometry parts from the 3D-CAD model, update them by right-clicking and selecting **Update**.



See Also

Rolling Back
Suppressing Features
Understanding Feature States
Modifying a 3D-CAD Model
What Is a Design Parameter?

Importing Geometry into 3D-CAD

3D-CAD allows you to import geometry from many different formats into STAR-CCM+ and modify it ahead of running a simulation. This ability facilitates the simulation process by allowing you to defeature a complicated geometry, extract an internal or external volume, and set up boundary names on the CAD model. This chapter includes information on the following topics:

See Also

Removing Unwanted Features
Working with Volume Extraction
Setting Face and Edge Names

Importing Geometry into a 3D-CAD Model

3D-CAD supports the following CAD file formats:

- .x_t/.x_b - Parasolid Transmit files
- .iges/.igs - IGES files
- .step/.stp - STEP files
- .vda - VDA files
- .prt/.asm - Pro/ENGINEER files
- .model/.exp/.session - CATIA V4 files
- .catpart/.catproduct - CATIA V5 files
- .sldprt/.sldasm - SolidWorks files
- .prt - Unigraphics NX files
- .3dm - Rhinoceros 3D Model files

You can import solid and sheet (zero thickness) bodies into 3D-CAD. During the import process, you are presented with the option to import sheet bodies.

Importing Geometry into 3D-CAD

An active 3D-CAD session is required to import geometry into 3D-CAD. If no session is active, you can create one, or edit an existing 3D-CAD model.

1. Right-click on the **3D-CAD Model** node at the top of the model tree and select **Import > CAD Model**



The *Open* dialog appears.



2. Select the CAD file to import.
3. If you want to import sheet bodies into 3D-CAD, activate the *Import Sheet Bodies* option. Deactivate this option if you want to ignore sheet bodies and only import solid bodies.
4. If you want to check the validity of the geometry, activate the *Check Body Validity* option. See Checking Geometry Validity.
5. Click **Open** to begin the import process.

The geometry is imported into the 3D-CAD model and one or more body nodes are created under the **Bodies** manager node. The import process also adds an **ImportCad** node to the 3D-CAD feature tree.



To undo the import:

- Delete the **ImportCad** feature.

Once the geometry is imported, you can modify it using the tools available in 3D-CAD.

See Also

What Is a 3D-CAD Model

Modifying Imported Geometry

You can modify geometries that you import into 3D-CAD in the following ways:

- Adding features to the model by creating more geometries.
- Extracting an internal or external volume in order to analyze the internal or external flow.

- Defeaturing the geometry by deleting bodies, and using extruded and revolved cuts.
- Imprinting a profile on the surface using the imprint operation.
- Renaming faces so that they are defined as separate surfaces when the model is used in the simulation.

See Also

What Is Volume Extraction?
Removing Unwanted Features
Imprinting Bodies
Setting Face and Edge Names

Using 3D Curves in 3D-CAD

A 3D curve is a sketch entity that is not constrained to any 2D plane. A simple example of a 3D spiral is shown below.

You can use 3D curves in 3D-CAD to perform the following operations:

- Extrusions: You can use a closed 3D curve to form a solid body, or an open 3D curve to form a sheet body.
- Sweeps: You can use 3D curves both as profiles or as path curves while creating a Sweep feature. For a 3D curve to be used as a profile, it can be closed (to form a solid body), or open (to form a sheet body). For it to be used as a path, it must be an open curve.
- Lofts: You can use two or more 3D curves (or a combination of 3D curves and planar sketches) to create a Loft feature. For a 3D curve to be used as a profile, it must be closed (to form a solid body), or open (to form a sheet body). For it to be used as a guide wire, it must be an open curve. A Pierce constraint can be used to connect profiles to guide wires—profiles and guide curves must intersect. See Pierce Constraint.

You cannot use 3D curves with the Revolve feature, nor can you edit 3D curves in 3D-CAD. Furthermore, composite curves that contain sharp edges are not supported.

An example of a sweep using a 3D curve is shown below; the 3D curve is used as the path for the sweep feature.

Importing 3D Curves from a Spreadsheet

You require an active 3D-CAD session to import a 3D curve into 3D-CAD. If no session is active, you can create one, or edit an existing 3D-CAD model. Spreadsheets containing 3D curves must be in the correct format and the units of the coordinate points must be in meters.

To import 3D curves:

1. Right-click the **3D-CAD Model** node at the top of the model tree and select **Import > 3D Curve**.



   The *Open* dialog appears.

2. Select the `.csv` spreadsheet containing the coordinates of the 3D curve.
3. Click **OK**.

   The *Import 3D Curve Options* dialog appears.



4. Choose one of the following options:
   - To import the curve as a single spline, select **Import as 3D spline**. If you intend to use a 3D curve in a 3D-CAD operation (such as **Extrude**, **Loft**, and **Sweep**), use this option.

- To import the curve as a succession of straight lines, using the coordinates as end-points, select**Import as 3D polyline**. Be aware that the resulting polyline does not define a smooth curve.

5. If you want 3D-CAD to close the imported curve, activate the **Close the curve** option.
6. In the **Coordinate System** drop-down menu, select the coordinate system in which the points are defined. Currently only Cartesian and Cylindrical coordinate systems are supported; the menu lists the systems available when the dialog is opened.
7. Click **OK** to import the curve.

The 3D curve is imported into the 3D-CAD model and a **3DCurve** node is added to the 3D-CAD feature tree.



Importing Multiple 3D Curves from a Single Spreadsheet

You can import multiple 3D curves that are stored in a single file. 3D-CAD looks for identical vertices in the spreadsheet and classifies everything in between these two points as a single 3D curve. As a result, each 3D curve is closed. You can have a blank line separating each 3D curve in the spreadsheet.

In the following example, the 3D curves that make up a propeller blade are imported into 3D-CAD from a single `.csv` file.



The 3D curves are imported as separate features.



You can then use these imported profiles to create a loft of the propeller blade.



3D Curve Spreadsheet Format

The spreadsheet containing the 3D curve coordinate points must be in comma-separated values (`.csv`) format.

Text is not supported in this file, meaning that headers, descriptions, or units must not be included in these files. A maximum of 8 decimal places per coordinate is supported. There is no limit to the number of rows (coordinate lines) per file. The coordinate points must be in meters. If your geometry uses a different system of units, convert the coordinate points into meters using your original CAD package.

An example file in Cartesian coordinates is shown below:

```
0.01817128,0.01377324,-0.01098872
0.04074841,0.02954630,-0.01947049
0.07237303,0.04867008,-0.02909696
0.12597960,0.07179400,-0.03726481
0.16208829,0.08217564,-0.04081605
0.22814133,0.09360534,-0.04507754
0.28460602,0.09557598,-0.04720828
```

Each row represents a point on the curve, and must contain three comma-separated values.

- When using Cartesian coordinates, the format must be [X, Y, Z]
- When using Cylindrical coordinates, the format must be [R, q, Z]

It is possible to import 2D profiles and sketches using this method by setting one of the axes to a constant for all points.

Exporting Geometry from 3D-CAD

3D-CAD allows you to export solid and sheet bodies to a selection of neutral CAD formats. You can choose to export the whole model, or only selected bodies.

To export a 3D-CAD geometry:

1. Choose the bodies to export:

   - If you want to export the entire 3D-CAD model, right-click the root node in the 3D-CAD feature tree and select **Export...**



   - If you want to export selected bodies only, multi-select the bodies within the **Bodies** node, then right-click on a selection and select **Export...**



The *Save* dialog appears.



2. Navigate to the directory where you want to save the file.
3. In the *File Name* field, enter a name for the file.
4. In the *Files of Type* field, select an export file format:
   - Parasolid Text (`.x_t`)
   - Parasolid Binary (`.x_b`)

- Initial Graphics Exchange Specification (`.iges`)[1]
- Standardized Exchange of Product (`.step`)[1]

5. Choose the export options. See 3D-CAD Export Options.
6. Click **Save**.

The file is exported to the chosen location. If you renamed a body during the 3D-CAD session, its name is preserved in the exported file.

1. Exporting to STEP and IGES formats requires a CAD Exchange licence.

3D-CAD Export Options

When exporting a model or body in 3D-CAD, you have a variety of options to choose from, depending on the format you choose.

These options are listed below:

Export Solid Bodies	Specifies whether to export solid bodies.	
	Activated	Includes solid bodies in the export.
	Deactivated	Excludes solid bodies in the export.
Export Sheet Bodies	Specifies whether to export sheet bodies.	
	Activated	Includes sheet bodies in the export.
	Deactivated	Excludes sheet bodies in the export.
Save Solids as Faces (IGES)[1]	Specifies how the solid geometry is saved.	
	Activated	Saves all faces of solid bodies as Trimmed Surface Entities (Type 144). The solid is represented by an Associativity Instance Entity (Type 402).
	Deactivated	Saves the geometry as a Manifold B-Rep Solid (Type 186).
Convert Analytics to NURBS	Specifies whether to convert analytic surfaces (such as cylinders, cones, and planes) to NURBS or not.	

(IGES/ STEP)[1]	**Activated**	Analytic surfaces are converted to NURBS surfaces.
	Deactivated	Analytic surfaces remain the same.
STEP Format[1]	Specifies which version of STEP to export to.	
	AP 203	Saves using Application Protocol 203.
	AP 214	Saves using Application Protocol 214.

1. Exporting to STEP and IGES formats requires a CAD Exchange licence.

Working with 3D-CAD Models

With 3D-CAD, you can create models and perform a full CFD simulation all within the user-friendly environment of STAR-CCM+. The fully integrated modeler makes using 3D-CAD models in a simulation quick and easy, and enhances the process of performing design and optimization studies.

To use a 3D-CAD model in a simulation, the first step is to convert it to geometry parts. This action brings the CAD data into the simulation, where it can be assigned to regions and boundaries.

Following this action, it is possible to modify the 3D-CAD model and run the same simulation for the altered geometry. These processes can be automated using the scripting functionality of STAR-CCM+.

Creating Geometry Parts from 3D-CAD

To use a 3D-CAD model in a simulation, it must first be used to create geometry parts. This process is described below and can be performed on either the whole 3D-CAD model, or on one or more of the model bodies.

To create geometry parts:

1. Exit 3D-CAD.
2. Select the **3D-CAD Model** node or one or more **Body** nodes, right-click, and select **New Geometry Part**.



The *Part Creation Options* dialog appears.



3. Specify the parameters for the new part. See Part Creation Options Dialog and Tessellation Density.
4. Click **OK**.

When geometry parts are created from the whole 3D-CAD model, all the bodies that are contained within the model are shown under the **Geometry** > **Parts** manager node. When particular bodies are selected, only these bodies are shown under the **Parts** node.



If you renamed any of the model faces during the 3D-CAD creation process, these faces are defined as separate surfaces under the **Surfaces** node for the respective body part. In the example below, two faces on **Body 1** were renamed `Inlet` and `Outlet`, respectively. When the body was converted, these faces were used to create separate surfaces. All other faces are included in the **Default** surface definition. To rename faces in the 3D-CAD model, edit the model. When done, update the geometry parts that are affected.



See Also

What Is a Geometry Part?
What Is a 3D-CAD Model
Updating 3D-CAD Parts

Part Creation Options Dialog

When creating parts, the following options are available:

Parts Body Type	This option allows you to specify what type of body is used to create the parts.	
	Solid and Sheet Body	Both solid and sheet bodies are used to create parts.
	Solid Body	Only solid bodies are used to create parts.

	Sheet Body	Only sheet bodies are used to create parts.
Mark Feature Edges	This option allows you to specify which of the mesh edges that lie along CAD edges are marked as feature edges. CAD edges are the edges or curves that belong to the underlying CAD model; mesh edges are the line segments that bound the mesh faces.	
	No CAD edges	Does not mark any mesh edges as feature edges.
	Sharp CAD edges	Marks any mesh edge that lies along a sufficiently sharp CAD edge (based on the *Sharp Edge Angle* value) as a feature edge. For any given CAD edge, either all or none of the mesh edges that lie along the CAD edge are marked as feature edges.
	All CAD edges	Marks all mesh edges that lie along any CAD edges as features edges.
Sharp Edge Angle (deg)	Allows feature curves to be created based on a supplied angle. The default angle is 30 degrees.	
Create Part Contacts from Coincident Faces	Activating this option causes STAR-CCM+ to search for coincident faces in the 3D-CAD geometry. For each pair found, part contact data is created and stored in the *Contacts* property of each part. See Creating Part Contacts From Coincident Faces.	
Coincident Face Tolerance	When the *Create Part Contacts from Coincident Faces* option is activated, part contact data is generated for parts that share coincident faces that are within the *Coincident Face Tolerance* distance from their neighbors. This option is intended to close small, erroneous gaps in imported geometries. Do not change this value specifically to create contacts or interfaces between parts that have large gaps between them. If parts have large gaps between them, move one part closer to the other using a translation operation.	
Tessellation Density	This option determines the level of surface refinement that STAR-CCM+ uses when creating the geometry part, and affects the degree to which surface curvature is preserved in the created part. In	

	general, it is recommended that you start by using one of the coarse options, then review the surface triangulation to see if the geometry part needs to be retessellated with a higher density. See Retessellating a Part. For geometries with high amounts of curvature, and for which a higher surface fidelity is desired, one of the finer options is recommended. The tessellation density option has no effect on planar surfaces—the same triangulation is obtained for these surfaces regardless of the settings used.	
	Very Coarse	Uses the least number of triangles to discretize a curved surface.
	Coarse	Uses fewer triangles to discretize a curved surface.
	Medium	Uses the default number of triangles to discretize a curved surface.
	Fine	Uses more triangles to discretize a curved surface.
	Very Fine	Uses the most number of triangles to discretize a curved surface.
Show Detailed Tessellation Parameters	Activating this option displays detailed tessellation Parameters within the *Part Creation Options* dialog. You can then edit these Detailed Tessellation Parameters as required.	

Modifying a 3D-CAD Model

As the 3D-CAD modeler is fully integrated into the STAR-CCM+ environment, it is quick and easy to modify the 3D-CAD model before and after the simulation is run.

If a solution has been found and you wish to rerun the simulation for a slightly different geometry, you can modify the 3D-CAD model and rerun the simulation. In this case, STAR-CCM+ automatically maps the previous solution onto the new geometry, decreasing the simulation run-time.

You can modify a 3D-CAD model by:

- Editing it explicitly
- Changing the value of a design parameter

The first method activates 3D-CAD and gives you access to the full range of CAD tools to edit the geometry.

The second method allows you to modify the model geometry without activating 3D-CAD. To use this method, you must prepare design parameters in advance. Design parameters appear under the **Design Parameter** node for the 3D-CAD model in the main simulation tree once you close 3D-CAD. By changing the value of a design parameter and updating the model, you can change the associated dimension in the 3D-CAD model.



For the changes to take effect in the rest of the simulation, perform the following steps:

- Update the affected geometry parts
- Regenerate the surface and volume mesh

The region and boundary data is mapped onto the updated geometry and you can proceed with the analysis. It is good practice to check that the parts, regions, and boundaries are all associated with the correct sections of the geometry following a modification to the 3D-CAD model.

When the original 3D-CAD model is used to create geometry parts, a single part surface, named **Default,** represents all unnamed faces.

If you:

1. rename or split the **Default** part surface into new part surfaces,
2. then modify the geometry in 3D-CAD,
3. and update the geometry parts,

the CAD surfaces and boundaries that are associated with each part surface are maintained as much as possible.

When you name surfaces on a 3D-CAD model, then use that 3D-CAD model to create geometry parts, individual part surfaces are created for each of the named 3D-CAD surfaces. If you edit the part in 3D-CAD, then update the geometry parts, the part surfaces, and associated boundaries, are maintained.

See Also

What Is a 3D-CAD Model
What Is a Design Parameter?
Updating the 3D-CAD Model
Updating 3D-CAD Parts
What Are Leaf-Level Parts?

The Turbo Slicing Tool

The turbo slicing tool extracts the primary gas path from the extracted fluid volume of a turbomachinery blade row.

With the inputs that you provide, this tool generates a neutral surface (three-dimensional sheet body) through a single blade that extends from the blade edges to the inlet, outlet, hub, and shroud. You can then use this generated surface to slice the fluid volume and isolate the selected blade for analysis.



When simulating a single blade, you can use a much finer mesh which affords you a more detailed look at the flow near the blade.

Identifying Suitable Geometries for the Turbo Slicing Tool

This section provides guidelines for identifying suitable geometries for use with the Turbo Slice tool.

Typically, you start with the raw CAD geometry of the turbomachinery. Currently, the Turbo Slice tool is most suited for axial geometries, such as axial-flow turbomachinery, propellers, and fans.

Your aim is to extract the fluid domain around a single blade. To achieve this, you must create a fluid domain for the entire geometry, then proceed to focusing on a single blade. Although the Turbo Slice tool is designed to operate on one blade row at a time, you can use the tool multiple times to create slice features for multi-blade row machines.

Creating a Suitable Fluid Domain for a Turbomachinery Analysis

Before using the Turbo Slicing Surface tool, make sure that you have an adequate fluid domain.

In 3D-CAD, there are multiple methods you can use to create the fluid domain, including the Extract Volume feature and Boolean operations. The following example uses the Subtract feature in 3D-CAD.

To create a fluid volume for the turbomachinery geometry:

1. Launch 3D-CAD and import the turbomachinery geometry.
2. Make sure that the geometry is aligned with the default axes.

   If necessary, use the **Transform** operations to reposition the geometry. See Transforming a Body.
3. Create a sketch on one of the primary sketch planes, ideally on the plane that runs through the center of the geometry.

4. Draw the profile of the fluid domain, making sure to align it with the geometry.

   Create the domain so that the domain boundaries coincide with the extents to which you want the neutral surface to extend to.

Note: If your axial geometry does not have a hub, you can mimic one by positioning the domain wall in the place that the hub would occupy. The neutral surface will extend to the axis of rotation.

5. Revolve the sketch 360 degrees about the centerline of the geometry.



6. Set the *Body Interaction* parameter in the *Revolve* panel to **None**.

   By default this option is set to **Merge**.
7. Subtract the original geometry from the fluid domain. See Performing Boolean Operations.

   By default, this operation consumes the original geometry. If you want to retain the original geometry, activate the **Keep Tool Bodies** option in the *Subtract Bodies* panel.
8. If you retained the original geometry, hide it.

The fluid domain is complete.



Generating the Neutral Surface

3D-CAD provides you with the tools to extract the fluid domain around a single blade.

To extract the blade passage fluid volume:

1. Make sure that your geometry has a suitable fluid domain. See Creating a Suitable Fluid Domain for a Turbomachinery Analysis.
2. Place the (**Create Turbo Slicing Surface**) icon on the toolbar. See Customizing a Toolbar.
3. Click (**Create Turbo Slicing Surface**).

   The *Turbo Slicing Surface* panel appears.
4. Use the **Select From Scene Tool** to select the faces that make up a single blade. See The Select From Scene Tool

   To help with the selection, you can hide any obstructing faces. See Hiding Faces.

   In some cases, small patches can cause problems. In these cases, omit the small patches from the blade selection. If patches exist, it is not necessary to fully capture the blade leading and trailing edge faces as the neutral surface is an approximate mid-surface.



5. Select the faces that make up the hub.

   It is not necessary to select all the faces that make up the hub. It is recommended you select the hub faces that envelop the selected blade. Forward and aft faces are not necessary. Do not select faces that are coincident with upstream or downstream boundaries.



6. Select the faces that make up the shroud.

   In cases without a shroud, for example open propellers, the shroud faces are the outer boundary of the extracted fluid domain.

   If you hid any faces on the domain, restore them by right-clicking on a blank area in the scene and selecting **Restore Hidden Faces**. In many cases, the shroud is the domain wall.



7. Select the faces that correspond to the inlet.

   The inlet must be planar.



8. Select the faces that correspond to the outlet.
   The outlet must be planar. It is also good practice for the outlet to be parallel to the inlet.



If the neutral surface can be calculated, a preview appears.



9. Set the turbo slicing surface parameters so that the neutral surface fits through the center of the blade. Also make sure that the areas where the neutral surface meet the inlet and outlet are normal to these boundaries.

   The surface must fully intersect all domain walls so that you can use it to slice the single blade passage from the full fluid domain.

   See Turbo Slicing Surface Panel for a definition of the parameters.



10. Click **OK**.

The neutral surface is created and appears as a sheet body under the **Bodies** node of the feature tree.

The Turbo Slicing Surface Panel

The options to control the neutral surface parameters are found in this panel.

Blade Faces	Specifies the faces that make up the blade through which the neutral surface must pass. The neutral surface is created based on the selected blade surfaces. The best practice is to select all the faces that form a loop around the blade surfaces. If the blades contain cooling features, you can ignore these. If the blade contains small patches, you are advised to omit these from the blade selection, as they can cause problems.
Hub Faces	Specifies the faces that make up the hub. Select the faces that surround the blade of interest.
Shroud Faces	Specifies the faces that make up the shroud. Select the faces that surround the blade of interest.
Inlet faces	Specifies the faces that make up the inlet.

Outlet Faces	Specifies the faces that make up the outlet.
Slicing Surface Parameter: Surface Resolution	Specifies the spanwise (hub to shroud) resolution of the grid that is used to define the neutral surface.
Leading Edge/Inlet Parameter: Transition Shift	Specifies the transition of the neutral surface from the inlet to the leading edge of the selected blade. The neutral surface is constrained to be normal to the inlet plane. This parameter defines the non-dimensional location between the inlet and the blade leading edge where the sheet transitions to satisfy the normal constraint. The valid range is from 0.0 to 1.0, from the leading edge to the inlet. The smaller the value, the sharper the transition.
Trailing Edge/Outlet Parameter: Transition Shift	Specifies the transition of the neutral surface from the outlet to the trailing edge of the selected blade. The neutral surface is constrained to be normal to the outlet plane. This parameter defines the non-dimensional location between the outlet and the trailing edge of the blade where the sheet transitions to satisfy the normal constraint. The valid range is from 0.0 to 1.0, from the trailing edge to the outlet. The smaller the value, the sharper the transition.

Extracting the Fluid Volume Around a Single Blade

After generating the neutral surface, you can use this to extract the fluid volume around a single blade.

Use the generated neutral surface to slice the fluid domain:

1. Rotate the neutral surface so that it is equally spaced between two blades. See Transforming a Body.

   A formula to calculate the exact amount to rotate the neutral surface is provided below:

   $$\frac{(360 / N)}{2}$$

   where N is the total number of blades in your geometry.

   Set *Action* to **Copy + Rotate Bodies**. The default option rotates the body, however you want to keep the original position for when you rotate the neutral surface in the opposite direction.



2. Repeat the previous step and rotate the original neutral surface in the opposite direction by the same amount.

   Leave *Action* set to **Rotate Bodies**.



3. Slice the fluid domain using the two neutral surfaces. See Performing Boolean Operations.

   Set *Slice by:* to **Face** and *Tool Type* to **Sheet Body**.
4. Use the **Select From Scene** tool to select the two neutral surfaces from the scene. See The Select from Scene Tool.
5. Delete the two neutral surfaces as well as original fluid domain - these items are no longer required.

The extracted fluid domain is complete.


STAR-CCM+
Z

Importing Geometry

One of the possible starting points for the meshing process is to import surface data.

There are three different types of surface data that you can import into STAR-CCM+:

- Surface meshes
- Neutral format data
- Native CAD data

Additionally, you can import XML data files that contain assembly data and references to component files based on the above formats.

Note: To import native CAD files, you require an additional CAD Exchange Reader license. To import JTOpen files, you require an additional JTOpen Reader license.

For each of the above three file types, the option to import the data as regions or to import the data as parts is provided during the import process. Progress that is made during the translation/import process can be reviewed and the resulting imported surface can be checked for errors if necessary. If problems are found, they can generally be fixed using the surface repair tools.

If the surface data resulted in regions and boundaries being created then they can be scaled, rotated, translated, or reflected in order to position them properly relative to one another. Surface data that is imported as geometry parts can similarly be manipulated before being assigned to regions.

Additional region-based manipulations are possible in order to make them conform with the requirements for multi-region meshing.

Surface data can be imported to complement existing volume meshes so that only part of the simulation domain requires preparation for volume meshing.

Importing Directly from the CAD Program

You can import CAD geometry into STAR-CCM+ directly from your CAD package by using one of the following supported CAD clients:

- STAR-Works
- STAR-CAT5
- STAR-NX
- STAR-Inventor
- STAR-Pro/E
- STAR-Creo
- STAR-SpaceClaim

Note: To use this feature, you require an additional CAD client license. STAR-SpaceClaim also requires the SpaceClaim Data Exchange II module and license.

What Is a Surface Mesh?

A surface mesh is a discrete representation of the geometry of the individual regions that are used for the volume mesh generation.

It is made up of faces (triangles) and vertices. Additionally, feature curves are included in surfaces that are created in some pre-processing packages. Internal feature surfaces such as baffles or interfaces are allowed.

The ideal surface mesh has the following properties:

- Closed - Contains no free edges or mismatches
- Manifold - Edges are shared between no more than two triangles
- Non-intersecting - The surface does not self-intersect

The quality of triangle faces is also important. The ideal surface contains:

- Equilateral sized triangles
- Gradual variation in triangle size from one to the next
- No sharp angles, surface folds, or other such characteristics within the triangle proximity sphere

Surface meshes traditionally originate from CAD packages directly or through third-party pre-processing/surfacing software, and are saved in specific formats. Consider carefully the method that you use to generate the surface mesh, as it directly influences the quality of the resulting volume mesh and the effort it takes to get to this step. In cases where a bad quality starting surface is the only option available, the surface wrapper and/or surface remesher tools can be used to improve further the surface as defined above. In cases where the import surface contains quads or other elements with more than three sides, the face is automatically triangulated.

What Is Neutral Format Data?

Neutral format data files allow CAD model data to be exchanged between different programs while still maintaining the model integrity.

They contain an exact description of the geometry of the individual regions that are used for the volume mesh generation. Surface meshes, on the other hand, contain a discrete representation. Depending on the format, the neutral data usually contains some kind of boundary representation (B-REP) model which provides a continuous mathematical description of the component, from which a discrete representation can be created. Provided the correct preparation is made, this continuous mathematical description can be used for CAD projection of surface vertices.

What Is Native CAD Data?

Native CAD data files are created with parametric based CAD programs and contain all the relevant solid model information for that program.

The model information can be read for each of the supported file types and converted into individual regions that can be used for surface and volume meshing. An additional reader license, specific to each CAD program format, is required in order to read the native files. Contact your CD-adapco representative for more information on how to obtain the appropriate license.

Provided the correct preparation is made, the underlying CAD model can be used for CAD projection of surface vertices.

What Is XML Data?

XML files are top-level data files that typically describe the assembly structure and components of a product.

It contains references to other files that contain the actual surface data for parts. Currently, two different implementations are supported - a basic version for Medina files and a more advanced PLM-based implementation suitable for other file formats such as JTOpen.

XML data provides a convenient method of importing hundreds of files into STAR-CCM+ using a common set of options for each, while still maintaining the structural hierarchy of components within the original design. This method allows you to import assemblies into STAR-CCM+ and retain the assembly structure in the geometry parts tree. Similarly, you can export data from the parts tree to a PLM-based XML file. This capability is useful for exporting many parts, or for preserving the organization of the parts for import into another simulation. Meta data information, such as materials, are currently ignored.

Importing Surface Data into STAR-CCM+

All surface data (except for XML-based imports) for use in the meshing pipeline is imported into STAR-CCM+ using the **File** > **Import** > **Import Surface Mesh...** menu item:



Alternatively, the (**Import Surface**) toolbar button can be used instead.

The standard file import *Open* dialog panel appears.



To import the file or files, select them and click the **Open** button. The *Import Surface Options* dialog appears. An example is shown below:



Depending on whether a surface file, neutral file or native CAD file format is selected, STAR-CCM+ prompts for various information to assist with the import process and better prepare the starting surface. By default, all surfaces are imported as parts, but the *Import Mode* option can be toggled to allow surfaces to be imported as regions instead if desired.

Additional information pertaining to each file format is provided in the format and import reference. Once the options have been selected, the **Open** button can be pressed to import the indicated file.

When importing native CAD or neutral format data, whenever STAR-CCM+ marks a mesh edge as a feature edge, it records which CAD edge the feature edge lies along. This association of feature edges to CAD edges can be used later. For example, when the surface wrapper or surface remesher splits a feature edge that is known to lie along a

CAD edge, it projects the split point to the CAD edge. (This projection requires you to activate the CAD Projection option.)

The mark feature edges option that is listed above controls this marking of edges. If no mesh edges are marked as feature edges during the import stage, it is not possible to recover the associativity with underlying CAD edges.

Import Mode

The *Import Mode* option determines whether the surface data file or files are imported as regions or as geometry parts.

In general, choose a mode that is based on the nature of the import data and what you want to do with it. For example, if you only have a few bodies and do not want to maintain hierarchal assembly of parts then importing the data as regions is probably more efficient. However, if you have dozens of bodies and want to organize them then in assemblies for future reference then using the parts option would be better. Also, to use closed bodies as volumetric sources, import them using the parts option.

Boundary Mode

When you select the **Create new Region** option for *Import Mode*, the boundary mode determines:

- How different surface patches that have been identified in the surface file are imported into STAR-CCM+
- How these patches are converted into boundaries

A patch is a part of the surface that has been tagged in some way to be different from the rest. For example, it has been assigned a different cell type or family name. Up to three different boundary mode options are presented:

- **One boundary per cell type/property ID/solid** (default) - This option results in a different boundary being created for each cell type, property ID or solid depending on the input file type. For the database and cell/vertex/input formats, the **Combine boundaries by name** option is on by default. This option is also available to allow cell tables with the same name to be automatically combined during import if desired.

- **One boundary per group** - This option results in a different boundary being created for each group. A group specifically refers to cell table grouping that was defined previously in the pro-STAR or pro-surf pre-processing software from CD-adapco.
- **One boundary for all cell types/property ID/solid** - This option results in only one boundary being created regardless of the initial number of cell types/property ID/solid definitions.

In all the above situations, the original patch definition in the import surface can be used later to edit the boundary definitions further with the split boundary by patch option.

Additionally, when importing pro-STAR database or cell/vertex/input files, if the **One boundary per cell type** option is selected for **Boundary Mode** then the **Combine Boundaries by name** toggle is available. Activating this option combines any boundaries that have the same name. Boundaries can have the same name when two or more cell types in pro-STAR have the same name.

STAR-CCM+ automatically assigns boundary types when a mesh is imported.

Region Mode

The region mode determines how the regions are created based on the input data.

The option is displayed when importing surfaces as regions and under the following conditions:

- Two or more surface mesh files are selected.
- One or more neutral format or native CAD data files are selected.

When importing multiple surface mesh files, two different options are provided for the region mode:

- **One region per file** - This option results in one region being created for each separate input file. The region name that is used matches the file from which it came.
- **One region for all files** (default) - This option only creates one region for all the files. In this instance, an additional option to **Combine patches across files** is provided to group any duplicate patch definitions in each file into the same boundary. (This grouping does not involve deleting one of the patches.) The default setting does not combine the patches. The patch combine option is not provided in the case of STL files.

When importing neutral format or native CAD data files, the options change to be as follows:

- **One region per body** - This option results in one region being created for each body that is contained in the input file.
- **One region for all bodies** - This option creates one region for all bodies that are contained in the file.

In addition to the above, if multiple files are selected then the **One region per file** option is also available.

Part Mode

The **Part Mode** option is available when importing surface mesh files and determines the way geometry parts are created based on the underlying surface description.

The option is displayed when importing surfaces as parts and under the following conditions:

- A single STL (`.stl`), database (`.dbs`), or JTOpen (`.jt`) file type is selected.
- Two or more surface mesh files are selected.

Three different options are provided for the part mode depending the above selections:

- **One part** - This option results in only one part being created from the surface file (applicable when a single STL, database, or JTOpen file has been selected).
- **One part per file** - This option results in a new part being created for each input file (applicable when two or more surface mesh files have been selected).
- **One part per body** - This option results in a new part for each underlying body that is contained in surface mesh file or files. The exact definition of “body” is given below for each surface mesh file type.

The definition of the word “body” varies according to the surface mesh file type. The following is a summary of the terms that are used in the import dialog:

Body Name	File Format Applicable To
Cell Type	pro-STAR (`.inp`), database (`.dbs`)

Solids	STL (`.stl`)
Patches	Felisa Front Surface (`.fro`)
Shapes	JTOpen (`.jt`)
Property ID	Nastran (`.nas`), Patran (`.pat`)

For example, importing a surface from a database file includes the **One part per Cell Type** option for **Part Mode**. When you select this option, each cell type in the surface mesh results in a separate part being generated.

The selection that is made for **Part Mode** also influences the state of the Surface Mode option. Typically, the surface mode options are only available when the **One part** or **One part per file** option has been selected for **Part Mode**.

Surface Mode

The **Surface Mode** option is available when importing surface mesh files and determines the way individual part surface nodes are created for a geometry part.

The option is displayed when importing surfaces as geometry parts. Where the **Part Mode** option is also available then the part mode choice influences the state of the surface mode selection. Generally, the surface mode choice is only available when the import results in one part being created for each file.

Depending the surface mesh file type being imported, up to four different options are available:

- **One surface per body** - This option results in a surface node for each body that is contained in the file.
- **One surface per group** - This option results in a surface node for each group. The definition of a group depends on the file type.
- **One surface for all bodies** - This option results in only one surface node being created for the file.
- **One Surface per color** - This option results in a surface node for each color. This option is only available for STL binary files.

The definition of a *body* varies according to the file type and is described in the part mode section.

Additionally, for some surface mesh files, if the **One surface per cell type** option is selected for **Surface Mode** then the **Merge cell types by**

name toggle is available. Activating this option combines any surfaces that have the same name. Surfaces can have the same name when two or more cell types in pro-STAR have the same name.

Units

A choice of centimeters (cm), feet, (ft), inches, (in), kilometers, (km), meters (m) (default), miles (mi), millimeters (mm), micrometers (um), and yards (yd) is allowed.

The selection determines the scale factor that is applied to the import surface to make it conform to SI (meters). For example, selecting millimeters automatically scales the vertices by 0.001. If a unit other than one listed was used for the surface file, select the meters option (equivalent to no scaling). After the import process is complete, you can apply a user-defined scaling factor to scale the mesh.

If a unit other than meters (m) was selected, then the option to **Set preferred units for length** is available (off by default). This option can be used to set the selected unit as the default for all mesh reference and input values that involve length. Leaving the option off maintains meters as the default unit for length.

Merge Vertices

When importing either STL or JTOpen files, this option allows the vertices for each surface element to be automatically merged or not.

The default is to merge the vertices. In addition, the vertex merge tolerance value (default 1.0e-06 m) can be set to accommodate different scale specifications in the imported file.

This option only operates on individual parts. It does not attempt to merge coincident vertices between different parts.

Database ID

When a pro-STAR database file is imported, the database ID option allows specific surface geometries to be selected, when the file contains more than one.

The database ID contains a reference ID number and an optional text summary describing the contents.

Import Type

This option appears when importing geometry from file types that support both CAD and tessellated imports--such as CATIA V5 and SolidWorks formats.

Two options are available:

- **CAD Geometry**- Imports the CAD geometry.
- **Tessellation** - Imports the original surface mesh data from the CAD package. The tessellation density cannot be modified after import. Also note that if this option is selected, any options available on the CAD geometry (such as CAD-based boolean operations and aligned meshing) do not apply.

Mark Feature Edges

This option allows you to specify which of the mesh edges that lie along CAD edges are marked as feature edges.

CAD edges are the edges or curves that belong to the underlying CAD model; mesh edges are the line segments that bound the mesh faces.

This option applies only to those mesh edges that lie along edges of the CAD model - it does not mark as feature edges any mesh edges that do not lie along CAD edges. The available settings are:

- No CAD edges - Do not mark any mesh edges as feature edges.
- Sharp CAD edges - Mark as a feature edge any mesh edge that lies along a sufficiently sharp CAD edge (based on the **Sharp Edge Angle** value, default 30 degrees). For any given CAD edge, either all or none of the mesh edges that lie along it are marked as feature edges.
- All CAD edges - Mark as features edges all mesh edges that lie along any CAD edges.

Marking feature edges in this way activates CAD projection back to the underlying CAD edge when changes are made to the mesh discretization of those edges.

Sharp Edge Angle

Creates feature edges from edges of attached faces that make an angle sharper than the *Sharp Edge Angle* value.

The default *Sharp Edge Angle* is 30 degrees.

Merge Parts by Name

This option is available when the **Import Mode** is set to create new geometry parts.

Any bodies in the imported assembly with the same name are merged into one part automatically when this option is activated. If this option is deactivated, each body becomes its own part in STAR-CCM+ regardless of name.

Create Part Contacts from Coincident Entities

This option is available when the **Import Mode** is set to create new geometry parts.

When activated, the importer searches all of the imported CAD faces for pairs of coincident faces that are within the specified *Coincident Face Tolerance*. For each pair found, part contact data is created and stored in the *Contacts* property of each part. Further information is provided in the section, Creating Part Contacts From Coincident Entities.

Coincident Face Tolerance

When the *Create Part Contacts from Coincident Entities* option is activated, the option to specify the *Coincident Face Tolerance* becomes available.

The *Coincident Face Tolerance* specifies the maximum distance at which coincident faces are imprinted. The default value is 1.0E-5. This option is intended to close small, erroneous gaps in imported geometries. Do not use this option to create contacts or interfaces between parts that have large gaps between them. If parts have large gaps between them, move one part closer to the other using a translation operation.

Check Parts' Validity

The Check Parts' Validity option is available when importing any surface mesh file type. When this option is activated, the validity of the surface mesh import is examined.

A simple message is displayed in the output window which states if the surface mesh is valid. For example: `Body 2 is not valid`. If the surface mesh is not valid, you can check the quality thresholds and fix the validity issues using the surface repair tool. See Repair Surface Mode.

Tessellation Density

The **Tessellation Density** determines the level of surface refinement that STAR-CCM+ uses when importing geometry.

This option affects the degree to which surface curvature is preserved in the imported model, that is, a fine option uses more triangles to discretize a curved surface, while a coarse option uses less. In general, it is recommended you start by importing using one of the coarse options. After importing, review the surface triangulation and, if necessary, retessellate the geometry using a higher density. For geometries with high amounts of curvature, it is recommended to use one of the finer tessellation options if a higher surface remeshing fidelity is desired. The tessellation density option has no effect on planar surfaces - the same triangulation is obtained for these surfaces regardless of the setting used.

Selecting the **User-Defined Tessellation Density** option automatically activates the **Show Detailed Tessellation Parameters** option. The Detailed Tessellation Parameters are then displayed within the *Import Surface Options* dialog. You can edit these parameters as required.



Activating the **Show Detailed Tessellation Parameters** option is also available when each of the other **Tessellation Density** options are selected.

The **Distance-biased** option contains a pre-configured set of tessellation parameters that have shown to produce good results for a wide variety of geometry models.

When selecting the **User-Defined** option, the tessellation parameter values are copied from the default **Medium Tessellation Density** setting. If you select any **Tessellation Density** option other than user-defined before selecting user-defined, the tessellation parameter values displayed are those of the previously selected **Tessellation Density**. This allows you to make adjustments to the tessellation parameter values for each of the **Tessellation Density** options.

For geometry parts of 3D-CAD, CAD and shape types, it is possible to retessellate the parts using a different tessellation density before they are used in the meshing pipeline.

Show Detailed Tessellation Parameters

Activating the **Show Detailed Tessellation Parameters** option displays detailed tessellation parameters within the *Import Surface Options* dialog. You can then edit these parameters as required.

All distance and length tessellation parameters are specified as a percentage of the length of the diagonal of the bounding box in which the parts or bodies are contained.

Curve Chord Tolerance	The maximum allowed distance between an edge and the chords that are used to approximate it. 
Curve Chord Angle	The maximum angle (in degrees) that is allowed between a chord

	and its originating curve. This angle is calculated as the sum of the two angles that are formed between a chord and the curve tangents. 
Surface Plane Tolerance	The maximum allowed distance between a tessellation facet and the surface. All points of the tessellation facet lie on the surface being tessellated. 

Surface Plane Angle	The maximum allowed sum of the two angles (in degrees) between the normal of a facet and two face normals, which are measured at any two positions on the boundary of the facet. 
Minimum Facet Width	Minimum length of any of the sides of the facet. The value specifies a width below which Parasolid disregards supplied curve, surface, and planar tolerance values. You can produce facets with dimensions smaller than this width.
Maximum Facet Width	Maximum length of any of the sides of the facet.
Curve Chord Max	The maximum allowed length of a chord that is used to approximate a portion of an edge.

Sewing Tolerance

For IGES files, it is possible for gaps to exist in the sheet bodies during translation. The sewing tolerance is used to create solid bodies after the translation process.

Gaps that exist between each sheet are sewn together based on the supplied sewing tolerance value (1e-5 m by default). If the sheet bodies are not sewn together properly, then each sheet is treated as a separate

body. If the **Region mode** is set to **One region per body**, multiple regions are produced. Additionally, if the translated surface is used as the starting point for the surface wrapper, then the surface quality at edges and corners where sheets meet is not as good as when the surfaces are closed to form a solid body.

See Also

Multi-Pass Sewing

Multi-Pass Sewing

For IGES files, it is possible for gaps to exist in the sheet bodies during translation.

The Multi-Pass Sewing option, if enabled, sews together gaps that exist between each sheet, based on the sewing tolerance, before and after importing the geometry into STAR-CCM+. If this function is disabled, sewing only takes place once after the geometry is imported. This procedure aims to create solid bodies with improved surface quality at edges and corners where the sheets meet.

See Also

Sewing Tolerance

Open Geometry Scene After Import

A new geometry scene can be created automatically containing the imported surfaces.

This option is on by default.

Run Diagnostics After Import

The surface quality mesh diagnostics report for the imported surface is produced in the *Output* window when this option is turned on for the import of surface meshes.

Elements to Import

For Medina files, this option allows you to select what mesh (surface, volume, or both) to import into STAR-CCM+.

Importing Surfaces Using an XML File

XML-based data files can be used to import many parts with a given hierarchy, such as a CAD assembly.

When this method is used, STAR-CCM+ automatically imports the individual parts that are specified in the XML file and arranges them according to the data in the file. The procedure for XML import is similar to that described above but it uses a different starting point, namely the **Import XML** toolbar button. By default, the button is not available on the interface. You can add it to any suitable toolbar or a new toolbar (as desired).

To add the **Import XML** button to a toolbar:

1. Open the *Customize Toolbar*s dialog by selecting **Window > Toolbars > Customize...** from the main STAR-CCM+ display.

   The *Customize Toolbars* dialog appears.

2. In the dialog, locate the **File > Import XML...** node.



3. Drag and drop the **Import XML...** node to an existing toolbar (such as the Mesh Generation one) or use the **New Toolbar...** button to create a new empty toolbar beforehand.

   The button is now available on the main display.



4. Press the **Import XML...** button.

   The *Open* dialog appears.

5. In the dialog, set the **Files of Type** option to the desired file type (either `.xml` for Medina files or `.plmxml` for other types) and locate the import file.



6. Press **Open** to start the import process.

   The *Import Surface Options* dialog appears:



This dialog provides options to merge vertices (with a vertex merge tolerance), and to set the units for each imported reference file. The merge vertices option is applied to STL and JTOpen format files as they are imported.

7. If appropriate, make the changes that you want.
8. Click **OK**.

Each part that is listed in the XML file is imported one at a time using default options depending on the format. If multiple instances of the same part are listed, then a separate part is created for each. Any part files that cannot be located are listed at the end of the process.

XML (MEDINA format) files are only supported in versions of STAR-CCM+ running on Linux. PLMXML files are supported in both Windows and Linux versions of STAR-CCM+.

Importing Multiple Surface Data Files

STAR-CCM+ allows multiple surface input files of the same format to be imported simultaneously using the **File > Import Surface** option.

Multiple files are selected by pressing the **Ctrl** or **Shift** key at the same time as selecting the file. In addition to the import options provided for single file input, the **Region Mode** option can further influence how regions are created when importing multiple surface data files as regions.

When importing surface mesh files, regardless of the **Region Mode** option that is selected, the imported surface contains two surfaces where one part of the geometry region joins the next. Additional advice on what to do about duplicate surfaces at mesh interfaces is given in the section on multi-region meshing.

Importing a series of surface mesh files one at a time produces the same result as importing the files simultaneously using the **One region for all files** option without merging the patches (or vertices in the case of STL files). This option is effectively the safest import option as no assumptions are made about how the surfaces are joined.

Note the following:

- When you import multiple pro-STAR database files, it is assumed that the surface in each file is saved with the same database ID value.
- It is assumed that the same units have been used for all surface input files.
- Two or more surface data files of differing formats cannot be imported at the same time.

When you import multiple neutral format or native CAD files, regardless of the **Region Mode** option that is selected, the imported data contains two boundary surfaces at common interfaces where one geometry region joins the next. Additional advice on what to do about duplicate surfaces at mesh interfaces is given in the section on multi-region meshing.

Note the following:

- Importing a series of neutral format or native CAD data files one at a time produces the same result as importing the files simultaneously using either the **One region per file** or **One region per body** option for **Region Mode**. This option is the safest import option as no assumptions are made about how the faces are joined.

- Two or more neutral data files of differing formats or two or more native CAD data files of differing formats cannot be imported at the same time.

Reviewing the Surface Data Import Report

A summary of the surface statistics for the imported file is reported in the *Output* window.

A typical example for when surface mesh files are imported is shown below:



When you import surface mesh files, review the total number of faces, vertices, and edges for the surface to make sure that they are what you expect. `Special Edges` refer to line cells, which are used to represent feature curves.

A more extensive report is obtained when running the mesh diagnostics option and reviewing the report for the surface mesh. When you import translated neutral format or native CAD data, the face quality for the translated neutral format surface can be low since only a basic triangulation of the surface is performed.

Using Surface Data Import Representations

When importing surfaces as regions, an **Import** representation is created under the **Representations** node in the simulation tree.



The import node can be used to:

- Export the surface.
- Replace the existing surface with a new one.
- Repair the surface manually.

If you are replacing the existing surface, all existing parameter settings are maintained.

When a surface is imported as a part, a **Geometry** representation is created instead:



The geometry node is simply a place holder in this instance to indicate that a surface has been imported as a part. As such it has no properties.

Checking the Surface Validity

The validity of the surface can be checked using one of two different methods:

- Using the surface preparation options dialog.
- Using thresholds.

Three different derived part threshold scalars can be used to check the quality. The checks are:

- Free edges
- Non-manifold edges
- Pierced Faces

Free edges refer to cell faces that are not connected to a neighboring face. A non-manifold edge refers to a cell face edge that is connected to two (or more) other edges. A pierced face is one that has one or more edges of another surface cell piercing it.

To create the above threshold checks, do the following:

- In the **Input Parts** list, specify the boundary surfaces that you want to check. It is not sufficient to include only the region part.
- Select the appropriate *Scalar* for the desired surface check (either **Free Edges**, **Non-manifold Edges, or Pierced Faces**).
- Set the **Extraction Mode** to **All Above**.
- Set the **threshold value** to 0 and press the **Create** button.

A new derived part is created containing the cell faces next to the problem areas which allows you to investigate the problems further.

If you have any doubts over the validity of the starting surface, then run these checks as a precaution. Fix any problems that are found: the volume meshers cannot proceed until all problems are fixed. Currently, free edges can be fixed by using the hole filler, the edge zipper, or the manual fixing tools, depending the nature of the problem. A non-manifold edge is sometimes valid depending on whether the surface cells causing the problem are supposed to be a part of the geometry or not. For example, consider a geometry that contains a baffle surface that is joined to the exterior surface. This geometry is not a problem as long as the surface is converted to an interface before proceeding with the volume meshing. If the non-manifold edge is not expected, then you can delete the faces that are causing the problem. Hand fixing tools to

remove pierced faces are now supported in STAR-CCM+ and as such can be fixed directly within the software itself.

Understanding Geometry Parts

You can use the geometry parts manager node to setup meshing procedures.

Using the parts-based meshing approach allows a more flexible, re-playable, pipeline of operations, which is particularly useful for any or all of the following scenarios:

- Where you want to mesh multiple parts in the same region with different types of volume mesh.
- Where a hierarchical assembly of CAD assemblies and parts exist that must be maintained within the STAR-CCM+ environment.
- Where you want the control over the ability to repair or swap out individual parts without having to work with the entire surface.
- Where you want to define a template `.sim` file, containing meshing and modeling settings so that geometries can be imported later.
- Where you want to use a non-standard volumetric control shape that a closed volume part defines.
- Where you want to use the 3D-CAD modeler within STAR-CCM+ to design and parameterize a geometry so that an optimal flow solution is obtained.

This section covers how geometry parts are created or imported, how to organize them and how to apply tags and filters to them. Geometry parts can be repaired and manipulated in a similar way to surface meshes. Once parts have been organized and prepared, they are assigned to regions and an initial surface mesh is created by initializing the meshing models.

Although the parts-based meshing approach is recommended, you do not have to use parts to generate a mesh. You can import parts, assign parts to regions, then generate mesh at the regions level. However, if you are working with assemblies that contain dozens or hundreds of

parts, using the parts-based meshing approach is beneficial. Also, if you intend to use the 3D-CAD modeler to design the geometry from scratch or modify an existing CAD model, using geometry parts is a natural part of the process of going from CAD to flow solution.

See Also

Introducing 3D-CAD Modeling
What Are Tags?
What Are Filters
Repairing Part Surfaces and Curves
Manipulating Parts
Assigning Geometry Parts to Regions
Initializing the Mesh
Preparing Surfaces

What Is a Geometry Part?

A geometry part represents an object (leaf level part) or collection of objects (composite part) that can be used as the input geometry for the meshing tools.

These geometry parts can originate from several different sources:

- Leaf Level parts
  - conversion of 3D-CAD bodies
  - imported surfaces
  - creation of shape geometries within STAR-CCM+
  - results of mesh operations
- Composite Parts
  - new empty composite parts
  - from two or more existing parts

The representation of a single part (generated during import and called **Body 1**) in the explorer pane is shown below:



Leaf-level parts are the most fundamental kind of part represents a surface and curve set, which themselves contain faces and edges respectively. A surface can contain one or more sets of faces and a curve

can contain one or more sets of edges. These features are represented below for the **Body 1** part:



A geometry part can also refer to an assembly (composite) of parts or any combination of parts and assemblies. Assemblies can also contain subassemblies, which themselves can contain further subassemblies, and so on, for which the final level is one or more leaf-level parts. An example composite part, with five leaf-level parts, is shown below:



A more complex example, containing leaf-level parts, assemblies, and subassemblies is shown here:



There is no limit to the number of leaf-level and composite parts that are allowed in STAR-CCM+. Within a given composite part, the leaf-level parts must have unique names although these names can be reused in other composite parts.

In order to help identify and sort parts, they can be tagged such that a filter can be applied to them when they must be selected for a given meshing operation. Leaf-level parts can also be rearranged so that they can belong to composite parts and composite parts can be exploded so that they no longer form part of an assembly. This feature makes the organization and managing of parts a lot easier than if they are represented only as regions.

When you examine geometry parts in a scene, you can select them by type using a filter in the Selection toolbar.

Geometry parts provide a more efficient way of organizing data that are eventually used for the meshing tools. Not all parts are required for a given meshing process: in order to represent the domain of interest, specific geometry combinations can be built up instead, by selectively assigning parts to regions.

What Are Leaf-Level Parts?

Leaf-level parts are parts that represent an object, such as a three-dimensional shape, which is formed from a collection of surfaces and curves.

Different operations can be carried out on parts to modify them, for example combining two or more parts. The following types of leaf-level parts are supported in STAR-CCM+:

- 3D-CAD parts that originated from the 3D-CAD modeler
- CAD parts that originate from a CAD import
- Surface mesh parts that originate from non-CAD import formats
- Shape parts that can be created from within the **Parts** node itself.
- Mesh Operation parts that are the result of mesh operations.

What Are 3D CAD Parts?

3D-CAD parts originate from being directly converted from the 3D-CAD modeler.

Light gray part icons, such as this one , represent these parts. They still retain a connection to the originating 3D-CAD body such that if a design parameter is changed then the 3D-CAD part can be updated to match the design change that was made. These parts have properties and an extended right-click menu.



3D-CAD Parts Properties

See Common Properties of Leaf-Level Parts

What Are CAD Parts?

CAD parts originate from native CAD and neutral format CAD data.

Blue part icons, such as this one , represent these parts. Once imported into the geometry part node, they are considered static

topologies although they can be transformed and duplicated if necessary. CAD parts share the common properties for leaf-level parts and are subject to the same set of operations on their right-click menu.

CAD Parts Properties

See Common Properties of Leaf-Level Parts

What Are Shape Parts?

Shape parts are simple shape geometries (block, cone, cylinder, or sphere) that can be created directly within the **Parts** node without requiring use of the 3D-CAD modeler.

These simple parts can be used to provide exterior boundary domains for other parts, or to provide the input shapes for volumetric controls. They could also be used to generate a basic mesh directly.

The blue icons , , and , represent shape parts. Shape parts dimensions can be edited at any time. They expose the common properties for leaf-level parts along with other properties specific to their geometrical definition. Details on creating these parts are provided in the section, Creating Simple Shape Parts.

Block Part

This part is a regular six-sided block that can be created in a chosen coordinate system. Geometric properties are provided on the part node in addition to the common set for leaf-level parts.



Additional Geometric Properties

Corner 1	Coordinates of the first corner of the block.
Corner 2	Coordinates of the diagonally opposite corner of the block.
Coordinate System	Coordinate system in which the corner coordinates are defined.

Cone Part

This part is a cone that is defined by the center coordinates of the start and end faces and their respective radii. It can be created in a chosen coordinate system. Geometric properties are provided on the part node in addition to the common set for leaf-level parts.



Additional Geometric Properties

Start Radius	Radius of the face whose center is the start coordinate.
End Radius	Radius of the face whose center is the end coordinate.
Start Coordinate	Face center coordinate of the start face.
End Coordinate	Face center coordinate of the end face.
Coordinate System	Coordinate system in which the start and end coordinates are defined.

Cylinder Part

This part is a cylinder that is defined by the center coordinates of the start and end faces and a constant radius. It can be created in a chosen coordinate system. Geometric properties are provided on the part node in addition to the common set for leaf-level parts.



Additional Geometric Properties

Radius	The constant radius of the cylinder.
Start Coordinate	Face center coordinate of the start face.
End Coordinate	Face center coordinate of the end face.
Coordinate System	Coordinate system in which the start and end coordinates are defined.

Sphere Part

This part is a sphere that is defined by an origin and radius, and can be created in a chosen coordinate system. Geometric properties are provided on the part node in addition to the common set for leaf-level parts.



Additional Geometric Properties

Origin	Origin of the sphere.
Radius	Radius of the sphere.
Coordinate System	Coordinate system in which the origin is defined.

What Are Mesh Parts?

Mesh parts are parts that originate from surface mesh data such as STL, NASTRAN, and PATRAN.

Green part icons, such as this one , represents these parts. The tessellation density of these parts cannot be modified after import, although these parts can still be modified manually using the surface repair tool.

Any combination of part types is allowed in STAR-CCM+ and each can be assigned to the same or different regions for meshing purposes. A summary example of all the different kinds of leaf-level parts is shown below:



Mesh Parts Properties

See Common Properties of Leaf-Level Parts

What Are Mesh Operation Parts?

Mesh Operation parts are parts that result from executing mesh operations.

The Mesh Operation Part icons are identical to the specific mesh operation icon that they are created from. If you modify the Mesh Operation, re-execute the Mesh Operation to update the Mesh Operation Part.

Mesh Parts Properties

Operation	Specifies the type of mesh operation that is used to create this part.

See Common Properties of Leaf-Level Parts

Common Properties of Leaf Level Parts

When selected, leaf-level parts expose the following properties and expert properties in the *Properties* window.

Common Part Properties

Region	The region to which this part has been assigned (if any). By default, the *Region* property is **None** meaning that the part has not yet been assigned to a region in readiness for meshing.

Expert Properties

Tags	List of tags (if any) that have been assigned to this part. These tags are used with filters in part selection operations. Click **[]** to edit the list directly.
Metadata	Metadata is user-defined data that is stored as key-value pairs. This field can be populated automatically for some import formats and can be used with filters for part selection operations. Click **{}** to edit the list directly.
Index	Part index number.
Contacts	This property is used to specify which other parts it is in contact with. Contacts are always reciprocal and can be converted into interfaces automatically when generating a new region during region assignment.
Descriptions	Provides a list of descriptions that are available for the part. For more information, see Mesh Descriptions.
Face count	The number of faces present on the part surfaces.

What Are Composite Parts?

Geometry parts, once converted, imported, or created within STAR-CCM+, can be reorganized if necessary to provide a more coherent tree structure.

Composite parts allow other parts to be grouped into assemblies and subassemblies for easier organization. Exploding a composite part allows the children of the part to become parent nodes instead.

What Is a Part Surface?

A part surface is a collection of patches on a geometry part.

Patches on a part surface are composed of a collection of mesh faces. A patch can only belong to one part surface, although a part surface can contain more than one patch. However, it is possible to manipulate which part surface a patch belongs to using the split by patch feature.

The **Parts** > **[Part]** > **Surfaces** node is composed of one or more **[Part Surface]** nodes which make up the shape of the leaf-level part geometry. In order to be valid for use with the meshing tools, each part must be composed of at least one part surface.

To see the part surfaces that belong to a part, open the **[Part]** > **Surfaces** node, as shown below:



Each part surface node is represented using the icon. To view the properties of a part surface in the *Properties* window, select a **[Part Surface]** node:



Part Surface Properties

All part surface nodes have the following properties:

- tags specifier
- boundary specifier
- metadata specifier

What Is a Part Curve?

A part curve is a collection of feature edges on a geometry part.

The **Parts > [Part] > Curves** node is composed of one or more **[Part Curve]** nodes which make up the sharp edges and features of the leaf-

level part geometry. The choice of whether the part curve contains feature edges is made during the import process and is arbitrary depending on the requirements of the geometry.

To see the part curves that belong to a part, open the **[Part] > Curves** node, as shown below:



Each part curve node is represented using the icon. The part curves can be manipulated in different ways so that the assignment of feature edges to part curves can be made easier if necessary.

The properties of a part curve can be seen by selecting the **Curves > [Part Curve]** node:



Part Curve Properties

All part curve nodes have the following properties:

- Tags specifier
- Feature curve specifier

Mark Part Curves Options Dialog

Mark Sharp Edges	When activated, edges that occur between two faces that meet at an angle equal to or less than the specified Sharp Edge Angle are marked as part curves.
Sharp Edge Angle	The angle above which all edges are marked as part curves. If you set the angle to 50, then any angle above 50 is marked as a part curve and any angle below 50 is not.
Mark Free Edges	When activated, all free edges are marked as part curves. See What Are Free Edges?

Mark Non-manifold Edges	When activated, all non-manifold edges are marked as part curves. See What Are Non-Manifold Edges?
Mark Patch Perimeters	Marks all edges that surround a patch (a collection of mesh faces on a part or region) as part curves.
Mark Part Surface Perimeters	Marks all edges that surround a part surface (a collection of patches on a geometry part) as part curves.
Mark Boundary Perimeters	Marks all edges that surround a boundary as part curves.
Automatically Fix Errors	Automatically fixes minor errors related to face quality, face proximity, and self-intersection.
Maintain Existing Edges	Edges that are already defined as part curves are maintained.

Creating and Importing Parts

Geometry parts can be created in the following ways

Geometry parts can be created in the following ways:

- Converting 3D-CAD models to parts
- Importing surfaces as parts
- Importing multiple parts using an XML file
- Creating simple shapes within STAR-CCM+

Creating and Updating 3D-CAD Parts

3D-CAD parts are created when a 3D-CAD model is converted to a geometry part.

3D-CAD parts are created when a 3D-CAD model is converted to a geometry part. The 3D-CAD model must first be created using the tools or import options that are provided in the 3D-CAD module, for which full documentation is provided in the section, Using 3D-CAD. If the 3D-CAD model is modified after its corresponding 3D-CAD part has been created, it will be necessary to update the 3D-CAD part to reflect the changes.

Updating 3D-CAD Parts

Update leaf-level parts can be created from a 3D-CAD model.

Update leaf-level parts that are created from a 3D-CAD model when either of the following occurs:

- the originating 3D-CAD model was edited
- the design parameters on which the original 3D-CAD model bodies are based on change for any reason.

In the latter instance, 3D-CAD model parts are flagged with the “out-of-date” badge once the 3D-CAD model bodies have been updated themselves after the design parameter change has been made. Examples of two out-of-date flagged parts are shown below:



In the above example, the two bodies (**Body 1** and **Body 2**) belonging to the **3D-CAD Model 1** node were updated after the design parameters **h1** and **r1** were changed.

In order to obtain the correct geometry, update **Body 1** and **Body 2** in the **Parts** node.

1. This update can be done by selecting the parts node of interest, right-clicking, and selecting **Update** from the pop-up menu:



The geometry part is now updated with the correct dimensions and retessellated. The “out-of-date” warning badge is removed, indicating that the part is now up-to-date.



Multiple part selections can be updated using this method.

The initial surface representation for meshing (if it exists) must be reinitialized after updating 3D-CAD parts. Stale 3D-CAD parts and the initial surface representation (if present) are automatically updated on invoking either of the surface or volume meshing pipelines.

Importing Surfaces as Geometry Parts

The option to create a geometry part is selected by default when importing surface data.

Selecting the **Create New Part** option within the *Import Mode* group box presents the import options relevant to the file type being imported. An example for importing CAD data is shown below:



Depending on the type of surface file, the import process results in one or more leaf-level parts being created. Composite parts (assemblies) are also created when the original file supported and used this option. CAD parts are created when the original files contain neutral format data or native CAD data. Mesh parts are created when the original files contain surface mesh data.

The import parts process reads all CAD parts before tessellation begins. If a surface on a part fails to tessellate (due to invalidities in the CAD model), the tessellation process stops for that part and continues for the next part. All parts are imported into STAR-CCM+, but those parts which failed to tessellate completely do not display correctly.

The Mark Feature Edges and Tessellation Density options are described within the section, Importing Surface Data into STAR-CCM+. When the **Create Part Contacts from Coincident Faces** option is activated, part contact data is generated for parts that share coincident faces that are within the *Coincident Face Tolerance* distance from their neighbors. Do

not use this option to create contacts or interfaces between parts that have large gaps between them. If parts have large gaps between them, move one part closer to the other using a translation operation. If the **Merge Parts by Name** option is selected, any two bodies in the imported assembly with the same name are merged into one part automatically. If the option is deactivated, each body in the assembly becomes its own part in STAR-CCM+ regardless of its name.

To manage the hierarchy of various components to better suit the analysis in hand, geometry parts can be reorganized once they have been imported. Once organized, the meshing models can be initialized to provide a starting initial surface representation that is based on the currently selected geometry parts. In general, all the native CAD file formats plus the Parasolid Transmit format can maintain the same assembly hierarchy as the original CAD program that produced the geometry in the first place. Additionally, the part names are also carried over for these formats and, under certain circumstances, the names that are attached or tagged to faces are also translated. IGES data does not carry part names but STEP data will under certain circumstances. In this respect, the Parasolid Transmit file is the better export option out of all of the neutral file formats in terms of maintaining the organizational data structure.

The import of a geometry part also creates a **Geometry** representation in the **Representation** node. This node has no properties and is simply a place holder for indicating that geometry parts exist.

Creating Simple Shape Parts

Shape parts are normally created at any point after a geometry part has been imported so that the exterior domain can be sized using the imported part as a reference.

They can be added at any time before meshing initialization is performed.

To create a shape part, right-click the **Geometry** > **Parts** node and select **New Shape Part** from the pop-up menu:



Four different shape parts are listed:

- Block
- Cone
- Cylinder
- Sphere

Select the option that corresponds with the shape that you want to create. If a valid scene is open, then the *edit* tab appears in the explorer pane allowing the properties of the shape to be entered, along with other applicable options. The graphical point tool is also displayed allowing you to set the coordinate values graphically if preferred.

Once all the properties are specified, pressing the **Create** button generates the part. You can create additional parts of the same shape by changing the input values and pressing **Create** again. A new leaf-level part with CAD association is added to the **Parts** node for each new shape that is created.

Note: Only simple shape parts that are created using STAR-CCM+ v9.06 or later have CAD association. However, if you use STAR-CCM+ v9.06 to change the parameters of a simple shape part created in STAR-CCM+ v9.04 or earlier, the simple shape part is given CAD association.

The shape part can be re-edited at any time by selecting the shape part node, right-clicking, and selecting **Edit Part in Current Scene**:



Alternatively, the values and options that are used for each shape can be edited directly using the *Properties* window for each node, as described in the section on each part shape. If a simple shape part is modified using surface repair, the CAD association is updated.

If a valid scene is not open when selecting the shape part to create, then a normal dialog appears instead of the *edit* tab. The appropriate properties and selections as described above are prompted for and the shape can be created as normal.

Creating a Block

This section describes how to create a block part.

To create a block part:

1. Rright-click the **Parts** node and select **New Shape Part** > **Block** from the pop-up menu:



The *Create Block Part* window appears in the *edit* tab in the explorer pane:



2. To define the block opposite corners, supply the maximum and minimum X, Y and Z (or equivalent) coordinates in the current **Coordinate System**.

   These coordinates can be entered manually into the window or by using the graphical based point tool, which is automatically displayed when the *edit* tab is opened. The spherical balls in the center of each block face can be dragged using the mouse to extend the block to the appropriate dimensions, as shown below:



3. Right-clicking anywhere on the magenta colored box and moving vertically resizes all the block dimensions.
4. Clicking the box using the middle mouse button translates the box in any direction while maintaining its current size.
5. To generate the part shape, click **Create** once the coordinates have been set. Additional block parts can be created by changing the coordinate values and clicking **Create** again.
6. Click **Close** when all operations are complete.

   For each block created, a new part node appears in the **Parts** folder:



7. Selecting the block part node exposes its properties in the *Properties* window.

By default, block parts are composed of one surface node and one curve node, both of which can be manipulated if necessary.

Creating a Cone

This section describes how to create a cone part.

To create a cone part:

1. Rright-click the **Parts** node and select **New Shape Part > Cone** from the pop-up menu:



The *Create Cone Part* window appears in the *edit* tab in the explorer pane:



Two separate locations (Start Circular Disc and End Circular Disc) and two radii (**Start Radius** and **End Radius**) based on either the global cartesian or a local coordinate system must be defined. These values can be entered manually into the window or by using the graphical based point tool, which is automatically displayed when the *edit* tab is opened. The spherical balls at the center of each disc face can be dragged using the mouse to extend the cone to the appropriate dimensions, as shown below:



2. Once the coordinates and radii have been set, press the **Create** button to generate the part shape at the specified tessellation density. Additional cone parts can be created by changing the coordinate values and pressing **Create** again.
3. Press the **Close** button when all operations are complete.

   For each cone created, a new part node appears in the **Parts** folder:



4. Selecting the cone part node exposes its properties in the *Properties* window.

By default, cone parts are composed of one surface node and one curve node, both of which can be manipulated if necessary.

Creating a Cylinder

This section describes how to create a cylinder part.

To create a cylinder part:

1. Right-click the **Parts** node and select **New Shape Part > Cylinder** from the pop-up menu:



The *Create Cylinder Part* window appears in the *edit* tab in the explorer pane:



Two separate locations (Start Coordinate and End Coordinate) and the **Cylinder Radius** must be supplied based on either the global cartesian or a local coordinate system. These values can be entered manually into the window or using the graphical based point tool, which is automatically displayed when the *edit* tab is opened. The spherical balls at the center of each disc face can be dragged using the mouse to extend the cylinder to the appropriate dimensions, as shown below:



To generate the part shape at the specified tessellation density:

2. Click **Create** once the coordinates and radius have been set. Additional cylinder parts can be created by changing the coordinate values and clicking **Create** again.
3. Click **Close** button when all operations are complete.

For each cylinder created, a new part node appears in the **Parts** folder:



Selecting the cylinder part node exposes its properties in the *Properties* window. By default, cylinder parts are composed of one surface node and one curve node, both of which can be manipulated if necessary.

Creating a Sphere

This section describes how to create a sphere part.

To create a sphere part:

1. Right-click the **Parts** node and select **New Shape Part > Sphere** from the pop-up menu:



The *Create Sphere Part* window appears in the *edit* tab in the explorer pane:



The **Origin** of the sphere and the **Radius** must be defined, based on either the global cartesian or a local coordinate system. The point tool allows the origin to be selected graphically by dragging the ball to the appropriate location, or alternatively, the coordinates can be entered directly into the window. An example of a sphere shape part during the creation process is shown below:



2. Once the coordinates and radius have been set, press the **Create** button to generate the part shape at the specified tessellation density. Additional spherical parts can be created by changing the coordinate values and pressing **Create** again.
3. Press the **Close** button when all operations are complete.

For each sphere created, a new part node appears in the **Parts** folder:



Selecting the sphere part node exposes its properties in the *Properties* window. By default, sphere parts are composed of one surface node and one curve node, both of which can be manipulated if necessary.

Organizing Assemblies Using Composite Parts

Composite parts allow other parts to be grouped into assemblies and subassemblies for easier organization.

Exploding a composite part reverses this process. These operations are described as follows:

- Creating composite parts (assemblies)
- Exploding composite parts

Creating Composite Parts (Assemblies)

In order to help organize leaf-level parts or other composite (subassembly) parts, you can create a composite part.

Two general methods to create composite parts are:

- create a new empty composite part and move existing leaf-level parts or other composite parts into it afterwards
- create a composite part from two or more existing parts

Creating a New Empty Composite Part

To create a new empty composite part (that does not contain any other parts):

1. Right-click the **Parts** node and select **New Composite Part** from the pop-up menu, as shown below:



A new composite part node appears, with the default name **Composite**:



2. Leaf-level or other composite parts can be moved into the new composite part by selecting them and dragging and dropping them into the composite part node.

This results in the parts becoming children of the new composite part, as shown in the arrangement below:



Any number of new composite parts can be created using this method.

Creating Composite Parts from Other Parts

A new composite part can be created directly from two or more existing parts so that the selected parts automatically become children of the new composite part.

For example, to make two leaf-level parts the children of a new composite part:

3. Select the leaf-level parts, right-click, and select **Composite** from the pop-up menu:



A new node with the default name **Composite** is created containing the selected leaf-level parts as children:



Similarly, a new composite part can be created from two existing composite parts, such that they automatically become children or subassemblies of the new composite part.

4. As for leaf-level parts, select the two composite part nodes, right-click, and select **Composite** in the pop-up menu:



In this instance, composite parts **Composite** and **Composite 1** are rearranged to become subassemblies of new part **Composite 2**:



Exploding Composite Parts

A composite part can be exploded such that it is removed from the tree and any child parts of the node are moved up one level.

To explode a composite part:

1. Select the composite node, right-click, and select **Explode** from the pop-up menu:



The selected composite part node is removed and the child nodes are positioned at the next level up:



If the part contains no child nodes, the composite part node is removed.

Using Part Contact Data to Create Interfaces

A geometry part has a *Contacts* property that provides a list of all other parts with which the part is in contact. This part contact data can be used to generate interfaces for meshing and physics coupling automatically without having to create and assign coincident faces manually.

The two key steps in successfully obtaining interfaces are to:

- Generate part contact data using an appropriate method.
- Assign parts to regions in the correct manner.

Using this procedure is useful when you have lots of parts that are in contact with each other and want to produce a conformal mesh at the interface of each part. This procedure requires you to use the surface remesher with the tetrahedral or polyhedral volume meshers.

Generating Part Contact Data

Part contact data can be obtained through the following routes:

- For CAD parts, if the original parts are imprinted on their adjacent neighbors before being imported, select the **Create Part Contacts from Coincident Entities** option on the *Surface Import Dialog*.

  This option causes STAR-CCM+ to detect the coincident faces that are within the specified *Coincident Face Tolerance* automatically and create corresponding part contacts. For 3D-CAD models, this option is used internally when converting them to 3D-CAD parts. See the section, Creating Part Contacts From Coincident Entities for more information.

- The imprint mode of the Surface Preparation Tool can be used to generate coincident part surfaces and associated part contacts.

  This facility is available for all part types, including surface mesh parts and inherited part surfaces. An example of the procedure is given below in section, Imprinting Parts on their Neighbors.

- The **Find Part/Part Contact...** menu action can be used to run a proximity analysis that detects adjacent parts.

  This action does not imprint neighboring parts: imprinting happens during the surface and volume meshing process, which can only take place after the parts have been assigned to regions. Further information is given in section, Finding Part Contacts Automatically.

- It is also possible to create part contact data manually, either by using the **Create Part/Part Contact** menu item or by populating node properties directly.

A part contact browser is provided to give easy visualization of the parts that are in contact.

Assigning Parts to Regions

It is crucial that all parts are assigned to new regions in a single operation.

To assign parts to regions:

1. Select all parts in the object tree using `<Ctrl><Click>` or `<Shift><Click>`.
2. Right-click and choose **Assign Parts to Regions...**

   The *Assign Parts to Regions* dialog appears.
3. To identify which options to choose for **Region Mode** and **Boundary Mode** on this dialog, use the information in the table below. The **Create Interfaces from Contacts** option must always be activated.

Table of Valid Methods and Options

The table below summarizes the routes for generating part contact data, and gives the valid options that can then be used when assigning parts to regions using the *Assign Parts to Regions* dialog.

Status of original parts	Relevant Routes for Creating Part Contact Data	Valid Part Types for Route	Valid Options for Region Assignment	
			Region Mode	Boundary Mode
Neighboring parts that are already imprinted	Activate the **Create Part Contacts from Coincident Entities** option when importing. For 3D-CAD parts, this option is always	CAD and 3D-CAD only	One region for all parts	One boundary per part surface
			One region per part	One boundary for all part surfaces
			One region per part	One boundary

	activated internally when converting models to parts			per part surface
Neighboring parts that are not imprinted	Manually imprint parts using the Imprint mode of the Surface Preparation Tool	All part types	One region for all parts	One boundary per part surface
			One region per part	One boundary for all part surfaces
			One region per part	One boundary per part surface
	Select the **Find Part/ Part Contacts...** menu action	All part types	One region per part	One boundary for all part surfaces
	Manually define part contact data	All part types	One region per part	One boundary for all part surfaces

Example Process

The steps below illustrate the process of how interfaces are generated from part contact data.

(1) Part contact information is stored on each body



(2) Parts are assigned to new regions (in a single operation)


Star 1
Geometry
3D-CAD Models
Parts
Body 1
Body 2
Tags
Filters
Update
Highlight
Composite
Transform
Duplicate
Combine
Boolean
Split Non-Contiguous
Split by Surface Topology
Repair Surface...
Re-tessellate...
Assign Parts to Regions...
Set Region
Apply Tag
Statistics...
Multiple Objects - Pr
Properties
Region
Tags
Expert
Contacts
Face count
Meta Data
Assign Parts to Regions
Show All
Filter by Path
Parts
Body 1
Body 2
Create a Region for Each Part
Select...
Create One Boundary for All Part Surfaces
Boundary
Create One Feature Curve for All Part Curves
Feature Curve
Create Interfaces From Contacts
Apply
Close
Help

The result of this operation is the creation of two new regions and an interface for the contact between the regions.



(3) The final mesh contains a conformal interface where regions are in contact



See Also

Surface, Contact, and Curve Edit Controls

Creating Part Contacts from Coincident Entities

STAR-CCM+ can automatically create part contacts from pairs of coincident faces when importing CAD models, provided the **Create Part Contacts from Coincident Entities** option is selected on the *Import Surface Options* dialog.

This option is exposed when the import mode is set to create geometry parts.

The operation searches all of the imported CAD bodies for pairs of coincident faces that are within the specified *Coincident Face Tolerance*. For each pair found, part contact data is created and stored in the *Contacts* property of each part. A strict definition of coincidence is adhered to, in that the faces belonging to each pair must have an identical surface definition and identical trimming loops. They must overlap exactly without any offset unless they are imprinted.

This same operation is applied when converting 3D-CAD models to geometry parts, and when faces are imprinted using the single or multi-part imprint tool.

The example below shows a 3D-CAD model where the cylinder has been imprinted on the surface of the block. As the imprint operation results in a pair of coincident faces, STAR-CCM+ creates part contact data when the 3D-CAD model is converted to a geometry part.



Geometry Parts: standard view



Split By Patch view on larger block shows the imprinted face

The advantage in using this method to generate the part contact data is that each part surface can be mapped to its own boundary during region assignment. A separate boundary is created for the interface. The **One boundary per part surface** option gives this result.

Imprinting Parts on Their Neighbors

In some cases, the imported data results in multiple geometry parts that have multiple part surface per part, but the parts have not been imprinted on their neighbors. For these cases, it is possible to use the single or multi-part imprint mode of the surface preparation tool to generate separate part surfaces at the point of contact.

Part contact data is created automatically during the imprint process: you can review the contact data using the part contact browser. Following this procedure means that the **One boundary per part surface** option can be used during region assignment.

As an example of this procedure, consider the five parts that are shown below which contact at various points. On import, no coincident face pairs were present in the model.



1. It is good practice to check that any surface repairs are complete before imprinting.
2. All parts are selected, and the surface preparation tool is launched using the right-click menu option, **Repair Surface...**



3. In the *Surface Repair Options* dialog, to include all **Geometry Parts** that are selected, click **OK**.
4. To change the selection of parts, click **Select**.



5. In the *Select Objects* dialog, click the button to the left of each part to select or deselect that part, then click **OK**.



6. In the *edit* panel that appears for the surface preparation tool, the **Mode** must be changed to **Merge/Imprint**.



7. As multiple parts are present, the *Imprint Mode* is set to **Multi-Part Imprint**. In this mode, an automatic option is provided to find pairs of contacting parts (within the *Part Browser* group box).
8. The **Tolerance** value must be set to an appropriate value for the model, that is, a value less than the minimum expected gap thickness between non-contacting parts.
   For this example, the value is reduced to `0.0001 m`.

9. The **Find Pairs** button is clicked to execute the automatic search.



10. The part pairs that are found can be reviewed using the **Browse Pairs** navigation buttons, and if necessary, the button can be used to delete the pair.
11. Assuming all pairs have been correctly identified, the **Imprint All** button is pressed to begin the imprint operation.

    Information on the number of new interfaces that are created is shown in the *Output* window, and a final update is made to the *Graphics* window showing the newly created interfaces.

    For this example, the final view in the *Graphics* window after the imprint operation shows the creation of five new interfaces.



12. When you click **Close** to exit the Surface Preparation Tool, the *Contacts* property of each body is populated with the results of the imprint operation. **Body 4**, for example, is in contact with three other bodies.



The final stage is to assign all the bodies to regions in a single operation.

13. All parts are selected, and the **Assign Parts to Regions** option is chosen from the right-click menu.



14. In the *Assign Parts to Regions* dialog that appears, the following options are selected:

    - **Create a Region for Each Part**
    - **Create a Boundary for Each Part Surface**
    - The **Create Interfaces from Contacts** option is activated.



On clicking **Apply**, parts are assigned to new regions and faces to new boundaries. Interfaces are created from their own boundaries and shown in the **Interfaces** manager node.



You can see that separate boundaries have been created for the portions of the original surface that were separated out during the imprint process. Any other separate part surfaces in the original imported surface would also have received their own boundary.

Finding Part Contacts Automatically

Part contacts can be found automatically between two or more parts that are based on a user supplied tolerance value. This feature is useful in circumstances where several parts exist, some of which can be touching two or more other parts.

To find part contacts automatically:

1. Select a starting part node, right-click, and select **Find Part/Part Contacts...** from the pop-up menu:



The *Part/Part Contact Search Options* dialog appears:



The original part is already selected and any other parts that are in contact with the first part are also available to select.

2. You can specify the **Tolerance** value in any allowable units.
3. Clicking **OK** starts the process and generate part contacts for each part combination that have surfaces that are within the tolerance value of each other.

   The *Output* window indicates the total number of part contacts produced. Each part node is updated with the appropriate part contact information:



In the above example, the **Body 2** part is added as a contact for the **Body 1** part. Selecting the **Body 2** part node would show **Body 1** as a contact. The part contacts list can be edited if necessary.

Creating Part Contacts Manually

This section describes how to create part contact data manually.

There are two ways to create part contact data manually:

- Using the action **Create Part/Part Contact** on the right-click menu for parts.
- Editing the *Contacts* property directly.

Using the Menu Action

Part contacts can be created manually between two or more parts. This operation is useful when the parts have surfaces that are in contact with one another and require an interface for the physical coupling once the mesh has been built.

To create the part contact manually:

1. Select the part nodes of interest, right-click, and select **Create Part/ Part Contact** from the pop-up menu:



The part contact is created and updated in the *Properties* window for each part. For example:



In the above example, the **Body 2** part is added as a contact for the **Body 1** part. Selecting the **Body 2** part node would show **Body 1** as a contact. The part contacts list can be edited if necessary.

This method does not check whether the parts are really in contact with one another. Each part is added to the *Contacts* list of all the other parts that are initially selected.

Editing the Part Contacts List

The *Contacts* property for each part node can be used to either define new part contacts or edit existing ones for the part.

2. Selecting the part node shows the current contacts list:



3. Clicking the *Contacts* property opens the *Select Objects* dialog:



4. The parts that are in contact with the originating part node can be selected from the list on the left-hand side and the **Close** button pressed.

   The *Contacts* list for the part node is updated accordingly:



Always define part contacts in pairs such that each part has a reciprocal member. If the part is removed from the list, also edit the reciprocal part to remove the partner node.

Browsing Part Contact Data

The *Part/Part Contact Browser* allows you to scan through the part contact data and see the contacting parts in the *Graphics* window.

This operation automatically updates the view to hide parts that are not part of a contact list.

To start the browser:

1. Right-click a part (or selection of parts) and choose **Browse Part/Part Contacts...**



The *Part/Part Contact Browser Options* dialog, which is a form of the *Select Objects* dialog, appears.

2. In this dialog, choose all of the parts for which you want to visualize part contact data and click **OK**.



The **edit** panel appears showing the interface for the contacts browser.



Parts that are selected in the *Part/Part Contact Browser Options* dialog are listed in the menu within the *Source Parts* group box: this menu can be

used to jump directly to a particular source part. The navigation buttons beneath the menu can be used to step through the list of parts in sequence, and also provide buttons to jump to the start or end of the list.

The menu within the *Target Parts* group box lists all parts that are in contact with the source part shown in the *Source Parts* group box above. At the same time, the *Graphics* window is updated to show the actual pair of parts, source, and target, that are selected in their respective menus. The source part is colored solid blue, the target part by solid red, and other parts in contact with the source part is shown in a transparent shade of green.

In the example scene below, the blue part is the source part, the red part is the target, and the other transparent parts are in contact with the source part. These other parts appear in the menu list under *Target parts*.



3. The link between a particular pair of parts can be broken using the **Remove this contact** button; links between all parts can be broken using the **Remove all contacts** button.

   Part contact data on parts is updated accordingly.

4. Ticking the **Show all sources** checkbox results in all parts that are contained within the menu of the *Source Parts* group box being displayed in a transparent manner in the *Graphics* window.

   A lighter level of transparency is used for those parts that are not involved in the current source/target pairing.

   An example is shown below.



5. Exit the *Part/Part Contact Browser* by clicking **Close**.

   The associated scene in the *Graphics* window is closed automatically.

Updating Part Interfaces

Updating part interfaces between two regions becomes necessary when one or more interfacing parts have been modified (for example it has been rotated), or has changed region.

In these cases, the original interface that existed between the two regions is now obsolete. If a part (or part surface) has been modified, perform an imprint before updating the interfaces.

Interfaces can be updated using the following methods:

- Updating the interfaces from the Parts manager node.
- Updating the interfaces from the Part/Assembly node.

In both cases, STAR-CCM+ checks part contacts, ensuring that appropriate interfaces exist where needed, and that interfaces that are no longer valid are removed.

Updating Part Interfaces Using the Parts Manager Node

If a geometry has changed, for example as a part has moved from one region to another, the part interfaces must be updated.

There are two possible scenarios where interfaces must be updated:

- The regions/boundaries that parts or part surfaces are associated with have changed.
- A geometry part or part surface has been modified (for example, the geometry has been rotated). In this case, perform an imprint before updating the interfaces.

To update the interfaces using the **Parts** manager node:

1. Right-click the **Parts** manager node and select **Update Interfaces**.



Updating the interfaces using the **Parts** manager node ensures that all existing parts have up-to-date interfaces. This method is the recommended way to ensure that all interfaces in the case are up-to-date.

Updating Part Interfaces Using the Part/Assembly Node

If a geometry has changed, such as a part (or assembly) has moved from one region to another, the part interfaces must be updated.

There are two possible scenarios that require you to update the interfaces:

- The part (or assembly) has changed region or boundary.
- A geometry part or part surface has been modified (for example, the geometry has been rotated). In this case, perform an imprint before updating the interfaces.

To update the interfaces of a specific part or assembly:

1. Right-click the node of the desired part and select **Update Interfaces**.



Only the interfaces affecting the selected part is checked and updated using this method.

Managing Geometry Parts

Tags and filters can be used to together to make finding and selecting geometry parts easier.

Once a tag has been created, it can be applied to one or more leaf-level or composite parts to flag them as being different from other parts. In order to distinguish one tag from another, the icon that is used for the tag can also be replaced using any suitable image. A filter can then be created based on the tag or alternatively be specified metadata properties. The filter can then be used to limit the parts that are available for selection in various dialogs and tools within STAR-CCM+.

What Are Tags?

Tags can be used to help organize parts in the model by allowing tagged parts to be filtered when selecting parts at a later stage.

A green tick-mark by default, as shown here ✓, represents a tag by default. An example of a tag that is applied to two different geometry parts is shown below:



Tags can be added to geometry parts, part surfaces, and part curves and can include a user-specified image for the icon. Once a tag has been created and applied, it can be used as part of a filter for when selecting parts, such as when parts are assigned to regions.

See Also

Understanding Geometry Parts

Creating User Tags

Any number of user tags can be created for flagging geometry parts.

To create a user tag:

1. Right-click the **Tags** manager node and select **New Tag** from the pop-up menu:



A new node that is called **User Tag** appears:



You can rename the user tag so that it has a more descriptive and unique name before it is applied.

2. To rename the user tag, right-click the tag node and select **Rename...** from the pop-up menu:



3. To apply the name change to the node, enter the new name in the inline editor and press `<Enter>`.

A green tick mark ✓ represents a user tag by default. In order to help differentiate one tag from another, a custom icon can be used instead if desired, which is especially useful once they have been applied to a part.

4. To use a custom icon, right-click the tag node and select **Set Custom Icon...** from the pop-up menu:



The *Open* dialog appears allowing the selection of the bitmap icon image:



5. The file selector lets you locate the icon image to use, and preview the selected file in the dialog.

   The allowable image file types are Joint Photographic Experts Group (.jpg, .jpeg), Graphics Interchange Format (.gif) and Portable Network Graphics (.png).
6. Once the file has been selected, press the **Open** button to import it.

The image is reduced to 16x16 pixels and displayed next to the tag node.

For example:



The tag custom icon is stored as part of the sim file so the original image used is not maintained on the system.

Applying a Tag to a Part

A tag can be associated or applied to a part once it has been created.

Each part, whether a leaf-level or composite, can have two or more different user tags applied to it if necessary. There are three different methods of assigning or applying tags to parts:

- dragging and dropping the tag node onto the part node
- dragging and dropping the part node onto the tag node
- applying the tag through the part node right-click options
- editing the part node properties

Each of these methods is discussed below.

1. The simplest and easiest way of applying a user tag to a part is to drag and drop the tag node onto the part node, or the part node onto the tag node.

   A miniature version of the tag icon is placed next to the part node indicating that it has been tagged, as shown below:



2. You can repeat this procedure as many times as necessary for each user tag and part node in order to establish the required tagging.

Note: You can also use the above procedure to create tags on **Geometry** > **Contacts** nodes.

Alternatively, the user tag can be applied directly from the part node pop-up menu.

3. Right-click the part node and select **Apply Tag** and then the name of the tag (**User Tag** for example):



This results in the tag being applied to the **Body 1** part. The tag icon appears next to the part node.

4. If the **Apply Tag** > **New...** option is selected for a part, then a new user tag is automatically created and applied to the part at the same time.

For example:



A new user tag is created (called **User Tag 2** in this instance) and applied to the part:



If no tags have been defined so far, then the **Apply Tag** menu option has no further selections.

5. Simply selecting **Apply Tag** behaves the exact same way as when the **New...** option is selected, as described above.

For composite parts, you can tag all the child parts that belong to the composite (including other composite parts).

6. Right-click the node and select **Apply tags to child parts** from the pop-up menu:



The *Select tags* dialog appears:



7. Select the tags of interest from the list and click **Apply**.

   The child part nodes (including the original composite node) are updated with the appropriate tag icon. For example:



Currently, tag filters ignore tagged composite parts: only the tagged child leaf-level parts are listed when a tag filter is used.

The user tags that are specified for a part can be reviewed or modified by selecting the part node and looking at the *Tags* property. For example:



In the above example, no tags have been applied for the **Body 1** part. You can also see tags by looking for any tag icons that are included next to the part node icon.

To add a tag to the part node:

8. Click the *Tags* property and use the object selector.

The part node is then updated with the appropriate tag icon and the tag name is added to the *Tags* property list:



Applying Tags to Part Surfaces

The application of tags to part surfaces is similar to applying them to parts.

A brief overview of applying tags to part surfaces is provided here for completeness.

As for parts, the methods of assigning or applying tags to part surfaces are as follows:

- dragging and dropping a tag node onto a part surface node
- dragging and dropping a part surface node onto a tag node
- applying the tag through the part surface node right-click options

- editing the part surface node properties

The simplest and easiest way of applying a user tag to a surface is to drag and drop the tag node onto the part surface node. A miniature version of the tag icon is placed next to the face node indicating that it has been tagged.

Alternatively, the user tag can be applied directly from the part surface node pop-up menu.

1. Right-click the surface node and select **Apply Tag** and then the name of the tag (**User Tag** for example):



This results in the tag being applied to the **Faces** surface. The tag icon appears next to the surface icon, as shown below:



2. To review the user tags that are specified for the surface, select the surface node and look at the *Tags* property.
3. To modify the selections, activate the object selector by clicking the **[]** part of the *Tags* property.

Applying Tags to Part Curves

The application of tags to part curves is similar to applying them to parts.

A brief overview of applying tags to part curves is provided here for completeness.

As for parts, the methods of assigning or applying tags to part curves are as follows:

- dragging and dropping a tag node onto a part curve node
- dragging and dropping a part curve node onto a tag node
- applying the tag through the part curve node right-click options

- editing the part curve node properties

The simplest and easiest way of applying a user tag to a curve is to drag and drop the tag node onto the part edge node. A miniature version of the tag icon is placed next to the node indicating that it has been tagged.

Alternatively, the user tag can be applied directly from the part curve node pop-up menu.

1. Right-click the curve node and select **Apply Tag** and then the name of the tag (**User Tag** for example):



This results in the tag being applied to the **Edges** curve. The tag icon appears next the curve icon, as shown below:



2. To review the user tags that are specified for each curve, select the curve node and look at the *Tags* property.
3. To modify the selections, activate the object selector by clicking the **[]** part of the *Tags* property.

Removing a Tag from a Part

A user tag can be removed from a part node using the *Tags* property that belongs to the node or by using the part tags manager for composite parts.

Selecting the part node displays the tags that have been applied. For example:



In the above case, a single tag that is called **User Tag** has been assigned to the **Body 1** part node.

To remove this tag:

1. Click the *Tags* property so that the object selector is displayed.

   The properties for the part node are then updated and the tag icon is removed from the part node:



Managing Tags on Children of Composite Parts

For composite parts, you can manage the tags that are applied to child parts.

For composite parts, you can manage the tags that are applied to child parts:

1. Right-click the **Geometry** > **Parts** > **Composite** node and select **Manage Tags on Children...**



The *Edit tags on assembly* dialog opens. This dialog can be used to remove or apply tags to all of the child parts. For tags that have only been applied to some of the parts in the composite part, the checkbox appears grayed out.


Edit tags on assembly
Tag
Tag 2
OK
Cancel
Help

Deleting a Tag

A tag node can be deleted in a similar way as other objects.

To delete a tag node:

1. Right-click the node and select **Delete Tag...** from the pop-up menu.

   For example:



All parts and filters that are associated with the tag are also updated so that the link is removed from each.

What Is Metadata?

The metadata specifier can be used to associate a key and value with the part.

A key is typically a property of the part, such as “material”, “color” or “grade”. The value is the respective quantity or identifier, such as “steel”, “blue” or “2”. The value can be either a string (word) or an integer. In order to create a valid metadata input, both a key and a value must be specified. Each geometry part can have one or more metadata that are specified for it.

Metadata must be entered using the format `{key:value}`. For example:

```
{material:steel}
```

If more than one metadata specifier is required for the part or part surface, then a comma can be used to separate them. For example:

```
{material:steel,color:blue}
```

The *Properties* window lets you enter the metadata key and value for the part:



The specifier can be entered directly in the window (within the {} braces) or the ... button can be pressed to open the property customizer window. The metadata specifier can then be entered. For example:



Clicking **OK** applies the metadata text and reorganizes it if necessary to conform to the requirements of STAR-CCM+. The metadata can be reviewed afterwards in the part *Properties* window:



Notice that STAR-CCM+ automatically adds quotation marks `''` to `material`, `steel`, and `grade`, but not to the `2`. When quotation marks are not applied, this means that the `2` is an integer value. When quotation marks are applied to the value, for example `{'grade':'2'}`, this means that the `2` is a string. The same holds true for decimal values `{'grade':2.0}` and boolean values `{'conductive':true}`, where the quotation marks around the value are omitted.

Once the metadata information is entered, it can be used as part of a filter or series of filters in order to help identify and select parts at a later stage. The metadata specifier has no influence on the simulation and is provided as an organizational feature only.

Metadata on Part Surfaces

Part surfaces can also have metadata.

The procedure for setting up metadata is the same as for parts.



Manipulating Parts

STAR-CCM+ allows leaf-level parts to be manipulated in several different ways once they have been imported or created.

Note: The manipulation options available are dependent on the part type. For example, the splitting non-contiguous parts options are only available for mesh parts.

For a reference page that contains a short description of each right-click part operation, visit the Part Operations Reference Page.

See Also

Understanding Geometry Parts

Retessellating a Part

The surfaces of a geometry part are represented using triangles which come from a process that is known as tessellation.

For 3D-CAD, imported CAD, and shape part types:

1. The triangle density that is used to tessellate a surface can be modified using the **Re-tessellate...** option on the right-click menu for the part.



Selecting this option opens the *Tessellation options* dialog. In addition to providing control over the tessellation density, the dialog also provides control over the marking of CAD edges as features edges. These options behave in an identical manner to the options that are provided on the *Surface Import Options* dialog. See the sections on Tessellation Density and Mark Feature Edges for details.



It is possible to see the triangles that are used to represent the surfaces by choosing to display the mesh in a geometry scene for which the Geometry representation is active.

Transforming a Part

Five different types of transformations are provided to transform parts.

The five different transformations that are allowed are:

- Coordinate Transform
- Scaling
- Translation
- Rotation
- Reflection

These options appear when right-clicking a part node and selecting **Transform** from the pop-up menu.



The transform options are useful in situations where parts are not aligned correctly or have been saved at a different scale to other parts. Transform operations can be performed on one or more parts simultaneously. When transforming a composite part, the selected operation affects all leaf-level parts that are contained within the composite.

For 3D-CAD parts, transforms can be performed on the original 3D-CAD bodies using the 3D-CAD solid modeler and then updated

accordingly. Any transforms made to 3D-CAD geometry parts must be repeated if their underlying 3D-CAD model is modified.

Transforming Parts to a Coordinate System

Transforming a part to a coordinate system moves a part from a position relative to a source coordinate system, to a position relative to a destination coordinate system.

This transform is useful for both translating and rotating a part in the same operation, and maintains the exact shape of the part during the transformation.

For example, the highlighted corner in the following image is at the origin of the source coordinate system.



Transforming to the destination coordinate system has the following effect.



The highlighted corner is now at the origin of the destination coordinate system, and the shape of the part is maintained.

To transform a part to a coordinate system:

1. Right-click the part node and select **Transform > Coordinate Transform...**



The *Coordinate Transform* dialog opens.



2. Select the source coordinate system from the **From Coordinate System** menu and the destination coordinate system from the **To Coordinate System** menu.
3. Click **OK** to perform the transform.

If you select a spherical or cylindrical coordinate system, the transform uses a Cartesian coordinate system with the same origin as the selected system. This correspondence prevents distortion of the part when transforming to a non-Cartesian coordinate system.

Scaling a Part

You can scale a part if it uses units that are different to other parts you have already imported.

Scaling can take place in any or all of the three coordinate directions and can be used to increase or decrease the size of a part.

To scale a part:

1. Right-click the part node and select **Transform > Scale...** from the pop-up menu:



The *Scale Parts* dialog appears:



By default, the **Scale all dimensions uniformly** checkbox is ticked, which means that all dimensions are scaled equally. In this mode, the scale factor can be entered into the **X** text box.

2. To scale each coordinate direction independently, deselect the **Scale all dimensions uniformly** checkbox.
3. The coordinate system in which the scaling takes place can be selected from the **Coordinate System** menu.
4. Click **Apply** to carry out the operation, which is shown on the display if a valid scene is open.
5. Additional scaling operations on the part can be performed as described above and clicking **Close** when complete.

If a valid scene is open, the display is updated accordingly.

Translating a Part

Translating a part can be used to rearrange the geometry so that it fits together as intended.

This problem can arise when different parts are built in different coordinate systems. Translation can take place in any of the three coordinate directions.

To translate a part:

1. Right-click the part node and select **Transform** > **Translate...** from the pop-up menu:



The *Translate Parts* dialog appears:



2. The *Translation Vector* component values can be entered using either positive or negative values, along with the desired units.
3. The choice of **Coordinate System** is also allowed.
4. Click **Apply** to carry out the operation, which is shown on the display if a valid scene is open.

The operation can be repeated as many times as desired.

Rotating a Part

You can rotate a part to reorient it with other static components.

Rotation is allowed in any axis direction.

To rotate a part:

1. Right-click the part node and select **Transform** > **Rotate...** from the pop-up menu:



The *Rotate Parts* dialog appears:



2. The **Angle** of rotation can be entered (in radians by default) along with the *Axis Vector* for the rotation.
3. The choice of **Coordinate System** can also be made.
4. Click **Apply** to carry out the operation, which is shown on the display if a valid scene is open.

Reflecting a Part

Reflecting a part can be used to reposition it according to its mirror reflection in a given axis.

This transformation can be useful when the original part is a half model of the geometry. You can create the full model by first duplicating the part and then reflecting the copy (or original).

To reflect a part:

1. Right-click the part node and select **Transform > Reflect...** from the pop-up menu:



The *Reflect Parts* dialog appears:



2. Enter the **Normal Vector** values and select the **Coordinate System** before clicking **Apply**.

   If a valid scene is open, the part display is updated.

3. Clicking **Apply** a second time without changing the vector component values returns the part to its original position.

Duplicating a Part

You can duplicate a part in instances where only a half, quarter, or some fraction of a symmetrical or axis-symmetric model has been imported and the full geometry is desired instead.

Duplicating a part produces an identical copy of it with the same properties as the original, other than region and contacts specification, which are left null. The duplicate also occupies the same location as the original until it can be transformed. If a 3D-CAD part is duplicated, the duplicate copy is a leaf-level CAD part rather than a 3D-CAD part. That is, the link to the originating 3D-CAD model is lost.

To duplicate a part:

1. Right-click the part node and select **Duplicate** in the pop-up menu:



If the source part includes any repairs, a copy of the part is made and retessellated (although the level of tessellation is retained). The new part node appears next to the original:


Servers
Star 1
simulation
scene/plot
Star 1
Geometry
Parts
Body 1
Body 1 2
Body 2
Tags
Filters
Continua
Regions
Scenes
Representations

Combining Parts

Two or more parts can be combined to create a single mesh part.

To combine parts:

1. Select the part nodes of interest, right-click and select **Combine** from the pop-up menu:



The first selected part is used as the base for all the other parts. All part surfaces and curves belonging to each part are grouped in the base part, as shown below:



As this operation results in a single mesh part, it is not possible to recontract CAD parts from the single part. It is possible, however, to construct new mesh parts from one or more of the part surfaces contained within the combined part.

Two or more composite parts can also be combined as described above. For example:



All part bodies belonging to each composite are grouped into the first composite part selected:



Combining two composite parts in this way does not result in child CAD or 3D-CAD parts being converted to mesh parts. They retain their original type. You can undo the above process by first making a new empty composite part and then dragging and dropping the original part bodies back into the composite node.

Performing Boolean Operations

This section describes the basic Boolean operations that can be performed on geometry parts.

The following basic Boolean operations can be performed on geometry parts:

- Unite
- Intersect
- Subtract
- Imprint

Boolean operations are carried out by selecting composite and leaf-level part nodes under the **Geometry** > **Parts** manager node, then choosing one of the **Boolean** options in the right-click menu. Following a Boolean operation, the original parts remain unchanged, and the resulting geometry is added to the tree under a new leaf-level node that is called **Unite**, **Intersect**, or **Subtract**. The type of leaf-level part (CAD or mesh) created following a Boolean operation depends on the nature of the input parts and, where appropriate, which Boolean mode is selected.



Selecting the Boolean Mode

Boolean operations on geometry parts can be carried out using either the CAD data, if available, or the discretized surfaces of the parts.

The method that is used is known as the Boolean mode. If at least one of the selected parts is a mesh or shape part, including those parts within any selected composites, STAR-CCM+ uses the discretized surface (tessellated surface mesh) for all selected parts to carry out the operation. This mode is referred to as the **Discrete** Boolean mode, and in this mode the operation is carried out based on the current tessellation properties of each part. Any surface preparation that has been carried out on the input parts is incorporated into the final part, which is of Mesh type.

If the selected parts are all CAD-based parts (including 3D-CAD), you can choose between **CAD** or **Discrete** Boolean mode when the operation is executed. When the **CAD** option is selected, the operation is carried out based on the original CAD data and the resulting part is of CAD type. If you have carried out any surface preparation on the CAD parts, such as surface repair or altering the tessellation density, these changes only affect the discrete surfaces of the part, not the CAD data. Therefore, to preserve the effects of the surface preparation in the resulting part, select the **Discrete** Boolean mode.

To select a Boolean mode:

1. Choose the CAD parts to include in the Boolean operation from the simulation tree.
2. Right-click one of the nodes and select either the **Unite**, **Intersect**, **Subtract**, or **Imprint** option.
3. In the dialog that appears, set the **Boolean Mode** to either **CAD** or **Discrete** using the drop-down menu.

   The dialog for the subtract operation also includes a **Target** option.



Uniting Parts

The Boolean operation **Unite Parts** joins all the selected parts into one single part, the surface of which is formed from segments of the constituent part surfaces.

This operation differs from the Combine operation, which combines the selected parts into one part, but retains the part surfaces as separate, whole surfaces. To carry out a unite operation successfully, the selected parts must be in contact with each other, or overlap. If the parts do not meet this criterion, the part that the Boolean operation creates contains no surface data.

A unite operation can be performed on a combination of leaf-level and/or composite parts, or on a single composite part that contains at least two other parts. When a unite operation is carried out using parts from different levels in the tree, the resulting part is added to the tree at the same level as the highest selected part.

To unite parts:

1. Select the appropriate parts under the **Geometry** > **Parts** node, using **<Ctrl><Click>** to select multiple parts.
2. Right-click one of the selected parts and choose **Boolean** > **Unite Parts**.



3. If the *Boolean Mode* dialog appears, select a Boolean mode using the drop-down menu and click **OK**.

A **Unite** node is added to the tree for part that the Unite operation creates.



Intersecting Parts

The Boolean operation **Intersect Parts** finds the common volume that the selected parts share and creates a single part that is based on this volume.

The surface of the part is formed from segments of the constituent part surfaces. To carry out an intersect operation successfully, all the selected parts must share a common volume. If the parts do not share a common volume, the part that the Boolean operation creates contains no surface data.

An intersect operation can be performed on a combination of leaf-level and/or composite parts, or on a single composite part that contains at least two other parts. When an intersect operation is carried out using parts from different levels in the tree, the resulting part is added to the tree at the same level as the highest selected part.

To intersect parts:

1. Select the appropriate parts under the **Geometry** > **Parts** node, using **<Ctrl><Click>** to select multiple parts.

2. Right-click one of the selected parts and choose **Boolean** > **Intersect Parts**.



3. If the *Boolean Mode* dialog appears, select a Boolean mode using the drop-down menu and click **OK**.

An **Intersect** node is added to the tree for part that the Intersect operation creates.



Subtracting Parts

The Boolean operation **Subtract Parts** subtracts a defined shape from the volume of a target part.

Any number of parts can be used to define the shape to subtract. The surface of the resulting part is formed from segments of the constituent part surfaces.

A subtract operation can be performed on a combination of leaf-level and/or composite parts, or on a single composite part that contains at least two other parts. When a subtract operation is carried out using parts from different levels in the tree, the resulting part is added to the tree at the same level as the highest selected part.

To perform a subtract operation:

1. Select the relevant parts under the **Geometry** > **Parts** node in any order; use **<Ctrl><Click>** to select multiple parts.
2. Right-click one of the selected parts and choose **Boolean** > **Subtract Parts**.



The *Subtract parts* dialog appears.



The **Target Part** menu can be used to select the part from which the other parts are subtracted. The dialog can include a **Boolean Mode** menu, from which you can select a Boolean mode for the operation.

3. When the appropriate selections have been made, click **OK** to proceed with the operation.

A **Subtract** node is added to the tree for the part that the Subtract operation creates.



Imprinting Parts

The Boolean operation **Imprint Parts** merges coincident surfaces from the selected parts together.

When coincident surfaces are not identical (in shape or size), the edges of the smaller surface are used to split the larger surface. By imprinting parts, you can successfully generate interfaces between multiple connected parts. The interfaces are essential for conformal meshing.

An imprint operation can only be performed on a combination of leaf-level and/or composite parts, or on a single composite part that contains at least two other parts.

To perform an imprint operation:

1. select the relevant parts under the **Geometry** > **Parts** node in any order; use **<Ctrl><Click>** to select multiple parts.
2. Right-click one of the selected parts and choose **Boolean** > **Imprint Parts**.



When CAD data is available, the *Imprint Parts* dialog appears.



The **Boolean Mode** menu allows you to select what type of imprint to perform:

- **Discrete** - The imprint is carried out using the tessellated surface mesh of the parts.

Enter the tolerance for the imprint (with units) in the **Tolerance** text box. Parts that are in contact within this tolerance are imprinted together, and new part surfaces are created. This operation performs the same action as the Multi-Part Merge/Imprint Tool.

- **CAD** - The imprint is carried out using the original CAD data of the parts.

  If you select this option, you do not need to provide a tolerance. New surface patches are created for the imprinted surfaces--no new part surfaces are created.

If one or more of the selected parts have no CAD data available, the **Boolean Mode** menu locks to **Discrete**, as shown below:



3. Enter the tolerance for the imprint in the **Tolerance** text box, as described above.
4. Click **OK** to perform the imprint operation.

This operation modifies the existing parts and cannot be undone.

Splitting Non Contiguous Parts

A non-contiguous mesh part (that is, a mesh part that has faces that are not connected to other faces in the part) can be split into its individual component parts if necessary.

To split non-contiguous parts:

1. Right-click the part node and select **Split Non-Contiguous** from the pop-up menu:



The splitting operation proceeds and produces an additional part node for each contiguous component:



The output from a successful split operation is shown below:



Attempting to split a part that is already contiguous has no effect.

Splitting a Part by Surface Topology

The split by surface topology method can be used to split and combine parts that can be defined as distinct geometrical regions.

It is similar to the method used to split regions. The operation is based on the topology of closed, manifold, non-intersecting surfaces that the parts define. When assessing the surfaces, STAR-CCM+ considers the surfaces that are contained within one part to be a single surface. You do not need to combine or split any surfaces before carrying out the operation.

The technique allows you to:

- Separate two or more closed surfaces that are not in contact into two or more parts.

- Separate two or more closed surfaces that share a common interface.

- Create new parts from internal voids.



- Redefine the surfaces of parts that contain closed surfaces or other whole parts.



- Combine parts with shared part surfaces.

When this technique is used for multiple parts, STAR-CCM+ internally combines all the parts and treats them as a single part. If a part is not modified in the operation, it is returned as a CAD part with the same name as it had previously. If a part is modified, for example it gained extra surfaces from other selected parts, it is returned as a Mesh part.

To split a part by surface topology:

1. Right-click the part and select **Split by Surface Topology**.
2. If you wish to include several parts in the operation, select the parts using **<Ctrl><Click>**.
3. If you want to split parts by surface topology but combine parts that share a part surface, in the *Split Parts Options* dialog that appears, activate *Combine Parts With Shared Part Surfaces*.

Using Split by Surface Topology

An example where a part contains an inner void and one wishes to create another part from the inner surface is illustrated below.

4. Identify parts that can be split into one or more separate parts.

   In the figure below, the single part **Body 1** contains a cylindrical void.



5. Right-click the part that you want to split and choose **Split by Surface Topology**:



6. In the *Split Parts Options* dialog, click **OK**.



The result can be viewed in a geometry scene as shown below:



If separate part surfaces surrounded the inner void in the original part, these surfaces would be replicated in the splitting process. The source part (**Body 1** in the example above) is removed in the splitting process.

Replacing a Part

You can replace a single CAD or mesh part at any time. Replacing a part lets you provide new geometrical input for the meshing pipeline without affecting any of the other data input required to generate a mesh.

Only one part node can be replaced at a time. The operation does not apply to shape parts.

To replace a part:

1. Right-click the part node and select **Replace Part...** from the pop-up menu:



The *Open* dialog appears:



2. Select the replacement part file and click **Open** to import the data.

   When importing IGES files, to account for possible gaps that can arise in the sheet bodies during the translation, the following additional options are available:

   - Sewing Tolerance
   - Multi-Pass Sewing

The replacement part must be of the same type as the leaf-level node part that it is replacing. That is, a CAD part can only be replaced with a part originating from native CAD or neutral format data. A mesh part can only be replaced with a part originating from a surface mesh data file. The file types available in the *Open* dialog reflect this rule.

Also note that any manipulations or manual repairs that have been made to the original part are lost when the part is replaced.

Manipulating Part Surfaces

STAR-CCM+ allows leaf-level part surfaces to be manipulated in several different ways.

The following operations are possible depending on the part type:

- creating a part from a part surface
- creating an empty part surface
- combining part surfaces
- splitting non-contiguous part surfaces
- splitting a part surface by part curves
- splitting a part surface by angle
- splitting a part surface by patches
- splitting a part surface by contacts

In order to make part surfaces easier to search for, it is also possible to apply user tags to them.

Creating a New Part from a Part Surface

Leaf-level parts can have one or more of the surfaces describing the part geometry be moved into a new part if necessary.

This operation would be useful for extracting a fluid volume from an enclosing solid part, for example.

To create a part from a surface node:

1. Right-click the surface node or nodes belonging to the same part and select **Create New Part From Surfaces** in the pop-up menu:



A new part with the default name **Part** is created using the surface node geometry. The surfaces retain the same names:



The original surface nodes that are selected are removed from the originating part. This operation only applies to surfaces belonging to a single part.

Creating an Empty Part Surface

In STAR-CCM+, you can create empty part surfaces for an existing part. When using surface repair features, you can then assign faces to the empty part surfaces.

To create an empty part surface for any part:

1. Expand the **Geometry** > **Parts** > **[Part]** node.
2. Right-click the **[Part]** > **Surfaces** node and select **New Part Surface**.
   An empty **Surfaces** > **[Part Surface]** node appears.

Combining Part Surfaces

Two or more leaf-level part surfaces can be combined into a single part surface if necessary.

To combine the surfaces:

1. Select two or more surface nodes belonging to the same part, right-click and select **Combine** from the pop-up menu:



The surface nodes are combined into a single node (the first selected):



Note the following:

Boundary assignments on the secondary surface nodes are lost when they are different from the target (first selected) surface node.

Tag information relating to the secondary surface nodes are lost when they are different from the target (first selected) surface node.

Splitting Non Contiguous Part Surfaces

A part can have one or more of its faces split based on the non-contiguous nature of the surface.

This operation splits the selected surfaces according to whether they are connected to their neighbors or not.

To split a surface:

1. Right-click the surface node and select **Split Non-Contiguous** from the pop-up menu:



Any non-contiguous surfaces are split to form new surface nodes, as shown below:



Any tags belonging to the original surface nodes are not copied to the split nodes. The new surfaces are also not assigned to any boundary.

If the selected surfaces are contiguous, no action is taken.

Splitting a Part Surface by Part Curves

A part can have one or more of its surfaces split based on the part curves corresponding to the parent part.

To split a surface using curves:

1. Right-click the surface node and select **Split By Part Curves...** from the pop-up menu:



In the *Split Part Surfaces by Part Curves* dialog, the originally selected surface nodes are already selected.

2. To select the appropriate part curves for the surface click the **Select** button to the right of **Part Curves**:



3. In the *Select Objects* dialog, select the part curves that you require.



4. Click **Apply** to carry out the operation, and **Close** to close the *Split Part Surfaces by Part Curves* dialog.

Any new surfaces that the process creates are updated in the simulation tree:



Any tags belonging to the original surface nodes are not copied to the split nodes. The new surfaces are also not assigned to any boundary.

If no curves correspond with the selected surfaces, no action is taken.

Splitting a Part Surface by Angle

A part can have one or more of its surfaces split based on an angle.

To split a surface using an angle:

1. Right-click the surface nodes and select **Split By Angle...** from the pop-up menu:



The *Split Part Surfaces by Angle* dialog appears:



2. You can select more surfaces to split by clicking **Select** button to the right of **Part Surfaces**.
3. Enter the **Angle** value (in degrees) and click **Apply** to carry out the operation.

The new surface nodes are shown in the simulation tree:



Any tags belonging to the original surface nodes are not copied to the split nodes. The new surfaces are also not assigned to any boundary.

If the supplied angle value is too large, the split operation results in no additional surface nodes being created.

Splitting a Part Surface by Patch

You can split one or more surface nodes of a leaf-level part using the underlying patch information that is stored as part of the geometry.

To split a surface using patches:

1. Right-click the surface node and select **Split By Patch...**



Note: You must have a scene open to use this feature.

The *edit* tab opens in the explorer pane. The *Patch Selection Control* tools are not available until at least one *Part Surface Patch* is selected.



Here, you can edit the patches belonging to the selected surface node:

2. Select the part surface patches from the list box or from the geometry using the mouse in the current scene display.

   To select multiple patches:

   - Press the `<Ctrl>` key and click on multiple part surface patches, or
   - use the Patch Selection Control tools, or
   - use the zone selection tool from the Selection toolbar

   Selecting a patch results in it being highlighted in the scene window.



3. Enter the new **Part Surface Name**.



4. Click **Create** to create a surface node that is based on the selected patch information.
5. Repeat these steps for all required patches.
6. Click **Close** when done.

The new surface node for the part appears in the simulation tree.



Note: Any tags that belong to the original surface nodes are not copied to the split nodes. The new surfaces are also not assigned to any boundary.

Splitting a Part Surface by Contacts

You can split one or more part surfaces by the contacts that the part surfaces make.

The **Split by Contact** function only works with strong contacts that are created in the **Root Description**. For example, parts that are imprinted in 3D-CAD have strong contacts. If contacts are created in any other downstream description, **Split by Contact** will not work.

When parts are imprinted, the contacts that are created exist between the whole part surfaces that contain the contacting faces. However, when the part surface contact is selected, only the contacting faces are highlighted in the Graphics window:



Splitting part surfaces by contacts recognizes the contacting faces on coincident surfaces and specifies these collections of faces as independent part surfaces.

To split part surfaces using contacts:

1. Select at least one part surface node to split by contact.
2. Right-click the part surface node and select **Split By Contact**.
3. In the **Split By Contact Options** dialog, select an option. Create One Part Surface Per Contacting:

   - **Patch**

     A separate part surface is created from faces that belong to each patch within the contacting part surfaces.

     For example, if there are six patches contacting a part surface:



the part surface is split to create six corresponding part surfaces:



- **Part Surface**

  A separate part surface is created from faces that belong to each contacting part surface.

  For example, if a part contains three part surfaces:



  the contacting part surface is split to create three corresponding part surfaces:



- **Part**

  A separate part surface is created from faces that belong to each part within the contacting part surfaces.

  For example, if one part surface is contacted by one part:



the part surface is split to create one corresponding part surface:



In each example, specific contacting part surfaces are created that contain the exact faces that are in contact from each original patch, part surface, or part. The **Geometry > Contacts** update to show the specific part surfaces that form the contacts.

Selecting Patches for Part Surfaces

The *Patch Selection Control* tools include grow and shrink tools which facilitate the selection of patches for splitting.

These *Patch Selection Control* tools are only available for parts. There are no tools included for region boundary based patch splitting.

There are five *Patch Selection Control* tool buttons which are accessible once you select the Split by Patch option. These buttons become active once a *Part Surface Patch* has been selected from the *Part Surface Patches* group box or from the Geometry in the *Graphics* window. If all patches for the part surface are selected, then only the **Patch Clear** and **Patch Shrink** buttons are available.



The button descriptions are as follows:

Grow Selected Patches	Grows the selection one patch in each direction based on the perimeter vertices surrounding the initial patch selection.
Shrink Selected Patches	Shrinks the selection one patch in each direction based on the perimeter vertices of the selected patches.

	If all patches are selected then a shrink may not be possible. At least one free patch perimeter edge is required to start shrinking from. The result of growing a patch and then shrinking it does not always result in the original patch selection.
Multi-Grow Selected Patches	Opens a dialog which gives you access to a variety of patch multiple grow options.
Invert Selected Patches	Selects the inverse of the **Part Surface Patches** that are selected.
Clear Selected Patches	Clears any selection of **Part Surface Patches**.

Patch multiple grow options dialog

The default patch multiple grow options are shown in the dialog below:



Descriptions of each of the patch multiple grow options are specified below:

Number of Steps	Specifies the number of steps the patch selection is grown for.	
	Infinite	Grows the patch selection infinitely. Performing an infinite grow without including some reasonable Break Options can potentially take a long time.
	Finite	Allows you to specify a finite number of patch growth steps.
Grow Options	Allows you the option to grow your selection based on growth from edges or vertex.	
	Edge Based Growth	Grows the selected patches from the outer edges of each selected patch.
	Vertex Based Growth	Grows the selected patches from the outer vertex of each selected patch.
Patch Perimeter Break Options	Allows you to specify the type of feature which the patch perimeter grows to before breaking the selection. You can set the *Patch Perimeter Break options* to any or all of the following options:	
	Break at Feature Curves	Breaks the patch selection when the patch perimeter reaches a feature curve.
	Break at Non-manifold edges	Breaks the patch selection when the patch perimeter reaches a non-manifold edge.

	Break at Sharp Angles	Breaks the patch selection when the patch perimeter reaches edges between faces that meet at this angle or sharper.
	Angle	Becomes active when the *Break at Sharp Angles* option is activated and allows you to specify the sharp angle. The units of the sharp angle are specified in degrees only.

The break options provided are similar to those in surface repair. However, the break options criteria is only assessed on patch perimeters and not on patch interiors. For example, if a patch contains an interior sharp angle (as opposed to a sharp angle at the patch perimeter) then this sharp angle is ignored and the patch will always be completely grown into. Similarly, feature curves and non-manifold edges that are contained within a patch are also ignored. It is not possible to grow into only part of a patch - only whole patches are grown into, regardless of any break option features within that patch.

Smart Selection of Patches

Smart selection of patches lets you quickly grow a patch until it reaches a natural break point in the geometry. You can use smart selection to grow a patch until it reaches a natural break by double-clicking on a patch.

To combine a smart selection with any patches you have already selected, press and hold **<Ctrl>** before you perform the smart selection.

Manipulating Part Curves

STAR-CCM+ allows leaf-level part curves to be manipulated in several different ways.

The following operations are possible:

- computing part curves
- combining part curves
- splitting non-contiguous part curves.

In order to make part curves easier to search for, it is also possible to apply user tags to them.

Computing Part Curves

You can recompute part curves at any time. This feature lets you modify the curve definition using a number of different options, such as when more curves are needed.

To compute part curves:

1. Right-click the **Curves** manager node for the leaf-level part and select **Compute Part Curves...** from the pop-up menu:



The *Mark Part Curves Options* dialog appears:



2. Select the part surfaces for which the curves are required.
3. Select the required Part Curve Options which are used for the curve computation.
4. Clicking **Apply** proceeds with the operation.
5. Once it is complete, click **Close**.

The new part curves are stored in a node that is called **Part Curve**:



When computing new part curves, existing part curves in other nodes corresponding to the selected options are moved into the new part curve node.

Combining Part Curves

Two or more part curves can be combined into a single part curve if necessary.

To combine the curves:

1. Select two or more edge nodes belonging to the same part, right-click and select **Combine** from the pop-up menu:



The curve nodes are combined into the single node (the first selected):



Note the following:

- Tag properties relating to the secondary curve nodes are lost when they are different from the target (first selected) curve node.
- Feature curve assignments on the secondary curve nodes are lost when they are different from the target curve node.

Splitting Non Contiguous Part Curves

A part can have one or more of its curves split based on the non-contiguous nature of the curves.

This feature splits up the selected curves depending on whether they are connected to their neighbors or not.

To split a curve:

1. Right-click the curve node and select **Split Non-Contiguous** from the pop-up menu:



Any non-contiguous curves are split to form new curve nodes, as shown below:



Any tags belonging to the original curve nodes are not copied to the split nodes. The new curves are also not assigned to any feature curves.

If the selected curves are contiguous, no action is taken.

Repairing Part Surfaces and Curves

Repairs can be made to leaf-level parts at any time before surface initialization. They can be required to close a surface or fix some other non-valid entity that the translation process created.

The tools available to do these repairs are ostensibly the same as the tools provided for surfaces within regions and boundaries under the Representations node. All of these tools are contained within the panels of the Surface Preparation Tool.

The main differences are:

- Surfaces in the display are colored as part-specific items.
- The Part Surface/Part Curve Edit tool replaces the Boundary/ Feature Curve Edit tool.

If the Surface Preparation Tool is activated on a composite part, all of its child nodes are automatically selected for inclusion in the diagnostic process that precedes entry to repair mode.

See Also

Understanding Geometry Parts

Activating the Surface Preparation Tool for Parts

This section describes how to activate the Surface Preparation Tool.

To activate the Surface Preparation Tool:

1. Select the appropriate part nodes (which can include composite parts).
2. Right-click the selection and select **Repair Surface...** from the pop-up menu:



The *Surface Preparation Options* dialog appears:



3. Optional: To select specific parts using their name or other attributes such as tags or metadata, click the **Select** button to the right of **Geometry Parts**, then use the filter options within the *Select Objects* dialog.
4. Click **OK**.

The *edit* tab appears containing the surface preparation tool and its several modes of operation, identified as follows:

- surface diagnostics
- surface repair
- feature curve repair
- face offset tool
- connected surfaces edit
- merge/imprint tool
- face intersection tool

If the part is retessellated, any repairs that are made to part surfaces are lost. The lost repairs include the outcome of any imprint operations where more faces were created at the interface between two parts. Therefore, complete retessellation before making any manual changes with the surface repair tools.

Using Alternative Coloring for Part Surfaces

This section describes the options that are available for coloring part surfaces.

In addition to coloring surfaces by Default Color, Patch and CAD Association, the Display Control panel for surface repairs on parts provides the following alternative options:

- **Part Surface**: faces are colored according to the part surface to which they belong.
- **Part Boundary**: faces are colored according to the boundary to which they belong. This option is valid only if the part has been assigned to a region.
- **Part**: faces are colored according to the part to which they belong.

Moving Part Surfaces and Part Curves

This section describes how to move part surfaces and part curves.

You can use the *Part Surface/Part Curve Edit* panel to move faces into an existing or new part surface, and curves into an existing or new part curve. Details are given in the section, Surface and Curve Edit Controls.

Exporting Parts

This section describes how to export the data that is contained in the geometry parts tree.

Data that is contained in the geometry parts tree can be exported in two ways:

- By exporting individual leaf-level parts
- Using PLMXML export.

The PLMXML method can be used to export all the parts that are contained within a composite part, while retaining the groupings and hierarchy defined in the object tree. This method is useful if you have many parts to export and/or you wish to retain the organizational structure of the parts.

Exporting a Part

This section describes how to export leaf-level parts.

Leaf-level parts can be exported as surface mesh files using any of the following formats:

- pro-STAR database file (.dbs)
- Stereolithography (.stl)
- NASTRAN shells (.nas)
- PATRAN shells (.pat)
- pro-STAR shell input (.cel/.vrt/.inp)
- FELISA Front Surface (.fro)
- Medina Binary Input File (.bif)
- JT Open (.jt)

When exporting 3D-CAD, CAD, and shape parts, the current tessellation that is specified for the surface is used.

3D-CAD parts can also be exported as Parasolid Transmit files from within the 3D-CAD modeler.

To export a part:

1. Right-click the part node and select **Export...** from the pop-up menu:



The *Save* dialog appears. You can only save one part at a time using this method.



2. Specify a directory for the file.
3. In the *File Name* filed, enter a suitable name.
4. From the *Files of Type* field, select the file format.
5. Click **Save**.

Exporting Parts Using PLMXML

The PLMXML export option allows you to export large numbers of parts in a single operation.

The PLMXML export option creates a `.plmxml` file to retain information on how the parts are organized in the tree. When the PLMXML file is imported into another simulation, the parts are automatically imported and organized according to the structure in the PLMXML file.

Before performing PLMXML export, ensure that the geometry parts contained in the tree are organized as necessary. The PLMXML export is only available for composite files. You can export a single composite part or multiple composite parts to a single PLMXML file.

To export the parts:

1. Right-click the composite node and select **Export PLMXML**.



The *Save* dialog appears.



2. From the **Save In** drop-down menu, specify the save directory.
3. From the **Part Type** drop-down menu, select a desired file format for the export:
   - Database (.dbs)
   - Stereolithography (.stl)
   - Nastran (.nas)
   - Patran (.pat)
   - Shell input file (.inp)
   - FELISA Front Surface (.fro)
   - Medina Binary Input (.bif)
   - JT Open (.jt)
4. In the *File Name* field, enter a suitable name.
5. Click **Save**.

All leaf-level parts are exported into a single file of the selected format. You can monitor the export progress in the *Output* window.
The `.plmxml` file is also created containing the relevant structure, names, and file references for the exported parts.

To import the parts into another simulation, follow the procedure for importing XML files.

Note: You cannot import `.xml` and `.bif` files in versions of STAR-CCM+ that are running on Windows operating systems.

Getting the Face Count of a Part

This section describes how to obtain the number of faces of a geometry part.

To obtain the number of faces of a geometry part:

1. Right-click on that part and select **Statistics**.



The *Statistics* dialog appears with a *Face count* property.



2. To check the face count for a particular description, use the drop-down list at the top of the dialog.

If the selected description does not exist for that part, the face count is zero.

This same feature works for part surfaces, part curves, and composite parts.

Editing Imported Parts in 3D CAD

This section describes how to modify your imported geometry parts using 3D-CAD.

You can modify your imported geometry parts in STAR-CCM+ using 3D-CAD. This feature is useful if you do not have access to a commercial CAD package and you want to modify the CAD parts. You can carry out any 3D-CAD operation in context with your geometry. You can also rename faces on your geometry as an alternative to splitting surfaces by patch.

Note: This feature only supports parts that contain CAD data (such as CAD parts and 3D-CAD parts). Shape parts (such as Boxes, Spheres, and Cylinders) and mesh parts are not supported as they contain no CAD data.

Caution: CAD parts that you import into 3D-CAD are converted to 3D-CAD parts. This process is irreversible. This operation does not affect your original CAD geometry files.

To edit a geometry part in 3D-CAD:

1. Right-click the desired part in the **Geometry** > **Parts** node and select **Edit in 3D-CAD**.



Note: You can only edit leaf-level parts in 3D-CAD.

STAR-CCM+ launches 3D-CAD and transfers your geometry over.



A 3D-CAD model is created for parts that you edit in 3D-CAD.



2. In 3D-CAD, perform the modifications that you want.

   In the example, to increase the surface area in contact with the cooling, the height of the cooling fins is modified and more vertical cooling channels are added.



3. Click **Close 3D-CAD** when done.

   Modified parts have new icons to show that they have been modified and require an update.



4. Right-click the modified parts and select **Update**.

Note: Once you edit a part, you cannot select to edit it with any other parts. You can only edit one 3D-CAD part at a time.

Part Operations Reference Page

This section contains a list of available operations that appear when right-clicking a part.

Each operation contains a short description and a link to the relevant section in the documentation.

Update	If the selected part has been modified in 3D-CAD after it has been created, it will be necessary to update it to reflect the changes. This operation is only available for parts that are created in 3D-CAD.	
Highlight	The selected part becomes highlighted in the *Graphics* window.	
Composite	The selected parts are added to a composite (assembly) part. Useful for grouping parts together and for organizing the **Parts** manager node.	
Explode	Explodes the composite part, returning the child parts to their original level in the **Parts** manager node. This operation is only available for composite parts.	
Transform	The following Transformation operations are available; each operation launches its respective dialog:	
	Coordinate Transform...	Maps the selected part from the current coordinate system, to a coordinate system of your choice. By using this operation, you can translate and rotate a part at the same time while maintaining the exact shape of the part.
	Scale...	Allows you to modify the size of the selected part. Scaling can take place in any or all of the three coordinate directions and can be

		used to increase or decrease the size of a part. Scaling a part is useful if it uses units different to other parts you have already imported.
	Translate...	Allows you to move the selected part in any of the three coordinate directions. Translating is useful if different parts are built in different coordinate systems.
	Rotate...	Allows you to rotate the selected part) about a specified axis vector.
	Reflect...	Allows you to mirror the selected part in a specified axis. This option can be useful when the original part is a half model of the geometry. You can create the full model by first duplicating the part and then reflecting the copy (or original).
Duplicate	Produces an identical copy of the selected part with the same properties as the original, other than region and contacts specification, which are left null.	
Combine	Combines the selected parts to form a single mesh part. The first selected part is the one that is used as the base for all the other parts. All part surfaces and curves belonging to each part are grouped in the base part. This operation is irreversible.	
Boolean	The following Boolean operations are available:	

	Unite Parts	Joins all the selected parts into one single part. This operation differs from the combine operation, which combines the selected parts into one part, but retains the part surfaces as separate, whole surfaces. This operation consumes all selected parts.
	Intersect Parts	Finds the common volume that the selected parts share and creates a single part that is based on this volume. If the parts do not share a common volume, the part that the Boolean operation creates contains no surface data. This operation consumes all selected parts.
	Subtract Parts	Subtracts the defined shape from the volume of a target part. Any number of parts can define the shape. This operation consumes all selected parts.
Split Non-Contiguous	Splits non-contiguous mesh parts (that is, parts that have faces that are not connected to other faces in the part). Each part can be split into its individual component parts.	
Split by Surface Topology	Can be used to split and combine parts that can be defined as distinct geometrical regions. When this technique is used for multiple parts, STAR-CCM+ internally combines all the parts and treats them as a single part. The option is	

	available to combine parts with shared part surfaces.	
Replace Part...	This operation replaces a 3D-CAD, CAD, or mesh part in order to provide new geometrical input for the meshing pipeline (without affecting any of the other data input required to generate a mesh). Only one part node can be replaced at a time. The operation does not apply to shape parts.	
Create Mesh Operation	Mesh operations provide a way to define certain actions on geometry parts once, and to repeat those actions easily whenever their input parts or properties are changed. The result of a mesh operation is a new geometry part, but the input parts are not removed in the process; they remain intact as before. The following operations are available:	
	Extract Volume	This operation finds enclosed volumes within the selected part and create a part to fill the volume. This operation is similar to the split by surface topology action. The key difference that the extract volume operation does not remove the input parts.
	Fill Holes	This operation can be used to close any arbitrary shaped hole using a planar surface. This option is similar to the hole filling tool available for feature curves.
	Unite Parts	This operation performs a boolean unite on all selected parts that are in contact with each

		other. Bodies that overlap can also be united. This operation is similar to the boolean unite action that can be performed on geometry parts. The key difference that the mesh operation does not remove the input parts.
	Intersect Parts	This operation performs a boolean intersect on the selected parts. This operation is similar to the intersect parts action that can be performed manually on geometry parts. The key difference that the intersect parts operation does not remove the input parts.
	Subtract Parts	In a subtract operation, the volume of one or more geometry parts is removed from the target geometry part. This operation is similar to the boolean subtract action that can be performed manually on geometry parts. The key difference that the input parts are not removed for the subtract operation.
Repair Surface...	Activates the Surface Preparation tool where the part can be cleaned up in preparation for CFD analysis.	

Edit in 3D-CAD	Transfers the CAD part to 3D-CAD where you can work with it. Only parts with CAD data can be transferred to 3D-CAD.
Re-tessellate...	Modifies the tessellation of the selected part. The surfaces of a geometry part are represented using triangles which come from a process that is known as tessellation. Selecting this operation launches the *Tessellation options* dialog.
Assign Parts to Regions...	Launches the *Assign Parts to Regions* dialog, where the selected part can be added to a new or existing region. The region, boundaries, and feature curves that result from this operation can all be controlled in the dialog.
Update Interfaces	Checks for any changes to the part interfaces on the selected part and update them if necessary. Invalid or obsolete interfaces are removed.
Set Region	Allows you to add the selected part to an existing region. The available regions appear in the drop-down list.
Apply Tag	Applies a created tag to the selected part. You can apply more than one tag to each part, whether a leaf-level or composite, if necessary.
Statistics...	Displays the number of cell faces, in the initial mesh representation, that make up the selected part.
Find Part/Part Contacts...	Automatically finds part contacts between the selected parts (two or more parts) based on a specified tolerance value. This operation is useful in circumstances where several parts exist, some of which can be touching two or more other parts.
Browse Part/Part Contacts...	Allows you to scan through the part contact data using the *Part/Part Contact Browser* dialog, and view the contacting parts in the *Graphics* window. This operation automatically updates the view to hide parts that are not part of a contact list.
Create Part/Part Contact	Allows you to create part contacts manually between the selected parts (two or more parts). This operation is useful when the parts have

	surfaces that are in contact with one another and require an interface for the physical coupling once the mesh has been built.
Export...	Exports the selected part as a surface mesh file. The supported formats are listed below: • pro-STAR database file (.dbs) • Stereolithography (.stl) • NASTRAN shells (.nas) • PATRAN shells (.pat) • pro-STAR shell input (.cel/.vrt/.inp) • FELISA Front Surface (.fro)
Rename...	Allows you to rename the selected part.
Delete	Deletes the selected part from the **Parts** manager node.
Edit...	Opens the table editor, which combines the functionality of the simulation tree and *Properties* window. It displays the selected part, all of its child nodes, and their properties in a single dialog.

Preparing Surfaces

The *Surface Preparation* tool provides diagnostic checks for assessing the validity of surfaces, and a range of tools by which surfaces can be repaired.

It is possible that the surfaces of your geometry are in a suitable condition for surface meshing. Among other things, there can be faces on the surface that intersect other faces, or edges that have two or more shared faces. It is good practice to check the validity of the surfaces on your geometry before you attempt to generate a surface mesh.

Other types of preparation can be required on the surfaces of your geometry. For example, if you have multiple parts that are to form multiple regions in the computational domain, you may have to imprint the surface of one part on the surface of another part at the interface between them. You can also imprint surfaces through the *Surface Preparation* tool. The different modes that are provided in the *Surface Preparation* tool are summarized in the table below.

Repair Modes

Mode	Description	Diagnostics
Repair Surface Mesh	Provides a range of surface repair actions to correct errors discovered during the diagnostic check. See section, Using the Repair Surface Mode.	A range of checks are provided to assess the validity of the surface. See section, Using Surface Diagnostics.
Repair Features	Provides actions with which to modify the definition of feature curves. See section, Using the Repair Features Mode.	Checks are provided to detect errors in feature curves. See section, Using Feature Diagnostics.

Mode	Description	Diagnostics
Offset Faces/Edges	Lets you inflate surfaces such as thin baffles, and lets you translate surfaces by a specified distance. See section, Using the Offset Faces/Edges Mode.	None
Edit Connected Surfaces	Identifies surfaces that are not connected to the rest of the model; these surfaces can be deleted. See section, Using the Edit Connected Surfaces Mode.	Check for disconnected surfaces
Merge/Imprint	Provides imprinting tools for surfaces between adjacent parts. See section, Using the Merge and Imprint Mode.	Multi-part imprint can automatically find overlapping pairs of surfaces
Intersect/Boolean Faces	Finds the common surface from sets of intersecting faces. See section, Using the Intersect/Boolean Faces Mode.	None

When you have launched the *Surface Preparation* tool, you can switch between these modes using the **Mode** drop-down menu that is shown below.



The repair tools provide the following capabilities:

- Repair surface mesh
  - Diagnose and locate surface issues
  - Delete and create faces and edges
  - Collapse and smooth vertices
  - Fill and patch holes
  - Split and swap edges
  - Zip edges
  - Remesh faces locally, with or without feature retention
  - Project and transform vertices
  - Merge duplicate faces and close vertices
  - Automatic repair to fix issues relating to poor quality, close proximity, and pierced faces
- Repair features
  - Flag and unflag feature curves
  - Fix feature curves
  - Mark feature curves
- Offset faces/edges
  - Extrude edges to create a two-dimensional surface
  - Inflate thin shell and baffle surfaces with a thickness
  - Translate surfaces
- Isolate connected surfaces
- Merge and imprint surfaces
- Intersect or boolean faces

The tools in the *Surface Preparation* tool can operate on either geometry parts or surface representations. If there are two or more surface representations, fixing a surface only applies to the chosen representation.

After launching the *Surface Preparation* tool, the *Surface Repair Options* dialog appears, in which you select the parts or regions that will form the active model during the surface preparation session. You can change the parts or regions in the active model during the session by clicking the **Reset Parts...** button.

In addition to the actions provided by the various surface repair modes, be aware of some auxiliary controls that are provided across all modes. These controls are listed as follows:

- Selection controls
- Display controls
- Surface and curve edit options
- Set controls
- Undo and redo
- Keyboard shortcuts

Some of these controls appear in their own panel within the dialog. You can change the order in which these panels are arranged.

See Also

Using the Properties Window
Changing The Position Of Panels
Changing The Parts In The Active Model
What Are Representations?
Understanding Geometry Parts
Feature Diagnostics
Using Threshold Diagnostics

Launching The Surface Preparation Tool

The *Surface Preparation tool* can be launched in different ways.

To launch the *Surface Preparation tool*:

1. Click (**Launch Surface Repair**) in the **Mesh Generation** toolbar.
   - If you have geometry parts, but no representations, the tool automatically launches with the geometry parts available for repair.
   - If you have one representation, but no geometry parts, the tool automatically launches with the representation available for repair.
   - In all other cases, you can choose either **Repair Part(s)...** or **Repair Region(s)** from the pop-up menu.



Choosing to repair parts takes geometry parts as the input for the active model. If you choose to repair regions, also select the correct surface representation to work with from the submenu.



These representations are contained within the **Representations** manager node of the simulation tree.



2. Alternatively, right-click a geometry part or a surface representation and select **Repair Surface...**:



For both methods, the *Surface Repair Options* dialog appears with **Geometry Parts** or **Regions** displayed, depending on how you launched the tool:



In the *Surface Repair Options* dialog, to include only the **Geometry Parts** that are specified in the active model when surface preparation is launched, click **OK**. To change the selection of parts or regions that are specified, click **Select**.



In the *Select Objects* dialog, click the **Select Part** button to the left of each part that you require, then click **OK**.

If you launch the tool from the toolbar, the parts or regions that are currently displayed in the *Graphics* window are automatically selected in the dialog. You can control which parts appear in the *Graphics* window by either adding or removing from a geometry displayer, or by showing or hiding parts. If there is only one part or region in the simulation, it is automatically selected, regardless of whether it is displayed in the *Graphics* window or not.

If you launch the tool by right-clicking in the simulation tree, the parts that are selected in the tree are automatically selected in the dialog. You can select multiple parts by using the `<Ctrl>` `<Click>` method.

Use the **Faces to show on startup** option to choose what faces are shown in the *Graphics* window when the dialog opens. The three options are:

3. **No Faces** -- no faces are displayed on startup; you can select these using display controls, or by clicking the categories of errors that are discovered during any diagnostics check. This option can be useful for large models where plotting is slow.
4. **All Faces** -- all faces in the model are displayed.
5. **Currently Visible Faces** -- only the faces that are shown in the current scene are displayed.

After the parts for the active model have been selected, STAR-CCM+ will display the *Surface Preparation* in-place dialog. It is initially set to mode **Repair Surface Mesh**.


Servers
Star 1
simulation
scene/plot
edit
Mode
Repair Surface Mesh
Surface Diagnostics
Input
Active Model
[]
#
Display
Threshold
Count
Manage Thresholds...
Reset View?
Reset Displayed?
Surface Repair
Selection Control
Display Control
Part Surface/Part Curve Edit
Set Control
Reset Parts...
Close
Help

If you intend to check the validity of the surface and repair any errors, then your next step is to run the diagnostics checks for the surface repair mode.

Note: To prevent accidental closure of the *Surface Preparation* in-place dialog when repair operations are in progress, STAR-CCM+ temporarily deactivates the **Close** button. The **Close** button is available again when the repair operations are complete or after you click **(abort)**.

Face Orientation

When you launch the *Surface Preparation* tool, STAR-CCM+ does a topological orientation on the faces you selected. That is, it makes sure that all the faces are consistently oriented in the same direction, whether pointing inwards or outwards. STAR-CCM+ picks the orientation belonging to the largest number of faces and orients them all in that direction.

See Also

Display Controls
Diagnostics
Using Threshold Diagnostics

Changing The Parts In The Active Model

The active model is the parts or regions that are available for repair during the surface preparation session.

These parts or regions are selected in the *Surface Repair Options* dialog after launching the *Surface Preparation* tool. Clicking the **Reset Parts...** or **Reset Regions...** button changes the parts or regions in the active model during the surface preparation session. This feature allows you to simplify a complex geometry by working with smaller groups of parts or regions for diagnostics and repair.

1. To change the parts or regions in the active model, click the

   Reset Parts... **Reset Parts...** or **Reset Regions...** button at the bottom of the *Surface Preparation* tool.

   This button opens the *Surface Repair Options* dialog which has similar options to the dialog that appears when you launch the *Surface Preparation* tool. The only exception is the **Faces to show on startup** menu, which allows you to display **All Faces** or **No Faces** only. The **Currently Visible Faces** option is not available, as changing the parts means that the currently visible faces are unlikely to be in the

new active model. The default **Geometry Description** is set to the description that is currently being repaired.



2. You can change the parts or regions in the active model by clicking the **Select** button. In the *Select Objects* dialog, to select or deselect a part or region, click the button to the left of that item as shown:



Changing the parts in the active model makes the following changes to the surface preparation session:

- All face, edge and/or vertex selections are deselected.

- The surface diagnostics and feature diagnostic results are deselected and the *Diagnostics* group box returns to its default state. However, the state of the diagnostic checkboxes is maintained.
- The Edit Connected Surfaces diagnostics are deselected.
- All user-defined sets are deleted and entities in the default set are deselected.
- The Undo/Redo stacks are deselected.
- The single and multi-parts Merge/Imprint pairs lists are deselected.
- Faces selected for the Intersect/Boolean Faces Tool are deselected.
- Faces selected for Find Leak Path are deselected.
- The list of part surfaces and part curves available in the Part Surface/Part Curve Edit group box is automatically updated to match the new parts in the active model. The same is true for regions, regarding boundaries and feature curves.

When projecting vertices, the *Project Vertices* tool temporarily replaces the *Surface Preparation* tool. Therefore, you are unable to change the parts in the active model while projecting vertices, as the **Reset Parts...** button is unavailable.

See Also

Using The Leak Path Mode
Intersect Boolean Faces Mode
Undo And Redo
Set Controls
Feature Diagnostics
Using Threshold Diagnostics
Edit Connected Surfaces Mode
Diagnostics
Surface, Contact, and Curve Edit Controls
Merge And Imprint Mode
Projecting Vertices
Preparing Surfaces

Changing The Position Of Panels

The *Surface Preparation tool* contains a number of subpanels for the various groups of controls that are provided.

You can change the order of these panels using a drag-and-drop action as follows:

1. Place the mouse pointer over the title bar of the panel you wish to move.



2. Click and hold the left mouse button. The icon changes as shown below.



3. Move the mouse pointer to the new location within the panel. The icon changes to indicate that the pointer is hovering over a valid location.



4. Release the left mouse button. The panel is moved to the new location.



The sequence of panels, and their state (whether open or collapsed) is set uniquely for each mode in the *Surface Preparation* tool. When you close the tool, the sequence and states are saved for all modes,

and are reapplied whenever you open the tool in subsequent sessions of STAR-CCM+.

See Also

Preparing Surfaces

Diagnostics

The *Diagnostics* group box in the **Repair Surface Mesh** and **Repair Features** modes tool provides a summary of the surface errors in the active model.

Shown below are the *Surface Diagnostics* and *Feature Diagnostics*. These lists are displayed after you run the diagnostics.

Query diagnostics are also available, which provide information about selected vertices, edges and faces. Query diagnostics differ from surface and feature diagnostics in that they do not report errors.

Surface Diagnostics



Feature Diagnostics



The colored boxes in the *Diagnostics* group box report the number of times that a category of error is found in the active model. The same colors are used in the active scene to highlight the faces and edges on which the errors were found. Clicking a box displays the relevant faces, edges, and vertices in the active scene, so that you can identify the location and extent of an error.

As errors are fixed, the number in each category is updated to indicate the remaining number of items for which that category of error is found. It is common for the number of one error type to increase while fixing another error type. For example, deleting faces to remove a pierced surface temporarily increases the number of free edges.

Browsing Errors

The **Browse** buttons display and select the faces, edges, and vertices that are associated with each category of error in the diagnostics list. The selected items are displayed in the active scene that is shown in the *Graphics* window. You can use these buttons to browse through groups of problem items and decide how best to fix them.

View Options

While browsing, two options control the viewing of errors:

- **Reset View** resets the view to zoom in on the displayed faces, edges, and vertices

  If you deactivate the **Reset View** option, the view remains the same
- **Reset Displayed** resets the displayed faces, edges, and vertices to the error and two neighbors

  If you deactivate the **Reset Displayed** option, the displayed entities remain the same.

See Also

Changing The Parts In The Active Model

What Are Pierced Faces

A pierced face is a face that is intersected by one or more edges of another face.



If an edge pierces a face, the surface diagnostics color the face red in the *Graphics* window. If an edge is piercing a face, the surface diagnostics color the edge red.

An understanding of how the calculation of self-intersection works is required when you are diagnosing multiple parts.

Fixing Pierced Faces

You can fix pierced faces by deleting one set of erroneous faces, using the surface repair tools. This operation typically leaves free edges, which you must also fix.

See Also

Repair Surface Mode
Repair Surface Mode
What Are Free Edges

What Is Face Quality

Face quality is a measure of similarity between a face and the ideal face shape, which is an equilateral triangle.

The surface diagnostics calculate face quality as the ratio of the in-circle radius to the circum-circle radius, multiplied by a factor of 2. The in-circle of a face is the largest circle that fits inside the face, while the circum-circle is the circle that passes through the three points of the face. The diagram below illustrates these definitions for an ideal face.



Face quality ranges from 0 to 1, with 0 being a line (that is, no in-circle radius) and 1 being an equilateral triangle (that is, in-circle radius is half of circum-circle radius).

In **Surface Repair** mode, activation of the **Poor Quality Face** diagnostic marks poor quality faces (below 0.01 by default) in yellow.

Fixing Poor Quality Faces

There are instances when you do not need to fix poor quality faces, depending on how the surface is used. For example, if you use the surface with the surface wrapper or surface remesher, poor quality faces do not usually cause a problem.

However, if you use the surface directly for volume meshing, poor quality faces lead to poor quality or invalid volume cells. In this case, you can improve face quality by locally remeshing the poor quality faces, using the surface repair tools.

See Also

Surface Remesher
Surface Wrapper
Repair Surface Mode
Using The Face Quality Selection Mode

What Is Face Proximity

Face proximity is the ratio of the distance between a face and its nearest neighbor, to the size of the face.

Face proximity ranges from 0 to 1, with low values indicating one of the following problems:



- Folds in the surface, for example, “V” or “Z” shaped kinks in the surface causing localized overlaps
- Surfaces in close proximity, for example, two surfaces that are almost touching, such as either side of a thin baffle

The following illustration shows the first type of proximity issue:

In **Surface Repair** mode, activation of the **Close Proximity Faces** diagnostic marks close proximity faces (below 0.05 by default) in orange.

An understanding of how the calculation of proximity works is required when you are diagnosing multiple parts.

Fixing Close Proximity Faces

Depending on the nature of the problem, you can fix close proximity faces by moving the faces further apart, or by smoothing the mesh locally, using the surface repair tools.

There are instances when you do not need to fix close proximity faces. For example, when you expect close proximity faces due to the nature of the geometry. However, pay attention to these areas during surface and volume meshing, in case mesh quality problems arise.

See Also

Repair Surface Mode

What Are Free Edges

A free edge is an edge joined to only one face. Two joined faces share two vertices for the common edge.

For example, edges around a surface hole are marked as free edges, as is illustrated below:



The surface diagnostics color **Free Edges** green in the *Graphics* window.

Double-clicking a free edge selects all connected free edges. This feature is useful when fixing free edges using the fill holes tool.

An understanding of how the free edges are calculated is required when you are diagnosing multiple parts.

Fixing Free Edges

You can fix free edges by filling or patching the hole using the surface repair tools.

There are instances when you do not need to fix free edges. For example, a free-floating baffle has the outside edges marked as free edges, until you convert the baffle to an interface. As long as you are aware of the error, and convert the baffle to an interface before surface and volume meshing, you can ignore the free edges.

See Also

Converting Boundaries to Interfaces

Filling Holes
Launching The Surface Preparation Tool
Entity Selection

What Are Non Manifold Edges

A non-manifold edge is an edge that two or more other faces share.

For example, an interior surface joined to an exterior surface has non-manifold edges where the interior and exterior faces share an edge, as illustrated below:



The surface diagnostics color **Non-manifold Edges** blue in the *Graphics* window.

An understanding of how non-manifold edges are calculated is required when you are diagnosing multiple parts.

Fixing Non-Manifold Edges

You can fix non-manifold edges by deleting the faces that are causing the additional edges, using the surface repair tools.

There are instances when you do not need to fix non-manifold edges. For example, a baffle joined to the surface of a body has the shared edges marked as non-manifold edges, until you convert the baffle to an interface. As long as you are aware of the error, and convert the baffle to an interface before surface and volume meshing, you can ignore the non-manifold edges.

See Also

Repair Surface Mode
Converting Boundaries to Interfaces

What Are Non Manifold Vertices

A vertex is non-manifold if any of its connected faces have no other link with the rest of its connected faces, apart from the vertex.

For example, the diagnostics identify a non-manifold vertex when one surface joins another at a single vertex, as illustrated below.



The surface diagnostics color **Non-manifold Vertices** blue in the *Graphics* window.

Fixing Non-Manifold Vertices

If the geometry requires connected surfaces, you can fix non-manifold vertices by filling holes or zipping edges, using the surface repair tools.

If the geometry requires disconnected surfaces, you can fix non-manifold vertices by disconnecting the faces that are causing the non-manifold vertex, using the surface repair tools.

See Also

Repair Surface Mode

Understanding Surface Mesh Diagnostics For Multi Part Cases

Whether the surface belongs to single or multiple parts affects the diagnosis of non-manifold edges, proximity and self-intersection.

For example, if one of these errors is on an interface between parts, diagnosing a single part does not report the error. However, the diagnostics tool calculates free edges on a per-part basis, which is independent of the number of parts that are selected for diagnostics. See Launching the Surface Preparation Tool.

The main consideration in multiple parts cases is whether an interface exists between two neighboring parts. This section aims to provide an overview of what to expect when diagnosing multiple part cases, and how to recognize and fix surface errors. Some surface errors, such as free edges, do not affect the surface remesher. However, having any form of non-closed or non-manifold surface causes the volume meshers to produce an error message and does not generate a mesh.

This section uses a simple example, two cube surfaces, to highlight the different scenarios possible when dealing with multiple parts or parts. The diagram below illustrates the terminology:



The examples below focus on whether each cube is closed (that is, has a surface at the interface plane) and whether an interface or contact exists between the two cubes. Additionally, this section covers an example where two parts or parts share edges, due to merged vertices.

It is important to note, the surface diagnostics tool identifies errors regarding the active model and not on a per-part basis, except for free edges as noted earlier. Therefore, select all required parts in the Surface Repair Options dialog before staring surface repair. Otherwise, the diagnostics tool does not identify certain errors. See Diagnostics.

Two Connected Boundaries with No Interface

The diagram below illustrates this situation:

This situation is the simplest, where the two cubes are closed and no interface exists between the boundaries at the interface plane. The surface diagnostics do not report any free edges or non-manifold edges in this instance. However, the diagnostics tool identifies interface boundaries as being close and piercing one another. Attempting to volume mesh in this situation results in an error message reporting self-intersecting surfaces (providing you are not using per-part meshing).

To fix these surface errors, create an interface between the two boundaries at the interface plane. This fix allows the volume mesher to proceed and creates a conformal mesh between the two parts, if supported by the meshing model. See Creating a Boundary Interface.

Two Connected Boundaries with an Interface

The following diagram illustrates this situation:



This situation is the same as the previous situation, but with an in-place interface created between the boundaries at the interface plane. Surface diagnostics report no issues in this instance, and the meshers are able to create a volume mesh. This situation is the ideal situation for using the STAR-CCM+ meshing pipeline.

One Surface with No Interface

The diagram below illustrates this situation:

This section considers two situations. The first situation uses cubes without shared edges, that is, two disconnected cubes. The second situation uses cubes with shared edges, that is, two connected cubes.

In the first situation, the diagnostics tool marks the edges on the part containing the open surface as free edges. Attempting to volume mesh in this situation results in an error message reporting one part as open. To fix this error, close the open part by filling the hole, using the surface repair tools. Then create an interface between the two boundaries, which provides the ideal situation that is described above. The meshers are then able to create a volume mesh.

In the second situation, the diagnostics tool marks the edges on one part as free edges, since the surface only belongs to one part. Attempting to volume mesh in this situation produces the same error as above, reporting one part as open. To fix this error, close the open part by filling the hole, using the surface repair tools. See Repair Surface Mode.

However, this fix appears to swap one problem, free edges, for another, non-manifold edges. Therefore, create an interface between the two boundaries, which resolves the non-manifold edges, and allow volume meshing.

No Interface Boundaries, with No Interface/ Between Parts

The diagram below illustrates this situation:

In this situation, both cubes do not have an interface, and the surface diagnostics marks the open ends of both parts as free edges.

To fix this error, close each part separately by starting surface repair with only one part that is selected as the active model. Then create an interface between the two parts, which allows the volume mesher to run.

See Also

Creating a Boundary Interface
Repair Surface Mode
What Are Free Edges
What Are Pierced Faces
What Are Non Manifold Edges
What Is Face Proximity

Query Diagnostics

Query diagnostics provide information on the selected entity or entities in the *Graphics* window.

Currently, the following information is available:

- Shortest distances between entities
- Vertex coordinates of a selected vertex
- Edge dihedral angle
- Face area
- Face quality
- Face centroid
- Geometric range of selected entities
- List of part and part surfaces for the selected faces

Note: Query diagnostics are also available for region repair, although the output differs. For example, the Part and Part Surface list becomes the Region and Boundary list.

The query diagnostics button is located in the surface diagnostics group box.

The availability of the button depends on the type or combination of selected entities. Upon launching surface repair, the query diagnostics button is not available, as no entities are selected. Right-clicking the button when no entities are selected shows all queries that the feature can perform.



When you select a query diagnostic from the list, the button icon changes to represent the chosen option. The query diagnostics button only becomes active when the appropriate selections are made in the *Graphics* window.

If you have pre-selected entities, right-clicking the query diagnostics button only allows you to select options from the list that are appropriate for the current selection. For example, consider the selection that is shown below, which includes a group of vertices.



In this case, you can only select the **Geometric range of the selected vertices** option as it is the only query available when more than 2 vertices are selected.



If you have a mixed selection of entities (that is, a combination of faces, edges and vertices), the queries for individual types of entities are still available for selection, as well as the mixed options. In this case, any selections that do not apply to the chosen query are ignored when running the query.

Performing a query prints information to the *Output* window.



The units that are used in the output are the preferred units for the given quantity in the simulation. The preferred units for length and area can be different (for example, length in millimeters and area in square meters).

What Are Vertex Query Diagnostics

Vertex query diagnostics provide information specific to the vertex or vertices that are selected in the *Graphics* window.

To access the vertex query diagnostics, right-click the Query Diagnostics button and select the **Vertex** submenu. The available options are shown below:



Vertex Location in Lab or Local System

This query reports the coordinates and ID for the selected vertex/vertices. This query is available when you select a single vertex or multiple vertices.

By default, the location is defined relative to the laboratory coordinate system. An example report for a single vertex selection is provided below:



When selecting multiple vertices, the query reports the list of coordinates and ID's of the selected vertices. There is no correlation between the *Output* list order and the vertices selection order.

You can also choose to report the vertex/vertices location relative to a local coordinate system by setting the coordinate system of the view. One method of changing the view coordinate system is to click the **Save-Restore-Select views** button and select **View Coordinate System**, followed by your choice of system:



In this case, the query reports the vertex/vertices location relative to the selected local coordinate system and the laboratory coordinate system.



Distance Between Two Selected Vertices

This query reports the distance between two vertices. It also reports the vertex IDs, locations, and the distance between the selected vertices in the x-, y-, and z-direction. By default these components are displayed relative to the laboratory coordinate system. However, if the scene view is set to a local coordinate system, the components are displayed in the output window relative to both coordinate systems. This query is available when you select only two vertices.

An example report is provided below:

```
Output - Star 1
Vertex Distance Report
 Distance = 0.339408687794 (m)
Laboratory (m):
 Vertex 3 : [-0.0563497297057 -0.03 0]
 Vertex 8 : [0.06 0.0779860997883 0.3]
 Components (dx dy dz): (0.116349729706 0.107986099788 0.3)
Cartesian 1 (m):
 Vertex 3 : [0.224019789825 0.00728642763607 0.157484804765]
 Vertex 8 : [-0.0611299627533 0.0256098352269 -0.0256875967843]
 Components (dx dy dz): (-0.285149752578 0.0183234075908 -0.183172401549)
```

Geometric Range of the Selected Vertices

This query reports the minimum and maximum coordinate values (relative to the laboratory coordinate system) of a bounding-box that contains all selected vertices. It also reports the dimensions (change in coordinate value) of the bounding-box in the x-, y-, and z-direction. Select two or more vertices before this query becomes available.

An example report is provided below:

```
Output - Star 1
Vertex Geometric Range Report (m)
 Min. Global Value : [-0.155686 0.0982237 -0.120845]
 Max. Global Value : [-0.135411 0.1397 -0.0376828]
 Components : dx = 0.0202747
              dy = 0.0414763
              dz = 0.0831617
```

See Also

What Is a View?
Setting the Coordinate System of the View
Vertex Queries

What Are Edge Query Diagnostics

Edge query diagnostics provide information specific to the edges selected in the *Graphics* window.

To access the edge query diagnostics, right-click the Query Diagnostics button and select the **Edge** submenu. The available options are shown below:



Length of the Selected Edges

This query reports the length or total length of an edge or edges. For multiple edges, it also reports the minimum, maximum and average edge length. Select at least one edge before this query becomes available.

If only one edge is selected, the report appears as shown in the example below:



For multiple edges, the report appears as follows:

```
Output - Star 1
Edge Length Report (m)
 Total Edge Length = 0.130005
 Min. Edge Length  = 9.72937e-005
 Max. Edge Length  = 0.0278163
 Ave. Edge Length  = 0.014445
```

Dihedral Angle of the Selected Edges

This query reports the dihedral angle of an edge, which is the angle between its adjacent faces. For multiple edges, it reports the minimum, maximum and average dihedral angle, and the number of valid edges. Edges are considered invalid if they are free or non-manifold, as a dihedral angle cannot be calculated for such edges. Select at least one edge before this query becomes available.

If only one edge is selected, the report appears as shown in the example below:

```
Output - Star 1
Dihedral angle of edge 3175 is 0.0653701 (radian)
```

For multiple edges, the report appears as follows:

```
Output - Star 1
Edge Dihedral Angle Report (radian)
 No. of valid edges = 5
 Min. Angle = 1.49012e-008
 Max. Angle = 3.0761
 Ave. Angle = 0.637921
```

Geometric Range of the Selected Edges

This query reports the minimum and maximum coordinate values (relative to the laboratory coordinate system) of a bounding-box that contains all selected edges. It also reports the dimensions (change in coordinate value) of the bounding-box in the x-, y-, and z-direction. Select at least one edge before this query becomes available.

An example report is provided below:

```
Output - Star 1
Edge Geometric Range Report (m)
 Min. Global Value : [0.0093659 -0.03175 -0.0404585]
 Max. Global Value : [0.1905 0.243205 0.0345574]
 Components : dx = 0.181134
              dy = 0.274955
              dz = 0.0750159
```

See Also

What Are Non Manifold Edges
What Are Free Edges
Edge Queries

What Are Face Query Diagnostics

Face query diagnostics provide information specific to the faces selected in the *Graphics* window.

To access the face query diagnostics, right-click the **Query Diagnostics** button and select the **Face** submenu. The available options are shown below:



Area of the Selected Faces

This query reports the area a face or group of faces. For multiple faces, it also reports the minimum, maximum and average face area. Select at least one face before this query becomes available.

If only one face is selected, the report appears as shown in the example below:



For multiple faces, the report appears as follows:

```
Output - Star 1
Face Area Report (m^2)
 Total Face Area = 0.0887778
 Min. Face Area  = 3.11258e-009
 Max. Face Area  = 0.00576306
 Ave. Face Area  = 1.99995e-005
```

Centroid of the Selected Face

This query reports the location of the face centroid relative to the laboratory coordinate system. It also reports the face ID. This query is available when you select only one face.

An example report is provided below:

```
Output - Star 1
Face centroid location for face 19874 is (0.00137158 0.003175 -0.0175941) (m)
```

Normal of the Selected Faces

This query reports the face normal for a single selected face. This query is available when you select only one face.

The report appears as shown in the example below:

```
Output - Star 1
Face Normal Report (face 337)
 Face normal : [-0.000150901232119 0.000577854633815 0] (m^2)
 Normalized : [-0.252667304981 0.967553219722 0]
```

Quality of the Selected Faces

This query reports the face quality of the selected faces. For multiple faces, it also reports the minimum, maximum and average face quality. Select at least one face before this query becomes available.

If only one face is selected, the report appears as shown in the example below:

```
Output - Star 1
Face quality for face 44206 is 0.972888
```

For multiple faces, the report appears as follows:

```
Output - Star 1
Face Quality Report
 Min. Face Quality = 0.40796
 Max. Face Quality = 1
 Ave. Face Quality = 0.940224
```

Geometric Range of the Selected Faces

This query reports the minimum and maximum coordinate values (relative to the laboratory coordinate system) of a bounding-box that contains all selected faces. It also reports the dimensions (change in coordinate value) of the bounding-box in the x-, y-, and z-direction. Select at least one face before this query becomes available.

An example report is provided below:

```
Output - Star 1
Face Geometric Range Report (m)
 Min. Global Value : [-0.060325 0.003175 -0.047625]
 Max. Global Value : [0.060325 0.250825 0.0349248]
 Components : dx = 0.12065
              dy = 0.24765
              dz = 0.0825498
```

Part and Part Surface List for Selected Faces

This query reports the associated part and part surfaces for the selected faces. It also reports the number of selected faces on a given part surface from the total. To activate this query, select at least one face.

The following image shows an example of this report:

```
Output - sweepLoftCase
Part and Part Surface List Report
 Part: Main body -> Part Surface: "Loft" : (41 faces from 1015 selected)
 Part: Main body -> Part Surface: "Pipe" : (2496 faces from 2496 selected)
 Part: Main body -> Part Surface: "Outlet" : (9 faces from 14 selected)
 Part: Sleeve -> Part Surface: "Default" : (36 faces from 376 selected)
 Part: Body 4 -> Part Surface: "Default" : (1 face from 12 selected)
 Part: Body 5 -> Part Surface: "Default" : (2 faces from 12 selected)
```

If any of selected faces belong to a part contact, the query lists the part surface contacts at the end of the report:

```
Output - imprinted3DCADTestCase1
Part and Part Surface List Report
 Part: Body 1 -> Part Surface: "BoxSurface" : (4 faces selected from 10)
 Part: Body 1 -> Part Surface: "BoxContactFace" : (121 faces selected from 121)
 Part: Body 2 -> Part Surface: "TubeInlet" : (63 faces selected from 63)
 Part: Body 1 -> Part Surface Contact: "Box" : (1 face selected from 115)
 Part: Body 2 -> Part Surface Contact: "Tube" : (1 face selected from 115)
```

See Also

What Is Face Quality
Face Queries

What Are Mixed Query Diagnostics

Mixed query diagnostics provide information specific to the combination of entities that are selected in the *Graphics* window.

To access the mixed query diagnostics, right-click the Query Diagnostics button and select the **Mixed** submenu. The available options are shown below:



Distance Between Selected Vertex and Edges

This query finds which of the selected edges is closest to the selected vertex and reports the shortest distance between them. It also reports the closest edge ID, the vertex ID and location, and the maximum distance from the vertex to the edges in the x-, y-, and z-direction,

relative to the laboratory coordinate system. This query becomes available when you select only one vertex and at least one edge.

An example report is provided below:

```
Output - Star 1
Vertex-Edge Distance Report (m)
 Shortest distance = 0.0143231
 Source Vertex 21672 : [-0.138287 0.0417277 -0.0405851]
 Closest Edge 3857
 dx = 0.00102346
 dy = -0.014132
 dz = 0.00209533
```

Distance Between Selected Vertex and Faces

This query finds which of the selected faces is closest to the selected vertex and reports the shortest distance between them. It also reports the closest face ID, the vertex ID and location, and the maximum distance from the vertex to the faces in the x-, y-, and z-direction, relative to the laboratory coordinate system. The shortest distance is measured between the vertex and the closest point on the closest face (that is, not necessarily the closest vertex). Select only one vertex and at least one face before this query becomes available.

An example report is provided below:

```
Output - Star 1
Vertex-Face Distance Report (m)
 Shortest distance = 0.0471872
 Source Vertex 15183 : [-0.0196444 0.0310554 0.0333477]
 Closest Face 3774
 dx = -0.0406806
 dy = 0.00704464
 dz = -0.0228494
```

Distance Between Selected Edge and Faces

This query finds which of the selected faces is closest to the selected edge and reports the shortest distance between them. It also reports the closest face ID, the closest edge vertex ID and location, and the maximum distance from the edge to the faces in the x-, y-, and z-direction, relative to the laboratory coordinate system. The shortest distance is measured between the closest vertex on the edge and the closest point on the closest face (that is, not necessarily the closest face vertex). Select only one edge and at least one face before this query becomes available.

An example report is provided below:



Geometric Range of Selected Faces/Edges/Vertices

This query reports the minimum and maximum coordinate values (relative to the laboratory coordinate system) of a bounding-box that contains all selected faces, edges, and vertices. It also reports the dimensions (change in coordinate value) of the bounding-box in the x-, y-, and z-direction. To make this query available, make one of the following selections:

- At least one face and at least one vertex
- At least one face and at least one edge
- At least one vertex and at least one edge

An example report is provided below:



See Also

Mixed Entity Queries

Feature Diagnostics

The feature diagnostics tool finds errors on feature curves.

The tool identifies two types of feature error:

- errors in explicit feature curves (feature errors)
- errors in implicit feature curves

To run *Feature Diagnostics* in **Feature Repair** mode, click **Start Diagnostics...**.



Once the diagnostics are complete, the *Feature Diagnostics* group box displays a summary of feature errors.



The **Feature Diagnostics** tool does not always identify all feature errors, and can flag errors that are not problems.

After identifying errors, use the Fix All Feature Errors tool to fix them.

Feature Errors (Explicit Feature Curves)

Explicit feature errors are typically undesirable part curves mixed among the explicit feature curves. Explicit feature curves are created during surface import, or during manual computation from edge angles and patches. The following types of error can occur:

- An excess of feature curves converging on a point
- Disconnected or isolated curves
- Over-constrained curves close to one another
- Missing feature curves

Implicit Feature Errors

The software calculates implicit feature curves at the start of the meshing pipeline, using boundary perimeters, free edges, and non-manifold edges that exist on the surface. The features diagnostics tool can identify such errors ahead of time, so you can fix them on the surface before starting the meshing process.

See Also

Changing The Parts In The Active Model
Computing Feature Curves
Importing Geometry
Fixing All Feature Errors
Surface, Contact, and Curve Edit Controls
Preparing Surfaces

Repair Surface Mode

This section provides general guidelines on how to use the manual surface repair tools to identify, isolate, and fix errors on a surface.

It is important to repair the surface before moving onto the next stage in the meshing process. Note, it would not be feasible to document an exhaustive list of surface repair examples that use the tools. Only user experience can lead to effective and efficient use of the tools.

The surface repair operations can affect association with underlying CAD parts. In general, any surface repair operation that creates new faces causes a loss of CAD association for the new faces, and possibly the face neighbors. If you activate CAD projection, no projection takes place for these disassociated faces. Therefore, if CAD association is important, keep surface repair to a minimum. You can view the current CAD association for a given surface using the Color Faces display control.

The *Surface Repair* tool is a flexible and customizable threshold-based tool to identify, diagnose, select, and repair problem areas in surface meshes:



Activating any of the surface diagnostics options color-codes faces, edges, and vertices according to the diagnosed surface errors. In the example below, pierced faces appear red, and free edges appear green:



The tools do not allow you to create new vertices directly. If you want to fill many holes or zip many edges, use the hole filler and edge zipper tools respectively. For other cases, use the surface repair workflow to identify and repair issues.

What Is the Surface Repair Workflow?

In general, you can use the following steps to fix surface errors in surface repair mode:

- Use threshold diagnostics to identify problems
- Identify the problem area
- Isolate the problem area
- Select entities for repair
- Fix surface problems

See Also

What Are Pierced Faces
Zipping Edges
Filling Holes in the Surface

Launching The Surface Preparation Tool
Using Color Modes
Surface Mesher Controls and Values

Using Threshold Diagnostics

When you first start Surface Repair, no diagnostics are shown in the *Surface Diagnostics* panel.

To activate thresholds and to run diagnostics:

1. Click Manage Thresholds... (**Manage Thresholds**).

   The *Create/modify Thresholds* panel appears:



You can perform the following actions in the *Create/modify Thresholds* panel:

- Activating and Deactivating Diagnostics
- Creating and Modifying Non-Default Diagnostics
- Running Diagnostics
- Also, when using the panel it is important to understand the situations where the diagnostics can become out-of date

Activating and Deactivating Diagnostics

You can choose which diagnostic checks to run. For a new simulation, all diagnostics are activated. When loading a previously saved simulation, the activated diagnostics reflects the state when the simulation was saved.

2. To deactivate a diagnostic, deactivate the checkbox in the first column.



3. To activate or deactivate all thresholds, use the (**Enable All Thresholds**) checkbox in the table header.

   The pierced face and face quality diagnostics can take a long time on large parts. If you know that these types of errors do not exist, deactivate these checks. A diagnostic check is run immediately when you activate a diagnostic.

   STAR-CCM+ records the state of the *Surface Diagnostics* dialog when you close it, and is saved to the simulation file for a future session.

Creating and Modifying Non-Default Diagnostics

You can create and delete other non-default diagnostics and use these diagnostics to apply thresholds.

To create a threshold:

4. Click New Table Entry (**New Table Entry**).
5. Enter a suitable name. Blank names are automatically given a default value.
6. Choose from one of the following types:

- **Face Quality**: Selects faces according to the face quality measure.
- **Face Proximity**: Selects faces according to the face proximity measure.
- **Face Area**: Selects faces according to area.
- **Edge Length**: Selects edges according to length.
- **Edge Dihedral Angle**: Selects edges according to the angle between the face normals of the connected faces of the edge.

7. Select an operator type:
   - **Above**: The tool searches for all values higher than one you specify in the **Min** entry field.
   - **Below**: The tool searches for all values lower than one you specify in the **Min** entry field.
   - **Between**: The tool searches for all values higher than the one that you specify in the **Min** entry field but lower than the one you specify in the **Max** entry field.
   - **Not between**: The tool searches for all values lower than the one that you specify in the **Min** entry field but higher than the one you specify in the **Max** entry field.
8. Select a color from the tablet of provided colors.

The Diagnostic input area appears similar to the one shown below:



9. Click **Create** to create the Threshold.

   A new entry is added to the table:

✔	5	Non-manifold edges (Default)
☐	6	Non-manifold vertices (Default)
✔	7	Below Dihedral 0.03 radians

You cannot change the diagnostic type once it has been created. Instead, delete the diagnostic and create another one. To delete a diagnostic, click its name in the table and click Delete Selected (**Delete Selected**). The diagnostic is removed from the table. If it had been run before deletion, its results are removed also. Deleting a diagnostic cannot be undone.

You can only delete non-default diagnostics. Default diagnostics can be edited and deactivated but not deleted.

Running Diagnostics

Once you are satisfied with the selected diagnostics and their threshold settings:

10. Click **OK** to run the diagnostics tool.
    The panel closes and the Surface Diagnostics table is populated, ready for you to identify individual problem areas.



Out-Of-Date Diagnostics

STAR-CCM+ does not rerun all diagnostics every time you click **OK** in the *Create/modify Thresholds* panel.

A diagnostic is marked out-of-date if:

- You modified its threshold properties in the panel.
- You execute other diagnostics while it was not activated. This condition is true even if the other diagnostics are up-to-date.
- Once you deactivate a diagnostic and click **OK**, it is marked out-of-date.

See Also

Changing The Parts In The Active Model
Identifying A Problem Area
What Is Face Proximity
What Is Face Quality
What Is The Default Color Mode
Surface, Contact, and Curve Edit Controls
Preparing Surfaces

Identifying A Problem Area

After setting up and running the diagnostics checks, the results are displayed on the model.



A problem area is a collection of diagnostics that are connected to each other by one or more shared vertices. Faces, edges, or vertices can be identified in more than one diagnostic check. Each face, edge, or vertex is only displayed in the problem area collection that belongs to the last diagnostic check to identify it.

Note: If the surface is altered, the diagnostics and problem areas are automatically updated.

If more than 100 problem areas are identified, merges between two or more problem areas that were previously unconnected are no longer performed.

Smart selection (double clicking a face or edge that contains a diagnostic result) of what appears to be continuous problem areas no longer produces the expected selection result. Smart selections are limited to

the portions of the problem area that they belong to, even though visually they appear to be connected together.

The following message is displayed in the *Output* window:

```
Warning: Too many problem areas in surface, will not
attempt to merge neighboring areas.
```

To resolve this issue, either:

- Decrease the number of diagnostics that are activated. This action reduces the potential number of problem areas.
- Enter the *Manage Thresholds* panel and press **OK** to exit. The problem areas are automatically recalculated. Smart selection results are then as expected, regardless of the number of problem areas.

The **Display** checkbox can be deactivated for each diagnostic to allow you to inspect other diagnostics.



See Also

Using Threshold Diagnostics

Controlling the Display

After the diagnostics have been activated and run, the results can be modified to control the **Input** to the display.



Select one of the options:

- **Active Model**
- **Specified Parts / Specified Regions**
- **Displayed Set**

Active Model

1. Changing the *Input* to **Active Model** reports the diagnostics for the parts that are brought into surface repair initially.
   This setting is the default.

Specified Parts / Specified Regions

2. Changing the *Input* to **Specified Parts / Specified Regions** reports the diagnostics for the selected parts or regions.
3. To select a part or region, click the button next to **Specified Part(s):**

The diagnostics are reset to 0 until you select inputs from the selector tool.



The diagnostics are reported and displayed for the selected parts or regions only. The composite screenshot below shows the results of a diagnostic check for the PCB and Plate specified parts. Note the reduced diagnostic count.



Displayed Set

4. Changing the *Input* to **Displayed Set** reports the diagnostics only for the object that is shown in the scene.

   When all faces are shown, the results for **Displayed Set** and *Active Model* are identical.

5. More computations are performed when the *Input* is set to **Displayed Set** and faces in the display are changed, for example by clicking one of the (**Browse**) buttons.

6. If you have many faces, reduce your display set.

   The composite screenshot below shows a display set found when browsing through the Face Proximity diagnostic results. Note how the Count differs.



7. When the **Input** is set to **Displayed Set**, the **Reset Displayed** check box is deactivated.

Viewing a Selection Set

Once the diagnostics are run and the *Input* type is set, there are several ways that you can locate an error that is shown in the *Surface Diagnostics* panel.

1. Use the (**Browse**) buttons to cycle through the errors. The location of the issue is highlighted in the scene. If the *Reset View* option is activated, the display changes. If the *Reset Displayed* option is activated the subset of displayed faces changes. The composite screenshot above shows an example.
2. Click one of the *Surface Diagnostics* Count cells.

3. You can press <CTRL> and click several Count cells to combine the selection results.

   The affected locations are highlighted in the scene, even if the **Display** checkbox is deactivated for a clicked diagnostic.



Growing And Isolating A Problem Area

Once you have selected a suitable face, edge or vertex, you can isolate the error.

To isolate the error:

1. Press (**Grow Selected**), to select the area of interest.

This option grows the group of selected items, but does not affect the underlying displayed items.

2. Press (**Hide All Faces**), followed by **Show Selected Faces**.
   This option displays the selected entities in the *Graphics* window.

3. If necessary, press (**Grow Displayed Faces**).
   This option grows the displayed area.

Selecting Entities For Repair

You can select the faces, edges, and vertices that you wish to fix by using the mouse in the *Graphics* window.

Using the `<Ctrl>` key allows multiple selections. Select the face, edge, and vertex combinations that correspond with the required *Surface Repair* tool. You can deactivate all selections before selecting new items by using the (**Select None**) button, or clicking in any white space around the part.

Deleting Faces

This tool deletes the selected faces from the part. The tool becomes available when you select at least one face.

The following example deletes unwanted faces from the solid part of a geometry.



Clicking **Delete Faces** removes the extraneous faces, leaving the fluid volume.



Creating Faces From Vertices

This tool creates a face using the selected vertices. The tool becomes available when you select three vertices.

The following example fixes free edges by creating a face. The three vertices are selected.



Clicking **Create Faces from Vertices** creates a face that is defined by the selected vertices.



Collapsing Vertices

This tool collapses one vertex to the location of another.

The tool also collapses an edge or face to its mid-point. If you select an attached vertex with an edge or a face, the tool collapses the edge or face to the selected vertex. The tool becomes available when you select two vertices, an edge or a face (with an optional attached vertex for edges and faces).

The following example collapses two vertices, which fixes a poor quality face. You can encounter this scenario when fixing cracks on the surface using the zipping tool.



Using the **Collapse Vertices** tool removes the poor quality faces, without creating free edges.



Splitting Edges And Faces

This tool splits a face or faces attached to an edge in half, using the mid-point of the selected edges. The tool becomes available when you select at least one edge or face.

The following example shows the splitting of two adjacent faces by splitting the connecting edge in half.



Using the **Split Edges** tool in this manner has the following results.



Swapping Edges

This tool swaps the connected edge between two faces. The tool is available when you have selected only one edge.

The following example swaps edges to improve face quality.



In this instance, the **Swap Edges** tool improves the quality of the connected faces, by reducing face skew.



Smoothing Vertices

This tool locally smooths a selected vertex, or vertices to improve the quality of the surrounding faces.

Use this tool if you want to refine the quality of a surface mesh by improving the distribution of vertices on a surface or along a feature curve.

The smooth selected / attached vertices tool becomes available when you select at least one vertex, edge, or face.

To smooth the selected vertices:

1. Right-click (**Smooth selected / attached vertices**), and select **Options**.
2. Select the smoothing mode.

- Unconstrained
- Constrained

3. Set the **Number of Iterations** depending on the number of smoothing iterations that you want to perform.
4. Click **OK**.

Unconstrained Smoothing

Use this mode to smooth vertices arbitrarily on a surface.



The following *Smoothing Options* are available for unconstrained smoothing:

Preserve Feature Edges	Ensures that the tool does not move vertices belonging to a feature curve. However, where long chains of feature curve vertices are selected, it is possible that these vertices are smoothed along the length of the feature curve.
Reproject to initial surface	Performs an additional step after smoothing, which attempts to project the selected vertices back to the original surface.

The following example smooths the selected vertex to increase the quality of the attached faces.



Using the **Smooth Vertices** tool with an appropriate number of iterations has the following result.



Constrained Smoothing

Use this mode if you want to smooth vertices in a specified coordinate direction.

This action moves vertices along the edge to distribute the vertices more evenly. The vertices are smoothed to increase the quality of the faces surrounding the selected edge.



The following *Smoothing Direction* options are available for constrained smoothing:

Smooth in X direction	**Activated**	Smooths vertices in the X direction.
	Deactivated	Does not smooth vertices in the X direction.
Smooth in Y direction	**Activated**	Smooths vertices in the Y direction.
	Deactivated	Does not smooth vertices in the Y direction.
Smooth in Z direction	**Activated**	Smooths vertices in the Z direction.
	Deactivated	Does not smooth vertices in the Z direction.
Coordinate System	Allows you to specify the coordinate system that you require. Set to **Laboratory** by default.	

Note: Smoothing must be allowed in at least one coordinate direction. Therefore, it is not possible to deactivate all three smoothing directions.

The options to **Reproject to the initial surface**and **Preserve feature edges** are not available for constrained smoothing.

The following example shows an edge that is selected for constrained smoothing in the Y-direction.



The result after five iterations of smoothing shows that the vertices along the selected edge are smoothed in the Y-direction.



Projecting Vertices

This tool moves the selected vertex or vertices to a new location using one of five modes.

This tool is always available, as selections can be made in the Graphics window after the tool is activated. However, it is recommended that you select the required entities before activating the tool, as the *Selection Controls* are not available, which reduces the ease of selection.

Clicking **Project selected vertices** opens the *Project Vertices Options* in-place panel, which specifies the projection mode and options.



Select the desired projection mode and enter the appropriate specification options and values, before clicking **Apply** to execute the projection.

The tool remains open after projection so you can redo the operation with different values or specify a new projection.

When closing the dialog, values are only saved for the current **Projection mode** and for projections that have been applied.

You can only specify flat planes using this mode, regardless of the coordinate system type.

See Also

Changing The Parts In The Active Model

Projecting Vertices To A Plane

The **Project to a Plane** mode projects selected vertices by the shortest distance to an imaginary plane.

Origin and normal components define the plane. The pick operation is available to select origin and normal points interactively in the Graphics window. You can specify origin and normal coordinates in the following ways:

- Specify the plane origin and normal using a Vector.

  Either use one or both of the following methods:

  - Define the plane origin and normal using the plane tool.

    STAR-CCM+ extracts the plane origin and normal from the orientation of a plane that you define using the Plane Tool. The vertices are then projected to the plane that passes through the specified origin.

  - Define the plane origin and normal using Fx Section.

    The **Fx Section** button allows you to specify a plane section interactively in the *Graphics* window to represent the plane origin and normal.

  - Enter the plane origin and normal vector coordinates directly.

- Specify the plane origin and normal using Picked Points.

  The **Picked Points** operation allows you to select three points in a specified coordinate system to represent the plane normal. The three points all lie on the plane to which you project the vertices.

  When defining a plane with three points, all three points must be different and cannot be aligned. The plane cannot have zero area.

To Define the Plane Origin and Normal Using Fx Section:

1. Make sure that *Normal Specification* is set to **Vector**.
2. In the *Plane and Section Tools* group box, click **Fx Section.**



3. In the *Graphics* window, click once to begin your selection of the plane section.
4. Move the cursor until an appropriate plane section is defined, then click to confirm the selection.
5. If necessary, use the plane tool to refine the selection.

To Define the Plane Origin and Normal Using the Plane Tool:

6. In the *Plane and Section Tools* group box, activate *Display Plane Tool?*.



7. In the *Graphics* window, click and drag the Plane Tool to represent the plane origin and normal.
   You can also use the X, Y, and Z buttons to toggle the Plane Tool to lie on the X, Y, or Z axis.

To Define the Plane Origin and Normal Using Picked Points:

8. Define the Plane Origin.
9. Set *Normal Specification* to **Picked Points.**



10. To specify the plane points, either:
    a) Use the **Pick** buttons to define each plane point interactively in the *Graphics* window, or
    b) Enter each set of **Plane Point** coordinates directly.

Note: To set the plane origin coordinates as one of the plane normal coordinates, activate *Use plane origin for plane point #1.*

The diagrams below demonstrate an example of this method.

After selecting the vertices to be projected and entering the project vertices mode, the origin (yellow) and normal (red, green, and blue) points are selected.



After applying the operation to the geometry, the plane on which the vertices lie is coincident with the origin. The projected plane also maintains the normal specification of the selected normal points:



Note: The **Apply** button is only available if one or more vertices are selected for projection.

See Also

Picking Points To Define A Projection Plane Or Line

Projecting Vertices To A Line

The **Project to a Line** mode projects vertices by the shortest distance to an imaginary line.

You define the line with an origin and a direction vector, or two points in a specified coordinate system. The pick operation is available to select points interactively in the Graphics window, providing a simple intuitive method of defining a line. The **Apply** button is only available if one or more vertices have been selected for projection.


Servers
Star 1
simulation
scene/plot
edit
Project Vertices
Projection Mode
Project to a Line
Project to a Line
Line Specification
Pick End Points
Origin and Direction
Pick End Points
Line End Poir
X 0.1 m
Y 0.09166667 m
Z 0.3 m
Pick...
Line End Point #2
X -0.1 m
Y 0.09166667 m
Z 0.3 m
Pick...
Coordinate System
Laboratory
Apply
Return
Help

This diagram illustrates the **Line Origin**, **Line Direction** and **Coordinate System** specification:



This diagram illustrates the **Line End Point #** specification:



When defining a line with two points, both points must be different (that is, you cannot project to a zero length line).

You can only specify straight lines using this mode, regardless of the coordinate system type.

See Also

Picking Points To Define A Projection Plane Or Line

Projecting Vertices In A Coordinate System

The **Project in a Coordinate System** mode projects vertices to a single or double ordinate location in a given coordinate system.

Any unspecified ordinates maintain their original value. For example, you can project vertices to a specified X value in a Cartesian system or to a specified R-q value in a cylindrical system. The **Apply** button is only available if one or more vertices have been selected for projection.



This mode differs from the project to plane mode as projections to cylindrical and spherical planes are possible.

1. First select the **Coordinate System**.

Selecting the coordinate system (**Cartesian**, **Cylindrical, or Spherical**) updates the **Projection Ordinate(s)** menu with the appropriate coordinate directions.

2. You can then select the **Projection Ordinate(s)**, and enter the ordinate values in the desired units. The dialog accepts negative values for the radial direction in a cylindrical system, which projects vertices in the opposite direction.
3. The projection is performed when you click **Apply**.

The following example projects vertices in the radial direction of a cylindrical system. The projection reduces the sleeve outer radius value to produce a thinner sleeve. This example uses a cylindrical coordinate system that is defined at the axis of the sleeve, where the radial axis points outwards.

4. Select the vertices on the outer surface:



5. Open the *Project Vertices* in-place panel, and select the **Project in a Coordinate System** mode.
6. Select the cylindrical system for the **Coordinate System**, and enter the radial, R, **Projection Ordinate** value.



7. The projection is performed when you click **Apply**.
   The *Graphics* window displays the new result.



The tool maintains the original vertex selection, so you can repeat the process for different radial values.

Projecting Vertices To A Surface

The **Project to a Surface** mode projects one or more vertices to a set of faces.

The **Apply** button is only available if at least one vertex and one face have been selected for projection. By default, the tool projects using the shortest distance. However, projection in a specified direction is also possible. If projecting in a specified direction with two independent sets of faces selected, the tool projects to the nearest face, regardless of whether it is in the positive or negative direction. In all instances, the tool ignores vertices that belong to the selected surface.

The image below shows the default **Project to a Surface** mode. As you can see, projection using the shortest distance does not require a coordinate system.



If you select **Specified Direction** for **Projection Method**, the **Coordinate System** menu becomes available, along with the projection **Direction**.



1. First select the **Coordinate System**.

   Selecting the coordinate system (**Cartesian**, **Cylindrical, or Spherical**) updates the **Direction** menu with the appropriate coordinate directions.

2. The projection is performed when you click **Apply**.

   If projection is not possible in the specified direction, the vertex position remains unchanged and the *Output* window displays a warning message.

Projecting vertices onto a convex surface using the shortest distance method results in “edging”, where projected vertices congregate along the edges of the selected surface. To avoid this effect, project along a specified direction instead.

The following example projects vertices to a surface. The tool projects vertices on the inner surface of a sleeve onto the inner pipe surface, using the shortest distance method.

3. First, select the vertices on the sleeve inner surface, plus the faces on the pipe outer surface:



4. Open the *Project Vertices* in-place panel, and select the **Project to a Surface** mode:



5. The projection is performed when you click **Apply**.

   The *Graphics* window displays the new result:



The tool maintains the original vertex and face selections, so you can repeat the projection with alternative settings.

Projecting Vertices To An Edge

The **Project to an Edge** mode allows the projection of one or more vertices to the nearest edges.

The **Apply** button is only available if at least one vertex and one edge have been selected for projection. The tool ignores any vertices that belong to the selected edges.



The mode requires no additional options. The projection is performed when you click **Apply**.

The following example projects vertices to edges. The tool projects vertices on the inner edge of a sleeve onto the inner pipe edge.

The *Graphics* window displays the result.

1. First, select the vertices on the sleeve inner surface, plus the edge on the pipe outer surface.



2. Open the *Project Vertices* in-place panel, and select the **Project to an Edge** mode.
3. The projection is performed when you click **Apply**.



Picking Points To Define A Projection Plane Or Line

When using the Project to a Plane or Project to a Line modes, you can pick positions in the *Graphics* window to define points.

This feature provides a quick and intuitive way of defining a plane or line.

To begin a pick operation, click one of the **Pick** buttons in the *Project Vertices* in-place dialog. You can then pick vertices, edges, or faces in the *Graphics* window. Picking a vertex uses the exact position of the vertex. Picking an edge uses the selected position along the direction of the edge and picking a face uses the selected position on the surface of the face. The *Output* window reports whether a face, edge, or vertex has been selected and the coordinates relative to the global coordinate system.

When a pick is made, the coordinates of the point are added to the appropriate fields in the *Project Vertices* in-place dialog. You can manually change the coordinates of the point after picking to fine-tune its position. These coordinates are relative to the current coordinate system. Therefore, changing the coordinate system relocates the point to a position relative to the new coordinate system. The point is represented in the *Graphics* window with a color-coded sphere, which automatically updates when the point position changes.

To cancel the pick operation, press `Esc` on the keyboard, or click in the white space in the *Graphics* window.

This section provides the following examples:

- Projecting to a plane defined by picking three points
- Projecting to a line defined by picking end points

Picking Three Points to Define a Plane

This example illustrates picking three points to define a plane, in order to project a group of vertices to a surface.

1. First, select the required vertices using the *Selection Controls*, before activating the *Project Vertices* tool.



2. Define a plane by picking points on the required surface in the *Graphics* window.
   In this example, three vertices are selected, with the selection of **Plane Point #3** shown below:



The coordinates of each point are reported in the *Project Vertices* in-place dialog:



3. Clicking **Apply** projects the selected vertices to the user-defined plane.



Picking End Points to Define a Line

The following example illustrates picking end points to define a line, in order to project a group of vertices the line.

4. Select the required vertices using the *Selection Controls*, before activating the *Project Vertices* tool.



5. Define a line by picking the end points in the *Graphics* window. In this example, two vertices are selected:



The coordinates of each point are reported in the *Project Vertices* in-place dialog:



6. Clicking **Apply** projects the selected vertices to the user-defined line.



See Also

Projecting Vertices To A Plane
Projecting Vertices To A Line

Merging Vertices And Faces

This tool merges adjacent vertices within a given tolerance and merges duplicate faces (that is, faces that share three vertices).

The tool is always available, as you can use it on the active model without any selections.

Clicking **Merge Adjacent Vertices / Duplicate Faces** opens the *Merge Options* dialog.



The dialog specifies whether the tool merges adjacent vertices, duplicate faces or both. The default options merge adjacent vertices and duplicate faces with a tolerance of 0.0001m.

If you activate the **Merge Adjacent Vertices** option, the tool merges vertices first. If you activate the **Merge Duplicate Faces** option, the tool merges faces second. The tool deletes any collapsed faces or edges from the vertex merge operation in order to ensure a valid result.

Clicking **OK** starts the merge process using the options that are selected in the dialog.

Operate On Options

The *Operate On* options determine whether the tool merges the currently selected faces and vertices, or all faces and vertices in the active model. The active model is the faces and vertices belonging to the parts originally selected in the Surface Repair Options dialog. The default option is to operate on the selected set only.

When merging vertices using the selected set only, the tool removes all merged vertices from the current selection. Therefore, if you repeat the merge with an increased tolerance, the tool only uses unmerged vertices from the original selection. However, when merging duplicate faces, the tool retains all selected faces, apart from those faces that the merging process deleted.

When merging using the active model, the tool deactivates any existing vertex or face selections. When vertex merging is complete, the tool selects the master vertex for each successful merge. Similarly, when merging faces, the tool selects the master face for each successful merge. You can use this information to judge the success of the merge, and where it took place on the part.

See Also

Launching The Surface Preparation Tool

Merging Vertices

When merging vertices, review and change the tolerance value, if necessary.

Set the value to a reasonable value based on the maximum distance that separates the vertices you want to merge, plus a small tolerance (say 1/10th to 1/5th of the measured value). If the tolerance is too high, inadvertent merging of neighboring vertices can occur, especially when choosing to merge based on the active model. The safest way to merge vertices is to select the required vertices, and use the **Selected set only** option in the *Merge Options* dialog.

The following example is a part containing free edges, representing an unmerged portion of the surface:



1. Select the free edges by clicking the green box under *Surface Diagnostics*.



2. click **Select Attached Vertices** to select the vertices that are attached to the selected edges.



As an alternative method, you could use (**Select Zone**) to select the appropriate vertices.

3. To merge the vertices:
   a) Enter the tolerance in the *Merge Options* dialog.
   You can estimate the required tolerance value by measuring the distance between two vertices on either side of the gap.
   b) Deactivate the **Merge duplicate faces** option.
   c) Click **OK**.



The display updates with the merged result.



The tool has uniformly merged the vertices on one surface with the vertices on the other surface, resulting in a clean merge line. The *Output* window shows a summary of the number of vertices merged.



If you set the tolerance too low, either no or partial merging would take place. However, any unmerged vertices would remain in the selected set allowing for further merging with a revised tolerance value.

The vertex merging scheme uses a lowest vertex number approach, which means it always merges the highest vertex number (slave vertex) to the lower number (master vertex) within the specified tolerance distance.

Merging Duplicate Faces

A duplicate face is a face that shares three vertices as another face.

The **Repair Surface Diagnostics** tool flags these as close proximity faces and colors them orange. However, not all close proximity faces are necessarily duplicate faces.

If multiple duplicate faces exist for one face, the tool removes all duplicates using the merge process.

The vertex merge tool can occasionally create duplicate faces, especially when merging vertices from close proximity faces. Using the tool on the active model merges these additional duplicate faces.

The following example shows the **Merge Duplicate Faces** tool in use.



1. Select the close proximity faces by clicking the orange diagnostic box.



2. Click the merge button to open the dialog, and **OK** to proceed with the merging operation using the selected set.

   If any of the close proximity faces are duplicates, the merge removes the duplicate face. The *Output* window reports the number of duplicate faces deleted.



Filling Holes

This tool fills a hole with new faces using the selected edge loops.

For closed loops, select at least three edges for this tool to become available. For open loops, select at least two edges for this tool to become available.

The fill holes tool has three modes:

- Mixed - Fill a combination of holes that are defined by chains of selected edges.
- Closed loop(s) only — Fill a hole that is defined by a continuous chain of selected edges.
- Open loop(s) only — Fill a hole that is defined by a discontinuous chain of selected edges.

Filling Mixed Loop Holes

The mixed hole filler fills a combination of holes that are defined by continuous or discontinuous chains of selected edges.

This option is useful if both open and closed loop holes exist within the same surface.

To fill mixed loop holes:

1. Right-click (**Fill holes**) and select **Options**.
2. In the *Hole Filling Options* dialog, set *Fill Mode* to **Mixed.**.



3. Set the *Open Loop Closure Options* mode to one of the following:
   - Fill Each Open Loop Separately
   - Fill Between an Open Loop Pair

   See Filling Open Loop Holes
4. Set the Loop Interaction Options mode to one of the following:
   - **Fill All Holes Together**
   - **Fill Each Hole Separately**

   See Loop Interaction Options
5. Click **OK**.
6. In the *Graphics* window, select the edges that define the open and closed loops.
7. To fill the holes, click (**Fill holes**).

For example, the following surface contains open and closed loop holes.



Choosing the **Mixed** *Fill Mode* lets you fill all open and closed loop holes at the same time.



See Also

Loop Interaction Options
Filling Open Loop Holes

Filling Closed Loop Holes

The closed-loop hole filler fills holes that are defined by continuous loops of edges.

You typically use this option before extracting the internal fluid volume for a CFD analysis.

Note: When filling closed loops, the tool checks that the selected edges form a closed loop. The tool does nothing if it does not find the closed loop.

To fill a closed-loop hole:

1. Right-click (**Fill holes**) and select **Options**.
2. In the *Hole Filling Options* dialog, set *Fill Mode* to **Closed Loop(s) only**.



3. Set the Loop Interaction Options mode to one of the following:

- **Fill All Holes Together**
- **Fill Each Hole Separately**

See Loop Interaction Options.

4. Click **OK**.
5. In the Graphics window, select the edges that define the closed loop. Double-click a single free edge to automatically selects all the free edges connected to it.
6. To fill the hole, click (**Fill holes**) button.

The following example uses the fill holes tool to cap the hole in a solid component.



The image below shows the result of the hole fill operation. You can improve the poor quality faces by locally remeshing the surface. See Remeshing Faces.



See Also

Loop Interaction Options

Filling Open Loop Holes

The open-loop hole filler fills holes that are defined by discontinuous loops of edges.

Note: When filling open loops, the tool checks for closed loops in the selection and cancels the command when it finds one.

To fill an open-loop hole:

1. Right-click (**Fill holes**) and select **Options**.
2. In the *Hole Filling Options* dialog, set *Fill Mode* to **Open loop(s) only**.



3. Set the *Open Loop Closure Options* mode to one of the following:
   - Fill Each Open Loop Separately
   - Fill Between an Open Loop Pair
4. Set the Loop Interaction Options mode to one of the following:
   - **Fill All Holes Together**
   - **Fill Each Hole Separately**
5. Click **OK**.
6. In the *Graphics* window, select the edges that define the open loop(s).
7. To fill the hole(s), click (**Fill holes**).

Fill Each Open Loop Separately

This option individually fills the holes defined by each open loop of edges.

You can fill any number of open loops simultaneously.

Note: When filling separate open loops, the tool skips loops with collinear edges (straight lines) and loops with a single edge. Also, the tool checks for Y-junctions and cancels the command when it finds one.

For example, the following image shows two open loops.



8. Choosing to fill each open loop separately generates the following surface.



Fill Between an Open Loop Pair

This option fills the hole between a pair of open loops.

Note: When filling a pair of open loops, you can only select two open loops; the tool cancels the command if you select more or less. Also, the tool checks for Y-junctions and cancels the command when it finds one.

For example, the following image shows two open loops.



9. Choosing to fill between an open loop pair generates the following surface.



See Also

Loop Interaction Options

Maintaining Local Curvature

Selecting the *Maintain Local Curvature* option fills a hole with a surface that follows the surrounding curvature.

For example, the hole in the geometry below requires a filled surface that is not flat, but curves in each direction.



When the *Maintain Local Curvature* option is active, STAR-CCM+ generates the following surface.



For comparison, deactivating the *Maintain local curvature* option generates the following surface.



This feature produces the best results when faces surrounding the hole are associated with underlying CAD geometry. The CAD association reduces the amount of approximation that is required in generating the filled surface. If at least one of the surrounding faces has CAD association, then the filled surface also has CAD association, which allows you to project to CAD when surface meshing.

When the **Fill All Holes Together** mode is selected, the **Maintain Local Curvature** option can not be activated. Therefore to ensure curvature is maintained, set the *Loop Interaction Options* to Fill Each Hole Separately.

Select each hole, then click (**Fill holes**).

If for any reason the tool cannot maintain local curvature, it generates a filled surface without local curvature.

Using Feature Curves to Control the Filled Surface Shape

As a rule, mark sharp edges as feature curves before using the **Maintain Local Curvature** option. Applying feature curves to the edges surrounding the hole controls the construction of the filled surface. In the example below, maintaining local curvature of all surrounding edges generates an incorrect filled surface.



Filling the hole without marking any hole edges as feature curves generates the following filled surface. In this case, the tool maintains the curvature of the flat surfaces, which makes the filled surface initially tangential to each side. This set-up produces an incorrect shape.



Marking sharp edges as feature curves tells the hole filler to ignore local curvature around them. For this example, three sides of the hole (highlighted below) are marked as feature curves.



Running the hole filler with these edges marked as feature curves generates the correct surface.



See Also

Loop Interaction Options
Surface Mesher Controls and Values

Loop Interaction Options

You can fill all holes together, or separately.

- **Fill All Holes Together**

  If you have two loops, one inside the other, the **Fill All Holes Together** option will fill between the two loops. This option does not let you maintain local curvature.

- **Fill Each Hole Separately**

  If you have two loops, one inside the other, the **Fill Each Hole Separately** option will fill each loop separately and will not fill between the two loops. This option does let you maintain local curvature.

Zipping Edges

This tool locally zips faces together to remove free edges.

Select two or more edges before the tool becomes available. Right-clicking on the button provides access the same options as for zipping edges using feature curves.

The following example uses the zip tool to close a small crack on the surface mesh.



The tool projects the vertices on one side of the crack to the edge on the other side, and merges close vertices.



In some instances, the tool can produce poor quality faces, if it fails to merge close vertices due to tolerances.



You can fix these issues in several ways:

1. Select the short edge, and use the collapse vertices tool. This collapses vertices and remove the poor quality face from the surface mesh.
2. Select the close vertices and use the merge adjacent vertices tool, with an adequate tolerance specified.
3. Select two or more free edges on the geometry, then:

   1. In the *Selection Control* group box, click (**Grow selected**) once or twice to grow the selection of neighboring faces.

   2. Click (**Zip selected edges**).

   The selected faces are then locally surface remeshed. This eliminates the original faces and provides a high quality triangulation result. The advantage of this method over the other two is that it can be completely automated. There is no need to identify short edges or supply a merge tolerance.

See Also

Merging Vertices And Faces
Collapsing Vertices

Remeshing Faces

This tool locally remeshes the selected faces to improve the face quality.

The tool retains feature curves or edges that are flagged as feature curves on the surface. Select at least one face before this tool becomes available.

1. Right-click on the **Remesh Selected Faces** button and select **Options...** to open the *Remesh Faces Options* dialog.



2. **Specify target size** allows you to set the triangle size of the remeshed region. Deactivating the **Specify target size** option means that the surface remesher uses the local face edge lengths to determine the triangle size.
3. **Do Curvature Refinement** adds curvature-based refinement to the remeshing operation. You can specify the #*Pts/Circle* and *Minimum Size* values to improve surface definition on curved portions of the geometry.
   See Surface Mesher Controls and Values.

   The following example uses the local remesher to improve the quality of faces produced when filling holes.



Using the local remesher on the faces highlighted above has the following results.



See Also

Surface Mesher Controls and Values

Filling With A Polygonal Patch

This tool patches an arbitrary planar hole by clicking the surface surrounding the hole.

The tool is always available, as it does not require any selections.

Define the patch using the cursor in the *Graphics* window to click points around the patch perimeter. End the patch definition by right-clicking a point, or clicking the starting point.

The following example uses the **Polygonal Patch** tool to cover a hole before using the surface wrapper.



As the patch does not fill the hole exactly, the surface diagnostics report numerous errors. However, if you intend to use the surface wrapper, these errors are not an issue, as the surface wrapper only requires a closed surface.



Transforming Surface Entities

This tool transforms surface entities (vertices, edges, or faces) in a local coordinate system using one of five modes.

This tool is always available, as you can make selections in the *Graphics* window after the tool is activated. However, as the *Selection Controls* are not available within the tool, select the required entities before activating the tool.



To transform surface entities:

1. Click (**Transform vertices, edges and/or faces**).

   The **Transform Selection** in-place panel appears.
2. Select the required transformation from the *Transform Mode* menu.
3. Select the required vertices, edges, and/or faces.

   For faces, you can choose to transform a copy of the faces and keep the original faces in place. To do this, activate *Copy Selected Faces for Transform*.

   See Using the Copy Selected Faces for Transform Option
4. Enter the required input values and make any necessary selections.

   Transformation takes place relative to the current position of the selected vertices.
5. Click **Apply**.

The tool remains available after transformation so you can redo the operation with different values or specify a new transformation.

You can only use coordinate systems that were available when you opened the in-place panel. If you create a coordinate system while using the tool, it is not available until you exit and restart the **Transform vertices, edges and/or faces** tool. Also note, any values that you enter into the in-place panel are saved when you click **Return** and leave the panel.

Transforming Surface Entities to a Coordinate System

The **Coordinate Transform** mode transforms the selected surface entities from their position relative to a reference Cartesian coordinate system, to a new position relative to a target Cartesian coordinate system.

The diagram below illustrates the transformation for a single vertex. The ordinates of the position are the same in each Cartesian coordinate system.



- You can only use **Coordinate Transform** mode with a Cartesian coordinate system that was available when you opened the in-place panel.

  The tool does not list non-Cartesian coordinate systems in the drop-down menus.
- You can use **Coordinate Transform** mode to transform an entire part from one coordinate system to another, while maintaining the relative shape of the part.

Translating Surface Entities

The **Translate** mode translates vertices, edges, and/or faces along a specified vector, relative to a selected coordinate system, or a specified distance along the vertex normal.

Translating Surface Entities Along a Specified Vector

The image below shows the default **Translate** mode, with **Specify Vector** selected from the **Translate Mode** drop-down menu.



The **Translation Vector** components specify the desired translation. These components are relative to the **Reference Coordinate System**, which you select from the menu provided.

The **Reference Coordinate System** menu sets the coordinate system, with Cartesian, cylindrical, or spherical systems available. The **Translation Vector** component labels update to reflect the chosen type of coordinate system. The menu lists those systems that were available when you opened the dialog.

When translating vertices in a cylindrical or spherical coordinate system, the translation is relative to the component axes of these systems. For example, a translation in the R-direction moves the vertices in the radial direction, relative to the chosen coordinate system. Similarly, a translation in the q-direction rotates the vertices about the origin. If a vertex is at R=0, translating in the radial direction does not move the vertex. Also note that the tool permits negative radial translations.

Clicking **Apply** applies the translation, and report the result in the *Output* window.

The following example translates the selected vertices in the z-direction.



In this instance, the **Translation Vector** is `[0, 0, 0.05] m,m,m` to translate the vertices 0.05 m in the z-direction. Clicking **Apply** results in the following translation.



Translating Surface Entities Along a Vertex Normal

You can also translate vertices, edges, and/or faces in the normal direction by a user-defined distance. The normal is calculated as the average normal of all connected faces.

To translate vertices along a vertex normal:

- Select **Vertex Normal** from the **Translate Mode** drop-down menu.



The **Distance** specifies how far the selected vertices are translated. The **Distance** is relative to the current position of the vertices that are selected, or the vertices that are attached to the edges, and/or faces that are selected. Applying the same distance twice, doubles the translation.

Clicking **Apply** applies the translation, and reports the result in the *Output* window.

It is not always possible to calculate the vertex normal when one of the connected faces is of poor quality or if the connected faces converge to a sharp point. If the vertex normal cannot be calculated, the *Output* window reports:

```
Vertex normal calculation not possible for vertex XXXXX.
This vertex won't be translated.
```

The following example translates each selected vertex in its normal direction.



In this instance, as the surface is cylindrical, the calculated vertex normals offset the vertices, but maintain a circular shape.



See Also

Using the Copy Selected Faces for Transform Option

Rotating Surface Entities

The **Rotate** mode rotates the selected surface entities about an axis that is defined by vector components in the selected reference coordinate system. The rotation follows the right-hand rule.

To specify the angle in different units, enter the unit after the numerical value in the **Angle** entry box:



You can validate the angle specification by pressing `<Enter>` after entering the value and unit. If the unit specification is invalid, the **Angle** entry box appears red.

The **Axis Vector** component boxes define a vector about which the selected surface entities rotate. These components are relative to the

Reference Coordinate System, which you select from the menu provided.

The **Reference Coordinate System** menu sets the coordinate system, with Cartesian, cylindrical, or spherical systems available. The **Axis Vector** component labels update to reflect the chosen type of coordinate system. The menu lists those systems that were available when you opened the dialog.

Clicking **Apply** applies the rotation, and records the result in the *Output* window.

The following example rotates each selected vertex about the z-axis.



In this instance, the vertices defining the hexagonal volume rotate, but the vertices defining the cylindrical volume remain stationary. This effect can cause poor quality faces as the connectivity is maintained and faces become twisted.



Scaling Surface Entities

The **Scale** mode multiplies the position vector components of the selected entities by a user-defined scale factor.



Ticking the **Scale All Dimensions Uniformly** option, which is activated by default, grays out the second and third *Scale Factor* boxes. The factor that is given in the first box then applies to all coordinate directions.

The **Reference Coordinate System** menu sets the local coordinate system, with Cartesian, cylindrical, or spherical systems available. The **Scale Factor** component labels update to reflect the chosen type of coordinate system. The menu lists those systems that were available when you opened the dialog.

Clicking **Apply** applies the scaling, and records the result in the *Output* window.

The following example scales the selected vertices:



In this instance, the **Scale Factor** in each dimension is `[2.0, 0.5, 0.5]`. As the box is located as shown by the coordinate system, the box is twice as long in the x-direction and half as long in the y- and z-direction.



Reflecting Surface Entities

The **Reflect** mode reflects the selected entities about a plane with a user-defined normal vector.

The **Normal Vector** component boxes define a vector normal to the reflection plane. These components are relative to the **Reference Coordinate System**, which you select from the menu provided.

The **Reference Coordinate System** menu sets the local coordinate system, with Cartesian, cylindrical, or spherical systems available. The **Normal Vector** component labels update to reflect the chosen type of coordinate system. The menu lists those systems that were available when you opened the dialog.

Clicking **Apply** applies the reflection, and records the results in the *Output* window.

The following example reflects vertices about the X-Z plane of the laboratory coordinate system:



In this instance, the Y-axis is normal to the X-Z plane, so the entry for the **Normal Vector** is `[0, 1, 0]`. Clicking **Apply** results in the following reflection:



Using the Copy Selected Faces for Transform Option

Activating this option creates a copy of all faces that are selected and transforms the faces as specified.

In the example below, the highlighted faces are selected for a transform:



Note: At least one vertex, edge, or face must be selected to make the **Apply** button available. However, when *Copy Selected Faces for Transform* is activated, **Apply** is only available when faces are selected without edges or vertices.

When the following settings are applied:

- Transform Mode is set to Translate
- Copy Selected Faces for Transform is activated
- Translate Mode is set to Vertex Normal with a distance of 0.02 m

The result of the transform appears as follows:



Any feature curves that are attached to the selected faces are also copied. The selection set is also updated—the initial selection is lost and the latest selection contains only the copied faces and any attached feature curves.

Note: When faces or copies of faces are transformed, the transformed faces lose any CAD association.

The following example shows the result of a transform with the same settings as the example above, except *Copy Selected Faces for Transform* is deactivated:



Zero Translations

If **Transform Mode** is set to **Translate** and the translation parameters are specified as zero, the selected faces are copied and no transform takes place. Any CAD that is associated with the originally selected faces is maintained.

See Also

Translating Surface Entities

Auto Repairing

This tool automatically repairs errors reported by diagnostic checks.

This tool automatically repairs errors reported by the following diagnostics:

- Face quality
- Face proximity
- Self-intersections

The tool uses two fixing mechanisms.

- **Remesh** - Grows the problem area and surface remeshes, removing feature curves if necessary.
- **Patch** - Grows the problem area, deletes the local surface, and then invokes the hole filler.

First, the tool applies the remesher to faces identified as having a given error. If the remeshing succeeds in eliminating the issue or improving the quality, the tool takes no further action. If remeshing does not fix the error, the tool uses the patch option to remove the original error. You typically use the auto repair tool to resolve any errors that exist after surface wrapping.

Clicking **Auto-repair surface errors** starts the fixing process according to the above criteria. The tool only applies each fixing mechanism if you have run the diagnostic check for that particular type of error. For example, if you have not checked for face quality, the auto repair tool does not fix any poor quality faces.

The auto repair tool does not fix topological errors such as free edges, non-manifold edges, and non-manifold vertices. If these errors exist on the surface, you can fix them manually before using the auto repair process for the remaining issues. The auto-repair function is primarily designed for fixing single region or per-region cases that were surface wrapped. As such, the tool does not recognize interfaces between regions and can therefore potentially introduce errors on the faces if they are included in the active model.

Repair Features Mode

This section provides general guidelines on using the feature repair tools to create, diagnose, and fix feature curves, before running the surface remesher.

The tool can detect two types of feature errors: explicit and implicit. Feature errors that are not corrected can cause the meshers to produce poor quality surface and volume meshes.

Selecting **Repair Features** from the **Mode** menu provides access to the feature repair tools.



When the feature diagnostics are complete, feature curves with explicit errors are colored cyan, and feature curves with implicit errors are colored red.

The following steps are an example workflow for fixing feature errors:

- Mark features by edge angle, free edges, non-manifold edges, patch perimeters, and boundary perimeters.
- Start feature diagnostics to identify explicit feature errors (cyan) and implicit feature errors (red).
- Inspect the errors and mark valid features as non-erroneous.
- Fix all explicit feature errors.
- Inspect the feature curves and flag or unflag edges as feature curves, to preserve or ignore geometry details respectively.

You can select feature curves by using the browse buttons, or by using the mouse in the *Graphics* window. Selecting feature curves using the mouse allows greater accuracy when preserving or removing feature curves.

As you flag and unflag feature curves, the feature diagnostics automatically update and can consider a previous error valid. Similarly, it is possible for the addition of feature curves, with the intention of fixing larger errors, to increase the error count temporarily.

See Also

Diagnostics
Feature Diagnostics
Surface Remesher

Flagging Edges As Features

This tool creates part curves for the selected edges, in order to maintain geometry details when creating the surface mesh.

The tool becomes available when you select at least one edge.

1. You can assign edges to a different part curve using the *Part Surface / Part Curve Edit* tool.

   The tool adds marked feature curves to the **Surface Repair Edges** part curve, which is located in the simulation tree under the **Geometry** > **Part** > **geometry part** > **Curves** manager.

You typically use this tool when a required feature is missing, after having used the Mark Features tool to create feature curves automatically.

The following example flags the highlighted edge as a feature, in order to preserve the geometry detail during meshing.



2. Clicking **Flag edges as feature** creates a feature on the surface, which is marked with a thick black line.



See Also

Marking Feature Edges
Assigning Feature Edges To A Target Part Curve

Unflagging Edges As Features

This tool deletes the selected feature curves, in order to ignore geometry detail when surface remeshing.

The tool becomes available when you select at least one edge.

This example unflags all visible feature curves in the image below, as they do not represent important geometry details.



1. Select the unwanted feature curves using the mouse in the *Graphics* window.



2. Click the **Unflag feature edges** button to remove the selected feature curves from the surface.



See Also

Using The Select Curves Mode

Fixing All Feature Errors

This tool fixes all feature errors in explicit feature edges that the feature diagnostics identifies.

This tool systematically browses through each feature curve error and unflags the related feature curves. To make this tool available, run the feature diagnostics.

This tool only attempts to fix errors in explicit feature edges. Some errors in implicit feature edges are removed as a result of fixes to

neighboring explicit feature edges, but the tool does not fix these directly.

If implicit feature errors still exist after the tool has fixed the explicit feature errors, the cause is due to the underlying surface. In this instance, you can use the surface repair tools to fix surface errors, before repairing feature errors.

In this example, the cyan lines indicate explicit feature errors, as is typical when using feature diagnostics.



1. Clicking **Fix all feature errors** removes these explicit feature errors from the surface.



2. After fixing all errors automatically, inspect the surface and unflag any feature curves that you do not want to preserve. For example, in the image above it would be advisable to unflag all the visible feature curves, even though they are not errors.

See Also

Unflagging Edges As Features
Repair Surface Mode
Feature Diagnostics

Defining A Feature As Non Erroneous

This tool defines a feature error as not erroneous, and removes the error from the diagnostics summary.

The tool becomes available when you select at least one edge.

It is possible the feature diagnostics marks a part curve as an error, but you consider the part curve to be valid. In this instance, select the part curve and click **Not a Feature Error** to define the feature error as valid. The feature error summary reduces, and the feature diagnostics considers the curves valid for the current session.

Marking Feature Edges

This tool automatically creates feature curves that are based on edge angle, free edges, non-manifold edges, patches, and boundaries.

The tool is always available, as no previous selections are required. The tool adds marked feature curves to the **Surface Repair Edges** part curve, which is located in the simulation tree under the **Geometry** > **Part** > **geometry part** > **Curves** manager.

In most cases, automatically marking feature curves creates some erroneous and unwanted feature curves, which you can remove by using the fix all and unflag tools. For example, in the image above, it would be advisable to remove the features that are created where the branches of the manifold join.

Assign edges to a different part curve using the *Feature Curve Edit* tool.

1. Click **Mark Features** to open the *Mark Feature Edges Options* dialog:



The default options produce feature curves that are identified with a sharp edge angle of 31 degrees, and those curves on boundary perimeters in the active model. The active model is the parts that are selected in the Surface Repair Options dialog when starting surface repair.

2. Modify the feature edge options as required.
   The following example uses the Mark Features by Angle, etc tool with a sharp edge angle of 20 deg.



3. Click **OK** to create a part curve on all edges with an edge angle greater than the sharp edge angle.
   When complete, the *Output* window reports the number of part curves marked on the surface. In addition, the tool highlights all pre-existing and new part curves in the *Graphics* window.



See Also

Launching The Surface Preparation Tool
Unflagging Edges As Features
Fixing All Feature Errors
Using The Select Curves Mode

Offset Faces Edges Mode

The **Offset Faces/Edges** mode is used to inflate or translate faces or edges.

When inflating faces or edges, the selected entities are copied and offset a user-defined thickness. New faces are then created along the thickness to maintain connectivity between the original face or edge, and the new face or edge. When translating faces or edges, the selected entities are moved a user-defined distance and the connectivity with adjacent faces and edges is maintained.

Typical uses of the **Offset Faces/Edges** tool are illustrated in the following guided examples:

- Inflating a baffle surface
- Translating a surface
- Creating a zero-thickness baffle

Using **Offset Faces/Edges** mode is preferable to the wrapper inflation process, as it provides more control over the thickness and shape of baffle.

Selecting **Offset Faces/Edges** from the **Mode** menu provides access to the *Offset* group box.



1. Select the surface or edges to offset using the Selection and Display controls. With the selections made, you can then set the offset properties.

2. Alternatively, you can change the **Mode** to Repair Surface Mesh and fix the non-manifold edges with the surface repair tools.

   Offset Mode

   The *Offset Mode* determines whether you offset **Faces** or **Edges**.

   In **Faces** mode, **Offset Faces/Edges** inflates surfaces so they have a thickness, typically surfaces representing baffles or internal walls. It also translates or extrudes a surface, moving its position, while maintaining connectivity with adjacent surfaces.

   In **Edges** mode, **Offset Faces/Edges** inflates edges to create a two-dimensional surface, which can be used to create a zero-thickness baffle. It can also translate an edge, moving its position, while maintaining connectivity with adjacent edges and faces.

Inflate Offset Actions

The default options set *Offset Action* to **Bi-directional Inflate**, which inflates both sides equally, with a specified direction and thickness.

3. Setting the *Offset Action* to **Inflate** allows you to inflate one side with a specified direction and thickness.

   The diagram below illustrates both of these options:



In the above examples, the baffle surface is still non-manifold. However, the surface wrapper can account for this effect, and produce a manifold surface.

Translate Offset Action

4. Setting the *Offset Action* to **Translate** moves the selected surface or edge, with a specified direction and thickness, and maintains the existing connectivity.

   The diagram below illustrates this option:



Selecting a Direction Option

5. In each instance, you specify the direction of inflation using the **Direction** menu.

   Selecting **Face Normal** from the **Color Face by** menu in the *Display Control* group box colors positive face normals green, and negative face normals red. This feature is useful when choosing a **Direction**

   Four **Direction** options are available:

+ Normal Direction	Offsets the surface or edge in the positive normal direction. For edges, the average normal of the adjacent faces is used to determine the inflation direction. For **Bi-directional Inflate,** this option has no obvious effect, as the tool inflates in both the positive and negative direction.
- Normal Direction	Offsets the surface or edge in the negative normal direction. For **Bi-directional Inflate**, this option is not available, as it has the same effect as using the **+ Normal Direction** option.
Specified Direction	Offsets the surface or edge in a user-defined direction. This option requires **Direction Vector Components** in a given **Coordinate System** to specify the direction. This option is not available for cylindrical or spherical

	systems. You can only define a **Coordinate System** for cartesian systems.
Seed Point	Offsets the surface or edge in the normal direction, towards a user-defined seed point. In cases where there are multiple connected surfaces or edges, the tool uses a majority rule to determine the offset direction. This option is useful when using Bi-directional Inflation on surfaces or edges that have normals in various directions.

Setting a Thickness or Distance

6. With the direction option selected, you then set the **Thickness**, which defines the inflation thickness.
   a) Similarly, when using **Translate**, you set the **Distance**, which defines the translation distance.
   b) When using **Bi-directional Inflation**, the **Thickness** sets the total inflation thickness, meaning the tool inflates each side by half of this value.

Offsetting Faces

7. When you have provided all necessary settings, click **Offset** to carry out the operation.

See Also

Repair Surface Mode
Display Controls
Entity Selection

Inflating A Baffle Surface Using The Offset Tool

This example shows the inflation of a baffle surface using the **Bi-directional Inflation** option. Baffle surfaces can connect to the surrounding geometry, be free-standing, or intersect the geometry. In this case, the baffles intersect the geometry.



You first select the baffle surface using the Selection Controls. In this case, the **Select faces / edges on part surfaces / part curves** tool selects the three internal baffle surfaces.



You then select the required *Offset Action*. This example uses the **Bi-directional Inflation** option to inflate both sides of the baffle. The **Direction** is **+ Normal Direction**, which inflates the baffle in the normal direction of each face. The **Thickness** is an arbitrary value, with any allowable units.

To carry out the inflation, click **Offset**. The *Graphics* window updates with the modified surface.



If the starting surface contained feature curves, then these curves are also inflated and included on the new surface. Any new faces created by the process belong to the same surface as the originating face.

1.

See Also

Entity Selection

Translating A Surface Using The Offset Tool

This example shows the translation of a surface using the **Translate**option. This feature allows you to extrude or modify the surface geometry before surface meshing. In this case, the tool extrudes a manifold outlet pipe, increasing its length.



You first select the required surface using the Selection Controls. In this case, the **Select Zone** tool selects the ring of end faces:



You then select the required *Offset Action*, which is **Translate** in this case. The **Direction** is **+ Normal Direction**, although you could also use **Specified Direction** and **Seed Point**. The **Distance** is an arbitrary value, with any allowable units.

To carry out the translation, click **Offset**. The *Graphics* window updates with the modified surface.



The tool also translates any feature curves on the initial surface with the selected faces. The above translation does not create new faces.

1.

See Also

Entity Selection

Creating A Zero Thickness Baffle Using The Offset Tool

This example shows the inflation of edges using the **Offset Faces/Edges** tool in **Edges** mode. This feature can be used to create a zero-thickness baffle. In this case, a two-dimensional orifice is created in a section of cylindrical pipe.



You first select the required edges using the Selection Controls.



Then, select the required *Offset Mode*, which is **Edges** in this case. The *Offset Action* is set to **Inflation** so that the edge is inflated in only one direction. The **Direction** is set to **Seed Point**, with the point placed in the center of the pipe. These edges are inflated in a normal direction towards the center of the pipe.



To carry out the inflation, click **Offset**. The *Graphics* window updates displaying the zero-thickness baffle.



If a feature curve is selected, the corresponding inflated edge is also a feature curve. Also note that when inflating or translating an edge in a normal direction, the orientation of connected faces determines the direction. This orientation can cause unexpected results when inflating non-manifold edges, as the orientation of three or more edges is considered. This effect is illustrated in the example below. The blue edge is non-manifold and the pink edge is the result of an offset in the positive normal direction.



1.

See Also

Entity Selection

Edit Connected Surfaces Mode

The **Edit Connected Surfaces** mode identifies the connection or disconnection between geometry surfaces, and allows deletion of disconnected surfaces.

The tool can identify surfaces that must be retained, as they are important, while others can be unnecessary facets that you can delete.

You typically use **Edit Connected Surfaces** mode on a wrapped surface to remove geometric details, retained by the surface wrapper, which you do not require for the remaining meshing process. However, you can also use the **Edit Connected Surfaces** mode on imported, remeshed, or auto repaired surfaces.

This tool deletes connected surfaces as a whole. To delete individual faces, use the Surface Repair options in the **Repair Surface Mesh** mode.

Selecting **Edit Connected Surfaces** from the **Mode** menu provides access to the *Edit Connected Surfaces* group box.



1. To identify the connected and disconnected surfaces, click **Start Diagnostics**.

   When the diagnostics are complete, the *Edit Connected Surfaces* table populates with a summary of connected surfaces on the part:



The table contains a summary of the face count and area (m2) of each connected surface within the part. By default, the table sorts the data by descending face count.

2. You can resort the data by clicking the **Face Count** or **Area (m^2)** column headers.
   The tool also provides a summary of total **Surfaces**, **Faces, and Area** at the bottom of the group box.

   You can change the mode at any time to access one of the other tool sets. However, if you change the surface, you regenerate the connected surfaces information by clicking **Start Diagnostics**, when returning to **Edit Connected Surface** mode. Additionally, any changes that you make using **Edit Connected Surfaces** (such as deleting a surface) update the surface errors diagnostics in the surface repair tool.

   There are three different methods for using **Edit Connected Surfaces** mode:

   - Selecting surfaces in the part
   - Highlighting a selected face in the summary table

- Deleting unwanted surfaces

See Also

Changing The Parts In The Active Model
Using Threshold Diagnostics
Auto Repairing
Importing Surface Data into STAR-CCM+
Surface Remesher
Surface Wrapper

Selecting Surfaces Using The Summary Table

You can use the summary table to select part surfaces in the *Graphics* window, allowing you to locate and evaluate each connected surface.

To select surfaces using the summary table, you first generate the connected surface diagnostics. You can then activate the **Select Highlighted** option, which highlights selected table entries in the *Graphics* window.

Furthermore, two more options provide further help in recognizing and identifying connected surfaces:

- **Reset View** resets the view to zoom in on the displayed faces, edges, and vertices
- **Reset Displayed** resets the displayed faces, edges, and vertices to the error and two neighbors

1. If you use **Reset Displayed**, you can redisplay the entire part by clicking **Show All Faces** within *Display Controls*.
   The following example shows two highlighted table entries with the **Select Highlighted** option activated, which selects the corresponding surfaces.



The selected surfaces appear as follows:



2. To deselect faces, you can click anywhere in the white space in the scene or click **Select None** within *Selection Controls*.

Highlighting A Selected Face In The Summary Table

You can use the **Edit Connected Surfaces** tool to identify the table entry or entries that correspond to a selected face or faces.

To highlight a selected face in the summary table, you first generate the connected surface diagnostics. You can then select a face or faces of interest in the *Graphics* window, and click **Highlight Selected**. This option highlights the appropriate entry or entries in the summary table.

The following example shows a selected face on the part surface.



1. Clicking **Highlight Selected** highlights the table entry, indicating the connected surface group to which the face belongs.



2. If you also activated the **Select Highlighted** option, the entire connected surface would appear highlighted in the *Graphics* window.



Deleting Unwanted Surfaces Using The Summary Table

You can use the *Edit Connected Surfaces* tool to identify and delete unwanted surfaces.

To delete unwanted surface using the summary table, you first generate the connected surface diagnostics. You can then mark surfaces for deletion in the summary table using the following methods:

- Tick the appropriate checkbox in the left column
- Click and highlight the required table row (using `<Shift>` and `<Ctrl>` for multiple selections), and click **Highlighted**, next to **Toggle**, at the top of the table

You can also use the **All** button, next to **Toggle**, to toggle the current state of all surfaces. That is, any marked surfaces become unmarked, and any unmarked surfaces become marked. To mark all surfaces for deletion, you can tick the checkbox in the left-most column header, marking and unmarking all surfaces at once.

Marking a surface for deletion colors the surface red in the *Graphics* window. The tool also summarizes marked surfaces at the bottom of the group box.

1. Once you have marked all the desired surfaces, click **Delete Marked** to delete the surfaces.
   If you accidentally delete a surface, you can use undo.

   The following example shows two marked table entries, coloring the appropriate surfaces red.



The marked surfaces appear as follows in the *Graphics* window.



2. Clicking **Delete Marked** deletes the surfaces from the part.



Merge And Imprint Mode

The **Merge/Imprint** mode lets you imprint sets of faces and edges onto each other.

The tool also lets you merge two or more overlapping (co-planar) surfaces or sets of faces. When using **Merge/Imprint** on a single part, the result is either a common shared surface or no surface at all (that is, the common area is removed). If there are multiple parts, the tool creates part contacts for the shared surface.

Similarly, for feature curves belonging to the same parts, the tool can imprint two or more edges onto each other, effectively zipping them together. You can also imprint edges onto faces and imprint faces onto edges, depending on the requirement. No edge options are available when imprinting across multiple parts. In all instances, the faces or edges do not need to belong to a closed surface in order to use the tool.

Selecting **Merge/Imprint** from the **Mode** menu provides access to the *Merge/Imprint Tool* group box.



The *Merge/Imprint Tool* is applicable to two situations:

- When the surfaces belong to a single part
- When the surfaces belong to multiple parts

The details below are relevant to both situations. Review this information before using the tool for a specific example.

Merging and Imprinting Surfaces

You add the required surfaces to *Source* and *Destination* groups, so that the projection scheme used for merging and imprinting knows the direction in which to travel.

Note: If there are no vertices within the imprint tolerance that can project from one surface to the other, the imprint fails. To overcome this limitation, you can add vertices to the surface with the remesh or split faces tools.

In general, two scenarios exist when the merging and imprinting of overlapping surfaces is required:

- The perimeter or area of the overlap is the same
- The perimeter or area of the overlap is different

The following diagram illustrates an example where the overlapping boundaries have the same perimeter or area:

For cases where the two surfaces have a different perimeter or area, it is possible to select either surface as the source. The following diagram illustrates an example where the smaller surface is the source

Imprinting Edges to Faces

When imprinting edges onto faces for single parts, you use the same approach as described above. In this instance, the edges are typically the source, and the destination is one or more faces that are imprinted with the shape of the edge. Details of a procedure to use with the **Multi-Part Imprint** mode are given in Imprinting Baffle Edges onto Faces.

A typical example of imprinting edges onto a surface is a baffle surface perpendicular to an exterior surface. You can imprint one or more edges along the side of the baffle to the exterior surface, thus letting you convert the baffle to an interface for meshing purposes.

Example of Merging and Imprinting

The following example merges and imprints two closed surfaces together. Shown below are two closed bodies that are touching, but not merged.



As shown below, the result of the **Merge/Imprint** was to retriangulate both bodies, so that there is a one to one correspondence at the interface.



If both bodies belong to the same part, the tool provides the choice of keeping or deleting the shared surface. If both bodies belong to two different parts, the tool creates a part contact between the two surfaces.

Considerations When Using the Merge/Imprint Tool

The following conditions apply when using the **Merge/Imprint** tool for both single and multiple parts:

- The surfaces do not need to be unique to the overlapping area. For example, a closed body, consisting of one surface, can merge with the overlapping surface of another body.
- When merging faces with other faces, the tool does not generate feature curves. If feature curves are required, you create them manually.
- The surface faces and edges for single parts do not need to belong to a closed and manifold part. In other words, the tool allows use of planar and open surfaces.
- The quality of the original surface tessellation is not critical.
- You can use the surface remesher after merging to improve the quality of the resultant surface.

See Also

Changing The Parts In The Active Model
Surface Remesher
Converting Boundaries to Interfaces
Imprinting Baffle Edges Onto Faces
Splitting Edges And Faces
Remeshing Faces
Zipping Edges
Using Part Contact Data to Create Interfaces

Imprinting A Single Part

This options allows you to merge two or more overlapping (co-planar) surfaces from the same part.

The following image shows the **Merge/Imprint** tool in **Single Part Imprint** mode.



1. To use the **Merge/Imprint** tool in **Single Part Imprint** mode, first select **Faces** or **Edges** from the **Mode** menu for both the *Source* and *Destination*.
   The tool allows any combination of face and edge imprinting. You can select the source and destination faces and edges using the Selection Controls.



2. With the selections made, click + for either the source or destination to add the faces or edges to the group.
   In the *Graphics* window, faces or edges added to the **Source** appear green, and the faces or edges added to the **Destination** appear red.
3. You can remove faces or edges by selecting them in the display, and clicking − for the source or destination.
4. To remove all faces from the source or destination, click ✕.

5. If you select the opposite group of faces for source and destination, you can swap them by clicking .

When merging/imprinting faces to other faces, the common surface belongs to the destination surface by default. Options exist to delete the common surface (that is, the two bodies merge as one), or assign the common surface to a specified part surface.



Selecting **Specify Part Surface** activates the **Target Part Surface** option, allowing selection of existing part surfaces:



Clicking **New** creates new surface using the *Create New Surface* dialog:



6. Enter the **New Surface** name and click **OK** to define the new surface, which appears in the simulation tree.

   Merge Angle limits the operation to faces that are at a user-defined angle to each another. For example, the default angle, `45 deg`, means that the tool merges any source and destination faces that are less than 45 degrees to each another. This feature is useful when a curved surface touches a planar surface, such as a tire contact patch on a road surface. This option is only available when merging faces to other faces.

Tolerance limits the operation to faces and edges that are a user-defined distance from one another. The tool calculates the **Tolerance** value when merging edges to other edges (hence, it is not available).

7. Click **Merge/Imprint** to carry out the operation.

   When complete, the *Output* window shows feedback from the operation.



In situations where the merge was unsuccessful, the following error message appears.



In this instance, try reversing the selections that you made for the source and destination faces or edges. In cases where a relatively large distance separates the faces or edges, having finer triangulation for the items improves the merge, as increasing the tolerance value.

Imprinting Multiple Parts

This options allows you to merge two or more overlapping (co-planar) surfaces from multiple parts.

The following image shows the **Merge/Imprint** tool in **Multi-Part Imprint** mode.



The multi-part imprint mode differs from the single part mode in that the tool calculates source and destination pairs automatically. These calculated pairs appear in the *Part Browser*.

1. To use the **Merge/Imprint** tool in **Multi-Part Imprint** mode, first review and set the **Tolerance** value and any advanced parameters that you require.
   To access and edit advanced parameters, activate **Show Advanced Parameters**. These values affect how the tool finds pairs. For more details, see the single part imprint section.
2. Click **Find Pairs**.
   Finds all relevant part pairs that are based on the supplied values.
   a) If you enter an unsuitable *Tolerance* value resulting in an excessively long search, click **Stop** to stop the process and enter a new *Tolerance* value.

b) If the tool does not find all pairs, then change the *Merge Angle* (activate *Show Advanced Parameters*) or *Tolerance,* and click **Find Pairs** again to repeat the process.

Once **Find Pairs** is complete, the *Output* window shows a summary of all pairs that are found on a per-part basis. The *Output* window displays the text `MERGED` next to a previously merged pair. You can merge and imprint pairs of parts more than once by selecting source and destination faces manually, but the automatic find does not return previously merged pairs of parts



If multiple pairs are present, the **Browse Pairs** selector becomes active, allowing you to display each part pair. Clicking the forward or backward arrows cycles through each pair and displays them in the *Graphics* window.

The image below shows the **Merge/Imprint** tool after a complete search.



Each part pair appears green and red in the *Graphics* window, to indicate the source and destination surface respectively. Previously imprinted pairs remain in the browser, but are inactive (that is, they do not appear green or red) for the current *Surface Preparation* session.

3. If you wish to reverse the source and destination, click to swap them.
   You can also use the **Priority Indexing** button to define which part is the source and which part is the destination by comparing the **Priority Index** between the part pair. This option is useful for simulations with many parts.

The **Destination** and **Source** drop-down menus become available in the *Part Browser* after the search process is complete. The first drop-down menu is the **Destination** menu which contains all parts that form a pair with other parts. The second drop-down menu is the **Source** menu

which contains all parts that form a pair with the selected **Destination**part.



4. To remove pairs from the *Part Browser*, use the menus to select the appropriate part pair, and click to remove them for the current *Surface Preparation* session.
5. When you have identified all part pairs, click **Imprint All** to imprint all part pairs at once. Alternatively, click **Imprint Pair** to imprint the part pair that is selected as the **Destination** and **Source** in the *Part Browser*.

   When the default advanced parameters are maintained, the tool generates a new part contact with a conformal mesh for each imprinted part, which represents the common surface between the two parts. You can review the new surfaces in the *simulation* tab.



See Also

Using Part Contact Data to Create Interfaces
Imprinting A Single Part
Merge And Imprint Advanced Parameters

Imprinting Pairs More Than Once

In some circumstances, it is necessary to imprint a pair of parts more than once in order to apply different tolerance values to each imprint.

For example, in the geometry that is shown below, a large tolerance value that is applied to both ends of the top part results in the small cylinder being distorted after the imprint. Different imprint tolerances are required for each end of the top part.



To avoid a problem with the small end of the top part, set the *Tolerance* value to a small value initially for the automatic find and the first imprint. This setting results in a match as shown below.



Perform a second imprint with a larger *Tolerance* value. This time, however, the source and destination faces must be selected manually using the *Selection Controls*. Take care to avoid selecting faces around the site of the first imprint, as the larger tolerance value could affect the selection. A selection such as shown below would be acceptable for this example.



After the imprint operation, both ends of the top part will be connected to the base with appropriate matching areas. These interfaces are shaded green in the image below.



The **Imprint All** option only works for parts that have not yet been imprinted. Successive imprints such as that demonstrated here must be determined manually.

Editing Part Priorities

The **Edit Priorities** function defines the destination and source for each part pair by comparing the priority indices between part sets.

The part with the higher priority index becomes the destination (red), and the part with the lower priority becomes the source (green). This feature is useful when using merge/imprint on many parts, as it does not require manual assessment of each part pair.

Note the following considerations when editing part priorities:

- As **Priority Indexing** requires multiple parts, the feature is only available when in **Multi-Part Imprint** mode.
- Parts cannot belong to multiple part sets. Adding a part to a new part set removes it from the original part set.
- If you make a mistake when setting priorities, you can use swap to swap the destination and source part.

- Priority values must be greater than or equal to 1.

To edit part priorities, first click **Priority Indexing** within the *Part Browser* group box.



The *Merge/Imprint Priority Indexing* dialog appears. To create a **Priority** for a **Part Set**, click **Create new Set**, which adds a row to the table.



To add a **Part Set**, click the ellipsis (**...**) next to the row, which opens the Select Parts dialog. Select the parts that you want to add to the set and click **OK**.



Continue creating sets and adding parts as required.



To change the priority of a **Part Set** you can drag and drop the row higher or lower in the table, or change the value in the **Priority** column.



You can also delete **Part Sets** by ticking the appropriate checkbox and clicking **Delete Marked**.



Click **OK** to confirm the priority indexing. Browsing part pairs now define the part with the higher priority as the destination part.

Editing Part Priorities in the Simulation Tree

Editing part priorities using the *Merge/Imprint* tool adds a *Metadata* tag to the part properties.



When you are not using the *Surface Preparation tool*, you can manually edit the priority index of the part by changing the integer value in the *Metadata* property.

Imprinting Baffle Edges Onto Faces

It is possible to imprint the edges of a zero-thickness baffle part onto its adjacent faces while retaining the baffle as a separate part.

The zero-thickness can then be retained during meshing. This section describes a procedure for doing this imprint.

Take the following geometry as an example. There is a single zero-thickness baffle that is contained within a solid volume. The edges of this baffle are imprinted onto the cylinder so that the baffle can then be meshed as a conformal interface.



1. To begin this procedure, launch the *Surface Preparation tool* and enter the **Merge/Imprint** mode. Within the *Merge/Imprint Tool* panel, click the radio button for the **Multi-Part Imprint** mode.
2. Click the **Find Pairs** button within the *Parts Browser*. For this example, the source is initially discovered as the outer cylinder, and the destination as the baffle.
3. Click (**Swap**) to reverse this order, so that the baffle faces are set to be the source.
4. To select the faces on the baffle, click the green box within the *Source* group box.



5. Now deselect the source set by clicking the button.

The selection is not removed from the *Graphics* window: only the faces from the source set on the *Merge/Imprint Tool* are removed.

6. Within the *Source* group box, change the **Mode** to **Edges**.



7. Now move to the *Selection Controls* panel within the *Surface Preparation tool*, right-click (**Select Attached Edges**) and choose **Options...**



8. In the *Select Attached Edges Options* dialog, clear **Select All Edges** and tick **Select Perimeter Edges**.
9. Click **OK** to close the dialog.



10. Click (**Select Attached Edges**).
    The edges on the outside of the baffle surface are selected.

11. Click the button within the *Source* group box.
    The newly selected edges are added to the source set.
12. The source edges and destination faces are now fully set. To perform the imprint operation, click the **Imprint Pair** button.
    The imprinted edges are highlighted in the *Graphics* window on completion.



Using The Exact Face Match Method

Set the **Method** to the **Exact Face Match** option only when a part pair already has one-to-one correspondence between their coincident faces and vertices.

This conformal match between parts can occur in surface meshes exported from certain surface preparation software. Selecting this option speeds up the merge/imprint process, as the tool does not need to imprint and create identical faces on both surfaces. Instead, the tool

simply merges the coincident vertices and faces, and creates a part contact (interface for regions) between the two parts.

When using this option, set a small **Tolerance** value, as the two surfaces are coincident. This tolerance prevents the undesired merging of faces and vertices in thin parts.

As an example of surface mesh requirements, consider the two parts that are shown below. The first image shows the contacting area of a cylinder part:



This second image shows the conformal and coincident surface mesh on the contacting block part:



The contacting surface on the cylinder and the circular area on the block are identical.

Using The Conformal Resulting Mesh Type Option

Use this option to achieve a conformal mesh between contacting surfaces.

The following geometry shows two parts that have contacting surfaces between each other. To achieve a conformal mesh between the contacting surfaces, before you imprint the parts, set *Resulting Mesh Type* to **Conformal**.



After imprinting, the resultant surfaces are conformal, which means they have the same surface discretization.



By default, in-place part contacts are created between the contacting surfaces. The part contacts appear in the simulation tree:



By hiding the cylindrical surface, you can see that the surface discretization on the contacting surfaces is identical.



Using The Non Conformal Resulting Mesh Type

Non-conformal imprinting provides a way to imprint parts efficiently without spending time imprinting CAD parts conformally when it is unnecessary.

STAR-CCM+ allows you to imprint parts non-conformally. This process is useful when:

- a conformal interface is not necessary to achieve solution accuracy
- there are many contacting parts
- each part is within many pairs
- there are gaps between the CAD parts

During the process of non-conformal imprinting, pairs of split surfaces are created on contacting parts without altering the individual parts or creating shared faces.

The following geometry shows two parts that are in close proximity with a small gap between the surfaces.



To achieve a non-conformal mesh when imprinting these surfaces, before you imprint the parts, activate *Show Advanced Parameters* and set *Resulting Mesh Type* to **Non-Conformal.**



After imprinting the parts, the resultant surfaces are non-conformal, which means that the surface discretization on the imprinted surfaces does not match. Each part remains separate and there is no shared surface between each part. However, the original discretized surfaces are cut and a pair of split part surfaces are created which represent an in-place part contact between the surfaces.



Taking a closer look at the pair of split surfaces, you can see that the surface discretization on each split part surface is non-conformal:



The in-place part contact which represents the pair of split part surfaces appears in the simulation tree:



Using The Boolean Subtract Preprocess Option

This option subtracts the source part from the destination part before performing the imprint operation.

The geometry of the sources part is changed, but the geometry of the destination part is maintained.

Note: To perform a boolean subtraction, the input parts must be closed and manifold. Therefore, selecting the Preprocess option, Boolean Subtract, deactivates the face selection controls, as the tool can only imprint two complete parts. If the parts are not closed and manifold, there is no subtraction, but the parts are imprinted.

The following example shows two parts that penetrate one another.



Selecting the Boolean Subtract option and imprinting these two parts gives the following result, with the intersecting surface merged and imprinted.



By hiding the rectangular part, you can see that the rectangular part (destination) has been subtracted from the cylinder (source).



Using The Opposing Outward Face Orientation Option

Selecting the *Face Orientation* option, *Opposing Outward*, provides an imprinting mode which ensures that only faces that are facing each other with opposing outward normal directions are imprinted or merged.

When imprinting thin, penetrating parts, the imprinter can choose the wrong side of the destination part with which to imprint or merge the source part.

The following geometry shows two parts that intersect each other.



When these parts are imprinted with the *Opposing and Aligned* option selected, the surface of part 2 is imprinted onto the wrong side of part 1.



If these parts are imprinted with the *Opposing Outward* option selected, the surface of part 2 is imprinted onto the correct surface of part 1:



Merge And Imprint Advanced Parameters

Advanced parameters are available to set when merging and imprinting multiple parts using the surface repair panel.

Method	Specifies how faces are chosen and matched for imprinting.	
	Discrete Imprint	The imprint is carried out using the tessellated surface mesh of the parts. The original tesselated faces are cut to ensure that the imprinted faces match.
	Exact Face Match	The imprint is carried out using the tessellated surface mesh of the parts. Coincident vertices and faces are merged and a part contact is created between the two parts. Specify this option for part pairs that already have exact one-to-one correspondence between the faces and vertices. See Using the Exact Face Match Method.
Resulting Mesh Type	Determines if the imprinter merges the faces of the input parts.	
	Conformal	The imprinter merges faces of the imprinted parts to produce a conformal mesh on the output surface.
	Non-Conformal	The imprinter does not merge faces of the imprinted parts and a non-conformal mesh is

		produced on the output surface.
Preprocess	Allows you to specify if any boolean operations are performed before parts are imprinted.	
	No Preprocess	No boolean operations are performed before parts are imprinted.
	Boolean Subtract	A boolean subtract operation is performed before parts are imprinted.
	Intersect Surfaces (available for non-conformal imprinting only)	An intersect operation is performed to cut the existing surfaces before parts are imprinted. Coincident surfaces that lie inside other surfaces are also cut and automatically included in the resulting part surfaces.
Resulting Surfaces (available for conformal imprinting only)	Allows you to specify the surface that the resulting imprinted faces are added to. This option is only available for conformal imprinting as non-conformal imprinting creates new part surfaces by default. See Imprint Parts Resulting Part Surfaces.	
	Input Part Surfaces	The imprinted faces that are created are added to the part surface of the input part.
	New Part Surfaces	The imprinted faces that are created are added to a new part surface.
Face Orientation	Allows you to specify the orientation of the matching face normals.	
	Opposing and Aligned	Only target faces that are aligned and opposed outward

		from the normal of the matched source faces are merged.
	Opposing Outward	Only target faces that are opposed outward from the normal of the matched source faces are merged. See Using the Opposing Outward Face Orientation Option.
	Aligned Outward	Only target faces that are aligned outward from the normal of the matched source faces are merged.
Merge Angle	Specifies the maximum angle between the normals of the faces at which parts merge. Only pairs of faces which lie within the specified *Value* and have normal directions which differ by less than this *Merge Angle* are merged.	

See Also

Using The Opposing Outward Face Orientation Option
Imprint Parts Resulting Part Surfaces
Using The Exact Face Match Method
Imprinting Multiple Parts

Intersect Boolean Faces Mode

The *Intersect/Boolean Faces* tool allows you to intersect or perform boolean operations on two face sets.

A face set is a collection of faces. For face intersection, you can additionally use a user-defined plane instead of one of the face sets.

Setting **Mode** to **Intersect/Boolean Faces** provides access to the *Intersect/Boolean Faces Tool* group box.



This tool has the following modes:

- Intersect Faces
- Boolean Faces

Intersect Faces Mode

The **Intersect Faces** mode is similar to the boundary intersection feature, in that it intersects two or more separate surfaces to produce a merged result. Both tools produce identical results for a given face set.

The **Intersect Faces** mode provides the following functionality:

- Intersection of individual faces.
- Intersection of faces belonging to different surfaces.
- Ability to undo/redo the operation.

The following sections provide examples of using the **Intersect Faces** mode:

- Intersecting Two Face Sets

- Intersecting Faces with a Plane

Note: The tool does not create feature curves at the intersection plane of the two surfaces. If feature curves are required, create them manually.

Boolean Faces Mode

The **Boolean Faces** mode unites, intersects, and subtracts user-defined face sets.

To be a valid selection for a boolean operation, make sure that face sets meet the following requirements:

- Each face set is composed of single or multiple closed volumes that do not contain non-manifold edges or vertices.

  If the face set does not meet this requirement, the *Output* window indicates which face set is not closed or manifold.
- Face sets cannot contain faces that are directly involved in a part contact. To perform a boolean operation on such a face set, remove the part contact.

The tool aims to maintain part contacts that are indirectly involved in the boolean operation, such as a face in a part contact that is connected to a face in a face set.

The following section provides an example of using the **Boolean Faces** mode:

- Performing a Boolean on Faces

See Also

Changing The Parts In The Active Model
Performing A Boolean Operation On Faces

Intersecting Two Face Sets

Use this procedure to intersect two face sets.

1. Set **Intersect Mode** to **Intersect Faces**.
2. Select the faces or surfaces belonging to the first *Face Set* using the Selection Controls.
3. In the *Intersect/Boolean Faces Tool* , click + for *Face Set 1* to add the selected faces to the set.

4. Ensure that **Mode** is set to **Faces** and repeat this selection and adding process for *Face Set 2*. In the *Graphics* window, *Face Set 1* appears green and *Face Set 2* appears red.



5. You can remove faces or edges by selecting them in the display, and clicking for *Face Set 1* or *Face Set 2*.
6. To remove all faces from *Face Set 1* or *Face Set 2*, click .
7. Click **Intersect** to intersect the selected surfaces.

   The following example intersects the red and green surfaces.



After intersecting the surfaces, the tool retessellates the surfaces at the intersection plane and adds newly created faces to the nearest neighboring surface. The intersection operation also selects the edges at the join so that you can easily define them as feature curves.



See Also

Entity Selection

Intersecting Faces With A Plane

Use this feature to intersect geometry parts based on a plane that you specify.

To intersect a set of faces with a plane:

1. Set *Intersect Mode* to **Intersect Faces**.
2. Select the faces or surfaces that you want to intersect.
   a) Use the Selection Controls to select the appropriate faces.
   b) In the *Intersect/Boolean Faces Tool*, click + for *Face Set 1* to add the selected faces to the *Face Set*.
3. To define the intersection plane, set the **Mode** for *Face Set 2* to **Plane**.
4. Specify the Plane Origin and Plane Normal coordinates.
5. Set the **Coordinate System** to the system in which the origin and normal coordinates are defined.

6. Click **Intersect**.

See Also

Entity Selection

Specifying The Plane Origin

You can specify the Plane Origin coordinates using three possible methods.

Set specific coordinates in the Plane Specification group box

1. Enter Plane Origin X Y Z values.
2. In the *Graphics* window, review the location of the Plane Origin point to ensure that it is correct.

Define coordinates interactively in the Graphics window

3. Next to the Plane Origin X Y Z coordinates, click **Pick**.
4. In the *Graphics* window, click an area on the surface of the geometry that determines the location of the Plane Origin.

Define coordinates using the Display Plane Tool

5. In the *Plane and Section Tools* group box, activate *Display Plane Tool*.
6. In the *Graphics* window, click and drag the Plane Origin **Pick** tool until the required location is achieved.

   The Plane Origin coordinates can be set as one of three Plane Normal points when specifying the Plane Normal.

See Also

Specifying The Plane Normal
Specifying The Plane Normal

Specifying The Plane Normal

There are two Normal Specification options available to define the Plane Normal coordinates.

Vector



1. Set *Normal Specification* to **Vector**.
2. To specify the Plane Normal coordinates, either:
   - Enter Plane Normal X Y Z values, or
   - Define Plane Normal Vector coordinates using the Display Plane Tool.

     See Specifying the Plane Origin.

Picked points



3. Set *Normal Specification* to **Picked Points**.

4. To specify the Picked Points coordinates, either:
   - Specify three independent Plane Points coordinates that represent the Plane Normal. Using the same process as when specifying the Plane Origin, enter the Plane Points' X Y Z values directly, or define the Plane Points interactively in the Graphics window.
   - Specify three Plane Points coordinates using the Plane Origin coordinates for one of the Plane Points. Activate *Use plane origin for plane point #1*, then specify the remaining Plane Points' using the same process as for specifying the plane origin.

See Also

Specifying The Plane Origin
Specifying The Plane Origin

Performing A Boolean Operation On Faces

A boolean operation allows you to unite and intersect two face sets, or to subtract the source face set from the target face set.

Note: The remaining faces after a boolean operation belong to the original part and part surface.

To perform a boolean:

1. Set *Intersect Mode* to **Boolean Faces**.
2. Define the *Source* and *Target* face set with the selection buttons.

   See Intersect/Boolean Faces Mode for a list of requirements for the face sets.
3. Set **Boolean Type** to one of the following options:

   See Performing a Boolean Operation on Faces.
   - **Unite** — Removes any overlapping volumes between the face sets.
   - **Intersect** —Maintains only the overlapping volume between the face sets and removes all other volumes.
   - **Subtract** — Removes the source face set from the target face set.
     - To get the required subtracted volume, you can use the swap button to change the source and target selections.

- To subtract multiple source volumes from the target volume, first unite the source volumes before performing the subtract operation.



4. Click **Boolean** to perform the boolean operation.

Examples of Boolean Operations

These examples use the geometry parts shown in the following image to demonstrate the different boolean operations. These two spheres overlap so that they share a common volume. The green sphere is the source face set and the red sphere is the target face set.



Unite

Using the unite operation removes the common volume between the two face sets.



Intersect

Using the intersect operation leaves only the common volume between the two face sets.



Subtract

Using the subtract operation removes the source face set (green) from the target face set (red).



See Also

Intersect Boolean Faces Mode

Entity Selection

The buttons in the *Selection Control* tool allow you to select faces, edges, and vertices by specifying a zone, threshold, or surfaces, or by growing and shrinking from an existing selection.

These options are useful when you select many items.



The following list summarizes other methods of selecting entities:

- Right-click surfaces in the *Graphics* window to select associated entities.
- Click individual entities manually in the *Graphics* window to select them.
- Double-click an edge to select a continuous line of edges, whose face normals are similar to one another.
- After you have run surface diagnostics, double-click a free edge to select all connected free edges.
- Double-click a face to select a continuous surface that is constrained by feature curves or sharp edges.

Filtering Selection Entities

The **Faces**, **Edges, and Vertices** filters restrict the entities available for selection. The default options activate the **Faces** filter, meaning the mouse only selects faces, with multiple selection possible by pressing `<Ctrl>` and selecting more items. Activating the **Edges** and **Vertices** filters allows selection of edges and vertices respectively.

For example:

1. Only tick the checkbox for **Faces**.

The cursor and zone selection only selects faces. The box opposite each checkbox lists the number of selected entities for that filter. The state of the filters does not affect the grow and shrink options.

Clearing Selections

To deselect entities:

2. To deselect individual entities, click selected entities in the *Graphics* window, while pressing the `<Ctrl>` key.
3. To deselect all selections, click anywhere off the part in the *Graphics* window.
4. To deselect all faces, edges and vertices independently, click the clear filter buttons .
5. To deselect all entities, click the larger button.
   The filters do not affect the deselection of entities. That is, the selection entities cannot be "locked" by deactivating a filter. To store a selection for later use, you use sets.

See Also

Set Controls
What Are Free Edges
Selecting With A Zone
Selecting With A Specified Threshold
Selecting Surfaces And Curves With The Selection Controls
Selecting Surfaces Or Curves With The Graphics Window
Using The Chase Edge Selection Mode
Using The Select Curves Mode
Storing Sets

Selection Options

There are four selection options available for use with the zone, threshold, and surface selection tools.

You can access different selection options by right-clicking a selection control button. The options are:

- Select

- Add Select

- Deselect

- Subset Select

The following examples use the zone selection tool. However, the same applies for the threshold and surface selection tools. When using the *Select Part Surface/Curves* dialog, only edges that have been defined as part curves are selected.

Select

This option is the default selection option. Using **Select** deselects any previously selected entities in preference to the new selection. The example below illustrates the use of the select option.



Right-clicking to close the polygon deselects the original selection and the select the following faces.



Add Select

This option adds entities to an existing selection. The example below illustrates the use of this option.



Right-clicking to close the polygon selects the following faces.



Deselect Zone

This option deselects entities from an existing selection. The example below illustrates the use of the deselect option.



Right-clicking to close the polygon leaves the following faces selected.



Subset Select Zone

This option selects a new group of entities from an existing selection. The example below illustrates the use of the subset select option.



Right-clicking to close the polygon selects the following faces.



See Also

Creating A Part Curve During Surface Preparation
Selecting With A Zone
Selecting With A Specified Threshold
Selecting Surfaces And Curves With The Selection Controls
Selecting Surfaces Or Curves With The Graphics Window
Using The Select Curves Mode
Pop Up Menu Of The Surface Preparation Graphics Window

Selecting With A Zone

This tool selects entities on the displayed surface using a polygonal zone.

The filter options control which entities the tool selects.

The tool selects a face if all three vertices belonging to the face are within the zone, even if the zone does not completely encompass the face. The tool selects an edge if the two vertices belonging to the edge are within the zone. The tool selects a vertex when the vertex is within the zone.

To end the zone selection, right-click in the *Graphics* window or click the starting point.

Four selection options are available for the zone selection tool, which you can access by right-clicking the **Select zone** button.



See Also

Selection Options
Entity Selection
Using The Select Curves Mode

Selecting With A Specified Threshold

This tool selects entities on the displayed surface using the *Threshold Selection* dialog.

The filter options control which entities the tool selects.

Four selection options are available for use with the threshold tool, which you can access by right-clicking the **Select threshold** button.



Using a threshold selection option opens the *Threshold Selection* dialog, where you specify the **Threshold**, and enter **Minimum** and **Maximum** values.



See Also

Selection Options
Entity Selection

Using The Geometric Range Selection Mode

This section describes how to specify a threshold using a geometric range.

To activate the **Geometric Range** threshold:

1. Select the **Geometric Range** option from the **Select Threshold** drop-down menu.



The mode has four selection options available from the **Select** drop-down menu.



The defaults for each selection option use the global laboratory coordinate system. To change the coordinate system:

2. Use the drop-down menu that is labeled **Coordinate System**.



All properties that require specification of a point in 3-dimensional space are updated according to the chosen coordinate system. The tool accepts any available length unit, with different unit systems permitted in each field.

Min-Point, Max-Point

The **Min-Point, Max-Point** option specifies selection using minimum and maximum bounds, for each coordinate direction.



For linear values (X, Y, Z, and R), the order in which you specify the minimum and maximum bounds is not important. That is, a range from `10 mm` to `2 mm` selects the same entities as a range from `2 mm` to `10 mm`. For angular values (θ and φ) in cylindrical or spherical coordinate systems, the order of specification is important. The tool uses the right-hand rule to determine angular values in the positive rotation direction, relative to the axis of rotation. A range from `0 deg` to `90 deg` gives a different selection than a range from `90 deg` to `0 deg`. To illustrate, consider the following examples, which select faces in a cylindrical coordinate system:

- A range in θ from 0 deg (minimum) to 90 deg (maximum) results in the following selection:



- A range in θ from 90 deg (minimum) to 0 deg (maximum) results in a 270-deg selection:



The following table lists the valid range of values for angular coordinates:

System Type	Coordinate Direction	Valid Range
Cylindrical	θ	$-2\pi \le \theta \le 2\pi$
Spherical	θ	$0 \le \theta \le \pi$
	φ	$-2\pi \le \varphi \le 2\pi$

The minimum and maximum bounds have precise tolerances. If a vertex is marginally out of the geometric range, the selection does not include the associated faces and edges. You can overcome this issue in one of three ways:

3. Specify a slightly larger range, for example, -5 deg to 90 deg, instead of 0 deg to 90 deg.
   This workaround is not possible for θ in a spherical coordinate system, due to the valid range limits identified above.
4. Modify the local coordinate system in a manner that allows the range to include the required entities.
5. Use the Project Vertices tool to modify the location of vertices with respect to the chosen coordinate system.

Sphere Around Point

The **Sphere around Point** option selects entities that are included within a spherical threshold, which is defined by a point and radius.



6. Activating the **Include partially contained** option selects entities if any part of the entity is within the spherical threshold. If this option is deactivated, to tool only selects entities that are contained within the threshold.
7. The *Point* property specifies the centerpoint of the spherical threshold. The *Radius* specifies the radius, and hence extent of the spherical threshold.

   Provided below is an example of the spherical threshold selection option. The transparent sphere does not appear in the *Graphics* window; it has been added for illustrative purposes.



Box Around Point

The **Box around Point** option selects entities that are included within a box threshold, which is defined by a point and half-distance.



8. Activating the **Include partially contained** option selects entities if any part of the entity is within the box threshold. If this option is deactivated, the tool only selects entities that are contained within the threshold.

   The *Point* property specifies the centerpoint of the box threshold. The *Half-distance* specifies the extent of the threshold in each cartesian direction (x, y, and z), from the centerpoint. Therefore, a cube is always used as the threshold is defined by the same distance in all directions.

   Provided below is an example of the box threshold selection option. The transparent box does not appear in the *Graphics* window; it has been added for illustrative purposes.



Closest N Faces

The **Closest N faces** option selects the closest user-defined amount of faces to a specified point. To determine the closest faces, the code first determines the closest point on each face to the user-defined point and finds the distance between these points.

The Closest N faces are then identified as the first N faces with the shortest distance.



9. The *Point* property specifies the reference of the threshold. The *N* property specifies the number of selected faces closest to the specified point.

   Provided below is an example of the **Closest N faces** selection option. This example shows the closest 25 faces to the user-defined reference point. The sphere representing the point does not appear in the *Graphics* window; it has been added for illustrative purposes.



Using The Edge Length Selection Mode

The **Edge Length** threshold selects edges according to length, which is based on a user-defined range.

1. To activate the **Edge Length**selection mode, select the **Edge Length** option from the **Select Threshold** drop-down menu.



2. When using the **Edge Length** threshold, you specify the minimum and maximum edge lengths in the *Threshold Selection* dialog.



The tool accepts any available length unit, with different unit systems permitted in the minimum and maximum fields.

Four options are available from the **Select** drop-down menu, which alter the range definition.

- **Between Min and Max** - Selected edges have a length between the user-defined **Min Length** and **Max Length**.
- **Above Min** - Selected edges have a length above the user-defined **Min Length**, with the **Max Length** field deactivated.
- **Below Max** - Selected edges have a length below the user-defined **Max Length**, with the **Min Length** field deactivated.
- **Below Min and Above Max** - Selected edges have a length below the user-defined **Min Length** or above the **Max Length**.

Using The Dihedral Angle Selection Mode

The **Dihedral Angle** threshold selects edges according to the angle between the face normals of the faces that are connected to the edge.



1. To activate the **Dihedral Angle**selection mode, select the **Dihedral Angle** option from the **Select Threshold** drop-down menu.



2. When using the **Dihedral Angle** threshold, you specify the minimum and maximum angles in the *Threshold Selection* dialog.



The tool accepts any available angle unit, with different unit systems permitted in the minimum and maximum fields.

Four options are available from the **Select** drop-down menu, which alter the range definition.

- **Between Min and Max** - Selected edges have a dihedral angle between the user-defined **Min Angle** and **Max Angle**.
- **Above Min** - Selected edges have a dihedral angle above the user-defined **Min Angle**, with the **Max Angle** field deactivated.
- **Below Max** - Selected edges have a dihedral angle below the user-defined **Max Angle**, with the **Min Angle** field deactivated.
- **Below Min and Above Max** - Selected edges have a dihedral angle below the user-defined **Min Angle** or above the **Max Angle**.

See Also

Using The Chase Edge Selection Mode

Using The Face Area Selection Mode

The **Face Area** threshold selects faces according to area, which is based on a user-defined range.

1. To activate the **Face Area** selection mode, select the **Face Area** option from the **Select Threshold** drop-down menu.



2. When using the **Face Area** threshold, you specify the minimum and maximum face areas in the *Threshold Selection* dialog.



The tool accepts any available area unit, with different unit systems permitted in the minimum and maximum fields.

Four options are available from the **Select** drop-down menu, which alter the range definition.

- **Between Min and Max** - Selected faces have an area between the user-defined **Min Area** and **Max Area**.
- **Above Min** - Selected faces have an area above the user-defined **Min Area**, with the **Max Area** field deactivated.
- **Below Max** - Selected faces have an area below the user-defined **Max Area**, with the **Min Area** field deactivated.
- **Below Min and Above Max** - Selected faces have an area below the user-defined **Min Area** or above the **Max Area**.

Using The Face Quality Selection Mode

The **Face Quality** threshold selects faces according to face quality, which is based on a user-defined range.

1. To activate the **Face Quality** selection mode, select the **Face Quality** option from the **Select Threshold** drop-down menu.



2. When using the **Face Quality** threshold, you specify the minimum and maximum face quality in the *Threshold Selection* dialog.



Four options are available from the **Select** drop-down menu, which alter the range definition.

- **Between Min and Max** - Selected faces have a face quality between the user-defined **Min Quality** and **Max Quality**.
- **Above Min** - Selected faces have a face quality above the user-defined **Min Quality**, with the **Max Quality** field deactivated.
- **Below Max** - Selected faces have a face quality below the user-defined **Max Quality**, with the **Min Quality** field deactivated.
- **Below Min and Above Max** - Selected faces have face quality below the user-defined **Min Quality** or above the **Max Quality**.

See Also

What Is Face Quality

Using The CAD Association Selection Mode

The **CAD Association** threshold selects faces, edges, or vertices that are based upon CAD association.

1. To activate the **CAD Association** selection mode, select the **CAD Association** option from the **Select Threshold** drop-down menu.



2. Three specification options are available from the **Select** drop-down menu:

   - **Faces on CAD Face** selects all faces that have association with CAD faces. A CAD face is a surface on the part that is defined in the CAD model.
   - **Edges on CAD Edge** selects all edges that have association with CAD edges. A CAD edge is the perimeter of a CAD face. This option is useful for identifying which edges are used with the aligned meshing and geometric features when using the surface remesher.
   - **Vertices on CAD Vertex** selects all vertices that have association CAD vertices. A CAD vertex is the point where two CAD edges intersect.



Activating the **Select NOT Associated with CAD** option has the opposite effect and selects face, edges, or vertices that do not have CAD association.

As an example, consider the following surface, which has faces with CAD association highlighted in beige using the CAD association display option:



Using the **Faces on CAD Face** option and ticking the checkbox to select faces that do not have CAD association has the following effect:



The following image shows edges that have CAD association, which is used with the aligned meshing and geometric features:



See Also

Using Color Modes

Selecting Surfaces And Curves

For a surface, this tool selects the associated entities according to the faces, edges, and vertices filters. For a part curve, this tool selects the associated edges.

There are two approaches to selecting entities on surfaces:

- Using the Selection Controls
- Using the Graphics Window

Selecting Surfaces And Curves With The Selection Controls

This method of selecting entities allows you to pick surfaces from the object selection dialog.

To select surfaces and curves with the selection controls:

1. Select the appropriate entity filters which define the selection of associated faces, edges, and/or vertices.
   This action is not necessary if you are selecting part curves, as only edges are available.
2. Right click the **Select faces/edges on part surfaces/part curves** button and select one of the selection options.



3. In the *Select Part Surface/Curves* dialog, select the required items and click **Apply**.



For example, selecting this outlet boundary with only the Vertices filter activated has the following result.



See Also

Selection Options
Entity Selection

Selecting Surfaces Or Curves With The Graphics Window

This method of selecting entities allows you to pick surfaces and curves by right-clicking in the *Graphics* window.

To select surfaces or curves with the *Graphics* window:

1. Select the appropriate entity filters which define the selection of associated faces, edges, and/or vertices.
2. Right-click the required surface in the *Graphics* window and select one of the selection options.
   If no entities are currently selected, only the **Select** option is available.



Similarly, you can use the same selection options on feature curves.

Example of Selecting Entities on a Surface

The following image shows a surface for which the tool selects associated entities.



With the **Edges** and **Vertices** filters activated, the tool selects the following entities.



Example of Selecting Entities on a Curve

The following image shows a curve for which the tool selects associated entities.



With the **Vertices** filter activated, the tool selects the following entities.



See Also

Selection Options
Entity Selection

Growing And Shrinking The Selection

Use these tools to grow or shrink a selection of faces, edges, or vertices by one set of neighboring faces.

This tool grows the selection one set of neighboring faces from the pre-selected faces, edges, or vertices. The tool grows from any connected vertices in the selection. This button becomes available when you select at least one face, edge, or vertex.

This tool shrinks the pre-selected faces by one neighbor, which is based on the connected vertices of the faces. This button becomes available when you select at least one face.

Growing The Selection By Multiple Neighbors

This tool grows multiple sets of neighboring faces from the pre-selected faces, edges, or vertices.

Clicking this button opens the *Multiple grow options* dialog.



Number of Steps

This group box provides two options:

- **Infinite** grows faces until the selection includes all connected part faces.

- **Finite** grows faces by the supplied finite value, with growth based on connected edges or vertices, depending on the *Grow Options*

Grow Options

The *Grow Options* determine how the selection grows, be it edge-based or vertex-based.

- **Edge based Growth** selects faces with an edge connected to the currently selected faces. The tool does not include faces connected to the surface by a vertex.
- **Vertex based Growth** selects faces with a vertex connected to the currently selected faces. This results in a faster growth of faces per step, and allows selections on another surface that is connected by a single vertex.

Break Options

This group box provides the following options for constraining the growth extent:

- **Break at Feature Curves** breaks the growth of faces at feature curves on the part surface
- **Break at Non-manifold Edges** breaks the growth of faces at non-manifold edges on the part surface
- **Break at Sharp Angles** breaks the growth of faces at sharp angles in the surface, with a user-defined **Angle**

Selecting Attached Edges And Vertices

Use these tools to select edges or vertices that are attached to the currently displayed edges, vertices, or faces.

This tool selects the attached edges that are based on the currently displayed faces or vertices. Right-clicking this button provides options in the Select Attached Edges Options dialog.

This tool selects the attached vertices that are based on the currently displayed faces or edges. No other options are available for this tool.

Select Attached Edges Options

The *Select Attached Edges Options* dialog configures the behavior of the **Select attached edges** tool.



- **Select All Edges** is the default option, which selects all edges belonging to highlighted faces and vertices.

  Deactivating this option makes the **Select Perimeter Edges** and **Select Feature Edges** options available.
- **Select Perimeter Edges** selects the edges lying on the perimeter of a face or contiguous group of faces.
- **Select Feature Edges** selects attached edges that are marked as feature curves.

  It is possible to choose **Select Perimeter Edges** and **Select Feature Edges** together.

Inverting The Selection

This tool inverts the current selection of faces, edges, and vertices that are based on the active model (not the displayed model).

The filter options are taken into account when the inversion is performed. Suppose, for example, that you have selected a number of faces and edges, but only want to invert the selection of faces. You can deactivate the **Edges** filter checkbox, keep the **Faces** filter checkbox

activated, and click the **Invert Selection** button. All faces in the active model that were not previously selected are selected, and the faces that were previously selected are deselected.

This button is active whenever at least one face, edge, or vertex is selected. If nothing is selected before the inversion, the inverted set is empty.

As an example of this operation, consider the following selection that contains faces, edges, and vertices.



If all three filters for **Faces**, **Edges** and **Vertices** are ticked, and the **Invert Selection** button is clicked, the corresponding inverted selection includes faces, edges, and vertices.



If, however, only the **Faces** filter had been ticked, and the filters for **Edges** and **Vertices** had been deactivated, the inverted selection would only contain faces. The rest of the current selection would remain as before.



If all three filters are deactivated, no inversion is performed.

Using The Chase Edge Selection Mode

This tool selects edges manually using the shortest distance between two vertices.

This tool becomes available when you select the **Edges** filter.

Clicking **Start Chase Edge Selection** switches the *Selection Control* group box to *Chase Edge Selection* mode. Entering *Chase Edge Selection* mode maintains all previous selections while using the *Chase Edge Selection* tools.



The first vertex that is selected, indicated using a red ball, becomes the starting point for the selection process. Thereafter, selecting vertices on the surface selects the edges along the shortest path between the previous and current vertex. The red ball relocates to indicate the last vertex selected.

- **Follow Sharp Angle** selects a path that is defined by edges whose attached faces have a dihedral angle, in degrees, that is greater than the user-defined **Angle**.
- **Continue from Last** continues selecting edges from the last selected vertex. Deactivating **Continue from Last** only allows selections between the first and second selected vertices (that is, a non-continuous selection).
- The **Selection Mode** menu provides the following options:
  - **Add** includes the next set of edges to the current selection
  - **Remove** excludes the next set of edges from the current selection
- **Clear selected edges** deselects the edges that are selected within *Chase Edge Selection* mode. If you selected any edges before entering *Chase Edge Selection* mode, this action does not deselect them.
- The **Undo** and **Redo** buttons cycle backward and forward through the steps that are made in the selection process.
- **Close** completes the operation and returns *Selection Control* to its default mode. Leaving *Chase Edge Selection* mode maintains selections so you can perform further operations on the edges.

See Also

Using The Dihedral Angle Selection Mode
Entity Selection

Using The Select Curves Mode

This tool lets you quickly select the edges that have been marked as feature curves by Mark Features by Angle, etc or related tools.

This tool becomes available when you select the **Edges** filter.

Typical uses for the tool include unflagging feature edges and assigning feature edges to a new or existing part curve.

1. Clicking **Select curves** switches the *Selection Control* group box to *Edit Feature Curves* mode. This action deselects all previous face and vertex selections, and only maintains existing edge selections if they are feature curves.



Starting *Edit Feature Curves* mode also makes the part surface translucent and highlights feature curves yellow.



The tool has two selection modes:

- **Select Single Edge** - Selects only one feature edge at a time.



- **Select Connected Curves** - Selects all feature edges that are attached to the selected edge.



2. Activating the **Select Connected Curves** option allows the specification of two further options. These options determine where the feature edge selection stops:
   - **Break at Edge Junctions** - Stop the selection where three or more edges meet
   - **Break at Sharp Angles** - Stop the selection at a user-defined angle
3. Clicking multiple feature edges automatically adds to the current selection (that is, you do not need to press `<Ctrl>` for multiple selections).

The zone selection tool is also available, which selects all feature edges within a polygonal zone, which is defined by the cursor in the *Graphics* window. The standard selection options, select, add select, deselect, and subset select are accessed by right-clicking the button.

To deselect a selection, there are two available options:

4. Click a selected feature edge in the *Graphics* window to deselect the edge or connected edges. The settings in the *Edit Feature Curves* tool determine the edge or the connected edges.
5. Click **Select none** to deselect all selected edges.
6. Clicking **Close** exits *Edit Feature Curves* mode and maintains the feature edge selection.

See Also

Selection Options
Selecting With A Zone
Assigning Feature Edges To A Target Part Curve
Unflagging Edges As Features
Marking Feature Edges
Entity Selection

Using The Leak Path Mode

This tool selects faces using the shortest path between one source and multiple target faces.

This tool becomes available when you select the **Faces** filter.

1. Clicking **Find Leak Path** switches the *Selection Control* group box to *Find Leak Path Tool* mode.



2. To use the *Find Leak Path Tool*, first select a single source face from the displayed surfaces in the *Graphics* window, and click (**Add**) within the *Source Face* group box.
3. Then, select one or more target faces from the displayed surfaces in the *Graphics* window and click (**Add**) within the *Target Face(s)* group box.

4. You can remove faces from the *Source* and *Target* by selecting them in the *Graphics* window and clicking the appropriate **Remove selected faces from the merge set** button. Similarly, you can deselect all faces in the *Source* and *Target* by clicking the appropriate **Clear all faces from the merge set** button.
5. When you have defined the *Source* and *Target* faces, click **Find Leak Path** to perform leak path detection.

   Upon completion, the tool highlights all faces that lie on the shortest paths between the *Source Face* and the *Target Face(s)* in the *Graphics* window.

   In the example below, the green face is the source face and the two red faces are targets.



Clicking **Find Leak Path** selects and highlights all faces on the connected paths, along with the original source and target faces.



See Also

Changing The Parts In The Active Model

Display Controls

The *Display Controls* allow you to manipulate the displayed faces, in order to make selections.

At any time during surface repair, you can show or hide all faces at once. You can also show or hide a subset of faces based on the current selection or surfaces, or you can grow and shrink the displayed face by one neighbor. In addition, Color Faces by menu colors the selected faces by patch, surface, or part.



The following section provides a description of each display button, showing and hiding surfaces in the *Graphics* window and the available color modes.

See Also

Storing Sets

Using The Display Buttons

The following section details the use of each display control button.

Showing and Hiding All Faces

This tool shows all part faces. Any currently selected faces remain selected.

This tool hides all part faces. Any currently selected faces remain selected.

Showing and Hiding Selected Faces

This tool displays the faces currently selected in the *Graphics* window.

This tool hides the faces currently selected in the *Graphics* window. You can redisplay the hidden faces by growing from neighboring displayed faces, or by showing all faces.

Showing and Hiding Faces on Surfaces

This tool shows faces on a part surface using the *Show Part Surfaces* dialog.

This tool hides faces on a part surface using the *Hide Part Surfaces* dialog.

Clicking one of the above buttons opens a hierarchical object selection dialog. The following example shows the *Show Part Surfaces* dialog.



Growing and Shrinking Displayed Faces

This tool grows by one set of neighbor faces from the currently displayed faces.

This tool shrinks by one set of neighbor from the currently displayed faces, which are based on connected vertices.

Hiding And Showing Surfaces In The Graphics Window

In addition to using the dialog to hide and show faces on surfaces, you can right-click a part surface in the *Graphics* window and use the available options.

The right-click menu allows you to hide or show parts, part surfaces, and boundaries for the entity beneath the pointer. In order to make the boundary option available, the part surface must be assigned to a boundary.

The following options are available from the right-click menu:

- Hide
- Show
- Show Only

Hide

The **Hide** option removes all faces associated with the part, part surface, or boundary from the current display.



The effect of selecting **Hide** for the above **Outlet** part surface is shown below.



Show

The **Show** option adds all faces associated with the part, part surface, or boundary to the current display, including faces that are hidden by other display controls.



The effect of selecting **Show**for the above **Wall** part surface is shown below.



Any previously displayed faces remain in the display. This behavior is different to the show only option.

Show Only

The Show Only option shows only the selected part, part surface, or boundary in the current display, with all other faces removed.



The effect of selecting **Show Only** for the above **Outlet** part surface is shown below.



If a part or part surface is partially displayed (that is, some of its faces are hidden), clicking **Show Only** shows the currently displayed faces that belong to the part or part surface. This option does not show the hidden faces. To show the all hidden faces, you can click the **Show** option.

See Also

Assigning Part Surfaces to Boundaries
Pop Up Menu Of The Surface Preparation Graphics Window

Using Color Modes

Using the **Color Faces by** menu colors visible faces in the *Graphics* window.

Selecting an option from this menu over-rides any other display color options.



Twenty eight colors are available for use in the above modes. If there are more than twenty eight patches, parts, or surfaces, the colors are repeated. Selecting faces colors faces with the highlight color, irrespective of the **Color Faces by** option.

Note: When usingSTAR-CCM+ v9.06 or later with simulation files saved in STAR-CCM+ v9.04 or earlier, some non-default colors from the legacy simulation could appear differently.

See Also

Using The CAD Association Selection Mode

What Is The Default Color Mode

Selecting **Default Color** colors faces in the display with the default color, apart from erroneous faces that are colored by the surface diagnostics.

An example of the default color mode with surface diagnostics is shown below.



See Also

Using Threshold Diagnostics

What Is The Patch Color Mode

Selecting **Patch** colors faces according to their underlying patch definition. Faces that share a patch definition use the same color.

An example of the patch color mode is shown below.



What Is The CAD Association Color Mode

Selecting **CAD Association** colors faces according to whether they have CAD association or not.

Beige faces represent surfaces that have underlying CAD available for CAD projection during the meshing process. Dark gray surfaces represent surfaces with no CAD association. An example of the CAD association color mode is shown below.



See Also

Surface Mesher Controls and Values

What Is The Part Surface Color Mode

Selecting **Part Surface** colors faces according to the surface to which they belong.

An example of the part surface color mode is shown below.



What Is The Part Boundary Color Mode

Selecting **Part Boundary** colors faces according to the boundary to which they are assigned.

An example of the part boundary color mode is shown below.



What Is The Part Color Mode

Selecting **Part** colors faces according to the part to which they belong.

An example of the part color mode is shown below.



What Is The Face Normal Color Mode

Selecting **Face Normal** colors faces according to their topological orientation, that is, the direction of the face normals.

Green faces have a positive face normal and red faces have a negative face normal.

For a closed-body, the positive face normals (green) point outwards, and negative face normals (red) point inwards. If you import a surface representing the solid component and manually extract the wetted surface using surface preparation tools, the orientation is reversed as the face normals were pointing outwards from the solid component upon import.

Also, when in **Merge/Imprint** mode, this option over-rides the default colors. Therefore, the green and red faces do not identify the source and target faces, but identify the positive and negative face normals instead.

An example of the face normal color mode is shown below.


Geometry Scene 1
X
Z

Surface, Contact, and Curve Edit Controls

The *Part Surface, Part Surface Contact and Part Curve Edit* group box assigns faces to an existing or new part surface, or an existing part surface contact, and feature edges to an existing or new part curve.

Note: This tool is only available when using Parts-Based Meshing. There is no equivalent option when repairing regions.



The surface and feature diagnostics update following any action that is performed with the *Part Surface, Part Surface Contact and Part Curve Edit* controls. Review the updated diagnostics to identify and fix any errors that the assign action causes, before closing the *Surface Preparation tool*.

You can reverse the assigning of faces or curves using the Undo button. However, this option does not apply to the creation of part surfaces or curves, which you must manually delete from the simulation tree. You can move faces into and out of part contacts, however, you cannot create new part contacts within this tool. To create a part contact, return to the simulation tree. See Creating Part Contacts Manually.

Note: If you break a part surface contact by moving one of the contacting faces to another surface, both contacting faces move together. This creates free edges on the remaining part.

The *Part Surface, Part Surface Contact and Part Curve Edit* controls provide the following functionality:

See Also

Changing The Parts In The Active Model
Using Part Contact Data to Create Interfaces
Feature Diagnostics
Using Threshold Diagnostics

Creating A Part Surface During Surface Preparation

Follow this procedure to create a part surface without leaving the *Surface Preparation tool*.

1. In the *Part Surface, Part Surface Contact and Part Curve Edit* group box, select **Target Part Surface** and click **New**.
2. In the *Create New Part Surface* dialog, select the part to which the new surface belongs, and enter the name for the new surface.
3. Click **OK** to create and select the new surface as the target.

Assigning Faces to a Target Part Surface

All faces in the part belong to a part surface by default. This tool assigns faces to a different surface.

1. To assign faces to a different surface, in the *Part Surface, Part Surface Contact and Part Curve Edit* group box, select **Part Surface**.
2. Then, click the ellipsis (**...**) and select the target part surface from the **Select Object** dialog.
   Alternatively, right-click the target part surface in the *Graphics* window and select **Part Surfaces: > [Part Surface] > Set in Edit Panel**.

   Both methods have the same result, which is to select the surface in the *Part Surface, Part Surface Contact and Part Curve Edit* group box.

Note: Right-clicking on a baffle contact surface displays options differently from a non-baffle contact because a baffle contact is applied to a single part surface belonging to a single part.

3. Select the required faces in the *Graphics* window, and click **Modify** to assign the selected faces to the target part surface.

See Also

Pop Up Menu Of The Surface Preparation Graphics Window

Assigning Faces to a Target Part Surface Contact

This section explains how to modify the assignment of faces to part surface contacts in surface repair mode.

Within the *Part Surface, Part Surface Contact and Part Curve Edit* group box, you can use the **Part Surface Contact** option to assign faces to existing part surface contacts. This feature is useful in situations where you use surface repair tools for creating faces that belong to a contact.

1. In the *Part Surface, Part Surface Contact and Part Curve Edit* group box, select the target **Part Surface Contact**.

   Alternatively, in the *Graphics* window, right-click the part surface contact and select **Part Surface Contact: > [Contact/Contact] > Set in Edit Panel**.

Note: Only contacts that exist between the parts that are in the current surface repair mode are available. You cannot create a new contact using this tool. To create a new contact, return to the simulation tree. See Creating Part Contacts Manually.

2. In the *Graphics* window, select the faces which you want to add to the target part surface contact.
3. In the *Part Surface, Part Surface Contact and Part Curve Edit* group box, click **Modify**.

   The selected faces are added to the target part surface contact. However:

   - If the selected faces belong to a contact that is different from the target contact, the following warning is displayed in the *Output* window:

     ```
     Warning -- some faces were moved out of part
     contacts.
     ```

   - The faces that are added to the target part surface contact using this tool are assigned to a separate patch from those faces that already belong to the target part surface contact.

Creating A Part Curve During Surface Preparation

Follow this procedure to create a part curve without leaving the *Surface Preparation tool*.

1. To create a part curve without leaving the *Surface Preparation tool*, in the *Part Surface, Part Surface Contact and Part Curve Edit* group box, click **New**.
2. In the *Create New Part Curve* dialog, select the part to which the new curve belongs, and enter the name for the new curve.
3. Click **OK** to create and select the new part curve as the target.

See Also

Selection Options

Assigning Feature Edges To A Target Part Curve

Use this tool to assign feature curves to a different curve.

All feature curves that are created in the *Surface Preparation tool* belong to the **Surface Repair Edges** curve by default. This tool assigns feature curves to a different curve. You can only assign edges to a target part curve if you have marked them as feature edges. To mark feature edges, use the Flag Edges as Features tool.

1. To assign feature edges to a part curve, in the *Part Surface, Part Surface Contact and Part Curve Edit* group box, select **Part Curve**.
2. Then, click the ellipsis (**...**) and select the target part curve from the **Select Object** dialog.
   Alternatively, right-click the target part curve in the *Graphics* window and select **Part Curve: > [Part Curve] > Set in Edit Panel**.

   Both methods have the same result, which is to select the curve in the *Part Surface, Part Surface Contact and Part Curve Edit* group box.
3. Select the required feature edges in the *Graphics* window, and click **Modify** to assign the selected feature edges into the target part curve.

See Also

Using The Select Curves Mode
Flagging Edges As Features

Set Controls

A set is a collection of faces, edges, or vertices, which you can store, recall and browse, using the *Set Controls* group box.

This tool is useful when working with various sections of a geometry surface, as it allows you to retrieve a set of entities easily without the need for reselection.



Closing the *Surface Preparation tool* discards all currently stored sets (that is, sets are only available in the current session).

Creating a New Set

1. To create a set, click **New**, which opens the *Create new set* dialog.



2. Enter a unique name for the set and click **OK**.

   The new set becomes available for selection in the **Set**drop-down menu.



Deleting a Set

3. To delete a set, first select the set from the **Set** drop-down menu.



4. Click **Delete** to delete the set.
   This action only deletes the set definition, which does not remove the entities from the surface. Selecting the **Default** set makes the **Delete** button unavailable, as you cannot delete the Default set.

See Also

Changing The Parts In The Active Model
Entity Selection
Storing Sets

Storing Sets

Storing allows you to save a group of faces, edges, and vertices to a previously created or default set, which you can recall during the *Surface Preparation* session.

1. To store a set, select **Store** from within the *Action* group box.



In **Store** mode, the tool provides two options:

- **Selected** stores the entities that are selected with the Selection Controls
- **Displayed** stores the entities that are displayed with the Display Controls

2. Click **Store** to store the selected or displayed entities in the current set.
   Store has four options, which are available by right-clicking the button.



The button icon differs slightly depending on whether you have chosen the **Selected** or **Displayed** option.

Vertices are never displayed in the *Graphics* window, hence the **Store Displayed** option is unable to store vertices in a set.

3. To store vertices, select the required vertices and use the **Store Selected** option instead.
   Both the **Store Displayed** and **Store Selected** options can store faces and edges.

   Store

   This option stores the selected or displayed entities in the current set, replacing all existing entities in the set.

   Add Store

   This option adds the selected or displayed entities to the currently stored entities in the set.

   Remove from Set

   This option removes the selected or displayed entities from the currently stored entities in the set.

   Subset Store

   This option stores only those selected or displayed entities that already exist in the current set. Entities in the current set, that are not selected or displayed, are discarded.

See Also

Display Controls
Entity Selection
Recalling Sets

Set Controls
Recalling Sets

Recalling Sets

Recalling allows you to retrieve the collection of faces, edges, and vertices from a previously stored set.

1. To recall a set, select **Recall** from the *Action* group box and select the **Set** from the drop-down menu.



In **Recall** mode, the tool provides two options:

- **Selection** recalls the entities to the current selection
- **Display** recalls the entities to the current display

In addition, the **Faces**, **Edges, and Vertices** filters are available, which limit the type of entities recalled by the tool. The default options activate all filters, which recall all entities. Deactivating a filter omits that particular entity from the recall operation.

2. Click **Recall** to recall the selected set to the current selection or display, according to the activated filters.
   Recall has four options, which are available by right-clicking the button.



The button icon differs slightly, depending on whether you have chosen the **Selection** or **Display** option.

Recall

This option recalls the set, replacing all existing entities in the current selection or display.

Add Recall

This option adds entities in the set to the current selection or display.

Remove Recall

This option removes entities in the set from the current selection or display.

Subset Recall

This option recalls entities that are in both the set and current selection or display.

See Also

Storing Sets
Storing Sets

Browsing Sets

Browsing sets allows you to cycle and display groups of continuous entities within the currently selected set.

For example, in the diagram below, all the selected faces are in one set, but there are four groups of continuous faces.



1. To browse, first select a **Set** from the drop-down menu.



2. Then, select **Browse** from the *Action* group box.



The default options when browsing sets display the group of continuous entities, and two neighboring faces, in order to identify the surrounding area. Using the browse buttons also selects the group of continuous entities in the *Graphics* window.

While browsing, two options control the viewing of sets:

- **Reset View** resets the view to zoom in on the selected entities.

  If you deactivate the **Reset View** option, the view remains the same.

- **Reset Displayed** resets the displayed faces, edges, and vertices to the selected entities, and two neighbors.

  If you deactivate the **Reset Displayed** option, the displayed entities remain the same.

3. Use the browse buttons to cycle through and select continuous groups of entities within the set.

Undo And Redo

The **Undo** and **Redo** buttons at the bottom of the *Surface Preparation* tool allow operations that are made during surface preparation to be undone or redone.

For example, if you deleted a face by mistake, clicking **Undo** would undelete the face, and return the part to its previous state. A complete history of all operations in the current surface preparation session is available. The history is deleted when you close the dialog. You can step backwards or forwards as necessary in order to return to the desired state.

This tool undoes the previous surface preparation operation. Holding the cursor over the button indicates, in the tooltip, the operation that the tool undoes.

This tool redoes the previously undone operation. Holding the cursor over the button indicates, in the tooltip, the operation that the tool redoes.

See Also

Changing The Parts In The Active Model

Query Diagnostic Responses From Java Macros

This section provides example Java and reference material for accessing information from surface preparation query diagnostics.

A query diagnostic response is data that comes from the server when you run a query diagnostic. Using java macros and tags, you can access individual parts of a diagnostic response to produce more specific information.

This section has the following subsections:

Java Examples

The following examples show several uses of responses to obtain specific information from query diagnostics.

As prerequisites, this section contains the required import paths and basic java code.

These examples are summarized below:

- Integer — Accessing a vertex ID
- Integer Vector —Accessing the number of faces that belong to a boundary
- Double — Accessing the total length of selected edges
- Double Vector — Accessing the X, Y, and Z coordinates of a vertex
- Client-Server Objects — Access the query units, query coordinate system, and the part surface or boundaries to which a face belongs.

Import Paths

To access and use query responses, add the following import parts to your macro:

```
import star.base.neo.*;
import star.common.*;
import star.meshing.*;
```

Prerequisites

To use these examples, add the following java to your macro:

```
Simulation simulation_0 = getActiveSimulation();
PartSurfaceMeshWidget partSurfaceMeshWidget_0 =
  simulation_0.get(PartSurfaceMeshWidget.class);
SurfaceMeshWidgetQueryController
surfaceMeshWidgetQueryController_0 =
  partSurfaceMeshWidget_0.getControllers()
  .getController(
  SurfaceMeshWidgetQueryController.class
  );
```

Int Example

The following example illustrates the use of integers to access the ID of a selected vertex:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0.queryVertexXYZ();
int vertId = response.getInt("VertexId");
simulation_0.println(
  "Vertex ID: " + vertId
  );
```

IntVector Example

The following example illustrates the use of integer vectors to access the number of selected faces that belong to a boundary:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0
  .queryFaceRegionBoundaryList();
IntVector faceCount =
  response.getIntVector("SelectedBoundaryFaceCount");
simulation_0.println(
  "Face count: " + faceCount.getComponent(0)
  );
```

Double Example

The following example illustrates the use of doubles to access the total length of all selected edges:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0.queryEdgeLength();
simulation_0.println(
```

```
  "Total length: " + response.getDouble("TotalLength")
  );
```

DoubleVector Example

The following example illustrates the use of double vectors to access the X, Y, and Z coordinates of the selected vertex:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0.queryVertexXYZ();
DoubleVector labXYZ =
response.getDoubleVector("LabXYZ");
simulation_0.println(
  "Vertex Coordinates: " + labXYZ
  );
simulation_0.println(
  "Vertex X-Coordinate: " + labXYZ.toDoubleArray()[0]
  );
```

ClientServerObject Examples

The following examples illustrate the use of client-server objects.

Units

This example accesses the preferred units for the query:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0.queryVertexXYZ();
Units units =
  response.<Units>getObject(
  "LengthUnits",simulation_0.getObjectRegistry()
  );
simulation_0.println(
  "Units: " + units.getPresentationName()
  );
```

Coordinate System

This example accesses the local coordinate system for the query:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0.queryVertexXYZ();
CoordinateSystem csys =
  response.<CoordinateSystem>getObject(

"LocalCoordinateSystem",simulation_0.getObjectRegistry()
  );
simulation_0.println(
```

```
  "Local CSys: " + csys.getPresentationName()
  );
```

Part Surfaces or Boundaries

This example accesses the first part surface to which the selected faces belong:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0.
  queryFacePartSurfaceList();
Vector<PartSurface> partSurfaces =
  response.<PartSurface>getObjectVector(

"SelectedPartSurfaces",simulation_0.getObjectRegistry()
  );
simulation_0.println(
  "First part surface: " + partSurfaces.get(0)
  );
```

This example accesses the first boundary to which the selected faces belong:

```
NeoProperty response =
  surfaceMeshWidgetQueryController_0
  .queryFaceRegionBoundaryList();
Vector<Boundary> boundaries =
  response.<Boundary>getObjectVector(
  "SelectedBoundarySurfaces",
  simulation_0.getObjectRegistry()
  );
simulation_0.println(
  "First boundary: " + boundaries.get(0)
  );
```

Query Diagnostic Tag Names And Types Reference

This section contains reference material on the available tag names and corresponding type for each query diagnostic.

The queries are summarized below:

- **Vertex Queries**
  - Vertex Location
  - Vertex Distance
  - Vertex Range

- **Edge Queries**
  - Edge Length
  - Edge Dihedral Angle
  - Edge Geometric Range
- **Face Queries**
  - Face Area
  - Face Quality
  - Face Centroid
  - Face Geometric Range
  - Part Surface List
  - Region Boundary List
- **Mixed Entity Queries**
  - Vertex to Edge Distance
  - Vertex to Face Distance
  - Edge to Face Distance
  - Mixed Entity Geometric Range

Vertex Queries

This section contains reference material on the available tag names and corresponding type for all vertex-based query diagnostics.

Note: The same limitations and conditions apply to using vertex-based queries in macros as with using the GUI. See What Are Vertex Query Diagnostics? for more information.

Vertex Location

The following response properties are available for single vertex selections.

Tag Name	Type	Description
`VertexId`	`Int`	Vertex ID of the selected vertex
`LabXYZ`	`DoubleVector`	Coordinates of the selected vertex relative to the laboratory coordinate system

Tag Name	Type	Description
LocalXYZ	DoubleVector	Coordinates of the selected vertex relative to the local coordinate system
CoordinateSystem	ClientServerObject	Coordinate system for the query
LengthUnits	ClientServerObject	Preferred length units for the query
AngleUnits	ClientServerObject	Preferred angle units for the query

The following response properties are available for single and multiple vertex selections.

Tag Name	Type	Description
AllVertexIds	Int	Vertex ID's of the selected vertices
AllLabXYZCoords	DoubleVector	List of the coordinates of the selected vertices relative to the laboratory coordinate system
AllLocalXYZCoords	DoubleVector	List of the coordinates of the selected vertices relative to the local coordinate system (only available if you set a local coordinate system)

Vertex Distance

Tag Name	Type	Description
StartVertexId	Int	Vertex ID of the starting vertex
EndVertexId	Int	Vertex ID of the ending vertex
VertexDistance	Double	Distance between vertices

Tag Name	Type	Description
StartVertexLabCoordinates	DoubleVector	Coordinates of the starting vertex relative to the laboratory coordinate system
EndVertexLabCoordinates	DoubleVector	Coordinates of the ending vertex relative to the laboratory coordinate system
XYZComponents	DoubleVector	X, Y, and Z components of distance from minimum to maximum
LengthUnits	ClientServerObject	Preferred length units for the query
StartVertexLocalCoordinates	DoubleVector	Starting vertex local coordinates
EndVertexLocalCoordinates	DoubleVector	Ending vertex local coordinates
XYZLocalComponents	DoubleVector	XYZ local coordinate system components from minimum to maximum in preferred units
AngleUnits	ClientServerObject	Preferred angle units specified for query
LocalCoordinateSystem	ClientServerObject	Coordinate system specified for query

Vertex Range

Tag Name	Type	Description
LabMinRange	DoubleVector	Minimum range relative to the laboratory coordinate system
LabMaxRange	DoubleVector	Maximum range relative to the laboratory coordinate system

Tag Name	Type	Description
XYZComponents	DoubleVector	X, Y, and Z components of range from minimum to maximum
LengthUnits	ClientServerObject	Preferred length units for the query

See Also

What Are Vertex Query Diagnostics

Edge Queries

This section contains reference material on the available tag names and corresponding type for all edge-based query diagnostics.

Note: The same limitations and conditions apply to using edge-based queries in macros as with using the GUI. See What Are Edge Query Diagnostics? for more information.

Edge Length

Tag Name	Type	Description
EdgeId	Int	Edge ID of the selected edge
TotalLength	Double	Total length of all selected edges
MinLength	Double	Minimum length of all selected edges
MaxLength	Double	Maximum length of all selected edges
AveLength	Double	Average length of all selected edges
LengthUnits	ClientServerObject	Preferred length units for the query

Edge Dihedral Angle

Tag Name	Type	Description
EdgeId	Int	Edge ID of the selected edge
DihedralAngle	Double	Dihedral angle of the selected edge
ValidEdges	Int	Number of edges that are valid for the query
MinAngle	Double	Minimum dihedral angle of all selected edges
MaxAngle	Double	Maximum dihedral angle of all selected edges
AveAngle	Double	Average dihedral angle of all selected edges
AngleUnits	ClientServerObject	Preferred angle units for the query

Edge Geometric Range

Tag Name	Type	Description
LabMinRange	DoubleVector	Minimum range relative to the laboratory coordinate system
LabMaxRange	DoubleVector	Maximum range relative to the laboratory coordinate system
XYZComponents	DoubleVector	X, Y, and Z components of range from minimum to maximum
LengthUnits	ClientServerObject	Preferred length units for the query

See Also

What Are Edge Query Diagnostics

Face Queries

This section contains reference material on the available tag names and corresponding type for all face-based query diagnostics.

Note: The same limitations and conditions apply to using face-based queries in macros as with using the GUI. See What Are Face Query Diagnostics? for more information.

Face Area

Tag Name	Type	Description
`FaceId`	`Int`	Face ID of the selected face
`TotalFaceArea`	`Double`	Total area of all selected faces
`MinArea`	`Double`	Minimum area of all selected faces
`MaxArea`	`Double`	Maximum area of all selected faces
`AveArea`	`Double`	Average area of all selected faces
`AreaUnits`	`ClientServerObject`	Preferred area units for the query

Face Centroid

Tag Name	Type	Description
`FaceId`	`Int`	Face ID of the selected face
`FaceCentroid`	`DoubleVector`	Face centroid coordinates
`LengthUnits`	`ClientServerObject`	Preferred length units for the query

Face Normal

Tag Name	Type	Description
`FaceId`	`Int`	Face ID of the selected face
`FaceNormal`	`DoubleVector`	Face normal vector in preferred units
`NormalizedFaceNormal`	`DoubleVector`	Normalized face normal vector
`NormalUnits`	`ClientServerObject`	Preferred length units specified for query

Face Quality

Tag Name	Type	Description
`FaceId`	`Int`	Face ID of the selected face
`FaceQuality`	`Double`	Face quality of the selected face
`MinQuality`	`Double`	Minimum face quality of all selected faces
`MaxQuality`	`Double`	Maximum face quality of all selected faces
`AveQuality`	`Double`	Average face quality of all selected faces

Face Geometric Range

Tag Name	Type	Description
`LabMinRange`	`DoubleVector`	Minimum range relative to the laboratory coordinate system
`LabMaxRange`	`DoubleVector`	Maximum range relative to the laboratory coordinate system
`XYZComponents`	`DoubleVector`	X, Y, and Z components of range from minimum to maximum

Tag Name	Type	Description
LengthUnits	ClientServerObject	Preferred length units for the query

Part Surface List

Tag Name	Type	Description
SelectedPartSurfaces	ClientServerObject	Part surfaces to which the selected faces belong
SelectedPartSurfacesFaceCount	IntVector	Number of selected faces that belong to a part surface
SelectedContactSurfaces	ClientServerObject	Contact part surfaces to which the selected faces belong
SelectedContactFaceCount	IntVector	Number of selected faces that belong to a contact part surface

Region Boundary List

Tag Name	Type	Description
SelectedBoundarySurfaces	ClientServerObject	Boundaries to which the selected faces belong
SelectedBoundaryFaceCount	IntVector	Number of selected faces that belong to the boundaries
SelectedInterfaceSurfaces	ClientServerObject	Interfaces to which the selected faces belong
SelectedInterfaceFaceCount	IntVector	Number of selected faces that belong to an interface

See Also

What Are Face Query Diagnostics

Mixed Entity Queries

This section contains reference material on the available tag names and corresponding type for all mixed-entity query diagnostics.

Note: The same limitations and conditions apply to using mixed-entity queries in macros as with using the GUI. See What Are Mixed Query Diagnostics? for more information.

Vertex to Edge Distance

Tag Name	Type	Description
`SourceVertexId`	`Int`	Vertex ID of the selected vertex
`ClosestEdgeId`	`Int`	Edge ID of the closest edge
`ShortestDistance`	`Double`	Shortest distance from the selected vertex to the closest edge
`SourceVertexLabCoordinates`	`DoubleVector`	Coordinates of the selected vertex relative to the laboratory coordinate system
`XYZComponents`	`DoubleVector`	X, Y, and Z components of distance from minimum to maximum
`LengthUnits`	`ClientServerObject`	Preferred length units for the query

Vertex to Face Distance

Tag Name	Type	Description
`SourceVertexId`	`Int`	Vertex ID of the selected vertex
`ClosestFaceId`	`Int`	Face ID of the closest face
`ShortestDistance`	`Double`	Shortest distance from the selected vertex to the closest face

Tag Name	Type	Description
SourceVertexLabCoordinates	DoubleVector	Coordinates of the selected vertex relative to the laboratory coordinate system
XYZComponents	DoubleVector	X, Y, and Z components of distance from minimum to maximum
LengthUnits	ClientServerObject	Preferred length units for the query

Edge to Face Distance

Tag Name	Type	Description
SourceEdgeId	Int	Edge ID of the selected edge
ClosestFaceId	Int	Face ID of the closest face
ShortestDistance	Double	Shortest distance from the selected edge to the closest face
SourceVertexLabCoordinates	DoubleVector	Coordinates of the selected edge relative to the laboratory coordinate system
XYZComponents	DoubleVector	X, Y, and Z components of distance from minimum to maximum
LengthUnits	ClientServerObject	Preferred length units for the query

Mixed Entity Geometric Range

Tag Name	Type	Description
LabMinRange	DoubleVector	Minimum range relative to the

Tag Name	Type	Description
		laboratory coordinate system
`LabMaxRange`	`DoubleVector`	Maximum range relative to the laboratory coordinate system
`LabMaxRange`	`DoubleVector`	X, Y, and Z components of range from minimum to maximum
`LengthUnits`	`ClientServerObject`	Preferred length units for the query

See Also

What Are Mixed Query Diagnostics

Keyboard Shortcuts

The following shortcuts execute the surface preparation operations using the keyboard, instead of the GUI.

The mouse cursor must be in the *Graphics* window for the shortcuts to be available. If the cursor is in any other part of the interface, the keyboard shortcuts can have a different effect or no effect at all.

Currently, the Surface Repair keyboard shortcuts are available in any surface preparation mode, even if the button is not visible in the *Surface Repair* in-place dialog. For example, if you are in **Merge/Imprint** mode, you can still use `F` to flag edges as feature curves, even though the button is not available. This shortcut saves time by removing the need to change between modes.

Surface Diagnostics

Action	Keyboard Shortcut
Browse First Problem Area	`UP`
Browse Previous Problem Area	`LEFT`
Browse Next Problem Area	`RIGHT`
Browse Last Problem Area	`DOWN`
Run Selected Query Diagnostic	`Q`

Surface Repair

Action	Keyboard Shortcut
Delete Face	`D`
Create Face from Vertices	`V`
Collapse Vertices	`C`
Split Edge	`S`
Swap Edge	`W`
Smooth Vertices	`M`
Project Vertices to a Plane	`L`
Merge Vertices/Faces	`J`
Fill Holes	`H`
Zip Edges	`Z`
Remesh Faces	`R`

Action	Keyboard Shortcut
Fill Polygonal Patch	P
Flag Feature Edge	F
Unflag Feature Edge	U
Transform Vertices	K
Auto Repair Surface Errors	A
Undo	CTRL+Z
Redo	CTRL+Y

Feature Repair

Action	Keyboard Shortcut
Flag Feature Edge	F
Unflag Feature Edge	U
Fix All Feature Errors	E
Mark as Not a Feature Error	N
Mark Features by Angle	G

Offset Faces/Edges

Action	Keyboard Shortcut
Offset Faces/Edges	ALT+O

Edit Connected Surfaces

Action	Keyboard Shortcut
Mark Highlighted Surfaces in Table	Space
Mark All Surfaces in Table	T
Delete Marked Surfaces in Table	ALT+D
Highlight Surface of Selected Faces	ALT+H

Selection Control: Default Mode

Action	Keyboard Shortcut
Select None	SHIFT+N
Select/Deselect/Subset Zone	SHIFT+Z

Action	Keyboard Shortcut
Select/Deselect/Subset Threshold Range	SHIFT+R
Select/Deselect/Subset Surfaces	SHIFT+B
Grow selected	+
Shrink selected	-
Multi-grow selected	SHIFT+M
Select Attached Edges	SHIFT+E
Select Attached Vertices	SHIFT+V
Invert Selection	SHIFT+I
Start Chase Edges Mode	SHIFT+C
Start Select Curves Mode	SHIFT+F
Start Find Leak Path Mode	SHIFT+L
Toggle Face Selection	ALT+1
Toggle Edge Selection	ALT+2
Toggle Vertex Selection	ALT+3
Clear Face Selection	SHIFT+1
Clear Edge Selection	SHIFT+2
Clear Vertex Selection	SHIFT+3

Selection Control: Chase Edges Mode

Shortcuts when using **Chase Edges** mode over-ride the shortcuts in other modes.

Action	Keyboard Shortcut
Toggle Follow Sharp Angle	F
Toggle Continue From Last	C
Change Selection Mode to Add	A
Change Selection Mode to Remove	R
Clear Selected Edges	Esc
Undo	LEFT or BACKSPACE
Redo	RIGHT
Close and return to Default mode	L

Selection Control: Select Curves Mode

Shortcuts when using **Select Curves** mode over-ride the shortcuts in other modes.

Action	Keyboard Shortcut
Toggle Single Edge or Connected Curve	`S`
Toggle Break at Edge Junctions	`E`
Toggle Break at Sharp Angles	`A`
Select/Deselect/Subset Zone	`SHIFT+Z`
Clear Selection	`SHIFT+N`
Close and return to Default mode	`L`

Selection Control: Find Leak Path Mode

Shortcuts when using **Find Leak Path** mode over-ride the shortcuts in other modes.

Action	Keyboard Shortcut
Find Leak Path	`D`
Close and return to Default mode	`L`

Display Control

Action	Keyboard Shortcut
Show All	`CTRL+A`
Hide All	`CTRL+SHIFT+A`
Show Selected	`CTRL+F`
Hide Selected	`CTRL+SHIFT+F`
Show Faces on Surfaces	`CTRL+B`
Hide Faces on Surfaces	`CTRL+SHIFT+B`
Grow Displayed	`CTRL+'+'`
Shrink Displayed (see note below)	`CTRL+'-'`

On a Windows computer, the shrink displayed shortcut (CTRL + '-') does not always function when using the minus (-) button from the numeric row. In this instance, use the minus button from the numeric keypad instead.

Pop Up Menu Of The Surface Preparation Graphics Window

The following pop-up menu items are available when you right-click a part or part surface in the *Graphics* window, while using the surface preparation tool.

Hide	Hide the selected part or part surface from the display.
Show	Show the selected part or part surface in the display.
Show Only	Show only the selected part or part surface and hide all other parts or part surfaces.
Select	Select the current part, part surface, or part curve and deselect all existing selections.
Add Select	Add the current part, part surface, or part curve to the existing selection
Deselect	Remove the current part, part surface, or part curve from the selection (if it has been selected).
Subset Select	Select the current part, part surface, or part curve as a subset of the current selection. As such, the selected part, part surface, or part curve must already be included in the selection.
Set in Edit Panel	Select the current part surface or part curve in the Part Surface/Part Curve Edit group box.

See also: common menu items.

See Also

Assigning Faces to a Target Part Surface
Selection Options
Hiding And Showing Surfaces In The Graphics Window

Displaying the Geometry

Displaying the geometry allows visual inspection of the geometry surface before generating a surface and volume mesh.

The easiest way to display the geometry is with a geometry scene. By default, this scene contains a geometry and outline displayer that have the *Geometry* and *Outline* properties that are activated respectively. You can change the properties of these displayers to control the display. See Geometry Displayer Properties.

To create a geometry scene:

1. Right-click the **Scenes** node and select **New Scene** > **Geometry**.



The geometry scene appears in the *Graphics* window. By default, the scene displays all parts with a constant color.



You can use the various view manipulations to navigate around the scene by panning, rotating, and zooming.

After creating a geometry scene, you can identify and hide regions and boundaries to simplify the display.

See Also

Changing a Three-Dimensional Visualization Display
Managing Parts in Displayers
Highlighting an Object Using the Tree Node

Defining the Simulation Topology

An essential part of the process in setting up your simulation is to define the relationship between geometry parts and regions, boundaries, and interfaces.

Geometry parts are used only to define the faces, edges, and vertices that make up the surfaces of the model. The simulation domain on which mesh is generated, and for which physics is solved, is defined using regions, boundaries, and interfaces. These simulation domain objects can be created automatically using the hierarchy of geometry parts.

If you import geometry into STAR-CCM+, you can choose to interpret the geometry directly as regions, boundaries, and interfaces without going through geometry parts. This method was the default in earlier versions of STAR-CCM+. However, the recommended approach is now to use geometry parts initially and assign them to the simulation domain after making any required changes to surface or feature curve definitions.

Assigning Geometry Parts to Regions

The mesh continuum models cannot be used for geometry parts. In order to mesh geometry parts, they must be assigned to one or more regions, and the mesh is then generated for the region.

The following diagram illustrates the relationship between objects in the part tree and the ones in the region tree:



A part must be assigned to a region before its part surfaces can be assigned to boundaries, and its part curves to feature curves, within the region. The relationships do not need to be one-to-one; multiple parts can be assigned to the same boundary, and multiple part curves to the same feature curve. The following mapping would also be valid, for example:



It is important to note that part contact information can only be used to generate interfaces automatically if all contacting parts are assigned to new regions in the same operation.

See Also

Understanding Geometry Parts

Assigning Parts to Regions

Before a leaf-level part can be used in the meshing process, it must be assigned to a region.

A bi-directional relationship exists between leaf-level parts and regions such that you can think of it in terms of "this part belongs to this region" or alternatively, "this region is composed of these parts". Several methods exist to form these relationships:

- by using Set Region;
- by using the Assign Parts to Regions dialog; or
- by using the Parts Property in the Region Node

Any of these methods can be used to define the desired relationship between parts and regions. However, although both composite parts and leaf-level parts can be assigned to a region (using the first method that is described above), the relationship is ultimately between the leaf-level part and the region. Additionally, a leaf-level part can only belong to one region at a time. Therefore if a part is reassigned to a new region it is automatically removed from the original region part list.

Along with the part being assigned to the region, all surfaces and curves belonging to the part are also assigned to the region boundaries and features curves respectively.

Using Set Region

The Set Region option lets you assign parts to regions.

To assign a leaf-level part directly to an existing region, right-click the part node and select **Set Region** > **Existing**. For example:



Select the name of the chosen region from the object selector.



In the above example, the **Body 1** leaf-level part is assigned to **Region 1**. More than one leaf-level part can be selected and assigned to the same region at the same time using this method.

In a similar way to leaf-level parts, a composite part can be assigned to a region so that all of its child leaf-level parts are automatically linked to the region. This linking can be done by right-clicking on the composite part node and selecting **Set Region** > **Existing**, as shown below, and then selecting a region from the object selector.



This action results in all child leaf-level parts under the composite part being assigned to a region.

Setting a part to a region using this method automatically assigns all surfaces in that part to the `default` boundary in the region. If the default boundary does not exist or has been renamed, all part surfaces are assigned to the boundary with the lowest index number. The index number of each boundary is shown in the Boundary Properties window. If the name of a part surface matches the name of an existing boundary, the part surface is assigned to that boundary.

The region that a leaf-level part is assigned to can be reviewed or modified using the part node properties. Selecting the part node of interest shows the *Properties* for that node:



The *Region* property activates an object selector to make the region selection from the available list. New regions cannot be created using this method.

Using the Assign Parts to Regions Dialog

The Assign Parts to Regions dialog lets you assign parts to regions, boundaries, or feature curves.

The *Assign Parts to Regions* dialog can be used to assign leaf-level parts to regions. To bring up this dialog, right-click on the desired part node and select **Assign Parts to Regions...**



The *Assign Parts to Regions* dialog appears.



The dialog is split into two main sections. The top section shows all the current parts in the simulation. The bottom section contains all the options for setting the regions, boundaries, and feature curves.

To select or deselect a part, click the selection button to the left of it in the selection column.

Clicking the **Apply** button creates the region with the selected settings. Clicking **Close** closes the dialog.

The **Create Interfaces from Contacts** option (on by default) can be used to generate automatically interfaces that are based on part contacts of the selected parts. The valid settings for regions and boundaries depend on the method that is used to obtain the part contact data, as summarized in the Table of Valid Methods and Options.

The Set Region operation is independent of the **Assign Parts to Regions** operation.

Region Options

The region options control how the selected part is assigned to a region.

The available settings for assigning parts to regions are listed below:



Create a Region for Each Part: this option creates a different region for each selected part.

Assign All Parts to an Existing Region: this option activates an object selection dialog in single-select mode. This option does not appear if no regions exist in the simulation.

Create One Region for All Parts: this option is selected by default and assigns the selected parts to one region. Selecting this option makes a text box available for naming the region. Enter the name of an existing region (if applicable) or a region that you wish to create.

Boundary Options

The boundary options control the boundary creation of the selected parts.

The available settings for boundary creation are listed below:



Create One Boundary for Each Part Surface: this option creates different boundaries for each part surface.

Create One Boundary for Each Part: this option creates different boundaries for each part.

Create One Boundary for All Part Surfaces: this option assigns all the part surfaces to one boundary. Selecting this option makes a text box available for naming the boundary. Enter the name of an existing boundary (if applicable) or a boundary that you wish to create.

Feature Curve Options

The feature curve options control how the edges are assigned to feature curves.

The available settings for creating feature curves are listed below.



Create a Feature Curve for Each Part Curve: this option creates different feature curves for each part curve.

Create One Feature Curve for All Part Curves: this option is selected by default and assigns all selected part edges to one feature curve. Selecting this option makes a text box available for naming the feature curve. Enter the name of an existing feature curve (if applicable) or a feature curve that you wish to create.

Automatic Naming of Regions and Boundaries

Using the *Assign Parts to Regions* dialog automatically names the regions and boundaries that are created.

In the automatic naming process, the name that is given to a region or boundary differs depending on what settings are used. In the following example, both cases have the same part and part surface name: **Body 1** and **Inlet** respectively. In both cases, the selected boundary option is **Create a Boundary for Each Part Surface**.

Region mode:	One region for each part	One region for all parts
New region name:	Body 1	Region
New boundary name:	Inlet	Body 1.Inlet

Using the Parts Property in the Region Node

The parts that are assigned to a region can be set using the region node properties.

Selecting the region of interest displays the *Parts* property. For example:



Click the right half of the property to activate the object selector. Use the features of the object selector to access and select the names of the chosen parts.



After you close the object selector, the *Parts* property of the region shows the leaf-level parts that are assigned to the region node.

When regions are created from the **Regions** manager node, initially no parts are associated with them. When parts are associated using one of the other methods that are described in this section, the *Parts* property is automatically populated.

Assigning Part Surfaces to Boundaries

As with parts and regions, a bi-directional relationship exists between part surfaces and boundaries.

Individual part surfaces cannot be assigned to a boundary until the parent part has been assigned to a region first. There are two different methods by which the surface/boundary relationship can be established:

- by assigning part surfaces to a boundary directly; or
- by selecting part surfaces for a given boundary

Either or both methods can be used. A surface however can only belong to one boundary at a time so if it is reassigned to a new boundary it is automatically removed from the original boundary list. To unassign a surface from a boundary, remove it from the *Part Surfaces* properties in the **Regions** > **Region node** > **Boundaries** > **Boundary** node that contains the surface. This step is outlined in the second method above.

After a surface face is assigned to a boundary, its color in the current scene changes from a light silver color to a darker gray color. This change indicates which surfaces have been assigned and which have not.

See Also

Hiding And Showing Surfaces In The Graphics Window

Assigning Surfaces to Boundaries Using the Part Surface Nodes

The Set Boundary option lets you assign surfaces to boundaries.

To assign a surface face directly to a boundary, right-click the face node and select **Set Boundary** > **Existing**. For example:



Select the name of the chosen boundary from the object selector.

More than one face node can be selected and assigned to the same boundary using the above method.

A new region always contains one boundary by default, usually with the name **Default**. If a new boundary is preferred for the selected face, then the **Set Boundary > New...** option can be selected instead, as shown below:



A dialog appears with a text box for you to name the new boundary. The boundary name matches the selected face node by default.



After you click **OK**, a new boundary is created and is automatically assigned the face. The new boundary is shown below:



If two or more face nodes are selected for the new boundary option, then the boundary name is based on the keyword **Boundary** instead of the face node name.

If the region to which the parent part is assigned has no boundaries existing for it, only the **New** suboption is available.

Alternatively, the boundary that a face node is assigned to can be reviewed or modified using the face node properties. Selecting the face node of interest shows the properties for that node:



Click the right half of the property to activate the object selector, which only shows boundaries belonging to the region to which the parent leaf-level part is assigned. A new boundary cannot be created using this method.

Assigning Surfaces to Boundaries Using the Region Boundary Nodes

The part faces that are assigned to a boundary can be set using the boundary node *Properties* window.

Selecting the boundary node of interest displays the *Part Surfaces* property. For example:



Use the object selector to assign faces to the boundary.

When boundaries are created from the **Boundaries** manager node, initially no part surfaces are associated with them. When part surfaces are associated using another method, the *Part Surfaces* property is automatically updated for the boundary.

Assigning Part Curves to Feature Curves

Before you can assign individual part curves to a feature curve, you must assign the parent part to a region.

By default this action assigns all child curves to the first feature curve that belongs to the region. If necessary, modify the arrangement so that some edges belong to specific or new feature curves. As for parts and regions, a bi-directional relationship exists between part curves and feature curves. As such, there are two different methods by which the part curve/feature curve relationship can be established:

- By assigning part curves to a feature curve directly; and
- By selecting part curves for a given feature curve

Either or both methods can be used to define the desired arrangement between part curves and feature curves. However, a part curve can only belong to one feature curve at a time. Therefore if it is reassigned to a new feature curve, it is automatically removed from the original feature curve list. To unassign a part curve from a feature curve, remove it from the *Part Curves* properties in the **Regions** > **Region node** > **Feature Curve** > **Curve** node that contains the part curve. (See the second method above.)

Assigning Curves to Feature Curves Using the Part Curve Nodes

The feature curve to which an edge node is assigned can be reviewed or modified using the edge node properties.

Selecting the edge node of interest shows the *Properties* window for that node:



The *Feature Curve* property can be used to make the feature curve selection from the object selector. The object selector only shows feature curves belonging to the region to which the parent leaf-level part is assigned. If no feature curves exist, create one before proceeding with this step.

Assigning Curves to Feature Curves Using the Region Feature Curve Nodes

The part curves that are assigned to a region feature curve can be set using the feature curve node properties.

Selecting the feature curve node of interest display the *Part Curves* property. For example:



To assign part curves to a feature curve, click the right half of the *Part Curves* property to activate the object selector.

Select the part edges of interest and click **Close**. The *Part Curves* property now shows the part curves that are assigned to the feature curve node:



When new feature curves are created from the **Feature Curves** manager node, initially no part curves are associated with them. When part curves are associated using another method, the *Part Curves* property is automatically updated.

Transforming Regions

There are four types of transformations that you can apply to regions: Scale, Translate, Rotate, and Reflect.

The transformations are typically used to create separate regions when more than one mesh is imported, and these regions would not match spatially without some sort of transformation.

The transformations can be done on a single region or on multiple regions at the same time. The transformations that you apply to a region also affect the interfaces and boundaries that are attached to the region.

When regions have one of the following assigned to them–and no volume mesh, interfaces form between the surface meshes of the regions:

- geometry part
- imported part
- surface wrapper
- remeshed surface

Interfaces between surface meshes require physical contact between the parts that contain the regions, however, interfaces between volume meshes do not. Interfaces can exist between widely spaced regions of parts with volume meshes. For this reason, you can transform regions of parts with volume meshes without affecting the interfaces.

Requirements when Transforming Regions:

You can transform a region that is connected to another region through an interface without affecting that interface as long as both regions have a volume mesh.

When you transform regions that have volume meshes and are connected through an interface:

- There are no part contacts between the regions
- No errors occur
- The transformation is successful

If you attempt to transform a region that is connected to another region through an interface and one or both regions are assigned to parts without a volume mesh, an error occurs. The following error message is displayed:



When you transform regions that are assigned to parts without a volume mesh:

- Interfaces are created from part contacts between the surface meshes of the parts
- Part contacts can not be broken
- The transformation fails



When you transform regions and one or more (but not all) regions are assigned to parts without a volume mesh:

- The transformation is successful
- The interface is not maintained

To Meet Requirements when Transforming Regions:

- Create volume meshes for all regions that are connected through an interface.
- Translate all coincident parts with the region that you want to translate. Then imprint the parts and re-create the part contacts.

Scaling Regions

Scaling regions is especially useful when the mesh has been created in non-SI units.

If the region is displayed in an open *Graphics* window, the scaling is applied immediately to that window. The schematic below illustrates the scaling operation (with a scaling factor less than unity).



To scale a region:

1. Right-click the region of interest and select **Transform > Scale...** You can select multiple regions if you want to scale more than one region by the same amount.



The *Scale Regions* dialog appears.



2. By default the **Scale all dimensions** option is activated, which means that all dimensions are scaled equally. In this mode, whatever you type into the **X** field is also entered into the **Y** and **Z** fields. Deactivate the **Scale all dimensions** option to scale each coordinate direction independently. In both modes, typing a number greater than 1 expands the mesh. Numbers smaller than 1 shrink the mesh.
3. Click **Apply** to begin the scaling process.

   The status bar at the bottom of the simulation tree shows the progress. After the scaling is complete, the progress bar displays `Finished`, and a summary message is displayed in the *Output* window.



4. Click **Close** to close the dialog if no additional scaling is required.

Translating Regions

Translating regions is useful if you import multiple regions separately and you want to move them relative to each other.

In addition to changing the relative position of regions, you can use this feature to make sure that the region positions correspond to the position in space appropriate for plot data and tabular data sets. If the region is displayed in an open *Graphics* window, the translation is applied immediately to that window. The schematic below illustrates the translate operation.



To translate a region:

1. Right-click the region of interest. Select **Transform > Translate...** You can select multiple regions if you want to translate more than one region the same amount.



The *Translate Regions* dialog appears.

2. In the *Translation* fields, enter real numbers into the **X**, **Y**, and **Z** text boxes to define the translation vector. The units of each direction correspond to the model units. You can enter a negative value to move the region in the negative axis direction.

Note: An invalid entry in any of the fields causes the background of the field to turn red. Clicking **Apply** with an invalid field automatically adjusts the invalid value to `0.0`.

3. Click **Apply** to translate the selections.

The progress bar notifies you once the process is complete and a summary message is displayed in the *Output* window.



4. Click **Close**.

Rotating Regions

Region rotation can be used if the orientation of the imported mesh is not suitable for viewing, or if one region does not match the orientation of other regions.

When the region is rotated, the vector and tensor solution variables are rotated as well to keep them consistent. If the region is displayed in an open *Graphics* window, the rotation is applied immediately to that window. The rotation operation is illustrated below.



To rotate a region:

1. Right-click the region of interest and select **Transform > Rotate...** You can select multiple regions if you want to rotate more than one region the same amount.



The *Rotate Regions* dialog appears.



2. In the **Angle (degrees)** field, specify the amount of rotation. Rotation is specified using the right-hand rule around the axis of rotation. Negative entries are treated the same way as their positive counterparts.
3. Specify the axis of rotation using the *Axis Direction* and the *Axis Origin* components. For the example below, assuming you want to rotate around the red dashed axis, the *Axis Direction* components would be [1, 0, 0] and the *Axis Origin* components would be [0, 1, 0]:



Invalid entries cause the background of the field to turn red.

4. Click **Apply** to rotate the selections.

   The progress bar notifies you once the process is complete and a summary message is displayed in the *Output* window.



5. Click **Close**.

   Notes

   - Two-dimensional problems are treated in the same way as three-dimensional problems. Do not rotate the mesh off the X-Y plane.
   - When the mesh rotates, vector and tensor solution variables are rotated as well to keep them consistent. Specifically, velocity is rotated for flow models.
   - All active scenes update to reflect the new mesh configuration.

Reflect Regions

You can use the reflection option to reposition the region according to its mirror reflection about an axis or combination of axes.

If the region is displayed in an open *Graphics* window, the reflection is applied immediately to that window. The reflection operation is illustrated below:



To reflect a region:

1. Right-click the region of interest and select **Transform > Reflect.**



The *Reflect Regions* dialog appears.



2. In the *Normal Vector* component, specify the axis about which you want the region to reflect.
3. The **Coordinate System** menu can be used to specify the coordinate system for reflection.
4. Click **Apply**.

The progress bar notifies you once the process is complete and a summary message is displayed in the *Output* window.

5. Click **Close**.

Region Manipulation

To fix surface or volume meshes that are not suitable for obtaining a valid solution, choose from among various techniques for manipulating regions.

The region manipulation techniques are:

- Combining regions
- Splitting regions
- Creating new regions from boundaries (for surface meshes)

Also, the following common actions are available for region manipulation:

- Deleting regions
  - Deleting regions removes any associated Interfaces and reinstates the original boundaries of the remaining regions.
- Copying region properties
  - You only copy region type and continuum properties. Boundary conditions, values, or cells, are not copied.

Regions are represented in the STAR-CCM+ simulation under the **Regions** manager. The region properties are accessed by selecting the node. Right-clicking the node accesses the pop-up menu.


Servers
Star 1
simulation
scene/plot
Star 1
Geometry
Continua
Regions
Catlyst
Exhaust
Manifold
Muffler
Derived Parts

Region Combination

To have two or more computational domains share physical properties, combine regions.

There are two approaches that you can use to combine regions:

- Copy properties between regions and then create interfaces between adjacent regions.
- Combine the regions into one, as the following diagram illustrates, and then fuse or create interfaces between any adjacent boundaries.



If you do not fuse or create interfaces between adjacent boundaries, STAR-CCM+ solves the combined regions as if they are separate regions with the same properties. That is, there is no flow between the regions.

Note: It is more computationally efficient to fuse adjacent boundaries than to use interfaces.

You can combine multiple regions, providing:

- All regions share a surface or volume mesh representation.
- You associate all regions with the same mesh and physics continuum.

Initializing or running the solver reintersects interfaces. Therefore, combining regions before creating interfaces is more efficient, because the combination process resets interfaces and removes vertices from the intersection. Also, if these interfaces reference faces and vertices, the combination can remove them from other regions.

STAR-CCM+ allows combined regions to be non-contiguous. That is, there is no contact between the computational domains, but they share physical properties. The following diagram illustrates the combination of non-contiguous regions.



Combining Surface Meshed Regions

You can immediately combine two surface-meshed regions using the region pop-up menu.

Note: The name of the combined region is the same as the region that you select first.

To combine regions that have a surface mesh:

1. Multi-select two or more regions.
2. Right-click a selected region node and select **Combine...**



Once the combine is complete, the *Output* window displays a summary of the operation.



Combining Volume-Meshed Regions

You can combine two volume-meshed regions using the region pop-up menu.

Note: The name of the combined region is the first region in the *Combine Regions* dialog (sorted in alphabetic order).

This operation can require you to fuse adjacent boundaries.

To combine regions that have a volume mesh:

1. Multi-select two or more regions.
2. Right-click a selected region node and select **Combine...**



3. In the *Combine Regions* dialog, specify the region combination options.

You can also change the selected regions within the dialog.



4. Click **OK** to combine the regions.

Once the combine is complete, the *Output* window displays a summary of the operation.

```
Output - Star 1

Merging region Region_2 into region Region_1
Loading/configuring connectivity (old|new partitions: 1|1)
  Region_1 : 14598 cells, 97340 faces
  Region_2 : 5481 cells, 35667 faces
  Region_3 : 14497 cells, 96639 faces
Configuring finished
Resetting interface Body 2/Body 3
Resetting interface Body 1/Body 2
Started Parasolid modeler version 24.01.143
```

Combine Regions Dialog

The *Combine Regions* dialog provides properties for combining and fusing boundaries during the combination of volume meshes. This dialog also allows you to change the selected regions.

The following properties are available in the *Combine Regions* dialog:

Property Name	Property Description
Combine boundaries with similar names	When activated, this property combines boundaries that have the same base name. That is, boundaries that are the same up to the first space. For example, Symmetry 8 and Symmetry 9, but not Symmetry-8 and Symmetry-9. When deactivated, this property appends a space and an incremental digit to the name of one of two boundaries that share a name. For example, Symmetry and Symmetry become Symmetry and Symmetry 1.
Fuse adjacent boundaries	When activated, this property combines adjacent boundaries and replaces them with interior cell faces; thus removing the boundary.
Fuse conformal faces only	When activated, this property only combines adjacent boundaries that

Property Name	Property Description
	have cell faces with matching vertices (conformal).
Tolerance	Defines the maximum distance between fused boundaries as a percentage of the smallest edge length that is attached to each vertex. This property is only available when you select **Fuse adjacent boundaries**.

Region Splitting

Region splitting allows you to modify the simulation topology by separating regions using several characteristics. Thereafter, you can apply different continua and conditions to different parts of the model.

The following methods are available for splitting regions:

- Splitting by Surface Topology
- Splitting Non-Contiguous Regions
- Splitting Regions with a Field Function

Splitting By Surface Topology

Splitting by surface topology splits closed, semi-manifold surface meshes into individual regions.

The following diagrams illustrate two scenarios where the split by surface topology option is useful:

- When one or more closed surfaces exist within another closed surface and each share a common boundary.



- When two or more closed surfaces share a common boundary.



You typically find these types of geometry from CAD exported files that provide one region for all parts or bodies, with a common boundary between each.

If the regions have two adjacent boundaries, as opposed to one shared boundary, use non-contiguous splitting and manually create an in-place interface.

Note: Splitting by surface topology is only available when a surface mesh representation is present. Delete any volume mesh representations to access this option.

To split the geometry by surface topology:

1. Right-click on a region and select **Split by Surface Topology**.



2. In the *Split Region Options* dialog, specify the region split options



3. Click **OK** to split the regions.

   If you split invalid topologies, the resultant surface mesh requires repairing.

The split by surface topology operation copies shared boundaries to ensure that all regions are closed. The split also creates an in-place

interface between copied boundaries to create a conformal mesh between the regions.

Repairing Invalid Topology Splits

Invalid topologies are surfaces that are not closed. The split by topology operation can split these surfaces, but the resultant surface mesh requires repairing.

If you split an invalid topology, the Output window displays the following warning:

```
WARNING created invalid regions during split by topology
```

To fix the split by topology surface mesh:

1. Right-click the surface mesh representation and select **Repair Surface...**



2. In the *Surface Preparation* tool, use the Surface Diagnostics to identify surface mesh errors.



3. Use the Surface Repair tools to fix the surface mesh errors. The following tools are useful for fixing invalid topologies:
   - Zip selected edges
   - Delete faces
   - Fill holes

Split Region Options Dialog

The *Split Region Options* dialog provides a property that controls the region combination after the splitting process.

The following property is available in the *Split Region Options* dialog:

Property Name	Property Description
Combine Regions with Shared Boundaries?	When activated, this property combines regions that have the same boundary name after the split. This property is useful when many internal geometries exist that you want to belong to the same region and continuum.

Splitting Non Contiguous Regions

Non-contiguous region splitting creates regions when the original region contains surface or volume meshes that are spatially separated.

The following diagram illustrates non-contiguous region splitting:



The following list defines contiguous regions:

- For surface meshes, each face shares an edge with another face in the region.
- For volume meshes, each cell shares a face with another cell in the region.

To split non-contiguous regions:

1. Right-click a region and select **Split Non-Contiguous**.



2. In the *Split Non-Contiguous Regions* dialog, set the region splitting options.

   To change the selected regions, click the **Select** button.



3. Click **Apply** to split regions.

   If you attempt to split contiguous regions, the *Output* window displays a warning.

Split Non Contiguous Regions Dialog

The *Split Non-Contiguous Regions* dialog provides properties to control the splitting of non-contiguous regions. This dialog also allows you to change the selected regions.

The following properties are available in the *Split Non-Contiguous Regions* dialog:

Property Name	Property Description
Set Representation	Specifies the mesh representation to which the split applies. For example, if you want to split the remeshed surface, set *Set Representation* to **Remeshed Surface**.
Minimum Connected Face Area	Splits contiguous regions which have connected faces with a surface area below the user-defined tolerance. For example, if two regions share a face, you can split these regions by setting *Minimum Connected Face Area* to a larger value than the connected face surface area.

The following button is available in the *Split Non-Contiguous Regions* dialog:

Button Name	Button Description
Preview	Reports how many regions the split operation generates. This button is only available when you set *Select Representation* to a volume mesh.

Splitting Regions With A Field Function

Splitting regions by a field function splits a surface or volume mesh region according to a user-defined expression.

This operation does not split cells. Therefore, if the cell vertices are not co-planar, the split leaves a non-uniform join between the two regions. The split generates two boundaries between the regions, from which you can create interfaces.

Note: The naming conventions for new regions follow the ascending order of the number of cells in the region. For example, if the field function creates two new regions from Region 1, the new region with the least cells is Region 1 2, and the most cells is Region 1 3.

The following diagram illustrates field function region splitting:



To split a region with a field function:

1. Define a field function that produces the required split
2. Right-click a region and select **Split By Function...**



3. In the *Split Regions by Function* dialog, set **Function** to the required field function.

   To change the selected regions, click the **Select** button.



4. Click **Apply** to split the regions.

Example of a Field Function to Split a Region

This example splits a region into three as the following image illustrates.



You can split the region with the following field function:

```
($$)Position[2] <= 0.05)? 1:(($$Position[2] >= 0.2)? 2:0)
```

In this case, after the split:

- The cells where 0.05 < z < 0.2 belong to Region 1, as the field function does not affect them.
- The cells where z ≤ 0.05 belong to Region 1 2, as it is the created region with the least cells.
- The cells where z ≥ 0.2 belong to Region 1 3, as it is the created region with the most cells.

Split Regions By Function Dialog

The *Split Regions by Function* dialog allows you to select the field function for splitting. This dialog also allows you to change the selected regions.

The following property is available in the *Split Regions by Function* dialog:

Property Name	Property Description
Function	Specifies the field function that the split operation uses.

The following button is available in the *Split Regions by Function* dialog:

Button Name	Button Description
Preview	Reports how many regions the split operation generates.

Region Creation

You can create regions from existing boundaries to apply different physical properties to a portion of the fluid domain. Alternatively, you can create an empty region and assign parts to it.

Importing surfaces as regions automatically creates one or more regions from the surface topology and translation options. However, this region topology can be unsuitable for multi-region meshing and require redefinition by reallocating boundaries.

Additionally, importing surfaces as parts does not create any regions. Therefore, you can create regions so that you can allocate the appropriate parts before meshing.

When dealing with parts, you can also assign parts to regions automatically.

Creating Regions From Boundaries

Creating regions from boundaries moves boundaries from one region to a new region. This feature allows you to modify the simulation topology and apply different physical properties to part of the fluid domain.

Before creating a region from boundaries, ensure that the created region will not share a boundary with any other region. In such cases, convert the shared boundary to an interface. This conversion creates an identical boundary and defines an in-place interface to connect the boundaries. You can then use the duplicate boundary to define the new region.

The following diagram illustrates the creation of regions from boundaries:



To create a region from boundaries:

1. Multi-select the required boundaries

2. Right-click one of the selected boundaries and select **Create New Region from Boundaries**.



If any region is open after this procedure, you can use the hole filler to close the surface. If you use the hole filler on a shared boundary, use an in-place interface to connect the boundaries.

Creating An Empty Region

You can create an empty region and allocate parts to it for meshing.

An empty region is a region with no assigned parts. If a mesh continuum exists, creating an empty region assigns this continuum to the new region. Also the empty region has an empty wall boundary and feature curve so that you can assign part surfaces and part curves.

Note: The meshing models ignore empty regions.

To create an empty region:

1. Right-click the **Regions** manager and select **New**:



Regions Manager Properties

The Regions manager node provides information on the number of regions in the simulation.

The following property is available from the **Regions** manager:

Regions	The number of regions in the simulation (read-only). This property changes when you add, combine, split, or delete regions.

Regions Manager Pop Up Menu

Right-click the Regions manager node to add regions or group them.

The following options are available when you right-click the **Regions** manager:

New Group	Creates a group for categorizing regions by creating a group manager in the **Regions** manager.
Group By	Group regions by: • Physics Conditions • Physics Values • Mesh Continuum • Physics Continuum • Parts • Type
Ungroup	Reverses the grouping process.
New	Adds a region to the simulation.

Manipulating Boundaries

It is often convenient to modify boundaries in the simulation without having to modify the mesh and re-import it.

Boundaries are easily accessible since they are included among the objects represented in the simulation tree. The **Boundaries** manager node, which is subject to common menu items, contains all the boundaries in the region.



In STAR-CCM+ it is critical that boundaries be appropriately separated and demarcated for specific conditions and values to be assigned, for both meshing and analysis setup. Surface and volume meshes originating from outside STAR-CCM+ usually have boundaries already defined in some fashion.

In most instances, boundary manipulation is only allowed when a single mesh representation is present. This is so that boundary mappings from one mesh to another can stay consistent.

Combining Boundaries

Combining two or more boundaries into a single boundary is useful when the imported surface or volume mesh contains separate boundaries which should have identical types, conditions, and values.

Combining the separate boundaries allows the boundary conditions to be set once and simplifies the topology. Combining boundaries does not modify the faces of the selected boundaries, it simply groups them together into one boundary. This operation is illustrated schematically below.



The boundaries being combined must belong to the same region and cannot include boundaries generated from computed interfaces. The boundaries being combined do not have to be contiguous. If the original boundaries were not topologically contiguous (no face from one boundary shares an edge with any faces from another boundary), the process can be reversed using the split non-contiguous boundary action.

To combine boundaries together, open the **Regions** node and the **Boundaries** node in the region of interest. Select the boundaries of interest using a multiple selection technique and right-click on the selection. Select **Combine** from the pop-up menu.



This will immediately begin the combining process and display a progress bar at the bottom of the simulation window. Once the process is complete a message will be displayed in the *Output* window indicating which boundaries were merged into which.

The boundary that is clicked first is the one that all the others are merged into. The simulation tree will be updated once the combine is complete.



See Also

Synchronizing Multiple Objects

Fusing Boundaries

The fuse command can be used to combine a pair of boundaries from the same volume mesh region and replace them with interior faces.

The fuse operation is illustrated schematically below.



For the fuse command to be successful, the boundaries must be coincident in space, but not necessarily composed of conformal discretizations. If the two boundaries are completely conformal or can be matched exactly via the interface computation, the selected boundaries will be deleted, and the matched faces will be added to the interior of the region. If the faces cannot be completely matched, any faces that do match will be added to the interior of the region while the remaining faces will be left in the original boundaries.

The advantage of removing boundaries by fusing is that there is no additional overhead involved in the calculation (as might be incurred at an interface).

To fuse boundaries, open the **Regions** node and the **Boundaries** node in the region of interest. Select the boundaries you want to fuse using a multiple selection technique and right-click on the selection. Select **Fuse...** from the pop-up menu.



This opens the *Fuse Boundaries* dialog.



The originally highlighted boundaries will be automatically selected in the dialog with a ticked checkbox. If you wish to match only conformal faces (faces which have vertices at the same coordinates), select the **Fuse conformal faces only** option. Otherwise, the full interface intersection computation is performed on the selected boundaries.

The boundaries must be geometrically adjacent and the **Tolerance** in the previous dialog is a measure of close the boundaries must be to be considered suitable candidates for fusing. The tolerance represents the amount that each vertex is allowed to move during the intersection process, and is the same as the *Intersection* tolerance specified on direct interfaces.

After entering the tolerance, click the **OK** button. This will begin the fusing process and display a progress bar at the bottom of the simulation window. Once the process is complete a message will be displayed in the *Output* window indicating the results of the fusing process.



Since an interface intersection is done when fusing boundaries (to convert the boundary faces into interior faces) you will see several messages about faces and vertices changing.

The simulation tree will be updated once the fuse is complete, and if all of the faces were matched with no remainders, neither of the boundaries that were selected will exist after fusing.



In some situations the **Fuse...** menu item is grayed even though you have selected more than one boundary.



There could be four reasons for this:

- one or more of the boundaries are taking part in an interface;
- one or more of the boundaries were created by an interface;
- the selected boundaries belong to more than one region; or
- you don't have a volume mesh representation currently displayed

Merging Boundaries

Boundaries that belong to a surface mesh can be merged together if they are in the same region and have surfaces intersecting one another.

When two or more boundaries intersect then the merge operation will re-tessellate the surface so that all the triangles in the boundaries are joined together. In the case of overlapping/co-planar boundaries, a common shared boundary is created to replace the overlapping area, so as to facilitate the setup for multi-region meshing cases for example.

An alternative intersection process for surfaces, which can be in different regions, is provided in the intersect faces tool as part of surface repair. The surface repair tool can also be used to merge or imprint two overlapping boundaries.

Merging Intersecting Boundaries

The merge intersecting boundaries feature allows two or more intersecting surface mesh boundaries to be merged together so that the triangles making up the surface are joined together.

This type of merging is achieved by re-tessellating the locally intersected surface so that new vertices and triangles are introduced at the intersection. If the input boundaries fully intersect then the result is a closed (watertight) surface volume (but not necessarily manifold), as illustrated below:



In the above example, boundaries 1 and 2 are fully intersecting but not joined in any way. After the intersection process, both boundaries have been re-tessellated at the intersection plane, forming a closed volume. You should note that no boundaries are removed or created by this process, so any surface triangles that are not required after the intersection should be removed in order to leave a manifold surface.

New feature curves are however automatically created at the intersection.

To intersect two or more boundaries, open the **Boundaries** node in the region of interest and then select the boundaries which are to be intersected. Right-click on any of the boundaries and select **Intersect** from the pop-up menu:



The intersection process will proceed and the status will be reported in the *Output* window:



Once the process is done, the current scene will be updated with the new surface triangulation at the intersection of the selected boundaries. The simulation tree will also show the new feature curves created:



The following conditions should be noted when using the boundary merge intersection tool:

- The boundaries to be intersected must belong to the same region. If they belong to different regions, then you should combine the regions first before proceeding.
- The sequence in which the boundaries are selected is not important.
- The boundaries do not need to belong to a region that is closed and manifold. In other words, planar surfaces and “open” boundaries are allowed.
- The boundaries must intersect one another, otherwise nothing will happen.
- The quality of the original boundary tessellation is not critical.
- Important: The process is not reversible, so you should save the simulation first before proceeding.
- It is recommended that the surface remesher be used after a boundary intersection to improve the quality of the surface.

Splitting Boundaries

It is useful to split boundaries if the original mesh contains too few of them for modeling or analysis purposes.

Six different ways are provided to split boundaries into smaller pieces:

- splitting non-contiguous boundaries;
- splitting boundaries by angle;
- splitting by function;
- splitting by patches;
- splitting by feature curves; and
- splitting interactively

In all cases, the properties of the new boundaries will be the same as the original boundary.

See Also

Preparing Boundaries for the Generalized Cylinder Mesher

Splitting Non-Contiguous Boundaries

This function will split a surface or volume mesh boundary into its non-contiguous pieces.

The splitting of a non-contiguous boundary is shown schematically below. If the boundary is contiguous to begin with, no split will occur.



To split non-contiguous boundaries, open the **Regions** node and open the region of interest. Then right-click on the boundary or boundaries of

interest from the **Boundaries** node. The selected boundaries may be from different regions. Select **Split Non-Contiguous** from the pop-up menu.



The splitting process starts and a progress bar is displayed at the bottom of the simulation tree. Once the process completes successfully, a message will displayed in the *Output* window.



The simulation will be updated to reflect the new boundaries that were created. Their names will be derived from the names of the selected boundaries with an identifying digit appended.



If any of the boundaries in the selection are already contiguous, no changes will be made to the boundaries and a message to that effect will be displayed in the *Output* window.



Splitting by Angle

This function enables surface or volume mesh boundaries to be split according to the included angle between faces.

The operation of splitting by angle is illustrated schematically below.



This would be useful, for example, if four walls of a box had been created as a single boundary, but you would prefer to treat them separately to support different conditions. The separation will only occur if there is a contiguous surface that can be demarcated by the specified angle criterion.

This function will also automatically execute the splitting non-contiguous boundary action, which occurs even if no splitting occurs based on the angle criterion.

To split boundaries by angle, open the **Regions** node and open the region of interest. Then right-click on the boundary of interest from the **Boundaries** node. Select **Split By Angle...** from the pop-up menu.



This opens the *Split Boundaries by Angle* dialog:



The boundaries that had been selected in the simulation tree are selected, and the **Angle** defaults to 89 degrees. When the angle between the faces is acute (less than ninety degrees) then an **Angle** value of less than the internal angle should be specified in order to split the boundary. When the angle between the faces is obtuse (more than ninety degrees) then the specified **Angle** value should be less than the included angle between the boundary normals. For example, to split a boundary which has faces at 135 degrees to one another, the included angle is 45 degrees, so a value of 44 degrees or less needs to be used to split the boundary. This is illustrated below:



If more than one representation exists then the surface or volume to be operated upon should be selected using the **Representation** menu. Once a representation is chosen, the **Preview** button can be used to determine how many boundaries would be created based on the supplied angle. For example:



Here, the default value of 89 degrees would result in the **Wall** boundary in **Region 1** being split into 2 boundaries (the revised original plus a new one). This is a useful tool for determining the effect of the angle

value in cases where there is a large variance in angles between boundaries.

Pressing the **Apply** button will start the boundary splitting process. Should other representations exist for the simulation then they will have to be deleted first since the boundary correspondance between them will no longer be applicable. In this case, the *Ok to Delete Representations* dialog will appear:



On pressing the **OK** button, STAR-CCM+ will proceed to split the boundaries for the selected representation. A progress bar will be displayed at the bottom of the simulation tree. Once the process is complete messages in the *Output* window will indicate the new boundaries that were created:



The new boundaries will also appear on the simulation tree. The name(s) of the new boundaries are taken from the source boundary, with an identifying digit appended:



If the boundary could not be split based on the angle that you supplied, then a message to that effect will be displayed in the *Output* window indicating as such.

Splitting by Function

This operation splits volume mesh boundaries according to the rounded integral value of a specified field function.

The operation of splitting by function is illustrated schematically below.



In this example, the field function f has the value 0 or 1, so that two new boundaries are created. In this case, f might be defined as follows:

($$Centroid[0] < 10) ? 0 : 1

or even more simply

$$Centroid[0] < 10

both of which cause a step from $f = 0$ to $f = 1$ at the global x coordinate of 10.

In the above example, the field function has two discrete integer values. If it consisted of a range of real numbers, say $0.2 < f < 1.8$, the boundary will be split into three parts containing:

- cells in which $0.2 < f < 0.5$;
- cells in which $0.5 \leq f < 1.5$;
- cells in which $1.5 \leq f < 1.8$.

If the chosen field function has a large numeric range, then STAR-CCM+ will attempt to create as many new boundaries as there are rounded integral values given by the field function in the original boundary. In other words, if $0 < f < 100$, 100 new boundaries would be created. Therefore, to prevent the creation of an unreasonable number of new boundaries, you should ensure that the field function has a restricted range.

Suppose we wish to split two purple and yellow boundaries along the X direction, as shown in the screenshot below.



Use any technique to find the X location of the split - we will use a value of `0.28` for this case. Create a new user field function with the name **Split@X=0.28** using the **Tools > Field Functions** node and set the *Definition* to `($$Position[0]<0.28)? 0:1`.



Initialize the solution to ensure that the user-defined field function will hold valid values. If errors are reported in the *Output* window then these should be corrected.

Now open the **Regions** node and open the region of interest. Right-click on the boundary of interest from the **Boundaries** node and select **Split By Function...** from the pop-up menu.



This opens the *Split Boundaries by Function* dialog with the boundaries selected in the simulation tree marked with a ticked checkbox. Pull down the Function drop-down list and select the user-defined field function that you created earlier:



Click the **Apply** button to start the splitting function. A progress bar will be displayed at the bottom of the simulation tree. Once the process is complete a message will be displayed in the *Output* window indicating the names of the new boundaries and the number of faces in each.



The simulation tree will be updated with the new boundaries. The names of the new boundaries are taken from the original boundaries with an identifier appended.



The screenshot below shows the new split in the boundaries.



Note that the splitting process does not divide or cut boundary faces into two. The process simply divides the boundary up based on the existing face structure and retains the same total number of faces before and after.

Splitting by Patch

The split by patch function allows you to split a boundary based on the underlying patch definition of the imported surface geometry.

A patch represents a cell table, group, family or some other similar entity used in the original surface pre-processing software to differentiate one surface from another. This information is read into STAR-CCM+ as part of the surface import process and is maintained as part of the import representation. All surface file formats other than STL support patch definitions.

Patches are useful in STAR-CCM+ as they allow a quick and easy method of reconstructing the boundaries to match the original surface definitions should they not initially be imported this way. A boundary can contain multiple patches but a patch cannot be shared among two or

more boundaries. The patch definitions can change depending on how the boundaries are manipulated. For example, if a boundary is split up (based on an angle or feature curves etc), then each new boundary will have its own patch definition.

To split boundaries by patch:

1. Right-click on the boundary name or names (using multiple selection) in the **Boundaries** node and select **Split By Patch...**



Note: You must have a scene open to use this feature.

The *Split Boundary by Patch* window appears in the *edit* tab in the explorer pane.



Here, the patches that are contained in the import surface are listed numerically. The surfaces on the geometry in the scene window will also be color-coded to represent each patch number.

2. Select the patches from the list box or directly using the mouse in the scene display. To select multiple patches, press the `<Ctrl>` key, or use the zone selection tool from the Selection toolbar. Selecting a patch will result in it being highlighted in the scene window.



3. Enter the new **Boundary Name**.
4. Click **Create**. The relevant patch or patches will be removed from the dialog and scene so that only the remaining patches are shown.
5. Repeat the process as many times as required until only one patch is left.
6. Click **Close** when done.

Note: When only one patch remains, the **Create** button is disabled. The source name will be used for the remaining boundary.

The new boundaries appear in the simulation tree. In the following example, the boundaries were created based on 6 patches.



Splitting by Feature Curves

The split by feature curve function allows a boundary to be split based on feature curves that belong to the same surface mesh region.

In order for a boundary to be successfully split, the feature curve must completely separate the boundary into two or more pieces.

Feature curves are usually imported as part of the surface definition using any of the allowable surface file formats other than STL and PATRAN. Being able to split a boundary using feature curves is extremely useful as feature curves can be added to any part of the surface and allow you to form any arbitrary shape that you want. Once the boundary has been split, the features curves can be deleted if desired or left in as part of the definition if they are important for maintaining a surface feature or detail.

To split boundaries by feature curves, open the **Regions** node and then the region of interest. Right-click on the boundary name or names (using multiple selection) in the **Boundaries** node and then select **Split By Feature Curves...** from the pop-up menu.



The *Split Boundary by Feature Curves* dialog will appear, as shown below:



The boundary or boundaries originally chosen in the object tree will appear selected. Use the checkboxes to select the feature curves which will be used for splitting and then press **Apply** to start the splitting process. The progress will be displayed at the bottom of the simulation tree and the final number of split boundaries will be reported back in the *Output* window.

Once the process is complete, the simulation tree will be updated with the new boundaries.



The names of the new boundaries are taken from the original boundary or boundaries with an identifier appended. Note that the feature curves are not affected in any way by this process.

Splitting Interactively

The split interactively function allows a boundary to be split based on the triangulation of the surface mesh region.

When you split a boundary interactively, anywhere from one face to all the faces belonging to the boundary can be selected to belong to the new boundary. Only one boundary can be selected to be split at a time.

To split a boundary interactively, open the **Regions** node and then the region of interest. Right-click on the boundary name in the **Boundaries** node and then select **Split Interactively...** from the pop-up menu.



The *Split Boundary Manually* window will appear in the *edit* tab:



The scene will also display the model with all boundaries other than the one selected set to be transparent. By default, the mouse can be used to select faces in the displayed boundary, with multiple selections possible by pressing the Ctrl key while selecting additional faces. Double clicking on the displayed boundary will select all continuous faces constrained by feature curves or sharp edges. To unselect entities, you can click on them again (while pressing the Ctrl key) or click anywhere off the model.

The *Split Boundary Manually* tool provides access to five *Selection Control* buttons which can be used to further manipulate the face selections made by either growing from a face pre-selection or doing zone based selection/de-selections. These options are particularly useful when a large number of face selections need to be made which would otherwise be too time consuming to select individually. For more details on the buttons, refer to the selection control options for manual surface preparation, which work in a similar way.

An example model with a part of a boundary selected is shown below:



When you have selected the face or faces of interest, enter in the new **Boundary Name** and then press the **Create** button to make a new boundary.



The new boundary name will appear in the simulation tree.



If you do not specify a new name then the original boundary name with a numerical increment will be used instead. The newly created boundary will be removed from the original displayed boundary.

The process can be repeated as many times as necessary for a given boundary and the **Close** button pressed when the operations are complete. If all faces for a given boundary are selected to become a new boundary then you will be left with an empty original boundary which should be then deleted.

Creating New Boundaries

New empty boundaries can be created at any time so that part surfaces can be allocated to them as required in order to form regions for meshing purposes.

Note: Part surfaces cannot be allocated to a boundary until the part that the surface belongs to has been assigned to the parent region of the boundary.

To create a new boundary, right-click on the **Boundaries** manager node and select **New** from the pop-up menu:



The simulation tree will be updated with a new boundary node called **Boundary 1** by default:



By default, no part surfaces are assigned to this boundary but these can be manually selected. The default boundary type is also set to wall but this can be changed at anytime thereafter.

Projecting Boundaries

Boundaries belonging to a volume mesh can be projected in order to modify an existing boundary surface to conform to a given plane.

The following operations are possible:

- projecting boundaries to a plane (2D and 3D meshes); and
- projecting boundaries to an axis (2D meshes only)

Projecting Boundaries to a Plane

The project to plane function can be used to modify a non-planar boundary on a 3D or 2D volume mesh and make it planar.

This projection is shown schematically below:



A typical application would be to make an undulating boundary more planar so that it could better be used for either the creation of an interface or for extracting a two-dimensional mesh from a three-dimensional mesh.

The project to plane option should be used with caution since it has the potential for adversely affecting the quality of the adjacent cells. In fact, it is possible to severely distort a boundary (and the adjacent cells) by specifying the wrong plane for projection as illustrated below:



To project a boundary, open the region of interest in the **Regions** node. Then right-click on the boundary of interest from the **Boundaries** node. Select **Project to Plane...** from the pop-up menu:



This opens the *Project Boundaries to Plane* dialog with the selected boundaries already marked with a ticked checkbox.



The plane is defined by specifying the length vector, in the *Plane* group box, that is normal to the plane onto which you want the projection to take place. The **Normal X**, **Normal Y** and **Normal Z** inputs specify the direction of the vector and the **Distance** input specifies the length of the vector. The example above specifies a plane that is parallel to the X-Z plane and that is 0.5 meters from the origin. The parameters are illustrated for a general case below.



Click the **Apply** button to start the projection process. A progress bar will be displayed at the bottom of the simulation tree. Once the process has completed successfully, a message will be displayed in the *Output* window indicating the boundaries that were projected and the plane that was used.



Projecting Boundaries to an Axis

The primary application for projecting boundaries to an axis is to ensure that the correct mesh definition is prescribed for an axis-symmetric type analysis.

The project to axis option can be used on a 2D mesh containing an axis type boundary (for axis-symmetric cases). It is similar to the project to a plane option in that it can take non-planar vertices on an axis boundary and automatically project them to the y=0 axis.

To project a boundary to the axis, the mesh must first have been converted to 2D and an axis type boundary defined. Selecting the axis

boundary node and right-clicking will give access to the **Project to Axis...** option, as shown below:



The *Repair Axis Boundary* dialog will appear:



Clicking the **OK** button will proceed with the operation and project the boundary vertices to y=0 in the global coordinate system. Once the process is complete, a confirmation message will appear in the *Output* window:



Note that this operation is not reversible so it is advisable to save the simulation beforehand.

Selecting Boundaries

A dialog is available for selecting boundaries for various operations.

For most of your work with boundaries, STAR-CCM+ opens a standard object selector dialog. It allows selection of multiple boundaries from various regions.

Copying Boundary Properties

If you have two (or more) boundaries that you want to have exactly the same *Type* property, but you have chosen not to combine them, then you can copy that property from one boundary to another.

This feature works on boundaries for the *Type* property only. No other properties, such as interfaces or faces, are copied between boundaries using this operation. This operation does not work with interface boundaries.

Properties of the Boundaries Manager

The Boundaries manager node provides information on the number of boundaries in the simulation.

Boundaries	The number of boundaries in the simulation (read-only). This property changes when you add, combine, split, or delete boundaries.

Pop-Up Menu of the Boundaries Manager

Right-click the Boundaries manager node to add boundaries or group them.

These are the pop-up menu items available when you right-click the **Boundaries** manager node.

New Group	Facilitates the grouping of objects by creating a new group node within the manager node.
Group By	Allows you to group the objects according to any of the properties of the objects in the manager node, listed in a submenu of this menu item.
Ungroup	Reverses the grouping process.
New	Adds a new boundary to the simulation.

Performing Surface Boolean Operations

To combine intersecting surfaces, use boolean operations.

Surface Boolean operations can be performed between two or more closed, manifold surface type regions that intersect in order to create a single region representing the modified surface based on the Boolean type. The new surface is also closed and manifold and can be used as part of the meshing process or as the starting point for additional Boolean operations.

This feature is useful when several separate components have been imported but need to be joined together without any loss of definition occurring, as may happen if the surface wrapper were used to wrap the surfaces into one. Note however that the feature should not be used as a substitute for CAD Boolean operations, which should always be done in preference to doing them once the surface has been tessellated.

Three types of surface Boolean operations are allowed:

- unite;
- intersect; and
- subtract

Each input region can contain one or more boundaries but depending on the Boolean type selected and the shape of the geometries, existing boundaries may be deleted in the resultant surface. Feature curves will also be created automatically to represent any new features that are generated as part of the process. Any pre-existing feature curves that belonged entirely to surfaces that are removed are also removed.

The quality of the input surfaces is not critical as long as the individual regions are closed and manifold. However, non-manifold baffles are permitted, provided the baffle boundary has been converted to an interface. The tessellated surfaces generated during the import of neutral and native CAD data files are adequate starting points in this

respect. The new boundary triangles that are added as part of the process are usually of relatively poor quality, so it is recommended that the surface remesher be used to improve the overall quality before performing the volume meshing step.

Note that Boolean surface operations are only possible when a single surface mesh representation is present. If additional representations exist then they must first be deleted prior to carrying out the operation.

What are Surface Boolean Operations?

A surface Boolean operation allows individual closed surface volumes that are not connected but intersect one another in some fashion to be joined into one closed volume surface.

Three different types of surface Boolean operations are allowed:

- unite;
- intersect; and
- subtract

All three Boolean types work in a similar manner and produce a closed, manifold, single region surface as the output.

If the input surfaces are not intersecting (for example, as in a box within a box) then a Boolean operation should not be performed and an alternative approach is required to setup up the regions for meshing.

In situations where two closed surface volumes don't intersect but have co-planar boundaries then the surface Boolean operations should not be used. If one or more of the surfaces are not closed then the boundary merge operation should be used instead, which allows unclosed surfaces to be intersected with one another.

Boolean Unite

This operation will join together two or more intersecting closed volume surfaces into one closed, manifold surface.

The Boolean unite operation is illustrated below:



Each surface to be united should belong to a different region. After the Boolean unite, the resulting surface will be stored in the first region selected for the operation. Boundaries and feature curves not required

in the output surface are automatically deleted and new curves are included for new features that are created during the process.

Boolean Intersect

This operation will extract the common volume that is shared between two or more closed volume surfaces.

The Boolean intersect operation is illustrated below:



Each surface to be intersected should belong to a different region. After the Boolean intersect, the resulting surface will be stored in the first region selected for the operation. Boundaries and feature curves not required in the output surface are automatically deleted and new curves are included for new features that are created during the process.

Boolean Subtract

The Boolean subtract operation will remove one or more intersecting closed volume surfaces from another closed volume surface.

For the simple situation where two bodies intersect one another, then either body 1 can be subtracted from body 2 or vice versa, which can lead to two different results. The choice of which body to subtract from is provided during the process. An illustration of both situations is provided below:



In the above example, the second body (region 2) is subtracted from the first body (region 1), leaving two large volumes in region 1. In the example below, the first body (region 1) is subtracted from the second body (region 2), resulting in a smaller set of volumes belonging to region 2.



Each surface to be subtracted should belong to a different region. Boundaries and feature curves not required in the output surface are automatically deleted and new curves are included for new features that are created during the process.

Applying a Surface Boolean

Surface Boolean operations are possible once certain prerequisites are met.

Surface Boolean operations can be carried out once a surface mesh has been imported and at least two regions exist, each containing a closed, manifold surface definition. Regions can be manipulated so that the correct arrangement can be obtained. The surfaces also need to intersect in some way so if the regions are not correctly positioned then they can be transformed so that they are properly aligned.

If any of the surfaces are not closed then the Boolean process will not proceed and the offending region or regions will be reported. You should check and repair the surface before attempting the process again.

Important: Since the Boolean operation is not reversible then its recommended that you save the simulation before proceeding with the steps outlined below.

The three different Boolean types will be demonstrated here on a simple box-cylinder geometry intersection, shown below:



In the above example, two separate bodies were created in the STAR-Design program and exported using the Parasolid Transmit format. They were then imported into STAR-CCM+ using the default organizational and tessellation options, which resulted in two separate regions, each with one boundary.

In the simulation tree, click on the first region, press and hold the CTRL button and then select the second and any additional regions to be included in the operation. Right-click on any of the selected regions and select the Boolean option. In the menu sub-options, three different selections will be listed:

- unite regions;
- intersect regions; or
- subtract...

Select the option that corresponds with the operation you wish to perform. These different selections are shown below:



Once the Boolean operation is done, the simulation tree will be updated with the new region and boundary definition. A successful Boolean operation on intersecting regions will always result in a single region definition, with additional new feature curves being created as needed. Boundaries with common names in the original regions will typically be re-named with an ascending numerical suffix:



If the original regions do not intersect, then the following behavior will be observed and reported back in the simulation tree:

- unite - combines the regions into the first one selected
- intersect - deletes the regions and boundaries selected
- subtract - leaves the region selected to subtract from

Performing a Boolean Unite

This operation lets you unite surfaces to combine regions.

Right-click on the selected regions you want to perform the Boolean unite operation on and select **Boolean** > **Unite Regions** from the pop-up menu. This is shown below:



Based upon the starting surface geometry, the Boolean operation will proceed to unite the surfaces together and combine the regions into the first one selected. The surface representation will automatically be updated with the new surface and displayed in the current scene:



In the above example, and more clearly in the picture below, the intersecting faces of the square block have been re-tessellated to join the cylinder at the intersection. The portion of the cylinder inside the block has also been removed.



Feature curves with the node name **Edges from Boolean unite** will have been included in the geometry at the intersection of the block and cylinder.

The following feedback from the entire process will be seen in the *Output* window:



Performing a Boolean Intersect

This operation lets you intersect surfaces to combine regions.

Right-click on the selected regions you want to perform the Boolean operation on and select **Boolean > Intersect Regions** from the pop-up menu. This is shown below:



Based upon the starting surface geometry, the Boolean operation will proceed to intersect the surfaces together and combine the regions into the first one selected. The surface representation will automatically be updated with the new surface and displayed in the current scene:



As shown above, the intersection of a block and a tube will result in the interior tube surface being retained, with the ends being capped off to form a closed volume. Feature curves with the node name **Edges from Boolean intersect** will have automatically been included for the ends of the cylinder.

Feedback from the process will be seen in the *Output* window, as shown below:



Performing a Boolean Subtract

This operation lets you subtract one region from another.

Right-click on the selected regions you want to perform the Boolean operation on and select **Boolean > Subtract...** from the pop-up menu. This is shown below:



The *Subtract Regions* dialog will appear as shown below:



Select the appropriate region from the pull-down menu for the **Select Source Region**. This will be the region from which the other regions will be subtracted. Press **OK** to continue.

Based upon the starting surface geometry, the Boolean operation will proceed to subtract the appropriate region or regions from the source region and save the result. The surface representation will automatically be updated with the new surface and displayed in the current scene:



As shown above, the subtraction of the cylinder from the block region will result in a closed block surface with a hole going through it. Feature curves with the node name **Edges from Boolean subtract** will have been included in the geometry at the intersection of the block and cylinder.

The following feedback from the entire process will be seen in the *Output* window:



If **Body 2** had been selected as the source region then the result would be different, as shown here:



In this instance, the subtraction of the block from the cylinder has resulted in two closed cylinder surfaces, both of which belong to the same region.

Working with Interfaces

Interfaces are used in STAR-CCM+ to connect regions in a simulation.

Mass, momentum, energy, or any other simulation quantity cannot pass from one region to another unless an interface is in place between them. For example, in the simulation below, a contact interface is present between the solid and fluid regions so that the heat flux can pass between them.



There are two categories of interface:

- between boundary faces
- between region volumes.

Interfaces between boundary faces include contact interfaces between solid and fluid regions, internal interfaces between fluid regions, and baffle interfaces representing thin-walled separators. Interfaces that are formed between region volumes are the heat exchanger and overset mesh interfaces. A complete list of interface types is given below.

Periodic and repeating topologies can be applied to interfaces formed between boundaries. Use the interface topology property to obtain this kind of geometric transformation. Fully-developed flow and heat transfer models can be created from a fully-developed interface with a periodic topology.

Interfaces between boundaries are formed from either direct matching of faces or indirect coupling. When interfaces that use direct matching are initialized, the solver intersects faces on one side of the interface with faces on the other side. This intersection is not required if a conformal mesh is present across the interface. A conformal mesh is where a one-to-one match exists between faces and vertices on either side of the interface. This interface can be generated during the meshing process if you make appropriate selections in the meshing pipeline.

Interfaces with direct matching have a one-to-one connectivity between the intersected faces on the two associated boundaries. Hence, the interfaces with direct matching are said to have imprinted connectivity.

Indirect coupling is available for the mapped contact and mixing plane interface types. The mapped contact interface is useful for conjugate heat transfer simulations where the mesh in the solid region may be coarser than the mesh in the fluid region. Mixing plane interfaces are useful for turbomachinery calculations in steady state. Mapped contact interfaces are said to have mapped connectivity, while mixing plane interfaces are said to have connect average connectivity.

Any interfaces that are created during the simulation setup process are accessible via the **Interfaces** manager node. The **Interfaces** node has properties and a pop-up menu.



Pop-Up Menu of the Interfaces Manager

The following pop-up menu items available when you right-click the **Interfaces** manager node.

Initialize	Populates the interface boundary nodes with faces from the original boundaries.
Reset All	Removes initialization from all child interfaces. Any vertices that are removed are reported in the output window for each region.
New Group	Facilitates the grouping of objects by creating a group node within the manager node. Interface nodes can then be dragged and dropped into the new subfolder.
Group By	Facilitates the grouping of objects by moving interfaces with common properties into a property category selected in the submenu. If you then change the property for a given interface node then it does not automatically move to the appropriate subfolder.
Ungroup	Reverses the grouping process.

Interface Manager Properties

Use the properties below to adjust the specifications of the **Interface** manager.

Verbosity	Controls the level of detail output at the interface initialization. When the option is enabled, more detail is printed to the output window.
Interfaces	The number of interfaces that are created in the simulation (Read Only)

Interface Types

The interface type determines the behavior of the interface with respect to the modeling of flow and heat transfer. Interfaces can connect individual boundaries (faces) or entire regions (volumes). You can create the following interface types:

Type	Connects:	Topology	Connectivity
Baffle	Boundaries	Direct Interface	Imprinted
Porous Baffle	Boundaries	Direct Interface	Imprinted
Contact	Boundaries	Direct Interface	Imprinted
Fan	Boundaries	Direct Interface	Imprinted
Fluid Film	Boundaries	Direct Interface	Imprinted
Fully Developed	Boundaries	Direct Interface	Imprinted
Internal	Boundaries	Direct Interface	Imprinted
Mapped Contact	Boundaries	Indirect Interface	Mapped
Mixing Plane	Boundaries	Indirect Interface	Connect Average
Heat Exchanger	Regions	Heat Exchanger Topology	Modeled
Overset Mesh	Regions	Overset Mesh Topology	Modeled

Working with Interfaces between Regions

Interfaces are created between region volumes to set up an overset mesh or to model heat transfer in a dual stream heat exchanger.

To create an interface between regions:

1. Select the two regions that define the interface (press and hold the Ctrl key to select the second node)
2. Right-click and select the **Create Interface** pop-up menu item, and make a selection from one of the options in the submenu.



The appropriate heat exchanger or overset mesh interface node is added to the tree.

Working with Interfaces between Boundaries

Various types are available for direct and indirect interfaces.

There are two categories for interfaces between boundaries:

- Direct Interface - boundary faces are intersected or imprinted (to obtain a conformal mesh)
- Indirect Interface - boundary faces are coupled using mapping or mixing

The following interface types are available for Direct Interfaces:

- Baffle
- Porous-baffle
- Contact
- Fan
- Fluid Film
- Fully developed
- Internal

The following interface types are available for Indirect Interfaces:

- Mapped Contact
- Mixing Plane Interface

Creating a Boundary Interface

An interface can be created between boundaries to produce a conformal mesh interface (as part of the meshing pipeline), or to maintain physics continuity between regions or in the same region.

Do not confuse the act of creating an interface with converting boundaries to an interface, which is a similar but different operation used primarily for meshing purposes. In this latter instance, a single (usually common) boundary belonging to a surface mesh is duplicated and then an interface made between the original and the copy for the primary purpose of defining multiple regions for meshing.

To create an interface between boundaries, select the two boundaries that define the interface. The first boundary selected is termed the master boundary; the second is termed the slave boundary.

Once the boundaries have been selected (press and hold the Ctrl key to select the second node), right-click and select **Create Interface**.



This action creates an interface between the two selected boundaries. Two new nodes appear in the simulation tree (under the **Boundaries** node for each region) indicating which boundary nodes are linked:



An interface node also appears in the **Interfaces** manager node, which has a default name, **Interface**, plus an identifying suffix. Selecting the interface node highlights the interface in the Graphics window.



To check the number of faces associated with each interface sub-node, open the subnodes of the Representations node and select the corresponding interface node:



The number of *Faces*, as shown in the *Properties* window, is zero due to the face that the intersection process has not yet been performed. The nodes appear grayed out, indicating that they are not yet active.

See Also

What Is a Generalized Cylinder?

Intersecting an Interface

The intersection process occurs the first time the solution is initialized for the volume mesh.

This process involves the appropriate mesh calculations required to define the interface fully. For large problems with complex physics, it is generally advantageous (especially with regard to memory usage) to check the interface intersection after activating the Space model but before activating other models.

Intersection populates the interface boundary nodes with faces from the original boundaries. The faces belonging to the original boundary are removed. For a perfect intersection, where the matching areas of the original boundaries are the same, then the boundaries are left with zero faces. The boundary nodes however are not removed as the intersection process is reversible.

If the intersection process is not perfect or if the overlapping areas of the original boundaries are different, then there are some faces remaining in the original boundary nodes. Depending on the situation, examine these left over faces. If they are acceptable, then assign the appropriate boundary types and conditions to them. If they are not acceptable then the intersecting tolerance can be changed and the interface reset before trying the intersecting process again.

Once the interfaces have been created or reset, the intersection process can be performed in one of two ways:

- by manually initializing the solution
- by using the **Interfaces** manager node

Right-clicking on the **Interfaces** node give access to the **Initialize** option for all interfaces at the same time:



Alternatively, each interface can be initialized separately by selecting the interface node, right-clicking and selecting **Initialize**. For example:



In the case of repeating interfaces used for sliding mesh calculations, the reintersection is done dynamically.

The results of the intersection process are displayed in the *Output* window along with messages related to the other initialization processes:

```
Output - Star 1

Updating interface Interface1 between Porous:Interface1 and Inlet:Copy of Interface1 tolerance 0.02
  Created 634 intersection faces (conformal match)

Updating interface Interface2 between Outlet:Interface2 and Porous:Copy of Interface2 tolerance 0.02
  Created 973 intersection faces (conformal match)

 Time to intersect          (s): 0.05
```

If the match between the two boundaries is conformal (that is, the master and slave boundary faces are all aligned and no additional intersections are required to produce a one-to-one connection) then this is indicated as such in the output, as shown above. If intersection of faces is required then the % difference in face area between the master and slave boundaries are shown instead. If the boundary perimeters are similar then only a small difference in intersection area is expected.

Additional information can be displayed in the output window by changing the interface verbosity.

On selecting the nodes of the interface boundaries within the **Representations** node, there is a non-zero value for the number of faces in the *Properties* window:



The interface boundary node is active.

The intersection process increases the number of faces that exist at the boundaries making up the interface. This is due to the process of subdividing the boundaries on each side so that a conformal, one-to-one connection is formed between each face making up the interface. An example before and after illustration showing the interface intersections on a periodic interface is shown below:



The intersection process may also introduce additional vertices into the model. This is not permanent and is undone when either the interface is reset or deleted. If the mesh is exported then the additional face divisions and vertices are not included. That is, the mesh is exported as if the interfaces did not exist.

If the simulation has been set up using the parallel server option, then the intersection process is done in parallel, resulting in a faster match being performed. Do not use this option if you have several interface pairs to intersect.

Adjusting the Interface Intersection Tolerance

The intersection tolerance limits the distance vertices on the slave boundary can be moved to match vertices, edges and faces on the master boundary for all direct type interfaces.

The tolerance value is specified as a fraction (from 0 to 1) of the length of the smallest edge attached to the vertex. The default value is 0.05 (which is 5% of the smallest attached edge length).

For each vertex on the slave boundary, the tolerance value is used to define the radius of a cylindrical volume within which the vertex moves during intersection. This cylindrical volume is centered on a projection vector that has the vertex as its origin, and whose direction is computed from the normals of adjacent faces. The length of the cylinder is ten times its radius.

The *Tolerance* value can be accessed in the *Properties* window for the **Physical Values > Intersection** node for the interface, available once a physics continuum has been created:



The example below shows a case where the interface creation was not 100% successful, as it left unmatched faces (shown in gray) in the interface:



In this case increasing the tolerance value corrected the fault and allowed all of the faces to be matched:



Before changing the tolerance value, reset the interface first. You can then attempt the interface intersection again.

Inserting a Shell Region at an In-Place Interface

In order to account for conduction across the width and height of a baffle, you can replace the baffle interface with a shell region.

When you use a solid baffle to model a thin solid plate, the flow goes around the baffle. Heat can conduct through the baffle, experiencing some heat resistance, but heat cannot conduct along the baffle. If you expect a large temperature gradient along the baffle, then neglecting heat conduction along the baffle can result in a significant modeling error. The baffle can also have a large heat capacity, which is important in some unsteady cases.

The replacement shell region interfaces to the solution domain on either side of the original baffle. A baffle interface models heat conduction through the thickness of the baffle only; a shell region models heat conduction in all directions.

A solid shell region is associated with a shell continuum. The models that are activated for that continuum define the behavior of the solid shell. The same physics continuum defines the material models for the shell.

The Insert Shell Region Interface option allows you to select internal interfaces and insert a shell region between the two original boundaries. Using this option is much quicker and easier than creating the shell region and interfaces manually. This option is available only for in-place interfaces.

To insert a shell region at an in-place interface:

1. Right-click the interface and choose **Insert Shell Region Interface**.



The shell region is created and the interface is replaced with two new contact interfaces that connect the shell region to the parent boundaries on either side.



Examining Interfaces

Interface creation normally proceeds without any problems if the underlying mesh is valid and the interfaces have been defined properly.

The most common problems associated with interfaces are:

- cells with negative volume; and
- unsuccessful intersection of interfaces

Cells with Negative Volume

One cell with a negative volume can keep the simulation from iterating. While negative-volume cells can have a variety of causes, an interface might be the cause if it was created from two incorrect or misplaced boundaries.

To detect this before running the simulation, obtain a report of the negative cell volumes after creating the interface, even if you generated such a report earlier.

Unsuccessful Intersection of Interfaces

If you attempt to solve a simulation with few or no intersections in an interface, errors are likely. Therefore, when setting up simulations with interfaces, check the interface intersection as early as possible for problems such as:

- Too few or zero intersections in an interface caused by specifying the axis of rotational periodics incorrectly.
- A boundary may unexpectedly be left without any faces after the intersection process; therefore check the boundary icons.
- The absence of any intersections in an interface usually results from the chosen boundaries not being geometrically beside each other. For details on detecting and solving , read about examining the intersection report.

To test for gaps or holes in the mesh, display the interface boundaries.

The various methods of checking the successful intersection of an interface include:

- checking the boundary icons
- examining the intersection report
- displaying the interface boundaries
- displaying remainder boundaries.

Examining an interface is important as it determines the integrity of the connection between or across the mesh from one boundary to another. Incomplete interfaces result in a loss of continuum information being transferred, resulting in a less accurate solution.

The typical causes of the intersection process failing are:

- selecting the wrong boundary pairs
- forgetting to set the rotational axis direction or setting the wrong direction for axisymmetric cases
- incorrect tolerance setting
- the boundaries are not coincident in space nor aligned correctly.

Checking the Boundary Icons

After a successful intersection process, the interface boundary icons under the **Representation** node become active.

An active interface boundary icon has an appropriate color applied to it. If an interface boundary contains zero faces after intersection then the icon under the **Representations** node remains grayed out, indicating that the intersection process was unsuccessful. The number of *Faces* in the *Properties* window also remains at zero, as shown below:



Additionally, check the intersection report for error messages and attempt to display the interface boundaries to further confirm that the interface intersection was not done.

Note: The boundary and interface icons under the original **Regions** node are not grayed out regardless of whether they contain faces or not. The intersection process affects the **Representations** node icons.

Examining the Intersection Report

The validity of the interface is only checked during the intersection process and not during the original selection of the boundaries making up the interface.

During the validity check, if no intersections can be found for the two interface boundaries then a warning message is displayed in the output window during the initialization process:



The likely cause in this instance is that you either selected the wrong boundary pairs or did not set the axis of rotation for an axisymmetric case. It's possible that the boundaries are not coincident in space within the specified tolerance. Changing the verbosity of the interface output may provide additional clues as to what is going on.

It is recommended that you delete or reset the interface and attempt to recreate them once you have diagnosed the original cause of the problem.

Displaying the Interface Boundaries

The quickest and easiest way to check if an interface was intersected is to display the interface faces as part of a scene.

There are three ways to show an interface in an existing scene:

- highlight the interface by selecting the appropriate boundary node in the simulation tree;
- right-click on the interface node and select highlight; and
- isolate the interface parts within the scene

Highlighting the interface node is very easy and gives a quick indication as to the state of the node. For example, the intersection process, the

interface nodes under **Regions** have an active icon but selecting it does not highlight anything in the current scene. This is because no faces have been associated with the interface. Once a successful intersection process has been carried out, selecting the icon now highlights the interface part on the display.

An example interface node under **Regions** is shown below:



Selecting the interface node highlights the part on the display, as shown below:



If after selecting the node, no highlighted faces are seen then it means that the interface creation process was not successful.

Similar to the above, right-clicking on the interface node itself and selecting the **Highlight** option produces the same result. For example:



Alternatively, the interface part can be isolated on the display to give a more permanent image for inspection. To do this, right-click on the **Parts** node for the current **Geometry** displayer and select **Edit**in the pop-up menu:



The *Edit* dialog appears:



Using this dialog, isolate the interface boundaries (not the original boundaries - the interface ones have the topology type as part of the name by default) on the right-hand side and press **Apply**. The interface boundaries are shown on the display:



The faces can now be inspected for any unexpected holes or gaps. The color of the interface can also be used to judge if the correct interface type has been selected. The following list indicates the color used by each type:

Interface Type	Color
Baffle	Brown
Contact	Green
Fan	Light brown
Fully developed	Purple
Internal	Yellow
Mixing plane	Blue
Porous-baffle	Orange

Displaying Remainder Boundaries

If the original boundaries from which the interface is created do not completely match, the intersection process leaves some boundary faces in the original boundaries. These are called remainder boundaries.

There are two instances when remainder boundaries occur:

- The two original boundaries have different perimeters; and
- The two original boundaries have the same perimeters, but are discretized to different resolutions.

In cases where the boundaries have different perimeters (typically when one has a larger area than the other) then a portion of original boundaries remain after the intersection process has taken place. An example would be when one cylinder of a smaller radius is joined to another cylinder with a larger radius:



In this instance, the boundaries share the same common area in the middle but the larger cylinder has remainder boundaries, as highlighted in blue. This is further demonstrated below in an exploded view of the two regions:



The remainder boundary faces remain a part of the original boundary node. It is important that the appropriate boundary type be set for them to ensure the correct modeling characteristics.

Where the boundary perimeters are identical but have different mesh resolutions (as would happen when joining a tetrahedral type mesh to a polyhedral type one for example), then small portions of the original boundaries remain, as highlighted by the blue faces below:



In this instance, specifying the correct boundary type for the remainder boundaries is less important as their impact is usually marginal.

To visualize the boundary remainder faces, follow the same procedure as used for displaying interface boundaries, but instead of displaying the interface boundaries, select one or both of the original boundaries.

See Also

Extrusion Parameters Properties

Specifying the Transformation and Axis Properties for Periodic Interfaces

If the periodic interface type is selected for *Topology*, then specify the periodicity and axis of rotation to allow the meshing tools to generate a conformal mesh interface for the periodic boundaries. This information is also required for the intersection process running the analysis.

The **Periodic Transformation** node can be used to specify the periodicity and axis of rotation:



The *Periodicity* option refers to the relative location to one another of the two boundaries making up the interface. By default, this is set to **Rotational + Translational**, indicating that both a rotational plus translational component is expected to "match" the slave boundary to the master. A rotational only or translational only option is also available, as shown below:



Select the appropriate periodicity type for the interfaced. Note: The rotational and/or translational components are calculated automatically.

Specify the axis of rotation of the periodic interface. Two options are available, namely using the reference axis for the region that contains the periodic boundaries or specifying a predefined axis for just that interface.



By default, the *Axis of rotation specification* is set to **Use region's reference axis**, meaning that the axis specified in the **Region node > Physics Values > Axis node** are used. The properties of this node are shown below. When the periodic interface is applied to boundaries from two separate regions, the region axis from the master boundary (Boundary-0) is used.



To access the **Physics Values** node, select a space model for the physics continuum that is linked to the region.

If the**Specify axis** option is selected instead then a new node called **Axis** appears under the **Periodic Transformation** node allowing the specification of an alternate coordinate system. This is shown below:



In both the above instances, enter the axis of rotation (as specified by the *Direction* components) plus the *Origin* coordinates. Both relate to the translational and/or rotational location of the interface boundaries.

Once the periodic transformation properties have been supplied, the surface remesher and volume meshers can be run. These assess the information entered for the periodic interface and determine whether a conformal mesh can be produced or not. This is indicated at the start of the output log from each mesher. If successful, the calculated rotation angle and/or translational distance is displayed in the *Properties* panel:



These values are automatically “locked” for the interface node to prevent recalculation during further meshing operations.To change the values, untick the *Locked* option. The angle and translation is recalculated during meshing.

Failure to properly specify the correct periodic properties causes the conformal mesh calculation to fail as shown below:

```
Output - Star 1
Initializing meshing pipeline...
  No initialization required.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1
Creating conformal mesh on periodic interface:
   Warning: Unable to make conformal mesh for interface Periodic 1
*Warning* - 1 out of 1 periodic interfaces are not conformal.
Surface Mesh is closed.
Surface Mesh is manifold.
----------------------------------------
Executing Surface Remesher
----------------------------------------
```

The meshing process still continues but ultimately leads to a non-conformal mesh being produced during intersection. If a conformal volume mesh is critical then review the periodic transformation property values for errors carefully. It is also recommended that a conformal surface mesh is produced first to guarantee that the volume mesh is also conformal for the periodic interface. That is, run the surface remesher first with the interface and properties defined fully before running the volume mesher.

Converting Boundaries to Interfaces

Converting boundaries to interfaces can solve two problems: insufficient demarcation of regions and surfaces that lack required thickness.

During mesh generation in STAR-CCM+, the original imported geometry may be lacking sufficient boundaries for the regions to be properly demarcated. In particular, it is a frequent occurrence that a shared boundary is used for inter-region surfaces whereas the requirements for multi-region meshing impose that two separate boundaries must exist in order for the regions to be properly defined.

Another situation is the existence of baffle surfaces within the geometry that are treated as infinitely thin surfaces that impede the flow and/or transfer heat along their length. These are typically imported as interior boundaries but is normally flagged as being non-manifold if used as part of the meshing process with STAR-CCM+. To meet the topology requirements of the meshing tools they have to be duplicated so that in a topological sense they have thickness.

Converting a boundary to an interface allows the requirements for the above situations to be met.

When a boundary is converted to an interface, two separate actions occur:

1. a copy of the original boundary or boundaries (with the same surface and perimeter triangulation) is made; and
2. an in-place interface is automatically created linking the original boundary or boundaries with the copy or copies

The copy of the original boundary can then be used as part of the definition of a new region, for example, by using the create new region from boundaries option. This process also automatically moves the interface node copy into the appropriate region. In the case of baffle surfaces, the perimeters of the two boundaries are merged so as to form an infinitely thin closed volume. If the original baffle boundary was next to an exterior surface then it is important that the baffle boundary be merged with the exterior surface converting it to an interface. The surface repair tools can be used for this purpose.

Note that converting a boundary to an interface is different from simply creating an interface between two preexisting boundaries, which corresponds with the second step described above.

Converting to an Interface

Before proceeding with converting a boundary to an interface, it is recommended that you combine any boundaries that belong to the same interregion surface into one. This simplifies the setup process thereafter by reducing the number of interfaces created.

An illustration showing the before and after boundary setup when a boundary has been converted to an interface is given below:



To convert to an interface, select the boundary or boundaries to be converted and then right-click. In the popup menu that appears, select **Convert to Interface(s)**.



This begins the interface conversion process. Once this is complete, the simulation tree is updated with the new boundary and interface information.



As can be seen, a copy of the original **Surface 1** boundary has been made, in addition to a pair of interface nodes linking the two boundaries. The boundary copy can now be used as part of a new region definition or simply left in the region to act as a baffle.

A new interface node with the same name as the original boundary is shown in the**Interface** node in the simulation tree. Selecting the interface node highlights the interface in the graphics window.



Examine the *Properties* window for the new interface node, for the default *Topology* and *Type* settings:



For normal meshing purposes, do not change these settings. For baffle surfaces change the *Type* property to **Baffle** or **Porous Baffle**, depending on the required flow properties.

Accessing an Interface from an Interface Boundary

In some simulations, such as with conjugate heat transfer, you might have a large number of boundaries and interfaces. STAR-CCM+ lets you access the properties of an interface directly from any of its interface boundaries, rather than having to scroll through the object tree.

Right-click the node of the interface boundary and select **Edit Interface** from the pop-up menu.



The table editor for the parent interface appears.



Deleting Interfaces

Deleting an Interface restores the boundaries to their state prior to the interface being generated.

When an interface is deleted, any new vertices or faces generated as part of the intersection process are removed. If you want to maintain the interface but delete the interface intersections then reset the interface instead. There is no undo operation, so save the simulation file performing this step if you are unsure of whether you need to delete the interface.

Boundary Interface Types Reference

There are different types of interface topologies that are used in STAR-CCM+.

On a broad characterization, all interfaces fall into one of the following two groups:

- Interfaces between regions
- Interfaces between boundaries

In general, any available interface type can be applied to an interface, regardless of the geometrical properties of the interface, such as whether the underlying boundaries are conformal or coincident in space. It is your responsibility to ensure that the interface type assigned to a specific interface makes sense from a physical modeling perspective.

The interface type must be compatible with the continua on either side of it, or STAR-CCM+ reports an error when the simulation is initialized. For instance, an internal interface can only be used between regions of the same continuum, and a contact interface can only be used between solid continua or a solid and a fluid continuum. Interfaces can also be used to model cyclic (periodic) repeats, conducting baffles and porous baffles.

Boundary interfaces can be one of the following:

- Direct Interface—the interface is intersected by default
- Indirect Interface—no intersection, simple association

The following types are available for direct matching interfaces with imprinted connectivity:

- Baffle
- Porous-baffle
- Contact
- Fan
- Fluid Film
- Fully developed
- Internal

There are two interface types available for interfaces with indirect coupling:

- Mapped Contact

- Mixing Plane Interface

Once an interface is created, all the compatible interface types, both direct and indirect, are available for the user to choose from.

What Is a Direct Interface?

A direct interface implies that the two boundaries composing the interface are directly joined, either permanently (for stationary cases) or temporarily for some time period (for cases with rigid body motion).

When used as part of the meshing process the two boundaries imprint on one another at the surface remeshing stage to ensure a conformal mesh across the interface where applicable. When applied to a volume mesh, the interface between the two boundaries is intersected. If the two boundaries are non-conformal, some new discretization is created. An explicit connection is created between cells on each side of the interface so that mass and/or heat transfer occurs across the intersected area.

When a direct interface is created from two boundaries, two new boundaries that are associated with that interface are automatically created as well. The faces of these new boundaries are composed of the intersection of the faces of the original boundaries. The original boundaries from which the interface is created are not destroyed; however, the surface area of each reduces by the amount it overlaps with the other. Direct interfaces can be created between boundaries with meshes that are either conformal or non-conformal and can belong to the same region or different regions depending on the type selected.



The three allowable direct type topologies are:

- in-place

- periodic
- repeating

The following types of direct interface are available:

- Baffle
- Porous-baffle
- Contact
- Fan
- Fluid Film
- Fully developed
- Internal

In-Place Interface

In an in-place interface, there is no physical separation in space between the two boundaries comprising the interface.

Furthermore, there is no requirement that the two boundaries have the same perimeter. After the imprinting (for surface remeshing), or intersection (for volume meshes used in an analysis) occurs the original boundaries maintain their boundary type and retain whatever surface area remains (if any).

The figure below shows different examples of in-place interfaces. The shaded red, green, and blue areas represent different mesh regions, such as porous/non-porous or fluid/solid combinations. The dotted black line in each case would represent the in-place interface that would be created between the coincident boundaries for each region. Where the areas of the original boundaries differ then remainder boundaries are produced.



When the in-place interface type is used with the surface remesher, an attempt is made at producing a conformal mesh on each side of the interface. In other words, if successful, the two boundaries when mapped from one to the other match identically in terms of the vertex locations and face topologies. The matching results in a higher-quality volume mesh definition for the analysis stage.

The option to produce a conformal mesh is automatically initiated when an in-place type interface has been created and the surface remesher is used afterwards. If it is not successful, then the meshing process is terminated and an error message is produced indicating the nature of the problem.

Periodic Interface

The boundaries of a periodic interface are physically separated in space but can be mapped from one to the other through some constant rotation and/or translation. The periodic interface represents a cyclic repeat of information across the boundaries, so fluxes that cross one boundary are transformed and applied to the other.

This effectively approximates a repeating geometry. For example, you can model an infinitely long pipe by using a short segment of the pipe and applying a periodic interface between the two ends.

To define a periodic interface, set the perimeter of one boundary to be coincident with the transformed perimeter of the periodic matching boundary. The figures below show examples of periodic interfaces as applied to geometries that repeat in a translational and rotational manner. The green and black dashed lines in each case represent the location of the periodic interface boundary pair. For each pair, the shape and size of the boundaries match, and they are separated by a given vector or angle. The translation vector or rotation angle are calculated automatically by STAR-CCM+, and are properties of the interface **Periodic Transformation** node.





When you use a periodic interface with the surface remesher, tetrahedral, polyhedral, or thin meshers, the mesher produces a conformal mesh on each side of the interface. In other words, if successful, the two boundaries when mapped from one to the other

match identically in terms of the vertex locations and face topologies. This results in a higher-quality mesh definition for the analysis stage.

The option to produce a conformal mesh is automatically initiated when a periodic type interface has been created and one of the above meshing tools is used thereafter. Attempts to produce a conformal mesh are made both during the surface remeshing stage and during volume mesh generation. This means that if a prepared surface mesh is imported from another source and used directly for volume mesh generation (without first resurfacing) then a conformal periodic boundary pair is still achievable within the bounds of the scheme.

The conformal mesh mapping is limited to rotational, translational, or mixed (roto-translational) cases where a boundary pair (within the in-built tolerance) is possible. If a successful mapping is found, then the transformation matrix, number of vertex pairs and face mappings is displayed in the *Output* window:

```
Output - Star 1
Initializing meshing pipeline...
  No initialization required.
Executing surface meshers...
Preparing surface for surface meshing in regions: Body 1
Computing transformation between boundaries Side 1 [Periodic] and Side 2 [Periodic]...
Transformation found:
< <  1.00000e+00  0.00000e+00  0.00000e+00  0.00000e+00 >
  <  0.00000e+00  8.66026e-01  5.00000e-01  0.00000e+00 >
  <  0.00000e+00 -5.00000e-01  8.66025e-01  0.00000e+00 >
  <  0.00000e+00  0.00000e+00  0.00000e+00  1.00000e+00 > >
Creating periodic vertex and face pairs
Periodic boundaries are not topologically identical 1
Making conformal mesh between boundaries
253 vertices matched up
421 faces matched up
```

If a transformation matrix cannot be extracted or mesh conformity is not satisfied between the periodic boundary pair, then a warning message is printed in the *Output* window. For example:

```
Output - Star 1
Initializing meshing pipeline...
  No initialization required.
Executing surface meshers...
Preparing surface for surface meshing in regions: Body 1
Computing transformation between boundaries Boundary 1 [Periodic 1] and Side 1 [Periodic 1]...
WARNING: A valid transformation could not be extracted.
Surface Mesh is closed.
Surface Mesh is manifold.
Projecting to CAD geometry
```

In this instance, the meshing scheme still produces a mesh but reverts to being non-conformal, exactly as if the periodic boundary pair did not exist. The absence of a periodic match can also be confirmed by looking at the number of faces belonging to each boundary in the periodic

pairing using the **Representation** node for the mesh. If the number of faces is different, then the mapping was unsuccessful.

If the automatic surface repair option for the surface remesher has been activated, then periodic boundaries are excluded from this operation. Take care when modifying the surface mesh by hand once the surface remesher has been run as changes to boundaries belonging to the periodic pairing are not preserved when it comes to the volume meshing step of the process.

Repeating Interface

A repeating interface is a combination of in-place and periodic interfaces and is associated with periodic, sliding mesh calculations.

The two boundaries have the potential to be matched completely to form an internal interface, but the mesh motion results in one underlying boundary being indexed translationally and/or rotationally relative to the other. The parts that overlap are treated as in-place, with a periodic interface being created from the remaining boundary areas.

If the boundaries on either side have equal areas, select a repeating interface. If the areas are not equal, use a mixing plane interface instead.

The figure below shows an example of a repeating interface used to simulate rotor-stator interaction. Periodic interfaces (not labeled) cause infinite repeats of each blade row. The repeating interface between the blade row is shown in two sections: an in-place part and a periodic part. If the meshes are aligned, there is no boundary face area in the periodic part. As the mesh is moved dynamically, the mesh intersections are recomputed. More of the face area moves from the in-place part to the periodic part, until “snap-back” occurs such that the mesh returns to its starting position.



When the periodic interface type is used with the surface remesher, tetrahedral, polyhedral and/or prism mesher, then activate the per-region meshing option to mesh each region separately. A conformal mesh interface is not required in this instance since the mesh in each region is moving past one another. Alternatively, a different mesh continuum could also be used for each region.

Baffle Interface

A baffle interface physically represents one or more thin sheets of impermeable, conducting materials in a fluid.

It can be placed between regions of the same continuum or of different continua. It has properties, conditions, values, and a pop-up menu.



Modeling

- The baffle has infinitesimal thickness.
- The fluid next to a baffle treats the baffle as a wall. It applies the non-slip boundary condition in viscous flow, and the component of the velocity normal to the baffle is zero.
- A thermal resistance per unit area R for conjugate heat transfer can be specified, either as effective resistance or with multiple layers.

  The relation between the heat flux across the contact interface, q, and the temperature difference, ΔT, is given by Eqn. (4). To calculate the thermal resistance of the baffle, the following formula can be used:

  where k_i and Δx_i are the conductivity and thickness, respectively, of each of the n layers comprising the baffle.
- If it is not required that heat transfer through the baffle is modeled, do not create a baffle interface. The same functionality could be achieved by assigning wall boundary types (with adiabatic thermal conditions) to the boundaries in question.

Baffle Properties

You can use the properties that are listed below to adjust the specifications of a **Baffle** interface node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)	
Boundary-1	The location of the boundary on one side of the interface (Read Only)	

Type	Defines the type of interface. Must be set to `Baffle Interface`.	
Topology	Defines the connection type between the interfaces	
	In-place	Uses the in-place topology.
	Periodic	Uses the periodic topology.
	Repeating	Uses the repeating topology.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Imprinted** connectivity is possible.	

Baffle Conditions

You can use the properties that are listed below to adjust the specifications of a **Physics Condition** interface node.

Baffle Thermal Option	Specifies whether the baffle conducts energy.	
	Conducting	Thermal conduction occurs through the baffle. You would typically select this option when you use an energy model and a baffle or porous baffle interface. A **Thermal Resistance** interface value is added.
	Non-conducting	No thermal conduction occurs through the baffle.
Energy Source Option	Provides energy source options for the interface. It is activated when the **Baffle Thermal Option** is set to *Conducting*.	
	None	Do not specify an energy source.

	Heat Flux	Specifies a user-defined heat flux in W/m2.
	Heat Source	Specifies a user-defined heat source in W.

Baffle Values

You can use the properties that are listed below to adjust the specifications of a **Physics Values** interface node.

Thermal Resistance	The value of the resistance to conduction through the baffle. It is activated when the **Baffle Thermal Option** is set to **Conducting**. A **Thermal Resistance** boundary value is added to the object tree when the Thermal Specification boundary condition is set to Environment.	
	Dimensions	Specify the resistance in units of $m^2 \cdot K / W$ or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, **User Code**, or Multi-Layer Resistance.
Heat Flux	Specify the heat transferred per unit area. It is activated when the **Energy Source Option** is set to **Heat Flux**.	
	Dimensions	Specify the heat flux in units of W / m^2 or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Heat Source	Specify the heat transferred. It is activated when the **Energy Source Option** is set to **Heat Source**.	

	Dimensions	Specify the heat source in units of *W* or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Intersection	Added for all direct interfaces to control the tolerance. The tolerance is a fraction of the smallest cell face on the secondary boundary that the secondary boundary vertices are allowed to move during an intersection. If the intersection fails, the following message is displayed in the output window - `Warning: no intersections found`. The solution cannot proceed under these conditions.	

Baffle Values: Multi-Layer Resistance

This **Thermal Resistance**value is useful when, for example, you are simulating conduction through multiple materials that vary in conductivity. It is calculated using Eqn. (2).

This object is also available in the Thermal Resistance boundary value, which is added to the object tree when the Thermal Specification boundary condition is set to Environment.

Multi-Layer Resistance Properties

Number of layers	Specifies the number of layers within the baffle.	
Thickness of each layer	Sets a thickness for each of the layers. The number of comma-separated entries must match the number of layers. If the number of layers is decreased, the last value is dropped. If a layer is added, a zero value is appended by default.	
Conductivity of each layer	Specifies the conductivity through a series of comma-separated values equal to the number of layers. STAR-CCM+ works with discrepancies between these specifications as follows:	

	• If you make more conductivity value entries than layers, for example seven values and five layers, then STAR-CCM+ works with the first five values. • If you make fewer conductivity value entries than layers, for example three values and five layers, STAR-CCM+ interprets values for the remaining layers as zero conductivity.

Porous Baffle Interface

Physically, a porous baffle represents a porous membrane through which fluid passes and experiences a pressure drop.

Conductive heat transfer is modeled in the impermeable portion of the baffle. Porous baffles can be used to model perforated plates, thin screens, and wire screens. It has properties and a pop-up menu.



- Use this interface type between regions belonging to the same fluid continuum.
- When debugging it is appropriate to report mass flow rates and pressures on the interface boundaries. Reporting on interfaces is more representative than creating derived parts nearby.
- For multiphase flows using the VOF model, porous baffle interfaces have a conformal (that is, matching) mesh at locations where phases of different density are expected to cross the baffle. If this requirement is not observed, convergence problems could arise.
- Remember that a baffle is a “lumped” model; when modeling a baffle, you determine the bulk effects that the baffle has on the flow,

for example, pressure drop, not the fine details. If you want to compute the flow coming out from the little openings next to the baffle, model the actual geometry, with these small openings, and a fine mesh. If you are interested only on the bulk effect of the baffle downstream, then you can appropriately use the porous baffle interface in your model.

Modeling

- The baffle has infinitesimal thickness.
- If the porous baffle is thought of as a perforated surface, then the ratio of the area of the holes to the area of the surface, is called the porosity, χ. Wire screens would have a large value of χ, while something like a perforated plate would have a smaller value.
- Each interface has a direction from side 0 to side 1. When you create the interface, select side 0 first, then side 1. For a porous-baffle interface, the upstream is side 0. Therefore if you select the sides in the wrong order, the sign of the pressure jump is reversed.
- The viscous shear stress is computed as a linear combination of the shear stress in the fluid and the shear stress at the wall. The porosity, χ, provides the linear weighting.
- For heat transfer analysis, the porous baffle offers a finite thermal resistance. Conduction is modeled as a one-dimensional process and thermal inertia is neglected. A thermal resistance per unit area, R, for heat transfer can be specified. The relation between the heat flux across the contact interface, $\dot{q}''$, and the temperature difference, ΔT, is given by Eqn. (4). Estimate the thermal resistance using Eqn. (2).
- The total heat transfer across the porous baffle is computed as a linear combination of the heat transfer due to convection and diffusion in the fluid and the conduction through the baffle. The porosity, χ, provides the linear weighting factor.
- The pressure drop, Δp, is specified as follows:

$$\Delta p = -\rho\left(a\left|v_n\right| + \beta\right)v_n \tag{3}$$

  where:

  - v_n is the superficial velocity normal to the surface.
  - a is the user-specified constant **Porous Inertial Resistance**.
  - β is the user-specified constant **Porous Viscous Resistance**.
  - ρ is the fluid density at the interface.

- It is assumed that the direction of the flow is unchanged as it passes through the baffle.

Porous Baffle Properties

Use the properties that are listed below to adjust the specifications of a **Porous Baffle** interface node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)	
Boundary-1	The location of the boundary on one side of the interface (Read Only)	
Type	Defines the type of interface. Must be set to `Porous Baffle Interface`.	
Topology	Defines the connection type between the interfaces	
	In-place	Uses the in-place topology.
	Periodic	Uses the periodic topology.
	Repeating	Uses the repeating topology.
Intersection	Added for all direct interfaces to control the tolerance. The tolerance is a fraction of the smallest cell face on the secondary boundary that the secondary boundary vertices are allowed to move during an intersection. If the intersection fails, the following message is displayed in the output window - `Warning: no intersections found`. The solution cannot proceed under these conditions.	
Porous Inertial Resistance	The value of the coefficient a in Eqn. (3), set by the user. It can be positive or 0. The default value is 0. This value is input as a scalar profile.	
Porous Viscous Resistance	The value of the coefficient β in Eqn. (3), set by the user. It can be positive or 0. The default value is 0. This value is input as a scalar profile.	

Porosity	The factor χ in Eqn. (3) that is used in calculating heat transfer and optionally viscous shear. The value ranges from 0 to 1. If the porous baffle is thought of as a perforated surface, then χ may be interpreted as the ratio of the area of the holes to the area of the surface. This value is input as a scalar profile.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Imprinted** connectivity is possible.

Contact Interface

The contact interface type is used to join either of the following to permit conjugate heat transfer between the regions:

- Two solid regions, either from the same continuum or from different continua.
- A fluid and solid region.

It has properties, conditions, values, and a pop-up menu.



Modeling

- A thermal contact resistance per unit area, R, for conjugate heat transfer can be specified. The relation between the heat flux across the contact interface, $\dot{q}''$, and the temperature difference, ΔT, is given by:

$$\dot{q}'' = \frac{\Delta T}{R} \tag{4}$$

Contact Interface Properties

You can use the properties that are listed below to adjust the specifications of a **Contact Interface** node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)	
Boundary-1	The location of the boundary on one side of the interface (Read Only)	
Type	Defines the type of interface. Must be set to `Contact Interface`	
Topology	Defines the connection type between the interfaces	
	In-place	Uses the in-place topology.
	Periodic	Uses the periodic topology.
	Repeating	Uses the repeating topology.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Imprinted** connectivity is possible.	

Contact Interface Conditions

You can use the properties that are listed below to adjust the specifications of a **Physics Condition** interface node.

Energy Source Option	Provides energy source options for the interface.	
	None	Do not specify an energy source.
	Heat Flux	Specifies a user-defined heat flux in W/m^2.
	Heat Source	Specifies a user-defined heat source in W.
External Load	Specifies the method that is used to define the load. Only available for fluid-solid contact interfaces when the Solid Stress model is activated.	

	Force	This option adds a Force node to the **Values** node, which must be entered as a vector profile.
	Pressure	This option adds a Solid Stress Pressure node to the **Values** node, which must be entered as a scalar profile.
	Traction	This option adds a Traction node to the **Values** node, which must be entered as a vector profile.
Contact Impact	Specifies the type of contact between the two solid regions.	
	Frictionless	The contacting surfaces are free to slide over each other, however movement is prohibited normal to the interface. The shear stress at the contact interface is equal to zero, and the normal forces are equal and opposite. No additional node is added to the tree.
	Welded	The contacting surfaces undergo continuous displacements. If the solid regions are different materials, discontinuities in the strains are preserved across the material interface. No additional nodes are added to the tree.

Contact Interface Values

You can use the properties that are listed below to adjust the specifications of a **Physics Values** interface node.

Contact Resistance	The value of the resistance to conduction through the interface. It is entered as a scalar profile. If a field function defines the value and the Ignore Boundary Values property on the field function is ticked, the function is computed using data from the cell next to the boundary.	
	Dimensions	Specify the resistance in units of $m^2 \cdot K / W$or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Heat Flux	Specify the heat transferred per unit area. It is activated when the **Energy Source Option** is set to **Heat Flux**.	
	Dimensions	Specify the heat flux in units of W/m^2 or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Heat Source	Specify the heat transferred. It is activated when the **Energy Source Option** is set to **Heat Source**.	
	Dimensions	Specify the heat source in units of W or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.

Intersection	Added for all direct interfaces to control the tolerance. The tolerance is a fraction of the smallest cell face on the secondary boundary that the secondary boundary vertices are allowed to move during an intersection. If the intersection fails, the following message is displayed in the output window - `Warning: no intersections found`. The solution cannot proceed under these conditions.

Fluid Film Interface

The fluid film interface type is used to join a fluid and a fluid-film region.

It represents the contact surface between two fluids. It has properties, conditions, values, and a pop-up menu.



Modeling

- The equation for mass conservation is used to calculate the film thickness. The volume and surface area of the film are functions of the thickness. Contributions to mass source or sink come from droplet impingement, film stripping, mass transfer, or user-defined sources.
- The equation for momentum conservation is based on kinematic and dynamic conditions at the interface between the film and the surrounding fluid (film free surface).
- The equation for energy conservation is based on contributions to the energy source or sink from droplet impingement, film stripping, or user-defined sources.

Fluid Film Interface Properties

You can use the properties that are listed below to adjust the specifications of a **FluidFilm Interface** node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)	
Boundary-1	The location of the boundary on one side of the interface (Read Only)	
Type	Defines the type of interface. Must be set to `Fluid-Film Interface`	
Topology	Defines the connection type between the interfaces	
	In-place	Uses the in-place topology.
	Periodic	Uses the periodic topology.
	Repeating	Uses the repeating topology.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Imprinted** connectivity is possible.	

Fluid Film Interface Values

You can use the properties that are listed below to adjust the specifications of a **Physics Values** interface node.

Contact Resistance	The value of the resistance to conduction through the interface. It is entered as a scalar profile. If a field function defines the value and the Ignore Boundary Values property on the field function is ticked, the function is computed using data from the cell next to the boundary.	
	Dimensions	Specify the resistance in units of $m^2 \cdot K / W$ or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Intersection	Added for all direct interfaces to control the tolerance. The tolerance is a fraction of the smallest cell face on the secondary boundary	

	that the secondary boundary vertices are allowed to move during an intersection. If the intersection fails, the following message is displayed in the output window - `Warning: no intersections found`. The solution cannot proceed under these conditions.

Fully Developed Interface

A fully developed interface is used to specify the periodic pressure drop for periodic (cyclic) repeats.

Although it makes most physical sense that this type is applied to periodic interfaces, it is not restricted to periodic interfaces. It has properties, conditions, values, and a pop-up menu.



Modeling

- For both the Coupled and Segregated Flow models, the pressure jump Δp is specified at the interface.
- For both the Coupled and Segregated Flow models, the mass flow rate $\dot{m}$ may be specified instead of the pressure jump.
- Each interface has a direction from side 0 to side 1. When you create the interface, select side 0 first, then side 1. For a fully developed interface, the upstream is side 0. Therefore if you select the sides in the wrong order, the sign of the pressure jump is reversed.
- This interface type may only be used between regions belonging to the same fluid continuum.
- Fully developed periodic heat transfer cannot be accommodated. Therefore, if energy is activated in the simulation, the boundaries of the fully developed continuum must be adiabatic.
- If the mass flow rate is specified, there can be no inlets or outlets that are connected to the continuum containing the interface. In addition,

specifying the mass flow rate typically does not make sense for an in-place interface.

- When debugging it is appropriate to report mass flow rates and pressures on the interface boundaries. This is more representative than creating derived parts nearby.

Fully Developed Interface Properties

You can use the properties that are listed below to adjust the specifications of a **Fully-developed Interface** node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)	
Boundary-1	The location of the boundary on one side of the interface (Read Only)	
Type	Defines the type of interface. Must be set to `Fully-developed Interface`.	
Topology	Defines the connection type between the interfaces	
	In-place	Uses the in-place topology.
	Periodic	Uses the periodic topology.
	Repeating	Uses the repeating topology.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Imprinted** connectivity is possible.	

Fully Developed Interface Conditions

You can use the properties that are listed below to adjust the specifications of a **Physics Condition** node.

Fully Developed Flow Option	Specifies the flow when either the Coupled or Segregated Flow model is used.	
	Pressure Jump	Specifies pressure jump for the interface. A **Pressure Jump** interface value is added.

	Mass Flow Rate	Calculates the mass flow rate for the interface. When the Segregated Flow model is used, a **Mass Flow Rate** interface value is added. When the Coupled Flow model is used, this value and a **Mass Flow Rate Controls** interface value are added.
Fully Developed Energy Option	Specifies the energy option at the interface.	
	Periodic	No temperature or enthalpy discontinuity at the interface.
	Constant Temperature Walls	This option corresponds to the situation where the walls of the periodic domains are held at constant temperatures. An interface condition **Inflow Temperature Specification** is added. The interface values **Bulk Inflow Temperature** and **Temperature Reference** are added. If the wall temperature does vary, the scaling of the temperature profile is not correct, and an error is incurred.
	Constant Heat Flux Walls	This option corresponds to the situation where a constant heat flux is introduced at the

		boundaries of the periodic domain. An interface condition **Inflow Temperature Specification** is added. A **Bulk Inflow Temperature** interface value is added. If the heat flux does vary, the scaling of the temperature profile might be affected, introducing some additional error. Periodically varying heat fluxes might be acceptable, since the resulting temperature field is linear, allowing the principle of superposition to be applied.
Inflow Temperature Specification	Specifies the temperature option at the interface. Activated when the **Constant Temperature Walls**or **Constant Heat Flux Walls**is selected for **Fully Developed Energy Option**.	
	Bulk Mean Inflow Temperature	Specifies the bulk mean temperature at the inflow boundary. A **Bulk Inflow Temperature** interface value is added.
	Minimum Inflow Temperature	Specifies the lower limit of the temperature at the inflow boundary. A **Minimum Inflow Temperature** interface value is added.
	Maximum Inflow Temperature	Specifies the upper limit of the temperature at the inflow boundary. A

		Maximum Inflow Temperature interface value is added.

Fully Developed Interface Values

You can use the properties that are listed below to adjust the specifications of a **Physics Values** node.

Bulk Inflow Temperature	Specifies the bulk mean temperature at the inflow boundary.	
Temperature Reference	Activates a standard *Select Objects* dialog for selecting boundaries.	
Mass Flow Rate	The mass flow rate for the fully developed interface.	
Maximum Inflow Temperature	Specifies the upper limit of the temperature at the inflow boundary. For the **Constant Temperature Walls** condition, this value should be larger than the reference temperature.	
Mass Flow Rate Controls	Options for the *Mass Flow Rate*.	
	Computed Pressure Jump	The pressure jump across the interface (read-only).
	Verbose	Expert Property. Additional output is printed while the simulation is running when this option is activated.
	Update Frequency	Expert Property. Sets the number of iterations or time-steps that must occur before the value is updated.
	Under-Relaxation Factor	Expert Property. At each iteration, this property governs the extent to which the old solution is supplanted

		by the newly computed solution.
Minimum Inflow Temperature	Specifies the lower limit of the temperature at the inflow boundary. For the **Constant Temperature Walls** condition, this value should be smaller than the reference temperature.	
Pressure Jump	This interface value is only applicable when **Pressure Jump** is selected as the **Fully Developed Flow Option**.	
Intersection	Added for all direct interfaces to control the tolerance. The tolerance is a fraction of the smallest cell face on the secondary boundary that the secondary boundary vertices are allowed to move during an intersection. If the intersection fails, the following message is displayed in the output window - `Warning: no intersections found`. The solution cannot proceed under these conditions.	

Internal Interface

The internal interface joins two regions within the same continuum.

It can be used to combine together separate regions from the same continuum for in-place interfaces or periodic (cyclic) repeats for Periodic interfaces.

It is possible to combine regions of the same type in the same continuum using an internal interface. In certain cases, it is more efficient to use the fuse function to eliminate the boundaries altogether, resulting in a simpler topology without the interface. It has properties, values, and a pop-up menu.



Internal Interface Properties

You can use the properties that are listed below to adjust the specifications of an **Internal Interface** node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)	
Boundary-1	The location of the boundary on one side of the interface (Read Only)	
Type	Defines the type of interface. Must be set to `Internal Interface`.	
Topology	Defines the connection type between the interfaces	
	In-place	Uses the in-place topology.
	Periodic	Uses the periodic topology.
	Repeating	Uses the repeating topology.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Imprinted** connectivity is possible.	

Internal Interface Values

You can use the properties that are listed below to adjust the specifications of a **Physics Values** node.

Intersection	Added for all direct interfaces to control the tolerance. The tolerance is a fraction of the smallest cell face on the secondary boundary that the secondary boundary vertices are allowed to move during an intersection. If the intersection fails, the following message is displayed in the output window - `Warning: no intersections found`. The solution cannot proceed under these conditions.

What Is an Indirect Interface?

An indirect interface represents an association between the faces of the boundaries that compose it.

The indirect interface makes use of such associativity (rather than direct/imprinted connection) to model the transfer of mass, momentum, energy, and other physical quantities across the interface.

There are two interface types available for indirect interfaces:

- Mapped Contact Interface
- Mixing Plane Interface

Each interface type is relevant to different applications and they use different techniques to create the association between the interfacing boundaries.

Mapped Contact Interface

The Mapped Contact Interface is an indirect interface type for heat exchange at fluid/solid and solid/solid boundaries.

This interface type is a variation of contact interface type applicable to indirect interfaces. The mapped contact interface relies on an indirect association between faces of the interfacing boundaries, for data mappers, rather than an imprinted connection, for direct interfaces.

The Mapped Contact Interface is created on the volume meshes generated by the meshing process. When the interface is created from two boundaries, two new interface boundaries are automatically added. The original boundaries from which the interface is created are not destroyed.

Each face of the original boundary is linked with a group of faces on the opposing boundary using a proximity-based algorithm. Once the associativity of all the faces on the original boundary has been processed, the faces are split into an associated face-set and a disassociated face-set.

The associated face-set is transferred on to the interface boundary, while the unassociated face-set is retained on the original boundary. This is analogous to the creation of a direct interface, except that this process does not result in any new faces.

The main advantage of a Mapped Contact Interface is that it allows for a non-conformal mesh across the interface. This type of mesh is desirable in cases where the interfacing boundaries have different mesh resolutions. For instance, the heat conduction in the solid region requires a coarser mesh compared to resolving the boundary layer in the fluid domain. By precluding the need for imprinting, this approach retains the high quality of the originally created mesh and also reduces the resultant face count. A co-product of this approach is better accuracy

in computing the face gradients as well as better quality mesh for visualization. This approach is also better suited for cases where the geometric features are not equally resolved on either side of the interface. While imprinting may leave more non-intersected faces, this technique can find an associated face on the opposing boundary more robustly.

The Mapped Contact Interface has properties, conditions, values, and a pop-up menu.



Mapped Contact Interface Properties

You can use the properties that are listed below to adjust the specifications of a **Mapped Contact Interface** node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)
Boundary-1	The location of the boundary on one side of the interface (Read Only)
Type	Defines the type of interface. Must be set to `Mapped Contact Interface`
Topology	Defines the connection type between the interfaces. Only the **Indirect** type is possible.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Mapped** connectivity is possible.

Mapped Contact Interface Conditions

You can use the properties that are listed below to adjust the specifications of a **Physics Condition** interface node.

Energy Source Option	Provides energy source options for the interface.	
	None	Do not specify an energy source.

	Heat Flux	Specifies a user-defined heat flux in W/m^2.
	Heat Source	Specifies a user-defined heat source in W.
Explicit Energy Conservation	Allows you to turn off the explicit energy conservation for the boundary. Explicit energy conservation is enforced by default and is appropriate for most simulations. In some cases, enforcing explicit energy conservation can cause numerical instabilities. These instabilities require you to decrease the energy URF, which causes the simulation to take longer to converge. You can negate this effect by deactivating this setting at the appropriate boundaries. However, relaxing the requirement that energy be conserved can adversely affect solution accuracy in these areas.	
	Activated	Energy is conserved at the boundary. This option is the default.
	Deactivated	Energy is not conserved at the boundary.
External Load	Specifies the method that is used to define the load. Only available for fluid-solid contact interfaces when the Solid Stress model is activated.	
	Force	This option adds a **Force** node to the **Values** node, which must be entered as a vector profile.
	Pressure	This option adds a **Solid Stress Pressure** node to the **Values** node, which must be entered as a scalar profile.

	Traction	This option adds a **Traction** node to the **Values** node, which must be entered as a vector profile.
Contact Impact	Specifies the type of contact between the two solid regions.	
	Welded	The contacting surfaces undergo continuous displacements. If the solid regions are different materials, discontinuities in the strain are preserved across the material interface.
	Frictionless	The contacting surfaces are free to slide over each other, however movement is prohibited normal to the interface. The shear stress at the contact interface is equal to zero, and the normal forces are equal and opposite.

Mapped Contact Interface Values

You can use the properties that are listed below to adjust the specifications of a **Physics Values** interface node.

Contact Resistance	The value of the resistance to conduction through the interface. It is entered as a scalar profile. If a field function defines the value and the Ignore Boundary Values property on the field function is activated, the function is computed using data from the cell next to the boundary identified by the *Boundary-0* property of the interface.

	Dimensions	Specify the resistance in units of $m^2 \cdot K / W$ or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Heat Flux	Specify the heat transferred per unit area. It is activated when the **Energy Source Option** is set to **Heat Flux**.	
	Dimensions	Specify the heat flux in units of W/m^2 or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Heat Source	Specify the heat transferred. It is activated when the **Energy Source Option** is set to **Heat Source**.	
	Dimensions	Specify the heat source in units of W or similar. (Read Only)
	Method	Specify the input method type from **Constant**, **Field Function**, **Table**, or **User Code**.
Interpolation Stencil	Specifies the stencil that is used to map the values.	
	Compact Stencil	This scheme builds a connectivity map by finding the nearest neighbor of a face. The scheme includes a symmetrizing step which can add extra faces to the stencil (that is, if a face on

		boundary 0 sees a face on boundary 1, then the face on boundary 1 will also see the face on boundary 0). After the connectivity map is created, linear least squares is used to define the interpolation weights
	Extended Stencil	This scheme builds a connectivity map starting with the nearest neighbor, but also adds the neighbors of the nearest neighbor before performing the symmetrizing step. This stencil will be larger than the compact stencil in general, and thus more accurate, but may be slower due to increased stencil size.
Mapped Interface Tolerance	Activates the search algorithm and specifies tolerances that the algorithm uses to find the neighbors to the cells.	
	Proximity Check	When activated, the minimum proximity check occurs.
	Minimum Proximity	The relative distance (based on cell size) to the target cell. 1 is equivalent to 100%.
	Normal Directions check	When activated, the normal direction check occurs.
	Minimum Angle	The angle, in degrees, between the source and target face normals. When the

		minimum angle is 0, the two normals point toward each other, when 180, the two point in the same direction.

Mixing Plane Interface

The mixing-plane interface type is used to compute steady-state simulations of multi-stage turbomachinery flows.

This interface makes it possible to simulate the flow across stages with different pitches moving relative to one another. At the interface, circumferentially averaged flow field data is transferred in a conservative manner between two rotationally periodic regions. These regions must be of the same continuum, and aligned on the same axis. The mixing plane interface type is available on indirect interface topologies only.

The mixing-plane interface allows the transfer of mass, momentum, energy, and other conserved quantities across the interface using a circumferential averaging technique. The values to be transferred across the interface are averaged in concentric, circumferential rings, or bins, of uniform radial thickness, covering the entire span and pitch of the boundary. Due to the circumferential averaging, the final solution is characterized by radially varying profiles of flow field quantities that remain uniform across the pitch within any given circumferential bin.



Implicit vs. Explicit Mixing-Plane Solution Approaches

The solution across the mixing-plane can be computed using either an implicit or an explicit approach.

In the implicit mixing-plane approach, the cells on either side of the interface are directly connected through each circumferential bin. The approach introduces additional degrees of freedom to the linear system -- one additional degree of freedom per variable per circumferential bin. These degrees of freedom represent the conservation of fluxes across the faces composing either side of a given bin. In this way, solution of the linear system results in implicit updates to circumferentially averaged values on either side. The implicit mixing-plane approach is available with both Coupled and Segregated Flow models.

The explicit mixing-plane approach closely follows the procedure used in the Non-Reflecting Boundary Condition. Here, updated values on the interface boundaries are computed directly, based on those from the previous iteration. As in the Non-Reflecting Boundary Condition, all outgoing waves are allowed to pass through the interface without reflection. Circumferential averaging of conserved quantities, combined with piecewise constant interpolation between bins in the radial

direction, is used to prescribe the zero-th mode solution required by the non-reflecting boundary condition. The explicit mixing-plane approach is available in combination with the Coupled Flow model only. See [23], [24], [25] and [47].

The mixing plane interface has properties, conditions, and a pop-up menu.



Mixing Plane Interface Properties

You can use the properties that are listed below to adjust the specifications of a **Mixing Plane Interface** node.

Boundary-0	The location of the boundary on one side of the interface (Read Only)
Boundary-1	The location of the boundary on one side of the interface (Read Only)
Type	Defines the type of interface. Must be set to `Mixing Plane Interface`.
Topology	Defines the connection type between the interfaces. Only the **Indirect** type is possible.
Connectivity Option	Shows how the boundaries on two sides of the interface are connected. Only **Connect Average** connectivity is possible.

Mixing Plane Interface Conditions

You can use the properties that are listed below to adjust the specifications of a **Physics Condition** interface node. You can specify either explicit or implicit methods.

Connectivity Option	Provides energy source options for the interface.	
	Implicit	Uses the implicit mixing plane

		approach. This is the default.
	Explicit	Uses the explicit mixing plane approach.
Bin Coarsening Factor	This expert property guides the number of circumferential bins (that is, constant-thickness rings or bands) used to perform the circumferential averaging ("mixing-out") of flow field data across the implicit mixing plane interface. Increase the coarsening factor to produce fewer, coarser (wider) bin rings for use in cases with disparate mesh densities on either side of the interface. The actual number of bins is set by STAR-CCM+; the factor only has an effect if the boundary meshes do not initially yield an equal number of bins on both sides of the implicit mixing plane interface. The default value is 1.	

Interface Pop Up-Menu

The pop-up menu items are available when you right-click an interface node.

Initialize	Populates the interface boundary nodes with faces from the original boundaries.
Reset	Reverses the process of initialization. Any vertices that the process removes are reported in the output window for each region.
Insert Shell Region Interface	Replaces the interface with a shell region and creates two new contact interfaces that connect the shell region to the parent boundaries on either side. The new shell region can be associated with a shell continuum to model both in-place and out-of-plane heat conduction. Only available for in-place interfaces. See Inserting a Shell Region at an In-Place Interface.

Swap Boundaries	Exchanges the Boundary-0 value for Boundary-1. Only available for direct interfaces.

See also: common menu items.

Meshing

A mesh is the discretized representation of the computational domain, which the physics solvers use to provide a numerical solution.

STAR-CCM+ provides meshers and tools that you can use to generate a quality mesh for various geometries and applications.

STAR-CCM+ provides several meshing strategies that are suitable for different applications. To identify the correct strategy for your case, see Comparison of Meshing Strategies.

For conceptual information on mesh components, see Mesh Construction.

The following surface meshers are available:

- Surface Wrapper

  Wraps the initial surface to provide a closed and manifold surface mesh from complex geometry.

- Surface Remesher

  Remeshes the initial surface to provide a quality discretized mesh that is suitable for CFD.

The following volume meshers are available:

- Trimmed Mesher

  Generates a volume mesh by cutting a hexahedral template mesh with the geometry surface.

- Polyhedral Mesher

  Generates a volume mesh that is composed of polyhedral-shaped cells.

- Tetrahedral Mesher

  Generates a volume mesh that is composed of tetrahedral-shaped cells.

- Prism Layer Mesher

  Adds prismatic cell layers next to wall boundaries. The mesher projects the core mesh back to the wall boundaries to create prismatic cells.

- Advancing Layer Mesher

Creates a volume mesh composed of prismatic cell layers next to wall boundaries and a polyhedral mesh elsewhere. The mesher creates a surface mesh on the wall and projects it to create the prismatic cell layers.

- Extruder Mesher

  Generates an extruded mesh region from a boundary that one of the core volume meshers has meshed.

- Thin Mesher

  Generates a prismatic layered volume mesh for thin geometries.

- Generalized Cylinder Mesher

  Generates a volume mesh appropriate for elongated cylindrical regions.

The following structured mesher is available:

- Directed Meshing

  Generates a mesh by creating a two-dimensional patch mesh on a surface and sweeping it through the volume along a specified guide surface.

Meshing Overview

This section contains the following subsections:

- Comparison of Meshing Strategies

  STAR-CCM+ provides several meshing strategies that are suitable for different applications. Use this section to identify the correct strategy for your case.

- Mesh Construction

  The volume mesh in a simulation is the mathematical description of the space (or geometry) of the problem being solved. This section describes the entities that STAR-CCM+ uses to construct a mesh.

- Parallel Meshing

  Allows you to speed up volume mesh generation using multiple processors. Large parts are subdivided and meshed on several processors.

- Concurrent Meshing

  Allows you to speed up volume mesh generation using multiple processors. Individual parts are meshed on separate processors. This meshing mode is intended for meshing large multi-part assemblies of small parts in per-part mode.

- Base Size

  The **Base Size** is a characteristic dimension of the model that you set before using any relative values.

Comparison of Meshing Strategies

STAR-CCM+ provides several meshing strategies that are suitable for different applications.

Use this section to identify the correct strategy for your case.

The following meshing strategies are available:

- Unstructured Meshing:
  - Parts-Based Meshing

    Parts-based meshing detaches the meshing from the physics and provides a flexible and repeatable meshing pipeline. This strategy provides several advantages over region-based meshing (see Advantages of Parts-Based Meshing).
  - Region-Based Meshing

    Region-based meshing meshes the surface and volume at the region-level. Only use region-based meshing if you intend to use a mesher that is not available in parts-based meshing (see Table of Supported Meshers).
- Structured Meshing:
  - Directed Meshing

    Directed meshing generates high-quality, structured meshes on CAD geometry. It operates by sweeping a surface mesh from the surface of a CAD geometry through its volume.

Table of Supported Meshers

The following table summarizes the meshers that are currently available in parts- and region-based meshing.

Mesher	Region-Based Meshing	Parts-Based Meshing
Surface Wrapper	x	x
Surface Remesher	x	x
Automatic Surface Repair	x	x
Polyhedral Mesher	x	x
Tetrahedral Mesher	x	x

Trimmed Mesher	x	x
Prism Layer Mesher	x	x
Advancing Layer Mesher	x	
Extruder Mesher	x	
Thin Mesher	x	
Generalized Cylinder Mesher	x	

Advantages of Parts-Based Meshing

Parts-based meshing is the focus of all current and future meshing developments in STAR-CCM+. It uses familiar concepts—composites, parts, surfaces, and operations—that are common with CAD and CAE software.

This meshing strategy executes a pipeline of mesh operations on geometry parts. The pipeline provides the following advantages:

- Greater Flexibility: To process the CAD import to a volume mesh, you can define a complex workflow of mesh operations and execute all operations in one action.
- Better Repeatability: To make design changes, you can modify an input part and the changes are propagated through the pipeline to the volume mesh.
- Intuitive Workflow: To build complex linear workflows easily, the simulation tree displays the mesh operations in the sequential order that the pipeline executes them.

In addition to the pipeline, parts-based meshing has several changes in implementation that make it easier to use. The following list summarizes these advantages:

- Custom Mesh Controls: To refine or coarsen the mesh, you can create a single mesh control and apply it to several parts, surfaces, curves, or volumes at once. Previously, you changed the mesh values for each region, boundary, feature curve, volume source individually.
- Inherited Parts: To find objects in complex geometries and pipelines easily, you can categorize the input parts for a mesh operation separately from the part surfaces and curves that the operation generates.

- Pre-Execute Parts: To save time when generating the pipeline, you can identify the inherited parts and new part surfaces and curves without executing the mesh operation.
- Surface Wrapper Operation: To make it easier to use the surface wrapper without the other meshers, you use the surface wrapper as a separate mesh operation (instead of a surface mesher).
- Decoupling the meshing from the physics.
- Consistent feature set and usage across all surface and volume meshers.

For unstructured meshing, parts-based meshing is the recommend strategy, unless you intend to use a mesher that is not currently available (seeTable of Supported Meshers).

Mesh Construction

The volume mesh in a simulation is the mathematical description of the space (or geometry) of the problem being solved.

This section describes the entities that STAR-CCM+ uses to construct a mesh.

A mesh is composed of the following mesh entities:

- Vertex
- Face
- Cell

If generating a volume mesh within STAR-CCM+ starting from a surface mesh, then feature curve entities are also used.

Since STAR-CCM+ expects the mesh to be valid and of good quality, the details of the mesh are important until these aspects have been checked.

Vertex

A vertex is a point in space that you can define with a position vector.

A number of vertices can be used to define a feature curve or a face.



The position vector is an available field function that is composed of the vertex coordinates. This vector is quantitatively different from the centroid vector of a cell or a face.

Feature Curve

A feature curve at its most basic consists of two vertices such that it defines a line in two-dimensional space..



Feature curves are created based on an underlying surface mesh and are used to represent sharp edges and/or surface detail that you want to maintain during the meshing process. They are not used as part of the simulation analysis and are not present in the volume mesh.

Face

A face comprises an ordered collection of vertices such that they define a surface in three-dimensional space.

In STAR-CCM+, faces can have any number of vertices. Four or more faces can be used to define a 3D cell.



Make sure that faces are always simple polygons (that is, the edges do not cross each other) and non-concave.

Correct Face Orientation

You define the orientation of a face with the right-hand rule. In the above example, the face orientation is upwards (assuming you see vertex 2 as being further out of the screen than vertex 4).

When creating a mesh in your mesh generator, make sure that the mesh conforms to this rule.

Concave Face

Concavity can be calculated for any face with four or more sides and refers the amount that one part of the face “sees” another part of the same face. A non-concave cell is flat or planar. However, as the surface becomes warped or twisted, the face starts folding back on itself and produces concavity.

The concavity of a face can be determined by triangulating a face and then extend the normals of each triangle outward until they cross. The angle between the face normals is the concave angle, as illustrated below:



A planar face would have a concave angle of 0 degrees. Keep concave face angles as low as possible make sure that they do not exceed 89 degrees.

See Also

Zeroing Reconstruction

Cell

A cell is an ordered collection of faces that defines a closed volume in space.



In STAR-CCM+, cells can have any number of faces as long as there are enough faces to create a closed cell volume. The faces of a cell cannot intersect each other, except where they touch along the common edges.

In three dimensions, cells of arbitrary polyhedral shape are permitted. In two dimensions, cells of arbitrary polygons are permitted.

Parallel Meshing

You can use multiple cores to speed up volume mesh generation and create larger meshes than it is possible to create with a single core.

This feature is particularly useful when meshing large parts. Parallel meshes are conformal and of comparable quality to serial meshes. To use parallel meshing in a simulation, start the simulation in parallel mode.

The following meshing models support parallel meshing:

- Polyhedral Mesher (including the underlying tetrahedral mesher);
- Tetrahedral Mesher
- Prism Layer Mesher (excluding the prism layer subsurface generation)
- Trimmed Mesher

The following meshing models do not support parallel meshing:

- Thin Mesher
- Extruder Mesher
- Generalized Cylinder Mesher
- Advancing Layer Mesher

For an average mesh (ten million cells or less), it is best to use between four and eight cores with one to two million cells per core. Using a large number of cores for a relatively small mesh reduces scalability and can result in slower performance. For polyhedral and tetrahedral meshes, a minimum of 250 thousand cells are needed per core. However, for trimmed meshes, you can use a minimum of 100 thousand cells per core. If the cell count for each core in parallel falls below the minimum number of cells per core, the trimmed mesher automatically uses fewer cores than are available and re-distributes the workload between these cores. This behavior ensures that there are a suitable number of cells for each core. When the total number of cells that are meshed in parallel falls below the minimum required for each core, the trimmed mesher automatically falls back to serial.

For distributed memory machines, the mesh size is limited by the number of cores and the amount of memory that is available to each core. Ensure that you have enough cores—and memory available to each core—for your entire mesh. The following table shows general memory requirement guidelines when meshing in parallel using the

different mesh types. Memory requirement for some cases can fall outside these values.

Mesh Type	Minimum Gb/million cells	Recommended Gb/million cells
Polyhedral	-	1.00-1.50
Tetrahedral	0.08	0.10-0.20
Trimmed	0.60-0.80	1.00-1.50

Note: The *Use Parallel Meshing* property affects the selected automated mesh operation only; if multiple mesh operations are present, you can run each one independently in serial or parallel.

You can perform Parallel Meshing in:

- Parts-Based Meshing
- Region-Based Meshing

Creating a Parallel Mesh

Parallel meshing is useful for speeding up volume mesh generation by meshing large parts on several processors.

When meshing a large number of small parts using several processors, use Concurrent Meshing.

Before Creating a Parallel Mesh:

Create a new simulation, or load an existing simulation, in parallel mode.

See Starting in Parallel Mode

To Create a Parallel Mesh in Parts Based Meshing:

1. Create an Automated Mesh Operation.
2. Select the **Geometry > Operations > Automated Mesh** node and in the *Properties* window, activate *Enable Parallel Meshing*.



3. Set the mesh controls that you require.

See Setting Default Controls and Creating and Defining Custom Controls

4. Execute the mesh operation.

To Create a Parallel Mesh in Region Based Meshing:

5. Select the **Continua > Mesh** node and in the *Properties* window, activate *Enable Parallel Meshing*.



6. Set the mesh model properties that you require.

See Setting Mesh Model Properties

7. Generate the Volume Mesh.

Concurrent Meshing

You can use multiple processors to speed up volume mesh generation and mesh individual parts on separate processors.

This feature is useful when meshing large multi-part assemblies of relatively small parts in per-part mode.

Instead of meshing each individual part successively in serial, the concurrent mesher meshes separate parts simultaneously on different processors that are available.

Concurrent meshing behaves in a similar way to parallel meshing, however, the distribution of parts between processors is different:

- During concurrent meshing, each processor meshes a different part.
- During parallel meshing, parts are subdivided and meshed between different processors.

All except one of the processors that are available for concurrent meshing are assigned separate parts for meshing. This processor acts as a control node and delegates work to the other processors. For this reason, there is no improvement in speed-up when using fewer than three processors for concurrent meshing.

Meshes that are generated using concurrent meshing are non-conformal and are of comparable quality to meshes that are generated in serial.

The following meshing models support concurrent meshing:

- Surface Remesher
- Surface Wrapper
- Trimmed Mesher
- Polyhedral Mesher (including the underlying tetrahedral mesher)
- Tetrahedral Mesher
- Prism Layer Mesher

The following meshing models do not support parallel meshing:

- Thin Mesher
- Extruder Mesher
- Generalized Cylinder Mesher
- Advancing Layer Mesher

Creating a Concurrent Mesh Operation

Concurrent meshing is useful for speeding up volume mesh generation by meshing large multi-part assemblies of small parts on several processors.

When meshing large parts with several processors, use parallel meshing.

Note: Concurrent meshing is available in Parts-Based Meshing and not Region-Based Meshing.

Before Creating a Concurrent Mesh Operation:

1. Create a new simulation, or load an existing simulation, in parallel mode.
2. Make sure that there are three or more processors available.

To Create a Concurrent Mesh Operation:

3. Create an Automated Mesh operation with compatible meshing models.
4. Select the Automated Mesh node, then:
   - Activate *Per-Part Meshing*
   - Set *Mesher Execution Mode* to **Concurrent**



5. Set the mesh controls that you require.

See Setting Default Controls and Creating and Defining Custom Controls

6. Execute the mesh operation.

Note: If you attempt to abort the concurrent meshing process, the operation will only abort after all processors have finished meshing the parts that are currently assigned to them.

Base Size

The **Base Size** is a characteristic dimension of the model that you set before using any relative values.

As general examples, you can set the base size to the diameter of an inlet, the length of the fluid volume, or a size that is convenient for scaling other values.

To set the base size:

- Select the **Default Controls** > **Base Size** node and set the *Value* property.

  For regions, select the **Reference Values** > **Base Size** node



Relative and Absolute Sizes

To provide more flexibility, you can select to enter reference values for certain models using either relative or absolute sizes.

Relative values are set as a percentage of the base size while absolute values can be entered in any of the allowable length units. The advantage of allowing a relative size specification is that changing the base size scales both global and object specific sizes.

To set a property as relative or absolute:

- Select an appropriate node and set *Relative/Absolute* to either **Relative to base** or **Absolute**.



Parts-Based Meshing

Parts-based meshing is a meshing strategy in which you create one or more mesh operations to define the steps that STAR-CCM+ performs to generate the volume mesh.

Mesh operations are actions that you use to:

- prepare the initial geometry
- surface wrap geometry
- generate surface and volume meshes
- Parts-based meshing detaches the meshing process from physics modeling and provides a flexible and repeatable sequence of mesh operations. The sequence of mesh operations is called the meshing pipeline.

See Also

Meshing Overview
Comparison of Meshing Strategies
Advantages of Parts-Based Meshing

Preparing Geometry for Meshing

Before auto-meshing or wrapping geometry parts, perform mesh operations to prepare the geometry.

Before Preparing Geometry for Meshing:

1. Provide the geometry to use in the simulation. There are two options:
   - Import CAD data or surface mesh

     See Importing Geometry.
   - Create a model using 3D-CAD

     See Creating Geometry Parts from 3D-CAD.
2. If you create a model in 3D-CAD, convert the 3D-CAD model to geometry parts.
   a. Right-click the **Geometry > 3D-CAD Models > [3D-CAD Model]** node and select **New Geometry Part**.

      The *Parts Creation Options* dialog appears.

   b. Click **OK**.

To Prepare Geometry for Meshing:

1. Apply any necessary boolean operations.
   a. Right-click **Geometry > Operations > New** and select one of the following Boolean operations:
      - **Unite Parts**

        Merges multiple overlapping/coincident parts into a single part.

        See Uniting Parts.
      - **Intersect Parts**

        Creates a part from the common volume of the selected parts.

        See Intersecting Parts.
      - **Subtract Parts**

        Removes the volume of one part from another part.

        See Subtracting Parts.

b. Complete the properties of the [*Boolean operation*] dialog as required and click **OK**.

   Unexecuted Boolean operation nodes are now displayed under the **Geometry** > **Parts** and **Geometry** > **Operations** managers.

c. Right-click the **Operations** > **[Boolean operation]** node and select **Execute**.

2. If there are holes or openings in the geometry, perform the **Fill Holes** operation before extracting the volume.

   a. Right-click the **Geometry** > **Operations** node and select **New** > **Fill holes**.

   b. Complete the properties of the **Fill Holes** operation dialog as required and click **OK**.

      Unexecuted **Fill Holes** nodes are now displayed under the **Geometry** > **Parts** and **Geometry** > **Operations** managers.

   c. Right-click the **Operations** > **Fill Holes** node and select **Execute**.

      See Fill Holes Operations for details of the operation properties.

3. If the imported geometry part represents the solid surrounding a fluid, and the fluid is the part being modeled, use the **Extract Volume** operation to create the Fluid part.

   See Extract Volume.

4. If faces of adjacent parts are in contact, use the **Imprint** operation to ensure that common surfaces of each part retain their own faces. A conformal mesh can still be formed between them.

   See the Imprint Parts Operation.

5. Assign parts to regions.

   See Assigning Parts to Regions.

6. Set boundary types for each boundary.

At the end of this task, you have a geometry that is ready for surface and volume meshing. However, if the geometry is complex, use the Surface Wrapper before generating the surface and volume mesh.

Assigning Parts To Regions

Assigning parts to regions creates a relationship between the meshed part and the physics region.

This action allows STAR-CCM+ to convert the volume mesh to a finite volume representation that the solver can use to obtain a solution.

You can assign part-level objects (parts, part surfaces, and part curves) to their equivalent region-level objects (regions, boundaries, and feature curves). From parts to regions, these assignments do not need to be one-to-one. That is, you can assign multiple part-level objects to the same region-level objects. The following diagram illustrates the connection between parts and regions.



However, you cannot assign multiple region-level objects to a single part-level object. If you reassign any of these objects to a new destination, it is automatically removed from its original assignment.

Note: Part contacts only become interfaces when you assign all contacting parts in the same operation.

The following methods are available for assigning parts to regions:

Using the Assign Parts to Regions Dialog

This dialog assigns several part-level objects to region-level objects in one action. This dialog is the recommended method of assigning parts to regions.

To assign parts to regions using the *Assign Parts to Regions* dialog:

1. Right-click a **Geometry > Parts > Part** node and select **Assign Parts to Regions**.



2. In upper-half of the *Assign Parts to Regions* dialog, select the parts that you want to assign to regions.
3. In the lower part of the *Assign Parts to Regions* dialog, set the properties to define how the parts are assigned to regions. See Assign Parts to Regions Dialog.



4. Click **Apply** and then **Close** to exit the dialog.

Assign Parts to Regions Dialog

First drop-down menu	**Create a region for Each Part**	Creates and assigns different regions to each selected part.
	Assign All Parts to an Existing Region	Assigns all the selected parts to an existing region. Selecting this option activates the option to select the existing region.

	Create One Region for All Parts	Creates and assigns one region to all the selected parts. Selecting this option allows you to name the region. The default name is `Region`.
Second drop-down menu	**Create a Boundary for Each Part Surface**	Creates and assigns different boundaries to each part surface.
	Create a Boundary for Each Part	Creates and assigns one boundary to all parts surfaces within each part.
	Create One Boundary for All Part Surfaces	Assigns all the part surfaces to one boundary. Selecting this option allows you to name the boundary. The default name is `Boundary`.
Third drop-down menu	**Create a Feature Curve for Each Part Curve**	Creates and assigns different feature curves to each part curve.
	Create One Feature Curve for All Part Curves	Assigns all selected part curves to one feature curve. Selecting this option allows you to name the feature curve. The default name is `Feature Curve`.
Create Interfaces from Contacts	**Activated**	Automatically generates interfaces from part contacts. When using this option, the valid settings for **Region Mode** and **Boundary Mode** depend on how you created the part contacts (see Table of

		Valid Methods and Options).
	Deactivated	Does not create any interfaces from part contacts.

Using the Parts Context Menu

The parts context menu allows you to assign a single part-level object to a region-level object.

The following capabilities are available with the context menu:

- Assigning Part-Level Objects to Existing Region-Level Objects
- Creating and Assigning Part-Level Objects to New Region-Level Objects

Assigning Part-Level Objects to Existing Region-Level Objects

You can use the context menu to assign a part-level object to a previously created region-level object.

The following section assigns a part to an existing region. You can use a similar procedure for assigning part surfaces to boundaries and part curves to feature curves.

Note: To assign multiple parts to an existing region, you must first create a composite part. To assign multiple part surfaces to boundaries or part curves to feature curves, you can multi-select the objects.

Setting a part to a region with this method assigns all part surfaces with the following conventions:

- If the name of a part surface matches the name of an existing boundary, this method assigns the part surface that boundary.
- If a boundary has the name, `Default`, this method assigns all other part surface to this boundary.
- If the `Default` boundary does not exist, this method assigns all other parts surfaces to the boundary with the lowest index. You can identify the index number of each boundary in the *Properties* window.

Note: A part-level object cannot belong to multiple region-level objects. Therefore, if you reassign a part to a different region, it is automatically removed from the original assignment.

To assign a part to a region with the context menu:

1. Make sure that a region exists.
2. Right-click a part and select **Set Region > Existing**.



3. Make a selection in the object selector. See The Object Selection Dialog.



Creating and Assigning Part-Level Objects to New Region-Level Objects

You can use the context menu to create and assign part-level objects to region-level objects. This facility is only available for part surfaces and part curves.

This action creates and assigns a boundary with the same name as the part surface. If you select multiple part surfaces, the new boundary name is based on the keyword **Boundary** (for example, Boundary 1, Boundary 2).

The following section creates and assigns a part surface to a new boundary. You can use a similar procedure for assigning part curves to feature curves.

To assign a part surface to a new boundary with the context menu:

- Right-click the required part surface and select **Set Boundary** > **New**.



You can also multi-select multiple boundaries.

Using the Part Level Properties

The part-level properties can assign the selected part-level object to an existing region-level object only.

The following section assigns a part to an existing region. You can use a similar procedure for assigning part surfaces to boundaries and part curves to feature curves. However, you can only assign part surfaces and part curves once you have assigned their parent part to a region.

Note: A part-level object cannot belong to multiple region-level objects. Therefore, if you reassign a part to a different region, it is automatically removed from the original assignment.

To assign a part to an existing region with the part-level properties:

1. Select the required part node and click the right half of the *Region* property.



2. Make a selection in the object selector.



Using the Region Level Properties

The region-level properties can assign the selected region-level object to a part-level object.

The following section assigns a region to a part. You can use a similar procedure for assigning boundaries to part surfaces and feature curves to part curves. However, you can only assign boundaries and feature curves once you have assigned their parent region to a part.

Note: A part-level object cannot belong to multiple region-level objects. Therefore, if you reassign a part to a different region, it is automatically removed from the original assignment.

To assign a region to a part with the region-level properties:

1. Select the required region and click the ellipsis **(...)** next to the *Parts* property.



2. In the object selector, select the assigned part and click **OK**.

Surface Wrapping the Geometry

The Surface Wrapper generates a closed and manifold surface mesh from poor quality or complex CAD geometry.

Before Surface Wrapping the Geometry:

- Complete Preparing Geometry for Meshing.

  See Preparing Geometry for Meshing.

To Surface Wrap the Geometry:

1. Create a **Surface Wrapper** mesh operation.
   a. Right-click **Geometry** > **Operations** and select **New** > **Surface Wrapper**.
   b. Select the parts to wrap in the **Create Surface Wrapper Automated Mesh Operation** dialog. Click **OK**.

      Unexecuted **Surface Wrapper** nodes are now displayed under the **Geometry** > **Parts** and **Geometry** > **Operations** managers.
2. To exclude part surfaces from a Surface Wrapper operation, perform partial wrapping. See Performing Partial Wrapping.
3. Set the Surface Wrapper properties.

   Select the **Geometry** > **Operations** > **Surface Wrapper** node and activate the required properties in the **Surface Wrapper - Properties** window.

   See Surface Wrapper Properties.
4. If default or custom controls need further adjustment, expand the **Operations** > **Surface Wrapper** node and update accordingly.
5. Execute the Surface Wrapper and Assign Parts to Regions. There are three different procedures you can follow to do this, depending on the outcome you need.
   - If you want to create interfaces from contacts between input parts and the output surface wrapper part, execute the **Surface Wrapper** operation with **Create Weak Contacts with input parts** activated before moving on to assign the surface wrapper part to a region.
     1. In the *Surface Wrapper - Properties* window, activate **Create Weak Contacts**.

Note: To select this option, in the *Surface Wrapper - Properties* window, deactivate **Per-Part Meshing**.

2. Right click the **Geometry** > **Operations** > **Surface Wrapper** node and select **Execute**.

   In-place contacts are formed between Part Surfaces.

3. Assign the surface wrapper part to a region.

   The contacts are converted to interfaces. See Assigning Parts to Regions.

- If you do not want to create interfaces from contacts between input parts and the output surface wrapper part, you can:
  - execute the surface wrapper operation with **Create Weak Contacts**deactivated in the *Surface Wrapper - Properties* window, or
  - assign parts to regions before executing the surface wrapper operation

Note: If you do not need to create interfaces, you can execute the Surface Wrapper (with **Create Weak Contacts**deactivated) or Assign the surface wrapper part to a region in either order. If **Create Weak Contacts** is activated, assign parts to regions before executing the Surface Wrapper, otherwise interfaces are created. See Assigning Parts to Regions.

- If you activate the **Per-Part Meshing** option in the *Surface Wrapper - Properties* window, no output parts form when you execute the **Surface Wrapper** operation. Instead, a new description node appears under the **Descriptions** manager. This description is a new representation of the input part.

At the end of this task, you have a geometry that is ready for surface and volume meshing.

See Also

Surface Wrapper
Mesh Operations

Setting Surface Wrapper Properties

The surface wrapper properties control the input parts and features of the surface wrapper. Specifically, per-part meshing, curvature and proximity refinement, gap closure, and mesh alignment.

To set the surface wrapper properties:

- Select the **Operations** > **Surface Wrapper** node and set the properties in the *Properties* window.



Per-part Meshing	Activating **Per-Part Meshing** causes STAR-CCM+ to wrap each part individually. In this case, results are stored as a new description for each part: there are no new parts as is the case when wrapping all parts at once. Any interfaces previously defined will be ignored causing the mesh between connected parts to be non-conformal. This option is useful if the Trimmed Mesher or Thin Mesher can not produce a mesh at the interface. When the Per-Part Meshing box is deactivated all input parts will be wrapped at once and the

	wrapper results are stored in a new part. No new description is created.
Input Parts	Specifies the parts to wrap in this operation.
Curvature Refinement	Activating Curvature Refinement increases the number of faces based on surface curvature or deviation distance.
Proximity Refinement	Activating Proximity Refinement increases the number of faces based on the distance to nearby surfaces.
Gap Closure	Activating Gap Closure creates a **Surface Wrapper > Default Controls > Gap Closure** node. This node allows you to close holes in the surface that are either: smaller than a user-defined size, specified by seed points, or both.
Mesh Alignment	Mesh Alignment aligns initial wrapping mesh lines with a user-defined point.
Create Weak Contacts with Input Parts	Activating the Create Weak Contacts With Input Parts option when the *Per-Part Meshing* option is deactivated creates new weak contacts. The weak contacts that are created do not result in a conformal mesh. However, when you assign parts to regions, these contacts are converted into interfaces which the solver then uses to make the mesh conformal.
Verbose Output	Activating Verbose Output prints more verbose output to the Output window during Meshing.

See Also

Surface Wrapper
Surface Wrapper Properties
Surface Wrapper Expert Properties

Setting the Volume of Interest

The volume of interest control determines the volume that the surface wrapper extracts from the input parts.

To set the volume of interest:

1. Select the **Surface Wrapper > Default Controls > Volume of Interest** node and set *Method* to one of the following properties:

   - **Largest Internal** — Specifies that the wrapper extracts the largest internal, self-contained volume from the input geometry surfaces.
   - **External** — Specifies that the surface wrapper extracts the outer surfaces of all parts, that is, the external surfaces.
   - **Seed Point** — Specifies that a user-defined location determines the surface that the wrapper extracts. The resultant mesh is the surface that surrounds this location. See Creating and Defining a Seed Point.
   - **Nth Largest** — Specifies that the wrapper extracts a particular surface according to volume size. The surfaces are ranked from largest to smallest, and the user-defined Volume Number determines the output surface. See Defining the Volume Number.



Creating and Defining a Seed Point

Seed points define a location within the volume of interest for the seed point volume extraction method.

To create and define a seed point:

1. Make sure that you have a scene open.
   - If you do not have a scene open, you can only define the seed point using its properties.
2. Right-click the **Surface Wrapper** > **Volume of Interest** > **Seed Points** node and select **New**.



3. To set the location of the seed point:
   - Define the position manually using the *X*, *Y*, and *Z* positions.
   - Activate the **Display Point Tool** checkbox and define the position interactively in the *Graphics* window.



4. Click **Create** and then **Close**.

Defining the Volume Number

The volume number property defines the volume of interest, ranked from largest to smallest, for the nth largest volume extraction method.

To define the volume of interest:

1. Select **Surface Wrapper > Default Controls > Volume of Interest > Volume Number** node, and then set *Volume Number* to define the volume of interest.

   For example, setting *Volume Number* to `1` extracts the largest volume and `2` extracts the second largest volume.



See Also

Surface Wrapper
Surface Mesher Controls and Values

Detecting Leaks

Before running the surface wrapper, use the leak detection tool to identify whether there will be any holes in the wrapped surface. This tool also provides access to surface repair modes so that you can fix the holes.

To detect leaks:

1. Right-click the **Operations > Surface Wrapper** node and select **Run Leak Detection**.



2. In the *Locate Points* group box, set *Current Point* to **Source Point**.



3. To set the source point location:
   - Define the point manually using the *X*, *Y*, and *Z* positions.
   - Define the point interactively using the point tool.
4. Set *Current Point* to **Target Point** and set the location as in the previous step.
5. If you require several target points, click **New Target Point** and define them.
6. When you have defined all source and target points, click **Recompute Template and Paths**.



- If there are no leaks, the *Output* window reports the following:

```
No path found between source and target points.
These points will be on different sides of the
final wrapped surface.
```

- If there is a leak, a path from the source to the target point appears in the *Graphics* window.



7. Click **Select Faces Close to Leak Path** to highlight the faces that surround the hole.



8. Set **Mode** to **Repair Surface** or **Offset Faces** to fix any leaks that the tool detects.



See Also

Surface Wrapper
Repair Surface Mode
Offset Faces Edges Mode

Gap Closure

The surface wrapper can close holes in the input geometry during surface wrapping.

Three methods are provided to control the behavior of gap closure, which you select from the properties of the Gap Closure node.



Gap Closure Properties

The following options are available for the *Gap Closure - Properties Method*:

Size	Allows you to specify gap closure **Size type** as an **Absolute** value, or **Relative to base**. Any holes that are smaller than the specified gap closure size are closed during surface wrapping.
Seed Point	Allows you to specify the location of **Seed Points** within your geometry. Only volumes that contain **Seed Points** are closed during surface wrapping. Generally, gaps are closed at the smallest possible opening between the gap closure seed point and the volume of interest. If there are multiple small openings,

	the wrapper does not guarantee which it closes.
Both	Allows you to specify a **Gap Closure Size** and use gap closure **Seed Points**. During surface wrapping, any holes that are smaller than the **Gap Closure Size** are closed. Holes that are larger than the **Gap Closure Size** are also closed based on the gap closure **Seed Points**.

Preventing Surface Contact

Contact prevention ensures that selected close surfaces do not touch each other in the wrapped surface.

The following modes of contact prevention are available:

- One Group Contact Prevention
- Two Group Contact Prevention

Preventing Contact between One Group of Surfaces

One group contact prevention surface makes sure that all surfaces within the group do not touch.

To prevent contact between one group of surfaces:

1. Right-click the **Operations > Surface Wrapper > Contact Prevention** node and select **One Group Contact Prevention Set**.



2. To apply the contact prevention to part surfaces:

- If you want to define the surfaces manually, select the **Contact Prevention > Contact Prevention 1** node and set *Surfaces* to the required part surfaces.



- If you want to define the surfaces interactively, right-click the **Contact Prevention > Contact Prevention 1** node and select the part surfaces with the *Graphics* window and selection buttons.
  - This option is only available if you have a scene open.



3. Set the *Minimum Size* to the lowest edge size that the wrapper generates in the refinement.
   - This value must be greater than 0.



Preventing Contact between Two Groups of Surfaces

Two group contact prevention makes sure that surfaces between the two groups do not touch, but surfaces within groups can touch.

To prevent contact between two groups of surfaces:

1. Right-click the **Operations > Surface Wrapper > Contact Prevention** node and select **Two Group Contact Prevention Set**.



2. To apply the contact prevention to part surfaces:

- If you want to define the surfaces manually, select the **Contact Prevention > Contact Prevention 1** node and set *Surface set 1* to the required part surfaces.



- If you want to define the surfaces interactively, right-click the **Contact Prevention > Contact Prevention 1** node and select the part surfaces with the *Graphics* window and selection buttons.
  - This option is only available if you have a scene open.



3. Repeat the previous step for *Surface set 2*.
4. Set the *Minimum Size* to the lowest edge size that the wrapper generates in the refinement.
   - This value must be greater than 0.



Note: A Tag appears on any new part surface that is created during gap closure. The tag differentiates the surfaces that gap closure creates from the pre-existing surfaces. See What Are Tags?.

See Also

Surface Wrapper

Performing Partial Wrapping

Some geometries only require surface wrapping for a portion of their surface. The surface wrapper allows you to exclude part surfaces that you consider sufficiently well tessellated for subsequent mesh operations.

For example, the geometry below shows the original tessellated surface of an aircraft. The main aircraft body is closed and manifold, does not require surface wrapping, and contains many intricate details. However, the engines that are part of the same geometry, as highlighted below, require surface wrapping.



Wrapping the whole geometry is unnecessary and can defeature highly detailed part surfaces–unless the surface wrapper is excessively refined using custom controls and contact prevention. Instead, you can preserve specified input surfaces and perform partial wrapping.

Follow this procedure to exclude specific part surfaces on a geometry from a Surface Wrapper operation.

1. Create a **Surface Wrapper** operation and select all parts that represent the complete, closed, manifold body.
   Including part surfaces that are excluded from the wrap, ensures that all relevant wrapped and unwrapped part surfaces are included in the resulting Surface Wrapper part.
2. Select the **Geometry** > **Operations** > **Surface Wrapper** node.
   a) Activate *Perform Partial Wrapping*.
      A **Default Controls** > **Preserved Input Surfaces** node is created.
   b) Define any other *Surface Wrapper-Properties* that you require.
3. Select the **Default Controls** > **Preserved Input Surfaces** node and set *Excluded Surfaces* to all part surfaces that you want to exclude from the wrap.
4. To capture details in areas where wrapped and unwrapped part surfaces share a perimeter, or intersect each other:
   a) Refine the Surface Wrapper operation using one or more of the following methods:
      - set an appropriate base size
      - define custom controls
      - create a contact prevention
      - perform a boolean subtract operation

        This operation maintains existing perimeters–on which you can define a curve control, or contact prevention.
   b) Specify any other Default Controls that are required for the Surface Wrapper operation.
5. Execute the Surface Wrapper operation.
   The original surface is duplicated and the perimeter edges between the wrapped surface and the duplicated original are imprinted together. In areas that are not connected, the Surface Wrapper operation then fills in holes . As long as the surface wrapper has appropriate refinement, the resulting Surface Wrapper part surface is closed and manifold.

Note: The Partial Wrap functionality is a post-processing step and does not change the usual behavior or settings of the Surface Wrapper in any way.

Generating Surface and Volume Mesh

Meshing parts generates the surface and volume discretization of the imported CAD geometry. The solvers use this discretization to obtain a physical solution.

Before Generating Surface and Volume Mesh:

1. Complete Preparing Geometry for Meshing.

   See Preparing Geometry for Meshing.

2. If the CAD geometry is poor quality or complex, surface wrap the geometry.

   See Surface Wrapping the Geometry.

3. If part contacts occur between two part surfaces in the geometry and have not been defined automatically, they must be created.

   a. To browse or find part contacts automatically, multi-select the required parts, right-click, and choose **Browse Part/Part Contacts**, or **Find Part/Part Contacts.**

   b. To create part contacts manually, multi-select and right click the parts in contact, then choose **Create Part/Part Contact**.

      See Generating Part Contact Data.

Meshing aims to generate a conformal mesh at part contacts. Each part contact becomes an interface in the physics region.

To Generate Surface and Volume Mesh:

1. Create an automated mesh operation.

   a. Right-click the **Geometry** > **Operations** manager node and select **New** > **Automated Mesh**.

      The *Create Automated Mesh Operation* dialog appears.

   b. Select the parts to mesh, then activate the required Meshers and click **OK**.

      An unexecuted Automated Mesh node is displayed under the **Geometry** > **Operations**manager.

2. Execute surface meshers.

   a. Right-click the **Operations** > **Automated Mesh** node and select **Execute Surface Meshers**.

   b. To display the surface mesh, right click an empty area within a mesh scene and select **Apply representation** > **Mesh Scene 1**.

See Creating a Mesh Scene

3. Execute volume meshers.
   - Right-click the **Operations** > **Automated Mesh** node and select **Execute**.

Note: If you create, update, or delete part contacts, mesh operations that include these part contacts are also updated. The updated mesh operation icons appear out of date ( ) and must be re-executed.

Selecting Meshers

The selected meshers determine the type of surface and volume discretization that you generate. For example, polyhedral, tetrahedral, or trimmed volume meshes.

Selecting Meshers for a New Automated Mesh Operation

To select meshers for a new automated mesh operation:

1. Create an automated mesh operation.
2. In the *Create Automated Mesh Operation* dialog, select the required meshers.



3. To deselect meshers, deactivate the checkboxes in the *Enabled Meshers* group box.

Setting Default Controls

The default mesh controls define the global surface and volume mesh behavior.

To set the default mesh controls:

1. Select a node under the **Operations** > **Automated Mesh** > **Default Controls** manager and set its properties.



See Also

Extract Volume
Meshing the Surface
Meshing the Volume
Surface Mesher Controls and Values
Volume Mesher Controls and Values

Creating and Defining Custom Controls

Custom controls refine or coarsen the surface and volume mesh for parts, part surfaces, part curves, and specified volumes.

To create and define custom controls:

1. Right-click the **Operations** > **Automated Mesh** > **Custom Controls** node and select a control from the **New** submenu.



2. To define the applied part entities and control behavior, select the **Custom Controls** > **Control** node and set its properties.



3. To activate a control, select a node under the **Controls** manager and set its properties.



4. To define the custom control value, select a node under the **Values** manager and set its properties.



See Also

Surface Mesher Controls and Values
Volume Mesher Controls and Values

Meshing the Surface

Before meshing the volume, generate and check the surface mesh. The volume mesh quality is directly dependent on the surface mesh quality.

To mesh the surface:

1. Right-click the **Operations** > **Automated Mesh** node and select **Execute Surface Meshers**.



See Also

Meshing the Volume
Extract Volume
Surface Meshers

Meshing the Volume

The volume meshers generate the discretized representation of the geometry part that the solvers use to obtain a solution.

Before meshing the volume:

- Assign the parts to regions.

To mesh the volume:

1. Right-click the **Operations** > **Automated Mesh** node and select **Execute**.



See Also

Extract Volume
Surface Meshers

Mesh Refinement in a Specified Volume

STAR-CCM+ allows you to refine surface and volume meshes within a specific volume in order to capture localized effects. You define the specific volume using a geometry part, which you can create or import. The volume can be created to intersect specific volumetric parameters within; geometry parts, fluid domains or both.



To refine the mesh:

1. Prepare the Geometry Part that defines the volume.

   See Preparing the Geometry Part that Defines the Volume.
2. Create a Volumetric Control.

   See Creating a Volumetric Control.
3. Set Volumetric Control Values.

   See Editing Volumetric Control Values.

After running surface or volume meshers, the mesh within the volumetric parameters specified now appears refined.

Preparing the Geometry Part that Defines the Volume

The volume within which the mesh is refined is defined by a geometry part.

It is not necessary for this geometry part to belong to the geometry you are modeling. You can import or create a separate part.

To obtain a part that defines the volume, choose one of the following options:

- Create a simple shape part in STAR-CCM+

  This option is the simplest way to create a block, cone, cylinder or sphere for the volumetric control shape part.

  See Creating Simple Shape Parts.

- Import a CAD part into STAR-CCM+

  This option allows you to use your own CAD part for the volumetric control shape part.

  See Importing Geometry.

- Create a shape part in 3D-CAD and import it to **Geometry > Parts**

  This option is useful if you want to create a more complex shape part for the volumetric control, but do not already have a CAD part to import.

  See Using 3D-CAD and Importing Geometry.

Creating a Volumetric Control

Before Creating a Volumetric Control:

1. Prepare the Geometry Part that Defines the Volume.

   See Preparing the Geometry Part that Defines the Volume.

2. Create an Automated Mesh Operation.

   See Selecting Meshers.

To Create a Volumetric Control:

1. Right-click **Geometry > Operations > Automated Mesh > Custom Controls** and select **New > Volumetric Control**.



2. Select the new **Volumetric Control**node and click the area to the right of **Parts** in the *Volumetric Control - Properties* pane.

   The *Volumetric Control* dialog appears.

3. Select the shape part that defines the volume.



4. Click **OK**

   You now have a Volumetric Control which is ready for you to customize. Until you customize the Volumetric Control, the settings will remain the same as the Automated Mesh which contains the Volumetric Control.

   Proceed to Editing Volumetric Control Values.

Editing Volumetric Control Values

Volumetric Controls expose meshing values when you choose which meshing models to customize.

You can add custom values for the following meshing models:

- Surface Remesher
- Polyhedral Mesher
- Tetrahedral Mesher
- Trimmed Cell Mesher
- Prism Layer Mesher (option only available after selecting a volume mesher)

See Volumetric Controls and Values.

Note: The Thin Mesher, Extruder, Embedded Thin Mesher, and Advancing Layer Mesher are not yet available in Automated Mesh operations. To use these Meshers within a volumetric control use Region Based Meshing. See Setting Up a Volumetric Control.

Before Editing Volumetric Control Values:

Create a Volumetric Control for the Part that Defines the Volume.

See Creating a Volumetric Control

To Edit Volumetric Control Values:

1. Select a **Volumetric Control > Controls > [Mesher]** node and activate the **[Customize]** property in the *[Mesher] - Properties* window.

   This creates a new **Volumetric Control > Values > [Custom]** manager node.

2. Expand each of the **[Custom]** manager nodes to show the options available to customize.

The following nodes are displayed when the Surface Remesher, Polyhedral Mesher and Prism Layer Mesher have all of their customizable control options activated:



3. Select a node under a **[Custom]** manager node to display values you can edit in the *Properties* window.

   For example, selecting the **Number of Prism Layers** node displays the following in the *Number of Prism Layers - Properties* window:

4. Click the box displaying the value, then type a value in the box and press `<return>` to confirm your entry.

Editing Custom Sizes

Some Volumetric control Values have the option to set values as a percentage that is relative to the base size, or absolute values. See Relative and Absolute Sizes.

For example:

If you choose to customize the **Polyhedral Mesher** and select the **Custom Size** node, the**Size type** property can be set to:

- **Relative to base** - which creates a **Custom Size** > **Relative Size** node.
- **Absolute** - which creates a **Custom Size** > **Absolute Size** node.

Creating a Parts-Based Mesh with Baffles

Baffle contacts are zero-thickness surfaces that the meshers maintain as interfaces in the surface and volume meshes.

Defining baffle contacts maintains the shape of the zero-thickness surfaces in the mesh. The meshers assume that the baffles are two-sided interfaces and generate a conformal mesh on either side.



STAR-CCM+ only supports baffles that intersect a fluid volume. STAR-CCM+ does not currently support a line interface between a baffle and an adjacent solid part. Baffles are represented as interfaces within the volume mesh.

To Create a Parts-Based Mesh with Baffles:

1. In 3D-CAD, create solid and sheet bodies for the fluid volume and baffle parts, or import your geometry.

   See Using 3D-CAD or Importing Geometry.
2. If necessary, extract the fluid volume from the geometry.

   See Internal Volume Extraction.
3. Imprint the baffles onto the fluid volume body.



See Imprinting Bodies.

4. Close 3D-CAD.
5. In STAR-CCM+, create geometry parts from the CAD parts. Ensure that **Create Part Contacts from Coincident Entities** is activated.

   See Creating Geometry Parts from 3D-CAD.

   The surface tesselation is conformal around the baffle edges.



6. Create and execute an extract volume operation for all parts.



See Extract Volume and Executing Operations.

7. Create an automated mesh operation for the extract volume part.

   See Creating Mesh Operations and Selecting Meshers.

8. Set the mesh parameters that you require.

   See Setting Default Controls and Creating and Defining Custom Controls.

9. Assign the extract volume part to a region.



1. Set the following options:
   - Create one Region for All Parts
   - Create a Boundary for Each Part
2. Ensure that *Create Interfaces from Contacts* is activated.

See Assigning Parts to Regions.

10. Execute the automated mesh operation.

A conformal mesh is produced.



See Also

Surface Wrapper with Internal Baffles
Solid Bodies in 3D CAD
Sheet Bodies in 3D-CAD

Mesh Operations

Mesh operations perform actions on geometry parts that you can easily repeat whenever their input parts or properties are changed.

The result of a mesh operation is a new geometry part or description, but the input parts are not changed.

The following mesh operations are available:

- Automated Mesh
- Surface Wrapper
- Extract Volume
- Fill Holes
- Unite Parts
- Intersect Parts
- Subtract Parts
- Imprint Parts
- Priority Indexing

The **Operations** node holds all mesh operations that are defined in the simulation.



Creating Mesh Operations

Mesh operations are actions that you can perform on the geometry in a sequential pipeline. This pipeline allows for flexible and repeatable meshing processes.

To create a mesh operation:

1. You can create a mesh operation using two different methods. Either:
   - Right-click the **Operations** node and select the required operation from the **New** submenu.



- Select one or more geometry parts, right-click and select the required operation from the **Create Mesh Operation** submenu.



When selecting the **Extract Volume** operation, if you have a scene open, a panel appears which allows you to define the properties. For all other operation selections, a dialog appears which allows you to select the input parts.

When the mesh operation is created, an operation node and a leaf-level part node appear in the tree.



2. Set the operation properties:

- Select the operation node and in the *Properties* window, set the properties that the operation requires.
- For extract volume operations that are defined using the *Extract Volume* edit panel, you can also edit properties on the Extract Volume node afterwards.

Selecting Input Parts for Mesh Operations

Selecting individual parts

The individual geometry parts that are directly selected, along with part surfaces and part curves that comprise the geometry part, belong to the mesh operation.

Selecting composite parts

All parts that belong to the composite part, along with part surfaces and part curves that comprise the geometry part, are included in the mesh operation, even if they don't appear selected individually.

Executing Operations

You execute operations to perform a specified action on a geometry part. These executed operations take the geometry part from imported CAD to a discretized volume mesh.

When you create an operation, change its properties, or modify the input parts, a warning triangle on the operation and part indicates that the operation needs executing.



The following methods are available for executing operations:

Executing Individual Operations

You can execute individual mesh operations to perform a single action on a geometry part.

Note: Executing an individual operation also executes any out-of-date operations that the selected operation requires as an input.

To execute individual operations:

1. Right-click a mesh operation and select **Execute**.



Executing All Operations

You can execute all operations to run the entire meshing pipeline on the geometry parts.

This action executes the operations sequentially in the order that they appear under the **Operations** manager.

To execute all operations:

1. Right the **Geometry** > **Operations** node and select **Execute All**.



Executing Operations on Creation

When creating **Extract Volume**, **Fill Holes**, or **Boolean** mesh operations, you can activate the option to *Execute Operation on creation*. If this option is activated, the operation is executed when you click **OK**.



Copying Controls Between Mesh Operations

It is possible to copy and paste controls between mesh operations.

Input parts for the controls on the source mesh operation are only copied to the destination controls if the input parts exist in the destination mesh operation (except inherited parts).

Parts from the source mesh operation are converted to the appropriate inherited parts in the destination mesh operation. It is also possible for inherited parts to be converted into their appropriate parent parts, depending on which operations the controls are copied to and from.

To Copy Controls Between Mesh Operations:

1. Click and hold the source mesh operation **Control** node that you want to copy.
2. Drag and drop the **Control** node onto the corresponding **Controls** manager node in the destination mesh operation.



The **Controls** manager node in the destination mesh operation contains a copy of the **Control**. The Part Surface names in the destination mesh operation **Control** reflect the appropriate inherited part name.



See Also

Creating Mesh Operations
Mesh Operations Context Menus

Mesh Operations Context Menus

Mesh Operations Manager

Use the context menu for the mesh operations manager to create and execute mesh operations.

New	Generates a new mesh operation. The following operations are available: • Automated Mesh • Surface Wrapper • Extract Volume • Fill Holes

	• Unite Parts • Intersect Parts • Subtract Parts • Imprint Parts
Execute All	Executes all defined operations in a sequential order from top to bottom.

Operation Nodes

Use the context menu for individual operation nodes to edit and execute the selected operation.

Edit Operation in Current Scene	Opens the appropriate panel so that you can edit the operation properties interactively. The following panels are available: • Extract Volume • Fill Holes
Execute	Performs the mesh operation on the input parts.
Pre-Execute	Identifies the new and inherited parts, without executing the operation.

See Also

Inherited Parts

Mesh Operation Errors

Parts with Errors or Empty Parts

If any of the parts that are selected for an automated mesh operation or surface wrapper operation have errors or contain no surfaces (empty parts), an error message appears in the output window. The names of affected parts are included in the message.

- Multi-Part Meshing
  - If there are multiple parts in the mesh operation and all parts are empty, the following message is displayed:

    ```
    Parts do not contain any faces. Please check surface
    before meshing.
    ```

- If there is only one part in the mesh operation, the following message is displayed:

  ```
  Part: Body 2 does not contain any faces. Please
  check surface before meshing.
  ```

- Per-Part Meshing

  In Per-Part meshing where there are multiple parts and one or more parts are empty, the following message is displayed:

  ```
  WARNING: Part: Body 2 is empty and therefore cannot
  be meshed
  WARNING: Part: Body 3 is empty and therefore cannot
  be meshed
  ```

  Any parts that are not empty are meshed.

Resolving the Empty Faces

After locating the part or parts that are specified in the error message, you can reimport or re-create the mesh, then re-execute the mesh operation.

Extract Volume

The extract volume mesh operation finds enclosed volumes within selected parts and generates a part to fill the volume.

If parts are imprinted before extracting the volume, there are no solid parts to merge, so their tessellated surfaces are merged instead. The tessellated surface allows the volume to be extracted, as long as the combination of merged internal tessellated surfaces are closed.

The **Extract Volume** node represents this operation in the simulation tree.



Extract Volume Pop-up Menu

After executing the Extract Volume operation in automatic mode, when you right-click the **Extract Volume** node, the following options are available from the pop-up menu:

Convert to Manual Mode	The volumes that are extracted automatically are available to modify manually using Surface or Point mode. See Extract Volume Properties. You can then re-execute the Extract Volume operation.
Remove Unused Output Parts	When the extract volume operation is executed, the number of output parts can increase. This option removes all output parts that are left in the tree which are not needed and do not contain any mesh.

This mesh operation is not the same as extracting the volume in 3D-CAD.

You can extract the internal volume using a mesh operation when there are multiple parts that enclose the volume. Extracting the volume in 3D-CAD does not work when multiple parts enclose the volume.

See Also

Mesh Operations Context Menus
Extracting the Internal Volume

Extract Volume Properties

Use the **Extract Volume** properties to define the input parts, volume extraction mode, and output part display options. You can also identify one or more output parts from this operation.

Input Parts	Specifies the geometry parts from which the operation extracts an internal volume.	
Volume Mode	**Automatic**	Creates a **Volumes** sub-node which allows you to specify if a **Single** volume or **All** volumes are automatically extracted. Any surfaces that are internal to the final extract volume parts are marked as baffles.
	Manual	Creates a **Volumes** > **Extract Volume** node which allows you to set the extract volume *Mode* to **Surface** or **Point** which you can define using the **Extract Volume** > **Surface** or **Extract Volume** > **Seed Points** child node that is created.
Combine By Part Surfaces	When activated, any volumes with shared part surfaces are combined.	
Output Parts	The name of the output part that the operation generates (read-only).	

Surfaces Properties

The **Volumes** > **Extract Volume** > **Surfaces** sub-node *Properties* allow you to specify the surfaces:

Surfaces	Specifies the part surface that the **Surface** mode uses to extract an internal volume. For example, you can use the part surfaces from

	hole filling, where the ends of an inner cylinder are capped off.

Seed Points Properties

The **Volumes** > **Extract Volume** > **Seed Points** node itself has no properties although right-clicking this node allows you to create seed points that are displayed as sub-nodes. The **Seed Points** sub-nodes *Properties* allow you to specify the seed point coordinates:

Coordinates	Specifies the Cartesian coordinates of a seed point that the **Point** mode uses to extract an internal volume. These seed point coordinates are relative to the Laboratory coordinate system.

See Also

Mesh Operations
Inherited Parts

Fill Holes

The fill holes mesh operation closes arbitrary-shaped holes with a planar surface.

The **Fill Holes** node represents this operation in the simulation tree.



Example of the Fill Holes Operation

The following geometry has tubular void running through it. For this geometry, it is necessary to cap the end holes before it is possible to extract the internal volume.



Executing the fill holes mesh operation generates an output geometry part that contains two more surfaces — one for each of the holes. You can then use this new part with the original part as input parts, for the extract volume mesh operation.



See Also

Mesh Operations Context Menus

Fill Holes Panel

Use the *Fill Holes* panel to define the geometry parts and location of holes.

Input Parts	Specifies the geometry parts on which the operation fills holes.
End Surfaces	Specifies the location of the holes with part surfaces. Make sure that these part surfaces surround the holes.
End Curves	Specifies the location of holes using part curves. Make sure that these part surfaces surround the holes. It is not necessary to choose the specific part curves that surround the holes — any curves that do not surround holes are automatically ignored.

See Also

Mesh Operations

Fill Holes Properties

Use the **Fill Holes** properties to define the geometry parts and location of holes. You can also identify the output part from this operation.

Input Parts	Specifies the geometry parts on which the operation fills holes.
End Surfaces	Specifies the location of the holes with part surfaces. Make sure that these part surfaces surround the holes.
End Curves	Specifies the location of holes using part curves. Make sure that these part surfaces surround the holes. It is not necessary to choose the specific part curves that surround the holes — any curves that do not surround holes are automatically ignored.
Output Parts	The name of the output part that the operation generates (read-only).

See Also

Mesh Operations

Unite Parts

The unite mesh operation performs a boolean unite on all input parts that are in contact with or overlap each other.

Selecting the **Operations** > **Unite**node allows you to set properties for the unite mesh operation in the properties window.

Example of the Unite Parts Operation

The following diagram illustrates the behavior of the unite parts operation.



See Also

Boolean Mesh Operation Properties
Mesh Operations Context Menus

Intersect Parts

The intersect mesh operation performs a boolean intersect on all input parts that overlap each other.

Selecting the **Operations** > **Intersect** node allows you to set properties for the intersect mesh operation in the properties window.

Example of the Intersect Parts Operation

The following diagram illustrates the behavior of the intersect parts operation.



See Also

Boolean Mesh Operation Properties
Mesh Operations Context Menus

Subtract Parts

The subtract mesh operation removes geometry parts from the target geometry part.

This operation performs geometrical changes (as opposed to topological) to the target part.

Selecting the **Operations** > **Subtract**node allows you to set properties for the subtract mesh operation in the properties window.

Example of the Subtract Parts Operation

The following diagram illustrates the behavior of the subtract parts operation.



See Also

Boolean Mesh Operation Properties
Mesh Operations Context Menus

Imprint Parts

The Imprint operation imprints close surfaces within a user-defined tolerance. Part contacts are created between the surfaces that are coincident and you can choose to create a conformal or non-conformal surface mesh between parts.

Selecting the **Operations** > **Imprint** node allows you to set properties for the imprint operation in the properties window.



Example of the Imprint Parts Operation for a Conformal Mesh

The following geometry shows two parts that have contacting surfaces between each other. To achieve a conformal mesh between the contacting surfaces, before you imprint the parts, set *Resulting Mesh Type* to **Conformal**.



After imprinting, the resultant surfaces are conformal, which means they have the same surface discretization. By default, in-place part contacts are created between the contacting surfaces.



By hiding the cylindrical surface, you can see that the surface discretization on the contacting surfaces is identical.



See Also

Priority Indexing
Mesh Operations Context Menus

Non-Conformal Imprinting

STAR-CCM+ allows you to imprint parts non-conformally. Non-conformal imprinting provides a way to imprint parts efficiently without taking time to imprint CAD parts conformally when it is unnecessary.

This process is useful when:

- a conformal interface is not necessary to achieve solution accuracy
- there are many contacting parts
- each part is within many pairs
- there are gaps between the CAD parts

During the process of non-conformal imprinting, pairs of split surfaces are created on contacting parts without altering each individual part.

The following geometry shows two parts that are in close proximity with a small gap between the surfaces. To achieve a non-conformal mesh when imprinting these surfaces, before you imprint the parts, set *Resulting Mesh Type* to **Non-Conformal.**



After imprinting the parts, the resultant surfaces are non-conformal, which means that the surface discretization on the imprinted surfaces does not match. Each part remains separate and there is no shared surface between each part. However, the original discretized surfaces are cut and a pair of split part surfaces are created which represent an in-place part contact between the surfaces.



The pair of split part surfaces are displayed in the **[Part] > Surfaces** folder:



Taking a closer look at the pair of split surfaces, you can see that the surface discretization on each split part surface is non-conformal:



The in-place part contact which represents the pair of split part surfaces appears in the simulation tree:



Imprint Parts Resulting Part Surfaces

The assignment of imprinted faces to a part surface depends on the **Imprint** > **Controls** > **Resulting Part Surfaces** option that you select.

This option is only available for conformal imprinting as non-conformal imprinting creates new part surfaces by default.

There are two options available:

- Input Part Surfaces
- New Part Surfaces

Input Part Surfaces

When *Resulting Part Surfaces* is set to **Input Part Surfaces**, the imprinted faces that are created are added to the part surface of the input part.

During the imprint operation, part surfaces from source input parts are assigned to the coincident part surfaces on destination input parts. If the part surfaces on the source part line up well with the destination part surfaces, the imprinted faces are assigned to the correct part surface.



If the part surfaces on the source part do not line up well with the destination part surfaces, the imprint operation only preserves part surface distinctions that are on the destination side of the interface.


PS 1
PS 4
PS 2
PS 5
PS 3
Part A
Destination
Part B
Source
PS 1
PS 4
PS 4
or
PS 5
PS 2
PS 3
PS 5
Part A
Destination
Part B
Source



To specify which side of the interface is the destination, use priority indexing. See Priority Indexing.

New Part Surfaces

When *Resulting Part Surfaces* is set to **New Part Surfaces**, the imprinted faces that are created are added to new part surfaces. This behavior also reflects the part surface assignment that occurs by default during non-conformal imprinting.

One new part surface is created for each contiguous component of an imprint in a particular operation. Re-executing the operation reuses the part surfaces that are created from a previous execution whenever possible.



Imprint Properties

Use the **Imprint** properties to define the input parts, imprint tolerance, and whether to create mesh contacts or not.

Input Parts	Specifies the geometry parts that the operation imprints.	
Tolerance	Specifies the upper limit on the distance between part surfaces that the operation imprints. Any part surfaces that are not within this tolerance are not imprinted.	
Merge/Imprint Method	Specifies how faces are chosen and matched for imprinting.	
	Discrete Imprint	The imprint is carried out using the tessellated surface mesh of the parts. The original tesselated faces are cut to ensure

		that the imprinted faces match.
	Exact Face Match	The imprint is carried out using the tessellated surface mesh of the parts. Coincident vertices and faces are merged and a part contact is created between the two parts. Specify this option for part pairs that already have exact one-to-one correspondence between their faces and vertices. See Using the Exact Face Match Method.
Resulting Mesh Type	Determines if the imprinter merges the faces of the input parts.	
	Conformal	The imprinter merges faces of the imprinted parts to produce a conformal mesh on the output surface.
	Non-Conformal	The imprinter does not merge faces of the imprinted parts so a non-conformal mesh is produced on the output surface.

See Also

Mesh Operations

Imprint Controls

Candidate Pairs (available for conformal imprinting only)	Allows you to specify pairs of parts to imprint.

	All Close Pairs	All close pairs of parts/part surfaces are imprinted.
	Use Existing Contacts	Only pairs of parts/part contacts with existing contacts are imprinted.
Preprocess	Allows you to specify if any boolean operations are performed before parts are imprinted.	
	No Preprocess	No boolean operations are performed before parts are imprinted.
	Boolean Subtract	A boolean subtract operation is performed before parts are imprinted.
	Intersect Surfaces (available for non-conformal imprinting only)	An intersect operation is performed to cut the existing surfaces before parts are imprinted. Coincident surfaces that lie inside other surfaces are also cut and automatically included in the resulting part surfaces.
Resulting Part Surfaces (available for conformal imprinting only)	Allows you to specify the surface that the resulting imprinted faces are added to. This option is only available for conformal imprinting as non-conformal imprinting creates new part surfaces by default. See Imprint Parts Resulting Part Surfaces.	
	Input Part Surfaces	The imprinted faces that are created are added to the part surface of the input part.
	New Part Surfaces	The imprinted faces that are created are added to a new part surface.

Face Orientation	Allows you to specify the orientation of the matching face normals.	
	Opposing and Aligned	Target faces that are aligned or opposed outward from the normal of the matched source faces are merged.
	Opposing Outward	Only target faces that are opposed outward from the normal of the matched source faces are merged. See Using the Opposing Outward Face Orientation Option.
	Aligned Outward	Only target faces that are aligned outward from the normal of the matched source faces are merged.
Merge Angle	Specifies the maximum angle between the normals of the faces at which parts merge. Only pairs of faces which lie within the specified *Value* and have normal directions which differ by less than this *Merge Angle* are merged.	

Boolean Mesh Operation Properties

Some boolean mesh operations have common properties which are described in this section.

The properties below apply to the following boolean mesh operations:

- Unite Parts
- Intersect Parts
- Subtract Parts

Properties

Input Parts	Specifies the geometry parts that the boolean operation is performed on.
Target Part (Subtract Parts operation only)	Specifies the single geometry part from which all other input parts are subtracted.
Output Parts	The name of the output parts that the operation generates (read-only).

Expert Properties

Remesh Intersection Curve	When activated, this option locally remeshes faces near the intersection curve in the output parts. This option is deactivated by default. See Intersection Curve Remeshing.

The **Imprint Parts** boolean operation has a unique set of properties which are described in the Imprint Properties section. See Imprint Properties.

See Also

Mesh Operations Context Menus
Inherited Parts

Intersection Curve Remeshing

The *Remesh Intersection Curve* option is available when performing one of the following boolean mesh operations:

- Unite Parts
- Intersect Parts
- Subtract Parts

If the input parts for a mesh operation intersect each other, the mesh operation can produce poor quality faces around the intersection curve area on the surface of the output parts. The following image shows the output from a unite operation with *Remesh Intersection Curve* deactivated:



Activating the *Remesh Intersection Curve* option causes the surfaces of the output parts to be remeshed automaticaly after the mesh operation is executed.The following image shows the output from a unite operation with *Remesh Intersection Curve* activated:



Priority Indexing

When imprinting parts, at least one part is designated as the destination part and any other parts are designated as source parts. Priority

Indexing lets you control which parts are the destination or source parts. The part or parts with the lowest index value have the highest priority. So any parts with the priority index of 1 are the destination parts.

The features and precise location of the destination parts are maintained during the imprint operation. However, the features of the source parts are modified to fit the features of the destination parts.

The imprint operation moves vertices of the source parts onto vertices of the destination parts. Faces from the source part that are coincident with a destination part are deleted and replaced with faces of the destination part.

Imprint priority indexing can also be useful to control the division of shared faces into part surfaces.

Part surfaces of all faces in destination parts are preserved despite being re-tessellated. However, part surfaces in the source parts change since duplicate faces from source parts are deleted during imprinting.

The number of interfaces that exist after imprinting depends on the assignment of priority indexes.

The Effect of Priority Indexing on Part Shapes

When imprinting the parts displayed below, you can specify either the L-shaped part or the block part as the destination part. Setting the priority index number of a part to 1, defines that part as the destination part.



When the priority index number of the block part is set to 1, the sharp corner of the block is preserved during imprinting:



When the priority index number of the L-shaped part is set to 1, the curve of the L-shaped part is preserved during imprinting:



See Also

Imprint Parts

Imprinting Parts With Priority Indexing

Setting a priority index for an imprint operation allows you to have control over imprinting while maintaining the flexibility of performing operations as part of the parts-based meshing pipeline.

Before parts are assigned a priority index number, first create a priority set within the Imprint Operation and assign a part or parts to that priority set. You can then reorder priority sets within the simulation tree to change their priority index number. The higher in the tree a priority index is, the higher the priority it has.

A priority set has the following properties:

- **Parts**

- **Enable Imprint Priority Set**
- A read-only **Index** number



Priority sets are not shared between operations, but you can copy and paste them between operations.

Before creating an imprint operation with a priority index:

1. Import / create parts.
2. If necessary, split parts by patch. See Splitting by Patch.
3. Assign parts to regions. See Assigning Parts to Regions.
   - In the *Assign Parts to Regions* dialog, select:
     - **One region per part**
     - **One boundary per part surface**
     - **One feature curve for all part curves**

   Interfaces are created between part surfaces that are in close proximity.

To create an imprint operation with a priority index:

1. Create an imprint operation.
   a. Right-click **Geometry** > **Operations** and select **New** > **Imprint**.
   b. In the *Create Imprint Operation* dialog, select the parts to include in the imprint operation.
   c. Select the**Imprint** node and set the properties that you require. See Imprint Properties.
   d. Click **OK**.
2. Set the priority index of the parts you have selected to merge.
   a. Right-click **Geometry** > **Operations** > **Imprint** > **Priority Indexing** and select **New**.

      A new **Priority Indexing** > **Priority Set 1** node appears.
   b. Select the **Priority Set 1**node.
   c. In the *Priority Set 1 - Properties* window, set *Parts* to the destination part for the imprint operation.
   d. Repeat the previous step to create additional priority indexes for the source parts if you require them.
3. If necessary, reorder priority sets:
   a. Right-click the **Priority Indexing** folder and select **Reorder**.
   b. In the *Reorder Imprint Priorities* dialog, click and drag the priority sets then release them at a suitable point in the table. The part with Priority 1 is the destination part.



   c. When the priority sets are in the correct order, click **OK**.

Note: It is possible to place a single part, or composites containing it, in multiple priority sets. If conflict occurs between two priority sets that

contain the same part, the priority set with the highest index value is selected first.

4. Execute the **Imprint** Operation.
   a. Right-click the **Operations** > **Imprint** node and select **Execute**.
   b. To display the imprinted parts, right-click an empty space in the *Graphics* window and select **Apply Representation** > **Imprint**.

Note: Priority sets that are deactivated and priority sets that contain no parts are displayed with a gray icon.

See Also

Imprint Parts

Inherited Parts

Inherited parts are the input parts to an operation, which are categorized separately from any new surface and curves that the operation creates. This categorization makes it easier to find and apply the input parts to other operations and mesh controls.

The following operations generate inherited parts:

- Surface Wrapper (with the Per-Part Meshing property deactivated)
- Extract Volume
- Unite Parts
- Intersect Parts
- Subtract Parts

Example of Using Inherited Parts for Mesh Controls

The following example uses the output part from a subtract operation that has several inherited parts. The following image illustrates this example and how the inherited parts are organized.



When using this part as the input part for an automated mesh operation, the categorization of inherited parts makes it easier to find and apply custom controls to the sensor surfaces.



For comparison, the following image shows the categorization for the same subtract part with a flat list of surfaces and curves. This image shows that it is not obvious which parts are associated with the Sensor surfaces and Default curves.



Automated Mesh

The automated mesh operation generates surface and volumes meshes for its input parts. It manages all of the meshers, default controls, and custom controls.

The **Automated Mesh** node represents this operation in the simulation tree.



Create Automated Mesh Operation Dialog

Use the *Create Automated Mesh Operation* dialog to select the input parts and meshers for the automated mesh operation.

Parts	Specifies the geometry parts for which the operation generates a mesh.
Surface Meshers	Provides the available meshers for generating a surface mesh. Currently, the following surface meshers are available: • Surface Remesher
Optional Surface Meshers	Provides the optional meshers for generating a surface mesh. Currently, the following optional surface meshers are available: • Automatic Surface Repair
Volume Meshers	Provides the available meshers for generating a volume mesh. Currently, the following volume meshers are available: • Polyhedral Mesher • Tetrahedral Mesher • Trimmed Cell Mesher
Optional Volume Meshers	Provides the optional meshers for generating a volume mesh. Currently, the following optional volume meshers are available: • Prism Layer Mesher

Enabled Meshers	Shows the meshers that you have selected. Deactivating the check box for any of these meshers removes it from the mesh operation.

See Also

Mesh Operations
Extract Volume

Select Meshers Dialog

Use the *Select Meshers* dialog to redefine the meshers for the automated mesh operation.

Surface Meshers	Provides the available meshers for generating a surface mesh. Currently, the following surface meshers are available: • Surface Remesher
Optional Surface Meshers	Provides the optional meshers for generating a surface mesh. Currently, the following optional surface meshers are available: • Automatic Surface Repair
Volume Meshers	Provides the available meshers for generating a volume mesh. Currently, the following volume meshers are available: • Polyhedral Mesher • Tetrahedral Mesher • Trimmed Cell Mesher
Optional Volume Meshers	Provides the optional meshers for generating a volume mesh. Currently, the following optional volume meshers are available: • Prism Layer Mesher
Enabled Meshers	Shows the meshers that you have selected. Deactivating the check box for any of these meshers removes it from the mesh operation.

See Also

Mesh Operations

Extract Volume

Automated Mesh Properties

Use the **Automated Mesh** properties and expert properties to define the behavior of the operation.

Properties

Per-Part Meshing	**Activated**	Meshes each part individually, which ignores any interfaces between parts. Therefore the mesh between connected parts is not conformal. This option is useful if the trimmer or thin mesher cannot produce a mesh at the interface.
	Deactivated	Meshes all the input parts at once.
Mesher Execution Mode	Allows you to set the execution mode when running parallel servers.	
	Serial	Executes the meshers on the main processor only.
	Parallel	Executes the meshers on parallel processors, depending on how you launch the simulation. See Parallel Meshing. Parts are subdivided between processors.
	Concurrent	Executes the meshers on separate processors for each individual part. Only available to use with the surface remesher, surface wrapper, trimmed, polyhedral, and

		tetrahedral meshers - with or without the prism layer mesher. See Concurrent Meshing.
Input Parts	Specifies the geometry parts for which the operation generates a mesh.	

Expert Properties

Preserve Surface Perimeters	**None**	Maintains only surface that you have assigned to boundaries.
	Major Mesh Controls Surface Perimeters	Maintains the surface perimeters for surfaces to which you have applied custom surface controls for the prism layer mesher only.
	All Mesh Controls Surface Perimeters	Maintains the surface perimeters for all surfaces to which you have applied custom surface controls.
	All Part Surface Perimeters	Maintains all of the surface perimeters in the surface and volume mesh.
Verbose Output	**Activated**	Provides detailed meshing information in the *Output* window. This option is useful for debugging meshing problems.
	Deactivated	Provides the standard meshing information in the *Output* window.

See Also

Mesh Operations
Extract Volume

Automated Mesh Context Menu

Use the **Automate Mesh** context menu to redefine the selected meshers and execute the surface and volume meshers.

Select Meshers	Opens the *Select Meshers* dialog so that you can change the surface and volume meshers that the operation executes.
Execute	Performs the mesh operation on the input parts, which generates a discretized volume mesh.
Execute Surface Meshers	Generates the surface mesh only, so that you can inspect its quality before running the volume meshers.

See Also

Mesh Operations
Extract Volume

Mesh Descriptions

A mesh description is the different versions of a surface mesh that each geometry part has. You create and change descriptions by importing surfaces or executing operations.

The following list summarizes the main types of mesh descriptions:

- Root Descriptions — This description holds the surface mesh from the tessellation of the geometry parts.
- Operation Descriptions — This description holds the surface mesh that was created or changed as a result of a mesh operation.
- Latest Description — This description holds the surface mesh from the most recent mesh operation that you executed on the geometry parts. It appears along with the root description when the part is first imported or created.

The following meshing operations create or change operation descriptions:

- Automated Mesh
- Surface Wrapper (with the Per-Part Meshing property activated)
- Imprint Parts

Whenever you execute a mesh operation on a part, the operation acts on the last description in the tree which contains that part.

If you set a scene to display the latest description, and then you remesh the surface or perform an imprint operation, the scene automatically updates to display the latest description.

Note: Other mesh operations create new parts, not descriptions. If you want to see the result of these operations, modify the visible parts in the current scene.

Visualizing Mesh Descriptions

You can use descriptions to visualize the effect of a mesh operation on a geometry part.

To visualize a mesh description:

1. Make sure that you have a mesh scene open and the geometry part is visible in the scene.

The following image shows the surface tessellation of an imported CAD part. This description is the root description.



2. Drag and drop the required description node into the scene.



The following image shows the imprinted surface mesh.



Exporting Mesh Descriptions

You can export descriptions to the surface mesh in another STAR-CCM+ session or third-party software.

To export a mesh description:

1. Right-click a mesh description and select **Export**.



2. In the *Choose Parts To Export For This Description* dialog, select the parts that you want to export.



3. In the *Save* dialog, set the *File Name* and click **Save**.

Description Properties

The description node properties display information on the selected description.

Described Parts	The geometry parts to which the selected description applies (read-only).
Faces	The number of faces in the surface mesh that the selected description has (read-only).

Vertices	The number of vertices in the surface mesh that the selected description has (read-only).

Description Context Menu

The description context menus can repair and export the description surface mesh.

Repair Surface	Launches the *Surface Repair* tool so that you can manually fix surface mesh issues for the selected description.
Export	Exports the surface mesh for the selected description to one of the following file types: • .dbs — pro-STAR Surface Mesh • .stl — Stereolithography File • .nas — Nastran Shells • .pat — PATRAN Shells • .inp — pro-STAR Shell Input (cell/vertex) File • .fro — FELISA Front Surface • .jt — JT Open

See Also

Preparing Surfaces
Exporting Mesh Descriptions

Mesh Contacts

Mesh contacts define the connection between two part surfaces. The meshers aim to generate a conformal mesh at the contact, which then becomes an interface in the physics region.

When you use contacts with the surface remesher it aims to generate a conformal match between the surface meshes. If the mesher cannot generate a conformal surface mesh, it terminates the meshing process and produces an error message that describes the problem.

A conformal volume mesh between contacts produces a high-quality discretization for the analysis. However, you can produce valid results with a non-conformal mesh, as STAR-CCM+ can still accurately transfer data between the cells.

For the trimmed mesher, a conformal volume mesh is not always possible, as the surface mesh does not directly define the volume mesh cells.

For all other meshers, a conformal, remeshed surface produces a conformal volume mesh. This conformity is inherent as the surface mesh directly defines the shape of the volume mesh cells.

However, a contact between a polyhedral and trimmed mesh is never conformal due to the difference in shape between the two types of cells.

The following types of mesh contacts are available:

- In-Place Contact
- Periodic Contact
- Baffle Contact

The **Contacts** node holds all mesh operations that are defined in the simulation.



Creating Part Contacts

Use mesh contacts to define a connection between two part surfaces.

The meshers aim to generate a conformal mesh at the contact, which then becomes an interface in the physics region.

You can manually create the following types of part contacts:

Creating Weak In-Place Contacts

Weak in-place contacts model two part surfaces that are in direct contact with each, but did not become part contact automatically. These contacts do not create a conformal mesh, but become an in-place, internal interface when assigning parts to regions.

To create a non-conformal in-place contact:

1. Multi-select the contacting part surfaces.
2. Right-click one of the selected part surfaces and select **Create weak in-place contact**.



This action adds a part contact under the **Parts** manager.



See Also

In-Place Contact

Creating Baffle Contacts

Baffle contacts model a zero-thickness surface that the meshers maintain in the surface and volume mesh. When assigning parts to regions, the baffle contact becomes an in-place, baffle interface.

To create a baffle contact:

1. Right-click a part surface and select **Create baffle**.



This action adds a part contact under the **Parts** manager, which is where you can change the baffle contact properties.



See Also

Mesh Contacts
Baffle Contact
Baffle Contact Properties

Baffle Interface

Creating Periodic Contacts

Periodic contacts model two part surfaces that are spatially detached, but a rotation and/or translation can make them coincident. When assigning parts to regions, the baffle contact becomes a periodic, internal interface.

To create a periodic contact:

1. Multi-select the periodic part surface.
2. Right-click one of the selected part surfaces and select **Create Periodic**.



This action adds a part contact under the **Parts** manager, which is where you can change the periodic contact properties.



See Also

Mesh Contacts
Periodic Contact
Periodic Contact Properties
Periodic Interface

In-Place Contact

An in-place contact is a direct connection between two part surfaces.

There are two types of in-place contacts: weak and strong. A strong in-place contact allows you to generate a conformal mesh between the two surfaces; whereas a weak in-place contact does not. Both types of in-place contact become interfaces when you assign the parts to regions.

The following actions create strong in-place contacts:

- Import a surface mesh.
  - Make sure you activate the **Create Part Contact from Coincident Entities** checkbox.
- Create a geometry part from a 3D-CAD model.
  - Make sure you activate the **Create Part Contact from Coincident Entities** checkbox.
- Executing the extract volume or imprint mesh operations.

The following actions create weak in-place contacts:

- Manually creating or finding part/part contacts between parts.
- Manually creating a part contact between two surfaces.

- Executing the surface wrapper operation.

Examples of In-Place Contacts

The following examples show two types of in-place contacts. The yellow, dashed line highlights surfaces where the mesher aims to generate a conformal mesh.

This first diagram illustrates a contact between two part surfaces that match exactly.



This second diagram illustrates an in-place contact between two part surfaces that have different areas. In this case, the contact applies to both surfaces, but the mesher only aims to produce a conformal mesh between the areas of the part surfaces that are coincident.



In-Place Contact Properties

The following table summarizes the properties that are available for in-place contacts.

Tags	Applies a user-defined tag so that you can organize mesh contacts.
Metadata	Specifies the information that you can use to filter the mesh contacts, which are based on properties that are unrelated to the STAR-CCM+ simulation.

See Also

Mesh Contacts
What Are Tags?
What Is Metadata?

Periodic Contact

A periodic contact defines two part surfaces that are spatially detached, but a rotation and/or translation can make them coincident.

When you create a periodic contact, the surface that you select first is the master surface and the other surface is the slave. STAR-CCM+ calculates the rotation and/or translation as relative to the master surface.

To make sure that the mesher can generate a conformal mesh, the size and shape of the periodic surfaces must be identical. If the mesher cannot generate a conformal mesh, the meshers continue, but the *Output* window displays a warning message.

To make sure that the periodic surfaces maintain their original shape, they are excluded from the automatic surface repair operations. Similarly, any manual repairs that you make to the periodic surface after remeshing are not maintained during volume meshing.

Examples of Periodic Contacts

The following examples show two types of periodic contacts. The green and red dashed lines represent the master and slave surfaces respectively.

The first diagram illustrates a rotational periodic contact.



The second diagram illustrates a translational periodic contact.



Periodic Contact Properties

The following table summarizes the properties that are available for periodic contacts.

Tags	Applies a user-defined tag so that you can organize mesh contacts.
Surface 1	Specifies the master surface in the periodic contact, that is, the periodic transform is relative to this surface (read-only). This surface is the one that you selected first when creating the mesh contact.
Surface 2	Specifies the slave surface in the periodic contact (read-only). This surface is the one that you selected second when creating the mesh contact.
Metadata	Specifies the information that you can use to filter the mesh contacts, which are based on properties that are unrelated to the STAR-CCM+ simulation.

Periodic Transformation Properties

The following table summarizes the properties that are available for the **Periodic Transformation** node.

Periodicity	**Rotational**	Specifies that the periodic transformation only requires a rotation angle to match *Surface 1* to *Surface 2*. Therefore the translation is not calculated.
	Rotational +Translational	Specifies that the periodic transformation requires a rotation angle and translation to match *Surface 1* to *Surface 2*.
	Translation	Specifies that the periodic transformation only requires a translation to match *Surface 1* to *Surface 2*. Therefore

		the rotational angle is not calculated.
Axis of rotation specification	**Specify axis**	Specifies that **Axis** node defines the axis of rotation.
Rotation angle	The calculated rotation angle that matches the periodic contact between *Surface 1* to *Surface 2* (read-only).	
Translation	The calculated translation that matches the periodic contact between *Surface 1* to *Surface 2* (read-only).	
Locked	**Activated**	Does not recalculate the rotation angle and translation when you generate the mesh.
	Deactivated	Recalculates the rotation angle and translation when you generate the mesh.

Axis Properties

The following table summarizes the properties that are available for the **Axis** node.

Direction	Specifies the direction of the axis of rotation relative to the user-defined *Coordinate System*.
Coordinate System	Specifies the coordinate system to which all other properties of this node are relative.
Origin	Specifies the origin of the axis of rotation relative to the user-defined *Coordinate System*.

See Also

Mesh Contacts
What Are Tags?
What Is Metadata?

Baffle Contact

A baffle contact is a zero-thickness surface that the meshers maintain in the surface and volume meshes.

Defining a baffle contact maintains the shape of the zero-thickness surfaces in the mesh. The meshers assume that the baffle is a two-sided surface and generate a conformal mesh on either side.

Example of a Baffle Contact

The following example shows guide surfaces in a 90 degree bend. If the thickness of these guides is negligible, you can mesh them as zero-thickness baffles.



Baffle Contact Properties

The following table summarizes the properties that are available for baffle contacts.

Tags	Applies a user-defined tag so that you can organize mesh contacts.
Show baffle proxy surface	Activating this property shows the duplicated surface that defines the other side of zero-thickness baffle. This proxy surface is not

	required for meshing, but is required for modeling physics. This property also allows you to specify the positive direction of the baffle using the **Baffle Orientation Point** node.
Metadata	Specifies the information that you can use to filter the mesh contacts, which are based on properties that are unrelated to the STAR-CCM+ simulation.

Baffle Orientation Point Properties

The following table summarizes the properties that are available for the **Baffle Orientation Point** node.

Orientation Point	Specifies a point to which original baffle surface faces — the other side is the proxy surface. When you assign the part to a region, the original surface becomes a boundary and the proxy surface becomes a copy of that boundary.
Coordinate System	Specifies the coordinate system to which all other properties of this node are relative.

See Also

Mesh Contacts
What Are Tags?
What Is Metadata?

Custom Mesh Controls

Custom controls override any default controls for the surface and volume meshers. This facility allows you to refine or coarsen the mesh for part curves, part surfaces, geometry parts, and specified volumes.

The following types of custom mesh controls are available:

- Curve Controls
- Surface Controls
- Part Controls
- Volumetric Controls

Custom mesh controls are available for the **Surface Wrapper** and **Automated Mesh** operations.



See Also

Creating and Defining Custom Controls

Custom Mesh Control Properties

Use the custom mesh controls to define the input parts, part curves, or part surfaces, display the entities that belong to the control, and deactivate the control.

Parts	Specifies the geometry parts, part curves, or part surfaces to which the mesh control applies.
Part Curves	
Part Surfaces	
Apply Only to contacting area (Surface Control only)	Allows you to apply specific controls on contacting surfaces only. You can use this property to avoid splitting the contacting surface from the default surface. However, the surfaces must form a strong contact with a conformal interface.

Custom Mesh Control Expert Properties

Enable Control	Allows you to activate or deactivate the control without having to delete or recreate it.	
Controls Display Mode	**All**	Displays all parts, part curves, or part surfaces under the **Parts** manager.
	Customized	Displays only the parts, part curves, or part surfaces to which the mesh control applies under the **Parts** manager. This option quickly identifies the entities that the control affects.

See Also

Creating and Defining Custom Controls

Custom Mesh Control Context Menu

Custom Controls Manager

Use the context menu for the custom controls manager to create custom controls.

New	Generates a new mesh control. The following controls are available: • Curve Controls • Surface Controls • Part Controls • Volumetric Controls

Control Nodes

Use the context menu for individual controls to duplicate the selected control.

Duplicate	Creates an exact copy of the selected control.

See Also

Creating and Defining Custom Controls

Curve Control

Curve controls specify alternative surface mesh settings for part curves.

If a part curve belongs to several controls, the following hierarchy determines the curve control that takes precedence. This hierarchy depends on how you apply the part curve to the control.

- Selecting the individual part curves.

  Only the individual part curve that is directly selected belongs to the curve control.

  Any curve controls that are applied directly to an individual part curve take priority over curve controls that are applied to the geometry part or composite part.



- Selecting a geometry part that includes the part curves.

  All part curves that comprise the selected geometry part are included in the curve control, even if they do not appear selected individually.

  Any curve controls that are applied to geometry parts take priority over curve controls that are applied to composite parts.



- Selecting a composite part that includes the part curves.

  All part curves on all geometry parts that comprise the composite part are included in the curve control, even if they do not appear selected individually.

  However, curve controls that are applied to the geometry part or directly to an individual part curve, take priority over curve controls that are applied to the composite part. For composite parts within composite parts, curve controls that are applied to the lowest level composite part take priority.



In addition, if a part curve belongs to several controls at the same level in this hierarchy, the control that produces the smallest cell size takes precedence.

Example of a Curve Control

This example refines the mesh along a feature of interest. The following image shows a part curve that defines the edge of a step.



Refining the surface mesh on the curve by 20% gives a higher density of surface discretization. This dense surface mesh produces a dense volume mesh around the feature.



See Also

Creating and Defining Custom Controls
Custom Mesh Control Context Menu
Custom Mesh Control Properties

Curve Controls and Values

Use the curve controls and values to define the mesh controls that you want to modify.

You can modify the following mesh controls with a curve control:

- Minimum Surface Size
- Target Surface Size

The following section summarizes the properties for all available mesh controls. You can specify these properties as relative or absolute. See Relative or Absolute Sizes.

Curve Control Properties

Control	**Specify Custom Value**	Allows you to specify a value for this control

		that overrides the default control.
	Use Parent Value	Uses the same value as the default control.

See Also

Custom Mesh Controls
Creating and Defining Custom Controls
Surface Mesher Controls and Values

Surface Control

Surface controls specify alternative surface mesh and prism layer settings for part surfaces, geometry parts, and composite parts.

If a part surface belongs to several controls, the following hierarchy determines the surface control that takes precedence. This hierarchy depends on how you apply the part surface to the control.

- Selecting the individual part surfaces.

  Only the individual part surface that is directly selected belongs to the surface control



- Selecting a geometry part that includes the part surfaces.

  All part surfaces that comprise the selected geometry part are included in the surface control, even if they don't appear selected individually.



- Selecting a composite part that includes the part surfaces.

  All part surfaces on all geometry parts that comprise the composite part are included in the surface control, even if they don't appear selected individually.



In addition, if a part surface belongs to several controls at the same level in this hierarchy, the control that produces the smallest cell size takes precedence.

Example of a Surface Control

This example coarsens the mesh in the far field, which provides a more efficient solution by only solving detailed flow in areas of interest. The following image shows the original volume mesh.



Coarsening the surface mesh by five times on the far field surfaces gives a lower mesh density. This coarse surface mesh produces a coarse volume mesh.



See Also

Creating and Defining Custom Controls
Custom Mesh Control Context Menu
Custom Mesh Control Properties

Surface Growth Rate

Surface Growth Rate can be specified on a part or part surface to control the triangle growth rate on those surfaces.

This control allows you to match the surface growth rate on symmetry, inlet, outlet, and far field boundaries with the volume growth rate.

This function is only available in parts-based meshing.

The examples below show a cylinder intersected by a symmetry plane. The *Surface Growth Rate* is based on a curve control from the feature curve where the cylinder meets the symmetry plane.

The following example shows a constant *Surface Growth Rate* of 1.3 on all surfaces:



The example below shows the surface mesh with a *Surface Growth Rate* of 1.1 on the cylinder surface and 2.0 on the plane surface.



Surface Controls and Values

Use the surface controls and values to define the mesh controls that you want to modify.

You can modify the following mesh controls with a surface control:

- Minimum Surface Size
- Prism Layers
- Surface Curvature
- Surface Proximity
- Surface Growth Rate
- Surface Remeshing
- Target Surface Size
- Trimmer Surface Growth Rate

The following section summarizes the properties for all available mesh controls. You can specify these properties as relative or absolute.

Surface Control Properties

Control	**Specify Custom Value**	Allows you to specify a value for this control that overrides the default control.
	Use Parent Value	Uses the same value as the default control.
	Disable	Does not apply the selected mesh control to the surfaces. This option is only available for the Surface Curvature, Surface Proximity, Surface Remeshing, and Prism Layers controls.
	Enable	Applies the selected mesh control to the surfaces. This option is only available for the

		Surface Remeshing control.

See Also

Custom Mesh Controls
Creating and Defining Custom Controls
Surface Mesher Controls and Values

Part Control

Part controls specify alternative volume mesh settings for the polyhedral and tetrahedral meshers.

If a part belongs to several part controls, the control that produces the smallest cell size takes precedence.

Selecting a composite part selects all geometry parts that comprise the composite part, even if they do not appear selected.

Example of a Part Control

This example refines the core volume mesh in the center section of the following geometry.



Increasing the volume mesh density and growth rate increases the number of cells within the core mesh. In this case, the control does not affect the mesh density and growth in the adjacent parts or prism layers.



See Also

Creating and Defining Custom Controls
Custom Mesh Control Context Menu
Custom Mesh Control Properties

Part Controls and Values

Use the part controls and values to modify the mesh density for the polyhedral and tetrahedral meshers.

Part Control Properties

Mesh Density	**Specify Custom Value**	Allows you to specify a value for this control that overrides the default control.

	Use Parent Value	Uses the same value as the default control.

See Also

Custom Mesh Controls
Creating and Defining Custom Controls

Volumetric Control

Volumetric controls specify alternative mesh sizes and prism layer settings for the surface and volume mesher.

Generally, volumetric controls have common properties that specify an alternate face or cell size (see Common Volumetric Control). You can specify the face or cell size as relative or absolute.

However, for the prism layer mesher and trimmed mesher, the behavior of volumetric controls is different (see Prism Layer Mesher Volumetric Control and Trimmed Mesher Volumetric Control).

If a volume belongs to several volumetric controls, the control that produces the smallest cell size takes precedence.

Examples of Volumetric Controls

The following examples show 30% mesh refinement with a cylindrical volumetric control. The image below shows the initial geometry and the cylindrical volumetric control.



Surface Wrapper



Surface Remesher



Polyhedral Mesher



Trimmed Mesher



Tetrahedral Mesher



Prism Layer Mesher

This example is slightly different from the others in that the number of layers and prism layer thickness changes within the volumetric control.



Volumetric Controls and Values

Use the volumetric controls and values to define surface and volume mesh sizes within a specified volume.

Common Volumetric Control

These controls are common to most volumetric controls and specify alternative face or cell sizes for the surface or volume mesh respectively.

Common Volumetric Control Properties

Customize	**Activated**	Allows you to apply an alternative cell or face size.
	Deactivated	Uses the cell or face size that the default controls define.

Prism Layer Mesher Volumetric Control

Use these volumetric controls to specify an alternative number of prism layers, prism layer thickness, and prism layer stretching.

Prism Layer Mesher Volumetric Control Properties

Customize Number of Layers	**Activated**	Allows you to specify an alternative number of prismatic cell layers.
	Deactivated	Uses the number of prismatic cell layers that the default controls define.
Customize Total Thickness	**Activated**	Allows you to specify an alternative total thickness for all cell layers.
	Deactivated	Uses the total thickness that the default controls define.
Customize Stretching	**Activated**	Allows you to specify an alternative stretching value between prismatic cell layers.
	Deactivated	Uses the stretching value that the default controls define.

Trimmed Mesher Volumetric Control

Use this volumetric control to specify an alternative isotropic or anisotropic cell size for the trimmed mesher.

Trimmed Mesher Volumetric Control Properties

Customize isotropic size	**Activated**	Allows you to apply an alternative cell size that is the same in each coordinate direction.
	Deactivated	Uses the cell size that the default controls define.
Customize anisotropic size	**Activated**	Allows you to apply a different cell size in

		each coordinate direction.
	Deactivated	Uses the cell size that the default controls define.

Trimmed Mesher Anisotropic Size Value

Use this value to specify different cell sizes in each coordinate direction for the trimmed mesher. You can specify these cell sizes as relative or absolute. See Relative and Absolute Sizes.

The following image shows an example of anisotropic refinement in the y-direction.



Trimmed Mesher Anisotropic Size Value Properties

Custom X size	**Activated**	Specifies an alternative cell size in the x-direction relative to the trimmed mesher coordinate system.

	Deactivated	Uses the cell size that the custom size value specifies.
Custom Y size	**Activated**	Specifies an alternative cell size in the y-direction relative to the trimmed mesher coordinate system.
	Deactivated	Uses the cell size that the custom size value specifies.
Custom Z size	**Activated**	Specifies an alternative cell size in the z-direction relative to the trimmed mesher coordinate system.
	Deactivated	Uses the cell size that the custom size value specifies.

See Also

Custom Mesh Controls
Creating and Defining Custom Controls
Volume Mesher Controls and Values

Parts Meshes Continua

The parts meshes continua represents the volume mesh that you generated for the part as a continuum that the physics region can use.

Conceptually, the parts meshes continua is a bridge between the parts-based mesh and the physics regions.

The **Parts Meshes** node represents this part-based volume mesh.



Parts Meshes Continua Properties

Use the **Part Meshes** properties to view information on the parts meshes continua and define the method for mapping solution data between different volume meshes.

Properties

Interfaces	The interfaces (or mesh contacts) between parts that are defined (read-only).
Regions	The physics region to which the parts meshes continua applies (read-only).

Expert Properties

Interpolation Option	**Nearest neighbor**	Maps the solution data from one mesh to another by comparing the cell centroids in each mesh. That is, the new cell gets the solution data from an old cell whose cell centroid is closest to

		its own. This option is quick, but does not provide a smooth interpolation of results.
	Higher-order stencil	Maps the solution data from one mesh to another by using a stencil of cells and faces on the original mesh. A smooth interpolation is achieved by using a distance-weighted, least-square interpolation. This option is recommended for transient analyses.

See Also

Extract Volume

Region-Based Meshing

When using region-based meshing, the meshes that you generate are associated with the regions and continua that you have defined, and mesh controls relate to regions and boundaries.

This approach differs to parts-based meshing, where the meshes are associated with the geometry parts. Review the comparison of meshing strategies before deciding on the strategy that is appropriate for you.

Creating a Region-Based Mesh

Before creating a region-based mesh, complete the following prerequisite stages.

- Import or create a geometry as the basis for the computational domain.
- Repair the surface so that it forms a watertight outer boundary to the computational domain, or be prepared to use the surface wrapper when meshing.
- Define the simulation topology by assigning parts to regions and defining the appropriate regions and boundaries.

Creating Mesh Continua

This section outlines how to create mesh continua.

A mesh continuum is required on which to define the appropriate meshing methods.

See Mesh Continua.

Automatically Creating a Mesh Continuum During Import

A mesh continuum is automatically created when a surface mesh is imported.

When a new simulation is created, it does not contain any continua and the **Continuum** node is empty. As soon as a surface mesh is imported (using the **Create New Region** option for *Import Mode*), a mesh continuum is automatically created:



By default, this continuum is called **Mesh 1**, but can be renamed if necessary. All mesh regions that are created during the import process are assigned to this mesh continuum. Importing further surface files for a given simulation does not generate any extra mesh continua but you can manually create a new mesh continuum at any point in the process.

Manually Creating a Mesh Continuum

This section describes how to create a new mesh continuum manually.

A new mesh continuum can be created by right-clicking on the **Continua** node and selecting **New > Mesh Continuum**.



The new continuum is given the next logical sequential name such as **Mesh 2** but can be renamed if necessary. Manager nodes for **Models**, **Reference Values**, and **Volumetric Controls** are added as child nodes to the new continuum node. Once the new mesh continuum has been created, you can proceed to select the meshing models for it. If the mesh continuum is the first that is created, then any null regions that currently exist are automatically assigned to this continuum.

You would typically require a new mesh continuum in any of the following circumstances:

- You want to have a different volume mesh type for each region.
- You are using the trimmed cell mesher for a multi-region case.
- You want to use the surface wrapper for a multi-region case.

In each of the above circumstances, a conformal mesh interface between each region would not be possible during the meshing process.

Assigning Mesh Continua to Regions

Regions are automatically assigned to the default mesh continuum on surface mesh import. They can be reassigned in one of the following ways.

- Setting the continuum for the region using the region properties panel
- Dragging and dropping the continuum onto the region

Setting the Region Mesh Continuum

When the node for a region is selected, the **Properties** window shows the currently selected *Mesh Continuum*.



When more than one mesh continuum exists, the appropriate selection can be made from the menu for the highlighted region. Setting the option to None indicates that no mesh operations are performed for that region.

Dragging a Mesh Continuum Onto a Region

A mesh continuum can be assigned to a region by using the drag-and-drop feature in STAR-CCM+ to drag the mesh continuum onto the region.

Using Per-Region Meshing

This section describes how to activate per-region meshing.

By default, each mesh continuum applies its meshing models to all regions associated with the continuum, in a single operation. In general, this means that the meshing models account for any interfaces that exist between regions and create conformal mesh at the interfaces.

Meshing models that support conformal meshes across interfaces are:

- surface remesher
- tetrahedral mesher
- polyhedral mesher
- prism layer mesher

Meshing models that do not support conformal meshes are:

- surface wrapper
- trimmed cell mesher
- thin mesher

Specific situations where per-region meshing is essential are:

- When using a mesher that does not support conformal interfaces on a multi-region case.

  Meshers in this category are:

  - trimmed-cell mesher
  - thin mesher
  - surface wrapper
- When using overset meshes, and you want to mesh regions that occupy the same physical space using a single mesh continuum.
- For conjugate heat transfer cases where you want to have different cell densities on either side of the solid/fluid interface. The interface itself can be defined as a mapped interface, in order to improve the interpolation of solved quantities across the interface.

To activate per-region meshing:

1. Select the relevant mesh continuum node:



By default, the *Per-Region Meshing* option is turned off, meaning that any models for the continuum behave as described above.

2. In the *Properties* window, activate the **Per-Region Meshing** option:



The Mesh Conditions manager node is deactivated on interfaces when per-region meshing is activated.

The extruder mesher is not influenced in any way by the per-region meshing option.

Selecting Meshing Models

A meshing model in STAR-CCM+ determines the type of mesh that is used within a given mesh continuum. Models can be selected individually or together.

There are both surface meshing models and volume meshing models. Models can be selected individually or together depending on the desired action. Meshing model selections are executed using the (**Generate Surface Mesh**) and (**Generate Volume Mesh**) buttons in the **Mesh Generation** toolbar. Pressing the **Generate Volume Mesh** button also completes any selected surface meshing operations that

have not been performed so far, or that are out-of-date due to parameters that have been modified.

A continuum must contain at least one mesh model in order to be valid. If more than one meshing model is selected, then each operation is carried out in turn during the meshing process. The pipeline sequence in which meshing operations are carried out for a continuum is as follows:

1. Surface wrapper
2. Surface remesher
3. Volume mesher, which can be the:
   - Polyhedral Mesher
   - Tetrahedral Mesher
   - Trimmer Mesher
   - Thin Mesher
   - Advancing Layer Mesher
4. Generalized cylinder and/or embedded thin mesher (for polyhedral)
5. Prism layer mesher
6. Extrusion mesher

The Thin Mesher is a stand-alone volume mesher. If the Thin Mesher is activated, the pipeline can only be composed of the surface meshers plus the selected volume mesher.

For a given continuum, models are selected and changed using the mesh model selection dialog, which assists you with making the appropriate combination of models.

Using the Mesh Model Selection Dialog

The *Meshing Model Selection* dialog is used to select mesh models and/or change the existing model selection for a simulation.

To activate the dialog for a continuum, open the **Continua** node and right-click on the mesh continuum node name (say **Mesh 1**). In the context menu, select the item **Select Meshing Models...**



This same menu item is available in the pop-up menu of the **Models** manager node under the mesh continuum node.

The *Meshing Model Selection* dialog that appears can be used to either select new mesh models, review the existing models, or change the model selection.

Meshing model selection and deletion can only be done through this dialog:

- The dialog is similar to a wizard in that it can assist you through a fairly complex process selection network.
- It is unlike a wizard in that it does not have the traditional **Next** and **Back** buttons.
- It does not restrict you to answering questions in a particular order.

Selecting Mesh Models in the Dialog

The *Mesh 1 Model Selection* dialog appears as shown below when first opened for a mesh continuum that has not yet been defined.



The *Enabled Models* group box on the right-hand side of the panel is empty, indicating that no models have been selected. Two types of models can be selected, namely *Surface Mesh* and *Volume Mesh*. The **<Optional>** labels indicate that the selection of these items is not mandatory.

For example, if you wanted to select the surface meshing models, namely the surface remesher and surface wrapper, then click the checkbox next to each. The models are then activated and listed in the *Enabled Meshing models* box.



The *Mesh 1 Model Selection* dialog changes dynamically in response to the selections made. For instance, if the **Polyhedral Mesher** model option is selected, then the optional **Prism Layer Mesher**, **Embedded Thin Mesher**, **Generalized Cylinder**, and **Extruder** models then become available.



The embedded thin mesher and generalized cylinder meshers are specialized options that are only available when the polyhedral volume meshing option has been selected. These two options are not available for the trimmer and tetrahedral volume meshers, although a tetrahedral mesh option can be selected afterwards using the properties node for the thin mesher model.

Make the appropriate model selection for the continuum before pressing the **Close** button to complete the model setup. For regular use, the best practice is to follow a particular sequence when selecting your models.

The dialog is “live” in that the server is updated each time that you change the state of a checkbox or radio button. Changes are also reflected in the tree when you have the relevant nodes open.

Reviewing Selected Mesh Models

The *Meshing Model Selection* dialog can be used to view the models that have been activated in a continuum. If the dialog has been closed, it can be reopened by right-clicking on either the **Mesh** node or the **Models** node and choosing the **Select Meshing Models...** item in the context menu.

As shown below, the **<Optional>** models group box on the left-hand side of the dialog lists the models that have not been selected. The *Enabled Meshing models* group box on the right-hand side shows all the meshing models that have been activated. The options that are grayed out are required for other (non-grayed out) models and therefore cannot be deactivated without first deselecting the appropriate model.



Changing Mesh Model Selection

The meshing models that are selected for a continuum can be modified in order to change the characteristics of the resulting meshing during the mesh generation process. Optional models in the group box on the left-hand side can be activated at any time before mesh generation; conversely, activated models can be deselected and new optional models can be reselected instead.

Viewing the Mesh Models Using the Object Tree

The meshing models that have been activated for a given mesh continuum can be viewed in the object tree by opening the **Models** node.

The tree view expands as shown below:



Setting Mesh Model Properties

This section describes how to set the properties of a mesh model

To set the properties of a mesh model:

1. Open the **Models** node in the mesh continuum of interest and then click the desired model name.

The *Properties* of the particular model are accessible for viewing and editing.



In general, the default properties for each model are suited for most applications although some adjustment can be necessary for some examples depending on the complexity of the geometry. Additional

details on recommended changes to the properties are discussed in the relevant sections on surface and volume meshing.

For more information, see Surface Meshers Overview and Volume Meshers Overview.

Setting Mesh Controls and Values

This section describes how to set mesh controls and values.

In addition to the reference values on a mesh continuum, mesh values can also be set on the following simulation objects:

- Regions
- Boundaries
- Interfaces
- Feature Curves
- Volumetric Controls

Not all mesh values can be set on all of these objects. The particular values that can be set, and the objects on which they can be set, are determined by the meshing model.

In the absence of any custom mesh value, the reference value on the mesh continuum is used to construct the mesh. Custom mesh values override continuum mesh values according to the following order:

- Mesh values that are set on regions override equivalent values that are set on continua.
- Mesh values that are set on boundaries override equivalent values that are set on regions or continua.
- Mesh values that are set on interfaces override equivalent values that are set on the corresponding boundaries.
- Mesh values that are set on feature curves override equivalent values that are set on boundaries.

Mesh values that are set on volumetric controls also override equivalent values that set elsewhere, but the behavior in relation to boundaries depends on whether a surface or volume mesh is affected.

The general procedure for applying a custom mesh condition is as follows:

1. Expand the **Mesh Conditions** node of the relevant simulation object, for example, a boundary.

2. Some customization nodes have a drop-down menu of options; others have a checkbox. Make the required selection or activate the checkbox.

   In the example below, the option to **Specify Custom Values** for surface curvature is chosen.



3. If the customization requires any input, the **Mesh Values** folder is added to the parent object, and appropriate mesh value nodes are shown within it. Set the necessary inputs on the mesh value nodes.



Additional details on custom settings are discussed in the relevant sections on surface and volume meshing.

Initializing the Mesh

This section describes how to initialize the mesh.

Meshing initialization is required in two general situations:

- Whenever geometry parts are used to provide the starting surface for the meshing models. See What is a Geometry Part?.
- If design parameters relating to 3D-CAD leaf-level parts are changed, requiring a parts update to be performed, and in-turn, a reinitialization of the starting surface. See What Are Leaf-Level Parts? and Updating 3D-CAD Parts.

The option is available in the meshing toolbar (**Initialize Meshing** button) or by using the **Mesh** > **Initialize Meshing** menu item:



Before meshing initialization takes place, the following should have been completed:

- At least one mesh continuum model has been created.
- One or more leaf-level parts have been assigned to regions.

The meshing initialization option is available as soon as a mesh continuum has been created. However, using the option at this stage results in no initialization being performed. Only when at least one leaf-level part has been assigned to a region can meshing initialization proceed and result in the generation of an **Initial Surface** representation.

Clicking the initial surface node provides a summary of the surface:



The initial surface is the surface representation of whatever leaf-level part surfaces and curves that have currently been assigned to regions. If you change the part surfaces and/or curves that are associated with the region boundaries, then perform the initialize meshing operation again in order to update the meshing pipeline. The link between a part and a region is broken if the part is deleted. Since the initial surface representation is associated with a region, deleting a part does not delete the existing mesh for the region. If you use the initialize meshing option in this case, STAR-CCM+ does not clear the mesh.

If you forget to perform the meshing initialization step before surface or volume meshing, then is automatically done for you based on the current surface and curve assignments. Additionally, if design parameters relating to 3D-CAD parts have been changed since the last surface or volume mesh operation then the initial surface is automatically updated when reinvoking the meshing pipeline.

See Also

Understanding Geometry Parts

Generating the Surface Mesh

Once the surface mesh models and properties have been set up, the surface mesh generator can be launched.

To start the process:

1. Click the (**Generate Surface Mesh**) button in the toolbar or use the **Mesh** > **Generate Surface Mesh** menu item.



The text output from the selected meshing models is shown in the *Output* window. This output includes information on the time it took for each stage of the surface meshing pipeline (for example, surface remeshing), as well as the overall time for the whole process. The surface mesh generation process continues until the completed surface is formed or the mesh generation task is aborted.

Once the process is complete, the surface mesh representation can be visualized on the plot display in the same way as other mesh representations.

Generating the Volume Mesh

Once the volume mesh models and properties have been set up, the volume mesh generator can be launched.

To start the process:

1. Click the (**Generate Volume Mesh**) button in the meshing toolbar or use the **Mesh** > **Generate Volume Mesh** menu item.



The selected meshing models generate output that is shown in the *Output* window, and the progress of meshing is shown in the status bar. The output includes information on the time it took for each stage of the volume meshing pipeline to complete, as well as the overall time for the whole process.

The volume mesh generation process continues until the completed volume is formed or the mesh generation task is aborted.

Clearing the Mesh

This section describes how to clear the mesh.

Each surface and volume mesh representation that STAR-CCM+ generates is saved automatically as part of the simulation. When generating a new mesh, the old representation is overwritten as part of the process. Some of the reasons for wanting to clear the generated meshes and representations are:

- to reduce the size of the simulation file
- to start the meshing process over
- to modify the boundaries or patch definitions in the import surface

You can clear all the generated meshes by pressing the (**Clear Generated Meshes**) button in the meshing toolbar or use the **Mesh** > **Clear Generated Meshes** menu item.



The mesh representations nodes are also deleted and any scene displays revert to the import, geometry, or initial representation, which are unaffected by this action.

Replacing a Mesh Region

A region containing a volume mesh in STAR-CCM+ can be replaced by another topologically identical volume mesh at any point during the analysis.

One of the main reasons to replace a mesh is to improve the solution integrity and convergence characteristics by replacing a coarser mesh with a finer one in the middle of the analysis. The current solution is automatically mapped to the new mesh region in order to improve the restart conditions. The mesh replacement option differs from deleting and reimporting a new mesh with respect to this last point; no solution mapping is done in the latter case.

The assumptions that are made when replacing a mesh region with a new one are:

- the region topology is identical in both cases

- the number of boundaries are the same and use the same source names

The grid structure of the two regions does not have to be the same. In other words, a polyhedral type mesh can replace a tetrahedral one, or any other combination of cell shapes allowed.

Before replacing a mesh region, import the new volume mesh region into STAR-CCM+ and configure the boundaries to have the same number and source names as the original mesh region.

To replace a mesh region:

1. Open the region node of interest and right-click on the name. In the pop-up menu that appears, select the **Replace Mesh...** option.



The *Replace Volume Mesh* dialog appears.



2. Two options are provided for interpolation:

   - **Use Nearest Neighbor for interpolation**

     The cell/face value on the new mesh is assigned the value of the cell/face on the old mesh whose cell/face centroid is nearest to the cell/face centroid of the new mesh. This option does not provide smooth interpolations of the results from one mesh to the other.

   - **Use higher-order stencil for interpolation**

     The cell/face value on the new mesh is smoothly interpolated from a stencil of cells/faces on the old mesh using the distance-weighted, least square interpolation method. Generally speaking, the method would be preferred when replacing the mesh at some point during a transient analysis. It can also be useful for steady problems, but the requirement for smoothness is not as critical, as the simulation can continue to converge to a steady state solution after the mesh replacement. The technique is slightly more computationally expensive than the nearest neighbor mapping, but it can offer more fidelity when mapping with thin parts.

3. Using the pull down menus, select the region to replace and then the region that contains the new mesh. Press the **OK** button.

   The progress bar notifies you once the process is complete and a summary message is displayed in the *Output* window:



Solution data in the old region is mapped to the new one, and the old region is deleted. The simulation tree shows the updated region structure.



The physics and boundary conditions from the original region is copied across to the region that replaces it.

Feature Curves

Feature curves are an important part of the surface definition in order to produce a high-quality mesh (both surface and volume).

As part of the meshing process they are used to define edges and surface features from the surface geometry all the way through to the final volume mesh. If they are not included in the starting surface, then such features are not preserved. Feature curves can also be used to specify custom sizes on edges and features for the surface wrapper, surface remesher, and trimmer meshing models.

Feature curves can be imported from third party pre-processing packages into STAR-CCM+ using any of the allowable surface file formats other than STL and PATRAN. They can also be created automatically during the import of neutral data files such as IGES, STEP, and Parasolid Transmit and native CAD files such as PRO/E, CATIA, and SolidWorks.

Feature curves can be computed automatically after surface import based on a number of criteria. They can also be manually added and edited. Similar to other objects, feature curves can be combined, split, renamed and deleted.

What Are Feature Curves?

A feature curve in its most basic form is a two-dimensional cell with a vertex at each end.

They are exclusively used in surface meshes to indicate where sharp edges exist, and to define surface details to maintain in the final volume mesh.

When imported or created in STAR-CCM+, feature curves are listed under the **Feature Curves** node for each region. An example showing two sets of feature curves is shown below:



If the surface remesher is used to improve the starting surface quality, then including feature curves is critical to preserve sharp edges and/or surface detail (such as imprints). If the feature curves are not included, then edges are rounded off and surface detail is lost.

Including feature curves for edges when utilizing the surface wrapper is not as critical as the wrapper feature angle, which can be used to extract features to maintain. However, for features that have an angle less than the angle specified, or for surface details such as imprints, then it is recommended that feature curves be used to preserve the definition.

You can have as many feature curve entries as you want for the surface. However, it is typical that only the critical edges (say features over 45 degrees) are included in feature curves to limit the complexity of the model.

Computing Feature Curves

Feature curves can be automatically created for a given set of boundaries by marking edges using different criteria.

Curves can be created for a single boundary or multiple boundaries simultaneously. Thereafter, further feature curves can be manually added or the original set can be edited so that the exact edge curve definition is obtained.

Before feature curves can be created, a surface mesh type representation (such as import or initial surface) has to exist.

To create a new set of feature curves:

1. Right-click on the **Feature Curves** node for the desired region and select **New Feature Curve > Mark Edges...** from the pop-up menu:



This action opens the *Mark Feature Edge Options* dialog, which can be used in the same way as the object selection dialog to select the boundaries on which to mark feature edges:



2. First, select the boundaries that require feature curves.

   The *Feature Edge Options* criteria can then be reviewed and the appropriate selections for the given boundary set made. The following options are available:

   - **Mark sharp edges** - Allows feature curves to be created on edges whose attached faces make an angle that is sharper than the supplied Sharp edge angle value (31 degrees by default)
   - **Mark free edges** - Allows feature curves to be created for any free edges

- **Mark non-manifold edges** - Allows feature curves to be created for any non-manifold edges
- **Mark patch perimeters** - Allows feature curves to be created based on the underlying imported patch definition
- **Mark boundary perimeters** - Allows feature curves to be created based on the perimeters of each of the currently selected boundaries
- **Automatically Fix Feature Errors** - Fixes all feature errors in explicit feature edges. See Feature Diagnostics
- **Remark existing edges** - Moves any previously marked edges that fit the criterion into a new feature curve. When deactivated, any existing edges remain in their current feature curve

3. To mark feature edges, press the **Apply** button.

   The feature curve node name for the feature curves that are created is based on the method used. For example, if the **Mark sharp edges** option has been selected with the default angle criteria, then the feature curve node name would be **Sharp Edges (31 deg.)**. When more than one selection criteria are used, the generic node name **Feature Curve** is used.

   An example entry is shown below:



4. If you do not want to create feature curves automatically and prefer to add them manually, right-click on the **Feature Curves** node and select **New Feature Curve** > **Empty** from the context menu.

   This selection is shown below:



A new feature curve node (with the generic name **Feature Curve**) is created to which new feature curves can be added by hand.

Adding New Feature Curves

Feature curves can be added by hand to an existing feature curve node.

The node can be empty or contain feature curves that were previously imported or computed. Before feature curves can be created, a surface mesh type representation (such as import or initial surface) has to exist.

To add to an existing feature curve node:

1. Right-click and select the **Add New Edges in Current Scene** option in the pop-up menu, as shown below:



The *Add Edges to Feature Curve edit* tab appears in the explorer pane:



The existing feature curves belonging to the curve set are highlighted in yellow for easier viewing, as shown below:



2. To identify the feature curve start, click a vertex (a red ball appears) and then click a second vertex to identify the feature curve end (the red ball moves to this second location).

   By default the shortest path along the surface mesh is followed to create the feature curve. An example is shown below:



3. If you do not want the feature curve to continue from the last vertex that you selected, deactivate the **Continue from Last** option. You then have to select two new vertices for each set of feature curves you want to create.
4. If you want the feature curve to follow an existing surface feature edge, then turn on the **Follow Features** option and supply an edge **Angle** value (in degrees).
5. Once you have completed adding in feature curves, press the **Close** button to close the dialog and return to the normal scene view.

Any new feature curves that are created during the session are added to the existing feature curve node and displayer. Feature curves can only be added to a topologically connected surface: they cannot be made to span across one surface that is completely disconnected from another one.

Editing Feature Curves in the Current Scene

This section describes how to edit feature curves.

Feature curves can be edited by hand in the following ways:

- delete a part of or the whole of a feature curve

- assign a part of or the whole of a feature curve to a new feature curve node

Before feature curves can be created, a surface mesh type representation (such as import or initial surface) has to exist.

To edit an existing feature curve node:

1. Right-click and select the **Edit Existing Edges in Current Scene** option in the pop-up menu, as shown below:



The *Edit Feature Curve* dialog appears in the *edit* tab:



The existing feature curves in the feature curve set are highlighted in yellow for easier viewing. By default, clicking a highlighted feature curve only selects a single edge. An example is shown below:



2. If you want to select a set of continuous feature curves, then use the **Select Connected Edges** option, which allows features curves that share a common vertex to be selected in one go.

   An example is shown below:



3. When using the connected edges mode, you can specify further criteria to determine where to break the edge selection.

   The two options that are provided are:

   - **Break at Edge Junctions** - stop the selection at where three or more feature curves exist
   - **Break at Sharp Angles** - stop the selection at the specified **Angle** along the curve

4. The **Select Zone** and **Deselect Zone** buttons can also be used to select and unselect feature curves respectively, using a polygonal zone.
5. To clear an existing selection from any of the above described methods, press the **Clear Selection** button.
6. Once a set of feature curves has been selected, you can delete them by pressing the **Delete Selected** button or assign them to a new curve node by pressing **New Curve from Selected**.

   The new curve node name is based on the original one with an appropriate numerical suffix. These curves are colored red and can no longer be selected in the current operation.

7. Press **Close** to complete the process and return to the normal scene view.

Splitting Feature Curves

Before feature curves can be split, a surface mesh representation (such as import or initial surface) has to exist.

To split non-contiguous feature curves:

1. Right-click on the feature curve name and select **Split Non-Contiguous** from the pop-up menu:



More than one feature curve can be split (using multiple selection) at the same time.

2. If you want to split feature curves in some other way, then use the edit feature curves option instead, and manually select the curves. See Editing Feature Curves in the Current Scene.

Combining Feature Curves

Two or more feature curves in the same region can be combined into one in order to simplify the meshing setup.

Before feature curves can be combined, a surface mesh representation (such as import or initial surface) has to exist.

To combine the feature curves:

1. Select the names in the simulation tree (using multiple selection) and right-click. In the pop-up menu, select **Combine**, as shown below:



The name that is used for the combined feature curve is the first one selected.

Using Feature Curves to Repair a Surface

Two different tools are available to help with basic surface repair within STAR-CCM+, both of which utilize feature curves.

The two surface repair options are:

- filling holes
- zipping edges

The hole filler can be used to close any closed loop (both single and multiple loops) planar hole within a surface. The edge zipper can be used to join two mismatched surface edges that are within close proximity to one another, so that a closed surface is formed.

A surface mesh representation (such as import or initial surface) has to exist before feature curves can be used to fill holes or zip edges.

Filling Holes in the Surface

The hole filler can be used to close any arbitrary shaped hole using a planar surface.

Two different types of holes can be filled:

- Simple holes with a single connected loop of free edges
- Multi-loop holes with one loop inside the other

To fill a hole, first define feature curves in a continuous closed loop around the edge of the hole. The easiest way of doing to define the features is to use the **Mark Free Edges** option when computing feature curves, although they can be added manually if preferred. Only include the feature curves relating to the hole as all feature curves are used for hole filling, not just those features that are defined on a free edge. Use the non-contiguous split option to separate the feature curve sets beforehand if necessary or the edit tools to clean up the definition.

An example of a simple hole with the appropriate feature curve set is shown below:



To fill a hole:

1. Select the feature curve node name and right-click. In the pop-up menu select the **Fill All Holes** option:



The *Hole Filling Options* dialog appears:



2. The **Fill Mode** option determines the method that is used to fill multiple holes.

   Two options are provided:

   - **Fill each hole separately** - Treats each connected loop of free edges as a separate (simple) hole.
   - **Fill all holes together** - Looks at all the loops that are selected and fills in the void between them.

   The default operation is to treat each loop as a separate hole and fill it without reference to any other loops that have been selected. Each hole is filled with a planar surface using the least number of

triangular faces possible. An example result of filling a simple hole is shown here.



The output for a successful simple hole fill is shown here:



A new boundary representing the filled hole is also created in the simulation tree and a new displayer containing the object is added to the current scene. The default name of the new boundary is **Filled Holes**.



The use of the hole filler usually results in poor quality triangles being produced in the area of the gap. For this reason, it is recommended that you apply the manual surface remesher to the new boundary faces in order to improve the triangulation quality before generating a volume mesh.

In instances when the feature curves do not form a closed loop or have extra feature curves intersecting the loop, the hole filler fails and no surface is produced. An example output showing when the hole filler is unsuccessful is shown here:



3. In this instance, use the manual add and edit tools to clean up the feature curve definition so that a closed loop is created and then run the hole filler again.
4. To fill more complex holes (for example, holes that have a loop within a loop), first ensure that a connected loop of feature edges exists for each side of the hole.

   An example is shown below:



   Each loop of feature edges does not need to belong to the same node.
5. To fill the hole, select the feature curve nodes and right-click. In the context menu select the **Fill All Holes** option.

   The *Hole Filling Options* dialog appears:



6. In the dialog, select **Fill all holes together** for the **Fill Mode**, so that the gap between the loops is filled rather than each loop separately (which would result in overlapping surfaces).

An example result of filling a multi-loop hole is shown below:



As for simple holes, a new boundary node that is called **Filled Holes** is created containing the surface.

Zipping Edges

The edge zipper can be used to join mismatched surfaces together to form one continuous surface.

It does this by retriangulating the faces at the edges and adding new vertices so that a one to one connection is obtained.

Before using the edge zipper, make sure that feature curves are defined along both sets of edges that you wish to join. The easiest way of marking the edges is by using the **Mark Free Edges** option when computing feature curves, although they can be added manually if preferred. The feature curves used for the zipping process can belong to the same node or to different nodes, but have to belong to the same region. Extra feature curves not belonging to the free edges can also be included but they have no effect on the zipping process.

An example geometry with mismatched surfaces and the appropriate feature curve set is shown below:



To run the edge zipper:

1. Select the feature curve node name or names (using multiple selection) and right-click. In the pop-up menu select the **Zip Edges** option:



The *Edge Zipping Options* dialog appears:



The **Zip tolerance** can be set to determine the tolerance for locating vertices to form the join.

2. To perform the edge zipping operation, press the **OK** button.

   If new faces are required in order to zip the two surfaces, then they are added to the appropriate boundary node in the surface.

   Feature curves can also be deleted depending on the zipping triangulation.

   An example of a zipped surface is shown here.



The output from a successful edge zipping operation appears as follows:



The edge zipper often produces poor quality triangles being along the length of the surface mismatch. For this reason, it is recommended that you apply the surface remesher to the closed surface in order to improve the triangulation quality before generating a volume mesh.

Volumetric Controls

You use a volumetric control to specify the mesh density in a specific zone for both surface and volume meshes.

You can define the mesh refinement zone using volume shapes and geometry parts.

You can apply each volumetric control to any combination of meshing models. Therefore, you can set specific cell sizes within the zone for each mesh generation stage.

Volumetric controls can overlap and extend outside the region boundary definition. Volumetric controls can also overlap from one region to another, but the effect is only included if the region belongs to the same mesh continuum as the volumetric control. If two or more volumetric controls overlap, the smallest user-defined cell size takes priority.

The relevant growth rate for each core mesh model determines the transition in cell sizes from the volumetric control to the core mesh.

The following image illustrates volumetric controls. The blue cells are within a block-shaped volumetric control which specifies a cell size that is 25% of the base size. The red cells are within a cylindrical volumetric control with a refinement to 10% of the base size. The image also shows how growth rates influence cell size transition.



Note: Do not use the same geometry part for the volumetric control that you use as the bounding geometry for the meshing continuum. To make the volumetric control influence the boundaries, extend it a small distance beyond the original geometry.

Volumetric controls affect each meshing model in a different way:

- Surface Wrapper — Refines the surface size. If you use the wrapper scale factor, the surface size within the volumetric control is also scaled.
- Surface Remesher — Refines the surface size.
- Polyhedral Mesher — Polyhedral cell size is uniform within the volumetric control.
- Trimmed Mesher — Refines with either an isotropic or anisotropic cell size. If you specify both options, the smallest cell size in each coordinate direction takes priority.
- Tetrahedral Mesher — Tetrahedral cell size is uniform within the volumetric control.
- Prism Layer Mesher — Refines the number of prism layers, prism layer stretching, and prism layer thickness. To include the prism

layer, the extent of the volumetric control must exceed the boundary.

- Advancing Layer Mesher — Refines the polyhedral core cell size and the prismatic layer cell size. However, you cannot refine the number of layers, stretching factor or layer thickness.

The **Volumetric Controls** manager node, which has its own context menu, contains the volume control nodes for the mesh continuum.



See Also

Generalized Cylinder Mesher

Examples of Volumetric Control Refinement

The following examples show 30% mesh refinement with a cylindrical volumetric control.

The image below shows the initial geometry and the cylindrical volumetric control.



Surface Wrapper



Surface Remesher



Polyhedral Mesher



Trimmed Mesher



Tetrahedral Mesher



Prism Layer Mesher

This example is slightly different from the others in that the number of layers and prism layer thickness changes within the volumetric control.



Advancing Layer Mesher

With the Advancing Layer Mesher, a volumetric control can only change the cell size. The number of layers and layer thickness are not available.



Effect of Volumetric Control on Surface and Volume Meshers

The volumetric control is activated individually for each meshing model. This functionality can cause unexpected results if the volumetric control is activated on the volume mesher, but not the surface mesher.

For example, the following image shows the resultant mesh when a volumetric control is activated on the polyhedral mesher, but not the surface remesher.



In this case, the volume mesh is coarse near the surface, as the volumetric control does not refine the surface mesh.

Activating the volumetric control for both the polyhedral mesher and the surface remesher generates the following mesh.



This image shows that mesh refinement is observed up to and including the surface mesh.

Setting Up a Volumetric Control

A volumetric control refines cells within a user-defined volume shape or geometry part. You can set this refinement using either an absolute cell size or relative to the mesh base size.

You can create a volumetric control after importing a surface mesh as a region or creating a mesh continuum.

Note: A volumetric control cannot coarsen a mesh — that is, the smallest size takes precedence. Therefore, for the volumetric control to refine the mesh, the size that you specify must be smaller than the corresponding reference or boundary surface size.

To refine with a volumetric control:

1. Right-click the **Volumetric Controls** node and select **New**.



Initially, the volumetric control node is inactive , which means it does not have a geometry part or volume shape that is assigned to it.

2. To assign a geometry part, click the *Parts* property and select the geometry part from the object selection dialog.



3. To assign a volume shape, click the *Shapes* property and select a previously created volume shape from the dialog.



4. Under the **Volumetric Control** > **Mesh Conditions** node, select the required mesh models that you want the volumetric control to influence.

   The available meshers depend on the meshing models that you have selected. In this example, only the Surface Remesher, Trimmer, and Prism Layer Mesher are available.



This action adds the **Custom Size** node under the **Mesh Values** manager.

5. Select the **Mesh Values** > **Custom Size** node and define whether the refinement is relative or absolute.



6. Select the **Relative Size** or **Absolute Size** node and specify the level of refinement in the *Properties* window.



Setting the Trimmer Isotropic Size

The trimmer isotropic size is equivalent to the cell size for other core mesh models. It specifies the same cell size in each coordinate direction.

Note: The surface cell size is subject to the trimmer mesher surface size restriction. Therefore, the core mesh cell size can change inside the volumetric control as it approaches the surface.

To specify an isotropic cell size:

1. Select the **Volumetric Controls** > **Mesh Conditions** > **Trimmer** node and activate the *Customize isotopic size* property.



This action adds the **Custom Size** node under the **Mesh Values** manager.

2. Select the **Mesh Values** > **Custom Size** node and define whether the refinement is relative or absolute.



3. Select the **Relative Size** or **Absolute Size** node and specify the level of refinement in the *Properties* window.



Setting Trimmer Anisotropic Sizes

The trimmer anisotropic cell sizes allow you to specify a different cell size in each coordinate direction.

The coordinate directions are relative to the trimmer mesher coordinate system.

If you do not specify an anisotropic size for a coordinate direction, the volumetric control uses the isotropic size. If you also do not specify an isotropic cell size, the volumetric control uses the global reference values.

The trimmed mesher can only refine cells by splitting the original cell in half. Therefore, the volumetric control refines the cells to the next lowest allowable cell size. For example, if the core mesh cell size is 8 units, then the cells within the volumetric control can only be split to sizes of 4, 2, and 1 units. So if you specified a cell size of 3, you would actually get a cell size of 2 within the volumetric control.

Cells within an anisotropic volumetric control are not subject to the template growth rate settings. See Trimmed Mesher Properties. Anisotropic cells transition from small cells to large cells as fast as possible (equivalent to the fast growth rate setting).

If the boundary is aligned with the anisotropic coordinate directions, the volumetric control refines cells up to and including the boundary. However, if the surface is not aligned with the coordinate system, then anisotropic refinement reverts to isotropic refinement near the surface, as this behavior maintains good cell quality.

To specify anisotropic cell sizes:

1. Select the **Volumetric Control** > **Mesh Conditions** > **Trimmer** node and activate the *Customize anisotropic size* property.



This action adds the **Trimmer Anisotropic Size** node under the **Mesh Values** manager.

2. Select the **Mesh Values** > **Trimmer Anisotropic Size** node and define whether the refinement is relative or absolute.



3. Activate the following properties for cell sizes in the X-, Y-, and Z-direction respectively:
   - *Custom X size*
   - *Custom Y size*
   - *Custom Z size*



4. Select the sub-node and specify the level of refinement in the *Properties* window.



Example of Anisotropic Cell Sizes

The following example shows anisotropic refinement in the y-direction:



Setting Prism Layer Volumetric Controls

Applying a volumetric control to the prism layer model allows you to refine the number of prism layers, prism layer stretching, and prism layer thickness.

To specify prism layer volumetric controls:

1. Select the **Volumetric Control** > **Mesh Conditions** > **Prism Layer Mesher** node and activate the *Customize prism layer mesher* property.



2. Use the following nodes and their properties customize the prism layers within the volumetric control:
   - Number of Prism Layers
   - Prism Layer Stretching
   - Prism Layer Thickness

     For the prism layer thickness, you can use a relative or absolute value.

Volumetric Control Manager Context Menu

This section describes the context menu options that are available for the Volumetric Control Manager.

New	Executes the create volumetric control action.

Working with Volume Shapes

A volume shape is a closed geometric figure that can be used to specify a volumetric control for surface or volume mesh refinement or coarsening during the meshing process in STAR-CCM+.

The **Tools > Volume Shapes** node is the manager object for all volume shapes that have been created in the simulation.



It appears in the object tree as soon as you create a mesh continuum, and is empty to allow the creation of the first volume shape. The volume shapes manager can hold any number of volume shapes up to the memory limits of the computer.

Volume shapes allow any one of four different types of closed volumes to be defined using the context menu:

- Block
- Cone
- Cylinder
- Sphere

These volume shapes can be used to define volumetric controls for mesh continua in STAR-CCM+. The volume shapes can be created manually (using data values) or graphically using the point tool, in either the global coordinate system or any locally defined ones.

Context Menu of the Volume Shapes Manager

These options are the context menu items available when you right-click the **Volume Shapes** manager node.

New	Creates a new volumetric control of the type that is selected in the submenu.	
	Sphere	Executes the create sphere action to add a sphere to the **Volume Shapes** node.
	Block	Executes the create block action to add a block to the **Volume Shapes** node.
	Cylinder	Executes the create cylinder action to add a cylinder to the **Volume Shapes** node.
	Cone	Executes the create cone action to add a cone to the **Volume Shapes** node.

Context Menu of Individual Volume Shape

These options are the context menu items available when you right-click an individual volume shape node.

Highlight	Highlights the volume shape in the current scene. This option is intended for use when the default highlighting has been deactivated, and is only available if a scene is showing in the *Graphics* window.
Edit Volume Shape in Current Scene	Executes the edit action allowing you to set the properties of this object interactively. This option is only available if a scene is showing in the *Graphics* window.

Working with Blocks

The block volume shape is a rectangular volume that is aligned with either the global or a local xyz axis.

It is defined using the coordinates of two opposite corner points. Once defined, this volume shape has a node in the STAR-CCM+ object tree, has its own properties and shares a context menu with other types of volume shapes.



The block shape can be created either manually or graphically using the point tool. An example is shown below:



Creating a Block

On selecting the **Block** option in the **New Shape** pop up menu, the *Create Block Volume Shape* dialog appears.



To define the block, specify two separate locations (Maximum and Minimum Coordinates of the opposite corner points) based on either the global cartesian or a local coordinate system. These locations can be entered manually into the dialog, or the point tool can be used to drag the face center balls graphically to the appropriate locations. Pressing the **Create** button creates a new node that is called **Block 1** under the **Volume Shapes** node. This process can be repeated as many times as necessary to create further block shapes. Press the **Close** button to complete the operation.

Block Properties

You can use the properties that are listed below to adjust the specifications of the block.

Coordinate System	The coordinate system (either global or local) in which the volume shape is defined.

Corner1	A Cartesian position vector that defines the location of one corner. Note that a comma must separate the numeric values.
Corner2	A Cartesian position vector that defines the location of the other corner, which completes the object. Note that a comma must separate the numeric values.

Working with Cones

To define the cone volume shape, specify two radii and the xyz center points of the two cylindrical ends in either the global or local co-ordinate system.

Once created, this volume shape has a node in the STAR-CCM+ object tree, has its own properties and shares a context menu with other types of volume shapes.



The cone volume shape can be created either manually or graphically using the point tool. An example is shown below:



Creating a Cone

On selecting the **Cone**option in the **New Shape** pop up menu, the *Create Cone Volume Shape* dialog appears on the plot display.



To define a cone, specify two separate locations (Start Circular Disc and End Circular Disc) and two radii (**Start Radius** and **End Radius**) using either the global cartesian or a local coordinate system. These locations can be entered manually into the dialog, or the point tool can be used to drag the start and end points to the appropriate locations. Note that a representation of the cone can be seen by entering the **Start Radius** and **End Radius** values and pressing enter for each. Pressing the **Create** button creates a new node that is called **Cone 1** under the **Volume Shapes** node. This process can be repeated as many times as necessary to create further cone shapes. Press the **Close** button to complete the operation.

Cone Properties

You can use the properties that are listed below to adjust the specifications of the cone.

Coordinate System	The coordinate system (either global or local) in which the volume shape is defined.
Start Coordinate	A Cartesian position vector that defines the center of the base of the cone. Note that a comma must separate the numeric values.
Start Radius	Radius of the base of the cone.
End Coordinate	A Cartesian position vector that defines the center of the tip of the cone. Note that a comma must separate the numeric values.
End Radius	Radius of the tip of the cone.

Working with Cylinders

To define a cylinder volume shape, specify a radius and the xyz center points of the two cylindrical ends in either the global or local coordinate system.

Once created, this volume shape has a node in the STAR-CCM+ object tree, has its own properties and shares a context menu with other types of volume shapes.



The cylinder volume shape can be created either manually or graphically using the point tool. An example is shown below:



Creating a Cylinder

On selecting the **Cylinder** option in the **New Shape** pop up menu, the *Create Cylinder Volume Shape* dialog appears on the plot display.



To define the cylinder, specify two separate locations (Start Coordinate and End Coordinate) and a **Radius** that is based on either the global cartesian or a local coordinate system. The coordinates can be entered manually into the dialog, or the point tool can be used to drag the start and end points to the appropriate locations. You can see a representation of the cylinder by entering a **Radius** value and pressing Enter. Pressing the **Create** button creates a new node that is called **Cylinder 1** under the **Volume Shapes** node. This process can be repeated as many times as necessary to create further cylinder shapes. Press the **Close** button to complete the operation.

Cylinder Properties

You can use the properties that are listed below to adjust the specifications of the cylinder.

Coordinate System	The coordinate system (either global or local) in which the volume shape is defined.
Start Coordinate	A Cartesian position vector that defines the center of the start of the cylinder. Note that a comma must separate the numeric values.
End Coordinate	A Cartesian position vector that defines the center of the end of the cylinder. Note that a comma must separate the numeric values.
Radius	Radius of the cylinder.

Working with Spheres

To define a spherical volume shape, specify a radius and the xyz center point of the sphere in either the global or local coordinate system.

Once created, this volume shape has a node in the STAR-CCM+ object tree, has its own properties and shares a context menu with other types of volume shapes.



The sphere volume shape can be created either manually or graphically using the point tool. An example is shown below:



Creating a Sphere

On selecting the **Sphere** option in the **New Shape** pop up menu, the *Create Sphere Volume Shape* dialog appears on the plot display.



To define the sphere, specify the **Origin** of the sphere and the **Radius**, using either the global cartesian or a local coordinate system. The point tool allows the origin to be selected graphically by dragging the ball to the appropriate location, or alternatively, the coordinates can be entered directly into the dialog. Note that a representation of the sphere can be seen by entering a **Radius** value and pressing enter. Pressing the **Create** button creates a new node that is called **Sphere 1** under the **Volume Shapes** node. This process can be repeated as many times as necessary to create further spherical shapes. Press the **Close** button to complete the operation.

Sphere Properties

You can use the properties that are listed below to adjust the specifications of the sphere.

Coordinate System	The coordinate system (either global or local) in which the volume shape is defined.
Radius	The radius of the sphere.
Origin	A Cartesian position vector that defines the center of the sphere.

Mesh Continua

A mesh continuum is a computational domain for which a mesh is created using a collection of surface and volume meshing models.

The input geometry and associated regions define the extents and topology of the continuum. A mesh continuum node has its own properties and context menu.



An example showing two mesh continuum nodes with various meshing models activated in STAR-CCM+ is shown below.



A mesh continuum can be defined independently of both the physics and existing regions. The models that are activated in the continuum determine the characteristics of the type of mesh (surface and/or volume) that is created for the regions that are associated with that continuum during meshing.



During meshing, it is not necessary for every region to be associated with a continuum in order for the meshing process to proceed. Those regions that are not associated with a continuum are ignored and no mesh is created for them. By default, if no mesh continua exist then all regions are assigned to the first mesh continuum created. Thereafter, a null region is left unassigned if further mesh continua are created and

must be assigned manually. More than one region can be associated with the same continuum.

In addition to the models comprising the continuum, the mesh continuum also contains:

- reference values
- volumetric controls

Mesh Continuum Properties

This section describes the properties of the mesh continuum.

Per-Region Meshing	Activates per-region meshing for regions that are associated with this continuum.
Use Parallel Meshing	Activates parallel meshing for this continuum.
Interfaces	The interfaces that are associated with the continuum (read-only).
Regions	The regions that are associated with the continuum (read only).

Mesh Continuum Expert Properties

This section describes the export properties of the mesh continuum.

OOC translation	Determines the use of file I/O for meshing operations.	
	Activated	Allows file I/O to be used for meshing operations.
	Deactivated	File I/O is not used for meshing operations.
Verbose Output	Provides further output from the meshing models while they are running. This output can be useful for debugging problems when they occur.	
	Activated	Activates verbose output.
	Deactivated	Leaves output at its default.

Interpolation Option	Determines the method that is used to map the solution from an existing volume mesh to a newly generated volume mesh. These options are the same as the options provided for region mesh replacement.	
	Nearest neighbor	The cell/face value on the new mesh is assigned the value of the cell/face on the old mesh whose cell/face centroid is nearest to the cell/face centroid of the new mesh. Does not provide smooth interpolation of results.
	Higher-order stencil	The cell/face value on the new mesh is smoothly interpolated from a stencil of cells/faces on the old mesh. Smooth interpolation is based on a distance-weighted, least squared, interpolation scheme. Preferred method for transient analyses.

Mesh Continuum Pop-Up Menu

This section describes the pop-up menu for the mesh continuum.

Select Meshing Models...	Executes the select models action to allow selection of the models to be activated in this continuum.

Interpreting the Verbose Output

Activating the *Verbose Output* property for a mesh continua reports more detail on the tasks the surface and volume meshers perform when

generating a mesh. This information can be useful for debugging and evaluating the volume mesh.

However, note that the extra time taken to generate the verbose output can increase meshing CPU time by approximately 20%, which can be an issue for large meshes.

There are many differences between the normal and verbose output, which would not be feasible or necessary to fully document. Generally, the verbose output provides more information at each stage of the meshing process. For example, when fixing non-manifold baffle surfaces, the normal output reports the following:

```
     Checking baffles for containment
     Fixing non-manifold surface, CPU Time: 0.13, Wall
Time: 0.13, Memory: 215.89 MB
```

Whereas the verbose output reports more detail on what the mesher is doing to fix the non-manifold baffle surface, as shown below:

```
    Checking baffles for containment
    Fixing non-manifold surface
    found 1810 nonmanifold edges
    found 0 nonmanifold vertices
    created face pairs
    created vertex manifolds
    created 1810 new edges
    created 602 new vertices
    done, CPU Time: 0.16, Wall Time: 0.16, Memory:
216.92 MB
```

This type of information can be useful for identifying the cause of errors during the meshing process.

The verbose output also reports more information on the mesh statistics throughout the meshing process. For example, when generating a polyhedral mesh, the normal output reports the following:

```
    Estimating number of cells...
    estimated number of cells: 1468015
    done, CPU Time: 0.03, Wall Time: 0.03, Memory:
225.77 MB
```

Whereas the verbose output reports further statistics on the volume mesh, as shown below:

```
    Estimating number of cells...
    estimated number of cells:   1468015
    volume enclosed by surface:  3.36484e-003
    area of surface:             2.55161e-001
    average volume:              1.87295e-009
    average area:                2.73713e-006
```

```
    average edge length:         2.51419e-003
    done, CPU Time: 0.03, Wall Time: 0.03, Memory:
226.19 MB
```

The key difference between the normal and verbose output is the mesh statistics that are printed at the end of the report. This output includes quality and validity metrics for the volume mesh, as detailed below:

- Polyhedral quality

  This value refers to the cell quality of the polyhedral cells.

- Cell face validity

  This value refers to the face validity of the cells.

- Cell orthogonality

  This value refers to the angle between the normals of two faces in a cell that share an edge. The closer the cell edge angles are to 90 degrees, the more orthogonal the cell is.

- Cell skewness angle

  See cell skewness angle.

- Cell face planarity

  This value is a measure of the how planar the cell face is. A value of 1.0 indicates that face is planar.

- Cell face twist

  This value is a measure of the how twisted the cell face is. The closer the value is to 1.0, the more twisted a face is.

- Cell volume ratio

  This value refers to the volume change between neighboring cells.

Each volume mesh metric is reported in a tabular form, with the metric value range, number of cells and percentage. The minimum or maximum value in the volume mesh (depending on which is less desirable) is also reported for each metric. An example of polyhedral quality output is provided below:

Output - Star 1

```
-----------------------------+-----------+-----------
 Polyhedral quality (norm.)   |  # Cells  | % of Total
-----------------------------+-----------+-----------
  1.000000000 - 0.750000000   |     44500 |   18.944%
  0.750000000 - 0.500000000   |    115152 |   49.020%
  0.500000000 - 0.250000000   |     63546 |   27.052%
  0.250000000 - 0.100000000   |      9558 |    4.069%
  0.100000000 - 0.010000000   |      1291 |    0.550%
  0.010000000 - 0.001000000   |       857 |    0.365%
  0.001000000 - 0.000100000   |         3 |    0.001%
  0.000100000 - 0.000001000   |         0 |    0.000%
  0.000001000 - 0.000000001   |         0 |    0.000%
  0.000000001 - 0.000000000   |         0 |    0.000%
-----------------------------+-----------+-----------
 Minimum Polyhedral Quality: 4.069396e-004
```

Surface Meshers

STAR-CCM+ contains meshing tools which can be used to help prepare a high-quality surface mesh before a volume mesh can be generated.

Depending on the initial quality of the geometry, there are several surface meshing tools available:

- Surface Wrapper
- Surface Remesher
- Automatic Surface Repair

The surface wrapper is used when the imported geometry is of such poor quality (intersecting surfaces, holes, gaps, and so on) that it is not directly usable for volume meshing. In these cases carrying out surface repair and remeshing would typically involve more effort than any extra fidelity obtained would warrant. The surface wrapper is used with the feature extraction and feature editing tools to ensure a high-fidelity representation of the underlying geometry.

The result of a surface wrapping operation is a closed (watertight) surface that may require further remeshing to improve mesh quality.

The surface remesher is used to retriangulate a closed starting surface to obtain a high-quality surface mesh. It is used to improve surface meshes from the surface wrapper or to improve the quality of imported tessellated geometry. When used on imported geometry directly, surface remeshing is used with the surface repair tools such as the hole filling and edge zipping tools to ensure a closed surface. Many of these repair tools can be automated by using the automatic surface repair tool to minimize the manual effort required.

This section also discusses:

- When to use each model
- How to launch the surface mesh generator

Surface Meshers Overview

This section provides a comparison of the Surface Remesher and Surface Wrapper

To decide which surface meshing approach to use in preparation for volume meshing, weigh up the level of fidelity that is required against the acceptable level of manual effort.

For example, in cases where it is important that there is little or no loss of surface detail, the surface remesher would be used to maintain the original features. In cases where a significant level of de-featuring is required, and there are many gaps and mismatches between surfaces and parts, the surface wrapper would be used.

To decide which model or models are appropriate, use the following guidelines:

Instances when surface remesher would be used:

- When a high level of surface mesh fidelity to the underlying geometry is required
- When the imported surface is closed but composed of bad quality triangles (for example, STL data)
- When the imported surface is closed but does not contain the required triangulation resolution
- When the surface wrapper has been used, and the surface mesh needs further improvement

Instances when the surface wrapper would be used:

- When a high level of surface accuracy is not critical
- When the imported surface contains large number of gaps, holes and/or mismatches that cannot be easily fixed manually
- When the imported surface contains multiple volumes that intersect but are not joined
- When the imported geometry contains multiple overlapping surfaces that must be merged into one
- When the imported geometry contains too much detail, requiring de-featuring

See Also

Preparing Boundaries for the Generalized Cylinder Mesher

Surface Wrapper

The surface wrapper can be used to provide a closed, manifold, non-intersecting surface when starting from poor quality CAD data.

It is typically used when the imported geometry has:

- Multiple intersecting parts
- Missing data in the form of holes and gaps
- Surface mismatches
- Duplicate and internal surfaces
- Overly complex geometry with too much detail

The surface wrapper can:

- Provide a closed, manifold, triangulated surface definition
- Close holes, gaps, and mismatches
- "Skin" duplicate surfaces and remove internal features
- Simplify the surface geometry by removing unwanted detail
- Include baffle thickness inflation for desired internal features
- Include refinements that are based on curvature, proximity, and individual surfaces

The resulting surface quality from the surface wrapper is not optimal so it is used with the surface remesher to provide a high quality starting surface for the volume meshers.

An example result from the surface wrapper is shown below:

The surface wrapper controls can be set in three different groupings:

- Default controls - the global settings.
- Custom controls - settings that control the mesh density near certain part surfaces, feature curves, or volume parts.
- Contact protection - settings to prevent surfaces incorrectly connecting.

These groupings let you control the relative level of fidelity versus de-featuring, and control the distribution of surface mesh resolution.

The location of these controls varies in the GUI depending on whether you are using parts-based meshing or region-based meshing. For parts-based meshing, the controls are all under the **Operations** > **Surface Wrapper** node. For region-based meshing, the global controls are under the **Continua** > **Mesh** node. The custom controls are under the region node to which they pertain, and the contact protection controls are under the **Regions** > **Body** > **Mesh Values** node.

The leak detector option can be used to check if holes or gaps exist in the surface either before wrapping or after the surface mesh has been created.

Input part surfaces or are typically used as input parts to a single surface wrapper operation to generate a single wrapped surface. If you want to wrap some parts separately, then use a separate surface wrapper operation (PBM) or mesh continuum (RBM) for each group of surfaces. Alternatively, activate the per-part meshing (PBM) or per-region meshing (RBM) option to generate a separate wrapped surface for each part or region.

If you only want to wrap part of a geometry, you can activate the surface wrapper property, *Perform Partial Wrapping*. This property allows you to select the faces that you want to exclude from wrapping in a surface wrapper operation. Excluded faces are not wrapped but are still included in the final mesh. See Partial Wrapping.

See Also

Generating a Surface Mesh

Leak Detection

The surface wrapper leak detection tool is useful for determining whether "leaks" exist in the resultant wrapped surface that is created from an imported surface.

A leak is a gap or a hole that would allow a path to be traced from one volume through to another volume or the outside without crossing through a surface. Therefore, a fully enclosed surface has no leaks.

The leak detection tool can only be used if an import surface representation of some kind exists. The leak is detected in the octree version of the mesh that is built based on the supplied wrapper inputs. However, if a leak is found then the path is displayed using the original import surface representation (not the octree or wrapped mesh). Displaying the leak in this way allows repairs to be made to the surface if necessary to avoid the leak from occurring in the actual wrapped surface representation.

The leak detection tool is accessed by right-clicking on the **Surface Wrapper** model node in the simulation tree and selecting **Run leak detection...** in the pop-up menu:



The *Wrapper Leak Detection edit* tab appears, as shown below:



The *Wrapper Leak Detection* edit pane can be used in three different modes:

- **Locate Points** - Create and move source and target points for the leak detection calculation
- **Repair Surface** - Patch up the surface before and after leak detection
- **Offset faces** - Offset or translate surfaces

By default, the **Mode** is set to **Locate Points**, allowing you to define the location of the source and target points using either the *Point Location* entry boxes and/or the graphical point tool.

By default, there is a single source and target point. The source point (blue) would typically reside inside the geometry volume while the target point (red) would be anywhere outside (in another volume if necessary).

To change the location of either point:

- set the **Current Point** menu to the appropriate selection, or
- in the *Graphics* window, right-click on a point to select it.

Then, use the *Point Location* entry boxes and/or the graphical point tool to alter the coordinates of the point that is selected. The larger red ball indicates the currently active point on the display.

Once the locations of the source and target points have been set, the **Recompute Template + Paths** button can be pressed to invoke the leak detection process. When running leak detection for the first time on a new geometry, the method builds a template that is based on the supplied input parameters for the surface wrapper. The wrapper then checks to see if a leak would exist based on the template cell sizes. If a leak is found, then it is traced using a line from the source point to the target point, as shown below:



Where the leak path crosses through the surface indicates the location of the hole which must be fixed. A summary of the leak detection calculation is reported in the *Output* window. A typical example is shown below:

```
Output - Star 1
User set advanced options to:
Generating template mesh...
Checking input surface validity
   done, CPU Time: 0.01, Wall Time: 0.01, Memory: 72.00 MB
Computing Bounding Box
   done, CPU Time: 0.00, Wall Time: 0.00, Memory: 72.23 MB
Computing Entity Depth
   done, CPU Time: 0.11, Wall Time: 0.17, Memory: 72.51 MB
Generating Octree Mesh
   octree mesh statistics:
      number of leaf octants intersecting input surface:  10097
   done, CPU Time: 0.66, Wall Time: 0.75, Memory: 73.60 MB
Computing path...
Path found between source and target points.
These points will be on the same side of the final wrapped surface, connected by the path indicated.
```

If a leak is detected, then it reports back that the source and target points are on the same side of the wrapped surface and a leak path is shown. If there is no leak, then no leak trace is displayed and a message saying

that no path has been found between the two points appears in the *Output* window.



If multiple holes exist in the input surface, then a single target point only detects one of them. The addition of extra target points in different locations around the geometry or in different volumes helps detect holes that the wrapper did not fill in. However, it is not guaranteed that adding a target point results in a new hole being found since the point location affects the computed path. Extra target points can be included by pressing the **New Target Point** button. The **Current Point** menu can be used as before to select these points so that their location can be adjusted. Each leak path to a different target point is shown in a different color, even if it passes through the same hole. The **Delete Current Target Point** button can be used to remove a pre-existing target point from the calculation.

Changing the location of any of source or target points, or adding or deleting target points requires the recalculation of the leak path. To recalculate the path, press the **Recompute Paths** button, which uses the previously defined template as the basis of the calculation.

Once a leak path is identified, the faces near the leak path hole can be automatically selected by pressing the **Select Faces Close to Leak Path** button. The *Select Faces Close To Path Dialog* appears:



Specify the **Radius** value (in any desired units) to give the size of a sphere that moves along the path from the source to the target. The faces that fall within the sphere are the ones that are selected (primarily the faces next to the hole). Press the **OK** button proceed with the operation.

An example leak path face selection is shown below:



Being automatically able to select the faces next to the hole aids with the manual surface repair step that follows, as described below. Once a leak path has been detected, the *Wrapper Leak Detection* dialog can be switched to **Repair Surface** mode by using the **Mode** pull-down menu.


Servers
Star 1
simulation
scene/plot
edit
Wrapper Leak Detection
Mode:
Repair Surface
Repair Surface
Show Free Edges?
Selection Control
Faces
1
Edges
0
Vertices
0
Display Control
Color Faces by
Default Color
Boundary/Feature Curve Edit
Set Control
Leak Detection Actions
Recompute Template + Paths
Recompute Paths
Select Faces Close to Leak Path
Apply
Close
Help

Using this dialog, any holes in the surface can be repaired using the *Surface Repair* buttons, which are identical in behavior to the normal surface repair options. The **Show Free Edges** toggle can be used to highlight (using green lines) the holes within the geometry. Of particular importance in the **Surface Repair** options is the **Fill polygonal patch** button, which can be used to patch arbitrary holes in the surface simply by clicking the surface around the hole. This option does not fill the hole perfectly but covers it sufficiently for the surface wrapper not to detect a hole anymore and produces a uniform surface in its place. The Selection Control, Display Control, and Boundary/ Feature Curve Edit options in the dialog can be used the same way as for the normal surface repair dialog.

Note that if you imported the geometry as a part, filling holes at this stage only applies changes to the initial surface representation, whereas the geometry part representation remains unchanged. This fact can cause problems if you reinitialize the mesh, as the surface repairs are lost. To ensure that fixes to the surface are permanent, exit the leak detection tool, and use surface repair at the parts level to fix the underlying geometry. This recommendation does not apply if you imported the surface mesh as a region.

Once any holes have been filled or patched, any remaining leak paths can be found by pressing the **Recompute Paths** button with the existing source/target point combinations. Depending on the number of holes that are filled and their size, it can be necessary to recalculate the template by pressing the **Recompute Template + Paths** button. If there is any doubt whether leaks still exist, follow this approach as the safest route for leak detection. At any time during the surface repair operation, you can return to the **Locate Points** mode dialog and reconfigure the source and/or target points to assist leak detection to find holes. Once all leak holes are closed, you can proceed to wrap the surface and then use the surface remesher to improve the resultant quality.

To reiterate - the leak detection option is run before the wrapped surface is created (that is, on the original input surface). The leak detection tool accounts for the current wrapper settings and determines if any gaps in the surface would be closed based on the supplied settings. If the sizes specified are not sufficient, it reports back with the leak path through the hole, which must be fixed before progressing. Alternatively, the wrapper settings, in particular the gap closure size, can be modified so that the wrapper can automatically close the hole itself.

Contact Prevention

Contact Prevention avoids the merging of two or more part surfaces that are close to one another during surface wrapping.

This functionality is useful during meshing as it keeps certain surfaces separated, which would otherwise be amalgamated as one due to the surface mesh size settings.

To activate contact prevention for two or more surfaces when using a parts-based meshing approach, right-click **Operations** > **Surface Wrapper** > **Contact Prevention** and select **New**:



If you are using region-based rather than parts-based meshing, contact prevention is found under **Regions** > **[Body name]** > **Mesh Values** > **Contact Prevention**. It works with region boundaries rather than part surfaces, but otherwise looks and functions in an identical way.

Two options are available:

- One Group Contact Prevention Set
- Two Group Contact Prevention Set

Once you have defined a contact prevention set, the following properties are available:

- Search Floor

- Enable Contact Prevention Set
- Using Parts and Shapes to limit contact prevention

In addition to common menu items, the right-click menu of a contact prevention set contains actions for:

- Interactively selecting surfaces within the set
- Specifically highlighting surfaces when default highlighting has been deactivated.

One Group Contact Prevention Set allows individual part surfaces to be grouped, such that each part surface forms a contact prevention with every other part surface grouped in the set.

1. Selecting **One Group Contact Prevention Set** creates a node that is called **Contact Prevention 1**.

   The properties of this node are shown below for a parts-based meshing approach:



The part surfaces that form the contact prevention set can be selected using two different methods.

The first option uses the object selection dialog, which is accessible by clicking the ellipsis (**...**) next to *Surfaces* in the *Properties* window.



If you only have a few part surfaces in the model and they are easily identifiable, using the object selection dialog to select the part surfaces is the easiest and simplest method.

However, if you have many part surfaces or prefer to select part surfaces interactively in the *Graphics* window, use the *Select Part Surfaces Interactively* dialog.

2. To access the *Select Part Surfaces Interactively* dialog, first ensure that there is a scene open, as the dialog requires selection of part surfaces in the *Graphics* window.
3. Then, right-click the **Contact Prevention** node and select **Select Part Surfaces...**



The *Select Part Surfaces Interactively* in-place dialog opens:



4. Select the part surfaces for the contact prevention set using the mouse in the *Graphics* window (press `<Ctrl>` for multiple selections).

5. Once you have made the appropriate selections, click + to add the part surfaces to the list. Alternatively, right-click a selected part surface and select **Add Part Surface**.

   Selected part surfaces appear blue in the display, indicating they belong to a set.

6. If part surfaces must be removed, select the part surfaces in the list box or in the display and click − to remove them. Alternatively, right-clicking on a part surface in the display provides a **Remove Part Surface** option. Clicking ✕ removes all part surfaces from the list box.

   Select at least two part surfaces in order for the contact prevention to be valid.

7. Once the selections are complete, click **Apply** to apply the changes, and then **Close** to return to the simulation tree.

Two Group Contact Prevention

Two Group Contact Prevention Set allows part surfaces in one group to form contact prevention sets with individual part surfaces in the other group, but not with part surfaces in their own group.

1. Selecting **Two Group Contact Prevention Set** creates a node that is called **Contact Prevention 1**.

   The properties of this node are shown below:



2. Part surfaces for each group are selected in the same way as for a single contact prevention set.

If part surfaces are selected interactively, a list box for both the first and second group is provided:



3. Part surfaces can be added or removed from groups in the same way as for a single contact prevention set, but care should be taken that they are added to the correct group.
4. You can also use the *Close Part Surfaces* finder to identify part surfaces within a specified tolerance of the surfaces in the first group, and add them to the second group.
5. To use the *Close Part Surfaces* finder, first add one or more part surfaces to the first group.
6. Then enter the **Tolerance**, which specifies the search distance from the surfaces in the first group, and click **Find**.

   The tool finds all surfaces that have any part of their surface within the tolerance, and adds them to the drop-down menu.



   Browse buttons are available to inspect individually each part surface the tool finds.
7. To add a part surface to the second group, select the part surface from the drop-down menu and click . To add all part surfaces in the menu, click .

   Surfaces that are selected for the first group appear blue in the display, while surfaces in the second group appear orange.
8. Select at least one surface for both the first and second group in order for the contact prevention to be valid.

Search Floor

The *Search Floor* is a minimum distance between the part surfaces at which contact prevention is guaranteed.

1. As a guideline, set this value to be approximately the largest distance between the surfaces that you want to keep separated.

Although contact prevention occurs if the distance between the surfaces is less than the search distances, it is not guaranteed to work in this case.

Enable Contact Prevention Set

A contact prevention set node can be temporarily deactivated, while still keeping all previously specified inputs, by using the *Enable Contact Prevention Set* property:



By default, the *Enable Contact Prevention Set* option is on.

1. Turning off the option deactivates the contact prevention node and deactivates the settings from any future surface wrapping operation.

The node icon that is associated with the set is also grayed out, indicating the change in status.

Using Volume of Influence to Limit Contact Prevention

The expert property, *Volume of Influence* is used to define a volume within which contact prevention applies.

1. Click the ellipsis (**...**) to the right of the property to launch the *Volume of Influence* selection dialog.
2. Choose one or more parts that define the volume within which contact prevention is to apply and click **OK**.

   As an example, the dialog below selects **Block 1** to define the volume of influence.



When the surface wrapper runs, contact prevention is only applied to the area of participating surfaces that lie within the specified volume. The wrapper does not prevent contact for the remainder of these surfaces outside the specified volume. For this example, the volume of influence is applied to the contact prevention set created for the sensor boundary in the manifold surface wrapping tutorial. The resulting surface wrap is modified as shown below. The outline of the volume shape **Block 1** is also shown.



Highlighting Contact Prevention Part Surfaces

By default, the part surfaces that are set on a contact prevention set are highlighted in a scene when you select the node for the contact prevention set.

1. If you have deactivated this default highlighting, you can still choose to highlight the part surfaces by right-clicking on the contact prevention set node and choosing **Highlight** from the context menu.

This option is only available if a scene is showing in the *Graphics* window.

Surface Wrapper with Internal Baffles

The surface wrapping process automatically detects and inflates internal surfaces representing baffles.

The total thickness that is applied to each side of the baffle is equal to 0.25 times the smallest refinement size that is used in the surface wrapping process. This size is not necessarily equal to the minimum surface size specified in the **Default Controls** node (PBM) or **Reference**

Values node (RBM). If you want to have more control over the thickness of the baffles, then it is recommended that you inflate the surface using the offset surface preparation tool before proceeding with surface wrapping.

Note that internal surfaces (by definition) are disconnected from the main geometry and are such subject to the smallest disconnected number of faces setting. By default, the surface wrapper automatically removes any isolated internal surfaces with less than 1000 faces (unless all surfaces have less than the specified number of faces in which case nothing is deleted). This process removes any small "artificial" surfaces that do not represent any geometry, which simplifies the volume meshing process. However, in some cases using a coarse surface wrap size can inadvertently coarsen an internal baffle surface such that it has less than 1000 faces. In such cases, consider reducing the value or deactivating it altogether by using a value of 0.

Surface Wrapper Properties

This section describes the properties that you can set for the Surface Wrapper.

When the surface wrapper is activated, the complete list of model properties can be viewed by selecting the **Geometry** > **Operations** > **Surface Wrapper** node when working with Parts Based Meshing. If working with Region Based Meshing, the same properties are found by selecting the **Continua** > **Mesh** > **Models** > **Surface Wrapper** node (except for per-part meshing).



By default, the option for curvature refinement is on while per-part meshing, proximity refinement, gap closure, and mesh alignment are off. Clicking the appropriate box in the panel toggles the option on or off.

The surface curvature and surface proximity properties are common to several meshers, and as such are documented with the common surface mesher controls and values.

Per-Part Meshing (PBM only)	The equivalent functionality for Region Based Meshing is found under **Continua > Mesh**.	
	Activated	Meshes each part individually, which ignores any interfaces between parts. Therefore the mesh between connected parts is not conformal.

		This option is useful if the trimmer or thin mesher cannot produce a mesh at the interface.
	Deactivated	Meshes all the input parts at once.
Input Parts (PBM only)	Specifies the geometry parts for which the operation generates a mesh. The equivalent Regions input for Region Based Meshing is found under **Continua** > **Mesh**.	
Perform Gap Closure (PBM) / Do Gap Closure (RBM)	The **Perform Gap Closure** (PBM) / **Do Gap Closure** (RBM) option can automatically close holes that are larger than any of the specified reference surface sizes. However, do not use this option in tandem with the smallest wrapping volume specification, to prevent closure of internal geometry parts that have a small diameter.	
Perform Mesh Alignment (PBM) / Do Mesh Alignment (RBM)	Selecting the **Perform Mesh Alignment** (PBM) / **Do Wrapper Mesh Alignment** (RBM) option allows the mesh alignment location values to be set. That is, a wrapper mesh line can be "pinned" to a given value in any of the X, Y and/or Z directions.	
Create Weak Contacts With Input Parts (PBM only)	Activating the Create Weak Contacts With Input Parts option when the Per-Part Meshing option is deactivated creates new weak contacts. The weak contacts that are created do not result in a conformal mesh. However, when you assign parts to regions, these contacts are converted into interfaces which the solver then uses to make the mesh conformal.	
Perform Partial Wrapping (PBM only)	Allows you to create a **Partial Wrap Set** to exclude surfaces which do not require surface wrapping from being wrapped in a Surface Wrapper operation. Surfaces that are excluded from wrapping in the Surface Wrapper operation are copied from the original geometry and fixed to the wrapped surfaces. See Performing Partial Wrapping.	

Surface Wrapper Expert Properties

This section describes the expert properties that you can set for the Surface Wrapper.

Verbose Output	When using Parts Based Meshing, the surface wrapper expert property verbose output is available. Activating this option results in more verbose output being printed to the window during meshing.

Surface Remesher

In order to improve the overall quality of an existing surface and optimize it for the volume mesh models, the surface remesher can be used to retriangulate the surface.

The remeshing is primarily based on a target edge length that you supply and can also include feature refinement that is based on curvature and surface proximity. Localized refinement that is based on part surfaces or boundaries can also be included. Specific surfaces or boundaries can also be omitted from the process so that the original triangulation from the imported mesh can be preserved.

The surface remesher is typically used for remeshing surfaces that are output from the surface wrapper and STL type data. As well as improving the surface for the volume meshers it also aids the subsurface generator when the prism mesher option is selected. An example starting surface and the remeshed surface is shown here.



The properties of the surface remesher model can be changed to provide additional meshing control during the meshing process. The input values used for the surface remesher can be set on four different levels:

- Global level
- Part surface or boundary level
- Feature curve level
- Interface level

A volumetric control can be included as part of the surface remeshing operation. This option allows the surface target size to be set for any of the allowable shapes that encompass the import surface.

Aligned meshing can also be included to optimize the resulting surface triangulation further.

See Also

Generating a Surface Mesh

Aligned Meshing and Geometric Features

In the surface remesher, aligned meshing is a fully automated process which creates a templated mesh on surfaces with four corners.

Aligned meshing and retaining geometric features improve the surface mesh representation on CAD features, such as fillets and rounds. Other benefits include quicker CAD projection and a more efficient use of faces on the surface mesh.

For flat surfaces, the aligned mesher only produces an aligned mesh when the cells that are produced are of higher quality than when using the regular remesher. If the aligned mesher is likely to create poorly shaped triangles, it automatically reverts to the basic remesher output for that CAD surface. For curved surfaces, an effort is made to align the mesh with the curvature. For this reason, STAR-CCM+ often favors the aligned mesher while allowing for some degradation in the cell quality, as it captures the geometry more effectively.

In the **Surface Remesher** model expert properties, *Create aligned meshes* is activated by default. However, this feature is only active for cases which have CAD association, and where the Project to CAD property is activated (default).



Limitations of Aligned Meshing

- It is possible that simulation files saved in STAR-CCM+ 5.06 or older do not contain sufficient information for aligned meshing and retaining geometric features. To overcome this issue, either reimport the CAD file or re-tessellate the part.
- The *Create aligned meshes* option only works with faces and vertices on surfaces that have 100% CAD association.
- It is highly recommended that the *Retain Geometric Features* option remains activated. If you deactivate *Retain Geometric Features*, ensure that you mark the feature edges. If no feature edges are marked, the remesher has nothing with which to align the mesh.
- You can use one or more feature edges as boundaries for aligned meshing, as long as there is only one underlying CAD surface.
- For aligned meshing to work on a surface, the surface must have at least four CAD vertices, or four feature edges, surrounding the area of interest.
- Surfaces that are on or touching a periodic interface map are not aligned meshed, even when only one edge is shared. However, in-place maps between neighboring parts can be align meshed.
- There must be no inner loops on the surface.

- It is possible that surfaces that are connected to an interface between two parts only receive aligned meshing on one side of the interface and not the other.

Surface Remesher with Interfaces

In order to achieve a conformal mapped mesh at the interface, the surface remesher can be used with specific direct type interfaces.

This recommended procedure results in a better volume mesh definition for the analysis.

For an in-place type interface, a conformal mesh can be obtained for both single region and multi-region cases, as long as the same mesh continuum is used. For example, in a single region case, an in-place interface can be used to represent an internal baffle of some kind. For multi-region cases, the interface can represent a change of material, such as a fluid/porous interface or fluid/solid of some kind. Additionally, the surface remesher can now be used regardless of the perimeter shape of the boundaries at the interface as long as the interface boundaries are planar. In other words, the boundaries belonging to the interface do not need to have the same shape or area. An imprint method is automatically included as part of the surface remeshing process, meaning that the boundaries are first retriangulated to have a common perimeter before proceeding with the surface remeshing step.

When the surface remesher is used with a periodic interface, then it attempts to produce a conformal mapped mesh for the boundary pair. Both boundaries have to belong to the same region in this instance.

When using one of the above direct type interfaces for surface remeshing, then the interface node controls the settings and sizes that are common to both boundaries. Where remainder boundaries exist for in-place type interfaces (that is, those that belong to a boundary that is part of an interface but is not common to the boundary pair), the normal boundary node settings and sizes are used instead.

When visualizing the remeshed surface, include the interface boundary nodes that are connected to the boundary pair in the displayer parts selection so that the mesh triangulation can be seen on those faces. The interface boundary surfaces are also colored according to their type.

When using the repeating type interface, a conformal mesh interface is not required because the meshes in each region are sliding past each other. Here the per-region meshing option must be activated in order to ensure a good quality surface at the interface.

Additional notes on surface remeshing with interfaces are provided in the guidelines section.

Surface Remesher Guidelines

This section provides guidelines for using the Surface Remesher.

- It is not necessary for the import surface that the surface remesher acts upon to be closed. However, the surface remesher does not close this surface for you - where free edges occur in the starting surface, an equivalent free edge is present in the remeshed surface.
- Feature curves must be added to the import surface if you wish to maintain sharp edges.
- Internal feature surfaces such as baffles can be included in the import surface. Convert these surfaces to an interface before surface remeshing.
- In cases where proximity refinement is desired but the starting triangulation is very coarse (for example, STL data), a uniform triangulation in the proximity area may not be achieved. In this case, perform a surface remesh without proximity effects using an intermediate triangulation density, delete the import representation and rerun the surface remesher again with proximity effects included.
- With the surface curvature option turned off, a uniform surface triangulation that is based on the target surface size is obtained. However, it is recommended that you use surface curvature for most cases as the resultant surface definition is far superior to the definition obtained using uniform triangulation. An example comparison is shown here.



- When using in-place interfaces in a multi-region case, where the perimeter of the two interface boundaries does not match (leading to remainder boundaries), the imprinting process can fail.

  If this scenario occurs, try one of the following:

  - Remesh one or both boundaries separately to improve the surface triangulation in these areas before attempting to match them. To remesh the boundaries separately, either use surface remeshing with per-region meshing turned on, or use the remesh selected faces option for surface repair.
  - Reverse the order of the interface boundaries. Delete the interface and then select the two boundaries in the opposite order to that used for the original one before making a new interface.
  - When matching a curved surface to a planar one (for example, a sphere touching the side of a box), project part of the curved surface onto the planar surface first. A common area is then formed between the two surfaces that allows an interface to be created.
- As with all the meshing tools in STAR-CCM+, the surface remesher is subject to continual improvement from release to release in order to increase the overall quality of the surface triangulation produced. As such, an increase in the resultant face and vertex count can be seen for the same geometry when run in two different versions.

  A typical example of this situation is seen in highly curved geometries, where a better definition is achieved in the newer version than the old. The increase in face count can sometimes be undesirable (especially if there is a substantial increase). In this instance, relaxing the surface curvature requirement somewhat helps to reduce the overall face count. Depending on the geometry, further adjustments in user-specified sizes can be required in order to reproduce similar face counts (albeit with a higher quality) when comparing the newer version to the old.

Surface Remesher Properties

This section describes the properties that you can set for the Surface Remesher.

When the surface remesher has been activated, the properties are found under **Automated Mesh** > **Meshers** > **Surface Remesher** when using Parts Based Meshing (PBM), and under **Continua** > **Mesh** > **Models** > **Surface Remesher** when using Region Based Meshing (RBM).



The surface curvature and surface proximity properties are common to several meshers, and as such are documented with the common surface mesher controls and values. See Surface Mesher Controls and Values.

Perform Compatibility Refinement (PBM) / Do Compatibility Refinement (RBM)	Activating this option imposes a surface growth rate of 2 to limit the difference in face sizes across a gap. Therefore, if the gap between faces is small, the relative size of faces on either side of the gap is close to one another. This option is useful when using the thin mesher, and improves volume mesh quality, as face sizes on opposing sides are similar without explicitly stating a size.

	This option is similar to proximity refinement. Proximity refinement specifies a mesh size that is based on the gap size and the number of points inside the gap. Whereas, compatibility refinement ensures that the mesh size on either side of the gap is similar. In some cases, this results in fewer elements and a more efficient use of the surface remesher. For example, consider the surface mesh below. The curvature of the cylindrical pin reduces the mesh size near the pin: 
	Applying compatibility refinement has the following effect on the surface mesh opposite the cylindrical pin: 

	As can be seen above, the surface mesh in the region opposite to the cylindrical pin is refined, as the compatibility refinement makes the mesh size across the gap similar. Note that the surface growth rate also influences the mesh size, so the sizes are not identical. The refinement on the surface mesh in turn affects the volume mesh by creating refined and better quality cells in the gap between the pin and opposing surface. For comparison, the following image shows the surface mesh with only proximity refinement activated.
	 The mesh size on the opposing surface is independent to the cylindrical pin and the region is only refined due to the gap thickness.

Surface Remesher Expert Properties

This section describes the expert properties that you can set for the Surface Remesher.

Retain Geometric Features	Retaining geometric features aims to preserve all CAD edges when generating a surface mesh, which provides a better representation of the original CAD model. This option also makes the CAD projection process more

	efficient, as faces on the surface mesh do not span multiple CAD surfaces. The *Retain geometric features* option is a property of the **Surface Remesher** meshing model, and is activated by default. However, if geometric features cause poor face quality, the surface remesher is free to ignore them. The following examples illustrate the effect of the *Retain geometric features* option. The blue sections highlight CAD edges for comparison between the two surface meshes. The first image shows a surface mesh with *Retain geometric features* deactivated. 
	The second image shows a surface mesh with *Retain geometric features* activated.

	 You can identify which edges have CAD association by using the CAD Association option in the *Select Threshold* tool.
Create Aligned Meshes	Activating *Create aligned mesh* generates a surface mesh with face edges that aim to follow the curvature of rounded CAD features, such as fillets. The following issues can prevent the surface remesher from generating an aligned mesh: • An aligned mesh reduces the overall surface mesh quality. • The aligned meshing conditions are not met. • The underlying geometry is not suitable for aligned meshing. The *Create aligned mesh* option is a property of the **Surface Remesher** meshing model, and is activated by default. Use the checkbox to deactivate aligned meshing. The following example illustrates the effect of the *Create aligned mesh* option. The blue sections highlight CAD edges for comparison between this surface mesh and those in the previous section.



Currently, the following conditions apply to the *Create aligned meshing* option:

- Aligned sections must have 4 sides and have geometric features as closed boundaries, for example:



- There must be less than a 30% difference between the number of vertices that the surface remesher creates on opposing sides of an aligned section.
- The remeshed surface must have a relatively uniform surface mesh size. As a general guideline, the minimum surface size (without aligned meshing) must be greater than half of the maximum surface size.

	• Aligned meshes cannot lie on a periodic interface boundary, or be connected to one along an edge (in-place interfaces are supported). • Aligned sections cannot contain inner loops, such as an inlet duct on a filleted edge:
	
Minimum Face Quality	The *Minimum Face Quality* property for the surface remesher is a value that ranges from 0 to 1, with 0 being the worst and 1 being perfect. The quality of a triangle is calculated as the ratio of the triangle face area to the area of an equilateral triangle that would exactly fit inside the circumcircle of the triangle. The circumcircle is the circle which touches all three corner points of the original face. In the diagram below, the Minimum Face Quality is defined as: Minimum Face Quality = A / A_{equil}

	 The default value is `0.05`, and it can be changed with a slider. The surface remesher ensures that all the triangles in the resultant surface mesh have a quality equal to or greater than the specified value. This minimum quality is achieved by allowing the mesher to disregard feature edges, move feature vertices, or avoid projecting vertices if necessary. As a result, the surface remesher removes or improves faces whose quality is below the specified value. Setting the value to `0` deactivates the option. Recommended values are in the range `0.01` to `0.1`. Values cannot exceed 0.1.
Choose Field Function Refinement Table (PBM) / Choose a Mesh Size Table (RBM)	The **Choose Field Function Refinement Table / Choose a Mesh Size Table** property allows for the refinement of the mesh using field functions, based on an existing solution.

Automatic Surface Repair

The automatic surface repair tool provides an automatic procedure for correcting a range of geometric type problems that can exist in the remeshed surface once the surface remeshing process is complete.

Up to three different metrics are used to determine if automatic fixes are required:

- Pierced Faces (surface intersections including opposing intersections)
- Surface Proximity
- Surface Quality

The metric for pierced faces is always on and is invoked when the automatic surface repair option is activated. The metric for surface proximity is only used if the surface proximity property has been activated for the surface remesher. The metric for minimum quality is only used when the surface remesher has a minimum face quality value of greater than 0.

Within each metric, two different fixing mechanisms are used, arranged in best to worst order in terms of preserving the original geometry. The fixing mechanisms are:

1. **Remesh** - Grows out the problem area and then surface remeshes, removing features if necessary.
2. **Patch** - Grows out the problem area, deletes the local surface, and then invokes the hole filler.

In this instance, the remeshing option is first applied to a given problem area. If this option succeeds in eliminating the issue or improving the quality, no further action is taken and the next problem area is addressed. If the remeshing step did not resolve the problem, then the patch option is applied to the original problem definition instead.

During the auto repair fixing operation, it is possible for small disconnected parts of the surface mesh to be formed as a by-product. Two options are provided to allow you to delete these disconnected surfaces automatically:

- Connected surface count limit
- Connected surface size limit

These options can be accessed by selecting the **Reference Values > Automatic Surface Repair** node:



The connected surface count limit can be used to control the number of individual isolated or disconnected surfaces. For example, if you are only expecting to have one surface at the end of the resurfacing process, then a value of 1 can be specified. The surface with the largest face count or area is retained while all others are eliminated. Alternatively or in addition, the connected surface size limit can be used to specify the minimum face count and/or area that each disconnected surface requires. The above two options can be used separately or in tandem in order to control the resultant surface.

The auto repair option is typically used when performing surface remeshing on a wrapped surface. The auto repair option does not fix topological type problems such as free edges, non-manifold edges, and non-manifold vertices. Therefore, if these problems exist in the surface, they must be fixed by hand afterwards using the manual repair tools.

Auto Repair Properties

This section describes the properties that you can set for the Automatic Surface Repair mesher.

Connected Surface Count Limit	The connected surface count limit provides a method of controlling which connected surfaces are kept (or conversely, which disconnected surfaces are deleted) based on the total number of connected/disconnected surfaces. For example, if you are expecting to retain two surfaces at the end of the surface remeshing and auto repair process, then a limit of 2 can be specified. The two surfaces with the largest face count or area are then retained and all other surfaces deleted. The option to determine whether to use the face count or area specification is set using the *Connected surface count limit* property in the **Automatic Surface Repair** node: 

	By default, the **None** option is specified indicating that the feature is turned off. Selecting either the **Keep Largest Face Counts** or **Keep Largest Areas** option activates the feature and activates the **Maximum # of Connected Surfaces** node, as shown below:  The default value of `1e9` effectively means that all surfaces are retained by default. Entering a value of `1` for example would specify that only one surface is retained after the auto repair process. The *Maximum # of Connected Surfaces* property can also be set with a slider.
Connected Surface Size Limit	The connected surface size limit provides a method of controlling which connected surfaces are kept based on either the number of faces in the surface and/or the area of the

surface. For example, if you only want to maintain surfaces with more than 100 faces, then a minimum limit of 100 faces can be specified. The auto repair process deletes any resultant surface with 100 or less faces in it. The option to determine whether the face count and/or total area specification is used is specified by the *Connected surface size limit* property in the **Automatic Surface Repair** node:



By default, the **None** option is specified indicating that the feature is turned off. If either of the face count based options are selected, the **Minimum Connected Surface Face Count** node is activated, as shown below:



The default value for *Minimum Connected Surface Face Count* is 0, meaning that all surfaces are retained. If a value of 1000 is entered for example, then the auto repair process automatically removes any disconnected surface with 1000 or less faces in it.

If either of the area-based options are selected, then the **Minimum Connected Surface Area** node is activated, as shown below:



The default value for *Minimum Connected Surface Area* is 0, meaning that all surfaces are retained. If a value of 0.5 m^2 is entered for example, then the auto repair process automatically removes any disconnected surface with an area of less than 0.5 m^2. The value can be specified in any of the allowable units.

If both the area and face based limits are selected, both nodes (as discussed above) are activated and both an area and face based limit can be applied. In this instance, the disconnected surfaces are retained if they meet both criteria.

If both the connected surface count limit and connected surface size limit are used, then the latter is applied first. For example, if the

	maximum number of connected surfaces is set to 5, and the minimum connected surface face count is set to 500, surfaces with fewer than 500 faces are removed first. Thereafter, if there are more than 5 surfaces remaining then the 5 with the largest surfaces areas would be kept and the others removed.

Auto Repair Expert Properties

This section describes the expert properties that you can set for the Automatic Surface Repair mesher.

Minimum Face Quality	The equivalent property for Region Based Meshing is found as a node under **Reference Values > Automatic Surface Repair**. The *Minimum Face Quality* property for the auto surface repair is a value that ranges from 0 to 1, with 0 being the worst and 1 being perfect. The quality of a triangle is defined as 2*(r/R). R the radius of the circle that passes through the three corner points of the triangle and r is the radius of the circle that fits inside the triangle. The default value is `0.05`. The auto surface repair ensures that all triangles that the process produces have a quality equal to or better than the specified value. The mesher disregards feature edges where necessary, moves feature vertices, or avoids projecting vertices in order to remove or improve faces whose quality is below the specified value. Setting the value to `0` deactivates the option. Recommended values are in the range `0.01` to `0.1`. Values cannot exceed 0.1.

Surface Mesher Controls and Values

This section details the controls and values for all surface meshers.

Alignment Location

This control aligns the wrapper mesh lines relative to the global cartesian system. That is, you can direct the mesh lines along the x-, y- and/or z-direction.

This facility is useful when wrapping multiple regions; and the alignment of each wrapped mesh to its neighbor is important for merging and interfacing purposes.

Alignment Location Properties

Location	Specifies the location for aligning the wrapper mesh lines. The current scene highlights the current location of the mesh alignment point with a circular ball.

Automatic Surface Repair

See Automatic Surface Repair

Base Size

See Base Size

Basic Curvature

This value specifies the number of points around a circle on the surface mesh, which can refine the mesh on surface curvature. Increasing the value increases the relative refinement of the cells next to the curved surface.

The following diagram illustrates this value.



For pipes or tubes that have a relatively small diameter when compared to the rest of the mesh region, decrease this value to avoid over-refinement of the area.

Basic Curvature Properties

# Pts/circle	Specifies the number of triangles (for a surface) or cells (for a volume) are used to define a 360 degree cylindrical surface. Any other types of cylindrical surface use an equivalent number of triangles. For example, a half circle use half of this value and an ellipse would use more faces on tighter curvature. When used with the curvature deviation distance, this property is the minimum number of edges that are used around a complete circle.

CAD Projection

This control allows the mesher to project vertices back to the imported CAD surface during surface wrapping or surface remeshing. This capability results in a surface definition that is closer to the CAD geometry.

CAD projection is only available for the following CAD file types:

- Native CAD Formats
- Parasolid Transmit
- IGES
- STEP
- VDA

Projection back to an underlying CAD edge is possible only for feature edges that the surface import process defined.

For the wrapper, you can expect 20% or more CPU time for CAD projection. For the remesher, you can expect 300% or more CPU time.

Note: The surface wrapper cannot project some vertices, such as areas of gap closure or non-CAD data. The *Output* window provides a summary of the projected vertices at the end of the surface wrapping process.

CAD Projection Properties

Project to CAD	**Activated**	Projects the vertices to the original CAD surface.
	Deactivated	Does not project the vertices to the original CAD surface.

Curvature Deviation

The surface wrapper, surface remesher, and trimmed cell mesher use this distance to populate the surface mesh with an appropriate number of edges. This means that the diameter of a curve is accounted for in the mesh distribution.

An illustration of this distance is shown below:



To estimate an initial curvature deviation distance, measure the larger final mesh sizes on curved areas and multiply that distance by between 0.1 and 0.01.

Curvature Deviation Properties

Max # Pts/circle	Upper limit on the number of points that are permitted on a complete circle.

Curvature deviation distance	Maximum distance that is permitted between the center of any mesh edge and its associated input surface.

Edge Proximity

Edge Proximity allows you to specify a target minimum number of mesh faces that you require between opposing perimeter edges on a remeshed part surface.

To apply edge proximity refinement, STAR-CCM+ first calculates the distance between opposite perimeter edges on the input part surface. This distance is divided by the specified target number of faces. The result is then set as a target size for faces on the part surface. Assuming that this target size is greater than the Minimum Size, and less than the current target mesh size, the triangulation on the part surface shows refinement after remeshing.

If one part surface defines an entire geometry part, applying edge proximity has no effect since there are no surface edges present. You can either:

- Split the part surface appropriately and apply edge proximity to the required part surfaces.
- Use the Stop at Feature Edges option.

Edge Proximity is only available as a custom control for an automated mesh operation and is deactivated by default.

Edge Proximity Properties

Number of Faces	The target number of mesh faces on the surface between each given boundary.
Enable Search Ceiling	Allows you to specify a Search Ceiling distance. Surfaces with a proximity distance that is larger than this value are excluded from processing to avoid unnecessary refinement. Deactivated by default.
Search Floor	Surfaces with a proximity distance that is smaller than this value are excluded from processing. Therefore, unnecessary refinement is avoided across very thin patches or part surfaces and in sharp or tapered corners. Set to 0.0 (deactivated) by default.
Stop at Feature Edges	Determines if feature edges are added to the collection of edges for edge proximity

	computation. By default this is deactivated and part surface edges are used for edge proximity computation.	
	Activated	Feature edges are added to the collection of edges that are used for edge proximity computation.
	Deactivated	Feature edges are not added to the collection of edges that are used for edge proximity computation.
Stop at Patch Perimeters	Determines if patch perimeter edges are added to the collection of edges for edge proximity computation. By default this is deactivated and part surface edges are used for edge proximity computation.	
	Activated	Patch perimeter edges are added to the collection of edges for edge proximity computation.
	Deactivated	Patch perimeter edges are not added to the collection of edges for edge proximity computation.

Feature Angle

This value determines the surface angle where the wrapper generates a feature edge. Use this feature specification instead of generating feature curves in the input surface.

The surface wrapper produces a part curve (PBM)/feature curve (RBM) that is called **Edges from Wrapping** that contains any edges that did not exist in the starting import surface. In parts based meshing, the surface wrapper creates a tag (called **XYZ**) on the **Edges from Wrapping** part curve to differentiate the part curve that the wrapper creates from the pre-existing part curves.

Note: The wrapped surface retains any existing part curves (PBM)/feature curves (RBM) that you previously defined. However, they can be modified according to the new wrapped surface.

Feature Angle Properties

Angle	Specifies the minimum angle between faces that the surface wrapper defines as a feature.

Gap Closure Size

This value controls the filling of holes in the import surface during the surface wrapping process. You can specify the gap closure size as relative or absolute.

To close a hole, the specified gap closure size value must be bigger than the actual hole size. That is, the surface wrapper does not close holes with a dimension that is larger than the specified size.

The gap closure size cannot be smaller than the target surface size; the surface wrapper uses the larger of the two values to close gaps. This rule is also true when you use custom boundary values. This behavior prevents the surface wrapper from refining all of the surface down to the size of the smallest gap size.

Gap closure size works with the smallest wrapping volume option. The surface wrapper only closes gaps on surfaces that enclose a volume that is larger than the specified volume.

The surface wrapper creates a surface that is called **Gap Closure Faces**, which contains the faces that fill the hole.

Minimum Proximity

The minimum proximity specifies the minimum allowable distance between faces. Proximity is defined as a percentage of the average edge length of triangle that is checked.

Note: The default value of 0.05 is suitable for most cases. This means that any two faces cannot be closer than 5% of the average edge length for the initial face.

Minimum Proximity Properties

Minimum Proximity	Minimum allowable distance between two faces before STAR-CCM+ invokes automatic surface repair. A value of zero deactivates this property.

Minimum Quality

The minimum quality specifies the minimum allowable face quality.

Minimum Quality Properties

	Minimum allowable face quality before STAR-CCM+ invokes automatic surface repair. A value of zero effectively deactivates this property.

Minimum Surface Size

This value specifies the lower limit of edge lengths on the surface mesh. You can specify the minimum surface size as relative or absolute.

Decreasing the minimum surface size allows more refinement in regions where the mesher is trying to capture complex features. However, this refinement increases the CPU time and size of the surface mesh.

Maximum Surface Size

This value specifies the upper limit of edge lengths on the surface mesh. You can specify the maximum surface size as relative or absolute.

Increasing the maximum surface size reduces the number of faces in regions where there is little complexity on the surface. However, this coarsening can ignore features that you want to capture in the surface mesh.

Smallest Disconnected Surface

This control removes detached volumes from the surface mesh if their total number of faces is below this minimum.

Surface wrapping can generate many disconnected surfaces that you do not want in the surface mesh. This control allows you to remove these surfaces automatically.

Smallest Disconnected Surface Properties

Number of Faces	Specifies the minimum number of faces a wrapped volume can have before it is deleted from the surface mesh. If all disconnected surfaces have less than this number of faces, all surfaces are maintained. A value of zero effectively deactivates this property.

Smallest Wrapping Volume

This control ignores volumes in the wrapped surface if their volume is below this minimum. These volumes are not included in gap closure, or for the largest internal volume or the nth largest volume of interest options.

Smallest Wrapping Volume Properties

Smallest Wrapping Volume	Specifies the minimum volume that a wrapped volume can have before it is deleted from the surface mesh. A value of zero effectively deactivates this property.

Surface Curvature

This control includes cell refinement for the surface wrapper and surface remesher, using the minimum number of points around a circle (#Pts/circle) and the curvature deviation distance.

The number of points around a circle is always available, but you manually activate curvature deviation distance.

Surface Curvature Properties

Enable curvature deviation distance	**Activated**	Allows you to control surface curvature refinement with the curvature deviation distance control.
	Deactivated	Removes the curvature deviation distance control.

Surface Growth Rate

This control determines the rate at which face edge sizes vary from one face to its neighbor. The effect is only seen when there is transition from an area of refinement to the rest of the surface mesh.

As an example, the default value of 1.3 means that a face edge can be 1.3 times its neighbor length when growing from a short to a long edge length.

Surface Growth Rate Properties

Surface Growth Rate	Specifies the maximum rate of face edge transition. Decreasing this value generally produces a surface mesh with more faces.

Surface Proximity

This control provides mesh refinement on surfaces that are closed to one another. The increased number of faces in the surface mesh produces more volume mesh cells between the gap.

Surface Proximity Properties

# Points in a gap	Determines the local face size by dividing the distance from one face to another across a gap by this value. However, if the gap distance is less than the *Search Floor* or more than the *Search Ceiling*, there is no refinement.
Search Floor	Specifies the minimum gap size that the mesher includes in proximity, which prevents unwanted refinement. This property is useful when the surface contains overlapping faces or faces that are too close, due to complex detail or CAD artifacts.
Search Ceiling	Specifies the maximum distance that the mesher includes in proximity, which prevents unwanted refinement. This property is useful if a component has a variable thickness and you need an option to locally refine only those areas where the thickness is less than a specified value.

Surface Size

This control determines the method by which the surface size is defined. For the surface wrapper and surface remesher, it determines the face and edge sizes. For the trimmed cell mesher, it determines the edge size of the cells next to the surface.

Surface Size Properties

Relative/Absolute	**Relative to base**	Specifies that the surface sizes are a percentage of the Base Size.

	Absolute	Specifies that the surface sizes are absolute values.
Size Method	**Min and Target**	The mesher aims for face edges that meet the target size, regardless of the local face size of the input surface. Where there is mesh refinement, the face edges do not go below the minimum size.
	Min and Max	The mesher aims for faces edges that meet current local triangle size of the input surface. Where the local triangle size exceeds the minimum or maximum size, the mesh is refined or coarsened. Where there is mesh refinement, the face edges do not go below the minimum size.
	Min only	The mesher aims for faces edges that meet current local triangle size of the input surface. Where the local triangle size is less than the minimum size, the mesh is coarsened. Where there is mesh refinement, the face edges do not go below the minimum size.

Target Surface Size

This value specifies the edge length that the mesher aims for in the absence of any mesh refinement. You can specify the target surface size as relative or absolute.

Decreasing the target surface size generates more refined and detailed surface mesh. However, this refinement increases the CPU time and size of the surface mesh.

Volume of Interest

This control determines the volume that the surface wrapper extracts from the original geometry surface.

Volume of Interest Properties

Method	**Largest Internal**	Specifies that the wrapper extracts the largest internal, self-contained volume from the input geometry surfaces. In the following example, the largest internal volume is the red box, without the blue and green parts, as internal geometries within the volume are maintained.  To ignore internal volumes that you do not want in the surface mesh, use the smallest wrapping volume control.
	External	The *External* method specifies that the

		surface wrapper extracts the outer surfaces of all parts, that is, the external surfaces. The following example demonstrates this method:  In this example, all internal surfaces are omitted from the final wrapped surface. If they are part of the outer surface, the intersecting surfaces are included to form one volume.
	Seed Point	Specifies that a user-defined location determines the surface that the wrapper extracts. The resultant mesh is the surface that surrounds this location. In this example, two seed points are placed inside the blue and green volumes. The result from the wrapper is two closed surfaces that are shown to the right of the diagram.

		
	Nth Largest	Specifies that the wrapper extracts a particular surface according to volume size. The surfaces are ranked from largest to smallest, and the user-defined Volume Number determines the output surface. Geometries that intersect one another are considered as both a single volume plus individual volumes when assessing their volumetric size. The following example shows the second largest volume that is extracted from the input geometry, which is the red box minus all other parts.  To ignore internal volumes that you do not want in the surface mesh, use the smallest wrapping volume control.

Wrapper Scale Factor

This control scales all refinement sizes (for example, surface size and volumetric controls) by a common factor during the surface wrapping process.

This facility allows you to generate a wrapped surface with different controls to the other meshers, so that you can generate a finer surface mesh without running the wrapper independently. It also allows quick and easy alteration of all the surface wrapper controls.

The wrapper scale factor is not applied to the following controls:

- Smallest Wrapping Volume
- Smallest Disconnected Faces
- Gap Closure Size
- Contact Prevention
- Wrapper Feature Angle

Wrapper Scale Factor Properties

Percentage	Specifies the scale factor by which all surface wrapper controls are scaled. A higher or lower value results in a coarser or finer mesh respectively.

See Also

What Is The CAD Association Color Mode

Volume Meshers

STAR-CCM+ contains different types of meshing models that can be used to generate a volume mesh starting from a suitably prepared surface mesh.

The five types of meshing models are:

- Tetrahedral - tetrahedral cell shape based core mesh
- Polyhedral - arbitrary polyhedral cell shape based core mesh
- Trimmed - trimmed hexahedral cell shape based core mesh
- Thin Mesh - tetrahedral or polyhedral based prismatic thin mesh
- Advancing Layer Mesh - polyhedral core mesh, with in-built prismatic layers that advance inward from a polygonal surface mesh.

Prismatic near wall layers can be included with the first three types of mesh that are listed above by including the prism meshing model as part of the volume meshing process. For multi-region cases, a conformal mesh at the interface between one region and the next is currently possible for the tetrahedral and polyhedral mesh types as part of the meshing process. This option is currently not supported for the trimmed cell mesher.

Different mesh types can also be used for different mesh regions if desired by using a different mesh continuum for each region. A conformal mesh interface at inter-region boundaries is not supported in this case.

Volumetric controls in the shapes of bricks, spheres, cylinders, and cones can be used during the volume meshing of the core mesh to increase or decrease the local cell density. Arbitrary shaped volumes that are imported as parts can also be used to define the refinement volume. Feature refinement is automatically included in the case of the trimmer meshing model.

The extruder mesher can be used with the core mesh methods to extend the original domain by extruding a boundary. The extruder produces

orthogonal cell layers from the chosen boundaries. An example use of the extruder mesher would be to extend an inlet duct upstream.

The generalized cylinder mesher can be used with the polyhedral meshing model to generate a mesh on elements in the geometry that are roughly cylindrical. The generalized cylinder mesher generates extruded orthogonal cells along the length of cylindrical boundaries.

In cases where a previous solution exists and the mesh is modified, the solution is automatically mapped to the new mesh in order to restart the simulation.

Various field functions are created during volume meshing operations; some of which are documented in the appropriate reference section.

Once the volume mesh has been built, it can be checked for errors and exported to other packages if desired.

Volume Meshers Overview

This section introduces volume meshing.

Choosing a Volume Mesh Model

In general, the choice of which mesh type model to use for the core volume mesh is dependent on several factors. Some of these factors are:

- the turn around time available for building the mesh
- the desired solution accuracy and convergence rate
- the amount of memory available on your computer
- whether you are building a single or multi-region mesh
- the quality of the starting surface mesh
- whether the geometry is predominantly thin or not.

However, some of these factors are not mutually exclusive. For example, the tetrahedral mesher is the fastest volume mesh generator (in terms of number of cells that are generated per minute) and the one that uses the least amount of memory. However, in terms of solution quality, approximately five to eight times the number of cells are required to produce the same accuracy as the equivalent polyhedral or trimmed cell mesh. The additional solution time, memory, and increased time to convergence when solving on this type of mesh must be factored in against any time gains that are made during the meshing process.

In terms of general accuracy for a given number of cells, the trimmed and polyhedral cell type meshes produce more accurate solutions when compared to a tetrahedral mesh. Mesh independent solutions can be obtained for all three mesh types if the density is increased to the appropriate level.

If you are building a multi-region grid where a conformal mesh is required at interfaces, then use the tetrahedral or polyhedral meshers with a single mesh continuum. However, the trimmer model does not produce conformal mesh interfaces and a different continuum must be used for each region. Alternatively the per-part or per-region meshing option can be activated instead to apply the same continuum to all regions.

Both the tetrahedral and polyhedral volume meshers are directly dependent on the quality of the starting surface triangulation. In other words, a bad quality starting surface leads to a bad quality volume mesh. For complex meshing problems however, the overall cell quality of the polyhedral meshes can suffer when compared to a tetrahedral mesh starting from the same surface. The trimmer model is not directly

dependent on the surface quality of the starting surface and as such is more likely to produce a good quality mesh for most situations.

For geometries that are predominantly composed of thin surfaces, it is recommended that the thin mesher is used. The thin mesher produces a better quality prismatic type mesh than would be created using one of the other three volume meshers.

Polyhedral Mesher

Polyhedral meshes provide a balanced solution for complex mesh generation problems.

They are relatively easy and efficient to build, requiring no more surface preparation than the equivalent tetrahedral mesh. They also contain approximately five times fewer cells than a tetrahedral mesh for a given starting surface. Multi-region meshes with a conformal mesh interface are allowed.

The polyhedral meshing model utilizes an arbitrary polyhedral cell shape in order to build the core mesh. In STAR-CCM+, a special dualization scheme is used to create the polyhedral mesh from an underlying tetrahedral mesh, which is automatically created as part of the process. The polyhedral cells that are created typically have an average of 14 cell faces.

An example polyhedral core mesh is shown below:



The polyhedral core mesh density can be increased or decreased by using the volume mesh density factors. Volumetric controls that are chosen from a range of prescribed shapes can also be included to increase or decrease the mesh density locally. Global refinement can be applied to increase the density for every cell, and field function based refinement can be used to create local density variations which are based on a mesh size table. When the polyhedral mesher is used with a periodic interface, then it attempts to produce a conformal mesh for the boundary pair where possible.

For geometries containing cylindrical elements, the polyhedral mesher can be used with the generalized cylinder mesher, which generates extruded orthogonal cells along the length cylindrical sections. For geometries containing thin walled structures, the polyhedral mesher can be used with the embedded thin mesher, which generates prismatic type cells in areas that are considered to be thin.

Polyhedral meshes can normally be generated with free floating baffles. However, issues are sometimes seen when the baffles do not have prismatic layers. If you encounter such issues, activate the prism mesher for the baffles so that at least one layer of prism cells is generated to help alleviate problems with cell generation at baffle edges.

See Also

Generalized Cylinder Mesher

Polyhedral Mesher Expert Properties

This section describes the polyhedral mesher expert properties.

The polyhedral mesher expert properties that are common with the tetrahedral mesher are documented in the volume meshers control and values section.

Run Optimizer (polyhedral mesh only)	Performs additional optimization of vertices during the dualization process, resulting in higher-quality meshes. Although this property incurs an additional CPU overhead, you do not generally deactivate it.
Include Refinement (RBM)	Provides global refinement of the mesh by one or two refinement levels. When activated, you set the refinement level on the **Refinement Level** node under the **Mesh > Reference Values** manager.

Global Polyhedral Mesh Refinement

The *Include Refinement* option on the **Polyhedral Mesher** node provides a quick method to increase the cell count in a polyhedral mesh globally.

This option is only available for Region Based Meshing. Ticking this option adds the **Refinement Level** node that is shown below.



Two refinement levels are available:

- Refinement Level 1: Splits each polyhedral cell into approximately 6 or 7 new cells.
- Refinement Level 2: Adds another level; so splits each original polyhedral cell into approximately 40 - 50 new cells.

Therefore, the overall cell count of the polyhedral mesh increases by a factor of 6 - 7 for one refinement level, and by a factor of 40 - 50 for 2 refinement levels.

Refining a coarse mesh takes less computational time to increase cell count when compared to creating a fine mesh from the initial surface mesh.

The thickness and number of the prism layers remain unchanged by the global refinement of the polyhedral mesh.

Also, as global refinement splits the polyhedral cells, it does not affect the surface mesh. Therefore, you do not need to run the surface remesher with global refinement and the surface mesh is not as refined as the volume mesh.

As an example of this global refinement, observe the changes that result from one level of refinement, and then from two levels of refinement:

- Initial coarse mesh.



- Refinement level 1.



- Refinement level 2.



Behavior of Global Refinement with Field Function Refinement

This section summarizes the behavior of the polyhedral mesher when using global refinement with field function mesh refinement.

For one level of global refinement, consider a cell in the mesh where:

- FF — is the cell size that is prescribed by field function mesh refinement.
- C — is the cell size before refinement.

If FF < C, global refinement splits the cell according to the refinement level algorithm. The final cell size does not necessarily meet the requirements of the field function refinement.

For two levels of global refinement, consider the same cell where:

- FF — is the cell size that is prescribed by field function mesh refinement.
- c — is the cell size after one level of global refinement.

Similarly, if FF < c, global refinement splits the cell according to the refinement level algorithm. Again, the final cell size does not necessarily meet the requirements of the field function refinement.

The following example refines a mesh section to less than the original cell size, coarsens the rest of the mesh, and then applies global refinement. The original cell size within the mesh is approximately 0.07 m as shown below.



The applied field function refines the mesh to a cell size of 0.02 m in the center and coarsens the mesh to 0.12 m in the rest of the mesh.



The following image shows the mesh when only field function mesh refinement is applied. The coarse section does not have a consistent cell size of 0.12 due to growth rates.



The following image shows the mesh after applying field function refinement and global refinement. In the refined section, the mesher applies global refinement and the cells are approximately six or seven times smaller. In the coarse sections, the mesher does not apply global refinement and the cells still have an approximate cell size of 0.07 m.


Scene 1
STAR-CCM+
Cell Size
0.00
0.023
0.047
0.070
0.093
0.12
0.14

Tetrahedral Mesher

Tetrahedral meshes provide an efficient and simple solution for complex mesh generation problems.

Out of the three provided models, the tetrahedral is the fastest and uses the least amount of memory for a given number of cells. Multi-part or multi-region meshes with a conformal mesh interface are allowed.

The tetrahedral meshing model utilizes the tetrahedral cell shape in order to build the core mesh. In STAR-CCM+, the Delaunay method is used to construct the mesh, which iteratively inserts points into the domain, forming high-quality tetrahedra in the process. The resulting mesh strictly conforms to the triangulation of the surface at the domain boundary, so the quality of the original surface mesh must be good to ensure a good quality volume mesh.

An example tetrahedral core mesh is shown below:



The tetrahedral core mesh density can be increased or decreased by using the volume mesh density factors. Volumetric controls that are chosen from a range of prescribed shapes can also be included to increase or decrease the mesh density locally. When the tetrahedral mesher is used with a periodic interface, it attempts to produce a conformal mesh for the boundary pair where possible.

Tetrahedral Mesher Expert Properties

The tetrahedral mesher expert properties that are common with the polyhedral mesher are documented in the volume meshers control and values section.

Trimmed Mesher

The trimmed cell mesher provides a robust and efficient method of producing a high-quality grid for both simple and complex mesh generation problems.

It combines a number of highly desirable meshing attributes in a single meshing scheme:

- predominantly hexahedral mesh with minimal cell skewness
- refinement that is based upon surface mesh size and other user-defined refinement controls
- surface quality independence
- alignment with a user specified coordinate system.

The trimmed cell mesher is available for both the Parts Based Meshing (PBM) and Region Based Meshing (RBM) approaches.

By default, the trimmer meshing model utilizes a template mesh that is constructed from hexahedral cells from which it cuts or trims the core mesh using the starting input surface. The template mesh contains refinement that is based on the local surface mesh size and local refinement controls. Growth parameters can be used to control the transitioning of the mesh cell sizes from small to large both at the surface and far field. A maximum and/or minimum cell size can be supplied as well to control the upper and lower cell size bounds. The template can be aligned in any direction in a user specified cartesian coordinate system.

An additional feature useful in modeling external aerodynamic flows is the ability to refine cells in a wake region. This region is generally the volume of fluid at the rear of a moving body.

The resulting mesh is composed predominantly of hexahedral cells with trimmed cells next to the surface. Trimmed cells are polyhedral cells but can usually be recognized as hexahedral cells with one or more corners and/or edges that are cut off.

An example trimmed cell mesh is shown below:



The input values used for the trimmer model can be set on four different levels:

- global
- wake refinement region
- part surface or boundary
- part or feature curves

Volumetric controls can also be included to increase or decrease the mesh density locally in the template using a range of prescribed shapes to define the extents.

The current implementation of the **Trimmed Cell Mesher** (PBM) / **Trimmer** (RBM) model is restricted to one trimmed volume mesh per automated mesh operation (PBM) or mesh continuum (RBM). In other words, a different automated mesh operation or mesh continuum must be used for each trimmed volume mesh that you wish to generate. Alternatively, the per part (PBM) or per-region (RBM) meshing option can be activated to mesh parts or regions independently. At part or region interfaces a non-conformal interface is generated. Additionally,

interfaces can only be used with the trimmer meshing model when the boundaries in the interface belong to the same region. For example, an in-place interface could be used to model a baffle in the flow domain, or a periodic interface for periodic flow condition. In all instances, a conformal trimmed mesh across the interface is not guaranteed.

The trimmer model also has a built-in ceiling to control the size of the template cells next to surfaces. This ceiling is to avoid grossly de-featuring the model and generating bad quality cells. At present, the maximum surface template cell size is fixed to 150% of the average triangle edge length size in the local surface. This factor does not affect cells away from the surface.

Trimmed Mesher Properties

This section covers the trimmed mesher properties for both Parts Based Meshing (PBM) and Region Based Meshing (RBM).

Screenshots are of the Parts Based Meshing approach. Where the names of nodes or properties differ between the approaches, this is indicated by the acronym PBM for Parts Based Meshing and RBM for Region Based Meshing. When using Parts Based Meshing the properties are found under the **Geometry** > **Operations** > **Automated Mesh** > **Meshers** > **Trimmed Cell Mesher** node. When using Region Based Meshing the properties are found under the **Continua** > **Mesh** > **Models** > **Trimmer** node.

Coordinate System	The trimmer template mesh is aligned with the global cartesian laboratory coordinate system by default. If the imported surface is not aligned to this coordinate system, then it can be advantageous to realign the template using a user-defined cartesian coordinate system instead. Once a new coordinate system has been defined, it can be selected using the *Coordinate System* pull-down menu in the *Properties* dialog:

	 The trimmed cell mesh is now generated based on the new local coordinate system instead of the global one.
Perform Mesh Alignment (PBM) / Do Mesh Alignment (RBM)	The trimmer mesh alignment option is turned off by default, meaning that the template mesh is centered and aligned based on the global bounds and model alignment in the specified coordinate system. In certain circumstances, such as free surface calculations, it is desirable to align a plane of mesh faces with a physical plane corresponding to a salient feature such as an initial free surface. Such alignment is done, not only in order to make the physical model setup and initialization easier, but also to improve accuracy. Turning the *Perform mesh alignment* (PBM) / *Do mesh alignment* (RBM) option on allows the mesh location coordinates to be specified for the chosen trimmer coordinate

	system, so that the mesh lines coincide with that location: 
	With the perform mesh alignment option turned on, the **Mesh Alignment Location** node under **Default Controls** (PBM)/ **Reference Values** (RBM) can be used to specify a location in the currently specified coordinate system. The trimmer mesh is then aligned so that this location is on a cell corner vertex in the resultant mesh. Selecting the **Mesh Alignment Location** node allows the mesh alignment values to be set:

	 The *Location* coordinate can be specified in any of the allowable units. If mesh alignment is only required in one direction, leave the other two location coordinate values set to 0. The current location of the grid alignment point is highlighted in the scene using a circular ball once the **Mesh Alignment Location** has been set.
Template Mesh Type (RBM)	The *Template Mesh Type* property allows one of two different cell shapes to be used for the underlying template mesh. The two cell shapes are: • hexahedral • dodecahedra

These selections are shown below for the Region Based Meshing approach. Options are found under the **Mesh** > **Models** node for the appropriate Continuum:



By default, the *Template mesh type* is set to *Hexahedra* meaning that hexahedral shaped cells are used to construct the template. Changing the option to *Dodecahedra* results in a template that is built from regular rhombic dodecahedral cells. A regular rhombic dodecahedral cell has 12 faces, with each face being a rhombus of 4 equal length sides. For PBM, by default, the equivalent template mesh type property is not available to modify in the properties window as it is set to Hexahedra automatically.

To specify the template mesh type as Dodecahedra in PBM, run a macro such as:

```
AutoMeshOperation autoMeshOperation_0
=
((AutoMeshOperation)
simulation_0.get(MeshOperationManager
.class).getObject("Automated Mesh"));
```

	`TrimmerAutoMesher trimmerAutoMesher_0 = ((TrimmerAutoMesher) autoMeshOperation_0.getMeshers().getObject("Trimmed Cell Mesher")); trimmerAutoMesher_0.getTemplateMeshOption().setSelected(TemplateMeshOption.DODECAHEDRA);` The mesh alignment option cannot be used in conjunction with the dodecahedra template mesher.
Size Dependent Growth Rates (PBM) / Template Mesh Growth Type (RBM)	The *size dependent growth rates / template mesh growth type* property determines the growth rate options that are used to control the refinement next to surfaces and in the core mesh. Two options are provided: • **All cells** (PBM) / **Simple** (RBM) • **Small and large cells** (PBM) / **Two Level** (RBM) These selections are shown below for the Parts Based Meshing approach:



By default, the **Size dependent growth rates** (PBM) / **Template mesh growth type** (RBM) is set to **All Cells** (PBM) / **Simple** (RBM).

Trimmed Mesher Expert Properties

This section covers the trimmed mesher expert properties for both Parts Based Meshing (PBM) and Region Based Meshing (RBM).

Screenshots are of the Parts Based Meshing approach. Where the names of nodes or properties differ between the approaches, this is indicated by the acronym PBM for Parts Based Meshing and RBM for Region Based Meshing. When using Parts Based Meshing the expert properties are found under the **Geometry** > **Operations** > **Automated Mesh** > **Meshers** > **Trimmed Cell Mesher** node. When using Region Based Meshing the expert properties are found under the **Continua** > **Mesh** > **Models** > **Trimmer** node.

Choose Field Function Refinement Table (PBM) / Choose a Mesh Size Table (RBM)	The **Choose Field Function Refinement Table / Choose a Mesh Size Table** property allows for the refinement of the mesh using field functions, based on an existing solution.
Run Optimizer	The *Run Optimizer* option, which is off by default, allows additional optimization of vertices during the trimming process. This option should only be turned on if a poor quality mesh was originally obtained using the trimmer model and can help improve the resultant quality. An additional CPU overhead is introduced by using this option.

Setting Trimmer Default Controls (PBM) / Reference Values (RBM)

The default controls (PBM) or reference values (RBM) are the global level controls for the trimmer.

The following reference values are applicable for the trimmer:

- base size
- maximum cell size
- mesh alignment location
- surface size
- volume growth rate (PBM) / template growth rate (RBM)

These options can be set within the **Default Controls** node (PBM), or the **Reference Values** node (RBM) in the simulation tree:



Volume Growth Rate (PBM) / Template Growth Rate (RBM)

Growth rates control the rate at which cell sizes increase from one cell size to another within the trimmed cell mesh.

A fast growth rate increases cell sizes rapidly; a slow growth rate means that the mesher uses multiple cell layers to provide a gradual transition.

The growth rate defines cell size transition by specifying a minimum number of cell layers of each size. The following diagram illustrates this method:



Medium Growth Rate (2 Layers per Cell Size)



Slow Growth Rate (4 Layers per Cell Size)

To refine cell sizes and maintain connectivity, the mesher subdivides cells into eight smaller cells. In some cases, this method of splitting means that the mesher does not produce the minimum number of cell layers, so the mesh has extra layers. This is shown in the following example. Here, the mesh must meet the following criteria:

- Surface cell size of 1 unit.
- Core cell size of 8 units.
- Fast growth rate, so at least one layer of each cell size (1, 2, 4, and 8 units).

Also, this example uses a two-dimensional mesh, so larger cells split into four smaller cells.



- Image 1: The mesher generates a uniform mesh with cell size 8.
- Image 2: The mesher splits the beige cells in Image 1 to make the cell size closer to surface size on the red boundary.
- Image 3: The mesher performs a second split on the beige cells in Image 2 so that the cell sizes match the surface size.
- Image 4: The blue and green cells in Image 3 are cell layers of 1 and 2 units respectively; however there are no cell layers of 4 units. Therefore, the mesher splits the beige cells in Image 3 to create cells with a size of 4 units.

This final image shows that the splitting of larger cells can result in a number of cell layers that does not match the minimum.

The *Default Growth Rate* controls the rate of size change in the core volume mesh. You set this property for the automated mesh operation (PBM) or mesh continuum (RBM), within the **Default Controls** (PBM) or **Reference Values** (RBM) manager respectively. Two methods of control are available:

- **Volume Growth Rate** (PBM).
- **Template growth rate** (RBM).

The *Surface Growth Rate* (PBM) / *Boundary Growth Rate* (RBM) controls the rate of size changes between the cells that are next to surfaces and

the cells in the core. You can set this property at a global or local level on a part surface, boundary, or volumetric control.

This section contains the following topics:

Volume Growth Rate (PBM) / Template Growth Rate (RBM) Properties

The **Volume Growth Rate** node (PBM) or **Template Growth Rate** node (RBM) controls the stepping or transition of one cell size to the next within the template that is used for the trimmer.

When the **Trimmed Cell Mesher** (PBM) / **Trimmer**(RBM) node is selected, the *Properties* window displays the *Size dependent growth rates* (PBM) / *Template mesh growth type* (RBM) as**All cells** (PBM) /**Simple** (RBM). Maintaining this default setting then selecting the **Volume Growth Rate** node (PBM) displays the properties as shown below:



Selecting the **Template Growth Rate** node (RBM) with the default **Trimmed Cell Mesher** settings maintained displays similar properties.

Properties of the Volume Growth Rate (PBM) / Template Growth Rate (RBM)

Default Growth Rate	This option controls the stepping from one cell size to the next within the core mesh. This option can be thought of as non-feature, far field growth control for the core mesh.	
	Fast	Default setting, which means that the fastest transition from the smallest cell size to the largest is used. This

		setting usually equates to having only one layer of each cell size for each transition, which results in the least number of cells being created for the supplied input values.
	Medium	Gives a minimum of two equal sized cell layers per transition.
	Slow	Gives a minimum of four equal sized cell layers per transition.
	Very Slow	Gives a minimum of eight equal sized cell layers per transition.
	Custom	Adds a Custom Default Growth Rate node on which to set a user-defined number of cell layers per transition.
Surface / Boundary Growth Rate	This option determines the number of cell layers to grow next to surfaces in the template.	
	Disable/None	Default setting, which means that the default growth rate is also applied to cell layers next to surfaces.
	Fast	Sets one cell layer next to surfaces.
	Medium	Sets two cell layers next to surfaces.
	Slow	Sets four cell layers next to surfaces.
	Very Slow	Sets eight cell layers next to surfaces.
	Custom	Adds a Custom Boundary Growth Rate node on which to

		set a user-defined number of cell layers next to surfaces.

Default Growth Rate

This section describes the properties that you can set for the Default Growth Rate.

If the Size dependent growth rate (PBM) / Template mesh growth type (RBM) property is set to **Small and large cells** (PBM)/**Two Level** (RBM) for the **Trimmed Cell Mesher**(PBM) / **Trimmer** (RBM) meshing model, the **Volume Growth Rate**node (PBM) or**Template Growth Rate** node (RBM) has no properties of its own. Instead, it has two child nodes, one for the **Default Growth Rate**, and another for the **Surface Growth Rate** (PBM) or **Boundary Growth Rate (RBM)**.



Properties on these nodes give a finer level of control over cell size transitions than is possible with the simple template growth rate.

Properties of the Default Growth Rate

Small Cell Cutoff Size Type	Sets the size type for the cutoff size.	
	Relative to Base	Set the small cell cutoff size relative to the base size.
	Absolute	Set the small cell cutoff size as absolute value with length units.

Default Growth Rate	This option controls the stepping from one cell size to the next within the core mesh. This option can be thought of as non-feature, far field growth control for the core mesh.	
	Fast	Default setting, which means that the fastest transition from the smallest cell size to the largest is used. This setting usually equates to having only one layer of each cell size for each transition, which results in the least number of cells being created for the supplied input values.
	Medium	Gives a minimum of two equal sized cell layers per transition.
	Slow	Gives a minimum of four equal sized cell layers per transition.
	Very Slow	Gives a minimum of eight equal sized cell layers per transition.
	Custom	Adds a Custom Default Growth Rate node on which to set a user-defined number of cell layers per transition.
Small Cell Growth Rate	This growth rate determines the growth rate in the vicinity of small cells in the template that are less than or equal to the cutoff size.	
	Fast	Default setting, which means that the fastest transition from the smallest cell size to the largest is used. This setting usually equates to having only one

		layer of each cell size for each transition, which results in the least number of cells being created for the supplied input values.
	Medium	Gives a minimum of two equal sized cell layers per transition.
	Slow	Gives a minimum of four equal sized cell layers per transition.
	Very Slow	Gives a minimum of eight equal sized cell layers per transition.
	Custom	Adds a Custom Small Cell Growth Rate node on which to set a user-defined number of cell layers per transition.

Surface Growth Rate (PBM) / Boundary Growth Rate (RBM)

This section describes the properties of the Surface Growth Rate (PBM) / Boundary Growth Rate (RBM) node

The properties of the **Surface Growth Rate** (PBM) / **Boundary Growth Rate** (RBM) node give access to additional template growth controls for cells next to surfaces of the mesh. The **Surface Growth Rate** (PBM) / **Boundary Growth Rate** (RBM) node is only added under the **Surface Growth Rate** (PBM) / **Template Growth Rate** (RBM) node when the **Small and large cells** (PBM) / **Two Level** option is chosen on the **Trimmed Cell Mesher** (PBM) / **Trimmer** (RBM) meshing model.

Properties of the Surface Growth Rate (PBM) / Boundary Growth Rate (RBM)

Small Cell Cutoff Size Type	Sets the size type for the cutoff size.	
	Relative to base	Sets the small cell cutoff size type relative to the base size.

	Absolute	Sets the small cell cutoff size type as an absolute value with length units.
Surface Growth Rate (PBM) / **Boundary Growth Rate** (RBM)	This setting controls the number of cell layers to use in the transition from cells next to surfaces to cells in the core volume.	
	The trimmer **Surface / Boundary Growth Rate** option determines the number of cell layers to grow next to surfaces in the template.	
	Disable (PBM) / **None** (RBM)	Default setting, which means that the default growth rate is also applied to cell layers next to surfaces.
	Fast	Sets one cell layer next to surfaces.
	Medium	Sets two cell layers next to surfaces.
	Slow	Sets four cell layers next to surfaces.
	Very Slow	Sets eight cell layers next to surfaces.
	Custom	Adds a Custom Surface / Boundary Growth Rate node on which to set a user-defined number of cell layers per transition.
Small Cell Surface Growth Rate (PBM) / **Small Cell Boundary Growth Rate** (RBM)	This growth rate determines the growth rate in the vicinity of small cells in the template that are less than or equal to the cutoff size.	
	Fast	Default setting, which means that the fastest transition from the smallest cell size to the largest is used. This setting usually equates to having only one layer of each cell size

		for each transition, which results in the least number of cells being created for the supplied input values.
	Medium	Gives a minimum of two equal sized cell layers per transition.
	Slow	Gives a minimum of four equal sized cell layers per transition.
	Very Slow	Gives a minimum of eight equal sized cell layers per transition.
	Custom	Adds a Custom Small Cell Growth Rate node on which to set a user-defined number of cell layers per transition.

Relative or Absolute Cutoff Size

You can set a relative or absolute small cell cutoff size.

It is possible that refinements due to the initial surface mesh size, or volumetric controls can result in small cells whose size is less than the target size in the template. If the small and large cells growth rate (PBM), or the Two Level (RBM) property has been activated for the **Trimmed Cell Mesher** (PBM) / **Trimmer** (RBM) model, then a relative or absolute small cell cutoff size can be set for both the **Default Growth Rate** and the **Surface Growth Rate** (PBM) / **Boundary Growth Rate** (RBM). For cells smaller than the small cell cutoff size, the trimmer applies the small cell growth rate to create transition cell layers between the small cells and the rest of the mesh.

The small cell growth rate is set as a property of the nodes for **Default Growth Rate** or **Surface Growth Rate** (PBM) / **Boundary Growth Rate** (RBM). On surfaces or boundaries, the growth rate applies to transition cell layers that are grown from small cells on the surface. Elsewhere, the default small cell growth rate applies.

Custom Growth Rate

If the built-in choice of growth rates fails to provide sufficient control over the number of cell layers that are used for transition between sizes, a custom option can be selected instead.

This option adds a custom growth rate child node whose single property sets the number of cell layers to use in size transitions. An example of this node for a custom default growth rate in the **Reference Values** section is shown below:



Properties of the Custom Growth Rate

Number of Growth Layers	This property sets the number of cell layers to use in the transition from one cell size to another. In the case of a custom boundary growth rate, this property applies to the transition in size between cells on a surface and cells in the core.

Prism Layer Mesher

The prism layer mesh model is used with a core volume mesh to generate orthogonal prismatic cells next to wall surfaces or boundaries. This layer of cells is necessary to improve the accuracy of the flow solution.

A prism layer is defined in terms of:

- Its thickness
- The number of cell layers within it
- The size distribution of the layers
- The function that is used to generate the distribution -- either by geometric progression or hyperbolic tangent.

These general properties can be defined globally within an automated mesh operation or a mesh continuum, and can also be defined locally on part surfaces, regions, or boundaries, and within volumetric controls. An example volume mesh showing two prism layers next to the wall surface is shown below:



By default, prism layers are not created next to boundaries (or associated part surfaces) unless they are of the wall boundary type. This setting can be changed on a boundary by boundary, or surface by surface basis.

When the prism layer meshing model is activated, its default and expert properties can be accessed by selecting the **Prism Layer Mesher** node. This node can be found under the **Meshers** node for the **Automated Mesh** operation if using Parts Based Meshing, or under the **Models** node for the mesh continuum if using Region Based Meshing.



Why Is a Prism Layer Mesh Required?

Prism layers allow the solver to resolve near wall flow accurately, which is critical in determining not only the forces and heat transfer on walls, but also flow features such as separation. Separation in turn affects integral results such as drag or pressure drop. Accurate prediction of these flow features depends on resolving the velocity and temperature gradients normal to the wall. These gradients are much steeper in the

viscous sublayer of a turbulent boundary layer than would be implied by taking gradients from a coarse mesh. Using a prism layer mesh allows you to resolve the viscous sublayer directly if the turbulence model supports it (low y+ ~1). Alternatively, for coarser meshes it allows the code to fit a wall function more accurately (high y+ > 30).

The turbulence modeling approach that is used and the desired fidelity of the physics determine the thickness, number of layers and distribution of the prism layer mesh. Depending on the Reynolds number, a turbulent shear layer requires in excess of 10-20 cells in the cross-stream direction for accurate resolution of the turbulence flow profiles. To resolve the viscous sublayer (that is, low y+ wall treatment), more cells are required. If only the gross flow features (such as a first-order estimate of skin friction) are required, coarser, high y+ wall-function type meshes with just a few prism layers can yield acceptable results.

Prism layers do not only provide near wall mesh density, they also allow high-aspect-ratio cells to be used, thus providing better cross-stream resolution without incurring an excessive stream-wise resolution.

Prism layers also reduce numerical diffusion near the wall. Numerical diffusion is a discretization error that smears discontinuities and steep gradients in a finite volume advection scheme. Numerical diffusion is minimized when the flow is aligned with the mesh. The use of prism layers greatly enhances accuracy as a result.

How is the Prism Layer Generated?

Before the core volume mesh is created, a subsurface is generated at the specified prism layer thickness values, in effect “shrinking” (for internal flows) or “expanding” (for external flows) the starting surface. The core mesh is built using this subsurface. The prism layer mesh is then generated by extruding the cell faces from the core mesh back to the original starting surface.

When the prism mesher is used with a periodic interface, then it attempts to produce a conformal mesh for the boundary pair where possible. The inclusion of a periodic interface overrules the boundary type that is specified for the boundary pair in this instance. Instead, the interface level prism layer option determines whether a prism mesh is produced or not. By default, no prism layers are produced on interfaces.

See Also

Advancing Layer Mesher
Generalized Cylinder Mesher

Controlling the Construction of the Prism Layer

The properties of the prism layer mesher model control different aspects of prism layer construction.

They can be split into two categories: the properties that control the general prism layer construction (such as the number of layers); and the properties that govern the behavior of the mesher when constructing a prism layer is difficult.

The following properties that govern general prism layer construction:

- Setting the number of cell layers
- Setting the thickness of the prism layer
- Setting prism layer stretching
- Setting transition ratios between prism layers and trimmed mesh

Global Controls

This section introduces the properties that control the general construction of the prism layer.

Properties that control the general construction of the prism layer can be set globally using the nodes within the **Automated Mesh > Default Controls** node when using Parts Based Meshing. If using Region Based Meshing, the same controls are within the **Reference Values** node of the parent mesh continuum.

Number of Prism Layers

The number of prism layers parameter controls the number of cell layers that are generated within the prism layer on a part surface or boundary.

This number is set using the **Number of Prism Layers** node.



Number of Prism Layers Properties

Number of Prism Layers	This property sets the number of prism layers to grow unless over-ridden by local settings or quality limits. The default value for *Number of Prism Layers* is to 2. Any whole value equal to or greater than 1 can be entered. The example below shows a prism mesh with three layers.

	 When specifying custom values, this node is added to the **Custom Controls > [Surface] > Controls** node for Parts Based Meshing. When using Region Based Meshing it is added to the **Mesh Values** node of a region or boundary.

Prism Layer Total Thickness

This section describes the Prism Layer Total Thickness properties.

The prism layer thickness controls the total overall thickness of all the prism layers.



Prism Layer Total Thickness Properties

Prism Layer Total Thickness	Selecting the **Prism Layer Thickness** node allows the choice of a relative or absolute *Size type* to be made. Depending on the **Size type** selection, a **Relative Size** or **Absolute Size** can be set to provide the value of the thickness. The example below shows how the prism layer thickness is measured.

	 When specifying custom values, this node is added to the **Custom Controls > [Surface] > Controls** node for Parts Based Meshing. When using Region Based Meshing it is added to the **Mesh Values** node of a region or boundary.	
	Relative to Base	Sets the prism layer total thickness relative to the base size.
	Absolute	Sets the prism layer total thickness as absolute value with length units.

Prism Layer Stretching

Prism layer stretching sets the target growth rate of successive prism layers away from the wall.

Prism Layer Stretching Properties

Prism Layer Stretching	The default **Prism Layer Stretching** value is 1.5.

Maximum Core Prism Transition Ratio

When using the prism layer mesh model with the trimmer model, you can control the ratio between the trimmed and prism layer cell sizes.

This feature, along with the **Near Core Layer Aspect Ratio** option, is designed to help you generate smoother cell volume transitions between the trimmed and prism layer meshes. It is set by activating the checkbox for the **Limit cell size by prism layer thickness** in the *Properties* window of the **Maximum Core/Prism Transition Ratio** node.



Maximum Core/Prism Transition Ratio Properties

Maximum Size/ Thickness Ratio	STAR-CCM+ adds a **Size/Thickness Ratio** node, which lets you set the limiting value for the ratio between the two type of mesh cells. The default value is 5.0, that is, trimmed cells are refined so that they are at most five times the size of the neighboring prism layer cells. The minimum ratio value that you can enter is 2.0. Smaller values result in increased cell count. The effect of using the **Size/Thickness Ratio** option is shown below. The left image is with this feature turned off, and the right shows

	that the additional refinement has reduced the volume change between the trimmed and prism layer cells. 

Customizing Prism Layers

Custom controls can be set for particular part surfaces, individual regions, or boundaries.

Custom controls can be set for particular part surfaces under the **Automated Mesh > Custom Controls** node when using Parts Based Meshing. When using Region Based Meshing, the same controls can be set under individual regions or boundaries using the **Customize Prism Mesh** node that exists within the **Mesh Conditions** node for these objects.

Identical prism layer options are provided for part surfaces, regions, boundaries, and volumetric controls.

By default, prism layers are only generated on all wall boundaries or related part surfaces. They are not generated for any other type of boundary or related part surface unless you specifically activate this using custom controls. Activating a prism layer on a part surface or boundary exposes value nodes for the number of layers, total thickness, and layer distribution. For boundaries that lie on an interface, an additional control must be considered, which is the Interface Prism Layer Option that sets whether to grow prism layers from an interface

or not. Custom values can also be supplied based on volumetric controls.

Prism Layers (PBM) / Customize Prism Mesh (RBM)

Custom controls can be activated for particular part surfaces, regions, or boundaries.

The **Custom Controls** > **[Part Surface]** > **Controls** > **Prism Layers** node (PBM) or the **Mesh Conditions** > **Customize Prism Mesh** node (RBM) allow custom controls to be activated for particular part surfaces, regions, or boundaries. A single property then controls what happens on the surface regarding the prism layer.



If the *Specify Custom Values* option is selected, additional properties are activated in the **Prism Layers > Customize** (PBM) node or the **Mesh Values** node (RBM) which allow specific prism layer settings to be made.

Properties of the Customize Prism Mesh Node

Prism Layers (PBM) / Customize Prism Mesh (RBM)	Use this property to control the prism layer mesher on a part surface, region, or boundary.	
	Use Parent Values (PBM) / Use Default Values (RBM)	The default setting. Apply the global, higher-level settings, no local customization.
	Specify Custom Values	Choose to define custom controls for the part surface, region, or boundary. Appropriate properties or value nodes are added to the part surface, region, or boundary.
	Disable	Do not create prism layers on this part surface, region, or boundary.

Custom controls set at the local level supersede controls that are set at a global or higher level. Custom controls set on a region are applied to all boundaries within the region unless they have their own custom controls, which in turn supersede the region controls.

For multi-region interfaces, the interface prism layer option supersedes the region level setting for these boundaries and provides an additional level of control. For example, if the generation of prism layers has been deactivated for a region, prismatic layers can still be generated for boundary interfaces using this option using values that are specified at either the boundary or global level.

The region type specification, namely fluid, porous, or solid, has no influence on the behavior of the prism mesh settings.

Interface Prism Layer Option

The prism layer settings at the interface level allow you to decide whether prism layers are grown from the interface or not during the volume mesh generation process.

Certain applications such as those that include a porous or a rotating region do not require prism layers at the interface. Others, such as interfaces between fluid and solid regions do require a prism layer on one or both sides of the interface.

The option to activate them or not is provided via the **Interface Prism Layer Option** node, which is located within each interface node that exists. The default state for the *Grow Prisms from Interface* option is off. This option is shown below:



Custom prism values specified for a boundary belonging to an interface are only included when the Grow Prisms from Interface option has been activated.

Improving Prism Layer Construction in Gaps and Corners

This section introduces the controls that help you to improve prism layer construction.

For some geometries, it is not possible to generate a prism layer of acceptable quality on all surfaces, or at least not with the same thickness or number of layers. The prism layer mesher has several controls that determine whether prism layer thickness is reduced, or the number of layers are reduced when local conditions require it. The following controls are provided:

- Gap Fill Percentage
- Minimum Thickness
- Layer Reduction Percentage
- Boundary March Angle
- Concave and Convex Angle Limits
- Generating Standard Cells Only
- Improve Subsurface Quality

View Prism Layer Cells Using a Cell Set

STAR-CCM+ provides a field function that allows you to view cells that are generated by the Prism Layer Mesher. You can use the same function to view prism layers that the Advancing Layer Mesher model creates.

Use this function in combination with a cell set to define the selection of prism layer cells that you want to view in a displayer.



You can view prism layer cells as a cell set in cases which use the prism layer meshing model combined with the following volume meshing models:

- Polyhedral Mesher
- Generalized Cylinder Mesher (also with Polyhedral Mesher)
- Tetrahedral Mesher
- Trimmer Mesher
- Thin Mesher

At present it is not possible to use the prism layer cells field function if you use the following meshing models:

- Extruder Mesher
- Embedded Thin Mesher

Before Creating the Cell Set:

1. The mesh that you use for this feature must be created in STAR-CCM+ v8.06 or later. If necessary, regenerate the mesh in STAR-CCM+ v8.06 or later.
2. This feature does not support imported volume meshes. If necessary, regenerate any imported volume meshes using STAR-CCM+ v8.06 or later.

3. Ensure that your geometry parts are assigned to regions and that they are meshed using the Prism Layer Mesher model.

To create the cell set:

4. Right-click **Representations** > **Volume Mesh** > **Cell Sets**.
5. Select **New** > **Threshold**.



6. In the *Threshold* dialog that appears, set:
   1. *Scalar* to **Prism Layer Cells**.
   2. *Extraction Mode* to **All Above**.



The *Threshold* value is set to zero by default. However, when *Scalar* is set to **Prism Layer Cells**, there is no need to specify the *Threshold value* as it is automatically set to display all prism layer cells.

7. Set *Display* to **New Geometry Displayer**, then click **OK**.

Selecting **New Geometry Displayer** while a mesh scene is active automatically creates a cell set geometry displayer in the mesh scene.

A **Cell Set 1** node appears under the **Cell Sets** manager.



8. To determine the number of cells in the cell set, select the **Cell Sets** > **Cell Set** node.

   The number of cells in the cell set is displayed in the *Properties* window, to the right of the *Cells* property.



To view the cell set:

9. Select the **Scenes** > **Mesh Scene** node and click the *scene/plot* tab.
10. Expand the**Mesh Scene** > **Displayers**node and multi-select all nodes except the **Cell Set Geometry** node.
11. Right-click the selected nodes and select **Toggle Visibility**.
12. Return to the *simulation* tab.

The prism layer mesh cell set is displayed in the *Graphics* window.

To view a plane section of the cell set:

13. Create a plane section of the cell set.
14. Create a threshold part for the cell set, using the following settings:
    1. Set *Input Parts* to **plane section**
    2. Set *Scalar* to **Prism Layer Cells**
    3. Set *Extraction Mode* to **All Above**
    4. Set *Display* to **New Geometry Displayer**



You can use the resulting derived part in a scene displayer to visualize the cross-section of the prism layer.

Prism Layer Model Properties

This section describes the properties of the Prism Layer model.

Stretching Function	Selects the stretching function to use in calculating the thickness of each cell layer within the prism layer.	
	Geometric Progression	The thickness of each cell layer is calculated using a geometric progression that is based around a constant size ratio. Selecting this option, sets the **Default Controls > Prism Layer Stretching** property to 1.5 by default.
	Hyperbolic Tangent	The thickness of each cell layer is calculated using a hyperbolic tangent stretching function. Selecting this option, sets the **Default Controls > Prism Layer Stretching** property to 0.5 by default.
Distribution Mode/ Stretching Mode	Selects the distribution/stretching mode to use in specifying the relative size parameters that the stretching function requires.	
	Stretch Factor	The definition of the stretch factor depends on which stretching function is being used. Adds the **Prism Layer Stretching** node for

		global and custom settings.
	Wall Thickness	Sets the thickness of the near-wall cell layer. Thicknesses of the remaining layers are distributed according to a calculated stretch factor. Adds the Thickness of Near Wall Prism Layer node for global and custom settings.
	Thickness Ratio	Sets the ratio of the thickness of the cell layer furthest from the wall to the thickness of the cell layer nearest the wall. Adds the Prism Layer Thickness Ratio node for global and custom settings.
Gap Fill Percentage	When prism layers are generated in a narrow passage (or gap), it is possible that the prism layer could become larger than the gap. The gap fill percentage sets the maximum allowable proportion of a gap that the prism layer mesh can fill.	
Minimum Thickness Percentage	The minimum thickness percentage is used to calculate a threshold thickness below which the prism layer is retracted.	
Layer Reduction Percentage	Used to calculate a threshold thickness below which the number of cell layers is progressively reduced.	

Working with the Stretching Function

The stretching function determines the underlying formula that is used to generate the cell layer thickness distribution.

Two options are provided: geometric progression and hyperbolic tangent.

Geometric Progression

This default option distributes the cell layers using simple geometric relationships. The stretching factor between adjacent cell layers is always calculated as a constant value. An example of a prism layer that is constructed using 10 layers with geometric progression is given below. A **Prism Layer Stretching** factor of 1.5 (default) is used for the entire prism layer.



Hyperbolic Tangent

This option distributes the cell layers according to a one-parameter hyperbolic tangent stretching function:

$$s(n_i) = 1 + \frac{\tanh\left[F_s(n_i / N - 1)\right]}{\tanh(F_s)} \tag{5}$$

where N is the total number of layers, n_i is the current node, F_s is the overall stretching factor, and $s(n_i)$ is the distribution value for node n_i. Note that $0 \leq n_i \leq N$, so that for 5 layers, n_i takes the values 0, 1, 2, 3, 4, 5.

For more information, see Prism Layer Mesh Bibliography.

The relation between nodes and layers is depicted in the diagram below.



An overall **Prism Layer Stretching** factor of 1.5 for 10 layers gives a distribution as shown above. The size ratio between layers is no longer constant, but is modified to give a smoother transition between cell thicknesses across the whole prism layer. This option allows tighter spacing of near-wall prism layers while avoiding large spacing of the outer prism layers.

See Also

Advancing Layer Mesher Properties

Working with the Distribution/Stretching Mode

The distribution/stretching mode is a property of the **Prism Layer Mesher** node. It controls the method that is used to calculate the thickness distribution of multiple cell layers within the prism layer.

The mode cannot be set separately for local part surfaces, regions, and boundaries--it is only set once at a global level. However, the chosen option determines the type of node to appear in the **Default Controls** node of the **Automated Mesh** operation when using Parts Based Meshing. When using Region Based Meshing, it determines the type of node to appear in the **Reference Values** node of the mesh continuum.

These three options, and their corresponding value nodes, are as follows:

Distribution Mode Option	Corresponding Value Node
Stretch Factor	**Prism Layer Stretching**
Wall Thickness	**Thickness of Near Wall Prism Layer**
Thickness Ratio	**Prism Layer Thickness Ratio**

Stretch Factor

When a geometric progression is used for the stretching function, the *Stretch Factor* is the ratio of the thickness of a cell layer to the thickness of the preceding layer.

The minimum allowed value is 1.0, which results in uniformly spaced prism layers. Any value greater than 1.0 results in non-uniformly spaced prism layers with a low prism layer height nearer to the wall which increases progressively in each prism layer away from the wall.



For example, if the stretch factor is 2.0, and the thickness of the first prism layer next to the wall is 0.5 units, the thickness of the next layer would be 1 unit. Continuing to move away from the wall would give thicknesses of 2, 4, 8, and so on.

When the stretching function is set to hyperbolic tangent, the stretch factor becomes a parameter that applies within the stretching function. The actual size ratio between cell layers is not a constant, but is determined by the distribution function.

You can set *Prism Layer Stretching* values according to the following table:

Stretching Function	Valid Values
Geometric Progression	The default value is 1.5. A value between 1.2 through 1.5 is recommended, although any value >= 1.0 is valid.
Hyperbolic Tangent	The default value is 0.5 within default controls and 1.5 within custom surface controls. A value between 2.0 through 5.0 is recommended, although any value > 0.0 is valid.

In addition to the global control, local control is available.

When using Parts Based Meshing:

1. Set the **Custom Controls** > **[Surface]** > **Controls** > **Prism Layers** property to **Specify Custom Values**.
2. In the **Prism Layers** > **Customize** node, activate **Customize Distribution**.
3. Under the **Values** > **Custom Prism Values** node, a local value of the **Prism Layer Stretching** can be set on the chosen part surface.

When using Region Based Meshing:

4. Set the **[Boundary]** > **Mesh Conditions** > **Customize Prism Mesh** property to *Specify Custom Values*.
5. Under the **Mesh Values** node, a local value of the **Prism Layer Stretching** can be set on the chosen boundary.

See Also

Advancing Layer Mesher Properties

Wall Thickness

The *Wall Thickness* option sets the thickness of the layer that is immediately next to the wall boundary (called the near wall layer).

In order to calculate the appropriate thickness distribution, a stretch factor is automatically calculated and applied to the remaining prism layers (between the near wall layer and the start of the core mesh). The *Wall Thickness* value cannot be greater than the total prism layer thickness divided by the number of layers. If a larger value is supplied, then a linear distribution is used.



The initial wall thickness must always be greater than zero, regardless of the stretching function being used.

On selecting *Wall Thickness* for the distribution/stretching mode, the **Prism Layer Near Wall Thickness** node is added to the **Default Controls** node when using Parts Based Meshing. When using Region Based Meshing, a node **Thickness of Near Wall Prism Layer** is added under the **Reference Values** node of the mesh continuum. These nodes are used to set the global near wall layer thickness.

In addition to the global control, local control is available.

When using Parts Based Meshing:

1. Set the **Custom Controls > [Surface] > Controls > Prism Layers** property to **Specify Custom Values**.
2. In the **Prism Layers > Customize** node, activate **Customize Distribution**.
3. Under the **Values > Custom Prism Values** node, a local value of the **Prism Layer Near Wall Thickness** can be set on the chosen part surface.

When using Region Based Meshing:

4. Set the **[Boundary] > Mesh Conditions > Customize Prism Mesh** property to **Specify Custom Values**.
5. Under the **Mesh Values** node, a local value of the **Thickness of Near Wall Prism Layer** can be set on the chosen boundary.

Thickness Ratio

The *Thickness Ratio* is the ratio of the thickness of the prism layer next to the core mesh to the thickness of the prism layer next to the wall.

A stretch factor is automatically calculated for the intermediate layers. The minimum allowed value is 1.0.



The default *Prism Layer Thickness Ratio* value of *1.5* means that the inner most layer is 1.5 times the thickness of its outer most layer. Valid values of thickness ratio are given in the table below.

Stretching Function	Valid Values of Thickness Ratio
Geometric Progression	>= 1.0
Hyperbolic Tangent	> 1.0

In addition to the global control, local control is available.

When using Parts Based Meshing:

1. Set the **Custom Controls** > **[Surface]** > **Controls** > **Prism Layers** property to **Specify Custom Values**.
2. In the **Prism Layers** > **Customize** node, activate **Customize Distribution**.
3. Under the **Values** > **Custom Prism Values** node, a local value of the **Prism Layer Thickness Ratio** can be set on the chosen part surface.

When using Region Based Meshing:

4. Set the **[Boundary]** > **Mesh Conditions** > **Customize Prism Mesh** property to **Specify Custom Values**.
5. Under the **Mesh Values** node, a local value of the **Prism Layer Thickness Ratio** can be set on the chosen boundary.

What is the Gap Fill Percentage?

The Gap Fill Percentage property of the prism layer mesher controls the thickness of the subsurface (and hence the prism layer) where boundary surfaces are close to one another.

For example, two parallel planar boundaries that are close together. It essentially defines the amount of space in a gap that is occupied with a prism layer.

The thickness for each prism layer is calculated by multiplying the **Gap Fill Percentage** value (25% by default) by the intersection distance. This intersection distance is the normal distance between the two boundaries. If one of the boundaries does not have the prism layer option activated, then the maximum allowable prism layer thickness is 2 * Gap Fill Percentage * Intersection Distance.

The gap fill percentage value influences prism generation for surfaces that are perpendicular to one another, for example, a symmetry plane and a wall boundary. In this instance, the symmetry boundary is considered close to the wall leading to a reduction of the subsurface (and hence prism layer) thickness in the vicinity of the intersection of the boundaries. To minimize this effect, increase the gap fill percentage value.

What is the Minimum Thickness?

The Minimum Thickness is the smallest subsurface/prism layer thickness that is allowed below which the thickness is forced to zero.

It is specified as a percentage of the global prism layer thickness reference value.

In order to avoid self intersection or to meet other quality criteria, the subsurface sometimes has to be retracted for certain boundary surfaces (for example, because of proximity effects). If the retraction is severe enough to cause the thickness to fall below the *Minimum Thickness* value, the subsurface is locally set to zero. This approach avoids producing poor quality cells through squeezing several prism layers into a thin section.

What is Layer Reduction?

The Layer Reduction Percentage property controls the point at which the number of layers within a contracting prism layer is reduced.

When generating a prism layer mesh, it can become necessary for the mesher to reduce the overall prism layer thickness due to corners, narrowing gaps, curved surfaces, or other local geometry. In these cases, it is preferable to reduce the number of layers locally as the thickness decreases, maintaining a lower aspect ratio (and hence better quality) for the remaining prism layer cells. The Layer Reduction Percentage property controls the point at which the number of layers within a contracting prism layer is reduced.

By default, the *Layer Reduction Percentage* is specified as a percentage of the global prism layer thickness value. Specifying a value between 0% and 100% activates the feature and the layer reduction thickness is calculated as:

$$\text{Layer Reduction Thickness} = \frac{\text{Layer Reduction Percentage}}{100} \times \text{Prism Layer Thickness}$$

If a prism layer thickness has been specified on a part surface, region, or boundary, or through a volumetric control, then this local value is used instead of the global value.

Wherever the prism layer thickness falls below the Layer Reduction Thickness, prism layers are naturally thinned out as the prism layer thickness reduces. The default *Layer Reduction Percentage* value is 50%, which means that when the prism layer is reduced to half its target thickness, layers begin to be removed. Setting *Layer Reduction Percentage* to 0% deactivates the feature.

The range of influence of the layer reduction scheme on the prism layer mesh is illustrated below:



In the diagram that is shown above, the prism layers corresponding to the blue shaded areas would be subject to prism layer reduction based on the layer reduction thickness value as drawn. As this value increases, more of the prism layer is subject to the reduction process.

A comparison of a prism layer mesh without layer reduction and with layer reduction is shown below:



In the above example, no layer reduction (that is, layer reduction percentage = 0) has been applied. As the subsurface becomes thinner towards the angled corner, the prism layers progressively get thinner and eventually truncate where the subsurface retracts to zero height (determined by the minimum thickness value).

With a layer reduction value applied, the above situation can be avoided, as shown below:



In this case, the prism layers start truncating as you move towards the angled corner, reducing the number of layers that remain and increasing the overall quality.

Notes:

- The layer reduction feature has no effect where there is only one prism layer.
- If a layer reduction percentage value of 100% is used, then the entire prism layer mesh is subject to layer reduction. In general, values of 25 - 75% are reasonable to use for most cases.
- If the option is used with a volumetric control, then the prism layers that are associated with the control shape are also subject to layer reduction in the same way as described above.

Prism Layer Model Expert Properties

This section describes the expert properties of the Prism Layer model.

Boundary March Angle	The boundary march angle is used to control the behavior when the side of a prism layer lies on a boundary that does not have its own prism layer.

	The **Boundary March Angle** controls the way the subsurface scheme marches along a boundary that does not have a prism layer (also known as a no-subsurface boundary). The angle (in degrees) specifies the deviation from the normal that is allowed in order for the subsurface to be successfully generated along the no-subsurface boundary.  In the example that is shown, a subsurface is "marched" along the no-subsurface boundary (shown by the blue line) only if the *Boundary March Angle* is greater than the normal angle shown. If the normal angle is greater than the *Boundary March Angle*, then the subsurface is retracted, as shown by the magenta dotted line.
Concave Angle Limit	Sets a limiting angle that causes the prism layer to be retracted as it approaches the apex inside a wedge. The concave and convex angle limits result in the local retraction of the subsurface where the surface edge angle is less than the concave angle limit, or greater than the convex angle limit. The purpose of which is to improve the resultant prism layer mesh quality in these areas. For example, in the case of a car wheel meeting the ground, the concave angle limit is used to retract the subsurface where the wheel and the ground meet. Doing so avoids having non-orthogonal and possibly high aspect ratio prism cells at the edge intersection of the two surfaces. The concave angle is illustrated below:

	 In the above diagram, the black and blue boundaries represent surfaces from which subsurfaces are generated. The dotted magenta line represents the subsurface that would have been produced without the concave angle limit being activated. With the concave angle specified to a suitable value (that is, more than the geometry interior angle), the red line represents the retracted subsurface, which results in a higher-quality prism layer mesh.
Convex Angle Limit	Sets a limiting angle that causes the prism layer to be retracted as it approaches the apex outside a wedge. The convex angle can be used in situations such as at the trailing edge of a wing to retract the subsurface back to give an external angle is greater than the specified limit. An example of this situation is illustrated below:  In the above diagram, the blue boundary represents the boundary surface from which a subsurface is generated. The dotted magenta line represents the subsurface that would have been produced without the convex angle limit being activated. With the convex angle specified to a suitable angle (less than the

	exterior angle of the geometry), the red line represents the retracted subsurface, which results in a higher-quality prism layer mesh.
Near Core Layer Aspect Ratio	Sets an aspect ratio between the layer thickness and cell width of the prism layer closest to the core mesh, which improves the transition from the prism layer to the core mesh. The **Near Core Layer Aspect Ratio** controls the ratio between the cell height (layer thickness) and cell width of the prism cell layer closest to the core mesh. The cell faces on the core mesh (known as the subsurface) govern the prism layer cell width, and so the layer height is adjusted to set the aspect ratio. When the geometry surface is curved, the Surface Curvature control refines the mesh, which results in smaller cell widths across the curved portion. The near core layer aspect ratio becomes effective on these curved portions, as it prevents the mesher from generating long, thin cells with a high aspect ratio. This property only affects the thickness of the cell layer that is next to the core mesh. The thickness of the remaining layers is defined according to the stretching mode and function, and with any other active constraints. The near core layer aspect ratio is illustrated in the diagram below. The red area represents the prism layer closest to the core mesh (that is, the core mesh begins after the red line). The dashed magenta line shows the profile of the prism layer if no aspect ratio is applied. 

	As shown above, the aspect ratio is defined as the layer thickness divided by the cell width. For example, setting a value of 2 gives a layer thickness twice as large as the cell width. To deactivate the property, set a value of 0. Using an aspect ratio improves the transition from the prism cell layers to the core mesh, as the prism layer cells are a similar size to the core mesh cells. This option also prevents "tall" cells being generated, as the aspect ratio of prism layer cells is explicitly specified. Setting the near core layer aspect ratio influences the overall prism layer thickness, particularly on curved surfaces. However, note that the overall prism layer thickness always remains at or below the user-defined value. As an example, the following image shows a thick boundary layer without the near core layer aspect ratio specified. There are high aspect ratio cells in the prism layer due to the curvature of the geometry and associated increased surface mesh density. 
	If the near core layer aspect ratio is set to 1, the transition from the prism layer to the core volume mesh is greatly improved. The volume of the prism layer cells is then closer to the volume of the adjacent core mesh cells.

	 The overall prism layer thickness is reduced in this example. A more uniform thickness would be obtained by increasing the number cell layers, as shown below. You could also increase the specified aspect ratio. 
Generate Standard Cells Only	This option can be used with the tetrahedral meshing model to produce a volume mesh containing only pure tetrahedral and prismatic cells with triangular end faces. In order to produce a volume mesh containing pure tetrahedral and prismatic cells only, the generate standard cells only option can be used with the tetrahedral meshing model. Ensure that the number of prism layers that

	are used in the mesh is uniform throughout. That is, no custom number of prism layers have been set for part surfaces, regions, or boundaries. If these conditions are met, then the volume mesh that is produced can be exported to third-party packages that can only accept standard cell shapes. The default state for this option is off. Note: When using Parts Based Meshing, this option is only available to set when using a tetrahedral volume mesher.
Improve Subsurface Quality	This option causes retriangulation of boundary surfaces that do not have prism layers after the subsurface stage of the volume meshing process, and results in a better quality mesh. The improve subsurface quality option, which is on by default, improves the quality of boundaries that do not have prism layers (such as inlets, outlets, and interfaces). This option results in the retriangulation of such surfaces after the subsurface stage of the volume meshing process. With the option deactivated, such boundaries are smoothed as part of the subsurface process. This smoothing can lead to poor quality triangulation, especially when the prism layer thickness is large compared with the size of the boundary mesh triangles. Activating the option leads to better quality triangulation, which results in a higher-quality volume mesh being created. ***Note:*** This option is hardwired as on and cannot be turned off when using Parts Based Meshing. In this case, the option is not shown.

Prism Layer Mesh Bibliography

This section lists bibliographic references.

Vinokur, M. 1980. "On One-Dimensional Stretching Functions for Finite-Difference Calculations", NASA Contractor Report 3313.

Advancing Layer Mesher

The advancing layer mesher generates layers of prismatic cells around the surfaces of regions and fills the remaining void with polyhedral cells. This mesher first generates a polygonal surface mesh and then advances this mesh into the region volume to form cell layers. Some advantages of this approach are a conformal match on either side of an interface, and the ability to generate thicker and more uniform cell layers.

The extent and characteristics of the near-wall cell layers are user-defined. In this respect, the advancing layer mesher is a stand-alone volume mesher -- it is not the same as the prism layer model. For this reason, no other volume meshers are compulsory after selecting the advancing layer meshing model.

The prismatic cell layers help capture the boundary layer, turbulence effects, and heat transfer near wall boundaries. Prismatic cell layers can also be projected on non-extrusion boundaries such as inlets, outlets, and symmetry planes. The polyhedral core mesh provides a balanced solution for meshing complex geometries and is relatively quick to generate. This approach makes the advancing layer mesher suitable for various applications.

An example of an advancing layer mesh is shown below. For illustration, the prismatic layers are colored blue and the core mesh is colored pink.



The advancing layer mesher is represented in the STAR-CCM+ mesh continuum by the **Advancing Layer Mesher** node. The advancing layer mesher properties and expert properties are accessed by selecting the node.



See Also

Advancing Layer Mesher Properties
Prism Layer Mesher

Advancing Layer Mesher Properties

Stretching Function	Selects the stretching function to use in calculating the thickness of each cell layer within the prism layer.	
	Geometric Progression	The thickness of each cell layer is calculated using a geometric progression that is based around a constant size ratio.
	Hyperbolic Tangent	The thickness of each cell layer is calculated using a hyperbolic tangent stretching function.
Stretching Mode	Selects the stretching mode to use in specifying the relative size parameters that the stretching function requires.	
	Stretch Factor	The definition of the stretch factor depends on which stretching function is being used. Adds the **Prism Layer Stretching** node for global and custom settings.
	Wall Thickness	Sets the thickness of the near-wall cell layer. Thicknesses of the remaining layers are distributed according to a calculated stretch factor. Adds the **Thickness of Near Wall Prism Layer** node for global and custom settings.
	Thickness Ratio	Sets the ratio of the thickness of the cell layer furthest from the wall to the thickness of

		the cell layer nearest the wall. Adds the **Prism Layer Thickness Ratio** node for global and custom settings.
Run Optimizer	Performs additional optimization of vertices during the meshing process.	
	Activated	Activate the optimizer during meshing
	Deactivated	Deactivate the optimizer during meshing

Advancing Layer Mesher Expert Properties

The following table contains expert properties that are specific to the advancing layer mesher model. Common volume mesher expert properties are also available.

Enable Patcher	Patch the surface to smooth cell faces on the surface, while maintaining geometry features. However, this property can generate high aspect ratio cells along sharp edges.	
	Activated	Activates the patcher during meshing
	Deactivated	Deactivates the patcher during meshing
Surface Element Type	Specifies the shape of the faces on the surface.	
	Triangle	Creates the surface mesh using triangles.
	Quad	Creates the surface mesh using quadrilaterals.
	Polygon	Creates the surface mesh using polygons.
Generate Standard Cells Only	Lets you choose if only standard cells are produced or not. This property is only valid when *Surface Element Type* is set to **Triangle** or **Quad**.	

	Activated	The mesh that is generated can only contain cells that are tetrahedral, hexahedral, wedges, and pyramids.
	Deactivated	The mesh that is generated can contain polyhedral cells.

See Also

Advancing Layer Mesher
Working with the Stretching Function
Stretch Factor
Patching the Surface

Selecting the Advancing Layer Meshing Model

To select the advancing layer meshing model:

1. Right-click the **Continua** node and select **New > Mesh Continuum**.
2. Right-click the **Mesh** node and select **Select Meshing Models...**
3. In the *Mesh Model Selection* dialog, select **Advancing Layer Mesher** from the *Volume Mesh* group box.



How Is an Advancing Layer Mesh Generated?

The generation of an advancing layer mesh is outlined in the steps below:

1. Generate a surface mesh from the initial triangulated surface and specify the surface type element shape.
2. Optionally, patch the surface to maintain features while providing a smooth surface.
3. Extrude the surface mesh to create prismatic cell layers.
4. Generate a polyhedral core mesh to the remaining volume.
5. Assemble the core and prismatic cell layers.
6. If required, you can view the prism layer cells as a cell set using a field function.

Generating a Surface Mesh

In the first step, the advancing layer mesher generates a surface mesh from the triangulated surface (provided by the Surface Remesher or Surface Wrapper).

An example of a remeshed surface using triangles is shown below.



The advancing layer mesher takes this surface and generates the surface mesh using triangles, quadrilaterals, or polygons as specified in the surface element type property.

Note: If the quadrilateral or polygonal surfaces that are created contain quality issues–such as intersections and face proximity, the advancing layer mesher uses the original triangulated surface mesh on the affected patches. This behavior can cause some boundaries to remain triangular when the surface element type property is set as **quadrilaterals** or **polygons**. In this situation, the following message appears in the *Output* window: `Warning: initial triangular surface has quality issues. Converted quad/polygon surface will contain triangles.`

An example of a polygon surface mesh is shown below.



An example of a quadrilateral surface mesh is shown below.



The advancing layer mesher can create a conformal match between cells on either side of an interface. For this to succeed, layer settings on surfaces that meet at the interface perimeter must be identical. You can choose to deactivate layer growth on the interface itself, as shown in the following image. The conformal mesh applies to periodic as well as direct interfaces.



Note: To guarantee conformality ensure that the layer settings (thickness, number of layers, and stretching) are all the same.



Deactivating Polygonal Surface Meshing

You can choose to deactivate the polygonal surface generation, so that the existing surface mesh is used to create the prismatic cell layers. The surface could have a polygonal surface mesh from a previous advancing layer mesh, or a triangulated surface from the Surface Remesher, Surface Wrapper, or third-party software. Reasons for deactivating the polygonal surface mesher include mapping to third-party codes that cannot interpolate data from polygons, or maintaining an existing surface mesh. An example of prismatic cell layers that are generated from a triangular surface mesh is shown below.



To deactivate the polygonal surface mesh generation:

- Change the *Surface element type* property of the Advancing Layer Mesher node to *Triangle*.

See Also

Surface Remesher
Surface Wrapper

Patching the Surface

In order to generate more uniform cell faces while maintaining the features of the original geometry surface, the advancing layer mesher can patch the surface mesh.

This process of patching involves identifying features on the surface and creating a rectangular face from adjacent polygonal faces, as illustrated in the image below.



The effect on the volume mesh is to create triangular-prism-shaped cells in the corners of the region. An example of volume mesh cells that are generated from a patched surface are shown in the image below.



Another benefit of patching the surface is that a more uniform cell thickness is obtained when generating the prismatic layers. The disadvantage to patching is that there is a high volume ratio between the triangular corner cell and the adjacent prismatic layer cell.

The image below shows the resultant volume mesh from a patched surface. This image illustrates the high volume ratio between the triangular prism cell in the corner of a geometry and adjacent prismatic cell.



By comparison, the volume mesh that is generated without surface patching is shown below. This image shows the improvement in volume ratio for the cells in the corner.



To activate surface patching:

- Activate the *Enable Patcher* property of the **Advancing Layer Mesher** node.

See Also

Advancing Layer Mesher Properties

Generating the Prismatic Cell Layers

Prismatic cell layers are generated from the polygonal surface mesh. This approach gives an advantage over the prism layer mesher as the prismatic cell layers are less dependent on the original surface topology.

A more uniform thickness is therefore possible across all prismatic cell layers.



In some cases, the first layer next to the boundary may not have a uniform thickness. This scenario occurs when the advancing layer mesher has to collapse faces and edges in order to fit the required number of layers. In the example that is shown below, four prism layers is not sufficient for the mesher to maintain a uniform thickness in the corner of the volume mesh.



Increasing the number of prism layers to six, so that the first prism layer is thinner, results in a uniform first layer thickness.



Generating the Polyhedral Core Mesh

Once the prismatic cell layers have been generated, the advancing layer mesher proceeds to fill the remaining region volume with polyhedral cells.

The advancing layer mesher uses the last prismatic cell layer as a basis for the polyhedral core mesh. This approach ensures a one-to-one correlation between adjacent cell faces in the core mesh and cell layers.



Assembling the Volume Mesh

The final stage in generating an advancing layer mesh is to assemble the layered and core mesh.

The mesh is then improved and optimized using cell quality metrics.



What Are the Differences between the Advancing Layer and Prism Layer Meshers?

The advancing layer and prism layer meshers both produce prismatic cell layers near wall boundaries. However, there are several differences between the meshers that make them suitable for different applications.

The advancing layer mesher generates a polygonal surface mesh from the initial triangulated surface. The prismatic cell layers are then generated by extruding the polygonal surface into the region volume. This approach relaxes the dependency between the subsurface and region topology, and allows the advancing layer mesher to generate a thicker, more uniform layer than the prism layer mesher can achieve.

The advancing layer mesher also collapses and split faces and edges, which helps to improve mesh quality -- particularly when reducing cell skewness.

The example below illustrates the difference between a prism layer mesh (left) and advancing layer mesh (right).



Another difference between the prism layer and advancing layer meshers is the way that the two meshers behave in narrow gaps. In order to maintain some polyhedral cells between the opposing prism layers, the prism layer mesher can gradually reduce the layer thickness and remove layers. The advancing layer mesher instead removes cell layers without any reduction in thickness. An example where it could be advantageous to use the advancing layer mesher is when using a non-hybrid wall function turbulence model. The advancing layer mesher maintains the first layer thickness to give more control of the y+ value, whereas the prism layer mesher reduces the first layer thickness and the y+ is changed.

This difference between the two meshers is illustrated in the example below. The left-hand region has been meshed using the polyhedral and prism layer meshers, and the right-hand region has been meshed using the advancing layer mesher. The prism layer mesher maintains three prism layers in the gap by reducing the layer thickness. In contrast, the advancing layer mesher drops to one prism layer in the gap, but maintains the first layer thickness. For clarity, the prism layer is shown in blue.



Another current limitation of the advancing layer mesher is that volumetric sources can only be used to refine the core polyhedral mesh. That is, they have no effect on the prismatic cell layers. Therefore, you cannot modify the base size, number of layers, layer thickness, or stretching factor for the prismatic cell layers. However, you can reduce the cell size in the prismatic layers by applying a volumetric source to the Surface Remesher.

Extruder Mesher

The extruder mesher can be included with the core mesh when generating a volume mesh in order to produce orthogonal extruded cells for specific boundary surfaces. The extruder mesh is typically used for inlet and outlet boundaries to extend the volume mesh beyond the original dimensions of the starting surface, so that a more representative computational domain is obtained.

However, it can also be used for any boundary type and for combinations of adjacent boundaries, such as expanding the walls of a box in all outward directions.

The extruder mesher can affect some types of interfaces, such as periodic interfaces, if these lie on the extruded boundaries. To avoid any potential issues, it is recommended that you create these types of interfaces after the volume mesh has been generated.

Note: The extruder meshing model cannot be used on an imported volume mesh in order to extend the boundaries. In this instance, the original volume mesh surface can be extracted and a new mesh can be generated using the pipeline in STAR-CCM+, with the extruder model activated to extend the domain.

The extruder meshing model can be selected for any mesh continuum once a core volume mesh model has been selected.

See Also

Extracting a Surface Mesh

What is the Extruder Meshing Model?

The extruder meshing model, when activated and enabled for a specific boundary, performs an additional volume meshing step once the core mesh and, optionally, the prism layer mesh, has been generated.

Based on supplied values for the number of layers, extrusion magnitude and stretching, the volume mesh is extruded from the specified boundary either in the normal or specified direction, extending the original domain. An example extrusion for an inlet boundary in the normal direction is shown below:



If the prism layer mesher is included in the volume meshing pipeline, then prism layers are also produced for the adjacent boundaries by virtue of the extrusion, as per the example above. By default the extruded mesh belongs to the same region as the originating boundary, but can be made into a new region if desired. The boundary types that are used for the extruded volume are set to match the types of the boundaries that they have been extruded from.

The extruder meshing model can be activated for any of the core mesh types and enabled for any boundary. Care must be taken however that the extrusion volume does not interfere with the existing mesh by intersecting it in any way. Although such intersection does not invalidate the computational domain, it is unphysical and it presents issues in post processing the analysis results. The properties of the extruder mesher can be changed to alter the behavior of extrusions on neighboring boundaries.

The extruder mesher differs from the prism layer mesher in that the latter cannot increase the domain extents beyond the initial surfaces. The prism mesher is primarily used to ensure suitable near wall mesh resolution for turbulence and heat transfer modeling.

See Also

Volume Mesher Controls and Values

Extruder Mesher Properties

Frozen Boundaries	The *Frozen Boundaries* option, which is on by default, determines the behavior of the extruder mesher when extruding a group of boundaries that do not lie on the same plane. In the image below, neighboring boundaries are extruded in their respective normal directions with the *Frozen Boundaries* option activated. The result is a change to the topology of the computational domain. 
	If this behavior is undesirable, deactivate the *Frozen Boundaries* option to allow the automatic creation of additional mesh blocks to retain the original topology. In this example case, the result would be a mesh topology as shown below:

	
	The *Frozen Boundaries* option is turned on by default as it produces the fastest extrusion of all selected boundaries, since a check for new boundaries is not required after each extrusion. Turning off the option can significantly slow down the extrusion process in some cases.
Check Validity	The check validity option, which is on by default, checks the validity of the resultant extruded mesh and takes corrective action if problems are found. It is recommended that this option is left on when using the extruder model.

See Also

Extrusion Type Properties
Extrusion Parameters Properties

Extrusion Type Properties

When the Extruder Meshing Model has been selected, an additional **Extrusion Type** node is added to the **Mesh Conditions** manager node on each boundary.

The *Extrusion Type* property of this node is used to activate the extruder on that boundary.



Extrusion Type	**None**	The default value of the property indicating that the feature is deactivated.
	Constant Rate	The mesh is extruded in a specified direction

		with constant thickness layers.
	Constant Rate Normal	The mesh is extruded in the normal direction with constant thickness layers.
	Hyperbolic Tangent	The **Hyperbolic Tangent** option uses a hyperbolic stretching law to compute the cell distribution growth in the direction of the extrusion. This option is geared specifically for aerospace and other applications where the boundary layer resolution is considered critical.
	Symmetric Hyperbolic Tangent	The **Symmetric Hyperbolic Tangent** option grows the mesh hyperbolically from both ends.
	Outward Subsurface	The **Outward Subsurface** option provides an alternative scheme that is based on the subsurface generation method. In general, the other extrusion methods simply extrude the mesh in a normal direction to the local face by the specified distance, which can sometimes lead to problem cells in some areas. The outward subsurface method looks at the surrounding faces and

		extrudes in a direction that is less likely to cause problems. Additionally, it can partially or wholly retract the extrusion in areas where performing the full extrusion would lead to errors in the resultant mesh. This method is suited to complex meshes that require multiple boundary extrusions to generate a solid exterior or interior domain for the purposes of including heat transfer as part of the solution. The frozen boundaries extruder property must be turned on in order to use the outward subsurface extruder option.

See Also

Extruder Mesher Properties
Extrusion Parameters Properties

Extrusion Parameters Properties

All five extrusion types require that the extrusion parameters be set in order to obtain a successful extrusion.

Additionally, when selecting the **Constant rate**, **Hyperbolic tangent** or **Symmetric hyperbolic tangent** options, the direction of the extrusion in the supplied coordinate system must also be specified.

By default, the *Extrusion Type* option is set to **None**, indicating no boundary extrusion. Selecting any of the other options from the pull-down menu activates an additional property node in the **Mesh Values**

node which allows the extrusion parameters to be set for the boundary. This node is shown below:



Extrusion parameters can only be set at the boundary level for each boundary and therefore other local or global settings do not affect the extrusion. If the boundary belongs to an interface, then only remainder boundaries that are not participating in the interface are extruded. If there are no remainder boundaries, then the extruder settings have no effect on the meshing process for that boundary.

Extrusion Parameter Properties

Stretching	The meaning of the stretching value depends on the type of extrusion selected. When selecting a constant rate extrusion, the *Stretching* parameter controls the relative

	length of the extruded cells in the direction of extrusion. The parameter is the ratio of the size of the end most layer to the base or root layer. A value above 1 produces clustering towards the originating boundary while a value below 1 produces clustering away from it. When a hyperbolic tangent type of extrusion option has been chosen, the *Stretching* parameter is a non-dimensional clustering factor. If a value close to zero is used, then a near equal layer spacing is achieved while values above zero result in a clustering of the cell layers towards the originating root boundary. When using the symmetric option, the layers are also clustered towards the end boundary. Any value of *Stretching* greater than zero is allowed. The default value of 1.0 produces different results depending on the extrusion type selected.
Specify a New Region	The *Specify a New Region* option determines where the extruded mesh goes. The options appear in the *Properties* window for the **Extrusion Parameters** node in the form of a drop-down menu. When a new extruded region is created, it contains the following boundaries: • The base boundary (for any portion of the original boundary that does not fall within the internal interface between the original region and the extruded region. In most cases it will contain zero faces) • The top boundary (a copy of the extruded boundary, displaced according to the extrusion law and extrusion magnitude) • One or more side boundaries, each of them named **extension of {Boundary Name}** • An Internal Interface Boundary New regions are not automatically displayed as part of the model, so to view the result you must add any new region parts to the current displayer

	None	Creates the extrusion mesh as part of the same region as its source boundary. This option is the default selection.
	Yes, with Wall Boundaries	Sets the extrusion mesh as part of a new region that is called **Extrusion from {Boundary Name}**
	Yes, with Extended Type Boundaries	Sets the extrusion mesh as part of a new region which is an extension of the source region with the same extended boundary conditions. The new region that is created is called **Extrusion from {Boundary Name}.** This setting moves the original boundary type from the extruded boundary to the extrusion top. The original extruded boundary itself becomes a **Wall** boundary (as its extrusion base boundary counterpart). Note: The extruder can only generate a new region when the source boundary has a surface representation (that is, from the surface remesher or wrapper). Therefore, generate a surface mesh before

		specifying a new region.
Coordinate System	Specifies the coordinate system to use for the Constant rate, Hyperbolic tangent, or Symmetric hyperbolic tangent extrusion types. then the coordinate system in which the extrusion is performed must also be specified. Extrusions can be carried out in cartesian, cylindrical, or spherical type coordinate systems.	
Number of Layers	Determines the number of cells in the direction of extrusion. The default value is `1`, which is the minimum. There is no maximum limit to the number of layers allowed.	
Direction Magnitude	The total distance in each coordinate direction by which the mesh is extruded in the specified coordinate system. The direction magnitude value can be set for each direction using any of the allowable units and must be greater than 0. When the normal constant rate or outward subsurface extrusion type has been selected, the **Normal Extrusion Parameters** node is accessible instead: In this instance, specify the *Magnitude* in the normal direction using any of the allowable units.	
Use Average Normal	The *Use Average Normal* option of the **Normal Extrusion Parameters** node forces the extruder to compute an average face normal across the whole boundary from which the extrusion is grown. Using a consistent extrusion normal direction helps to maintain a uniform extrusion along a single direction, and is useful for cases where there are small imperfections on an otherwise planar surface.	
	Backward Compatibility If you use this feature in simulation files that were created using versions of STAR-CCM+ older than 7.04, the extruder creates a single side boundary as it did in earlier versions of	

	STAR-CCM+. To change this behavior, two techniques are available: • Extrude the boundaries with the *Create a New Region* option deactivated. • Delete the new region, make sure that the *Create a New Region* option is active, and repeat the extrusion action. This step creates a separate region with the side boundaries automatically split. **Working with the Frozen Boundaries Option** Normally you can deactivate the frozen boundaries option. However, if the *Create a New Region* option is active for two or more adjacent boundaries, the extruder works as if frozen boundaries were activated.

See Also

Displaying Remainder Boundaries
Extrusion Type Properties
Extruder Mesher Properties

Thin Mesher

The thin meshers are used to generate a prismatic type volume mesh for thin regions. However, the prismatic type mesh is only produced in portions of the geometry that the mesher recognizes as thin. Thick or bulk portions of the same geometry are meshed with polyhedral cells.

The different portions of mesh are joined continuously without abrupt changes in the mesh between the bulk and thin portions of the geometry.

There are two implementations within STAR-CCM+ for similar but differing applications:

- **Thin Mesher**
- **Embedded Thin Mesher**

Thin Mesher

Select the **Thin Mesher** for geometries that have no other region surrounding them and are predominantly thin or have thin structures included in them.

The **Thin Mesher** is used to generate a prismatic type mesh throughout the thin portions of the geometry instead of using other core volume meshers.

Embedded Thin Mesher

Select the **Embedded Thin Mesher** for geometries that are entirely contained within another region and are predominantly thin or have thin structures included in them.

Note: If necessary, to determine which parts of the geometry are considered thin, set the *Customize Thickness Threshold* expert property. Specifying a custom number of thin layers forces STAR-CCM+ to apply the **Embedded Thin Mesher**.

Using the Thin Meshers

The **Thin Mesher** is typically used for thin plate geometries. In these geometries, good quality cells are required to capture the solid material thickness adequately. Heat transfer analysis can then be performed between any geometries that are in contact with each other.

For example, the **Thin Mesher** can be used during the mesh generation stage of an analysis involving the heat transfer between the thin plates of a car chassis.

Using the same example, if the car chassis is contained in a region representing (say) air, which is also included as part of the analysis, use the **Embedded Thin Mesher**.

Determining if Geometry is Thin or Bulk

The thin mesher meshes as much of a region as possible using prism cells. For any areas where the thin mesher cannot produce prism cells, the thin mesher produces polyhedral cells. The thin mesher does not send any cells to the polyhedral mesher.

The embedded thin mesher automatically determines if a region is thin or bulk. The thin regions of the geometry are sent to the thin mesher and the bulk regions are sent to the polyhedral mesher.

However, you can manually specify an area as thin or bulk by customizing thin mesher parameters, or setting a specific value for the *Customize Thickness Threshold* expert property.

To specify a region as thin:

- Expand the **[region] > Mesh Conditions > Customize Thin Mesher Parameters**node.
- Set *Customize Thin Mesher Parameters* to **Specify Custom Values**.
- If the default number of layers is not sufficient, select the **[Region] > Mesh Values > Thin Mesher Layers** node and update the *Number of Layers*.

Setting this property forces the mesher to recognize the region as thin, regardless of setting any custom values that are set.

To specify a region as bulk:

- Expand the **[region] > Mesh Conditions > Customize Thin Mesher Parameters**node.
- Set *Customize Thin Mesher Parameters* to **Disable**.

Setting this property causes the mesher to recognize the region as bulk. A mesh is produced using the polyhedral volume mesher.

Note: If you deactivate the thin mesher, ensure that there is an alternative volume mesher available.

See Also

Thin Mesh Expert Properties
Customizing Thin Mesher Region Conditions and Values

Troubleshooting the Thin Meshers

The following section describes issues that can occur due to known behavior of the thin meshers.

Multiple Boundaries in a Thin Geometry

Prism cells are not created on boundary surfaces when there is more than one boundary in a thin portion of geometry. Instead, prism layers are created inside the volume which avoids confusion between prism layers that belong to different boundaries.



To view the internal prism layers, take a cross section of the geometry.



If you want to maintain prism layers on the two boundaries in the thin portion of the geometry, either:

- Activate per-part meshing, or
- Create a separate mesh continuum for the thin portion that has two boundaries. For this mesh continuum, select the prism layer mesher along with a volume meshing model that is not thin.



Using a separate mesh continuum creates non-conformal mesh at the interface.

Collapsed Prism Cells at Thin-Thin Interface

When the side of a thin mesh meets the long edge of a thin mesh at an interface, the prism layers in the thin mesh can collapse on one side of the interface. View the internal volume mesh using a cross section.



To overcome this issue and produce a better quality mesh, either:

- specify smaller surface size locally on the interface,
- or, select a non-thin volume mesher along with the prism layer mesher for one of the thin parts.



Using a separate mesh continuum creates non-conformal mesh at the interface.

What Is the Thin Meshing Model?

When activated, the thin meshing model allows thin regions in the geometry to have a prismatic type volume mesh. Using this mesh improves the overall cell quality and reduces the cell count when compared with an equivalent tetrahedral or polyhedral type core mesh.

Using this type of mesh can aid the convergence of simulations involving conjugate heat transfer in thin geometries. The thin mesher can be activated independently of the other volume meshers or as part of an "embedded" process within the polyhedral volume mesher. The number of layers in the thin mesh can be set. Additionally, an option is provided to set the threshold size for thin/non-thin structures. A choice between triangular and polyhedral type thin mesh is also allowed.

An example thin mesh with two triangular based prismatic layers, generated using the default thin mesher, is shown below:



A good quality triangulated surface mesh is required as the input for the thin mesher. Failure to provide such as surface results in either a poor

quality thin mesh being produced or no mesh at all. Where the geometry is considered not to be thin by the meshing model, the prismatic mesh in the thin portion transitions naturally to a normal core type mesh in the non-thin section. The threshold for this transition occurs if the distance between two opposing surfaces (the base surface and the target surface) is greater than four times the average edge length of the base surface.

The properties of the thin mesh can be changed to provide more meshing controls during the mesh process. The input values can be set on two different levels, that are listed below in lowest to highest priority:

- Global Level
- Region Level

Both types of thin mesher can be used with any interfaces that exist between two or more regions. The default thin mesher ensures a conformal mesh between thin geometries in each region that are in contact with one another. However, in the case of the embedded thin mesher, a conformal mesh is guaranteed only when the bulk region entirely encloses the thin mesh region. If a boundary plane in the thin region is exposed on the side of the bulk domain, then a conformal mesh is not always possible.

There is no guarantee that a conformal mesh is produced between the thin regions and the surrounding mesh region. However, a mesh is at least guaranteed without needing to activate the per-region meshing option.

The thin meshers cannot currently be used with volumetric controls. No method of specifying the volumetric controls for the default thin mesher is provided in this instance. Since the default **Thin Mesher** is used instead of one of the normal core volume meshers, it cannot be used with the Prism Layer Mesher, Extruder, and Generalized Cylinder models.

However, as the **Embedded Thin Mesher** works with the **Polyhedral Mesher**, you can use the **Prism Layer Mesher** to generate a prism layer within thin sections of the mesh.



Prism meshes cannot grow at thin-thin interfaces because conformality cannot be maintained in such cases. This condition is enforced no matter what other parameter is used.

However, for thin-thin side interfaces, prism layers from the thin mesher are projected onto the interface to produce a conformal mesh.



However, the mesher still collapses prism layers when:

- The number of prism layers on each side are different.
- There are more than two sides that meet at one interface.

See Also

Thin Mesh Expert Properties

Thin Mesh Expert Properties

Polyhedral Cells Type	The polyhedral cells type option determines the type of prismatic cells the thin mesher generates. Two options are provided: • Triangular prisms • Polygonal prisms These options can be seen in the pull-down menu for *Polyhedral Cells Type* in the *Properties* window: 
	By default, the *Polyhedral Cells Type* is set to **Polygonal prisms**. This option means that a polyhedral based mesh is used. In thin regions,

	multi-sided prismatic cells are created. In non-thin regions, a regular polyhedral type mesh is produced. Changing the option to **Triangular prisms** results in a triangular based mesh being used instead. In regions that are determined to be thin, the cells are regular three sided prisms. In non-thin regions, a regular tetrahedral type mesh is produced. An example thin mesh using polygonal prisms is shown below:
	
Run Optimizer	This option is off by default and is not available for the Embedded Thin Mesher. The *Run Optimizer* option allows more optimization of the thin mesh to be performed when the polygonal prism mesh type has been set. A better quality mesh is then produced in the thin mesh regions at the cost of longer mesh generation time. The *Run Optimizer* option has no effect for the triangular prism mesh type.

Automatic Correction	This option is on by default. The *Automatic Correction* option allows the thin mesher model to use cells other than prisms in thin regions when a locally invalid mesh is originally obtained. To ensure that a continuous, valid mesh is obtained, leave this option on where possible.
Customize Thickness Threshold	The *Customize Thickness Threshold* option allows a thickness threshold to be specified. The geometry is considered thin below this threshold. By default, the thin mesher determines which parts of the geometry are thin and which are not, so that the appropriate mesh type can be used. If this default setting is not optimal, it can be over-ridden by turning on the customize thickness threshold option and supplying an appropriate cut-off size. Alternatively, you can force thin meshing for a region using the region level mesh conditions.
Customize Surface Size Ratio Threshold	When this option is selected, the thin mesher determines which parts of the geometry are thin based on the product of the local surface cell size and the thin solid surface size ratio. Any geometry that is thinner than both this size and the thickness threshold is meshed with prismatic layers. The thin solid surface size ratio can be set at the global level, and also for individual regions.
Optimization Cycles	Increasing this value, increases the quality of the volume mesh, at the expense of increased CPU time in generating the mesh. The default is set to 1 and the recommended range is from 1 through 8.
Quality Threshold	Allows you to specify the maximum cell quality value that you want the mesher to aim for. Increasing this value increases the cell quality, at the expense of increased CPU time in generating the mesh. The default is set to 0.4 and the recommended range is from 0.4 through 1.0.

Thin Mesher Layers

The **Thin Mesher Layers** node allows the specification of the number of prismatic layers the thin mesher creates in geometries that are considered to be thin. Selecting this node gives access to the *Number of Layers* property:



Thin Mesher Layers Properties

Number of Layers	The default value is 1. Any integer value greater than 0 can be specified. For a geometry that contains both thick and thin sections, a value between 1 through 5 is suggested in order to prevent low quality cells or meshing failure. A minimum of two layers are always produced when using polygonal prisms. No such restriction exists when using triangular prisms.

Thin Mesher Solid Thickness

The **Thin Solid Thickness**node allows the specification of the threshold thickness, below which the geometry is considered to be thin. For this node to be accessible, activate the customize thickness threshold model property. Selecting this node gives access to the *Thin Solid Thickness* size properties:



Thin Mesher Solid Thickness Properties

Size Type	Depending on the **Size type**selection, a **Relative Size** or **Absolute Size** can be set to determine the value of the thickness. The default value of `0.0 m` indicates that the thin mesher model determines the thin/non-thin thickness threshold itself. Specifying a value greater than zero tells the thin mesher to search through the geometry and determine the thin sections (based on the supplied value). Surfaces that do not meet the proximity

	threshold are considered non-thin and are meshed with a non-prismatic mesh type.	
	Relative to base	Set the size relative to the base size.
	Absolute	Set the size as absolute value with length units.

Thin Mesher Solid Surface Size Ratio



Thin Mesher Solid Surface Size Ratio Properties

Thin Solid Surface Size Ratio	The solid surface size ratio is multiplied by the local surface size to provide the thickness below which geometry is considered thin. Therefore, the surface size settings on a face determine whether to mesh with prismatic layers or not. This parameter is added when the customize surface size ratio threshold

	option is selected on either the thin mesher or embedded thin mesher models. This ratio is set to 4.0 by default.

See Also

Determining if Geometry is Thin or Bulk
What Is the Thin Meshing Model?
Customizing Thin Mesher Region Conditions and Values

Customizing Thin Mesher Region Conditions and Values

At the region level, three different parameters can be specified to control the behavior of the thin mesher when the Customize Thin Mesher Parameters option has been activated.

The three parameters are:

- Thin Mesher Layers
- Thin Solid Thickness
- Thin Mesher Solid Surface Size Ratio

Customize Thin Mesher Parameters

To customize thin mesher parameters at the region level, select the **Region > Mesh Conditions > Customize Thin Mesher Parameters** node. In the *Properties* window, select a customization option from the drop-down menu.



Customize Thin Mesher Parameters Properties

Customize Thin Mesher Parameters	**Use Default Values**	Specifies that the region uses the thin mesher settings that are specified at the global level. This option also allows the mesher to decide whether a region is thin or bulk.
	Specify Custom Values	Specifies that the region uses the thin mesher settings that are specified at the region level. This option also forces the mesher to consider the region as thin.
	Disable	For the embedded thin mesher, the region does not use any thin

		mesher settings; hence, the mesher considers the region as bulk. For the thin mesher, the mesher still generates prismatic cell layers, as there is no alternative volume meshing model available.

See Also

Determining if Geometry is Thin or Bulk
Thin Mesh Expert Properties

Generalized Cylinder Mesher

The generalized cylinder mesher is used with the polyhedral volume mesher to generate an extruded mesh along the length of parts that are considered to be generalized cylinders.

Using extruded prismatic cells reduces the overall cell count and improves the rate of convergence in some cases. In order to apply the generalized cylinder mesher, candidate cylinders must be defined in separate boundaries. A tool is provided with the mesher that can automatically detect which boundaries are suitable for generalized cylinder meshing, and the properties of the generalized cylinder model are set at the boundary level. The generalized cylinder mesher can be used with the prism layer mesher.

The generalized cylinder mesher is best suited to cases where there are roughly cylindrical sections of geometry, and the direction of flow is parallel to the vessel wall, for example flow down pipes. Such cases can be solved more efficiently by using cells that are oriented to the direction of the fluid flow. An example of a generalized cylinder mesh is shown below; some of the internal cells have been exposed.



Requirements

The mesher can automatically detect and mesh boundaries that are suitable for generalized cylinder meshing. However, the generalized cylinder mesher does not automatically mesh any boundaries that it considers unsuitable. If the generalized cylinder mesher does not automatically identify and mesh any boundaries, the following message is displayed in the output window:

```
Candidate cylinder(s) : 0
```

You can manually specify a boundary *Type* as a **Generalized Cylinder Extrusion Type**. See Customizing Generalized Cylinder Conditions and Values.

If you manually specify an unsuitable boundary from which a generalized cylinder mesh cannot be produced, the following message is displayed:



Volumetric controls that encapsulate one or both ends of the cylindrical section can indirectly affect the generalized cylinder mesher. When a custom surface size is specified in the volumetric control, this size determines the number of vertices on the cylinder end loops used by the generalized cylinder mesher.

In cases where the generalized cylinder mesher cannot create a mesh using layers of hexahedral cells, it reverts to using a polyhedral mesh instead.

To specify a cylindrical boundary manually, activate the *Generalized Cylinder Extrusion Type* mesh condition for that boundary. See Customizing Generalized Cylinder Conditions and Values.

See Also

Polyhedral Mesher
Prism Layer Mesher
Volumetric Controls

Generalized Cylinder Mesher Best Practices

What Is a Generalized Cylinder?

Generalized cylinders are defined as a bounding wall with a closed loop at each end, as shown below.



Candidate cylinders have the following characteristics:

- The cylinder is straight or follows a smooth curve.



- The cross-sectional shape must have smooth edges—for example, a square with rounded corners, but not a square.



- The size and shape along the length of the cylinder varies, but transitions must be smooth.



- Cylinders can have a sharp change in cross-section, providing you specify each cylindrical surface as a separate boundary.



- Cylinders can have a hollow centre, as long as the internal and external surfaces of the cylinder are specified as cylindrical boundaries.



- The closed loops that define the start and end of the cylinder are slightly non-planar.



See Also

What Are Interfaces?
Creating a Boundary Interface

Preparing Boundaries for the Generalized Cylinder Mesher

Before using the generalized cylinder mesher, generate a good quality surface mesh and specify each cylinder as an individual boundary.

If you do not prepare the simulation in this manner, the generalized cylinder mesher can generate a poor quality volume mesh, or fail. See Generalized Cylinder Mesher Best Practices.

The geometry that is shown in the example below contains two cylindrical pipes that are good candidates for generalized cylinder meshing. However, in this case, the geometry is defined as a single boundary, so the generalized cylinder mesher would not recognize the pipe sections.



So that the generalized cylinder mesher can recognize the pipes, they must be separated into individual boundaries by using one or more of the splitting tools available in STAR-CCM+.

The pipe sections are now defined as separate boundaries, and the generalized cylinder detection tool can be used to assess whether they can be meshed as generalized cylinders.



Currently, the generalized cylinder mesher does not automatically recognize cylindrical elements whose boundaries are interfaces. To use interfaces with a generalized cylinder mesh, generate the volume mesh first and then create interfaces.

Interfaces that exist on the cylindrical boundaries of generalized cylinder mesher regions cannot produce a conformal mesh with regions that are meshed using other volume meshers.

However, the generalized cylinder mesher forms a conformal interface automatically with any region which fills the hollow portion of the cylinder.

See Also

Surface Meshers Overview
Splitting Boundaries
Detecting Cylinders Automatically

Generalized Cylinder Mesher Best Practices

If you encounter problems when identifying cylindrical boundaries or generating a volume mesh, follow these guidelines.

Improve the Surface Definition

The generalized cylinder mesher requires a good surface definition in order to identify cylindrical boundaries. To improve surface definition before using the generalized cylinder mesher, reduce the surface size. If the meshers are still unable to produce a volume mesh, the underlying geometry is unsuitable for cylinder meshing.

Split Long Cylindrical Geometries Into Smaller Sections

If a cylindrical geometry is long and/or contains curved sections, split the cylinder into smaller sections with separate boundaries. This process speeds up the meshing process.



Set the Same Surface Size on all Boundaries

To produce a good quality surface mesh, the vertex density on the inlet and outlet boundaries must match the vertex density on the feature curves. If the Custom Surface Size property is activated on the inlet or outlet boundaries and there is a difference in surface size from the inlet or outlet boundaries to the cylinder boundaries, poor quality faces are created. To resolve this issue, either deactivate the custom surface size property on all boundaries, or set the custom surface size setting to the same value for each boundary.



Where the Radius of the Cylinder is not Constant, Make Sure that the Cylinder Length is at Least Three Times the Change in Radius

If the cylinder length is not sufficient, the generalized cylinder mesher can fail. If it is not possible to adapt the geometry appropriately, deactivate the generalized cylinder mesher for thin cylinders. For example, in the following geometry, the length of distance A is not three times the length of distance B. To mesh this geometry successfully, either increase the length of A, or deactivate the generalized cylinder mesher in the thin cylinder.



Make Sure That Each Cylindrical Boundary is a Complete Cylinder

If more than one boundary represents a cylinder, combine the boundaries to make a complete cylinder.

See Also

Generalized Cylinder Mesher

Detecting Cylinders Automatically

The generalized cylinder meshing model detects automatically which boundaries in the simulation are suitable for generalized cylinder meshing.

To detect cylindrical boundaries, first ensure that the generalized cylinder model has been selected, then right-click the **Models > Generalized Cylinder** node in the mesh continuum and select **Manage Cylinders**.



Selecting this option causes STAR-CCM+ to search for any boundaries that are roughly cylindrical. When the search is complete, boundaries that meet the criteria for generalized cylinder meshing are highlighted in the active scene (providing a scene is open) and listed in the *Select Cylinders* dialog. The dialog is used to activate or deactivate the

generalized cylinder mesher on the boundaries that are found during the search, by selecting or deselecting the appropriate boundary.



By looking at the highlighted boundaries in the active scene, you can decide which to mesh as generalized cylinders, then activate or deactivate the corresponding boundaries in the *Find Cylinders* dialogand click **OK**.

The following screenshot shows an occasion where the generalized cylinder search found a boundary that is not intended for generalized cylinder meshing, Chamber. Since it would be desirable to generate a generalized cylinder mesh on the InflowPipe and OutflowPipe boundaries, but not on the Chamber boundary, the mesher can be deactivated for the Chamber boundary.



When changes to the *Find Cylinders* dialog are complete, clicking **OK** activates the generalized cylinder mesher for the selected boundaries. This action sets the generalized cylinder extrusion type for each boundary to **Constant**, and the extrusion parameters for the mesh then have the default values. To select an alternative extrusion type or customize the extrusion parameters, changes must be made for each boundary individually. See Customizing Generalized Cylinder Conditions and Values.

The **Manage Cylinders** option can also be used to remove existing generalized cylinder settings from specific boundaries by deactivating the checkboxes in the *Find Cylinders* dialog for the desired boundaries, and clicking **OK**. Alternatively generalized cylinder settings can be removed from all boundaries by selecting **Clear all Generalized Cylinder Options**.

Detecting Cylindrical Boundaries on Hollow Cylinders

When using the generalized cylinder mesher with hollow cylinders, all external and internal cylindrical boundaries must be selected for the generalized cylinder mesher.



If only one internal or external boundary is selected on a hollow cylinder for the generalized cylinder mesher, the cylinder is meshed using the polyhedral volume mesher.



See Also

Preparing Boundaries for the Generalized Cylinder Mesher

Clearing All Generalized Cylinder Options

You can remove the generalized cylinder settings from all boundaries in the model by right-clicking on the **Models > Generalized Cylinder** node in the mesh continuum and selecting **Clear all Generalized Cylinder Options**.

This option resets the generalized cylinder extrusion type mesh condition for all boundaries to **None**, deactivating the generalized cylinder mesher.



Alternatively the generalized cylinder mesher can be deactivated for individual boundaries by using the **Manage Cylinders** option, or by changing the properties of the generalized cylinder extrusion type mesh condition for each boundary.

Customizing Generalized Cylinder Conditions and Values

Properties for the generalized cylinder meshing model are set at the boundary level - no mesh continuum, region, or interface properties are available.

The generalized cylinder mesh is activated or deactivated through the generalized cylinder extrusion type mesh condition. Once activated, the mesh can be controlled by setting the generalized cylinder parameters.

Generalized Cylinder Extrusion Type

The generalized cylinder extrusion type determines if a generalized cylinder mesh is generated on the boundary, and if so offers a choice for the type of extruded mesh. After the generalized cylinder meshing model is selected, every boundary then has a generalized cylinder extrusion type node.

The *Extrusion Law* property sets the extrusion type, and the options available are:

- none
- constant
- symmetric hyperbolic tangent.

When the generalized cylinder detection tool detects a cylinder, the generalized cylinder mesher is activated for that boundary by setting the *Extrusion Law* to **Constant**.



Generalized Cylinder Extrusion Type Properties

None	The default setting for the extrusion law is None. When this option is selected a generalized cylinder mesh is not generated on the boundary, and a polyhedral mesh is generated instead.
Constant	This option results in a mesh containing cells that are a constant length. The properties of the mesh are controlled through the generalized cylinder parameters node, which is added to the boundary under the **Mesh Values** node.
Symmetric Hyperbolic Tangent	Selecting **Symmetric Hyperbolic Tangent** allows a mesh to be generated that has shorter cells at the start and end of the cylinder, and longer cells in the middle. The properties of the mesh are controlled through the generalized cylinder parameters node, which is added to the boundary under the **Mesh Values** node.

Generalized Cylinder Parameters

The generalized cylinder parameters determine the characteristics of the mesh that is generated with the generalized cylinder mesher. This node is activated if the *Extrusion Law* property of the generalized cylinder extrusion type node is set to either **Constant** or **Symmetric Hyperbolic Tangent**. Depending on the selections that have been made, the properties available are:

- number of layers
- stretching.

Generalized Cylinder Parameters Properties

Number of Layers	The Number of Layers property determines how many cell layers are generated along the length of the extrusion. Only integer values can be used to set the number of layers, and the default value is 20. Increasing the number of layers can improve the mesh that is generated on cylindrical

	elements that have a higher degree of curvature.
Stretching Factor	The *Stretching* property determines the relative size of each cell layer along the length of the cylinder when the **Symmetric Hyperbolic Tangent** option has been selected for the extrusion law. The property can be used to generate longer cells in the middle of the cylinder, and shorter cells at either end. Increasing the value causes increase clustering of the cells towards the end of the cylinder. Any value above 1.0 introduces clustering. The default value property is `1.0`, which generates cell layers of equal size. Values of less than 1.0. although valid, have no effect on the mesh produced.

Volume Mesher Controls and Values

This section details the common controls and values for all volume meshers. Controls that are specific to one mesher are detailed in the properties for that section.

Common Expert Properties

The following expert properties are available for all volume meshers (excluding optional volume meshing models) except where specified:

Choose Field Function Refinement Table (PBM) / Choose a Mesh Size Table (RBM)	Specifies the table for refinement of the mesh using field functions, which are based upon an existing solution. This option is not available for the thin mesher.
Optimization Cycles	Specifies the number of optimization cycles that the mesher runs on the tetrahedral mesh (in the case of a polyhedral mesh, on the initial tetrahedral mesh before generating polyhedral cells). To improve mesh quality, use this property with the *Quality Threshold*, but with increased meshing time. The recommended range is between 1 (lowest quality) and 8 (highest quality). This option is not available for the trimmed cell mesher (PBM)/ trimmer (RBM).
Quality Threshold	Specifies the lowest allowable mesh quality; where 1 is the best mesh that the mesher can generate. Use this property with the *Optimization Cycles* to improve mesh quality, but with increased meshing time. The recommended range is between 0.4 (lowest quality) and 1 (highest quality). This option is not available for the trimmed cell mesher (PBM)/ trimmer (RBM).

When using either the tetrahedral, polyhedral, or advancing layer meshing models, the following expert property is also available:

Enable Growth Rate (PBM) / Enable Mesh Expansion Control (RBM)	Controls the growth rate by setting a limiting factor for the mesh growth and the speed of the growth. When activated, you set the mesh growth rate on the **Maximum Cell Size**

	(PBM) / **Tet/Poly Expansion Max Cell Size** (RBM) and **Volume Growth Rate** (PBM) / **Tet/Poly Expansion Rate**(RBM)**xxx**nodes under the **Default Controls**(PBM) **/ Reference Values** (RBM) manager.

Volume Growth Rate (Tetrahedral or Polyhedral Mesher)

When using either the tetrahedral or polyhedral meshing models, a volume growth rate can be implemented to control the density of the core mesh, both globally and on a local basis. Below is a screenshot of the plane section through a cube on which the volume growth rate has been used to decrease the density of the core mesh cells, away from the boundaries.



The **Enable Growth Rate** (PBM) / *Enable Mesh Expansion Control* (RBM) option on the **Tetrahedral** or **Polyhedral Mesher** node provides global control over the expansion of the mesh. Activating this option adds the **Maximum Cell Size** (PBM) / **Tet/Poly Expansion Max Cell Size**(RBM)and**Volume Growth Rate**(PBM)**/ Tet/Poly Expansion Rate** (RBM) nodes.

Optimization Cycles and Quality Threshold Properties

The Tetrahedral and Polyhedral Meshers provide two expert properties, Optimization Cycles and Quality Threshold, that you can use to improve the volume mesh quality. However, the increase in quality results in longer processing times during mesh generation.

The Polyhedral Mesher generates a mesh by first creating an initial tetrahedral mesh, which it then converts to a polyhedral mesh using a process that is called dualization.

The optimization process is carried out on the tetrahedral mesh for both meshers. The expert properties can be used to increase the number of optimization cycles and improve the quality of the tetrahedral mesh. If the mesh is polyhedral, improvements to the underlying tetrahedral mesh have a direct beneficial effect on the final polyhedral mesh quality.

The Quality Threshold limits the Tetrahedron Mean Ratio, which is defined as:

$$\text{Tetrahedral Mean Ratio} = \frac{12 \times (3v)^{2/3}}{\sum_i l_i^2}$$

where:

- v — Cell Volume
- l — Edge Length

The Optimization Cycles are processes that the mesher performs to improve the quality and suitability of the volume mesh for the geometry. The polyhedral mesher runs these cycles on the underlying tetrahedral mesh, which has a direct effect on the subsequent polyhedral mesh quality.

As a general guide, applying the highest-quality levels can increase the meshing time by a factor of 10 and the cell count by 15%.

Note: Increase the *Optimization Cycles* and *Quality Threshold* properties together. Changing only one of these properties provides little improvement to mesh quality.

To improve the mesh quality:

- Select the **Operations** > **Automated Mesh** > **Meshers** > **Tetrahedral/Polyhedral Mesher** node for Parts Based Meshing or the **Continua** > **Mesh** > **Models** > **Tetrahedral/Polyhedral Mesher** node for Region Based Meshing.
- Set the *Optimization Cycles* and *Quality Threshold* to one of the following recommendations:

Quality Level	Optimization Cycles	Quality Threshold
Default	1	0.4
1	2	0.6
2	4	0.8
3	8	1.0

- Generate a volume mesh.

Example of Mesh Quality Improvement

This example shows four volume meshes where the quality of the cells between the circular pegs improves with each level of mesh quality improvement.

One Optimization Cycle and 0.4 Quality Threshold (Default)



Two Optimization Cycles and 0.6 Quality Threshold



Four Optimization Cycles and 0.8 Quality Threshold



Eight Optimization Cycles and 1.0 Quality Threshold



Volume Growth Rate (PBM) / Tet/Poly Expansion Rate (RBM)

The **Default Controls > Volume Growth Rate** node (PBM) or the **Reference Values > Tet/Poly Expansion Rate** node (RBM) controls the growth rate of the volume mesh away from boundaries.



Volume Growth Rate Properties

Volume Growth Rate (PBM) / Tet/Poly Expansion Rate (RBM)	This property defines the growth of the mesh size. This value should be set to between 1.0 and 2.0. Expansion rate of 1.0 is associated with a mesh that has uniform density throughout the assigned region. Increasing the growth factor results in larger mesh sizes in the transition from boundary size to the size away from the boundary. That is, increasing the value closer to 2.0 results in mesh sizes that grow quicker with each layer of cells away from the boundary.

Maximum Cell Size (PBM) / Tet/Poly Expansion Max Cell Size (RBM)

The **Maximum Cell Size** (PBM) / **Tet/Poly Expansion Max Cell Size** (RBM) node limits the largest cell size so that the mesh is not too coarse in the center of the domain.

For a polyhedral mesh, the value that you set here is applied to an initial tetrahedral mesh that the polyhedral mesher generates. The tetrahedral mesh is then converted to a polyhedral mesh during a process that is known as dualization. For this reason, the maximum cell size limit does not apply directly to the polyhedral mesh. It indirectly affects the cell size through the intermediate tetrahedral mesh.

You have the option of setting the size type to absolute or relative to the base size in the **Properties** window of the node.



The maximum cell size parameter also controls the maximum allowable cell size in the template refinement scheme when used with the trimmer model.

Maximum Cell Size Properties

Size Type	Selecting the **Maximum Cell Size** node allows the choice of a relative or absolute *Size type* to be made.	
	Relative to Base	Set the maximum cell size relative to the base size.
	Absolute	Set the maximum cell size as absolute value with length units.

Mesh Density

When using the tetrahedral or polyhedral meshing models, a volume growth and/or density factor can be used to increase or decrease the mesh density, both globally and on a local basis. Select the **Default Controls** > **Mesh Density** node (PBM) or the **Reference Values Tet/Poly Density** node (RBM) to access the properties.



Mesh Density Properties

Density	The *Density* value (1.0 by default) can be used to change the overall density of the mesh everywhere in the bulk volume. Increasing the value to say 2.0 would approximately double the number of cells that are generated while decreasing it to 0.5 would approximately halve the number of cells.
Growth Factor	The *Growth Factor* can be used to increase or decrease the mesh density of the core mesh by changing the rate at which cells grow from coarse to fine areas. In other words, it affects the grading of the mesh in a localized fashion where finer surface cells exist. Increasing the value above the default of 1.0 would mean that the cell size would grow faster, resulting in

	fewer cells. Conversely, decreasing the value below 1.0 would make the cell size grow more slowly, resulting in more cells being created.

To demonstrate the effect of each parameter, consider the example of two opposing surfaces that are separated by a gap. If the surface mesh sizes are the same, then the growth factor has no effect and only the density factor could be used to change the volume mesh density. If the surface mesh sizes differ from one surface to the other, then one or both factors could be used to influence the resulting density.

Neither the volume mesh growth nor density factor affects the cells directly next to the starting surface. Neither parameter has priority over the other.

Volumetric Control Blending (PBM) / Tet/Poly Volume Blending (RBM)

The global volume blending factor can be used to control the mesh density transition when volumetric controls overlap or are in close proximity to the surface mesh boundary or interface. This blending factor avoids sharp transitions in mesh density in these areas which could lead to numerical instability during the analysis.

Select the **Default Controls > Volumetric Control Blending** node (PBM) or the **Reference Values > Tet/Poly Volume Blending** node (RBM) to access the properties.

When using Parts Based Meshing, this node appears in the tree only when at least one volumetric control has been activated.



Volume Control Blending Properties

Blending Factor	The default *Blending Factor* of `1.0` provides a reasonable transition in mesh density between the cell size that is used for the volumetric control and the surface mesh size. Decreasing the value (to say 0.5) provides a smoother transition by making the local mesh size closer to the boundary/interface surface meshing size. Such a value results in fewer cells in the core mesh (assuming that the volumetric control uses a cell size smaller than the boundary/interface triangle edge size). Increasing the value (to say 1.5) provides a sharper transition by making the local mesh size closer to the volumetric control mesh size. Such a value results in more cells in the core

	mesh (again assuming that the volumetric control uses a cell size smaller than the boundary/interface triangle edge size). Note that the volume blending factor has little or no effect on the core mesh if the volumetric control is far away from the surface boundaries. It can also be used to control the mesh density transition between two or more neighboring or overlapping volumetric controls which use different cell sizes.

See Also

What is the Extruder Meshing Model?

Mesh Refinement

STAR-CCM+ can refine the mesh by approximating the wake region or using solution-based approach. These features allow you to achieve efficient refinement by reducing the cell size in areas of interest.

This section provides information on the following methods of mesh refinement:

Wake Refinement

It is beneficial to refine the mesh downstream of a body as boundary layer separation causes unsteady and turbulent fluid phenomena. In external aerodynamic simulation, this separated region is known as the wake.

The wake region is formed by sweeping a surface in a user-defined direction and distance. A refinement size and growth parameter determine the cell sizes in the refinement zone.

The main advantage in using the wake refinement option rather than a volumetric control is that you define the refinement zone with a surface. Therefore, if the boundary surface changes, the wake refinement changes also. This characteristic is useful in parametric or batch studies where the same settings are applied to several geometries.

Example of Trimmed Mesh Wake Refinement



Refining the Mesh in the Wake

You refine the mesh in the wake to capture the turbulent and transient flow conditions that occur behind a body.

Note: You can create multiple **Wake Refinement Sets**. This facility is useful for marine simulations where the wake broadens downstream of the hull.

To activate the wake refinement option for a region:

- Right-click the **Mesh Values** > **Trimmer Wake Refinement** node and select **New**.



- Select the **Wake Refinement Set** node and set its properties to control the characteristics of the wake refinement



Wake Refinement Set Properties

Relative/Absolute	Specifies whether you define the cell size in the refined region as relative or absolute. The actual size value is set on the child node.
Boundaries/Feature Curves	Specifies the boundaries and feature curves that you use to define the shape of the wake refinement region.
Distance	Specifies the length by which the wake refinement region extends beyond the selected boundaries and feature curves.
Direction	Specifies the direction in which the wake refinement region extends beyond the selected boundaries and feature curves. Enter as vector physical quantity relative to the trimmer coordinate system.
Default Growth Rate	Sets the default growth rate within the wake region (see Setting Wake Refinement Growth Rates).
Boundary Growth Rate	Sets the boundary growth rate within the wake region (see Setting Wake Refinement Growth Rates).
Customize Anisotropic Size	When activated, this property allows you to set a custom cell size for each coordinate direction relative to the trimmer coordinate system.

Defining the Shape of the Wake Refinement

The shape of the wake refinement region depends on the shape of the boundary surfaces and feature curves that you define. For the wake refinement to be valid, select at least one boundary.

This section covers the interactive selection of boundaries in the *Graphics* window—the procedure for selecting feature curves is the same. Alternatively, you can select boundaries and feature curves using the object selection dialog.

To define the shape of the wake refinement:

1. Ensure that you have a scene open in the *Graphics* window that includes the region for which you are applying wake refinement.

2. Right-click the **Wake Refinement Set** node and select **Select Boundaries**.



3. In the *Graphics* window, select the required boundaries curves.
4. Right-click one of the selected boundaries and select **Add Boundary** to add the selections to the *Boundary list*.



5. Once you have added a boundary to the list, it appears blue in the *Graphics* window.



6. Click **Apply** and then **Close** to define the *Boundaries/Feature Curves* property.



To remove boundaries from the *Boundary list*:

7. In the *Graphics* window, right-click an added boundary and select **Remove Boundary**.



8. Click to remove all boundaries from the list.

Setting the Wake Refinement Size

The wake refinement size determines the cell size within the refined zone. This refinement size can be either absolute or relative to the base size.

Note: If the wake refinement size is larger than the original cell size, there is no refinement.

To set the wake refinement size:

1. Select the **Wake Refinement Set** node and set the *Relative/Absolute* property to define either a relative or absolute size.
2. For a relative cell size:
   - Select the **Wake Refinement Set** > **Relative Size** node and set *Percentage of Base*.
     - The *Absolute Size* displays the cell size as a read-only property.



3. For an absolute cell size:

- Select the **Wake Refinement Set** > **Absolute Size** node and set the *Value* property.



Setting an Anisotropic Size

Anisotropic sizes allow you to define a different cell size in each coordinate direction.

Note: These coordinate directions are relative to the trimmer coordinate system.

To define anisotropic cell sizes:

- Select the **Wake Refinement Set** node and activate the *Customize anisotropic size* property.



- Select the **Trimmer Anisotropic Size** node and set the properties to define which coordinate directions are anisotropic.



- Select the child nodes and set the relative or absolute cell sizes.



Trimmer Anisotropic Size Properties

Custom X size	Activate this property to define a relative or absolute X size.
Custom Y size	Activate this property to define a relative or absolute Y size.
Custom Z size	Activate this property to define a relative or absolute Z size.

Setting Wake Refinement Growth Rates for the Trimmed Mesher

The *Default Growth Rate* and *Boundary Growth Rate* control the cell size transition in the wake refinement region.

These options have the same behavior as their global equivalents, the Trimmer Default Growth Rate and the Trimmer Boundary Growth Rate.

The following example illustrates and explains growth rates:



1. Depending on which is smaller, the surface triangulation, custom boundary size, or custom feature curve size determine the initial cell size.
2. The *Wake Region Boundary Growth Rate* maintains the initial cell size for a number of layers that are adjacent the boundary surface or feature curve. The number of layers depends on other meshing properties and the extent of the wake zone, as the mesher aims for a smooth transition in cell sizes. As an example, setting *Boundary Growth Rate* to **Medium** maintains the initial trimmed cell size for at least four layers.
3. The *Wake Region Default Growth Rate* determines the speed of transition from the initial cell size to the wake refinement cell size. Thereafter, the mesher maintains the wake refinement cell size to the limit of the wake refinement zone.
4. Beyond the wake refinement region, the default growth rate determines the speed of transition from the wake refinement cell size to the global cell size.

To set the growth rates:

- Select the **Wake Refinement Set** node and set the *Boundary Growth Rate* and *Default Growth Rate* properties.


Servers
Star 1
simulation
scene/plot
Regions
Car
Boundaries
Feature Curves
Mesh Conditions
Mesh Values
Trimmer Wake Refinement
Wake Refinement Set 1
Physics Conditions
Physics Values
Derived Parts
Solution Histories
Wake Refinement Set 1 - Properties
Properties
Relative/Absolute
Relative to base
Boundaries/Feature Curves
[Car: Body]
Distance
1.5 m
Direction
[1.0, 0.0, 0.0]
Default Growth Rate
Medium
Boundary Growth Rate
None
Customize anisotropic size

Field Function Mesh Refinement

To achieve successively more accurate solutions, you can refine the mesh with respect to a solution quantity (such as turbulent kinetic energy). In STAR-CCM+, you achieve successive refinement using mesh refinement tables, which are generated from field functions.

Currently, the following meshing models support field function mesh refinement:

- Surface Remesher
- Trimmed Mesher
- Polyhedral Mesher
- Tetrahedral Mesher

Growth rates determine the transition of cell and face sizes from refined to unrefined zones. The following table details the growth rate properties for each meshing model:

Mesher	Growth Rate
Surface Remesher	Surface Growth Rate
Trimmed	Template Growth Rate
Polyhedral	Tet/Poly Density Growth Factor
Tetrahedral	Tet/Poly Density Growth Factor

If you apply other mesh controls—such as custom mesh values and volumetric controls—to a particular portion of the mesh, the smallest size takes priority.

The following steps outline the procedure for field function mesh refinement:

1. Obtain a solution in the usual manner on the initial mesh.
2. Create a field function to return a mesh size value with respect to a solution quantity.
3. Use this field function to fill an XYZ internal table with mesh sizes. The table properties define the regions and boundaries for which you require mesh refinement.
4. Assign the table as a property of the meshing model.
5. Run the surface mesher or volume mesher to produce the refined mesh.

6. If you require further stages of mesh refinement from new solutions, repopulate the internal table using the extract option.

As this feature provides absolute mesh sizes, you can also coarsen the mesh when the field function returns a mesh size that is larger than the original.

Mesh Size

Meters are the only length unit that field function mesh refinement recognizes—preferred units do not apply. Therefore, ensure that your field function returns the mesh size in meters.

The effect of the mesh size on the refinement is different for the trimmed, polyhedral, and tetrahedral mesher, due to the nature in which these meshes are generated:

- For the trimmed mesher, the mesh size is equivalent to the base size in the refinement zone. If the field function returns a value of zero, the mesh is not refined in that zone.
- For the polyhedral mesher, you can approximate the mesh size as follows:

  $$\text{Mesh Size} \approx 1.2 \times \text{Cell Volume}^{(1/3)}$$

  Therefore, if the field function returns a value of `1.2*pow($Volume, 1/3)`, the polyhedral mesher maintains a similar mesh size in that zone.
- For the tetrahedral mesher, you can approximate the mesh size as follows:

  $$\text{Mesh Size} \approx 2 \times \text{Cell Volume}^{(1/3)}$$

  Therefore, if the field function returns a value of `2*pow($Volume, 1/3)`, the tetrahedral mesher maintains a similar mesh size in that zone.

Examples of Field Function Mesh Refinement

The following examples show mesh refinement using the trimmed, polyhedral, and tetrahedral mesher. The field function refines the mesh in an area where the turbulence kinetic energy is greater than 4.0 J/kg:

- Trimmed Mesher:

  ```
  ($TurbulentKineticEnergy>4)? 0.1 : 0
  ```

- Polyhedral Mesher:

```
($TurbulentKineticEnergy>4)? 0.1 : (1.2*pow($Volume,
1/3))
```

- Tetrahedral Mesher:

```
($TurbulentKineticEnergy>4)? 0.1 : (2*pow($Volume,
1/3))
```

Trimmed Mesher

Before mesh refinement:



After mesh refinement:



Polyhedral Mesher

Before mesh refinement:



After mesh refinement:



Tetrahedral Mesher

Before mesh refinement:



After mesh refinement:



Creating a Field Function for Mesh Refinement

The first step in field function mesh refinement is to create a field function that returns the mesh size with respect to an appropriate solution quantity.

In this example, if the turbulent kinetic energy is greater than 4 J/kg, the mesh size is 0.1 m; otherwise it is 5 m.

Note: It is important to consider your initial mesh resolution when creating a mesh refinement field function. If your initial mesh is coarse, define a field function that ensures you get a smooth variation in cell sizes. A refinement field function that defines an abrupt change in cell sizes results in poor mesh grading and quality.

To create a field function:

1. Right-click the **Tools > Field Function** manager and select **New**.



2. Select the **Field Function** node and click the ellipsis next to the *Definition* property.
   - In this example, the field function name is **Mesh Size**.



3. In the *Field Function - Definition* dialog, enter a conditional expression that returns the mesh size with respect to a solution characteristic.



4. Click **OK** to save the field function definition.

Note: The mesh size units are always meters; regardless of the preferred length unit.

Next, use this field function to create an internal table for mesh refinement.

Creating an XYZ Internal Table for Mesh Refinement

You use the mesh refinement field function to populate an XYZ internal table of mesh sizes, with respect to each cell or face centroid. Populating the table applies the field function to the existing volume mesh and its solution data.

Note: If you are using the Trimmer meshing model, ensure that the coordinate systems of the Trimmer and the XYZ internal table are the same.

To create an XYZ internal table:

1. Right-click the **Tools > Tables** node and select **New Table > XYZ Internal Table**.



2. Select the **XYZ Internal Table** node and click the ellipsis next to the *Scalars* property.

In this example, the table name is **Mesh Refinement Table**.



3. In the *Scalars* object selection dialog, move your previously defined mesh refinement field function to the *Selected* box and click **OK**.



4. Similarly, click the ellipsis next to the *Parts* property and expand the regions node.
5. Select the regions and boundaries that you want to refine, and click **OK**.



6. Right-click the **XYZ Internal Table** node and select Extract to compute the field function for each of the selected region cells or boundary faces.

7. To display the contents of the mesh refinement table, right-click the table node and select Tabulate.

As shown in the following image, the mesh size is given for the cell and face centroids of each selected region and boundary.

Tabular Data

	Mesh Size	X (m)	Y (m)	Z (m)
187	5	0.675329983234	-36.531034987747	0.4111494543
188	0.1	0.675329983234	-36.371993026883	0.254308285842
189	0.1	0.675329983234	-36.359524521476	-0.001391675148
190	0.1	0.675329983234	-36.450970056558	0.026693820725
191	0.1	0.675329983234	-36.602629971532	0.164914876204
192	0.1	0.675329983234	-36.642705568733	0.261697093148
193	0.1	0.675329983234	-36.256248740624	0.081044657548
194	0.1	0.675329983234	-36.20712458638	0.121548854232
195	5	0.675329983234	-36.076168080704	0.005244237562
196	0.1	0.675329983234	-36.254069077788	-0.011615332361
197	5	0.675329983234	-50.594304324404	0.580038228804
198	5	0.74018119816	-50.613782345044	0.570898678579
199	5	0.675329983234	-50.519325994633	0.811923251723
200	5	0.683407646946	-50.517212626608	0.809935350978
201	5	0.675329983234	-50.574336904329	0.724442821671
202	5	0.683407646946	-50.550739657193	0.770534730738

Export... Close

Next, assign this table to the appropriate meshing model.

Applying the Table of Mesh Sizes

The final stage in field function mesh refinement is to apply the table of mesh sizes to the meshing model properties.

To assign the mesh refinement table to the meshing model:

1. Select the **Mesh** node and set *Choose a mesh size table* to the name of the table that you previously created.



You can now regenerate the mesh in the usual manner for surface and volume meshers.

Directed Meshing

Directed meshing is a method for generating high quality structured meshes on CAD geometries in STAR-CCM+.

It operates by sweeping a surface mesh from the surface of a CAD geometry through its volume.



The directed mesh feature gives you a great degree of control when generating a volume mesh. You can create hexahedral meshes for highly irregular geometries by first breaking them up into parts that you can

sweep between two surfaces. You can also create meshes that are flow-aligned, which result in better accuracy within solvers.

You can use the directed meshing feature on geometry parts whose surfaces are divided into the following categories:

- Source surfaces where you generate the initial 2D surface mesh
- Target surfaces where the directed mesh ends
- Guide surfaces along which the mesher sweeps the initial 2D mesh from the source surface

  The guide surfaces in the region that you intend to mesh with directed meshing must have complete CAD association.

Note: The part must be assigned to a region before executing the directed meshing operation. However, you can create a directed meshing operation before assigning the part to a region.



On the source surface, you can generate a mesh using two methods:

- Creating patches on the surface and defining quad cells within the patches. See Creating a Source Mesh Using Patches.
- Extracting a mesh from an existing volume mesh on which the surface lies. See Creating a Source Mesh Using an Existing Mesh.

Note: You cannot use the two methods in combination for the same source surface of a part.

Each source surface mesh is associated with a volume mesh distribution that is used to generate the final volume mesh. You can use a mesh that is created using the directed mesh tool in any STAR-CCM+ analysis, along with meshes created using other methods.

Currently, the smoothing algorithms within the mesher aim to replicate the source mesh onto the target mesh whenever possible. When source or target surfaces are non-planar, this process can generate poor quality and negative volume cells. If you encounter this issue with your surface geometry, generate a mesh using alternative surface and volume meshers. CD-adapco expects to address this issue in future versions of STAR-CCM+.

Mesh faces that span multiple CAD faces (spanning faces) are not allowed on the guide surface for directed meshing. However, STAR-CCM+ can perform directed meshing when there are spanning faces on source and target surfaces.

Instead of producing an invalid directed mesh for some geometries when the barycentric mapping fails, STAR-CCM+ improves the quality of the mesh by switching to an advanced mapping technique. When this switch happens, the following message appears in the output window: `Unable to use barycentric mapping. Switching to alternate method...`

If this message appears, the meshing time is likely to increase because the advanced technique is more computationally expensive.

Creating Directed Meshes in STAR-CCM+

The subsequent sections provide information on setting up and generating directed meshes in STAR-CCM+.

Preparing CAD Parts for Directed Meshing

This section provides information on preparing CAD parts for directed meshing.

You can only use directed meshing with geometry parts that have CAD association. You can either import CAD geometry or create geometry using the 3D-CAD modeler.

The following requirements must be met regardless of the source of the geometry parts:

- Each part must have a regular section profile along its length.

  There cannot be any features intruding into the sweep path between the source and target surfaces. For example, shapes such as Y-junctions cannot be meshed without being split into two mirrored parts.

- Each part must have at least three surfaces, representing the source, target, and guide surfaces.

  You can subdivide any part into the necessary surfaces using the split by angle or split by patch options.

- A logically rectangular guide surface with five or more CAD edges is allowed.

  Include CAD vertices in the distribution path even if the distribution does not land on those points. Directed meshing smooths over any CAD vertices that are not designated as a feature.

Note: In cases where there is a dramatic bending of the sweep path for a guide surface containing five or more CAD edges, create four-sided guide surfaces.

- Imprint contacting parts on each other and ensure that part contacts are correctly defined on each part.

  Use the CAD-based imprinter under Parts, or ensure that imprinting is active for features in the 3D-CAD modeler. Do not use the surface repair imprinter or the discrete imprinter under Parts as doing so disassociates the CAD from the surface.

- If a part is associated with a region that is meshed using a mesh continuum, then that part is not available for use within the directed meshing selection dialog.

Delete the region or disassociate the part with the region in order for it to be available.

Note: Changing a 3D-CAD model whose parts are used in a directed mesh operation can result in the directed mesh operation becoming invalid. In this case, apply an update and review the settings in the directed mesh operation.

Launching a Directed Mesh Operation

The Directed Mesh tool performs all the operations that are associated with directed meshing.

You can access directed meshing tools in the simulation tree or using the edit menu.

To create a Directed Mesh operation:

1. Right-click the **Geometry** > **Operations** node and select **New** > **Directed Mesh**.



2. In the *Create Directed Mesh Operation* dialog, select the parts that you want to use for the directed meshing operation.

   Doing so creates an **Operations** > **Directed Mesh** node.



Parts that are selected for directed meshing are listed under the **Directed Mesh > Connected Parts** node.

If you select multiple parts that form a contiguous set of bodies, STAR-CCM+ represents these parts by a single Connected Parts object under the **Directed Mesh > Connected Parts** node.



Note: The yellow warning triangles indicate that further action is needed before the operation is fully defined. They are not displayed once you generate the volume mesh.

3. You can edit the directed mesh in two ways:

   - Selecting a directed mesh child node in the simulation tree and editing the properties in the *Properties* window:



   - Right-clicking the **Directed Mesh** node and selecting **Edit**, then editing the child nodes in the *Directed Mesh* panel.


Operations
Directed Mesh
Source
Target S
Connect
Source
Mesh Distributions
Edit...
Execute
Delete
Rename...

Choosing Source and Target Surfaces

Before creating a surface mesh for directed meshing, identify the source and target surfaces for each part.

Connected parts must share a source surface and target surface. The same volume mesh distribution parameters are applied to all of the connected parts. If the connected parts are of different lengths, then STAR-CCM+ makes an estimate with regards to the distribution in the sweeping dimension.

Note: Naming the faces on either side of the interfaces in a stack of parts makes it easier to identify the source and target surfaces for directed meshing. See Meshing a Stack of Parts for additional information.

Selecting Source and Target Surfaces Using the Tree

This section describes how to select source and target surfaces.

To select a source surface:

1. Select the **Directed Mesh** > **Source Surfaces** node and in the *Properties* window, click to the right of *Source Surfaces*.



2. In the *Source Surfaces - Source Surfaces* dialog, select at least one surface to represent the source surface.

To select a target surface:

3. Select the **Directed Mesh** > **Target Surfaces** node and in the *Properties* window, click to the right of *Target Surfaces*.



4. In the *Target Surfaces - Target Surfaces* dialog, select at least one surface to represent the target surface.

To remove any surfaces from source or target surfaces:

5. Select the **Source Surfaces/Target Surfaces** node.
6. In the *Properties* window, click ...(custom editor) for *Source Surfaces/ Target Surfaces*.
7. In the object selection dialog, deselect any surfaces that are not required as a source or target surface.

Selecting Source and Target Surfaces Interactively

In the *Graphics* window, you can select surfaces on the geometry and add them to the source or target surfaces for the mesh.

To select the source and target surfaces interactively in the *Graphics* window:

1. Edit the **Operations** > **Directed Mesh** node.
2. Select the surfaces that represent the source or target surface:
   - To select a single surface in the *Graphics* window, click the surface that you require.
   - To select multiple surfaces in the *Graphics* window, there are several methods:

- Use `<Ctrl>` + click
- Double-click

  When you double-click a part surface, all contiguous part surfaces that meet within an angle of 30 degrees are selected.

- Use the rubberband select tool

  Activate the rubberband tool and select the required surfaces.

Note: If you are meshing a part collection with the directed mesher, STAR-CCM+ can automatically account for intermediate source and target surfaces that exist between the first source surfaces and the final target surfaces. This feature is useful when you mesh a stack of parts as you do not have to specify each contacting surface as either source or target throughout the stack. In this situation, only define the first source surfaces and final target surfaces at this stage.

3. In the *Graphics* window, right-click the selected surfaces. For a part collection, select either the first source surface or the final target surface.
4. In the pop-up menu, select the appropriate option.

   For surfaces on individual parts:

   - **Part Surfaces > Add to Source**
   - **Part Surfaces > Add to Target**

   For part collections, only define the first surfaces as source, and the final surfaces as target.



5. If you are meshing a part collection, right-click either the first source surfaces or the final target surfaces, and select **Part Surfaces > Autopopulate Intermediate Source/Target Surfaces.**

   Intermediate source and target surfaces are defined on contacting surfaces that exist between the first source surfaces and the final target surfaces.



To remove any surfaces from source, target, or intermediate surfaces:

6. Right-click the source or target surface that you want to remove. For intermediate source and target surfaces, right-click either a source or target surface.

7. Select the appropriate option:
   - **Part Surface (source) > Remove From Source**
   - **Part Surface (target) > Remove From Target**
   - **Part Surfaces > Remove Intermediate Source/Target Surfaces**

Setting Source and Target Parts as Periodic Interfaces

It is possible to use boundaries that are specified as periodic interfaces for source and target surfaces in directed meshing.

You can also specify surfaces that are already defined as source and target parts as periodic interfaces:

- Create the periodic contact between the Source and Target surfaces.
- Assign Parts to Regions.
- Verify that the Rotational and Translation properties are correct.
- Proceed to Meshing:
  - Creating a Source Mesh Using Patches
  - Creating a Source Mesh Using an Existing Mesh

Meshing a Stack of Parts

This section describes how to mesh a stack of parts.

When meshing a stack of parts, prepare the parts with a consistent naming convention so that you can identify each side of the interface throughout the stack.



When setting the Source Surfaces for the Directed Mesh operation, choose the faces on the upstream side of the initial part in the stack.



When setting the Target Surfaces for the Directed Mesh operation, choose the faces on the downstream side of the final part in the stack.



1. To identify source and target interfacing surfaces in the stack, use one of the following options:
   - Selecting Source And Target Surfaces Interactively

     Use this option to autopopulate intermediate source and target surfaces.
   - Selecting Source And Target Surfaces Using The Tree
2. After identifying the source and target surfaces, define a mesh by either:
   - Patch Meshing
   - A previously defined volume mesh

Meshing Multiple Parts in a Closed Loop

STAR-CCM+ can mesh two or more parts that form a closed loop, using a single Directed Mesh Operation.

You specify the source and target surfaces in the same way as for a stack of parts.

The examples below show the results of meshing multiple parts in a closed loop, using the directed mesher.



To produce a directed mesh in a multi-part closed loop:

1. Set the **Source Surfaces** and **Target Surfaces** for the directed mesh.

   At each boundary of coincident surfaces, ensure that there is one **Source Surface** and one **Target Surface**.





2. Define the source mesh by:
   - Using a **Patch Mesh**

     See Creating a Source Mesh Using Patches.
   - **Using Existing Mesh**

     See Creating a Source Mesh Using an Existing Mesh

   The directed mesh is displayed in the *Graphics* window:



To refine the mesh quality across each part:

3. Select the **Mesh Distributions > Volume distribution** node.
4. In the *Properties* window, update the *Number of Repetitions* to match the number of parts that compose the closed loop.
5. To increases the refinement, increase the **Volume distribution > Number of Layers** in the loop.

   The following model is set to `16` layers in `4` repetitions:


STAR-CCM+

Meshing a Single Part in a Closed Loop with a Baffle

This section describes how to mesh a single part in a closed loop with a baffle.

STAR-CCM+ can perform directed meshing on a single part in a closed loop using one baffle surface as the designated source and target part.





To mesh a single part in a closed loop with a baffle:

1. In 3D-CAD, imprint the solid body and sheet body that represent your closed loop part and baffle respectively. See Imprinting Bodies.

Note: Parts must be imprinted in 3D-CAD to make sure that the baffle part is recognized as two-sided when you set the source and target surfaces.

2. In STAR-CCM+, create the geometry parts.
3. Combine the solid body and sheet body parts.
4. Expand the **[Body]** > **Surfaces** node.
5. Rename the surface that represents the baffle.
6. Right-click the **[Baffle]** surface and select **Create Baffle**.
7. Launch a directed mesh operation.
8. Set Source and Target surfaces.
   - Set Source Surface to **[Baffle]**.
   - Set Target Surface to **[Baffle] 2**.
9. Define the source mesh by creating a source mesh using an existing mesh or creating a source mesh using patches.
10. Create a Volume distribution.
11. Generate the Directed Mesh Operation.

Requirements

- The body that is meshed can only contain one baffle.
- There can be no baffles within the path of the sweep between the source and target.
- The baffle surfaces must be designated as source and target surfaces. Partial baffles are not allowed.
- The baffle surface should not have any free edge.

What Is Patch Meshing?

One of the two methods for generating a 2D mesh on the source surface is to use a set of patches on the source surface.

A patch is a four-sided shape that is made up of four patch vertices. The patch curves that define a patch can be straight or curved, and can be derived from CAD feature curves on the original geometry parts.

You can create as many patches as you like on the source surfaces providing you meet the following conditions:

- There is at least one patch
- The patches cover the entire surface

Currently, it is not possible to specify the exact location of patch vertices that connect the patch curves using coordinate values.

For patch meshing, all patch curves must be part of a closed loop. Patch curves that are not on feature edges must also be shared between patches.

Patch curves are colored according to their status:

- Red indicates that the patch curve is not part of a closed loop; add additional curves to create a complete patch
- Yellow indicates that the curve is part of a closed loop but it is only associated with a single patch
- Green indicates that the curve is part of a valid patch

You can only create a patch mesh when all the patch curves are green.



Patch vertices (green points) connect patch curves. A patch curve is created from a patch vertex that has the “focus”, shown using red lines passing through the vertex point. For curved patch curves, blue construction points are used to guide the curve, but do not act as patch vertices.



Patch meshing uses a patch *edit* panel. The panel consists of two modes:

- Creating a patch topology
- Defining the quadrilateral mesh that STAR-CCM+ generates within patches.

Creating a Source Mesh Using Patches

This section describes how to create a source mesh using patches.

Create a patch mesh:

1. To enter the directed mesh edit panel, right-click **Operations** > **Directed Mesh** and select **Edit**.



2. In the Directed Mesh panel, right-click the **Source Meshes** node and select **New Source Mesh** > **Patch Mesh**.

Note: Make sure that source and target surfaces are already specified before creating the source mesh.

3. In the *Part Collection For This Source Mesh* dialog, select the connected parts on which you are going to create the patch mesh and click **OK**.

   Doing so opens an *Edit* panel that is used to create patch curves:



The *Graphics* window updates to show the selected parts. The source surface on which you create the patches appears gray and the remaining surfaces are translucent for reference. Feature curves appear yellow.



Creating the Patch Topology

You use the tools in the Patch Topology mode of the patch mesh *Edit* panel to create patches on the source surface.

By default, the patch mesh *Edit* panel opens in this mode.



The available buttons in the *Create/Edit Patch Entities* box are used to create, delete, and split patch curves. The buttons in this panel are "sticky" buttons. That is, once pressed, they stay activated until another button is pressed or the `<Esc>` key is pressed to cancel the action.

You can undo or redo any of these actions with the undo and redo buttons at the bottom of the panel.

Creating a Patch

This section describes how to create patch curves.

There are three ways to create patch curves using the patch entities box:

- Auto-populate feature edges with patch curves
- Create patch curve using straight edges
- Create patch curve using construction points

These three methods can be used alone or in combination.

Extracting Patch Curves from Feature Edges

To create a patch curve on the feature edges:

1. Click the (**Auto-populate feature edges with patch curves**) button to create patch curves on the feature edges of a source surface. It also places patch curve vertices at the locations of feature vertices.

   If the feature edge is an ellipse, STAR-CCM+ places two vertices on opposite sides of the curve, resulting in two patch curves.



Creating Patch Curves with Straight Edges

You can create a patch curve as a straight edge between two vertices.

To create a patch curve out of two vertices:

1. Click the (**Create patch curve**) button and move the mouse anywhere over the source surface.
2. Click the location of the surface where you want to put the first patch vertex point.

   A green dot represents the vertex point. Click again where you wish to put the second patch vertex. A red curve is drawn between the two points. The most recently placed vertex has the focus.



3. To close the loop, create three more patch curves.



Creating Patch Curves with Constructions Points

You can create a patch curve that is not straight.

To create a patch curve that is not straight:

1. Click the (**Create patch curve with construction points**) button and move the mouse anywhere over the source surface.
2. Click the location of the surface where you want to put the first patch vertex point.

   This vertex point is a normal patch vertex. Continue adding points along a curve you want to make. These subsequent points are in blue to represent the curve construction points.
3. To end the curve, use one of the following options:
   - Click a patch vertex (green point) from another curve.
   - Right-click the last construction point that you made for the curve. This point is then converted to a regular patch vertex.
   - Press `<Esc>` to convert the last construction point into a regular patch vertex.

   The following patches are created using the patch curve tools in combination with the split curve and move patch vertex features.



Changing the Patch Curve Focus

At some point after creating one or more patches, the focus is no longer in a usable location.

To change the focus of the patch curve:

1. Right-click the patch vertex where you want the new focus point to be.
2. To clear the focus, click anywhere outside of the source surface, and then click again on the patch vertex to set the new focus location.
3. Press `<Esc>` to cancel the patch creation process and then press the **Create Patch Curve** button again to start a new patch in any desired location.

Specifying Exact Patch Vertex Location

You can move patch vertices by defining the exact coordinates of the new location.

To specify an exact patch vertex location:

1. Press <Esc> to cancel any current patch operations. Doing so removes the focus.
2. Click any existing patch vertex.

   The *Patch Vertex Location Options* appears in the patch mesh *Edit* panel.



3. The Snap To panel identifies the location of the patch vertex:
   - Feature Vertex (grayed out if none exist in the source mesh)
   - Feature Edge
   - Surface
4. The Coordinates panel displays the current x, y, z coordinate location of the patch vertex. Enter the specific coordinates for the patch vertex.
5. Select the Coordinate System.
6. Click **Apply**.

   The patch vertex moves to the specified location.

Moving Patch Vertices

You can move patch vertices or curve construction points to new locations.

To move a patch vertex:

1. Press `<Esc>` to cancel any current patch operations. Doing so removes the focus.
2. Click and drag the patch vertex that you want to move to its new location.



Note: Moving a vertex to the same location as another vertex could invalidate the patch topology. STAR-CCM+ removes any zero-length patch curve.

Splitting and Deleting Patch Curves

If you want to create a patch curve that crosses another patch curve, split the curve first.

To split a curve:

1. Click the (**Split the selected patch curve**) button. The *Split Patch Curve Options* appears in the patch mesh *Edit* panel.



2. To split a curve at a picked location, click *At Picked Location* and then click the curve location where you want the split to happen.
3. To split a curve at a specified percentage, click *At Specified Percentage*. Enter the percentage and click **Apply**.



4. To split a curve into an equal number of splits, click *Equi-spaced Multiple Splits*. Enter the Number of Splits and press `<Enter>`.



Note: Splitting a curve into two usually results in a topologically incorrect set of patches. In the following example, splitting a shared curve results in five sided patches.



To delete a patch curve:

5. Click the (**Delete the selected patch curve**) button and click any curve that you want to delete.

Deleting a patch curve with construction points deletes the entire curve and not just the clicked segment. Removing a patch curve could cause the patch definition to become topologically incomplete.

Undo and Redo in Patch Topology Mode

When creating or editing the patch topology, you can undo and redo actions using the undo and redo buttons at the bottom of the panel.

The undo tool reverses the previous action

The redo tool repeats a previously undone action

Hovering the cursor over the button indicates the action that the button reverses or repeats.

Note: If you change between modes (from **Patch Topology** to **Patch Mesh**) or close the *edit* panel the undo and redo memory is lost.

Creating Quadrilateral Meshes on Patches

This section describes how to create quadrilateral meshes on patches.

The **Patch Mesh** mode in the *Edit* panel is activated once the patch topology is complete, that is, all the patch curves are green. The patch mesh mode is used to define the mesh properties for the patch curves, including the number of divisions on each mesh curve, stretching distribution and patch mesh smoothing. STAR-CCM+ applies the same mesh properties to a selected patch curve and its corresponding mapped curves. Mesh curves for which the mesh properties are defined appear in black. Once all curve series have a defined mesh distribution, STAR-CCM+ automatically constructs and displays the mesh.



To create a patch mesh:

1. In the *Edit* panel, select the Patch Mesh mode.



This mode contains the *Mesh Properties* box.



2. Select the patch curve for which you want to set mesh properties.

   STAR-CCM+ automatically selects all associated curves. Also, you can select more than one base curve, and STAR-CCM+ automatically includes all the associated curves.

3. In the *Mesh Properties* box, set the **Number of Divisions** to the desired mesh division amount.

   After entering this value, press `<Enter>`.

4. In the *Distribution* group box:

   1. Set the type of stretching for the mesh:

      - **Constant** — The space between divisions is equal.
      - **One Sided Geometric** — The space between divisions increases in one direction with a geometric function. Specify the spacing at the start of the stretching.
      - **One Sided Hyperbolic** — The space between divisions increases in one direction with a hyperbolic function. Specify the spacing at the start of the stretching.

- **Two Sided Hyperbolic** — The space between divisions increases in both directions with a hyperbolic function. Specify the spacing at the start and end of the stretching.

2. To reverse the stretching direction, activate the **Reverse Direction** checkbox.
3. Click **Apply** to apply the distribution settings and generate the patch mesh.

The direction of the stretching is shown in the *Graphics* window with a red arrow. The following is an example of using the distribution tool with the surface mesh.



5. In the *Smooth* group box:
   1. To smooth the patch mesh when the patch topology is distorted or uneven, set the **Number of Iterations**.
   2. Activate the **Apply Distribution after Smoothing** box when the smoothing process distorts the selected stretching function.
   3. Click **Apply** to apply the smoothing settings and generate the patch mesh.

STAR-CCM+ applies the smoothing after the patch mesh is generated. Patch vertices are moved from their original position so that they no longer lie on the patch curves. The original vertex locations of are still shown as are the original curves. Dashed lines represent the original curves.

Note: If you increase the number of iterations, you achieve greater smoothing at a greater computational effort. The value of smoothing iterations is absolute and not cumulative from previous iteration values.

In the following screenshot, the mesh on the right shows smoothing of the mesh on the left by moving the vertices.



If you want to change the patch layout after creating the patch mesh, change to *Patch Topology* mode. Any changes that you make to patch mesh properties are automatically applied to the mesh shown in the display.

A new **Patch Mesh** node is added under **Source Meshes** once you close the patch mesh *Edit* panel. If you right-click the new node and select **Edit...**, you can return to the *Edit* panel and review or edit the patch definition.


Directed Mesh
Source Surfaces
Target Surfaces
Connected Parts
Cylinder
Source Meshes
Patch Mesh
Mesh Distributions

Creating a Source Mesh Using an Existing Mesh

It is possible to use an existing volume mesh as the source surface mesh.

The resulting directed mesh is conformal at the interface between the two volume meshes.



You can map an unstructured mesh to the source surface.



The following requirements must be met before an existing mesh can be extracted for the directed meshing part:

- The directed meshing part and the existing volume mesh part must be CAD imprinted to create an in-place part contact.
- If you are connecting to multiple existing volume meshes, all the existing meshes must belong to the same mesh operation or mesh continuum.
- Where the existing volume mesh and directed mesh part surfaces are coincident, the part surface boundaries must match up and share the same edges. You can use the split by patch option to separate the necessary part surfaces.
- Create a separate part surface for each side of the interface.

  See Splitting A Part Surface By Patch

To create a source mesh using a previously defined volume mesh:

1. In the *Graphics* window select one or more source surfaces.
2. Right-click one of the source surfaces and select **Part Collection > New Source Mesh > Use Existing Mesh**.



Alternatively:

1. Right-click the **Source Meshes** node in the *Directed Mesh* panel and select **New Source Mesh** > **Use Existing Mesh**.



2. In the *Part-collection for this source mesh* dialog, select one or more parts for the source mesh.

A **Source Meshes** > **[existing mesh]** node appears. Each source mesh node is named after the contacting part collection which the source mesh is defined on.



The properties of the **[existing mesh]** node contain the Part Collections that you selected for this distribution.

Existing Mesh Properties

Source mesh from	Specifies the source mesh	
	Contacting Volume Mesh	Extracts the mesh from a volume mesh that contacts the specified source surfaces.
	Triangle Mesh	Extracts the latest triangle mesh of the part or parts being meshed. If the original triangulation is obtained, the source mesh that is extracted may be of poor quality if it contains long, narrow triangles. If the part is remeshed, the remeshed

		triangulation is obtained.
	Quadrilateral Mesh	Obtained from the extracted triangulation. The mesh is dependent upon the quality of the extracted triangulation. The quadrilateral mesh that is obtained is quad-dominant with a mix of quads and triangles. The quadrilateral mesh does not faithfully replicate the extracted triangle mesh density.
Part Surfaces	Displays the parts that are specified in the Connected Parts node.	

Note: If the Source mesh from option is changed to Quadrilateral mesh, the dirty flag is set on the Use Existing Node, the corresponding Part collection, and the Directed Mesh Operation.

Updating the Source Mesh can be accomplished using one of the following options:

- Right-click the **Use Existing** node and select **Update**.
- Right-click the **Part Collection** and select **Generate Volume Mesh**.
- Right-click the **Directed Mesh** operation and select **Execute**.

When using the Automatic source meshing, all parts in the source surfaces should have remeshed surfaces. If they do not, then the initial surface for any non-remeshed part is used as the source surface for all the parts in the collection.

Initial Source Surface



Extracted Triangle Mesh



Quad/Quad-Dominant Mesh obtained from triangulated mesh



Updated Volume Mesh



STAR-CCM+ maps the existing mesh to the source surface.


Geometry Scene 1
Mesh Scene 1
Y
X

Deleting a Source Mesh

This section describes how to delete a source mesh.

To delete a source mesh:

1. Right-click the source mesh in the *Graphics* window and select **Delete Source Mesh**.



Creating a Volume Distribution

A directed meshing volume mesh is generated using a volume mesh distribution.

Once you have specified a patch mesh or an existing mesh for the source surfaces, you assign a volume mesh distribution to those parts.

To create a volume distribution:

- Right-click the guide surface in the *Graphics* window and select **Part Collection > New Volume Distribution**.



Alternatively:

1. Right-click the **Mesh Distributions** node and select **New Volume Distribution**.



2. In the *Parts for creating a new mesh distribution* dialog, select the **Connected Parts** for which you want to apply the distribution. You can apply the same distribution to any number of connected parts, regardless of how the source surface mesh is defined.

A new **Volume Distribution** node is added under **Mesh Distributions**.

The properties of the **Volume Distribution** node contain the Part Collections that you selected for this distribution.

Volume Distribution Properties

Part Collections	Sets the part collection to which this distribution is applied.
Number of Repetitions	Sets the number of times the volume mesh distribution is repeated between the source and the target surface. This repetition can be useful when meshing stacks of parts.

To set the number of layers in the volume distribution, select the **Volume Distribution** > **Number of Layers** node.

Number of Layers	Sets the number of cell layers to create between the source and target surfaces in the volume mesh.

To set the stretching function within the volume distribution, select the **Volume Distribution** > **Stretching Function** node.

Stretching function	Sets the stretching function for the mesh layers.	
	Constant	Constant cell distribution
	One Sided Geometric	Geometric progression to compute cell thickness

	One Sided Hyperbolic	Hyperbolic based stretching to compute cell distribution
	Two Sided Hyperbolic	Hyperbolic stretching from both ends to compute cell distribution

Note: When working with stacks of parts and parts that are connected at the guide surface, STAR-CCM+ attempts to fit the requested distribution to the stack of parts. The requested distribution in the resulting mesh cannot always be guaranteed since divisions close to the end of the part are moved to the end of that part.

If the stretching function is set to One Sided Geometric or One Sided Hyperbolic, a **One Sided Options** node is added under **Volume Distribution > Stretching Function**. Select the stretching mode that you want to use in the *Properties* window of this node and its value in its sub-node.

One Sided Options Properties

Stretching mode	Specifies the type of parameter that calculates the stretch value.	
	Stretching	Ratio of spacing at layer n and layer n+1
	Initial Wall Thickness	Spacing at the first layer of cells
	Thickness Ratio	Ratio of spacing between first layer and last layer of cells
Reverse Direction	When activated, reverses the direction of the stretching function.	

If the stretching function is set to Two Sided Hyperbolic, the **Two Sided** node is added under **Volume Distribution > Stretching Function**. Set the initial wall spacing at the source and target surfaces in the *Properties* window of this node.

Two Sided Properties

Spacing at Source	Sets the initial wall spacing at the source surface.

Spacing at Target	Sets the initial wall spacing at the target surface.
Reverse Direction	When activated, reverses the direction of the stretching function.

Deleting a Volume Distribution

This section describes how to delete a volume distribution.

To delete a volume distribution:

1. Right-click the guide surface in the *Graphics* window and select **Part Collection > Delete Volume Distribution**.



Alternatively:

2. Right-click the **Mesh Distributions > Volume Distribution** node and select **Delete**.

Generating the Directed Volume Mesh

You can generate the volume mesh for a single part or for several parts at the same time in the Directed Mesh panel.

Assign parts to a region before generating a volume mesh.

For a single part:

1. Right-click a specific part under the **Connected Parts** node and select **Generate Volume Mesh**.

For all the parts:

2. Choose one of the following options:
   - Right-click the **Directed Mesh** node and select **Execute** to generate volume meshes for all connected parts at the same time.



   - Click the (**Generate Volume Mesh**) button in the toolbar.

     This option generates all meshes that are defined for the regions, that is, directed meshes and those set using the standard STAR-CCM+ mesh continuum models. The directed meshes are generated after any region meshes.

The display is updated with the volume mesh.



The volume mesh is automatically transferred to the regions associated with the directed meshing parts. STAR-CCM+ also creates a **Parts Meshes** mesh continuum, and a basic physics continuum.



Note: If you clear the mesh or change the parts to regions assignment, use the **Generate Volume Mesh** button to re-execute the directed mesh operation. Doing so updates the regions and the mesh appropriately.

Checking the Volume Mesh

This section describes the tools available for checking the validity and quality of the volume mesh, and for diagnosing problem areas. Such problems can then be fixed, either inside STAR-CCM+ or in the originating mesh generator.

When using an external mesh generator, it is important to ensure that the volume mesh file or files were imported properly, by reviewing the import progress report and visually inspecting the mesh.

After the mesh has been successfully imported or generated, check the coordinate extents of the mesh to verify that the imported mesh is scaled to the correct dimensions. Incorrect scaling of the mesh adversely affects the solution. The mesh diagnostic report allows you to verify the mesh scaling by checking the mesh extents. If there is a problem, STAR-CCM+ provides functionality to rescale the entire mesh accordingly.

A common symptom of an invalid mesh is zero or negative cell volumes. In some cases negative volume cells are due to incorrect face orientation, which must be addressed in the mesh generator. It is often worth verifying that there are no negative cell volumes before any interfaces are created, in order to remove the effects of the interfaces on the mesh. In addition, checks for near-zero area faces are also explained.

Even though a mesh is valid and the solution can be successfully initialized with an appropriate Space model, poor mesh quality can negatively impact the solution. The direct measures of mesh quality in STAR-CCM+ are:

- cell skewness angle;
- boundary skewness angle;
- face validity metric;
- cell quality metric; and
- volume change metric

A description of how to check the cell quality metric is provided, which is also applicable to checking and visualizing the other cell metrics. For

more detailed diagnosis of mesh quality issues, the visualization tools in STAR-CCM+ can be used for identifying and visualizing bad cells.

In extreme cases, invalid cells can be removed so that you can continue running the analysis and obtain a solution.

What Is a Valid Mesh?

Fundamentally, a valid mesh is one that allows the simulation to be initialized and then run so that a solution is achieved. An invalid mesh is typically one for which no solution can be obtained or cannot even be initialized. As an example of the latter situation, a mesh that contains a negative volume cell would be considered invalid.

Based on that definition, the determination of whether a mesh is valid is not only dependent on the underlying mesh quality, but also the physics models in use with the mesh. For example, meshes with cells with high concave angles would not normally cause a problem, but in cases with high gradients or complex physics, a solution may not be obtainable. In this case, the mesh is considered invalid.

Other factors that can affect whether a mesh is valid include:

- choice of solver (coupled or segregated)
- under-relaxation or Courant number specification; and
- discretization scheme that is used

The approach that is used to define a CFD analysis is also important. For example, adding the physical models in one by one, iterating and restarting for each might lead to a valid solution. However, adding in all the physical models at once and attempting to run the simulation might cause the solution to diverge for a given mesh.

The mesh validity checks that are performed in STAR-CCM+ are:

- Unclosed cells;
- Invalid cell/vertex references;
- Zero area faces; and
- Zero or negative volume cells

What Is Mesh Quality?

In general, a poor quality volume mesh does not necessarily result in solver run time problems, but it does reduce the accuracy and efficiency of the solution obtained.

Consider the following global factors in assessing the overall quality of the mesh:

- mesh density - is the mesh sufficiently dense in high gradient areas to capture flow features accurately?
- mesh distribution - is the distribution of the mesh between high and low gradient areas optimal?
- near wall layers - does the mesh contain near wall layers for calculating turbulence and heat transfer effects?

On a local (cell to cell) level, consider the following factors:

- aspect ratio - is the aspect ratio reasonable for the expected flow regime?
- y+ range - are the y+ values reasonable for the turbulence model selected?
- volume ratio - is the ratio of a cell volume to its neighbor reasonable?
- skewness angle - are the cell skewness angles extreme?

There are many measures of cell quality, some of which are suited to specific cell shapes. The measures currently available in STAR-CCM+ are:

- the cell and boundary skewness angles;
- face validity metric;
- cell quality metric; and
- volume change metric

For fluid flow calculations, grid quality has most impact on the calculation of the convective and diffusive fluxes:

- To obtain convective fluxes with optimum accuracy and robustness, the line connecting the centroids of neighboring cells must ideally pass as closely as possible to the centroid of the face separating the cells. A two-dimensional example of this problem is as follows:



In this example, refining the mesh to create lower aspect ratio cells near the corner would reduce the problem.

- To obtain diffusive fluxes with optimum accuracy and robustness, the line connecting the centroids of neighboring cells must ideally be as close to orthogonal as possible to the face separating the cells. This angle is the cell skewness angle in STAR-CCM+. A two-dimensional example with high skewness angles is shown below:

In this example, refining the mesh to create lower aspect ratio cells would reduce the problem.

Running Mesh Diagnostics

This section describes how to run a mesh diagnostic report and how to interpret the output.

A mesh diagnostics report is a useful way of checking the validity of the mesh and getting statistics on the mesh quality. Two different types of reports are possible, a compact one with basic information on the model and a more extensive full report with more information on the boundaries. It is also possible to produce a diagnostics report for a surface mesh.

Running the Mesh Diagnostics Report

A mesh diagnostics report is a useful way of checking the validity of the mesh.

To run a mesh diagnostics report for a volume mesh, select the **Mesh > Diagnostics...** menu option.



The *Mesh Diagnostics* dialog then opens with all regions selected.



Deactivate the checkbox of any regions for which you do not want a report. The **Representation** selection lists all the current representations available that a diagnostics report can be produced for. If a volume mesh exists, the **Volume Mesh** option is selected by default, but a surface mesh representation can be selected instead. Select a **Report Type**, either **Compact** or **Full**.

Pressing the **OK** button generates the specified report and displays it in the *Output* window of the current simulation. The information that is displayed is dependent on whether the compact or full option is selected for **Report Type**. For a surface mesh representation, a surface mesh report is produced. In all instances information summarizing the entity count, mesh extents and mesh quality is produced.

If a Space model has not been selected before running the mesh diagnostics report for a volume mesh, then a three-dimensional physics model is automatically created in the **Continua** node for the simulation.

Interpreting the Compact Report

A compact mesh report for a volume mesh consists of statistics that are split up into statistics for each mesh region, and overall statistics.

For each region, the following subsections exist:

- Entity count
- Mesh extents
- Mesh validity check
- Face validity statistics
- Volume change statistics

The overall section also consists of the above items for all regions, minus the mesh extents summary. An example compact output for a single region is shown below:

```
Output - Star 1
Loading module: MetricsModel
----------------------------------------------------
--- Computing statistics in Region: Region 1
----------------------------------------------------
-> ENTITY COUNT:
   # Cells: 53192
   # Faces: 276966
   # Verts: 226439
-> EXTENTS:
   x: [ 0.0000e+00,  4.3180e-01] m
   y: [-1.0155e-01,  0.0000e+00] m
   z: [-1.0160e-01,  1.0160e-01] m
-> MESH VALIDITY:
   Mesh is topologically valid and has no negative volume cells.
-> FACE VALIDITY STATISTICS:
Minimum Face Validity: 1.000000e+00
Maximum Face Validity: 1.000000e+00
        Face Validity <  0.50          0    0.000%
0.50 <= Face Validity <  0.60          0    0.000%
0.60 <= Face Validity <  0.70          0    0.000%
0.70 <= Face Validity <  0.80          0    0.000%
0.80 <= Face Validity <  0.90          0    0.000%
0.90 <= Face Validity <  0.95          0    0.000%
0.95 <= Face Validity <  1.00          0    0.000%
1.00 <= Face Validity              53192   100.000%
-> VOLUME CHANGE STATISTICS:
Minimum Volume Change: 3.741970e-03
Maximum Volume Change: 1.000000e+00
         Volume Change <  0e+00          0    0.000%
0e+00 <= Volume Change <  1e-06          0    0.000%
1e-06 <= Volume Change <  1e-05          0    0.000%
1e-05 <= Volume Change <  1e-04          0    0.000%
1e-04 <= Volume Change <  1e-03          0    0.000%
1e-03 <= Volume Change <  1e-02         15    0.028%
1e-02 <= Volume Change <  1e-01       1963    3.690%
1e-01 <= Volume Change <= 1e+00      51214   96.281%
----------------------------------------------------
--- Computing statistics in Region: Region 2
----------------------------------------------------
```

For a compact report, the entity count includes a summary of the number of cells, faces, and vertices for the region. A more extensive

breakdown of the different cell shapes used is provided in the full report.

The mesh extents are always reported in meters based on the global cartesian system. STAR-CCM+ computes all physical quantities in SI and therefore assumes that the model is dimensioned correctly in meters. If a mesh was created using some other unit and was not properly scaled before import, it will result in incorrect physics being simulated. The mesh extents are originally reported at the end of the import report in the *Output* window for the simulation. Verify the coordinate extents for each region to make sure that they correspond with the expected dimensional range. If the extents are incorrect, you can:

- scale it in the original meshing pre-processor and reimport;
- scale the entire mesh in STAR-CCM+; or
- scale a specific region in STAR-CCM+

Once you have scaled the mesh or region, its worth running the mesh diagnostics report and checking the mesh extents again.

Interpreting the Full Report

The full diagnostics report is structured in a similar way to the compact report and contains all the information the compact report produces.

In addition however, the full report includes the following extra items for each region and the overall summary:

- Boundary details
- Entity count per cell shape
- Volume range
- Maximum interior skewness angle

The boundary details can be further split up as follows:

- Number of faces
- Boundary extent
- Surface area
- Maximum boundary skewness angle

An example output of the full report with the boundary information is shown below:

```
Output - Star 1

Boundaries of region Region 1:
Boundary Interface: 18568 faces (378 quadrilateral, 18190 polygonal)
 Extents:
  x: [0.0000000000e+00, 4.3180000000e-01] m
  y: [-1.0155037851e-01, 0.0000000000e+00] m
  z: [-1.0160000000e-01, 1.0160000000e-01] m
 surface area: 2.2217960656e-01 m^2
 maximum boundary skewness angle = 5.2533542633e+01 deg
Boundary Inlet: 60 faces (23 quadrilateral, 37 polygonal)
 Extents:
  x: [4.3180000000e-01, 4.3180000000e-01] m
  y: [-1.8945986797e-02, 0.0000000000e+00] m
  z: [1.9050000000e-02, 5.7150000000e-02] m
 surface area: 5.6100578513e-04 m^2
 maximum boundary skewness angle = 3.7604660034e+01 deg
Boundary Outlet: 61 faces (26 quadrilateral, 35 polygonal)
 Extents:
  x: [4.3180000000e-01, 4.3180000000e-01] m
  y: [-1.8958928660e-02, 0.0000000000e+00] m
  z: [-5.7150000000e-02, -1.9050000000e-02] m
 surface area: 5.6184205459e-04 m^2
 maximum boundary skewness angle = 3.4550983429e+01 deg
Boundary Symmetry: 3108 faces (803 quadrilateral, 2305 polygonal)
 Extents:
  x: [0.0000000000e+00, 4.3180000000e-01] m
  y: [-3.9837506798e-17, 0.0000000000e+00] m
  z: [-1.0160000000e-01, 1.0160000000e-01] m
 surface area: 3.0161229894e-02 m^2
 maximum boundary skewness angle = 5.7758831024e+01 deg
Boundary Interface [In-place 1]: 0 faces
```

An example output of the full report with the other additional information is shown below

```
Output - Star 1
Region Region 1:
     430 hexahedral cells
      57 wedge cells
   52705 polyhedral cells
   53192 cells total
  276966 interior faces (9562 triangular, 120452 quadrilateral, 146952 polygonal)
  226439 vertices
 Extents:
  x: [0.0000000000e+00, 4.3180000000e-01] m
  y: [-1.0155037851e-01, 0.0000000000e+00] m
  z: [-1.0160000000e-01, 1.0160000000e-01] m
 Maximum interior cell index delta: 459, average: 1.4594911650e+02
 Maximum cell face index delta: 2230, average: 1.2235293277e+03
 Volume range: [1.9417767394e-10, 5.1441674032e-07] m^3
 Minimum distance between centroids of neighbor cells = 6.1174513819e-04
Calculating skewness angle in Region 1
 Maximum interior skewness angle = 7.4360519409e+01 deg
```

Interpreting the Report for a Surface Mesh

The mesh diagnostics report for any of the surface type representations (import, wrapped surface, remeshed surface, or extracted surface) consists of sections for each boundary and each region.

Each of these sections has subsections for the following:

- Triangles, vertices, and marked feature edges count
- Boundary or region extents
- Boundary or region triangle quality
- Edge length statistics

Triangle quality is listed in 10 separate bands from 0 to 1, with 1 being the best quality. The edge length statistics are useful for gauging initial base values to be used as input for the surface wrapper and/or surface remesher tools.

An example output is shown below:

Output - **Star 1**

```
***** Mesh report for representation: Import
---------------------------------------------------
Boundary Wall Boundaries:
    4438 triangles
    2221 vertices
     165 marked feature edges
Extents:
  x: [-2.7559058873e-01, 1.0254813695e-02] m
  y: [-1.1218700822e-01, 1.1219020904e-01] m
  z: [-8.2640535249e-02, 1.4221760232e-01] m
Boundary triangle quality:
 Face Quality    Faces    Percent
  0.0 - 0.1          0       0.0
  0.1 - 0.2          0       0.0
  0.2 - 0.3          0       0.0
  0.3 - 0.4          0       0.0
  0.4 - 0.5          1       0.0
  0.5 - 0.6          1       0.0
  0.6 - 0.7          0       0.0
  0.7 - 0.8          3       0.1
  0.8 - 0.9         24       0.5
  0.9 - 1.0       4409      99.3
 Worst quality is 4.9e-01 on face 4438
 Edge lengths:
  Min: 5.4e-03
  Avg: 1.1e-02
  Max: 1.8e-02
---------------------------------------------------
Region Region 1:
    4438 triangles
    2221 vertices
     165 marked feature edges
Extents:
  x: [-2.7559058873e-01, 1.0254813695e-02] m
  y: [-1.1218700822e-01, 1.1219020904e-01] m
  z: [-8.2640535249e-02, 1.4221760232e-01] m
Region triangle quality:
 Face Quality    Faces    Percent
  0.0 - 0.1          0       0.0
  0.1 - 0.2          0       0.0
  0.2 - 0.3          0       0.0
  0.3 - 0.4          0       0.0
  0.4 - 0.5          1       0.0
  0.5 - 0.6          1       0.0
  0.6 - 0.7          0       0.0
  0.7 - 0.8          3       0.1
  0.8 - 0.9         24       0.5
  0.9 - 1.0       4409      99.3
 Worst quality is 4.9e-01 on face 4438
 Edge lengths:
  Min: 5.4e-03
  Avg: 1.1e-02
  Max: 1.8e-02
```

Checking The Mesh Validity

The mesh validity is automatically checked whenever a mesh diagnostics report is produced or when the simulation is initialized for the first time.

STAR-CCM+ performs the following mesh validity checks:

- Unclosed cells;
- Invalid cell/vertex references;
- Zero area faces; and
- Zero or negative volume cells.

If any of the above errors are found, except for near-zero area faces, then the mesh is invalid and must be fixed or regenerated in order to continue with the simulation. Alternatively, if the problem area is next to a surface (most issues usually are) improve the surface mesh in that area by hand before regenerating the volume mesh.

What Are Unclosed Cells

A topologically closed volume cell consists of edges that are connected to exactly two faces, and faces with normals pointing outwards. If a cell is missing a face, or the outward normals are inconsistent, it is considered unclosed and invalidates the mesh.

Sometimes this problem originates from the mesh exporter/importer or from the mesh generator itself. Fix any unclosed cells in the mesh by either deleting the cell and replacing it with a valid cell, or orienting the face normals in the outward direction.

Unclosed cells (as defined above) can be created during interface generation process as the intersecting vertices are not included in the vertex lists of the adjacent interior faces. In most instances, unclosed faces that are reported after interface generation are not a problem.

Unclosed cells can be visualized by creating a threshold part for cells with an unmatched edge count greater than zero. To create a new threshold part, set the Scalar field function to `Unmatched Edge Count`, set the Extraction Mode to **All Above** and enter a **Threshold Value** of `0`.

What Are Invalid Cell Vertex References

Invalid cell or vertex references usually originate from the mesh exporter/importer or the mesh generation process and must be fixed before proceeding any further. Deleting the cell and replacing it is usually the quickest and easiest solution to this problem.

STAR-CCM+ stores the mesh connectivity as an array of cells that are adjacent to each face together with a list of vertices belonging to the face. In a valid mesh, each cell interior face belongs to exactly two cells. Invalid cell and vertex references arise when the arrays storing the cell connectivity and vertices have entries that make reference to faces or vertices that fall outside the expected range of values.

Checking For Zero Area Faces

It is possible to generate a volume mesh containing cells that have faces of zero area magnitude. This situation would occur if all the vertices belonging to a face were coincident or collinear.

Since most calculations performed on faces involve flux calculations, a zero area face simply results in a zero value for the relevant flux. However, some calculations require a face unit normal, which is defined as the face area vector normalized by its magnitude. For this reason, STAR-CCM+ limits the face area magnitude that is used to a finite small value (less than 10^{-10} for a single-precision simulation) to prevent a divide-by-zero error. Zero area faces therefore do not usually cause a problem and can be ignored.

A face area calculation is performed whenever a mesh diagnostics report is produced or when the simulation is initialized. If any zero area faces are found, then the following message is displayed:



If desired, you can identify and visualize the cells that have near zero, or zero face areas.

Checking For Zero Or Negative Volume Cells

For a mesh to be valid for finite volume calculations, each cell must have a positive volume, otherwise the initialization process fails, and the simulation cannot be run. Negative volume cells can occur in a number of different ways but must be fixed or removed from the mesh in order to proceed any further.

1. To check if the volume mesh has zero or negative volume cells, you can either run the mesh diagnostics report or initialize the solution. An example report output for a simulation containing negative volume cells is shown below:



The total number of zero and negative volume cells are reported, along with a message reporting that the mesh is invalid. During the initialization step, an *Exception* dialog is displayed if there are zero or negative volume cells in the mesh:



It is recommended that you check for zero or negative volume cells by running the mesh diagnostics report immediately after importing or generating a volume mesh in STAR-CCM+. Do so avoids wasting time setting up a simulation for a mesh that is not valid.

Cells that have zero or negative volumes can be identified and visualized using threshold values.

How Do Negative Cell Volumes Occur

Negative volume cells occur for a number of different reasons.

These reasons include:

- Incorrect face orientation, which occurs if the pre-processing package used to create the mesh writes out the vertices using a left-handed rule, rather than the right-handed rule that STAR-CCM+ expects. In this situation, every cell in the region has a negative volume reported. Some tools, such as pro-STAR, have operations that can correct this problem and turn left-handed cells into right-handed ones.
- Badly formed cells. For hexahedral cells, the faces of high-aspect ratio cells can be highly warped leading to poorly defined face and cell centroids. For polyhedral shapes, negative volume cells can occur due to the dualization of poor quality tetrahedral cells, resulting in badly formed, and sometimes twisted shapes. Once you have identified these cells, go back to the pre-processor to repair them or remove them directly in STAR-CCM+.
- Interfaces that are created from incorrect or misplaced boundaries. Make sure that you check for negative volumes after you create interfaces in STAR-CCM+.

Using The Connected Cells Field Function

When the volume mesh generation is complete, STAR-CCM+ exposes a field function, **Connected Cells**, that has the same value for any pair of cells that share a face (interior or interface).

To check the connectivity of the volume mesh, use this field function in scalar scenes, threshold parts, reports, and other analysis tools. Initialize the solution when you want to check connectivity across interfaces between regions, or within a region that is divided into separate section by baffles.

This function could be useful when you want to:

- Check for any small non-contiguous cell islands in a single region which can cause the solvers problems. To remove problem cells, use the remove invalid cells option.
- Check that all pieces of a multi-region simulation are thermally connected. Interfaces must be intersected by initializing the solution.

Checking the Skewness Angle

The Cell skewness angle and boundary skewness angle are important measures of cell quality. This section describes the meaning of the terms and how to find the worst (highest) values in the mesh diagnostic report.

If the values for some cells are beyond desired limits, you can identify and visualize the bad cells before returning to the meshing pre-processor to correct the problems.

To view the maximum skewness angles, run the mesh diagnostic report using the full report type option. The cell skewness angle is reported for each mesh region:

```
Output - Star 1

Region Region 1:
    1027 tetrahedral cells
   23317 hexahedral cells
    6285 wedge cells
      31 pyramid cells
    6293 polyhedral cells
   36953 cells total
  106020 interior faces (9688 triangular, 90425 quadrilateral, 5907 polygonal)
   43669 vertices
 Extents:
  x: [7.4997538600e-05, 3.8107493000e-01] m
  y: [2.7863855055e-05, 6.6675762000e-02] m
  z: [1.4901161460e-08, 1.7780002540e-01] m
 Maximum interior cell index delta: 1217, average: 4.1822294850e+02
 Maximum cell face index delta: 3504, average: 1.5822524558e+03
 Volume range: [2.9906282312e-11, 4.7686967264e-06] m^3
 Minimum distance between centroids of neighbor cells = 2.7314738334e-04
 Maximum interior skewness angle = 8.0427368164e+01 deg

Face validity:
Minimum Face Validity: 1.0000000000e+00
Maximum Face Validity: 1.0000000000e+00
        Face Validity <  0.50          0    0.000%
0.50 <= Face Validity <  0.60          0    0.000%
```

The boundary skewness angle is reported for each boundary:

```
Output - Star 1

Boundaries of region Region 1:
Boundary Symmetry Plane: 2397 faces (104 triangular, 1977 quadrilateral, 316 polygonal)
 Extents:
  x: [7.4997538600e-05, 3.8107493000e-01] m
  y: [6.6675762000e-02, 6.6675762000e-02] m
  z: [1.4901161460e-08, 1.7780000000e-01] m
 surface area: 3.2238636166e-02 m^2
 maximum boundary skewness angle = 6.6457313538e+01 deg
Boundary Wall: 7620 faces (1910 triangular, 5147 quadrilateral, 563 polygonal)
 Extents:
  x: [7.4997538600e-05, 3.8107493000e-01] m
```

Scroll through the report and look for any `Maximum interior skewness angle` and `Maximum boundary skewness angle` values of more than 85°. If you find any large skewness angles, identify and visualize the bad cells before taking corrective action.

See Also

Omitting Secondary Gradients

What Is Cell Skewness Angle?

This skewness measure is designed to reflect whether the cells on either side of a face are formed in such a way as to permit diffusion of quantities without these quantities becoming unbounded.

To understand this measure, consider the following two-dimensional diagram of a face and the cell centroids on either side of the face:



The skewness angle θ is the angle between the face area vector **a**(face normal) and the vector connecting the two cell centroids, **ds**. An angle of zero indicates a perfectly orthogonal mesh. Variables that are stored at interior faces in STAR-CCM+ cannot be displayed; instead, the worst skewness angle for all the faces of a cell are stored in that cell. Thus, two cells often have the same skewness angle corresponding to the face that they share.

Skewness angles of 90° or greater, which can occur in concave cells where the centroid of both cells lies on the same side of the boundary

face, typically result in solver convergence issues. Problems result because the diffusion term formulation for transported scalar variables contains the dot product **a** · **ds** in the denominator, and this dot product is zero when the angle between **a** and **ds** is 90º. Even though STAR-CCM+ takes care to avoid such a divide-by-zero error, the accuracy of the diffusion calculation is reduced and there is a loss of robustness in such cases.

To reduce the impact on robustness, avoid skewness angles greater than 85º, as reported in the mesh diagnostics full report.

What Is Boundary Skewness Angle

The boundary face skewness can be as important as cell skewness. Boundary face skewness is defined as the angle between the area vector and the vector connecting the cell centroid and the boundary face centroid.



This angle is important since the dot product of these two vectors appears in a denominator in a component of the equation for diffusive fluxes at a boundary face.

In turbulent flows where wall functions are used, the equation for diffusive fluxes at a boundary is not used to compute diffusion fluxes at walls, so boundary face skewness become less important. However, it is important for laminar flows and for heat transfer in solids.

To reduce the impact on robustness, avoid boundary skewness angles greater than 85 degrees, as reported in the mesh diagnostics full report.

Checking The Face Validity, Cell Quality, And Volume Change Cell Metrics

Other than the skewness angle, three more cell metric checks are available to determine the overall quality of the volume mesh.

The metrics are:

- face validity
- cell quality
- volume change

In order to create and visualize the face validity and volume change metrics first initialize the solution with the appropriate physical models activated. For the above instances, you need only have the **Three Dimensional** *Space* model selected before doing the initialization.

Each of the above metrics can be identified and visualized using a threshold as part of geometry or scalar scene.

Face Validity Cell Metric

The face validity is an area-weighted measure of the correctness of the face normals relative to their attached cell centroid.



In a good quality cell, the face normals point outwards, away from the cell centroid. In a cell with bad face validity, one or more of the face

normals point inwards, towards the cell centroid. A face validity of 1.0 means that all face normals are correctly pointing away from the cell centroid. Values below 1.0 mean that some of the cell faces have normals pointing inward towards the cell centroid, indicating some form of concavity. Values of below 0.5 signify a negative volume cell.

A minimum required value for the face validity metric can be entered on the dialog for removing invalid cells.

Cell Quality Metric

The cell quality metric algorithm is based on a hybrid of the Gauss and least-squares methods for cell gradient calculation methods. It is a function not only of the relative geometric distribution of the cell centroids of the face neighbor cells, but also of the orientation of the cell faces. Generally, flat cells with highly non-orthogonal faces have a low cell quality.



A cell with a quality of 1.0 is considered perfect. A cubic cell is an example of such a perfect cell, but other polyhedral cell shapes can theoretically also attain this level of perfection. A degenerate cell has a cell quality approaching zero.

Depending on the physics that has been selected for the analysis, the cell quality of a cell can be fairly low and still provide a valid solution. However, poor cell quality is likely to affect both the robustness and accuracy of the solution. Cell quality is only defined for three-dimensional meshes.

A minimum required value for the cell quality metric can be entered on the dialog for removing invalid cells.

Volume Change Metric

The volume change metric describes the ratio of the volume of a cell to that of its largest neighbor.



A value of 1.0 indicates that the cell has a volume equal to or higher than its neighbors. As the cell volume decreases relative to its neighbors (such as a sliver or flat cell), then the volume change metric can be used to flag this cell. A large jump in volume from one cell to another can cause potential inaccuracies and instability in the solvers. Consider cells with a volume change value of 1e-05 or below to be suspect and investigate these cells further.

A minimum required value for the volume change metric can be entered on the dialog for removing invalid cells.

Identifying And Visualizing Bad Cells

To identify and visualize cells in STAR-CCM+ according to a given criteria, such as one of the quality metrics, use threshold values.

Two main types of threshold are available to achieve this goal:

- Derived part threshold
- Cell set threshold

Two techniques are available when using derived part thresholds to identify poor quality cells in the mesh:

- Identification and visualization using cell metrics. This method can be used for finding cells with negative volumes, high cell skewness angles, and low values of face validity, cell quality, and volume change.
- Identification and visualization using boundary metrics. This method is useful for finding cells with zero area faces and high boundary skewness angles.

The cell-set threshold options allow multiple threshold criteria to be defined for identifying cells, as well as providing tools for manipulating the cell collections.

Other than a tool to remove invalid cells, STAR-CCM+ does not currently have any volume mesh fixing tools to correct problems found using the threshold checks. However, it is useful to extract information such as the cell index ID or the location of the cell centroid to help identify the cells in the original pre-processing package. An example way of creating a XY plot showing bad cells that have been identified using a threshold is presented.

Identification And Visualization Based On Cell Metrics

Using cell-based threshold metrics, the bad cells in the interior of the mesh can be identified and isolated for further investigation.

This technique is applicable for finding cells that are based on:

- negative volumes;
- cell skewness;
- face validity metric;
- cell quality metric; and

- volume change metric

The method that is used to identify and visualize cells is identical for all the above scalars and is interchangeable once a threshold has been created. That is, the same threshold can be used to investigate any or all of the scalars one at a time by changing the properties of the node.

To access the above cell metrics as part of a threshold, first initialize the solution with the appropriate physical model mesh metrics activated. In this case, to access the cell metrics, select the **Three Dimensional** *Space* model before doing the initialization. If you do not perform this step, only the cell quality metric is then available by default.

A geometry scene must exist or one must be created in order to display the model. Simplifying the model to show only the outline of the geometry can assist the visualization of threshold cells. To display only the outline, edit the **Geometry 1** node for the scene so that it contains no parts, or hide all the parts. Alternatively, you can create a new displayer.

To create a threshold, right-click on the **Derived Parts** node and select **New Part > Threshold...** from the context menu.



The *Create Threshold* dialog appears, which allows the selection of the input parts, threshold scalar, extraction mode, value range, and display options.



To identify interior cells (as opposed to boundaries), the **Input Parts** list must contain the mesh region or regions nodes of interest. To edit the list, click the **Input Parts** list box to open the *Select Objects* parts editor:



Once you have selected the desired region or regions, click the **Close** button. The threshold scalar field function can be selected next. Click the *Scalar* pull-down menu, and scroll through the list to select the one of interest.



To check for one of the cell metrics that were listed earlier, choose from the scalars *Cell Quality*, *Volume*, *Skewness Angle*, *Face Validity*, or *Volume Change*.

To get an idea of the scalar value range, the *Scalar Range* **Query** button can be pressed. The *Extraction Mode* option can then be used to select how you want to identify the cells for the scalar threshold. Typically, you would want to select cells either above or below a given value but you can also select cells within a range or outside (below min and above max) a range. Use the **All Below** option here and a **threshold value** of `0.1`, meaning that all cells with a cell quality below 0.1 are identified.

The last item to decide is how the new threshold node is displayed. You can select not to create a displayer, to create a new geometry, scalar, or vector displayer, or add to an existing displayer. When selecting the last option, you can decide which pre-existing displayer the new node is added to. Use this option here and add the threshold part node to the **Geometry 1** displayer.



Click the **Create** button to create the threshold part. This may take several seconds depending on the number of regions and type of scalar selected. Press the **Close** button on the dialog once the process is complete. A new threshold node is created in the **Derived Parts** node. To view the new threshold node properties, click the node:



In the above case, the current geometry scene is updated with the results of the threshold check. With the **threshold** node selected, the cells are shown highlighted in the display:



You can rotate, translate, and zoom the display in order to investigate the poor quality cells. Thereafter, the *Properties* of the threshold node can be changed directly in order to update the display. Alternatively, you can right-click on the **threshold** node and select **Edit Part in Current Scene** in order to return to the original dialog used to create the threshold.



Identification And Visualization Based On Boundary Metrics

Boundary metrics can be used to identify bad cells on the boundaries.

This technique is applicable for finding boundary cells which have:

- zero area faces; and
- poor boundary skewness angles;

To identify the boundary surfaces using the boundary metrics that are listed above, follow the same procedure that was used for identifying and visualizing cells using cell metrics. However, for the **Input Parts** list in the *Create Threshold* dialog, select the desired boundaries instead of regions to operate on in the *Select Objects* editor. Additionally, to check for zero-area faces, use the Area: Magnitude option for *Scalar*.

This results in the threshold boundary surfaces being displayed at the end. However, this type of display is limited for the following reasons:

- Displaying the value of the scalar quantity on the boundary does not indicate the shape of the cell.
- If the quantity of interest is the face area magnitude and the face areas are almost zero, then the polygon representing the face is typically hard to locate and view.

To visualize the cells next to the threshold surfaces, create a new cell surface derived part using the threshold by right-clicking on the

Derived Parts node. Select **New Part > Cell Surface...** from the context menu.



The *Create Cell Surface* dialog appears, which allows the selection of the input parts and display options for the new part.



The **Input Parts** list in this case contain the threshold parts previously created. To select the nodes of interest, click the **Input Parts** list box and use the *Select Objects* parts editor:



Press the **Close** button on the editor when done. To display the new part, select an appropriate displayer. The easiest method is to add the new part to an existing geometry displayer.



To create the new part, click the **Create** button, and then **Close** to close the dialog. A new **Derived Parts** node that is called cell surface appears in the simulation tree:



The geometry scene is also updated with the results of the cell surface creation which uses green colored cells by default.



The full shape of the boundary cells can now be seen and inspected further by using the normal zoom, translate, and rotate options.

Creating Threshold XY Plots

A XY scatter plot of the cell or boundary thresholds against the pro-STAR cell index can be made in order to visualize the distribution of cell quality in the mesh.

An appropriate cell metric or boundary metric based threshold must be made first before you can analyze the bad cells in the mesh.

To create an XY plot, first right-click on the **Plots** node and select `New Plot > XY Plot`.



A new **XY Plot 1** node is created with the properties shown below:



Edit the *Parts* list in order to add the cell or boundary-based threshold. To open the *Select Objects* in-place dialog, click the *[] Parts* property, and select the threshold part that was created earlier. Add it to the Selected list by pressing the > button.



To close the editor, press the **Close** button. The cells in the threshold part are used as the basis to extract data for the XY plot, which means that only data for the bad cells are shown. Next, set the properties for the X and Y axes, both of which use scalars.

Open the **XY Plot 1** node and click the **X Type** node in order to view its properties. Set the *Type* property to **Scalar** by using the pull-down menu, as shown below.



Setting the *Type* to be a scalar, allows the use of pro-STAR cell index numbers for the X-axis. Open the **X Type** node and click the **Scalar** node underneath. In the *Properties* window for the node, set *Scalar* to be **Prostar Cell Id** by using the pull-down menu.



The X-axis of the plot then uses the pro-STAR cell index values from the threshold part, which are typically anywhere in the range of 1 to the highest cell number in the mesh.

The Y-axis scalar can now be set. For this example, use the same scalar as used for the cell threshold that you previously defined, namely cell quality. Open the **Y Types** node and click the **Y Type 1 > Scalar**node. In the Properties window, select **Cell Quality** for the *Scalar* option.



The XY plot now has sufficient information to update the scene with the data points.



The original derived part threshold properties can be used to control the number of cells and range that are used in the graph. The plot is useful to determine clustering or banding of data points for different cell and boundary metrics for the mesh.

Removing Invalid Cells

Poor quality cells can cause the solution to diverge or compromise the accuracy of the results. STAR-CCM+ provides a tool for removing poor quality mesh cells based on specified mesh criteria.

In STAR-CCM+, the **Removing Invalid Cells** tool allows you to remove mesh cells from the volume mesh region definition based on one or more of the following mesh criteria:

- Minimum Face Validity
- Minimum Cell Quality
- Minimum Volume Change
- Minimum Contiguous Cells
- Minimum Connected Face Area
- Minimum Volume
- Minimum/Maximum Function

STAR-CCM+ moves all the removed cells into a separate region that is not used in the analysis. Symmetry plane boundaries are automatically added to the neighboring cell faces of cells that have been removed, minimizing the impact of these removed areas on the solution.

To remove invalid cells:

1. Right-click the region or regions of interest and select **Remove Invalid Cells...**



The *Remove Invalid Cells* dialog appears.



2. In the *Remove Invalid Cells* dialog, select the relevant regions.
3. In the *Mesh Quality Metrics* group box, enable the criteria of interest and enter appropriate values.

   The following values are suggested guidelines:

   - **Minimum Face Validity:** 0.95 to 0.51 (default 0.51)
   - **Minimum Cell Quality:** 1e-05 to 1e-12 (default 1e-08)

     As the **Minimum Cell Quality** metric can sometimes flag valid trimmed cells (especially if they are anisotropic), you are recommended to use a low value.
   - **Minimum Volume Change:** 1e-05 to 1e-12 (default 1e-10)
   - **Minimum Contiguous Cells:** 1 to 1000000 (default 1)

- The **Minimum Connected Face Area** is dependent on the case dimensions and can be estimated based on the average mesh cell size.

If you change the metric values then they are automatically preserved as part of the STAR-CCM+ workspace and will be used instead of the default values in future sessions. To restore the default factory values, click **Reset Defaults**.

4. To preview the number of cells that each criteria would flag for removal, click **Update Counts**.
5. Optionally, you can click **Create Cell Set** to create a cell set with the cells that will be removed with the current settings.

   Any cell set will be added under the **Representations** > **Volume Mesh** > **Cell Sets** node.
6. When you are satisfied with the criteria, click **OK** to remove the cells. If you do not want to remove cells, click **Cancel**.

   The *Output* window displays the result:



Cells that do not meet the provided criteria are moved into a new region called **Cells deleted from [Region 1]**. The new region node appears in the simulation tree.



New symmetry plane boundaries with the names **internal-0** and **internal-1** are created in each region accordingly. The new region has no associated physics continuum, so it is not necessary to delete it in order to run the analysis.

Minimum Contiguous Cells

The minimum contiguous cells metric can be used to identify discontiguous regions of cells (as determined by the minimum connected face area metric). These cell regions can either pre-exist in the volume mesh or be created when bad quality cells are removed based on the other three invalid cell criteria. Discontiguous fluid cells would cause the solver to fail under most circumstances.

A value in the range 1 - 1000000 can be entered, where the value indicates the number of cells in the discontiguous region. On utilizing the feature, first the other criteria are applied and any invalid cells are flagged. Then any discontiguous regions of unflagged cells are checked and compared to the supplied metric value, as well as the supplied minimum connected face area value. If the number of cells is less than the amount specified, and/or the minimum area criteria is not met, then the region is also marked for removal.

Minimum Connected Face Area

The minimum connected face area is the smallest area that is allowed between two connected cells, below which the cells would be considered discontiguous and subject to the minimum contiguous cells check criteria.

It is sometimes possible during meshing and cell removal that is based on the other criteria, that two cells are left connected but with an extremely small face area between the two. Such scenarios can lead to the flux calculation becoming zero and the solution becoming unstable. Although topologically connected, if the face area between the two cells is below the specified value, then the cells are no longer considered to be connected. In this case, remove the smaller group of cells using the above contiguous cell check.

Minimum Volume

Use the minimum volume metric to remove all the cells with volume below a specified value. You can specify any value above zero.

Minimum/Maximum Function

The minimum/maximum function metric allows you to remove the cells for which a selected field function has a value below/above a specified limit. To specify the field function, click ... (**Custom Editor**).

Volume Meshing Field Functions Reference

Several field functions are made available in STAR-CCM+ as a result of mesh model selections or meshing operations.

This section provides details on some of these field functions.

Area	The **Area** field function displays the area of the cell faces. Area can be shown as either a Magnitude or as Laboratory coordinate values.	
	Magnitude	Displays the absolute value of area.
	Laboratory	Displays the value of area normal to the i, j, or k directions.
Boundary Index	This field function displays the index number that is associated with the boundaries. Each boundary has a unique boundary index that distinguishes it.	
Boundary Sliver Cell Indicator	This field function indicates cells that are "boundary slivers". Boundary slivers are flat cells on the outer boundaries of the solution domain that can be removed without affecting the topology of the mesh. A value of 0 indicates that the cell is not a boundary sliver cell, and a value of 1 indicates that the cell is a boundary sliver cell. (Available when the Cell Quality Remediation model is used). For more information, see the Remedying Cell Quality Field Functions Reference page.	
Cell Aspect Ratio	This field function is used to identify regular polygonal cells from skewed or stretched cells. A value of 1 indicates that the cell is a regular polygon, and a value close to 0 would represent a highly stretched cell. For more information, see the Modeling Space Field Functions Reference page.	

Cell Index	This field function displays the index for each cell. The index is only unique within the portion of the mesh for a given region on the local partition. So if you have R regions that are partitioned across P processors there are R*P cells with index 0 (assuming enough cells to allocate at least one per region per partition). It is not an ID, it is an offset.	
Cell Quality	Cell quality is a function that is determined from the relative geometric distribution of the cell centroids of the face neighbor cells and of the orientation of the cell faces. Generally, flat cells with highly non-orthogonal faces have a low cell quality. A cell with a quality of 1.0 is considered perfect. A degenerate cell has a cell quality approaching 0. Depending on the physics that has been selected, the cell quality of a cell can be fairly low and still provide a valid solution. However, poor cell quality is likely to affect both the robustness and accuracy of the solution. Cell quality is only defined for three-dimensional meshes.	
Centroid	The centroid is the location vector of the center of mass of the volume (or shell). The center of mass can be displayed as a Magnitude or by Laboratory coordinates. For more information, see the Modeling Space Field Functions Reference page.	
	Magnitude	Displays the absolute location of the center of mass.
	Laboratory	Displays the location of the center of mass in either the X, Y, or Z direction.
Chevron Quality Indicator	This field function identifies chevron cells. Chevron cells are pairs of thin slender cells which meet at a common face at an angle such that the line joining the cell centers does not pass through the common face. Chevron cells are marked as 1.0 and every other cell as 0. For	

	more information, see the Modeling Space Field Functions Reference page.	
Element Type	The element type identifies the specific star cell topology for each cell. Each number corresponds to a different type of topology, as shown below:	
	One Dimensional	
	3	STAR_LINE_CELL
	21	STAR_QUADRATIC_EDGE
	Two Dimensional	
	5	STAR_TRI_CELL
	7	STAR_POLYGON_CELL
	9	STAR_QUAD_CELL
	22	STAR_QUADRATIC_TRIANGLE
	23	STAR_QUADRATIC_QUAD
	Three Dimensional	
	10	STAR_TET_CELL
	12	STAR_HEX_CELL
	13	STAR_WEDGE_CELL
	14	STAR_PYRAMID_CELL
	24	STAR_QUADRATIC_TETRA
	25	STAR_QUADRATIC_HEXAHEDRON
	26	STAR_QUADRATIC_WEDGE
	29	STAR_TRIQUADRATIC_HEXAHEDRON
	41	STAR_UNIDENTIFIED_CELL

	42	STAR_POLYHEDRON
Face Index	Face index displays the partition-local index for each boundary face.	
Face Type	Face type displays the number of vertices per face. For example, a triangular face has a Face Type of 3 and a hexagonal face has a Face Type of 6.	
Face Validity	Face validity is an area-weighted measure of the correctness of the face normals relative to their attached cell centroid. In a good quality cell, the face normals point away from the cell centroid. In a cell with bad face validity, one or more face normals point towards the attached cell centroid. A face validity of 1.0 means that all face normals are correctly pointing away from the centroid. Values below 1.0 mean that some cell face normals point towards the centroid, indicating some form of concavity. Values of below 0.5 signify a negative volume cell.	
Interface Parent Face Index	This field function is the offset of the parent face for each direct boundary interface boundary face. Each face on a direct boundary interface boundary is a piece of an intersection of one face from each of the contributing boundaries. This index identifies those faces from the contributing boundaries.	
Least Squares Quality	Least squares quality is an indicator of the quality of a cell that is based on the physical location of a cell centroid relative to the cell centroid locations of its face-neighbors. (Available when the Cell Quality Remediation model is used). For more information, see the Remedying Cell Quality Field Functions Reference page.	
Partition	The primary partition that the cell/face/vertex is on. When a simulation is run in parallel, the regions are decomposed across the number of	

	processors such that each cell is allocated to a primary partition.	
Periodic Index	Repeating interfaces can require multiple passes to intersect all faces. For example, a rotor-stator (sliding interface) case intersects once in situ, rotates/translates one side by the periodic transform and intersects again, and repeats until all faces are consumed. For a given intersected face, this value reports the number of steps that were required until it was matched. So it indicates how many multiples of the periodic transform must be applied to align it with the other side of the interface.	
Position	This field function displays the position of the vertices relative to the origin. Position can be displayed as a Magnitude or as Laboratory coordinates. For a cell, the position is returned as the average of the positions of the vertices of the cell. For more information, see the Modeling Space Field Functions Reference page.	
	Magnitude	Displays the absolute position.
	Laboratory	Displays the position in either the X, Y, or Z direction.
Prostar Face Id	When a face on an imported STAR-CD mesh has been assigned a boundary, this field function returns the index of the original pro-STAR boundary. In pro-STAR, this number would be plotted in response to the `NUMB,BOUND,ON` command.	
Prostar Cell Id	This field function contains the original pro-STAR cell number for the imported STAR-CD mesh. In pro-STAR, this number would be plotted in response to the `NUMB,CELL,ON` command.	
Prostar Cell Type	This field function contains the pro-STAR cell table index number for the imported STAR-CD mesh. In pro-STAR, this number would be	

	plotted in response to the `NUMB,TYPE,ON` command.
Radial Coordinate	If an axis is defined, this field function provides the distance from the axis.
Region Index	This field function contains the region index. Each region has a unique region index that distinguishes it from other regions.
Skewness Angle	This field function contains the skewness angle of the cells. It indicates whether the cells on either side of a face are formed in such a way as to permit diffusion of quantities without these quantities becoming unbounded. To reduce the impact on robustness, avoid skewness angles greater than 85 degrees.
Unmatched Edge Count	This field function displays any unclosed cells. A topologically closed volume cell consists of edges that are connected to exactly two faces, and faces with normals pointing outwards. If a cell is missing a face or the outward normals are not consistent, then it is considered unclosed and invalidates the mesh.
Volume	This field function contains the volume of each cell.
Volume Change	This field function describes the ratio of the volume of a cell to that of its largest neighbor. A value of 1.0 indicates that the cell has a volume equal to or higher than its neighbors. A large jump in volume from one cell to another can cause potential inaccuracies and instability in the solvers. Investigate cells with a volume change value of 1e-05 or below as they can be considered suspect.
Zero Face Area Indicator	This field function has the value of unity if the area of any of the faces approaches zero; otherwise it has a value of zero. It is used to diagnose potential mesh problems. (Available when the Cell Quality Remediation model is used). For more information, see the

	Remedying Cell Quality Field Functions Reference page.

Working With Representations

This section provides information on representations as a useful part of the meshing process.

This section describes what representations are, how they are used as part of the meshing process and how they can be visualized. Replacing, exporting and extracting a surface mesh are also discussed. The creation of cell sets for volume mesh representations is also presented.

What Are Representations?

Representations help you keep track of the surface and volume meshes that STAR-CCM+ uses as part of the meshing and simulation setup process. Several different representations are available.

The types of representations are:

- **import** - imported surface mesh (using the regions import mode option)
- **geometry** - geometry surface (using the parts import mode option)
- **initial surface** - initial surface from geometry parts
- **wrapped surface** - surface wrapper mesh
- **remeshed surface** - surface remesher mesh
- **volume mesh** - volume mesh (with optional prismatic layers)
- **extracted surface** - surface boundary mesh that is extracted from a volume mesh
- **proxy** - placeholder that you apply to scenes, report, and plots so that you can control the representation of these items from one place
- **recorded solution** - copy of the volume mesh that is stored on an external solution history file (containing selected solution data)
- **overset mesh** - contains the background mesh and all overset region cells, regardless of whether they are active or inactive

The representations can be seen by clicking the **Representations** node in the simulation tree. For example:



A representation node is added whenever an operation results in the creation or import of that kind of mesh. For example, when a new simulation is made, no representations exist. When a surface mesh is imported using the regions import mode option, then an **Import**node appears in the simulation tree. Similarly, importing a surface as a part results in a **Geometry** node being created instead. If a volume mesh is imported or generated, then a **Volume Mesh** representation node is shown. Except for the **Geometry** node (which simply acts as a place holder), the properties of the mesh can be seen by clicking the node and viewing the *Properties* for it. For example, the total number of cells, interior faces and vertices for a volume mesh region can be seen by selecting the region name of interest:



Only one of each kind of mesh representation can exist for a given simulation although multiple extracted surfaces are allowed. The result of a second import operation adds the properties of the new mesh to the existing node. The result of generating a new mesh replaces the properties of the existing node. Representation nodes can be deleted like any other object, but at least one representation must always remain. All meshes (and appropriate representations) generated within STAR-CCM+ can be cleared in one go if preferable.

See Also

Preparing Surfaces

Displaying Representations

The new surface or volume mesh representations that are generated using the meshing tools are not automatically displayed in the plot window once the meshing process is complete.

The mesh representations can be visualized using three different methods:

- Selecting the appropriate representation from the scene displayer pull-down menu in the *Properties* window.
- Selecting the appropriate representation by right-clicking the graphics display.
- Dragging and dropping the representation node object onto the graphics display.

Only the representations that have so far been imported or generated are listed in the menus.

Displaying Representations using the Scene Displayer

Each scene displayer node provides access to the mesh representations that currently exist.

To change the current representation, select the scene displayer node of interest (for example, **Geometry 1**) and then click the *Representation* option in the expert *Properties* window. The currently available mesh representations are listed in the pull-down menu.



Select the appropriate mesh representation from the list and the scene updates with the new mesh. Any other scene displayer nodes (for example, **Outline**) for the same scene does not have the same representation selected for them - update these using the same technique as described above.

Displaying Representations via the Graphics Display

The pop-up menu in each scene display provides access to the mesh representations that currently exist.

To change the current representation, right-click any portion of the display that the model does not occupy and select **Apply Representation** in the pop-up menu that appears. The currently available mesh representations are listed.



Select the appropriate mesh representation from the list and the scene is updated with the new mesh. Any other scene displayer nodes belonging to the scene (for example, **Outline**) have the same representation selected for them automatically using this method.

Displaying Representations by Dragging and Dropping

In order to update the scene with the new mesh representation, each mesh node in the **Representation** node can be dragged and dropped onto the graphics display scene.

An example of using the drag-and-drop feature is illustrated below:



All displayers belonging to the scene are automatically updated to have the same representation using this method.

Using Proxy Representations

A proxy representation acts as a placeholder and allows you to control the representation of multiple objects (scenes, reports, and plots) from a single location.

Often, you have multiple plots and reports in your simulation. If you want to apply the same representation to all these objects, you could manually edit each object in turn and change the representation. Instead, by using a proxy representation you only need to set the representation once in the various objects. You can then control these representations using the proxy representation.

To create a proxy representation:

1. Right-click on the **Representation** node in the object tree and select **New Proxy Representation**.



A proxy representation, **Proxy**, appears in the object tree.



Note: You can create and use as many proxy representations as you want.

2. Navigate to the objects where you want to change the representation. Set the *Representation* to the **Proxy** representation.



3. Return to the **Representation** > **Proxy** node. In the *Properties* window, select the representation you want to apply to the objects. In the example, a solution history representation is selected.



The representation you select in the proxy is applied to every object that is assigned the proxy representation. You can quickly change the representation that is applied to the object by changing the proxy representation.

Replacing a Surface Mesh

Any of the five types of surface mesh representations, namely import, initial surface, remeshed surface, wrapped surface and extracted surface, can be replaced at any given time by a new surface read in from an external file.

There are two main reasons for wanting to replace a surface mesh:

1. The new surface contains fixes performed by some other third-party pre-processing software; and
2. The new surface contains a geometrical design change from the original

Replacing a surface does not disrupt the meshing pipeline and is a convenient method for performing multiple mesh runs on different surfaces since all parameter settings and values are preserved.

The currently allowable surface mesh replacement formats are:

- .dbs - pro-STAR surface database mesh file
- .inp - pro-STAR cell/vertex input file
- .nas - NASTRAN shell file
- .pat - PATRAN shell file
- .fro - FELISA front surface file

To replace a surface, right-click on the surface mesh name and select **Replace Surface...** from the pop-up menu. An example for the **Import** node is shown below:



The *Open* dialog appears.



Select the surface file that you want to import and click **Open** to import the surface. You are then prompted for the units for the new surface file before it is imported. If the underlying patch definitions for the original and replacement surface are the same, then the current boundary list is maintained. If extra patches exist in the new surface, then new boundaries are created to represent them. In instances where patches are missing from the new surface, deselect existing boundaries that become empty.

When using the .dbs or .inp formats for replacing a mesh, an option to **Match Boundaries by Name** is provided in the interim dialog. This option is turned off by default, meaning that no attempt is made to match old and new boundaries based on their names but instead their cell types are used. Turning on this option matches patches to boundaries based on their names instead of cell types. The process of matching boundaries by name is case-sensitive.

If you are replacing a surface mesh with a surface of .inp format that was exported from an existing STAR-CCM+ simulation, underscores replace any spaces in boundary names in the .inp file. As a result, STAR-CCM+ may be unable to match boundaries by name when the surface mesh is replaced. Under these circumstances, it is recommended that you export the surface mesh to the .dbs format instead.

To replace the volume mesh, use the import option in the **File** menu to reimport a whole new mesh or alternatively you can replace a single mesh region.

Preparing a Surface Mesh

Any of the surface mesh representations can be manually prepared using the Surface Preparation tool that is contained within STAR-CCM+.

Exporting a Surface Mesh

Any of the five types of surface mesh representations, namely import, initial surface, remeshed surface, wrapped surface and extracted surface can be exported from STAR-CCM+ for further manipulation.

The currently allowable export formats for surfaces are:

- .dbs - pro-STAR surface database mesh file
- .inp - pro-STAR cell/vertex input file
- .nas - NASTRAN shell file
- .pat - PATRAN shell file
- .stl - Stereolithography file
- .fro - FELISA front surface file
- .bgp - XML boundary grouping

All the above formats support single or multiple region definitions. All the formats other than PATRAN and STL support feature curves. It is possible to export the entire representation, single regions within the representation, or boundaries belonging to regions within the representation.

Exporting the Complete Surface

You can export the complete surface representation to one of multiple file formats.

To export the complete surface representation:

1. Right-click on the representation node of interest and select **Export Surface...** from the pop-up menu:



The *Save* dialog appears:



2. Select the file type from the **Files of Type** pull-down menu.

3. Enter the **File Name** and click **Save**.

   Details particular to the file format may be prompted for prior to the file being saved. In the case of Nastran, the default option is to export the geometry using the **Extended** file format, which provides greater node precision (G16.9) as compared to the **Standard** option (E8.3). A loss in surface quality is therefore encountered with the latter.

   An export only format (XML Boundary Grouping) is allowed here if you wish to reconstruct the regions/boundaries hierarchy established in STAR-CCM+ when the surface patch data (exported using any of the other allowable formats) is imported into a third-party package that also supports the XML boundary grouping format.

Exporting Individual Regions

As an alternative to exporting the entire surface, you can export individual regions belonging to each surface type representation.

1. To export a region, right-click the region of interest and select **Export Region...** from the pop-up menu:



The *Save* dialog appears as before, allowing just the surface mesh for the selected region to be saved in the desired format. Only one region can be selected at a time for this export process.

Exporting Boundaries

Boundaries belonging to regions within each surface representation can be exported, either singly, or with others selected at the same time. It is also possible to export multiple boundaries from different regions in a single operation.

1. To export one or more boundaries, first select their nodes, right-click, and choose **Export Boundaries...**



The *Save* dialog appears as for other export operations, allowing the selected boundary or boundaries to be saved in the desired format. The *File Name* field is automatically populated with the name of the selected boundary (or the first boundary if more than one is selected).

Not all formats preserve boundary names: only the two pro-STAR formats and the NASTRAN format retain actual names (with duplicate names receiving an additional integer index); PATRAN and FELISA formats retain individual boundary definitions without the names; the Stereolithography format combines all boundaries into one. Region definitions are not retained in any format.

Extracting a Surface Mesh

This facility extracts a surface mesh from a volume mesh representation so that you can generate a new volume mesh using the STAR-CCM+ meshing tools.

Surface mesh extraction is useful when you have a volume mesh (for example, a tetrahedral mesh) from another package, but want to remesh it using the different mesh types in STAR-CCM+. This can result in a lower cell count or better cell quality (or usually both).

Extracting a surface mesh maintains the boundary and interface definitions. However, exporting a surface mesh representation does not preserve boundary types, except from name association, and never preserves interfaces. Therefore, after importing an extracted surface, check the surface mesh, redefine interfaces and possibly redefine boundary types. If you are importing the surface mesh as a part, you may also need to merge/imprint coincident surfaces, so that interfaces are conformal. Also note, volume meshes do not store feature edges from the surface mesh. Therefore, when extracting a surface, redefine the feature edges on the extracted surface mesh.

Before extracting a surface mesh, a volume mesh representation must exist (generated when importing a volume mesh). To extract a surface mesh from a volume mesh:

1. Right-click the **Volume Mesh** node and select **Extract boundary surface** from the pop-up menu.



2. In the *Extract Boundary Surface* dialog, select the required region(s).



When the process is complete, a new representation called **Extracted Surface** appears in the simulation tree. Each extracted surface representation contains a closed surface mesh of the selected region(s)



If you extract multiple regions, multiple nodes appear in the **Representations** node with numerically incremented names.

The extracted surface can now be used as the initial surface mesh for a new meshing process within STAR-CCM+. If other surface mesh representations exist, delete them so that the meshing pipeline uses the correct surface.

To access the full range of boundary manipulation options after extracting a surface mesh, create a mesh continuum manually if one does not exist before performing the operation.

See Also

Extruder Mesher

Using Cell Sets

The purpose of a cell set is to allow you to identify collections of cells that are useful for visualizing the model, such as during post processing or mesh checking.

STAR-CCM+ allows cell sets to be created once you have a 3D volume mesh representation. Any number of cell sets can be created for a given volume mesh, with each one containing anywhere from zero to the total number of cells in the model.

A cell set is similar to a derived part in many respects except that it is more versatile and can be derived from more than one condition. For example, you can create a cell set based on a combination of a threshold and a zone in order to obtain the cell grouping you want. For this reason, it is a powerful tool for isolating and storing cell-based parts for further use.

All operations relating to cell sets take place via the **Cell Sets** node, which appears under the **Finite Volume Regions** node in the **Volume Mesh** representation in the simulation tree:



A new cell set can be created based on a cell list, a scalar threshold or a zone. Once a cell set has been created, it can be modified if desired using

a number of global options such as fill, clear, grow and shrink. Boolean operations are also possible to add to a set, take a subset of a set, remove from a set and replace a set altogether, using the same options as used to create the original set.

Cell set operations only allow the groupings of cells that exist in the model. No cells can be deleted or created by this process. The cell set operations described here are also currently only applicable to 3D meshes - if you have a 2D mesh then a cell set cannot be displayed.

Each cell set created within STAR-CCM+ has a corresponding field function whose *Function Name* is a contracted form of the cell set name. This field function can be used within other field function definitions to detect whether a cell belongs to the specified cell set or not. This feature is useful when applying region sources to a particular set of cells, for example.

Creating a New Cell Set

Generate a new cell set before any further manipulation can be done on the set or the set contents displayed.

By default a 3D volume mesh representation has no cell sets defined for it. If you have a 2D volume mesh, then a cell set cannot be displayed for it.

Four methods are available for creating a cell set:

- Using a cell list;
- Using a scalar threshold;
- Using a zone; and
- Using a geometry range

Each new cell set created is based on all the cells in the model, regardless of any previously generated sets. For the cell list and threshold options, a dialog window appears with the appropriate options to allow you to create the cell set, while the zone and geometric range options allow you to select the zone or block of cells of interest manually. In all four instances however, you are prompted as how to add the new cell set part to the current scene in order for it to be displayed.

To view the GUI options available for creating a cell set, right-click on the **Volume Mesh > Cell Sets** node and select the **New** option:



The **Empty** cell set option allows you to create an empty cell set node, which can be modified afterwards to include cells using any combination of the tools provided.

Once a cell set has been created, a new sequentially numbered node, starting from **Cell Set 1**, is created under the **Cell Sets** node. Clicking the node provides a summary of the number of cells that are contained in the set:



A field function is also created for each cell set by default, which appears under the **Tools > Field Functions** manager node. The field function has the same name as the cell set by default. For example:



This cell set field function can be useful is situations such as when a region needs to be split based on an existing collection of cells. Once the original cell set is deleted, the associated field function is also deleted.

The new cell set creation tools only allow a single action to create the original cell set. If you want to apply additional criteria to manipulate the set further, then modify the cell set using the range of global and Boolean tools available.

Cell set nodes can be renamed and deleted in the same way as other nodes. They can also be copied and pasted to perform Boolean operations between two different sets.

Creating a New Cell Set Using a Cell List

The cell list option for creating a cell set allows individual cell ID's or offsets to be used to identify cells. In general, these options are included for developers to debug problems as the ID's and offsets are not generally available or known to the average user.

Cell ID's are unique values assigned to each cell and correspond with the pro-STAR cell ID index. Cell offsets are used when the mesh has been partitioned for a parallel analysis.

To access the cell list options, right-click on the **Cell Sets** node and select **New > Cell List...** from the pop-up menu:



The *Cell List* dialog appears.



In the dialog, select the region of interest using the **Region** pull-down menu and then the method that is used to set the **Id Type**, namely **Cell Id** or **Cell Offset**. The individual cell ID's or offset values can then be entered into the white box, separated by spaces.

Then select the *Display* option to determine how the new cell set part is added to the current scene. Click **OK** to create the cell set based on the supplied information or **Cancel** to void the operation.

The *Cell List* dialog is automatically closed once the operation is complete. If you want to apply additional criteria to manipulate the set further, then modify the cell set using the range of global and Boolean tools available.

Creating a New Cell Set Using a Threshold

The threshold option allows cell sets to be created based on any of the normally available scalars such as centroid geometric range, face area magnitudes and volumes, as well as post processing items such as pressure, temperature and velocity etc, should they be available.

To access the threshold options, right-click on the **Cell Sets** node and select **New > Threshold...** from the pop-up menu.

The *Threshold* dialog appears:



First select the **Scalar** field function of interest using the pull-down menu. If desired, click **Query** to query the range of field values in the entire cell set. Set the **Extraction Mode** to determine the method that is used to extract the cells for the threshold. Enter the appropriate minimum/maximum or threshold value depending on the option selected.

Then select the *Display* option to determine how the new cell set part is added to the current scene. Click **OK** to create the cell set based on the supplied information or **Cancel** to void the operation.

The *Threshold* dialog is automatically closed once the operation is complete. If you want to apply additional criteria to manipulate the set further, modify the cell set using the range of global and Boolean tools available.

Creating a New Cell Set Using a Zone

The zone option allows cell sets to be created based on cells that are contained fully within an arbitrary closed loop polygon. As such, it is a very flexible and quick method of isolating cells with just a few mouse clicks.

Note: A valid scene must be displayed in order for this option to be active.

To access the zone options, right-click the **Cell Sets** node and select **New > Zone...** from the pop-up menu:



This activates the mouse cross-hair selection tool, shown below:



Click the scene to indicate the starting point of the polygon zone and then continue to click around the model to create a closed loop zone of the area of interest. An example is shown below.



To complete the zone, you can either right mouse-click or click back on the starting point of the zone. If at anytime you want to discontinue the zone, press the Escape key. Press the D key to undo the previous zone point selection.

On the successful completion of a zone, the *Zone* dialog appears to indicate the display choices:



Then select the *Display* option to determine how the new cell set part is added to the current scene. Click **OK** to create the cell set. If the appropriate display options are selected, the cell set automatically appears within the scene:



If you want to apply additional criteria to manipulate the set further, modify the cell set using the range of global and Boolean tools available.

Creating a New Cell Set Using a Geometric Range

The geometric range option allows cell sets to be created based on cells that are contained fully within a Cartesian based block, extending in the x, y and z directions.

If a valid scene is currently displayed, the minimum and maximum geometric coordinates can be set either graphically or manually, otherwise they can be entered manually by hand if a scene is not open.

To access the geometric range options, right-click on the **Cell Sets** node and select **New > Geometric Range...** from the pop-up menu:



The *Cell Set Geometric Range* dialog appears:



The dialog can be used to enter the maximum and minimum coordinate values in each of the three cartesian directions. Alternatively, the graphical display tool, which is automatically displayed in the scene with the bounding box extending to the model maximum and minimum coordinate directions, can be used to specify the geometric range:



Click and drag on the spherical balls to extend or reduce the box size in each direction. Right-clicking and dragging vertically on the box allow you to resize all the coordinates of the box simultaneously. Clicking using the middle mouse button on the box allows it to be translated in each direction. An example showing a reduced box size in the vicinity of the valve region of the model is shown below:



Then select the *Display* option on the dialog to determine how the new cell set part is added to the current scene. Click **Create** to create the cell set based on the supplied information. The new cell set is then displayed depending on the *Display* options selected:



Additional cell sets can be created by repeating the above operation. Press the **Close** button when you are done.

Creating a New Empty Cell Set

The advantage of using an empty cell set is that you are initially not presented with any options for how the set is created or displayed as part of the scene. You can then modify the cell set using the range of global and Boolean tools available.

An empty cell set node, that is, one containing no cells, can be created by right-clicking the **Cell Sets** node and selecting **New** > **Empty** from the pop-up menu:



Modifying an Existing Cell Set

Once a cell set has been created, it can be easily modified using a range of global and/or Boolean operations in order to manipulate the cell collection further.

In general, global actions are those which allow a single operation on the cell set resulting in collecting all the cells, removing all the cells or growing or shrinking the cell set in some way. Boolean operations allow the existing cell set to be modified using the same modes used for the creation of the original cell set, namely based on a cell list, zone or threshold.

The global actions that you can perform on a cell set are:

- fill - Replace the existing cell set with all the cells in the model
- clear - Remove all cells from the cell set
- invert - Invert the cell set to give all remaining cells

- grow - Grow the existing cell set by one layer based on exposed faces
- grow across interfaces - Grow cells across an interface
- shrink - Shrink the existing cell set by one layer based on internal faces
- shrink All - Shrink the existing cell set by one layer based on all faces

The Boolean operations that are allowed are:

- subset - Create a set based on the existing set only
- add - Add to the existing set based on all the cells in the model
- remove - Remove cells from the existing set
- replace - Create a set based on all the cells in the model

Boolean operations between different cell sets are also possible using the copy and paste options.

To access the global and Boolean operations, right-click on the cell set node as shown below:



You can also edit the cell set in the current scene, which provides all the above operations in one single in-place dialog. The advantage of this method is that multiple operations can be performed on the cell set without the dialog being closed each time.

Using the Fill Option

The fill option replaces a previous collection of cells with all the cells that currently exist in the model.

To access the cell set fill option, right-click the desired cell set node and select **Fill** from the pop-up menu:



An example before and after snapshot of the fill operation is shown below:



In this example, the original cell set contained a portion of the model cells. After the fill operation, the cell set contains all the model cells.

Using the Clear Option

The clear option can be used to clear all the cells from the cell set. The previous collection of cells is removed leaving the cell set empty.

To access the cell set clear option, right-click the desired cell set node and select **Clear** from the pop-up menu:



An example before and after snapshot of the clear operation is shown below:



In this example, the original cell set contained a portion of the model cells. After the clear operation, the cell set contains no cells.

Using the Invert Option

The invert option deletes the existing cells from the set and replaces them with all other cells in the model, that is, those cells that were previously outside the set.

To access the cell set invert option, right-click the desired cell set node and select **Invert** from the pop-up menu:



An example before and after snapshot of the invert operation is shown below:



In this example, the original cell set contained a portion of the model cells. After the invert operation, the cell set contains all remaining cells.

Using the Grow Option

The grow option can be used to expand the current cell set by a single layer of cells in all directions.

To access the cell set grow option, right-click on the desired cell set node and select **Grow** from the pop-up menu:



An example before and after snapshot of the grow operation is shown below:





In this example, the original cell set contained a portion of the model cells. After the grow operation, the cell set has expanded to include the nearest cell neighbors based on the faces of the original cell set.

Using the Grow Across Interfaces Option

The `Grow Across Interfaces` option can be used with the grow option to expand the current cell set by one layer across an interface.

A single layer of cells is grown from the interface faces of the existing cell set cells and added to the set from another region or from the opposite side of a baffle. Before this operation can be used however, the simulation must be initialized.

To access the cell set grow across interfaces option, right-click on the desired cell set node and select **Grow Across Interfaces** from the pop-up menu:



Wherever a cell face in the cell set belongs to an interface then the opposite cell connected to that face is grown. Thereafter, the normal grow option can be used to continue growing the cells.

Using the Shrink Option

The shrink option can be used to reduce the current cell set by one layer, removing cells that have a face exposed in the interior of the model. Cells with faces that are on a boundary are ignored.

To access the cell set shrink option, right-click the desired cell set node and select **Shrink** from the pop-up menu:



An example before and after snapshot of the shrink operation is shown below:



In this example, the original cell set contained a portion of model cells. After the shrink operation, the cell set has been reduced in size by one neighbor layer where the neighbor cells have a face exposed in the model interior.

Using the Shrink All Option

The shrink all option can be used to reduce the current cell set by one layer, removing cells that have a face exposed in the interior of the model or on a boundary.

To access the cell set shrink all option, right-click on the desired cell set node and select **Shrink All** from the pop-up menu:



An example before and after snapshot of the shrink all operation is shown below:



In this example, the original cell set contained a portion of the model cells. After the shrink all operation, the cell set has been reduced in size by one neighbor layer where the neighbor cells have a face exposed on the boundary or the model interior.

Using the Subset Option

The subset option allows you to create a set of cells based on the existing cell set only, rather than the entire model. Cells that are not contained in the original cell set are ignored.

To access the cell set subset options, right-click the desired cell set node and select **Subset** from the pop-up menu:



Four different subset options are allowed:

- cell list - Create a subset based on cell index ID's or offsets.
- threshold - Create a subset based on a scalar threshold.
- zone - Create a subset based on a zone.
- geometric range - Create a subset based on a geometric range.

An example before and after snapshot of a subset operation using a zone is shown below:



In this example, the original cell set contained a portion of the model cells. The subset operation using a zone was used to select the lower half of the set. Any cells in the model not contained in the original cell set were ignored in the new cell set.

Using the Add Option

The add option allows you to add cells to the existing cell set based on all the cells in the model.

To access the cell set add option, right-click the desired cell set node and select **Add** from the pop-up menu:



Four different add options are allowed:

- cell list - Add to the set based on cell index ID's or offsets.
- threshold - Add to the set based on a scalar threshold.
- zone - Add to the set based on a zone.
- geometric range - Add to the set based on a geometric range.

An example before and after snapshot of an add operation using a zone is shown below:



In this example, the original cell set contained a portion of the model cells. The add operation using a zone was used to select an additional area below the original to add to the set.

Using the Remove Option

The remove option allows you to remove cells from the existing cell set.

To access the cell set remove option, right-click on the desired cell set node and select **Remove** from the pop-up menu:



Four different remove options are allowed:

- cell list - Remove cells from the set based on cell index ID's or offsets.
- threshold - Remove cells from the set based on a scalar threshold.
- zone - Remove cells from the set based on a zone.
- geometric range - Remove cells based on a geometric range.

An example before and after snapshot of a remove operation using a zone is shown below:



In this example, the original cell set contained a portion of the model cells. The remove operation using a zone was used to select the left half of the set, which was removed, leaving the right half in the set.

Using the Replace Option

The replace option allows you to replace the existing cell set with a new one based on the cells in the entire model. Cells that are contained in the original set can be reselected if desired.

To access the cell set replace option, right-click on the desired cell set node and select **Replace** from the pop-up menu:



Four different replace options are allowed:

- cell list - Replace cells based on cell index ID's or offsets.
- threshold - Replace cells based on a scalar threshold.
- zone - Replace cells based on a zone.
- geometric range - Replaced cells based on a geometric range.

An example before and after snapshot of a replace operation using a zone is shown below:



In this example, the original cell set contained a portion of the model cells. The replace operation using a zone was used to select a new set containing some cells below the original set, as well as the right half portion of the original set.

Editing a Cell Set in the Current Scene

In order to allow multiple set operations to be performed without the dialog closing each time, a cell set can be edited using an in-place dialog as part of the current scene. The same global and Boolean operations as described above are allowed.

To edit the cell set in the current scene, right-click the cell set node and select **Edit Cell Set in Current Scene...** from the pop-up window:



The in-place *Edit Cell Set* dialog opens:



To perform Boolean operations, such as Subset, Add, Remove or Replace, select the appropriate **Boolean Mode** option. The **Filter** option determines whether a Cell List, Threshold or Zone is used. Depending on the latter option, the *Filter Options* groupbox displays the appropriate selections to complete the operation. Press **Apply** to perform the change.

To perform global operations, the Fill, Clear, Grow, Shrink and Shrink All buttons can be pressed.

Copying and Pasting Cell Sets

Boolean operations between two cell sets can be performed by copying and pasting from one cell set to another.

The following options are allowed when you paste a cell set:

- Replace cells - Replace the target cell set with cells from the copied set
- Subset cells - Form a new set of the common cells between the copied cell set and the target one

- Add cells - Add the cells from the copied set to the target set
- Remove cells - Remove common cells in the copied set from the target set

To copy a cell set, right-click on the cell set node and select **Copy**from the pop-up menu, as shown below:



The Boolean operation is performed when pasting to another cell set. Right-click on the target cell set node and select **Paste** from the pop-up menu, along with the desired operation as described above:



The behavior of the Boolean operations between cell sets is the same as previously described for individual sets. However, they provide a powerful manipulation tool to help isolate cells quickly for further analysis.

Mesh Visualization

You can visualize any mesh representation in STAR-CCM+. After you have imported or generated a mesh, you visualize it to check the cell quality and mesh validity.

You visualize the surface or volume mesh using displayers in scenes. These displayers have properties that allow you to change the mesh display characteristics. You can display either the wrapped surface, remeshed surface, or volume mesh in the scene by using mesh representations.

This section focuses on the display of surface and volume meshes only. However, you can also visualize solution data and mesh quality parameters on the mesh surfaces.

Displaying the Mesh

Displaying a mesh allows visual inspection of a surface or volume mesh for quality and validity.

You can display a mesh by using the toolbar buttons or by creating a mesh scene.

Using the Toolbar Buttons

The easiest way to activate and deactivate the mesh display is with the buttons in the **Visualization** toolbar:

- To display a mesh, click (**Show All Meshes**). Clicking this button once adds the mesh lines, and overrides the settings of all the displayers in the scene. Clicking it a second time reverts to the displayer settings.
- To hide the mesh, click (**Hide All Meshes**). Clicking this button once removes the mesh lines, and overrides the settings of all the displayers in the scene. Clicking it a second time reverts to the displayer settings.

Using a Mesh Scene

To display a mesh, you can also create a mesh scene. A mesh scene contains a mesh displayer that shows the geometry surface and mesh lines for the selected regions. By default, the mesh displayer colors the geometry surface by boundary type.

Creating a mesh scene displays the most recent mesh representation. For example, if you run the volume mesher and then create a mesh scene, the scene displays the volume mesh. To change the displayed mesh, select a new representation.

To create a mesh scene:

1. Right-click the **Scenes** node and select **New Scene** > **Mesh**.



The mesh scene appears in the *Graphics* window:



2. You can use the various view manipulations to navigate around the mesh by panning, rotating, and zooming.

By default, the mesh scene shows all geometry. You can identify and hide regions and boundaries to simplify the display.

You can change the displayer properties to alter the visualization characteristics. Some common examples are included in this section:

- Changing the Surface Opacity
- Displaying the Volume Mesh on a Plane Section

Changing the Surface Opacity

Changing the surface opacity allows you to view internal structures within a surface or volume mesh.

To change the surface opacity:

1. Make sure that you have a mesh scene open.
2. In the *Explorer Pane*, click **scene/plot**.
3. Select the **Displayer** > **Mesh** node and set the *Opacity* property.
   - A value of 1 gives a solid surface; lower values give more transparency



   - The surface and mesh lines appear transparent and reveal internal surfaces with darker shades.


Mesh Scene 1
Y
Z
X

Displaying the Volume Mesh on a Plane Section

You can visualize the volume mesh on a plane section to view the internal cell structure.

Note: As the plane section displays a two-dimensional cut through the volume mesh, some cells can appear incorrectly distorted.

To display the volume mesh on a plane:

1. Make sure that you have a mesh scene open.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.
   - You can also use a plane that you have created for other scenes.
3. In the *Create Section* in-place dialog, define the plane parameters.
4. In the *Explorer Pane*, click **scene/plot**.
5. Right-click the **Displayers** > **Mesh** > **Parts** node and select **Edit**.
6. In the object selection dialog, select the plane section.



The plane section displays the volume mesh where it cuts through the region.



You can also sweep planes to view the mesh throughout the entire volume in an animation, or create multiple sections to view the mesh on several pre-defined planes with the same orientation.

Example of Poor Surface Representation in a Volume Mesh

Plane sections are useful for viewing the internal cell structure around complex geometry to ensure that you have accurately captured the surface in the volume mesh.

The following image shows a volume mesh that does not accurately capture a complex surface.



By producing a finer mesh around the surface, you can improve the discretization. You can validate this improvement by displaying the new volume mesh on the same plane section.


Mesh Scene 1
Z
X

Comparing a Volume Mesh with the Geometry

By displaying two different mesh representations, you can compare the volume mesh with the geometry surface to determine the accuracy of the discretization.

In the following example, the original mesh scene shows the volume mesh on a plane section that cuts through the region.



To compare the volume mesh with the geometry:

1. Right-click the **Displayers** node and select **New Displayer > Geometry**.
   - Adding a displayer allows you to apply two different mesh representations.
2. Select the geometry displayer and set the properties to control the display.
   - To highlight the difference between the two mesh representations, this example sets the geometry displayer properties to:

- *Outline*: Activated
- *Line Width*: 3.0

3. Select the **Parts** node and set the *Parts* property as required.
   - To display the comparison on the same part, this example sets **Parts** to same plane section.
4. Finally, select the geometry displayer and set *Representation* to **Geometry**.
   - Now, the mesh displayer displays the volume mesh and the geometry displayer displays the original geometry surface.



This image shows that the volume mesh is a poor representation of the original geometry. By refining the surface size on the boundary, you can improve the accuracy of the discretization, but this results in a higher cell count.


Mesh Scene 1
Z
Y
X

Two Dimensional and Axisymmetric Meshes

The use of a one-cell-thick three-dimensional mesh is much less efficient than using a real two-dimensional mesh for two-dimensional and axisymmetric simulations. When circumstances allow, use real two-dimensional meshes.

There are several ways to obtain a 2D mesh in STAR-CCM+:

- Create a 2D mesh directly in STAR-CCM+.

  This option provides the most efficient and computationally least expensive method for creating a 2D mesh in STAR-CCM+. See Creating a Two-Dimensional Mesh.

- Create a 2D mesh in an external program and import it directly into STAR-CCM+.

  The Gridgen mesh generator, for example, supports 2D meshes.

- Create a 3D mesh and use the STAR-CCM+ option, **Mesh > Convert to 2D** to extract the cell faces that lie on the Z=0 plane into a 2D mesh.

  See Examples Of Two Dimensional Mesh Extraction.

  During the conversion process, the choice is given to either keep the original 3D mesh or delete it automatically and maintain only the derived 2D mesh. The choice of which region or regions to convert is also provided.

Creating a Two-Dimensional Mesh

Follow these instructions to produce a real two-dimensional mesh in STAR-CCM+ from selected two-dimensional part surfaces.

Creating a 2D mesh is more efficient and less computationally expensive than extracting a 2D mesh from a 3D mesh. 2D meshes are useful for some electric machine and aerospace industry analyses.

Before creating a 2D mesh:

1. The whole part surface on which a 2D mesh is required must be coincident with the Z=0 plane, within a tolerance of 1.0e-6. If necessary, transform the parent part so that the target part surface lies on the Z=0 plane. You can also use splitting tools to break out the target surface into its own part surface. See Manipulating Part Surfaces.

Note: For 2D axisymmetric analyses, the axis must lie on the X-axis and the mesh must also be above the local Y=0 plane. If necessary, transform the parts. See Transforming Parts to a Coordinate System.

To create a 2D mesh:

2. Badge the 2D parts that contain the Z=0 surfaces.
   See Identifying the Two-Dimensional Part Surfaces.
3. Assign parts to 2D regions. See Assigning Parts to Two-Dimensional Regions.

   Separate regions are created for 2D and 3D parts.

Note: Parts that are badged as 2D cannot belong to the same region as parts that are not badged.

4. Create an **Automated Mesh (2D)** operation.
   See Creating the Two-Dimensional Automated Mesh Operation.
5. Generate the volume mesh.

   The 2D mesh is generated only on the Z=0 surfaces that you identify using the **Badge For 2D Meshing** operation.

Identifying the Target Part Surfaces for 2D Meshing

You identify part surfaces on which a 2D mesh is required using the **Badge for 2D Meshing** operation. This operation identifies part surfaces that lie on the Z=0 plane. If any part surfaces lie partially on the Z=0 plane, the **Badge for 2D Meshing** operation fails.

To identify parts for 2D meshing:

1. Create a **Badge for 2D Meshing** operation.
   You can create any number of **Badge for 2D Meshing** operations, however, a part can only belong to one **Badge for 2D Meshing** operation.
   a) Right-click **Geometry** > **Operations** and select **New** > **Badge for 2D Meshing** .
   b) In the *Create Badge for 2D Meshing Operation* dialog, select the parts that are required to produce a 2D mesh. It is not necessary to select part surfaces, as STAR-CCM+ identifies the part surfaces that lie on the Z=0 plane automatically.
2. Execute the **Badge for 2D Meshing** operation.
   Parts and part surfaces are badged and available to select for **Automated Mesh (2D)** operations as follows:

   = Available to select as 2D region

   = Available to select as 2D boundary

Note: Deleting a **Badge for 2D Meshing** operation removes the badge from the input parts and part surfaces.

Understanding the Results of a Badge for 2D Meshing Operation

The following example shows how part surfaces on a simple three-dimensional block part are badged during a **Badge for 2D Meshing** operation.

On a simple block that is split into six part surfaces, only one part surface can wholly lie on the Z=0 plane. Only the part surface that lies on the Z=0 plane is badged as a 2D region.



The four part surfaces that are connected to the part surface on the Z=0 plane, are badged as boundaries. The edges of these part surfaces, that lie on the Z=0 plane, form the boundaries of the 2D region containing the final 2D mesh.



Any part surfaces that are disconnected from the Z=0 plane, are not badged.



Parts which contain badged part surfaces, also display a badge.

Assigning Parts to Two-Dimensional Regions

This section describes how to assign parts which are badged as two-dimensional to regions.

1. Select the **Geometry** > **Parts** that you are assigning to 2D regions.
2. Right-click one of the selected parts and select **Assign Parts to Regions**.
   Parts that are badged as 2D are assigned to different regions from parts that are not badged.
   - When parts that are not badged as 2D are assigned to regions, all part surfaces on a part are assigned to boundaries as specified in the *Assign Parts To Regions* dialog.

     See Assigning Parts To Regions
   - When parts that are badged as 2D are assigned to regions, only part surfaces that are badged as 2D boundaries, are assigned to boundaries.



Boundaries are not created for part surfaces that are disconnected from the 2D region. These disconnected boundaries are disregarded as they do not have any boundary conditions and they are not part of the final 2D mesh.

Creating the Two-Dimensional Automated Mesh Operation

The **Automated Mesh (2D)** operation generates the 2D mesh using the mesher that you choose. Only parts that are badged for 2D are available for selection as input parts for an **Automated Mesh (2D)** operation.

Before creating the **Automated Mesh (2D)** operation:

1. Identify and badge the 2D parts. See Identifying the Target Part Surfaces for 2D Meshing

To create the **Automated Mesh (2D)** operation:

2. Right-click **Geometry** > **Operations** and select **New** > **Automated Mesh (2D)**.
3. In the Create **Automated Mesh (2D) Operation** dialog:
   a) Select the parts from which to create the 2D mesh.
   b) Select a volume Mesher:
      - Polygonal Mesher
      - Quadrilateral Mesher
      - Triangular Mesher

      You can also select the optional Prism Layer Mesher.
4. Execute the **Automated Mesh (2D)** operation.

   Parts that are badged as 2D are meshed. Any part surfaces that are badged as 2D boundaries, create boundaries which appear as edges within the 2D mesh.



Note: You can copy default controls between 3D and 2D meshes.

Extracting The Two Dimensional Mesh

This section describes how to extract the two dimensional mesh from the surface of an existing three-dimensional mesh. This existing mesh can either be imported into STAR-CCM+ or generated using the three-dimensional volume meshers.

1. For the existing three-dimensional mesh, check that the boundary that you want to extract lies in the Z = 0 plane.

   You can check the alignment in a number of ways:

   - To ensure that the Z extent includes a zero, review the extents of the mesh by:
     - Looking at the results of the import process in the *Output* window
     - Running a mesh diagnostic report and looking at the report in the *Output* window

       See Running Mesh Diagnostics.
   - Create a new geometry scene with the default laboratory coordinate system included as part of the view. You can then verify visually the location of the Z = 0 plane. See Displaying the Geometry.

   The example below shows the default laboratory cartesian coordinate system included in the model display. As can be seen, one of the boundaries is aligned with the Z = 0 plane:



If the model is not oriented correctly, rotate and/or translate the model until the face of interest lies on the Z = 0 plane. If a boundary lies near the Z = 0 plane but not consistently on it, you can project the boundary vertices to that plane.

2. Once you are satisfied with the positioning of the boundary that you wish to extract, convert the mesh by selecting **Mesh > Convert To 2D...**.



The *Convert Regions to 2D* dialog opens:



3. Select the regions to convert. Turning the **Delete 3D regions after conversion** option on would automatically delete the original regions that are selected for conversion. The default option is not to delete the regions.

   The **Tolerance** value determines the allowable deviation from the Z=0 plane of the mesh boundary that you wish to convert. For example, if the selected 3D regions have boundaries that are not exactly on the Z=0 plane, then the **Tolerance** value can be increased so that a complete boundary plane is found.

4. Pressing **OK** starts the conversion process.

   Once the conversion is complete, the results are reported in the *Output* window where the names of the new regions and boundaries are listed along with the number of cells, faces, and vertices.



The simulation tree is also updated with the new region specification. Each converted region name is appended to include the 2D specifier, as shown below:



If the option to delete the original 3D regions was activated, then only the new 2D regions remain:



If STAR-CCM+ cannot find any faces in the z = 0 plane, then a message in the *Output* window reports that no suitable boundaries could be found:



Notes:

- If no boundary faces are found in the Z = 0 plane, no action is taken.
- Every distinct boundary with faces on the Z = 0 plane becomes a new, separate region.
- After extracting the two-dimensional mesh, you will need to reconstruct all Part Displayers.
- The extracted two-dimensional mesh is treated as being 1 meter in depth by the solver (1 radian for axisymmetric cases), even if you selected different units for length (such as feet or inches). Compensate manually for different depths.
- If only a partial two-dimensional mesh is created, you can delete the regions and try again with a larger tolerance (assuming the original 3D regions were not deleted).
- If the original 3D region has an associated physics continuum, then a new physics continuum is created for the 2D converted region. Further, if the 3D space model is active for the original physics continuum then the space model is automatically reset to be 2D in the new physics continuum for the two-dimensional region. All other models that are currently selected for the continuum are also carried over to the new continuum.

See Also

Projecting Boundaries to a Plane
Translating Regions
Rotating Regions

Examples of Two-Dimensional Mesh Extraction

The two-dimensional extraction feature is useful when the mesh generator does not support true two-dimensional mesh creation and you have:

- Two-dimensional meshes that are extruded into the third dimension to have one cell thickness. Use 2D mesh extraction to run such cases as true two-dimensional simulations in STAR-CCM+
- Axisymmetric meshes that are created as annular segments. Use 2D mesh extraction to run such cases as 2D axisymmetric simulations in STAR-CCM+
- Models with symmetry planes that are representative of the whole model. Use 2D mesh extraction to run simplified 2D simulations in STAR-CCM+to investigate parameters such as model selection and boundary conditions.

One Cell Thick 2D Extruded Meshes

The three-dimensional mesh that is shown below has been created to simulate flow over a cylinder. It is composed of one layer of hexahedral cells using symmetry boundaries in the spanwise direction.



The two-dimensional mesh that is shown below is the result of the `Mesh > Convert To 2D` command.



Axisymmetric Annular Segment

The image below shows a one-cell-thick hexahedral mesh with wedge-shaped cells next to the axis for simulating flow through an axisymmetric diffuser. The original model is three-dimensional with cyclic boundaries.



The axisymmetric version of the mesh is extracted from the surface mesh on one of the cyclic boundaries, and is composed of quadrilateral cells.



Symmetry Plane

The three-dimensional hexahedral mesh below simulates the flow over a wing that is attached to a symmetry plane.



The two-dimensional version that is shown below, composed of quadrilateral cells, has been extracted from the surface mesh on the symmetry plane to run a simplified simulation.



Scaling the Entire Mesh

Scaling the mesh is important as the units of the model must be in meters before running an analysis.

If the surface or volume mesh was imported using different units, then the process below must be followed.

To scale the entire mesh inside STAR-CCM+:

1. Use the **Mesh** > **Scale Mesh...** menu option:



   The *Scale Mesh* dialog opens with all regions selected by default.

2. Enter the scale factor into the **Scale Factor** text box. This value is used to multiply each dimension to scale the model.
   In the example below, the scale factor would convert a model that had been constructed in inches to meters.



3. Click **Apply**.

   The progress bar at the bottom of the simulation tree view indicates the progress:



Once the operation is complete, the summary information is written to the *Output* window. This output is useful to check that the scale factor has been applied to the appropriate regions:



All active scenes are updated to reflect the new mesh configuration.

Working with Baffles

In STAR-CCM+, there are two methods to accommodate 2D surfaces that are intended for use as thin-walled baffles. The first method is to create interfaces that are based on imported surfaces and so have a true zero-thickness baffle. The second method is to inflate the imported surface to create a baffle of finite physical thickness. Both of these methods are discussed in this section.

Generating Baffles with Zero Thickness

When correctly defined, zero thickness baffles are represented as an interface within the **Interfaces** manager node.

To define a zero thickness baffle:

1. Unless you intend to create a floating baffle, ensure that the baffle is properly connected to the surrounding geometry.
2. Create a separate boundary for the baffle and assign the baffle surface to it.
3. Convert the baffle boundary to an interface of the relevant type.
4. To distinguish which side of the baffle belongs to the original baffle boundary, use the offset surface mode of the *Surface Preparation Tool*.

   A section detailing the process for updating baffle geometry is also provided.

For more information, see the following sections.

Ensuring the baffle is a Connected Surface

The merge/imprint mode of the Surface Preparation tool can be used to connect a baffle surface into the surfaces surrounding it. See Imprinting Baffle Edges onto Faces .

In the example that is shown below, a two-dimensional surface is located within the duct, but is not connected to the base surface of the duct.



In the merge/imprint options, the edges along the lower side of the baffle are added to the *Source*, and the faces forming the base of the duct are added to the *Destination*. On executing the **Imprint** operation, the lower side of the baffle is stitched into the base of the duct.



See Also

Merge And Imprint Mode

Converting a Baffle Boundary to an Interface

Represent a zero-thickness baffle by an interface.

The surface that is used to create a baffle must be assigned to its own boundary within the region to which it belongs.



To convert this boundary to an interface:

1. Right-click the boundary representing the baffle and select **Convert to Interface(s)**.



Note that a mesh continuum must be present in the simulation, and the meshing pipeline must be initialized, before this option is available on a boundary. For more information, see Initializing the Mesh.

The result of the operation is the creation of an additional three boundaries, two of which form the opposite sides of the in-place interface. The original boundary, and the newly created copy, are used to govern any surface area that does not form part of the interface (and are therefore redundant in most cases).



In the example above, the boundaries **Baffle [Baffle]** and **Copy of Baffle [Baffle]** refer to the two sides of the interface. If necessary, different mesh or physical properties can be set for each side of the interface, for example, a prism layer on one side of the interface only. Finding out which boundary applies to which side of the baffle is described in the next step.

Distinguishing Between the Two Sides of a Baffle

The general rule for distinguishing between the boundaries on either side of a baffle is that the interface boundary with the same name as the original boundary has positive face normals.

This positive face normal is colored green when the baffle surfaces are viewed within the offset faces mode of the *Surface Preparation tool*.

To access this view:

1. Activate the *Surface Preparation tool* from the appropriate surface representation of the model.

   See Launching the Surface Preparation Tool.
2. Switch to the offset faces mode within the Surface Preparation tool.

   See Offset Faces/Edges Mode.
3. Activate the **Color Face Normals** checkbox.

In the example that is shown here, the **Baffle [Baffle]** boundary belongs to the side colored green, and the **Copy of Baffle [Baffle]** boundary belongs to the side colored red.



Updating Baffle Geometry

Changes that are made to the baffle geometry surface are automatically transferred to the fluid boundary, once you have connected the baffle to a surface, and converted it to an interface.

The baffle boundary must also be associated with a part surface, as illustrated in the image below. For more information, see Assigning Surfaces to Boundaries Using the Region Boundary Nodes.



The following example uses 3D-CAD to reduce the inner diameter of an orifice baffle within a pipe. If you are using an imported baffle, importing a new baffle surface as a geometry part has the same effect.



When a geometry part is changed in 3D-CAD, it must be updated before the changes are made to the geometry part.



Note that after updating or importing a part, any modifications to the surface using the surface preparation tools are over-written. Therefore, if you used the surface repair tools to connect the baffle to a surface, repeat the operation after updating the part.

1. Click **Initialize Mesh** to initialize the fluid region, which transfers the modified baffle surface to the fluid boundary.

   The updated baffle is then shown in the scene.



See Also

Ensuring the baffle is a Connected Surface
Converting a Baffle Boundary to an Interface
Introducing 3D-CAD Modeling

Generating Baffles with Finite Thickness

If the surface wrapper is used on an imported surface containing 2D baffle faces, and the procedure for generating zero-thickness baffles is not followed, then the surface wrapper automatically inflates the 2D faces. The inflation thickness is 0.25 times the smallest refinement size that is used in the surface wrapping process (not necessarily equal to the minimum surface size specified in the **Reference Values** node).

If the surface wrapper is not used, or you require more control over the thickness of the baffles, inflate the two-dimensional baffle surface using the Offset Faces mode of the *Surface Preparation tool*. For more information, see Offset Faces/Edges Mode.

In the case where the surface wrapper has not been used, it will be necessary to perform some manual cleanup after inflating the baffle to ensure that the resulting surfaces are properly connected. It may also be necessary to duplicate some faces so that a separate solid region can be created for the baffle if it is expected to participate in a conjugate heat transfer simulation.

Multi-Region Meshing

This section outlines the surface requirements for performing multi-region meshing, as well as step-by-step instructions when starting from surface imports that result in single region and multiple regions definitions. Additionally, guidelines are provided on what the best approaches are for exporting surface data from CAD packages when multiple regions are required.

Multi-Region Surface Requirements

For STAR-CCM+ to remesh surfaces successfully, and/or generate a volume mesh for a multi-region case with conformal mesh between regions, prepare the surface to adhere to the following conditions:

- each region must be closed, manifold, and non-intersecting
- boundaries that exist between regions must be grouped and flagged as interfaces

The above requirements are illustrated here:



The surface remesher can automatically produce a conformal mesh across an in-place interface containing mismatched boundaries. To do this, STAR-CCM+ imprints one boundary on the other so that a mapped boundary surface is produced, resulting in a conformal triangulation. For a case such as that above, use a single mesh continuum to mesh all regions in order to ensure that a conformal surface and volume mesh is produced at interfaces.

In instances when a conformal volume mesh is not required between each region, then the requirement to have interfaces and imprints defined is relaxed. In this case, each region can be meshed independently by using a different mesh continuum or by turning on per-region meshing. The meshes can later be joined back together either by using an in-place interface or by fusing the boundaries. In some respects, this approach is far easier than ensuring that the region interfaces and imprints are properly defined. However, the process

results in increased overhead at the simulation stage, due to the extra interface information that is stored and the extra face divisions in the boundary fusing process.

When dealing with multi-region topologies, there are many factors in the definition of the CAD data that affect the requirements for STAR-CCM+. These factors include:

- the format that is used to export the CAD data, for example, STL, IGES, VDA
- whether each part in the assembly is exported to the same file or to separate files
- whether duplicate surfaces exist at inter-region boundaries
- whether inter-region surfaces have been imprinted or not
- whether the assembly contains all relevant fluid/solid parts or only the solid material portion of the geometry.

For example, certain CAD packages such as SolidWorks allow you to save an assembly as a single STL file or each part as a separate STL file. Similarly, in most CAD packages, you can elect to open parts individually and save them as separate IGES files or other format. Alternatively, the entire assembly can be saved to a single file, which is usually the default (or only option allowed). Depending on the method that is used to save the import file or files, inter-region boundaries can have either single or duplicate surfaces, which do not necessarily have identical perimeter triangulation. Additionally, some CAD packages allow imprinting (whereby one part can be “stamped” onto another, leaving an edge definition that represents the common area) while others do not have this feature. Finally, the assembly typically only represents the solid material, rather than the fluid volume that is required for CFD simulations. Some CAD packages easily allow the fluid portion to be extracted and combined back into the assembly while others do not. In summary, many different factors influence the level of effort that is required to prepare the import geometry for STAR-CCM+.

In general, regardless of the factors that affect the multi-region geometry definition, STAR-CCM+ contains the required tools to prepare the surface for meshing. In particular, two operations specific to multi-region meshing are essential for manipulating improperly prepared surfaces in order to conform to the above requirements. They are:

1. create a (duplicate) interface boundary
2. create a new region from a collection of boundaries.

Being able to create a duplicate boundary at an inter-region interface allows two regions that were previously “joined” to be separated out,

allowing a new region to be created from individual boundaries. The surface boundary that was duplicated maintains an in-place interface definition as part of its make up so no further action is required once the region has been redefined.

The use of the in-place interface option is important for defining the surface correctly in cases where regions are defined at import but no connectivity information between regions is included. By including an in-place interface, conformal meshes at interface boundaries can be created for the surface and volume meshers.

Two different scenarios are discussed here showing how the initial CAD surface data can be manipulated in STAR-CCM+ so that a conformal volume mesh can be constructed. The two different scenarios are:

- the CAD import results in a single region
- the CAD import results in multiple regions.

See Also

Surface Remesher
In-Place Interface
Using Per-Region Meshing
Fusing Boundaries
Converting Boundaries to Interfaces
Creating Regions From Boundaries

Starting from a Single Region

The discussion here focuses on multi-region cases that start off as a single region as a result of the import process into STAR-CCM+.

For example, the following operations all result in the creation of a single region definition during import:

- Selecting the **One region for all files** option for **Region Mode** when importing multiple files at the same time
- Importing an assembly that has been saved to a single file.

Depending on the surface export options in the CAD package and the import options in STAR-CCM+, surfaces at inter-region boundaries are sometimes duplicated. This leads to two different scenarios:

- Inter-region boundaries are common/shared (and therefore have same perimeter triangulation)
- Inter-region boundaries are duplicated and can have dissimilar areas.

Common Inter-Region Boundaries

In cases with common inter-region boundaries, the split by surface topology option can be used to split regions into their constituent parts and generate the appropriate boundary interfaces to ensure a conformal mesh.

This option effectively performs steps 1 through 3 as described below and also allows regions that originally belong to the same boundary to be combined automatically, reducing the overall number. For more information, see Splitting by Surface Topology.

The above option is the recommended route for all such cases, from simple to complex, as no manual interaction is required. However, if you prefer to define the individual regions by hand, then follow the guidelines below to prepare the surface further:

1. Combine boundaries that belong to the same inter-region interface together into one boundary. Repeat as necessary for all inter-region boundaries.

   See Combining Boundaries.



2. Use the **Convert to Interface(s)** option to create a duplicate boundary for all the inter-region boundaries.

   See Converting to an Interface.



   A pair of in-place internal interfaces are generated at the same time linking the boundary pair.

3. For each new region in turn, select the boundaries that make up the region and select the **Create New Region from Boundaries** option.

   A new region is created. One half of the in-place internal interface is also transferred into the new region. Alternatively, if several new regions are required, the split non-contiguous region option can be used instead to perform the operation in one go. For more information, see Splitting Non-Contiguous Regions.



4. Continue with the meshing process as normal.

See Also

Creating Regions From Boundaries

Inter-Region Boundaries Are Duplicated

In cases where the inter-region boundaries are not common (that is, duplicated) and the perimeter triangulation and boundary areas may not be the same, follow the guidelines below:

1. Combine boundaries that belong to the same inter-region interface together into one boundary. Repeat as necessary for all inter-region boundaries.



2. Separate out each region either by using the **Create New Region from Boundaries** or the **Split Non-Contiguous Region** options, such that each region is composed of unique boundaries.



3. Create an in-place interface between each boundary pair (**Boundary 1** and **Boundary 6** in the above example) and then proceed with the meshing process as normal.

   In this instance, the surface remesher imprint process is used to produce the same triangulation for inter-region boundaries before proceeding with producing new triangulation for each region. A conformal mesh is produced at the inter-region boundary interface. The original import or initial surface representation is left unaffected by this process. For additional information on using the surface remesher with interfaces, see Surface Remesher with Interfaces.

See Also

Combining Boundaries
Creating Regions From Boundaries
Splitting Non Contiguous Regions
Surface Remesher

Starting From Multiple Regions

The discussion here focuses on multi-region cases that have individual regions that are defined during the import process in STAR-CCM+.

For example, the following operations all result in the creation of multi-region definitions during import:

- Selecting the **One region per file** option for **Region Mode** when importing multiple files at the same time
- Importing each file one at a time.

Surfaces at inter-region boundaries can have dissimilar triangulation, depending on how they were prepared. This leads to two different scenarios existing:

- Inter-region boundaries have the same triangulation or the same perimeter triangulation
- Inter-region boundaries do not have the same perimeter triangulation or do not have the same areas.

Regardless of which of the above scenarios you have, the same procedure is used to prepare the surface ready for meshing. Follow the guidelines below:

1. Combine boundaries that belong to the same inter-region interface together into one boundary. Repeat as necessary for all inter-region boundaries.



2. To create an interface between each inter-region boundary pair, use the in-place interface option.

3. Continue with the volume meshing process as normal.

   In this instance, the surface remesher imprint process is used to produce the same triangulation for inter-region boundaries before proceeding with producing new triangulation for each region. A conformal mesh is produced at the inter-region boundary interface. The original import or initial surface representation is left unaffected by this process.

See Also

Combining Boundaries
Surface Remesher
Surface Remesher with Interfaces

CAD Data Export Guidelines for Assemblies

The text below provides guidelines for the best approach for how to export multi-region cases from CAD packages and what issues are typically avoided or encountered when they are imported into STAR-CCM+.

Guidelines are provided for the following formats:

STL

Save each part in the assembly as a separate STL file but use the same edge length deviation and curvature angle settings for each. Import the separate STL files into STAR-CCM+ using the **One region per file** option which results in separate regions automatically being defined. Split the single boundary definition for each region up by angle and add feature curves as normal to define the edges. The perimeter triangulation at the inter-region boundaries will most likely be different (dependent on the CAD package) but a single mesh continuum with interfaces can still be used, resulting in a conformal mesh at the inter-region interfaces. If the surface triangulation is such that a single mesh continuum cannot correctly edge zip and swap out inter-region boundaries, then a different mesh continuum must be used for each region, resulting in a non-conformal volume mesh interface.

If you save the assembly to a single STL file, the import process usually results in a single region with a single boundary, regardless of the option that is selected for **Boundary Mode.** Few CAD programs support the multiple solid definition that STL allows. Split the single boundary up using the non-contiguous option and then split the regions up using the non-contiguous option. Add features curves as normal and continue with the meshing process for multiple regions.

See Also

Starting From Multiple Regions

NASTRAN/PATRAN

Where the CAD program can export tessellated surfaces as Nastran or Patran files, save the assembly to a single file if the inter-region boundaries are common, or have the same perimeter triangulation.

When this file is imported into STAR-CCM+, the boundaries can then be separated out into interface and non-interface groups and duplicate copies can be made of the boundaries at the inter-region interface. New regions can then be created and a single mesh continuum can then be used to generate the volume mesh for the entire geometry resulting in a conformal mesh at the boundary interface.

If the assembly contains duplicate inter-region boundaries with different triangulation on each, then save each part to a separate file if possible. When you import these files into STAR-CCM+, you are presented with the choice of whether to create one region per file (recommended), or one region for all files. Regardless of the choice that you make, individual regions often have dissimilar triangulation at the inter-region boundaries. Usually STAR-CCM+ can resolve different inter-region boundary triangulation by using the edge zipper to match the perimeter nodes and then deleting one of the boundaries to allow the creation of a conformal interface. However, depending on the variation in perimeter node positioning, the edge zipping can fail, in which case, use a different mesh continuum for each region, which results in a non-conformal volume mesh interface.

VDA

If using VDA to export from the CAD package, save the assembly to a single file and then use the pro-surf surfacing preparation tool to read in the geometry and triangulate the surface. At inter-region boundaries, pro-surf automatically detects surfaces that are within close proximity and merges them, resulting in a common boundary definition at the interface. It also provides the option for generating feature curves when exporting the triangulated surface to the pro-STAR database file. When this database is imported into STAR-CCM+, the boundaries can then be separated out to interface and non-interface groups and duplicate copies can be made of the boundaries at the inter-region interface. New regions can then be created which have the same triangulation for the boundaries at the interface. A single mesh continuum can be used to volume mesh the entire geometry, resulting in a conformal mesh at the inter-region boundary interface.

If you save the individual parts in the assembly as separate VDA files, then use a separate pro-surf session to triangulate each component and save each to a different database file (or to a different ID within the same file). When imported into STAR-CCM+, you are given the choice of whether to create one region per file (recommended) or one region for all files. Regardless of the choice that you make, individual regions often have dissimilar triangulation at the inter-region boundaries. Usually STAR-CCM+ can resolve different inter-region boundary

triangulation by using the edge zipper to match the perimeter nodes and then deleting one of the boundaries to allow the creation of a conformal interface. However, depending on the variation in perimeter node positioning, the edge zipping can fail, in which case, use a different mesh continuum for each region, which results in a non-conformal volume mesh interface. For this reason, it is recommended that you use the first approach that is discussed above to simplify the meshing process and guarantee a successful conformal interface.

IGES/STEP

If using IGES or STEP to export from the CAD package, save the assembly to a single file. STAR-CCM+ can then be used to read this file and translate the surface. At inter-region boundaries, STAR-CCM+ automatically detects surfaces that are within close proximity and attempts to merge them, resulting in a common boundary at the interface (note that this process is not always successful). It also provides the option for generating feature curves. The boundaries can then be separated out to interface and non-interface groups and duplicate copies are made of the boundaries at the inter-region interface. New regions can then be created which have the same triangulation for the boundaries at the interface. A single mesh continuum can be used to volume mesh the entire geometry, resulting in a conformal mesh at the inter-region boundary interface.

If you save the individual parts in the assembly as separate IGES or STEP files, then on importing them into STAR-CCM+ (either at the same time or one after the other), you are given the choice of whether to create one region per body, one region per file (recommended) or one region for all files. Regardless of the choice that you make, individual regions often have dissimilar triangulation at the inter-region boundaries. Usually STAR-CCM+ can resolve different inter-region boundary triangulation by using the edge zipper to match the perimeter nodes and then deleting one of the boundaries to allow the creation of a conformal interface. However, depending on the variation in perimeter node positioning, the edge zipping can fail, in which case, use a different mesh continuum for each region, which results in a non-conformal volume mesh interface. For this reason, it is recommended that you use the first approach that is discussed above to simplify the meshing process and guarantee a successful conformal interface.

It is recommended that you export parts as Parasolid Transmit files whenever possible, as this format generally results in a cleaner import into STAR-CCM+.

See Also

Parasolid Transmit

Parasolid Transmit

If using Parasolid Transmit to export from the CAD package, save the assembly to a single file. This file can then be imported into STAR-CCM+ directly, maintaining the names of the parts in the assembly. Refer to the importing IGES/STEP section for more information on working with multi-region geometries.

It is recommended you use Parasolid Transmit files where possible in preference to IGES and STEP data, as it results in a cleaner translation process on import to STAR-CCM+.

See Also

IGES/STEP

Modeling Physics

This part of the documentation describes the models in STAR-CCM+ and how to use them. Separate chapters deal with:

- Space, time, and motion
- Materials
- Flow and energy
- Species
- Turbulence and transition
- Radiation
- Aeroacoustics
- Combustion
- Multiphase flow
- Dynamic fluid body interaction
- Harmonic balance
- Solid stress
- Electromagnetism
- Electrochemistry
- Casting

These chapters describe how to use the models and their associated solvers, describe the formulation of each model and provide troubleshooting pointers. Additional chapters describe:

- How cell quality can be remedied
- Generally how transport equations are solved

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Using tools in STAR-CCM+ to facilitate your work
- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases

- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format
- A guide to using STAR-View+, a viewer for three-dimensional STAR-CCM+ scenes

Setting Up Physics

This section provides information on how to set up physics models in STAR-CCM+. First, to represent the different materials present in your simulation, create physics continua. Then select the physics models for each physics continuum, and assign the physics continua to the appropriate regions in the CFD model. Afterward, proceed to set appropriate conditions and values for the boundaries in your model, and also implement changes to region or interface types.

Using Physics Continua

A physics continuum node, which has its own properties and pop-up menu, contains all of the model nodes in the object tree.



Generally, a physics continuum is defined as *continuous* and *nonspatial* such that no part of it is distinct or distinguishable from adjacent parts. This definition is also reflected in STAR-CCM+, which logically interprets a continuum as a collection of models that represents the substance (fluid or solid) being simulated.

In keeping with the nonspatial characteristic, you can define physics continua in STAR-CCM+ independently of the mesh and independently of existing regions. The models that are activated within a physics continuum determine the following:

- Whether the continuum is three-dimensional
- Whether it comprises a solid, liquid, or gas
- Whether it is stationary in time or transient
- The quantities that are defined within it

Certain models, when activated in a physics continuum, require the activation of other models in that continuum. For instance, once a physics continuum contains a liquid or a gas, it needs a flow model. Once it has a flow model, it needs a viscous model (inviscid, laminar, or turbulent). Once turbulence is activated within a fluid continuum, a turbulence model must be selected.

A simulation can contain many physics continua, each with its own name and properties. Each region must be associated with a physics continuum, and conversely each physics continuum can be associated with one or more regions.

In addition to the models comprising the physics continuum, the continuum also contains:

- Initial conditions
- Reference values

Copying Continua with Tables

STAR-CCM+ excludes an excessively large table from the process of copying a physics continuum; the rest of the continuum does get copied, however.

To copy a continuum with a table:

- Attempt to copy the continuum and paste it in its destination.
- If the table is too large, STAR-CCM+ generates a warning in the *Output* window.



- Export the table to a `.tbl` file using the right-click menu action.
- Import the table into the destination continuum.

Physics Continuum Properties

Interfaces	The interfaces that are associated with the continuum (read-only).
Regions	The regions that are associated with the continuum (read-only).

Physics Continuum Pop-Up Menu

Select Models...	Allows selection of the models to be activated in this continuum.

Automatically Creating a Continuum During Mesh Import

When a new simulation is created, it does not contain any continua, and the Continua node is empty. As soon as a volume mesh is imported, a physics continuum is automatically created. By default, this continuum is given the name `Physics 1`, but it can be easily renamed if necessary.

You can also manually create a physics continuum before a mesh is imported.

Manually Creating a Physics Continuum

A new continuum can be created by right-clicking the Continua node and selecting New > Physics Continuum.



The new continuum is given the next logical sequential name such as `Physics 2` but can be renamed if necessary. Manager nodes for Models, Reference Values, and Initial Conditions are added as child nodes to the new continuum node. Once the new physics continuum has been created, you can proceed to select the physics models for it.

Assigning Physics Continua to Regions

Generally, regions are automatically assigned to the default physics continuum (that is, the first physics continuum node that is found in the first node in the Continua manager node) on mesh import. They can be reassigned in one of two ways:

- Setting the continuum for the region using the region properties
- Dragging and dropping a continuum onto the region

Setting the Region Physics Continuum

When a specific region node is selected, the physics continuum to which the region is assigned is displayed in the *Properties* window for that region. If more than one physics continuum exists in the simulation, it is possible to select it in the Physics Continuum property.



Dropping a Physics Continuum Onto a Region

To assign a physics continuum to a region, use the drag-and-drop feature in STAR-CCM+. In this process, drag the continuum node onto the region node; doing the opposite does not work.

Selecting Physics Models

A physics model in STAR-CCM+ defines how a physical phenomenon in a continuum is represented. Essentially, physics models define the primary variables of the simulation (such as pressure, temperature, and velocity) and what mathematical formulation is used to generate the solution. An appropriate combination of models is necessary for the complete definition of a physics continuum.

Models in STAR-CCM+ have varying levels of complexity and functionality, but their major purpose is to work with solvers to obtain a solution and to help present the information to you. The tasks of a typical model include making relevant field functions available and placing initial conditions and reference values in its continuum.

In situations where transport equations are required, the model provides these equations to the solvers. Therefore model properties often include coefficients relevant to the transport equations, as well as parameters relevant to the discretization approach.

In any continuum, models are selected and changed using the Physics Model Selection dialog, which assists you in recommending models and in selecting the appropriate combination of models.

Using the Physics Model Selection Dialog

The *Physics Model Selection* dialog is used to select physics models and to change the model selection in a simulation. To activate the dialog in a physics continuum (say Continuum 1), open the Continua node and right-click on the Physics 1 node. In the pop-up menu, select the item Select models...



This same menu item is available in the pop-up menu of the Models manager node under the continuum node.

Model selection and deletion can only be done through this dialog:

- This dialog is similar to a wizard in that it can assist you through a fairly complex process selection network.

- It is unlike a wizard in that it does not have the traditional Next and Back buttons.
- It does not restrict you to answering questions in a particular order.
- Models not available with STAR-CCM+ Lite are grayed out.

In this section and the following ones, the descriptions and illustrations of model selection refer to physics models; however, the general principles apply to meshing models as well. The following screenshot shows the dialog part way through a selection process, where each significant aspect of the dialog can be seen.



Selecting Models in the Dialog

The *Physics Model Selection* dialog first appears as it looks below for a physics continuum that has not yet been defined.



The *Enabled Models* group box on the right-hand side of the dialog is empty, indicating that no models have yet been selected. Two models are available for selection, the Space model and the Material model. The <Optional> labels indicate that the selection of these items is not mandatory. The Auto-select recommended models checkbox provides an option for the automatic selection of recommended models.

The *Model Selection* dialog changes dynamically in response to previous selections. For instance if the Gas and Coupled Flow model options are selected, it appears as below. Now, you are prompted to select models for motion, space, Equation of State and other categories. The <Additional

selections are required> label tells you that the model selection is incomplete. An incomplete model selection means that the continuum cannot be initialized in its current state.



Trying to iterate when you have not completed the model selection results in an error similar to the one shown below.



The most effective way to make model selections is to keep making choices from the model family that appears at the top of the available models side of the dialog. This technique adds models to the *Enabled Models* group box in the same order that they were selected (by you or by STAR-CCM+when in auto select mode).

The dialog is “live” in that the server is updated each time that you change the state of a checkbox or radio button. If you have the relevant nodes open, the object tree reflects these updates.

Model Selection in STAR-CCM+ Lite

If you are running STAR-CCM+ Lite, some of the models are grayed out as shown below. These models are available for selection only in the full version of STAR-CCM+.



STAR-CCM+ Lite lets the *Physics Model Selection* dialog display any selected models, even ones that were activated in the full version of STAR-CCM+. However, STAR-CCM+ Lite cannot run the simulation. An *Information* dialog such as the one shown below appears when you attempt to run the case.



Using Auto-Select Recommended Models

The Auto-select recommended models checkbox is ticked by default, which indicates that when there is a choice of models, the recommended model is automatically selected. As an example of how this feature works, consider the continuum setup below, which indicates that a turbulent flow continuum is complete except for the choice of turbulence.



If you select K-Epsilon Turbulence, the *Physics Model Selection* dialog immediately appears as below, where two more recommended models have been added:

- Two-Layer All y+ Wall Treatment
- Realizable Two-Layer K-Epsilon

The reason that these two models were all added after a single selection is that they are the recommended submodels that were automatically selected.



If the Auto-select recommended models checkbox is deactivated before K-Epsilon Turbulence is selected, the *Model Selection* dialog appears as below. You now have to select among the choices of K-Epsilon Turbulence Models, and following that have to select a wall treatment model if necessary.



Therefore, if you want to select a model other than one that is recommended (say an alternative wall treatment model) there are two options:

- Use Auto-select recommended models, then change model selection to the required one.
- Deactivate the Auto-select recommended models checkbox and then make all the selections as required.

Viewing Selected Models

The *Physics Model Selection* dialog can be used to view the models that have been activated in a continuum.

As shown below, the *Optional Models* group box on the left-hand side of the dialog lists the models that have not been selected, but are not necessary to complete the continuum. These models would only be selected if you determined that the physics of the situation warranted them.

The *Enabled Models* group box on the right-hand side of the dialog shows all the models that have been activated. The selections that are grayed out cannot be deactivated, since other activated models (not grayed out) require them. Therefore, deactivate the models that are not grayed out first.



Changing Model Selection

The *Physics Model Selection* dialog can be used to change the selection of models in a continuum.

As shown below, the *Enabled Models* group box on the right-hand side of the dialog shows all the models that have been activated. The selections that are grayed out cannot be deactivated, since other activated models (not grayed out) require them. Therefore, deactivate the models that are not grayed out first. Once models have been deselected, they can be reselected as required.



For best results, deselect models by deactivating the checkbox of the model at the top of the *Enabled Models* group box. Continue until you "undo" the chain back to the point where you want to choose a different model.

Although you can deselect models in a different order, you could become uncertain about model dependencies, and about how to choose another set of models. If you encounter this situation, deselect everything and start again.

Viewing the Models Using the Object Tree

The models that have been activated can be viewed using the Model Selection dialog, which is the preferred method for changing the model selection. However, for setting model properties, or for simply viewing a list of active models, you can use the object tree. To view the models in a

physics continuum (say Physics 1), simply open the Physics 1 node within the Continua node, then open the Models node. The tree view expands as shown below.



Setting Model Properties

To set the properties of a model, first view the models in the continuum using the object tree, then click the selected model. This action causes the *Properties* and *Expert* properties to become available for inspection and editing.

Unless you are thoroughly familiar with the theoretical aspects of the model and the discretization techniques that are used in STAR-CCM+, do not change any property settings within the *Expert* category. The values in that

category reflect both the design of the model and the discretization approaches that have been optimized for accuracy and performance. Tampering with them can diminish the effectiveness of the model.



Setting Types

The simulation of continua requires a bridge between what is understood in the material world and the mathematical constructs of transport equation models. In the physical world, we understand concepts such as "walls" or "porous media". In the modeling world, we think in terms of "boundary conditions" or "volumetric source terms". These concepts are bridged in STAR-CCM+ by formally defining types, conditions, and values in such a way that there is no ambiguity. These three entities are used to provide to the continuum models complete information for dealing with boundaries, regions, and interfaces.

This section includes information about:

- Setting types
- Specific region types
- Specific boundary types

Separate sections cover the following topics in detail:

- Conditions and values
- Specific interface types

What Are Types?

In STAR-CCM+, types are assigned to regions, boundaries, and interfaces as a way of informing the models of the associated continua how these entities must be dealt with. As much as possible, the types are analogous to physical entities so that they are as unambiguous as possible. For instance, a typical boundary type is a wall, which is a concept that most people can understand and identify with. All models in STAR-CCM+ must know how to deal with the predefined types.

Setting Types

STAR-CCM+ initially assigns types during mesh import (based on information that is contained in the mesh being imported) or when a region or boundary is created from within the simulation. To change the type of a

region, boundary or interface, select the region, boundary, or interface in the object tree. In the *Properties* window, select one of the available types from the drop-down list corresponding to the Type property.



The models that are available in the continuum determine the list of available types. For instance, the Free Stream boundary type is only valid for ideal gas flow. So if the continuum does not contain the Ideal Gas Model, the Free Stream boundary type is not presented.

Region Types Reference

Regions are assigned types that models within their associated continuum use to determine what specific conditions and values are required.

STAR-CCM+ recognizes three different region types: fluid, porous, and solid. The availability of each type depends on the continuum that is associated with the region. Fluid and porous regions are available only if the continuum is a gas or liquid. Solid regions are available only if the continuum is a solid.



STAR-CCM+ initially assigns regions types upon mesh import (based on information that is contained in the mesh being imported). Region types can also be reassigned within STAR-CCM+.

Fluid

Regions that are assigned to gas or liquid continua are automatically assigned the fluid type. The physical models, such as flow, energy, and turbulence, make the appropriate modeling assumptions governing flow in these regions.

Application

- General fluid flow in the absence of porous media

Compatibility

- Available only for regions containing a fluid material

Porous

Regions that are assigned to gas or liquid continua can be designated with the porous type. The physical models, such as flow, energy, and turbulence, make the appropriate modeling assumptions governing porous media flow in these regions.

Applications

- Filters
- Tube bundles
- Flow straighteners

- Sponges
- Radiators

Compatibility

- Available only for regions containing a fluid material

Solid

Regions that are assigned to solid continua are automatically assigned the solid type. The physical models, such as energy and stress, make the appropriate modeling assumptions governing solids in these regions.

Applications

- Solid regions in which conductive heat transfer is to be modeled
- Solid regions in which stresses are to be modeled

Compatibility

- Available only for regions containing a solid material

Boundary Types Reference

Boundaries are assigned types that models within their associated continuum use to determine what conditions and values are required.

STAR-CCM+ recognizes nine different user selectable boundary types:

- Axis
- Flow split outlet
- Free-stream
- Mass flow inlet
- Pressure outlet
- Stagnation inlet
- Symmetry plane
- Velocity inlet
- Wall

The availability of these boundary types is dependent on the specific models that have been activated within the continuum that is associated with the boundary.

STAR-CCM+ initially assigns boundary types upon mesh import (based on information that is contained in the mesh being imported). These types can also be reassigned within STAR-CCM+.

Boundaries added through the creation of interfaces are assigned interface boundary types and names that you do not change.

Axis

The axis boundary represents the axis of a two-dimensional axisymmetric region.

Application

- The centerline of a pipe or other axisymmetric duct

Compatibility

- Available only for boundaries of regions containing the Axisymmetric model

Flow Split Outlet

The outlet boundary represents the outlet of a duct. When multiple flow split outlets bound a fluid continuum, you specify the fraction of the total mass flow through each of the outlet boundaries. The sum of all flow split fractions from all of the flow split outlet boundaries in a simulation must equal 1.0. When only one flow split boundary bounds the continuum, it is assumed that all the flow exits through this boundary.

The flow split can be specified by ratio or by mass flow rate. The pressure is not specified at a flow split outlet.

For more information, see Flow-Split Outlet.

Application

Incompressible internal flows, with multiple outlets, for which the flow split between the outlets is specified.

Compatibility

- Should not be used with compressible flows
- Cannot exist in the same continuum as pressure, stagnation and mass flow boundaries, or mixing-plane interfaces
- Not suitable for outflow boundaries at which recirculation (inflow) occurs

Free-Stream

The free-stream boundary represents the boundary of a bubble of fluid surrounding an object flying through space.

The free-stream boundary condition is based on extrapolation of Riemann invariants under the assumption of irrotational, quasi-1D flow in the boundary-normal direction. This condition generally applies to external flows when the boundary is placed sufficiently far from the body. Internal flows often have walls, possibly of irregular shape, immediately next to the flow boundary. These walls can produce boundary layers, vortices, or other multi-dimensional flow structures such that the irrotational, quasi-1D flow assumption breaks down.

For more information, see Free-Stream.

Application

- The free-stream or far-field of an external aerodynamic simulation
- If using this condition for internal flows, position the flow boundary in such a way that the irrotational, quasi-1D flow assumption is valid.

Compatibility

- Available only for boundaries of regions in which the ideal gas model is activated

Mass Flow Inlet

A mass flow inlet boundary represents an inlet for which the mass flow rate is known.

For more information, see Mass Flow Inlet.

Application

The inlet to a simulation of compressible internal flows for which the mass flow rate is known

Compatibility

- Should be used in combination with pressure outlet boundaries

Pressure Outlet

A pressure outlet boundary is a flow outlet boundary at which the pressure is specified.

For more information, see Pressure Outlet.

Applications

- Outflow of compressible and internal flows
- A plenum, for example, the boundary surrounding a jet emanating into a large chamber

Compatibility

- A pressure outlet boundary cannot exist in the same continuum as a flow split outlet.
- Although STAR-CCM+ does not prevent a pressure outlet boundary from being used as an inlet, a stagnation boundary is a better choice.

Stagnation Inlet

The stagnation inlet boundary is an inlet condition that is well-posed for compressible flows, although it is equally valid for incompressible flows. The stagnation conditions refer to the conditions in an imaginary plenum, far upstream, in which the flow is completely at rest. For incompressible flows, Bernoulli's equation is used to relate total pressure, static pressure, and velocity magnitude. For compressible ideal-gas flows, isentropic relations are used, and characteristic variables determine the propagation properties of the flow.

In compressible flow, the total pressure, total temperature, and flow direction are always used to obtain the inlet flow conditions. Confusion often arises, however, about the need to specify the static pressure at the inlet, termed supersonic static pressure in STAR-CCM+. The characteristics of the flow dictate that when the inlet flow is supersonic, all information comes from upstream conditions, including static pressure. This flow regime is termed *hyperbolic*, and the situation dictates that supersonic static pressure must be appropriately specified. When the inlet flow is subsonic, the flow regime is termed *elliptic*, which means that the static pressure is obtained from the downstream flow. In this case, the value that is specified for supersonic static pressure is ignored, and ideally, its value could be left at the default. This is not always advisable, as in some situations the inlet flow can momentarily become supersonic during the simulation iteration. If this occurs, and the supersonic static pressure is improperly specified, the solution might diverge. In these situations, it is recommended that isentropic relations be used to specify a value for supersonic static pressure that corresponds to a slightly supersonic Mach number.

As an example of using isentropic relations to set the supersonic static pressure, consider the isentropic relation for the ratio of static pressure to total pressure in an ideal gas:

$$\frac{P_t}{P} = \left\{1 + \frac{(\gamma - 1)}{2} M^2\right\}^{\gamma/(\gamma - 1)}$$

In this relation, P and P_t are the absolute static pressure and absolute total pressure, respectively, and M is the Mach number. To obtain the static pressure, one could set $M = 1.02$ for a mildly supersonic flow and use the known total pressure to compute the required value.

For more information, see Stagnation Inlet.

Applications

The inlet to a simulation of internal flows for which the stagnation conditions are known.

Compatibility

- Should be used in combination with pressure outlet boundaries
- A stagnation boundary cannot exist in the same continuum as a flow split outlet boundary.

Symmetry Plane

A symmetry plane boundary represents an imaginary plane of symmetry in the simulation. The solution that is obtained with a symmetry plane boundary is identical to the solution that would be obtained by mirroring the mesh about the symmetry plane (in half the resulting domain).

For more information, see Symmetry Plane.

Application

An imaginary plane of symmetry in a simulation that would be physically symmetrical if modeled in its entirety.

Velocity Inlet

A velocity inlet boundary represents the inlet of a duct at which the flow velocity is known.

For more information, see Velocity Inlet.

Application

The inlet to a simulation of incompressible internal flows.

Compatibility

- Can be used in combination with pressure outlet and flow split outlet boundaries
- Should not be used for compressible flow as this boundary does not prescribe the total energy of the system--which is important for compressible flow. (STAR-CCM+ does not, however, specifically prevent the use of velocity inlet boundaries for compressible flows).

Wall

A wall boundary represents an impermeable surface.

For more information, see Wall.

Applications

- An impenetrable boundary for inviscid flows
- An impenetrable, no-slip boundary for viscous flow simulations
- "Slip walls" representing an impenetrable but traction-free surface
- The boundary of a thermally active solid region at which the temperature or heat flux is specified

Interface Boundary Types

When a new interface is successfully intersected, new boundaries are created, appear in the simulation tree, and have types that are associated with them defined by the type of the interface. These boundary types cannot be changed except by changing the interface type. The possible generated boundary types are:

- Baffle boundary
- Contact interface boundary
- Fan interface boundary
- Fully developed interface boundary
- Internal interface boundary
- Mixing plane inflow boundary
- Mixing plane outflow boundary
- Porous baffle boundary

In addition, the name of the interface boundary cannot be changed except by changing the names of the original boundary and the interface.

Interface Boundary Names

Unlike the names of other boundary types, you cannot directly change the interface boundary names. STAR-CCM+ constructs them. The rule for these names is `<originalboundary> [<interface>]`.

For example, in the following screenshot, the selected nodes are of an interface, the original boundaries of that interface and the interface-type boundaries that were created along with the interface. The default name of each interface-type boundary contains the default names of the original boundary in its respective region and of the interface.



When the nodes of the interface and/or either original boundary is renamed, that name change is reflected in the corresponding nodes of the interface-type boundaries.



Link Redirection

The CD-adapco Documentation Group are migrating the STAR-CCM+ User Guide to a DITA/XML format. Migrating to a DITA/XML format provides a foundation from which CD-adapco can improve your experience with the User Guide in future releases.

In the STAR-CCM+ v9.06 release, some portions of the STAR-CCM+ User Guide are sourced from DITA/XML format and other portions are sourced from Adobe FrameMaker format. In this transition phase, there are no direct links between content in Adobe FrameMaker format and content in DITA/XML format.

In order to minimise the impact of the transition phase, this section contains a series of link redirections.

Using STAR-CCM+®

See Using STAR-CCM+

What Are Continua?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Continua?

What Are Models?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Models?

What Are Parts?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Parts?

What Are Simulations?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Simulations?

What Are Regions?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Regions?

What Are Boundaries?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Boundaries?

What Are Interfaces?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Interfaces?

What Are Solvers?

See Using STAR-CCM+ > What Are the Important Concepts? > What Are Solvers?

Working with Mixed or Double Precision

See Using STAR-CCM+ > Working with Mixed or Double Precision

Creating a New Simulation Interactively

See Using STAR-CCM+ > Working with Simulations > Creating a New Simulation Interactively

Saving Simulations to File

See Using STAR-CCM+ > Working with Simulations > Saving Simulations to File

Opening Files

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Using the Workspace > Opening Files

The Object Selection Dialog

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > The Object Selection Dialog

Single-Selection Mode

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > The Object Selection Dialog > Single-Selection Mode

Working with the Object Tree

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree

Copying Objects Using the Clipboard

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Copying Objects > Copying Objects Using the Clipboard

Copying Objects

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Copying Objects

Copying Properties

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Copying Properties

Deleting Objects

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Deleting Objects

Using Drag-and-Drop to Copy Objects

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Copying Objects > Using Drag-and-Drop to Copy Objects

Pasting Objects

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Pasting Objects

Common Pop-Up Menu Items

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Common Pop-Up Menu Items

Renaming Objects

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Renaming Objects

Using Drag-and-Drop

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Working with the Object Tree > Using Tree Views > Using Drag-and-Drop

Using a Property Customizer

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Using the Properties Window > Using a Property Customizer

Showing and Hiding the Description Area

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Using the Properties Window > Showing and Hiding the Description Area

Using the Part Selector

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Using the Properties Window > Using the Part Selector

Changing a Two-Dimensional Visualization Display

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Using the Graphics Window > Changing a Two-Dimensional Visualization Display

Changing a Three-Dimensional Visualization Display

See Using STAR-CCM+ > Using the STAR-CCM+ Workspace > Using the Graphics Window > Changing a Three-Dimensional Visualization Display

Viewing Two Scenes Side-by-Side

See Using STAR-CCM+® > Using the STAR-CCM+ Workspace > Using the Graphics Window > Splitting the Window > Viewing Two Scenes Side-by-Side

Renaming Parts in the Scene

See Using STAR-CCM+® > Using the STAR-CCM+ Workspace > Using the Graphics Window > Working with Graphics Windows > Renaming Parts in the Scene

Setting Conditions and Values

See Using STAR-CCM+ > Setting Conditions and Values

What Are Conditions?

See Using STAR-CCM+ > Setting Conditions and Values > What Are Conditions?

What Are Profiles?

See Using STAR-CCM+ > Setting Conditions and Values > What Are Profiles?

Using Scalar Profiles

See Using STAR-CCM+ > Setting Conditions and Values > What Are Profiles? > Using Scalar Profiles

Using Pressure Profiles

See Using STAR-CCM+ > Setting Conditions and Values > What Are Profiles? > Using Pressure Profiles

Using Scalar Array Profiles

See Using STAR-CCM+ > Setting Conditions and Values > What Are Profiles? > Using Scalar Array Profiles

Using Vector Profiles

See Using STAR-CCM+ > Setting Conditions and Values > What Are Profiles? > Using Vector Profiles

Using Tensor Profiles

See Using STAR-CCM+ > Setting Conditions and Values > What Are Profiles? > Using Tensor Profiles

What Are Values?

See Using STAR-CCM+ > Setting Conditions and Values > What Are Values?

Setting Conditions

See Using STAR-CCM+ > Setting Conditions and Values > Setting Conditions

Setting Profile Methods

See Using STAR-CCM+ > Setting Conditions and Values > Setting Profile Methods

Setting Values Using the Composite Vector Profile Method

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using the Composite Vector Profile Method

Using the Scalar Profile Quantity

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using the Composite Vector Profile Method > Using the Scalar Profile Quantity

Setting Values Using a Constant

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Constant

Setting Values Using a Composite Array

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Composite Array

Setting Values Using a Constant Array

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Constant Array

Setting Values Using a Field Function

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Field Function

Setting Values Using an Expression

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using an Expression

Setting Values Using Radial Equilibrium

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using Radial Equilibrium

Setting Values Using a Normal Distribution

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Normal Distribution

Setting Values Using a Normal Distribution Array

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Normal Distribution Array

Setting Values Using a Table

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table

Setting Vector Profiles Using a Table

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Setting Vector Profiles Using a Table

What Is Table (iteration)?

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > What Is Table (iteration)?

What Is Table (r)?

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > What Is Table (r)?

What Is Table (time)?

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > What Is Table (time)?

What Is Table (x,y,z)?

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > What Is Table (x,y,z)?

What Is Table (xyz, time)?

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > What Is Table (xyz, time)?

What Is Table (er)?

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > What Is Table (er)?

Example of Creating a Table for a Profile

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Example of Creating a Table for a Profile

Example of Extracting and Exporting a Table

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Example of Extracting and Exporting a Table

Example of Using Data with Sharp Discontinuities

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Example of Using Data with Sharp Discontinuities

Table Properties (Iteration, Radius, and Time)

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Table Properties Iteration Radius and Time

Properties of Table (x,y,z) and Table (xyz, time)

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Properties of Table (x,y,z) and Table (xyz, time)

Vector Table Properties (Iteration, Radius, and Time)

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Vector Table Properties Iteration Radius and Time

Properties of Vector Table (x,y,z) and Vector Table (xyz, time)

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Properties of Vector Table (x,y,z) and Vector Table (xyz, time)

Expert Properties (Vector Table (xyz, time) only)

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Table > Properties of Vector Table (x,y,z) and Vector Table (xyz, time) > Expert Properties (Vector Table (xyz, time) only)

Setting Values Using a Feature Curve

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Feature Curve

Setting Values Using a Uniform Distribution

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Uniform Distribution

Setting Values Using a Uniform Distribution Array

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Uniform Distribution Array

Setting Values Using User Code

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using User Code

Setting Values Using a Composite Symmetric Tensor

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Composite Symmetric Tensor

Setting Values Using a Composite Tensor

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Composite Tensor

Setting Values Using a Principal Tensor

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Principal Tensor

Setting Values Using an Axisymmetric Tensor

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using an Axisymmetric Tensor

Setting Values Using an Isotropic Tensor

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using an Isotropic Tensor

Setting Values Using a Polynomial

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Polynomial

Polynomial Properties

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Polynomial > Polynomial Properties

Example of Using a Polynomial

See Using STAR-CCM+ > Setting Conditions and Values > Setting Values Using a Polynomial > Example of Using a Polynomial

Copying Properties

See Using STAR-CCM+ > Setting Conditions and Values > Copying Properties

Defining Dimensionality

See Using STAR-CCM+ > Setting Conditions and Values > Defining Dimensionality

Setting Multiple Corresponding Conditions Simultaneously

See Using STAR-CCM+ > Setting Conditions and Values > Setting Multiple Corresponding Conditions Simultaneously

Setting Multiple Corresponding Values Simultaneously

See Using STAR-CCM+ > Setting Conditions and Values > Setting Multiple Corresponding Values Simultaneously

Setting Reference Values

See Using STAR-CCM+ > Setting Conditions and Values > Setting Reference Values

What Is Client-Server?

See Using STAR-CCM+ > Working With Clients and Servers > What Is Client-Server?

Running the Client in Batch

See Using STAR-CCM+ > Working With Clients and Servers > Running the Client in Batch

What Is a Parallel Server?

See Using STAR-CCM+ > Working With Parallel Servers > What Is a Parallel Server?

Using a Remote Multi-Core Workstation

See Using STAR-CCM+ > Working With Parallel Servers > Using a Remote Multi-Core Workstation

Setting Reference Values

See Using STAR-CCM+ > Setting Conditions and Values > Setting Reference Values

Modeling Geometry

See Modeling Geometry

Exporting the Coordinate System from a Sketch Plane

See Modeling Geometry > Using 3D-CAD > Sketches > Sketch Planes > Exporting the Coordinate System from a Sketch Plane

Importing Geometry into 3D-CAD

See Modeling Geometry > Using 3D-CAD > Importing Geometry into 3D-CAD

Exporting Geometry from 3D-CAD

See Modeling Geometry > Using 3D-CAD > Exporting Geometry from 3D-CAD

Importing Geometry

See Modeling Geometry > Importing Geometry

Importing Multiple Surface Data Files

See Modeling Geometry > Importing Geometry > Importing Multiple Surface Data Files

What Is a Geometry Part?

See Modeling Geometry > Understanding Geometry Parts > What Is a Geometry Part?

Imprinting Parts on Their Neighbors

See Modeling Geometry > Understanding Geometry Parts > Using Part Contact Data to Create Interfaces > Imprinting Parts on Their Neighbors

What Are CAD Parts?

See Modeling Geometry > Understanding Geometry Parts > What Is a Geometry Part? > What Are CAD Parts?

What Are Mesh Parts?

See Modeling Geometry > Understanding Geometry Parts > What Is a Geometry Part? > What Are Mesh Parts?

Using Part Contact Data to Create Interfaces

See Modeling Geometry > Understanding Geometry Parts > Using Part Contact Data to Create Interfaces

Assigning Geometry Parts to Regions

See Defining the Simulation Topology > Assigning Geometry Parts to Regions

Create One Region for All Parts

See Defining the Simulation Topology > Assigning Geometry Parts to Regions > Assigning Parts to Regions > Using Set Region > Using the Assign Parts to Regions Dialog > Region Options

Transforming Regions

See Defining the Simulation Topology > Transforming Regions

Rotating Regions

See Defining the Simulation Topology > Transforming Regions > Rotating Regions

Scaling Regions

See Defining the Simulation Topology > Transforming Regions > Scaling Regions

Translating Regions

See Defining the Simulation Topology > Transforming Regions > Translating Regions

Region Combination

See Defining the Simulation Topology > Region Manipulation > Region Combination

Region Splitting

See Defining the Simulation Topology > Region Manipulation > Region Splitting

Manipulating Boundaries

See Defining the Simulation Topology > Manipulating Boundaries

Fusing Boundaries

See Defining the Simulation Topology > Manipulating Boundaries > Fusing Boundaries

Creating a Boundary Interface

See Defining the Simulation Topology > Working with Interfaces > Working with Interfaces between Boundaries > Creating a Boundary Interface

Specifying the Transformation and Axis Properties for Periodic Interfaces

See Defining the Simulation Topology > Working with Interfaces > Working with Interfaces between Boundaries > Examining Interfaces > Specifying the Transformation and Axis Properties for Periodic Interfaces

Adjusting the Interface Intersection Tolerance

See Defining the Simulation Topology > Working with Interfaces > Working with Interfaces between Boundaries > Adjusting the Interface Intersection Tolerance

Intersecting an Interface

See Defining the Simulation Topology > Working with Interfaces > Working with Interfaces between Boundaries > Intersecting an Interface

Working with Interfaces

See Defining the Simulation Topology > Working with Interfaces

Boundary Interface Types Reference

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference

What Is a Direct Interface?

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface?

In-Place Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > In-Place Interface

Periodic Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Periodic Interface

Repeating Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Repeating Interface

Baffle Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Baffle Interface

Porous Baffle Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Porous Baffle Interface

Contact Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Contact Interface

Fluid Film Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Fluid Film Interface

Fully-Developed Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Fully Developed Interface

Internal Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is a Direct Interface? > Internal Interface

What Is an Indirect Interface?

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is an Indirect Interface?

Mapped Contact Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is an Indirect Interface? > Mapped Contact Interface

Mixing Plane Interface

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > What Is an Indirect Interface? > Mixing Plane Interface

Interface Pop-Up Menu

See Defining the Simulation Topology > Working with Interfaces > Boundary Interface Types Reference > Interface Pop Up-Menu

Meshing

See Meshing

Cell

See Meshing > Meshing Overview > Mesh Construction > Cell

Face

See Meshing > Meshing Overview > Mesh Construction > Face

Assigning Mesh Continua to Regions

See Meshing > Region-Based Meshing > Creating a Region-Based Mesh > Creating Mesh Continua > Assigning Mesh Continua to Regions

Manually Creating a Mesh Continuum

See Meshing > Region-Based Meshing > Creating a Region-Based Mesh > Creating Mesh Continua > Manually Creating a Mesh Continuum

Generating the Volume Mesh

See Meshing > Region-Based Meshing > Creating a Region-Based Mesh > Generating the Volume Mesh

Surface Wrapper

See Meshing > Surface Meshers > Surface Wrapper

Surface Remesher

See Meshing > Surface Meshers > Surface Remesher

Checking the Volume Mesh

See Meshing > Checking the Volume Mesh

What Is a Valid Mesh?

See Meshing > Checking the Volume Mesh > What Is a Valid Mesh?

What Is Mesh Quality?

See Meshing > Checking the Volume Mesh > What Is Mesh Quality?

Interpreting the Compact Report

See Meshing > Checking the Volume Mesh > Running Mesh Diagnostics > Interpreting the Compact Report

Checking For Zero or Negative Volume Cells

See Meshing > Checking the Volume Mesh > Checking The Mesh Validity > Checking For Zero Or Negative Volume Cells

Checking the Skewness Angle

See Meshing > Checking the Volume Mesh > Checking the Skewness Angle

What Is Cell Skewness Angle?

See Meshing > Checking the Volume Mesh > Checking the Skewness Angle > What Is Cell Skewness Angle?

Extruder Mesher

See Meshing > Volume Meshers > Extruder Mesher

What Are Representations?

See Meshing > Working With Representations > What Are Representations?

Creating a New Cell Set

See Meshing > Working With Representations > Using Cell Sets > Creating a New Cell Set

Using Cell Sets

See Meshing > Working With Representations > Using Cell Sets

Exporting a Surface Mesh

See Meshing > Working With Representations > Exporting a Surface Mesh

Two-Dimensional Meshes

See Meshing > Two Dimensional and Axisymmetric Meshes

Scaling the Entire Mesh

See Meshing > Scaling the Entire Mesh

Multi-Region Meshing

See Meshing > Multi-Region Meshing

Best Practices

See Best Practices

Guidelines for Aeroacoustics Calculations

See Best Practices > Guidelines for Aeroacoustics Calculations

Axial Fans

See Best Practices > Guidelines for Aeroacoustics Calculations > Axial Fans

Weak Coupling

See Best Practices > Modeling Fluid-Structure Interaction > What Is Fluid-Structure Interaction? > Classifications of FSI Coupling > Weak Coupling

Strong Coupling

See Best Practices > Modeling Fluid-Structure Interaction > What Is Fluid-Structure Interaction? > Classifications of FSI Coupling > Strong Coupling

Using Small Displacements for FSI

See Best Practices > Modeling Fluid-Structure Interaction > Modeling FSI Using Finite Volume Stress > Using Small Displacements for FSI

Using Large Displacements for FSI

See Best Practices > Modeling Fluid-Structure Interaction > Modeling FSI Using Finite Volume Stress > Using Large Displacements for FSI

Reference Configuration

See Best Practices > Modeling Fluid-Structure Interaction > What Is Fluid-Structure Interaction? > Mapping Between Structure and Fluid Grids

Modeling FSI Using Finite Volume Stress

See Best Practices > Modeling Fluid-Structure Interaction > Modeling FSI Using Finite Volume Stress

Migrating from Fluent

See Best Practices > Migrating from Other Codes > Migrating from Fluent

Running

See Running

What Are Initial Conditions?

See Running > Using Initial Conditions > What Are Initial Conditions?

Using Initial Conditions

See Running > Using Initial Conditions

Setting Initial Conditions

See Running > Using Initial Conditions > Setting Initial Conditions

Using Constant Values

See Running > Using Initial Conditions > Setting Initial Conditions > Using Constant Values

Using Field Function Values

See Running > Using Initial Conditions > Setting Initial Conditions > Using Field Function Values

Using Tabular Data Values

See Running > Using Initial Conditions > Setting Initial Conditions > Using Tabular Data Values

Setting Initial Conditions for a Particular Region

See Running > Using Initial Conditions > Setting Initial Conditions > Setting Initial Conditions for a Particular Region

Using Solvers

See Running > Using Solvers

Setting up Stopping Criteria

See Running > Setting Up Stopping Criteria

Setting the Maximum Number of Inner Iterations

See Running > Setting Up Stopping Criteria > Using Automatically Generated Stopping Criteria > Setting the Maximum Number of Inner Iterations

Setting the Maximum Physical Time

See Running > Setting Up Stopping Criteria > Using Automatically Generated Stopping Criteria > Setting the Maximum Physical Time

Obtaining a Solution Interactively

See Running > Obtaining a Solution Interactively

Initializing the Solution

See Running > Obtaining a Solution Interactively > Initializing the Solution

What Is Solution Initialization?

See Running > Obtaining a Solution Interactively > Initializing the Solution > What Is Solution Initialization?

Initializing Manually

See Running > Obtaining a Solution Interactively > Initializing the Solution > Initializing Manually

Clearing the Solution

See Running > Obtaining a Solution Interactively > Clearing the Solution

Clearing a Region's Solution Data

See Running > Obtaining a Solution Interactively > Clearing the Solution > Clearing a Region's Solution Data

What Is a Residual?

See Running > Obtaining a Solution Interactively > Watching Residuals

Stepping the Solution

See Running > Obtaining a Solution Interactively > Stepping the Solution

Automatically Saving at Regular Intervals

See Running > Obtaining a Solution Interactively > Saving Intermediate Solutions > Automatically Saving at Regular Intervals

Using the Java API

See Running > Scripting the Application > Using the Java API

What Are User-Coded Profiles?

See Running > Working with User Code > What Are User Functions?

Changing Boundary Conditions or Ramping Boundary Values

See Running > Troubleshooting the Solution > Getting Past a Poor Initial Solution > Changing Boundary Conditions or Ramping Boundary Values

Analyzing

See Analyzing

Working with Derived Parts

See Analyzing > Accessing Solution Data through Derived Parts > Working with Derived Parts

Working with Implicit Sections

See Analyzing > Accessing Solution Data through Derived Parts > Working with Implicit Sections

Using Coordinate Systems with Derived Parts

See Analyzing > Accessing Solution Data through Derived Parts > Working with Derived Parts > Using Coordinate Systems with Derived Parts

Exporting to an External File

See Analyzing > Accessing Solution Data through Derived Parts > Working with Derived Parts > Exporting to an External File

What Is a Derived Part?

See Analyzing > Accessing Solution Data through Derived Parts > What Is a Derived Part?

Creating a Point Part

See Analyzing > Accessing Solution Data through Derived Parts > Working with Probes > Defining a Point > Creating a Point Part

Defining an Arbitrary Probe

See Analyzing > Accessing Solution Data through Derived Parts > Working with Probes > Defining an Arbitrary Probe

Defining a Cell Surface

See Analyzing > Accessing Solution Data through Derived Parts > Working with Topological Parts > Defining a Cell Surface

Defining an Isosurface

See Analyzing > Accessing Solution Data through Derived Parts > Working with Scalar Value-Based Parts > Defining an Isosurface

Defining a Threshold

See Analyzing > Accessing Solution Data through Derived Parts > Working with Scalar Value-Based Parts > Defining a Threshold

Visualizing the Solution

See Analyzing > Visualizing the Solution

What Is a Scene?

See Analyzing > Visualizing the Solution > Working with Scenes > What Is a Scene?

Creating Images for an Animation

See Analyzing > Visualizing the Solution > Changing the Attributes of a Scene > Creating Images for an Animation

Managing Parts Through a Scene

See Analyzing > Visualizing the Solution > Working with Scenes > Managing Parts Through a Scene

Creating a Scalar Scene

See Analyzing > Visualizing the Solution > Working with Scenes > Creating a Scene > Creating a Scalar Scene

Creating a Vector Scene

See Analyzing > Visualizing the Solution > Working with Scenes > Creating a Scene > Creating a Vector Scene

Updating Scenes

See Analyzing > Visualizing the Solution > Changing the Attributes of a Scene > Updating Scenes

Using Displayers

See Analyzing > Visualizing the Solution > Using Displayers

What Is a Displayer?

See Analyzing > Visualizing the Solution > Using Displayers > What Is a Displayer?

Displaying Scalar Quantities

See Analyzing > Visualizing the Solution > Displaying Scalar Quantities

What Is the Point Size Scale Field?

See Analyzing > Visualizing the Solution > Displaying Scalar Quantities > Working with Point Scale Modes

Adding Parts to Displayers

See Analyzing > Visualizing the Solution > Using Displayers > Managing Parts in Displayers > Adding Parts to Displayers

Using Surface and Outline Displays

See Analyzing > Visualizing the Solution > Displaying Geometry and Mesh > Using Surface and Outline Displays

Managing Parts in Displayers

See Analyzing > Visualizing the Solution > Using Displayers > Managing Parts in Displayers

Animating Scenes

See Analyzing > Visualizing the Solution > Animating Scenes

Changing the Scalar Quantity

See Analyzing > Visualizing the Solution > Displaying Scalar Quantities > Changing the Scalar Quantity

Changing the Vector Quantity

See Analyzing > Visualizing the Solution > Displaying Vector Quantities > Changing the Vector Quantity

Controlling the Animation

See Analyzing > Visualizing the Solution > Animating Scenes > Controlling the Animation

Using Streamline Tracers

See Analyzing > Visualizing the Solution > Animating Scenes > Using Streamline Tracers

Recording the Animation onto Separate Files

See Analyzing > Visualizing the Solution > Animating Scenes > Recording the Animation onto Separate Files

Creating Geometry Displayers

See Analyzing > Visualizing the Solution > Displaying Geometry and Mesh > Creating Geometry Displayers

Creating Streamline Displayers

See Analyzing > Visualizing the Solution > Displaying Streamlines > Creating Streamline Displayers

Creating Vector Displayers

See Analyzing > Visualizing the Solution > Displaying Vector Quantities > Creating Vector Displayers

Understanding the Simulation History File

See Analyzing > Recording and Viewing Solutions > Understanding the Simulation History File

What Is a Report?

See Analyzing > Reporting Results > What Is a Report?

Running Reports

See Analyzing > Reporting Results > Running Reports

Using Reports as Field Functions

See Analyzing > Reporting Results > Using Reports as Field Functions

Using Coordinate Systems with Force and Moment Reports

See Analyzing > Reporting Results > Using Coordinate Systems with Force and Moment Reports

Adding Reports

See Analyzing > Reporting Results > Adding Reports

Surface Integral

See Analyzing > Reporting Results > Report Types Reference > Statistical Reports > Surface Integral

Sum

See Analyzing > Reporting Results > Report Types Reference > Statistical Reports > Sum

Rigid Body Angle

See Analyzing > Reporting Results > Report Types Reference > Specific Report Types > Rigid Body Angle

Rigid Body 1-DOF Angle

See Analyzing > Reporting Results > Report Types Reference > Specific Report Types > Rigid Body 1-DOF Angle

Li-Ion Cell

See Analyzing > Reporting Results > Report Types Reference > Specific Report Types > Li-Ion Cell

Heat Exchanger

See Analyzing > Reporting Results > Report Types Reference > Specific Report Types > Heat Exchanger

Generalized Vibratory Force

See Analyzing > Reporting Results > Report Types Reference > Specific Report Types > Generalized Vibratory Force

Force

See Analyzing > Reporting Results > Report Types Reference > Specific Report Types > Force

DarsCFD Reports

See Analyzing > Reporting Results > Report Types Reference > Statistical Reports > DarsCFD Reports

What Is a Monitor?

See Analyzing > Monitoring the Solution > What Is a Monitor?

What Is a Field Function Monitor?

See Analyzing > Monitoring the Solution > Using Field Function Monitors > What Is a Field Function Monitor?

Using Field Function Monitors

See Analyzing > Monitoring the Solution > Using Field Function Monitors

Monitor Properties

See Analyzing > Monitoring the Solution > What Is a Monitor? > Monitor Properties

Changing the Trigger Event and Sampling Frequency

See Analyzing > Monitoring the Solution > Customizing Monitors > Changing the Trigger Event and Sampling Frequency

Plotting Results

See Analyzing > Plotting Results

Customizing Plots

See Analyzing > Plotting Results > Customizing Plots

Trigger Properties

See Analyzing > Triggers and Update Events > Trigger Properties

Using Data Set Functions

See Analyzing > Using Data Set Functions

Exporting Fields

See Analyzing > Using Data Set Functions > What Are Surface Time Histories? > Exporting Fields

What Are Point Time Fourier Transforms?

See Analyzing > Using Data Set Functions > What Are Point Time Fourier Transforms?

Working with Surface Spectra Field Functions

See Analyzing > Using Data Set Functions > Post-Processing Data Set Functions > Post-Processing Surface Spectra > Working with Surface Spectra Field Functions

Wizard Guide

See Link: Wizard Guide

Using Overset Meshes

Overset meshes, also known as "Chimera" or overlapping meshes, are used to discretize a computational domain with several different meshes that overlap each other in an arbitrary manner. They are most useful in problems dealing with multiple or moving bodies, as well as optimization studies. In most cases, using overset meshes does not require any mesh modification after generating the initial mesh, thus offering greater flexibility over the standard meshing techniques.

Overview

Any study involving overset meshes has a background region enclosing the entire solution domain and one or more smaller regions containing the bodies within the domain. A schematic is shown below:



In an overset mesh, cells are grouped into active, inactive, or acceptor cells. Within active cells, discretized governing equations are solved. Within inactive cells, no equation is solved, however, these cells can become active if the overset region is moving. Acceptor cells separate active and inactive cells in the background region and are attached to the overset boundary in

the overset region. Acceptor cells are used to couple solutions on the two overlapping grids. Variable values at donor cells of one mesh express variable values at acceptor cells in the other mesh, through interpolation. The donor cells are the active cells from the other mesh that are nearest the acceptor cell. The available interpolation options are described in the coupling regions section.

The solution is computed for all active cells in all regions simultaneously, that is, the meshes are implicitly coupled. When a reference is made to the variable value in an acceptor cell of one mesh within the discretized equations, a blend of variable values at donor cells from the other mesh is used. This value is reflected directly in the coefficient matrix of the algebraic equation system. This tight coupling of the overset and background regions allows for a solution that is within an arbitrary low level of iteration errors. The rate of convergence of the iterative solution method is therefore similar to that of a single mesh of the same resolution. Refer to the overset mesh methodology section for further description.

Applications

Overset meshes are most applicable in the following cases:

- Parametric studies, such as identical bodies at different relative positions or different bodies in the same environment. Fewer meshes are generated for these cases when using overset meshes.

- Complicated body motions, such as those involving bodies that are close together, have intersecting paths, or have large motions. Below is an example of two cars overtaking the third. Overset meshes are used for the two colored cars in motion.



Current Implementation

The overset mesh feature is supported for use with the following:

- Segregated and coupled flow
- DFBI
- Discrete Element Method (DEM)
- Motion models and reference frames
- Single phase
- VOF Multiphase
- Lagrangian Multiphase
- Eulerian Multiphase
- Fluid Film
- Combustion

- Passive Scalar

Current implementation of the overset mesh feature has the following limitations. Some of these restrictions will be lifted in future releases.

- Mass conservation is not strictly imposed with the current interface treatment. The usual error is on the order of 0.1% or less, that is, the error does not affect the solution to the first 3 - 4 significant digits.
- Regions that are involved in creating an overset mesh must be assigned to the same physics continuum.
- When using overset mesh together with the fluid film model, certain requirements must be met. For more information, see Setting up Overset Mesh with Fluid Film.
- The feature is not supported for the following STAR-CCM+ models:
  - De-Fogging
  - Harmonic Balance
  - Grid Sequencing
  - Radiation
  - Batteries.

Current Overset Capabilities

- Overlapping of multiple overset meshes
- Overset regions can extend beyond the solution domain

Working with Overset Grids in STAR-CCM+

The subsequent sections provide information on setting up and using overset meshes in STAR-CCM+:

- Summary of the overset mesh user procedure
- Creating meshes for the background and overset regions
- Creating an Overset Mesh Interface
- Best Practices
- Overset mesh methodology
- Overset mesh field function
- Troubleshooting

Understanding the Overset Workflow

The general workflow for using an overset mesh involves the steps that are outlined below. Two tutorials, transonic flow over an airfoil and DFBI Lifeboat, using overset meshes are available in the tutorial guide.



Preparing Overset Regions

Using an overset mesh requires creating a background region, which is a closed surface solution domain, and one or more overset regions, which contain the physical bodies. More than one overset region can be used on top of the background region. These overset regions can also overlap each other.

Setting the Overset Mesh Boundary

The overset mesh is created using the Overset Mesh boundary, which is the outer boundary of the overset region that is expected to be coupled with the background mesh. If part of the overset region boundary coincides with the boundary of the background region, that part must be of the same boundary

type as the background region. An example is shown below. The ground boundary in this case is part of both the background and the overset regions. It is not considered to be part of the Overset Mesh boundary.



To assign an Overset Mesh boundary:

1. Navigate to the designated overset mesh boundary under the Boundaries node in the overset region.

2. In the *Properties* window, set the type to Overset Mesh.



Coupling Regions

The Overset Mesh Interface is used to couple the overset regions with the background region. If you want to define overset regions that overlap each other, create overset mesh interfaces between each set of overlapping overset regions. The interfaces for the overlapping regions are in addition to the overset mesh interface between the respective overset region and the background mesh.

To create the interface, the overset regions must have an Overset Mesh boundary defined. The option to create this interface only becomes available once a physics continuum is defined and assigned to both the background and overset region. This interface ensures that both regions use the same physics continuum.

You have two options for creating such interfaces:

- Creating Overset Mesh Interfaces
- Creating Overset Mesh Zero Gap Interfaces

Creating Overset Mesh Interfaces

Use this type of overset mesh interface for any overset mesh simulation except for a zero gap configuration.

To create the interface between the background and the overset region:

1. Select the background region and the overset regions. You can use `<Ctrl><Click>` to select both regions.
2. Right-click on the selected regions and choose Create Interface > Overset Mesh.



An Overset Mesh 1 node is added under Interfaces.

To create an interface between overset regions that overlap each other:

1. Select each pair of overlapping overset regions. You can use `<Ctrl><Click>` to select both regions.
2. Right-click on the selected regions and choose Create Interface > Overset Mesh.

Creating Overset Mesh Zero Gap Interfaces

Use this type of overset mesh interface if you expect a zero gap configuration to occur in your simulation. Zero gap configuration means that two regions, either one background region and one overset region, or two overset regions, approach each other very closely or make contact.

The following diagrams of a ball valve show an example of such a configuration:





The upper diagram shows the mesh in its initial state, that is, before the simulation starts. The overset region (grey) consists of the mesh of the ball. The overset region extends beyond the boundaries of the background region (blue). You can see a very small gap between the background region and the overset region. The distance between the background region and the overset region is less than two cell layers and the zero gap algorithm applies.

The lower diagram shows the mesh after the simulation started. You see that the cells (overset and background) in the gap and the overset cells outside of the background region are inactive.

When you create the overset mesh zero gap interface between two regions, a zero gap wall boundary is automatically created in each of the regions. During the simulation, if the Overset Assembly process detects that the distance between the two zero gap wall boundaries of both regions is less than two cell layers, the cells in the gap become inactive. Cell faces between active and inactive cells are placed within the zero gap wall boundaries.

To create an overset mesh zero gap interface:

1. Multi-select the background region and the overset regions.

2. Right-click on the selected regions and choose Create Interface > Overset Mesh > Overset Mesh ZeroGap.



Two ZeroGapWall boundaries are added under the Boundaries node and an Overset Mesh ZeroGap 1 node is added under the Interfaces node.



3. Select the Overset Mesh ZeroGap 1 node and set the Overset Mesh Interface Properties.

Overset Mesh Interface Properties

Verbosity

Specifies how much extra information is provided in the *Output* window while the simulation is running. The default is to provide no additional information (None). Increasing verbosity is helpful when you are troubleshooting problem cases.

Region-0

Identifies the Overset region (read-only).

Region-1

Identifies the Background region (read-only).

Close Proximity

Activate this option when there is a narrow gap between the moving body and a boundary, such as a wall. This option allows more active cells in the overset mesh between the moving body and the boundary. This feature helps to prevent mesh intersector failure when the overset region overlaps the physical boundaries of the background region.

When the moving body goes close to a boundary, the overlapping region of the overset mesh can extend over the boundary. Normally any cell in the overset mesh that has a vertex outside the boundary (and therefore outside the background region) is deactivated. However, when this option is activated, those cells remain active as long as the cell centroid is within the background region. The cell is deactivated only when the cell centroid is outside the boundary.

Alternate Hole Cutting

Determines whether the global or the layered hole-cutting approach is used. For more information, see Hole-cutting process and Which Hole-Cutting Approach should you use?.

- **Activated** — Activates the alternate hole-cutting process: the Global Approach.
- **Deactivated** — This is the default option. Activates the default hole-cutting process: the Layered Approach.

Type

Defines the interface type as Overset Mesh Interface.

Interpolation Options

STAR-CCM+ couples the solution in the overset and background regions by interpolating donor cell values to provide data for acceptor cells. The available interpolation options are found under the Interpolation Option node.



Interpolation Option Properties

Options	Provides options for the overset mesh interpolation.	
	Distance weighted	Weighted based on distances to the four nearest neighbors (Default).
	Linear	Linear for 2D and 3D meshes.
	Linear Quasi2D	Linear for 3D meshes with only one layer of cells.
	Least Square	Least square for 2D and 3D meshes.

Note: Linear interpolation is a more accurate but computationally more expensive option.

Dynamic Overset Behavior

You use the dynamic overset behavior for overset mesh applications where an automatic switch from a wall boundary type to an overset boundary type or vice versa is useful. Examples for such applications are:

- Closely counter-rotating objects such as screw pumps
- IC engine piston motion where the cylinder has a variable radius

The dynamic wall boundary is an option that is available for wall boundaries after an overset mesh interface is created. There are two main scenarios for this option:

- Overset simulations without background mesh
- Overset simulations with background mesh

Simulations without Background Mesh

When you enable the dynamic overset behavior to model a scenario without a background mesh, the hole-cutting process checks if a boundary face resides inside or outside of the other overset region. If the hole-cutting process detects that a boundary face lies within the other overset region, the cells adjacent to the boundary become acceptor cells. The previous wall boundary becomes an overset boundary or the overset boundary becomes again a wall boundary depending on position.

The following diagram shows two counter-rotating square objects. The objects are embedded within two overset regions. There is no background mesh present. The highlighted boundaries use the dynamic overset behavior.



The advantage of this functionality is the fact that you do not need to set up a background mesh.

Simulations with Background Mesh

When you enable the dynamic overset behavior to model a scenario with a background mesh, the hole-cutting process checks if the moving wall boundary of the overset region still coincides with the wall boundary of the

background mesh. If the hole-cutting process detects that a wall boundary face of the overset region no longer coincides with a wall boundary face of the background region, the cells adjacent to the boundary become acceptor cells. The previous wall boundary of the overset region becomes an overset boundary. If such an overset boundary moves such that it coincides with a background region wall boundary, it reverts to a wall boundary.

The following screenshot shows an example application: the translational motion of a simplifed piston within a cylinder of variable radius.



Here, the overset region (grey), that is, the piston is in its initial position. The side wall boundaries of this overset region coincide with the side wall boundaries of the background region (blue), which is the cylinder. You must set up the sliding side walls of the piston as a wall boundary (and not as an overset boundary), because if part of the overset region boundary coincides with the boundary of the background region, that part must be of the same boundary type as the background region. For more information, see Preparing Overset Regions.

The piston moves upward.The following screenshot shows the mesh after the piston moved upwards.



The upper part of the side boundary of the piston no longer coincides with the wall boundary of the background mesh. The dynamic overset behavior automatically detects this situation and changes the previous wall boundary to an overset boundary. The lower part of the side boundary that still coincides with the wall boundary of the background region remains a wall boundary.

Using Dynamic Overset Behavior

To use dynamic overset behavior:

1. Prepare the overset regions for dynamic overset behavior.

   See Preparing Overset Regions for Dynamic Overset Behavior

2. Multi-select the overset regions or the overset region and the background region and create an overset interface between them.

   If you set up a simulation using a background region, the sequence of the multi-select operation is important. First, select the overset region. Then, select the background region. The overset region must always be Region-0, the background region must always be Region-1.

   See Coupling Regions

3. For each wall boundary that uses dynamic overset behavior:

a. Select the Regions > [Region] > Boundaries > [Boundary] > Physics Conditions > Dynamic Overset Behaviour node.

b. Activate Dynamic Overset.



Preparing Overset Regions for Dynamic Overset Behavior

In contrast to Preparing Overset Regions, you do not prepare a background region and you do not assign an overset boundary type to the boundaries that use dynamic overset behavior. Instead, you apply the wall boundary type for those boundaries that use dynamic overset behavior.

To prepare the overset regions:

1. Select the mesh boundaries under the Boundaries node for each region for which you want to apply dynamic overset behavior.
2. Set Type to Wall for each of the boundaries.

Overset Topology

Overset topology describes how the regions for an overset mesh simulation are related to each other. The overset topology determines which region is the background region and which region is the overset region. This information is crucial for the hole-cutting process. For many situations, the algorithm can detect this region relation automatically. However, for some applications, such as a piston within a cylinder configuration, a more detailed specification is necessary.

Setting up Overset Topology

Automatic overset topology is activated by default. This option is applicable to most situations. Only for specific situations, the topology algorithm fails and you must select the appropriate overset topology method. The following additional overset topology methods are available:

- Piston

  Select this option to simulate a piston within a cylinder configuration.

- Direction

  This is an expert option. Select this option only, if you wish to simulate a situation where the overset boundary does not completely enclose the overset region and this boundary is curved strongly at the same time. The remaining parts of the enclosing surface consist of wall or symmetry boundaries coinciding with respective boundaries of the background mesh.

The following diagram shows an example of such a configuration:



Here, the cylinder is the background region (blue) where the side boundary is a wall boundary. The overset region (bronze) consists of a curved inner overset boundary and a wall boundary coinciding with the wall boundary of the background mesh. The hole-cutting algorithm requires the direction in which the overset boundary appears to be closed. For this example, this direction is the z-direction.

- Bounded

  This is an expert option. Select this option only if you wish to simulate a situation where you have meshless space within your overset region that is bounded by a wall boundary of the background region.

The following diagram shows an example (pigging case) of such a configuration:



The overset region (grey) slides inside the background region (blue) cylinder. The overset region is bounded by overset boundaries (red) and wall boundaries. The meshless space of the overset region is bounded by the wall boundary of the background region.

To set up overset topology:

1. Select the Regions > [Region] > Physics Conditions > Overset Topology node and set Option to the desired overset topology method.

   See Overset Topology Properties.



2. If you set Option to Piston or Direction, select the [Region] > Physics Values > Overset Direction Specification node and select Direction.

   See Overset Direction Specification Properties

Overset Topology Properties

Option

Specifies the method that is used to determine the overset topology of the regions.

Automatic

The algorithm detects automatically which region is background region and which region is overset region.

Piston

Specific option for applications where a piston moves within a cylinder. The cells underneath the piston are deactivated. An additional [Region] > Physics Values > Overset Direction Specification node is created where you specify the direction in which the background mesh cells are deactivated (see Overset Direction Specification Properties).

Direction

Specific expert option for applications where the overset boundary does not completely enclose the overset region and this boundary is curved strongly at the same time. An additional [Region] > Physics Values > Overset Direction Specification node is created where you specify the direction in which the overset boundary appears to be closed.

Bounded

Specific expert option for applications where the overset boundary contains meshless space that is bounded by a wall boundary of the background region.

Overset Direction Specification Properties

The Direction specification has two different definitions depending on which overset topology Option you selected in Step 1. of Setting up Overset Topology:

Piston Method Properties

Direction	Specify the direction underneath the piston. This part of the background mesh is no longer covered by the piston overset region. However, the cells of this background mesh portion must be deactivated. Normally, the hole-cutting process detects cells to be deactivated only when the background mesh is covered by an overset region. Consequently, the direction in which cells of the background mesh underneath the piston must be deactivated is specified. The direction is specified in the coordinate system selected in the associated expert property.

Piston Method Expert Properties

Coordinate System	Coordinate system in which Direction is specified.

Direction Method Properties

Direction	Specify the direction in which the overset boundary appears to be closed. The direction is specified in the coordinate system selected in the associated expert property.

Piston Method Expert Properties

Coordinate System	Coordinate system in which Direction is specified.

Setting up Overset Mesh with Fluid Film

For certain applications, it is useful to combine overset mesh with fluid film.

Examples of such applications are:

- automotive window wipers
- fuel flow on an engine cylinder with a moving piston

You can now model a fluid film shell region passing from one overset region to another or to a background region.

For the combination of overset mesh with liquid film, the following requirements apply:

- The overset model must be set up in a sliding configuration, that is, a boundary of the overset region must coincide with a boundary of the background region.
- You must create the shell regions on the sliding boundaries of the overset region and the background region.
- Each overset interface for shell regions must have a corresponding three-dimensional overset interface.

Before you begin to set up overset mesh with fluid film, generate a suitable mesh considering these requirements. The following diagram shows an example:



The small cube is modeled as the overset region (red) and slides along the side wall of the large cube, which is modeled as the background region (grey). The fluid forms over the entire side wall of the background region, which coincides with the sliding wall of the overset region. You create two shell regions for the fluid film: one on the sliding wall of the red cube and one on the side wall of the grey cube. By creating an overset interface between the two shell regions, the fluid film can pass from one shell region to the other.

To set up overset mesh with fluid film:

1. Set up the fluid film model.

See Selecting the Fluid Film Model

See Defining Fluid Film Phases

See Working with Fluid Film Phase Models

2. Set up the shell regions for the fluid film model. You must create one shell region on a boundary of the background region and one shell region on a boundary of the overset region.

   See Creating Shell Regions

   Fluid-Film interfaces and edges of the imprinted sliding overset boundary are created automatically as part of this process.



3. Set the boundary type of the imprinted sliding overset boundary edges to Overset:

- Select the Regions > [Overset Shell Region shell] > Boundaries > [Overset Boundary edges] node and set Type to Overset Mesh.



4. Set up the overset mesh boundary on the overset region.

   See Setting the Overset Mesh Boundary

5. In the overset region, select the sliding boundary and set its Type to match that of the contacting boundary within the background region.Usually, this boundary is a wall boundary.



6. Select the Regions > [Overset Shell Region shell] > Boundaries > [Overset Sliding Wall] > Physics Conditions > Reference Frame Specificiation node and set Option to Lab Frame.
7. Create the overset mesh interface between the background and the overset regions.

   See Creating Overset Mesh Interfaces

Note: Currently, overset mesh with fluid film is not compatible with zero gap interface creation.

8. Create an overset interface between the fluid film shell regions of the background region and of the overset region.
   a. Multi-select the shell regions of the background region and of the overset region.
   b. Right-click on any of the highlighted items and select Create Overset Mesh Interface.

   The fluid film can now pass from the shell region on the background region to the shell region on the overset region and vice versa.

Generating Suitable Meshes

A critical factor in successfully applying overset meshes is to have enough cells across the overlapping zone between background and overset regions. This overlapping zone must contain at least 4–5 cell layers in both background and overset meshes. To minimize discretization error, ensure that the cell sizes on both meshes are similar within the overlapping zone. The coarser of the two coupled meshes determines the error level in the interpolation. An example is shown below.



Overlapping zone

A successful coupling by use of an overset interface results in a “hole” being cut in the background mesh (hole-cutting process). During this hole-cutting process, the status of the involved cells is determined, that is, it is determined whether the cells are active or inactive. There are field functions available that allow the visualization of the cell status (see Overset Mesh Field Function Reference).

Using multiple overlapping overset meshes can reduce mesh cell count, as described in the following example:

Consider meshing a small gap, for example the gap between the rotor teeth of a gear pump, using multiple overset meshes overlapping each other. Due to the requirement that the overlapping zone must at least contain four cell layers in both background and overset meshes, you need at least eight cell layers in the gap (four cell layers from each overlapping region). If the two overlapping overset regions overlap such that the background mesh is completely covered, the background mesh does not have to fulfill the overlapping four cell layer requirement. Therefore, the cell size of the background mesh in the gap does not have to be as small as for the overset regions. However, you must ensure that the background mesh in this area is always covered completely by the overset regions. Then, you can use a coarser background mesh for those cells that are never activated. These cells are never visible when visualizing the active cells by using the Volume Mesh representation.

Special Criteria for Transient Problems

For transient problems with moving bodies, set the time step such that the movement of the overset region within one time step meets the following conditions:

- When the time integration is set to the 1st-order Euler scheme, the maximum movement is the smallest cell in the overlapping zone.
- When the time integration is set to the 2nd-order implicit time integration scheme, the maximum movement is half the smallest cell size in the overlapping zone.

If one inactive cell becomes an active cell without being an acceptor cell for at least one time step, there is no previous time step solution at that cell to use to compute the rate of change term. The 1st-order Euler scheme needs one previous time step value and the 2-order implicit scheme needs values from the previous two time steps.

Visualizing and Analyzing Solution

Overset meshes are automatically created while running a simulation. Two mesh representations are available to visualize the mesh: volume mesh and overset mesh. By default, the mesh scene is opened using the volume mesh representation. You can change the type in the *Properties* window of the Mesh node.



Below is an example of an Overset Mesh representation. This representation shows all background and overset region cells, whether they are active or inactive.



To see only the active cells, select the Volume Mesh representation. Visualization and reporting is done based on this representation.



You can visualize the solution on each mesh individually. The nodes of the background mesh most often do not coincide with the nodes of the overset mesh in the overlapping region. If the meshes are not fine enough, this discontinuity can create some visualization issues. One such issue is with contour lines. The discontinuity in the contour lines that are shown below is the result of discretization error and reduces as the mesh is refined.



The overlapping of active cells of the background and the overset region (overlapping zone) also affects the volume integrals. For the volume integrals, the volumes of all active cells are integrated, that is, the overlapping volumes of both the cells in the background and in the overset region are integrated. This way, the volume integrals introduce some inaccuracies for the whole solution domain because of the duplicate integral evaluation of the active cells in the overlapping zone.

This issue does not apply to surface integrals when walls or physical boundaries are either in a background or an overset region. However, when you have boundaries from the background mesh and the overset mesh overlapping each other, this results in overlapping active faces. In this case, you can get inaccuracies for the surface integrals because all active faces are integrated. There is also a duplicate integral evaluation in the overlapping zone of the boundaries.

Overset Mesh Best-Practice Guidelines

This section outlines some guidelines for setting up an appropriate mesh in the overlapping region of an overset mesh.

The *required* features of overset regions and meshes are:

- Between the wall and the overset boundary of an overset region, there must be at least 4–5 cells in both the overset and the background mesh.

To demonstrate this point, consider a cross section of a simplified blower model that is shown below:



The overset region includes the set of rotor blades. It extends to the top and bottom walls of the blower in the axial direction, but two cylindrical surfaces limit the region in the radial direction. The space between the leading edge of the blades and the outer cylinder is about twice as large as the space between the trailing edge and the inner cylinder. Based on the given requirement, you need a background mesh that is twice as fine in the inner overlapping zone as in the outer overlapping zone. The cell sizes can

be clearly seen when you zoom in on one of the blades. The background mesh is shown in blue while the overset mesh is in red. Only the active cells are shown.



- To couple the overset and the background meshes, there must be at least 4–5 cells in each mesh in the overlap region.

This requirement is illustrated below with a body moving through a channel with a small clearance at the top. The image below is zoomed in on the body. The background mesh is shown in red while the overset mesh is in blue.



In this example, the overset region is kept completely within the background region. This means that the overset region covers only a small distance above the body, limited by the width of the gap between the body and the channel wall. To account for this limitation, the background mesh

has a prism layer with five cells in it along the channel wall. The thickness of the prism is made equal to the width of the gap between the body and the channel wall.

In the overset mesh for this example, there is a prism layer along the overset boundary containing four cells, and a prism layer along the body wall containing two cells. As both prism layers are present in the gap region, there are enough overlapping cells to ensure successful coupling of the two meshes as the body moves along the channel.

The minimum number of 4–5 cells that are required to couple the overset and background regions refers to cell count in the direction normal to the boundary. Depending on the cell aspect ratio and the orientation of cells on the two meshes relative to each other, this requirement does not apply. This criterion is also not applicable if an overset region is not fully embedded in the background region, that is, the overset and background regions share a boundary. In the overset region for such a case, the boundary must be of the same type as in the background region, and not an Overset Mesh boundary.

If the mesh does not fulfill all of the required features, the initialization of the overset region interface fails. It is not always obvious where the problem arises. For a moving overset region where a new Overset Mesh Interface is created at every time step, it could be that one or more requirements are violated at only one particular time. The design of the overset and background meshes thus requires a careful assessment of all possible situations during the simulation.

Refer to the troubleshooting section for how to visualize the overlapping cell status.

In addition to the above requirements, the following mesh features are *desirable* for accuracy reasons:

- The cells in the overlapping region are of similar size on both the background and overset meshes.

  Interpolation errors are introduced by expressing field values at acceptor cells in one mesh using field values at donor cells in the other mesh. If cells in both sides of the overlapping region are of similar size, these errors are reduced to the same order of magnitude as other discretization errors.

  For example, the same interpolation method is used in overset meshes as in computing field values at cell faces when convective fluxes are approximated. If the cells are of the same size, shape and orientation, then the distance between the acceptor cell centroid and the nearest donor cell centroid is never larger than half the cell width. This distance is the distance between cell face center and the cell centroid. If the cells on one mesh are substantially larger, interpolation errors when using

donor cells from this mesh are larger than in the other direction. This error is not important when the variation of field values is small. However, this error significantly affects the solution when the overlapping zone is part of a region where there is a significant variation in the solution.

- Select the extent of the overset region such that variables do not vary much in the overlapping zone.

  If you cannot minimize this variation in the overlapping zone, refine the mesh to satisfy accuracy requirements.

Mesh adaptation and interpolation accuracy is demonstrated in the following example in which wine is poured from a bottle into a glass. The bottle is initially vertical and positioned to the right of the glass. The bottle is placed back into this position after the pouring process.



The cell size in the background mesh is adapted to the size of the mesh around the bottle in the overset region within the zone of its motion. The background mesh is two times coarser than the overset mesh around the bottle opening while it is away from the glass and no wine is being poured. At this stage, the velocities are small in this region with insignificant

interpolation errors. Once the wine is poured and the bottle is brought to a nearly horizontal position, the overset mesh overlaps the refined background mesh above the glass. The interpolation accuracy matches that of other discretization approximations.

Overlapping the Background Region

To simplify cases that deal with small gaps between walls, you can set the overset region to extend beyond the solution domain. Any mesh cells of an overset region that lie outside of the background region are inactive. These cells can only become active again if they enter the background region, which always represents the solution domain. If the overset region moves outside the solution domain, new acceptor cells are found that are inside the background region such that the coupling criteria are satisfied.

Note: Of the overset region, only boundaries of the type Overset Mesh or Symmetry Plane can protrude the background mesh. Wall boundaries are not allowed outside of the background mesh.

You can see this condition in the following screenshot that shows motion using an injector.



The overset region in this case is over the needle, which has part of its geometry outside of the nozzle. Those cells are inactive and are in red. The acceptor cells are updated to represent the cells that are inside the nozzle. These cells are shown in blue.

Refer to the troubleshooting section for further information on analyzing cell types.

Incompressible Flows Enclosed by Walls

To drive residuals to convergence for incompressible flows that are enclosed by walls, activate an additional source term for the pressure correction equation:

- Select the Continua > [Physics Continuum] > Models > Segregated Flow > Conservation Option node and set Option to Source Term.



This source term is required for such cases because the overset mesh method is not conservative and comes with an interpolation error that prevents residuals from converging. STAR-CCM+ automatically calculates the source term.

Caution: Use this option only for incompressible flows that are enclosed by walls. Activating this option in other situations can cause non-physical solutions.

Which Hole-Cutting Approach should you use?

When you create an overset interface, you can first use the default layered approach. This approach is computationally more efficient than the global approach and suitable for most cases. If the overset interface creation gives

overset intersector errors, you can activate the Alternate Hole Cutting to use the global hole-cutting approach (see Overset Mesh Interface Properties). The global approach is more robust when you run the simulation on many cores. The global approach is more sensitive to non-standard overset topologies where the overset boundary does not form a closed surface (see Overset Topology).

Overset Mesh Methodology

To establish the connectivity between the background and the overset regions, a two-step Overset Assembly process takes place:

1. Hole-cutting, which determines which cells are active, inactive, or acceptor cells.
2. Donor Search, which ensures that donor cells are found for each acceptor cell.

Hole-cutting process

The hole-cutting process is part of the coupling of the overset region with the background region through an overset interface. A successful coupling by use of an overset interface results in a "hole" being cut in the background mesh - the hole-cutting process. The hole-cutting process determines whether cells are active or inactive in the coupled simulation. Currently, STAR-CCM+ provides two methods for hole-cutting:

- Layered approach
- Global approach

Layered Approach

Along the boundary of the overset region, a layer of cells is identified that forms the donor cell layer. The cells in the background region in immediate vicinity to the donor cell layer become the acceptor cells. These layers of cells must form a watertight container around the overset region. The cells of the background mesh that are completely covered by the overset region cells inside this closed container become inactive.

Global Approach

For each cell of the background region, the algorithm verifies whether the cell centroid of that cell is outside or inside the overset region. If the cell centroid of a background region cell is inside the overset region, it becomes inactive.

Donor Search

For each acceptor cell, donor cells must be found. The set of donor cells depends on the interpolation option that is chosen and on the number of active cells in the donor region around the acceptor cell centroid.

The interpolation options are:

- Distance-weighted interpolation, where the interpolation factors are inversely proportional to the distance from acceptor to donor cell center, resulting in the closest cell giving the largest contribution. This involves 3 donor cells (in 2D) or 4 (in 3D).
- Linear interpolation using shape functions spanning a triangle (in 2D) or a tetrahedron (in 3D) defined by centroids of the donor cells.

  This option is more accurate but also more expensive. It is important in simulations involving moving meshes as it ensures that interpolation elements do not overlap. The choice of donor cells is not unique since the available donor cell centroids that enclose the acceptor cell centroid can define more than one triangle or tetrahedron. This way the interpolation is continuous as the acceptor cell centroid passes from one interpolation element to the next.
- Least squares interpolation.

  For a description of the least squares interpolation method, see Least Squares Interpolation.

The interpolation function is built directly into the coefficient matrix of the algebraic equation system. This approach ensures implicit coupling of the overset meshes. In the image below, two acceptor cells are shown using dashed lines, one in the background mesh and one in the overset mesh.



The fluxes through the cell face between the last active cell and the acceptor cell are approximated in the same way as between two active cells. However, whenever the variable value at the acceptor cell centroid (marked by the open symbols in the above figure) is referenced, the weighted variable values at the donor cells are substituted:

$$\phi_{\text{acceptor}} = \sum \alpha_i \phi_i \tag{6}$$

In this equation, α_i are the interpolation weighting factors, ϕ_i are the values of the dependent variable ϕ at donor cells N_i and subscript i runs over all donor nodes of an interpolation element (denoted by the green triangles in the figure). This way, the algebraic equation for the cell "C" in the above figure involves three neighbor cells from the same mesh (N_1 to N_3) and three cells from the overlapping mesh (N_4 to N_6). The coefficient matrix of each equation solved (both for the segregated and coupled solution method) is updated accordingly to ensure that equations can be solved up to the round-off level of residuals.

Overset Mesh Field Function Reference

The following primitive field functions are made available to the simulation with overset mesh. Both the volume mesh as well as the overset mesh representation can be used to display these field functions. The volume mesh representation displays only the active cells. If you want to see the acceptor cells, use the overset mesh representation. These field functions have a typical set of base properties.

ChimeraGridTestData	Used to check the interpolation error at the Overset Mesh Interface.
Overset Cell Status	Used internally as a flag in the solvers to identify the status of each cell during the solution process of the equations. This function provides more information about the result of the Overset Assembly process. The available values are: • 1 - inactive • 0 - active • -2 - acceptor • -1 - acceptor that is used as donor

Overset Cell Type	Used by the Overset Assembly process to identify the status of each cell. It gives more information about the cells with regard to the Overset Assembly process. Can be used to analyze errors. The available values are: • 0 - active • -1 - inactive cell that resides outside of the background mesh • -2 - inactive cell that is covered by another overset region. Consequently, this type can only occur for multiple overset meshes. • 1 - donor • 2 - active, intermediate cell layer used by the hole cutting process. • 3 - acceptor
Overset Error Status	Indicates the status of each cell that has an error. The available values are: • 0 - cell has no error • 5 - cell is an acceptor cell that has inactive donor cells. • 6 - cell is an acceptor cell that has zero donor cells. • 20 - cell is a donor cell, but is inactive. • 30 - an inactive face: that is, a face between an active cell and an inactive cell.

Troubleshooting Overset Mesh

Several troubleshooting features are available for analyzing the quality of the overset mesh.

You can generate the overset mesh without running the simulation by initializing the Overset Mesh Interface:

- Right-click the Interfaces node and select Initialize.



If the background and overset regions were not successfully coupled, the following error message appears.



You can check the cell status of the overset mesh by creating an Overset Cell Status scalar field. For a list of overset cell status values, see Overset Mesh Field Function Reference.

- Set the scalar mesh representation to Overset Mesh.



Check each region separately. It is recommended that you manually change the scalar bar range from -2.0 to 1.0. In this case, all active cells have a value of 0, inactive cells have a value of 1, and acceptor cells have a value of -2. Acceptor cells must separate the active cells from inactive cells; direct contact between active and inactive cells is not allowed.

Below is the scalar field for the background region. The overset region cells have been deactivated.



The overset region only includes mesh cells for which the governing equations are solved.



All other scalar field options available for mesh analysis can be found under field function reference.

Modeling Space, Time, and Motion

This section describes the following models in STAR-CCM+:

- Space
- Time
- Motion

Modeling Space

The primary function of the Space models in STAR-CCM+ is to provide methods for computing and accessing mesh metrics. Examples of mesh metrics include cell volume and centroid, face area and centroid, cell and face indexes, and skewness angle.

This section describes the four Space models in STAR-CCM+ and how to select one:

- The Axisymmetric Model
- The Shell Three-Dimensional Model
- The Two-Dimensional Model
- The Three-Dimensional Model

Reference information is also provided for field functions.

What Is an Axisymmetric Model?

The Axisymmetric model is designed to work on two-dimensional axisymmetric meshes. When using the model, orient the mesh such that the axis of rotation is at y = 0 in global coordinate space. No part of the mesh can be below y = 0, and the boundary edge that lies along the axis must be of type Axis.

For boundary conditions and reporting purposes, the mesh is assumed to be swept through an angle of 1 radian. For applications that use a mass flow inlet, the mass flow is therefore given in *kg/rad s*. Any volumetric or area quantities reported for the Axisymmetric model are assumed to be for a 1-radian sector. If you want the value of a quantity for the full 360^0 revolution, multiply reported quantities by 2π radians.

What Is a Shell Three-Dimensional Model?

The Shell Three-Dimensional model represents a volume in space whose thickness is sufficiently small that it can be considered as a surface with a single-cell thickness in the normal direction. This model must be associated with a shell region, and is required when:

- Modeling heat conduction within a solid shell; and,
- Modeling fluid films.

What Is a Two-Dimensional Model?

The Two-Dimensional model is designed to work on two-dimensional meshes. Activated this model only if the mesh is indeed two-dimensional. In this model, the mesh is assumed to have a unit depth (in SI units) so that any volumetric or area quantities reported for the two-dimensional model are assumed to be "per meter".

What Is a Three-Dimensional Model?

The Three-Dimensional model is designed to work on three-dimensional meshes. Activate this model only if the mesh is indeed three-dimensional. The use of a one-cell-thick three-dimensional mesh is much less efficient than using a true two-dimensional mesh for two-dimensional and axisymmetric simulations. A better approach is to extract a two-dimensional mesh.

Selecting a Space Model

To select a Space model, open the Continua node of the simulation tree, then right-click the Physics 1 node. Select the Select models... item from the pop-up menu.



This command opens the *Model Selection* dialog. Select the correct Space model for the simulation by clicking the relevant radio button.



The following model selection dialog shows the Space model activated.



What Is a Shell Region?

The Shell Three-Dimensional model is used with shell regions. Shell regions, which have properties and a pop-up menu, are created from existing surfaces (boundaries) contained within regular regions.



It is not possible to create a shell region directly from geometry parts or baffles. However, you can use the Insert Shell Region Interface option to create a shell region at an in-place (baffle) interface.

Shell Region Properties

Index	Index reference for shell region (read-only).
Continuum	Defines the fluid film phase or physics continuum that is associated with this region. Only physics continua where the Shell Three-Dimensional space model is selected is shown in the drop-down menu.
Source Boundaries	The boundaries that are used to create this shell region (read-only).
Parts	List of parts that are associated with this region. Editing the property accesses the in-place object selector to add or remove parts.
Type	Defines the type of the shell region depending on the region from which it was created: fluid, solid, or porous.

Shell Region Pop-Up Menu

Transform	Scales, translates, or rotates the selected regions and all associated boundaries. These actions operate on one or multiple regions, and if more than one region is selected, the transform action is applied to each region in sequence. Scale... Executes the scale region action. Translate... Executes the translate region action. Rotate... Executes the rotate region action.
Reorder Mesh	Reorders the internal references to cells in the mesh region in order to optimize the solver performance in terms of memory usage.

Solution	Applies operations in the submenu to the individual region.	
	Clear	Executes the clear region solution action.

See also: common menu items.

Creating a Shell Region

Shell regions are created from existing region boundaries.

- Right-click the region boundary or boundaries from which you wish to create a shell region and choose Create Shell Region from the pop-up menu.



Several objects are added to the simulation as a result of this operation:

- A single shell region is created when multi-selecting the contiguous boundaries belonging to a single region (as is shown in the preceding example). If the selected boundaries belong to different regions, several shell regions are created. These shell regions must be connected with interfaces.
- Interfaces between the surface of the shell region and the boundaries from which it was created.
- The corresponding interface boundaries in both the source region and the shell region.

These objects are shown in the example below:



In the example that was used above, the following shell region was created.



After creating the shell region, you will need to set its Continuum property to point to either of the following items:

- A separate physics continuum that uses the Shell Three-Dimensional model. Use this property assignment when modeling heat conduction in solid shells.
- A fluid film phase.

Setting an Axis

An axis is used to define an axis of rotation when the simulation requires it. The Axis node, which represents the axis, appears as:

- A region value when a space model has been selected for the associated physics continuum, and

- A boundary value on a non-slip wall boundary when the tangential velocity specification condition is set to local rotation rate.

The properties of the node define the axis.



This axis definition provides information that is required when STAR-CCM+ creates periodic interfaces for a region, and for transformations involving these interfaces. The tangential velocity specification at a boundary also uses the axis definition to specify the local axis and origin of wall rotation.

Axis Properties

Direction	A Cartesian unit vector specifying the direction of the rotation.
Origin	Cartesian position vector specifying the origin of the axis.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.

Space Field Functions Reference

The following field functions are made available to the simulation when any one of the Space models is used. They include a typical set of base properties.

Cell Aspect Ratio

This field function is used to identify regular polygonal cells from skewed or stretched cells. A value of 1 indicates that the cell is a regular polygon, and a value close to zero would represent a highly stretched cell. This ratio is defined as:

$$CAR = \frac{ND \times NF \times max(V, 0)}{\sum_f |\boldsymbol{A}_f| \sum_f \left| \frac{\boldsymbol{A}_f \bullet \boldsymbol{dx}}{|\boldsymbol{A}_f|} \right|}$$

where ND is the number of dimensions (three for a 3D mesh and two for a 2D mesh). NF is the number of faces of the cell. V is the volume of the cell. $\boldsymbol{A}_f$ is the cell face area vector and $\boldsymbol{dx}$ is the position vector from the face centroid to the cell centroid.

Centroid

The centroid is a property of a cell (or face in 2D). It is the location vector of the center of mass of the volume (or shell).

Chevron Quality Indicator

This field function is used to identify chevron cells. These cells are pairs of thin slender cells that share a common face whose angle is such that the line joining the cell centers does not pass through the common face. This indicator function marks chevron cells as 1 and every other cell as 0. A chevron quality norm c, as defined below, is used to determine whether a cell can be classified as a chevron cell or not. The chevron quality norm is computed using a vector $\boldsymbol{dp}$ that is projected from the face center to the line joining the cell centers:

$$\boldsymbol{dp} = \left(\frac{\boldsymbol{A_f} \bullet \boldsymbol{dx}}{\boldsymbol{A_f} \bullet \boldsymbol{ds}}\right)\boldsymbol{ds} - \boldsymbol{dx}$$

$$c = 1 - \frac{\boldsymbol{dp} \bullet \boldsymbol{dp}}{max(\boldsymbol{dp} \bullet \boldsymbol{dv})}$$

where:

- $\boldsymbol{A_f}$ is the face area vector
- $\boldsymbol{dx}$ is the vector from the face center to the cell center
- $\boldsymbol{ds}$ is the vector connecting the two cell centers
- $\boldsymbol{dv}$ is the vector from the face vertex to the face center.

The *max* function, in evaluating the chevron quality norm, is taken over all the face vertices.

The chevron quality norm is:

- 1 for orthogonal cells in which the line connecting the two cell centers passes through the face centroid
- 0 when the line passes through an edge of the face
- negative when the line is outside the face.

The chevron quality indicator is marked 1 when the chevron quality norm is negative.

Two more properties are provided:

- Chevron Threshold Angle—the maximum angle included by the two cells, beyond which the chevron quality indicator is not computed and is set to 0. You specify this value.
- Cell Aspect Ratio Tolerance—if the Cell Aspect Ratios of both the cells in question are greater than this specified tolerance, then the chevron indicator is not computed and is set to 1. This tolerance prevents STAR-CCM+ from marking two irregular sized cells connected by a small face.

Position

The position vector is defined for every vertex in the model--it is the coordinate of the point. For a cell, the position is returned as the average of the positions of the vertices of the cell.

Modeling Time

The primary function of the Time models in STAR-CCM+ is to provide solvers that control the iteration and/or unsteady time-stepping.

This section describes the Time models in STAR-CCM+:

- Steady
- Harmonic Balance
- PISO Unsteady
- Implicit Unsteady
- Explicit Unsteady

It also explains the following:

- How to select a Time model
- How to set the physical time-step
- How to choose between steady and unsteady time models
- How to choose a time-step

What Is the Steady Model?

The Steady model is used for all steady-state calculations. When this model is activated, the concept of a physical time-step is meaningless. For those objects that offer a choice between Iteration or Time-Step in a trigger (such as Monitors or Scenes), choose only Iteration.

What Is the Harmonic Balance Model?

The Harmonic Balance time model is always used with the Harmonic Balance method which is designed for solving periodic unsteady problems, where the solution repeats itself with a known frequency.

What Is the PISO Unsteady Model?

The PISO Unsteady time model is for use with the PISO algorithm, in transient cases with small time-steps. The PISO time model is currently only compatible with cases using the standard K-Epsilon model, a single phase, and no combustion.

What Is the Implicit Unsteady Model?

The Implicit Unsteady model is the only unsteady model available with the Segregated Flow and Segregated Fluid Energy models. It uses the Implicit Unsteady solver.

With the Coupled Flow and Coupled Energy models, the implicit unsteady approach is the alternative to the Explicit Unsteady one. The choice between these two approaches is based on the time scales of the phenomena of interest.

When this model is activated, objects that offer a choice between Iteration or Time-Step for a trigger (such as Monitors or Scenes), can be set to update at each time-step.

What Is the Implicit Unsteady Solver?

The Implicit Unsteady solver becomes active when the Implicit Unsteady model is active. Its primary function is to control the update at each physical time for the calculation. It also controls the time-step size.

Properties

Time-Step

The physical time-step. In addition to a constant, you can enter an expression directly. The default Time-Step is 0.001 s.

Temporal Discretization

Selects the temporal discretization scheme. For guidance on scheme selection, see the discussion on the transient term.

1st-order

Selects the first-order temporal scheme. When this scheme is selected together with a rigid body, morphing, or user-defined vertex motion model, the discretized grid flux is of first order (see Eqn. (17)).

2nd-order

Selects the second-order temporal scheme. When this scheme is selected together with a rigid body, morphing, or user-defined vertex motion model, the discretized grid flux is of second order (see Eqn. (18)).

Expert Properties

Freeze Time

Controls whether to advance time or fix it at the current time-step.

Activated

Time does not advance.

Deactivated

Time advances. This option is the default setting.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. Activating this option is primarily a debugging tool that is not normally required for simulation work.

Activated

The solver does not update the solution.

Deactivated

The solver updates the solution.

What Is the Explicit Unsteady Model?

The Explicit Unsteady model is available only with the Coupled Energy

model. It is only compatible with the Inviscid and Laminar viscous regime models. It uses the Explicit Unsteady solver.

The explicit approach is the alternative to the Implicit Unsteady model. The choice between these two approaches is based on the time scales of the phenomena of interest.

When this model is activated, objects that offer a choice between Iteration or Time-Step for a trigger (such as Monitors or Scenes), can be set to update at each time-step.

What Is the Explicit Unsteady Solver?

The Explicit Unsteady solver becomes active when the Explicit Unsteady model is active. Its primary function is to control the iterative update for the calculation.

Properties

None.

Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating. Activating this option is primarily a debugging tool that is not normally required for simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Selecting a Time Model

Before or after selecting the Equation of State model and the Viscous Regime, select the relevant Time model radio button.



Setting the Time-Step for Implicit Unsteady Simulations

To set the time-step for implicit unsteady simulations, select the Implicit Unsteady solver under the Solvers node. The physical time-step can be set using the Time-Step property in *Properties* window. In addition to a constant, you can enter an expression, such as `10 * sin(3.14 * $Time)`, directly.



How Do I Choose Between Steady and Unsteady?

In most situations, the physics being simulated indicates the choice between the Steady and Unsteady (Implicit or Explicit) models. Simulations that are time-dependent, and hence require the Unsteady model include:

- Time-varying boundary conditions
- Sliding or deforming mesh problems
- Free surface (VOF) problems
- Transient heat transfer

Other simulations are separate cases, namely those using harmonic balance.

In some situations, however, unsteadiness in the physical problem is not readily apparent. It can manifest itself in a failure to converge when running with the Steady model. One example is the solution of flow past a cylinder. At certain Reynolds numbers, the flow is unsteady because of vortex shedding, and it is incorrect to use a steady solution approach.

When physical instabilities do exist, the correct way to model them is to perform a time-dependent simulation. Choose an appropriate time-step that is based on a representative Strouhal number. Accumulate time-averaged statistics only after the transient simulation has achieved a periodic state.

Using the Steady model for unsteady physical phenomena is not recommended. Running an unsteady problem with a Steady model is equivalent to using an inaccurate time-step in an unsteady simulation.

How Do I Choose a Time-Step?

The choice of time-step is an engineering judgment (just like grid refinement). Be sure to understand the time scales of the physical phenomenon. The convective time scale (CFL) is important if, for example, you are tracking the convective dispersion of a scalar through the flow. If molecular diffusion is dominant, consider the Von Neumann number, $\mu\Delta t/(\Delta x)^2$. For Large Eddy Simulation (LES), consider the turbulent time scale. For vortex shedding (for example, with Detached Eddy Simulation, or DES) consider the Strouhal number.

Modeling Motion

In engineering, motion can be defined as the change in location of a body relative to a particular frame of reference. Such motion is imparted to any fluid in contact with the surfaces of the body, and must be considered in an analysis of the fluid behavior. Some examples of motion-induced flows are:

- Mixing vessels, where mixer blades drive the flow
- Turbomachinery
- Flexible diaphragms in medical applications
- Ship and propeller motion
- Fluid structure interaction

In STAR-CCM+, it is also possible to model the motion of rigid bodies in six degrees of freedom with the fluid system.

Understanding Categories of Motion

There are three broad categories for defining motion in STAR-CCM+:

1. The first category contains those methods that involve actual displacement of mesh vertices in real time: use these methods with a transient analysis.
2. The second category contains the method whereby an unsteady motion problem is converted to a steady-state problem by imposing a moving frame of reference on a static mesh. Rigid rotation and/or translation can be simulated in this way.
3. The third type of motion is a corollary of the real-time displacement, and applies only to harmonic balance flutter. The imposed modal displacement of the flutter boundary is applied to the frequency domain solver.

The specific methods available in each category are listed as follows:

Category 1: Mesh Displacement in Real Time

Several methods are available in this category:

- Rigid Motion
  - Rotation
  - Translation
  - Rotation and Translation
- Morphing (deforming mesh)

- Dynamic Fluid Body Interaction
  - DFBI Rotation and Translation
  - DFBI Embedded Rotation
  - DFBI Morphing
  - DFBI Superposed Rotation
- Solid Displacement
- User-Defined Vertex Motion

The space conservation law is enforced in non-rigid body motion types by solving for a volume conservation equation.

Category 2: Moving Reference Frame in Steady-State

The reference frame types that can be used in this category are:

- Rotating
- Rotating and Translating

Category 3: Harmonic Balance Flutter

There is only one option for this approach.

What Is a Motion?

In STAR-CCM+, motion is defined using motion objects. These objects are contained within the Motions manager node in the main simulation tree and encapsulate the different methods available for introducing motion. Setting

a region to use a motion object other than Stationary results in movement of the mesh vertices during the simulation. This manager node uses a pop-up menu to add motion objects.



As motion is defined relative to a reference frame, STAR-CCM+ provides reference frame objects to complete the motion specification for a region. These objects are contained within the Reference Frames manager node in the main simulation tree.

While the types of motion objects differ, they are all subject to common procedures for working with motions.

The currently available motions are:

- Stationary
- Rotation
- Translation
- Rotation and Translation
- Morphing
- DFBI Rotation and Translation
- DFBI Embedded Rotation
- DFBI Morphing
- DFBI Superposed Rotation
- Harmonic Balance Flutter
- Solid Displacement
- User-Defined Vertex

You can add more than one of any types other than Stationary. This feature is helpful when you want to retain more than one set of property specifications for a particular motion.



If there is more than one motion of a particular type in the simulation, then the second and subsequent ones have numbers at the end of the labels. (Assuming they have not been renamed).

Motions can be copied, renamed, or deleted.

What Is the Stationary Motion?

The Stationary motion is used for simulations that do not involve any kind of motion. The Stationary motion node, which is subject to common menu items, is the default motion in a STAR-CCM+ simulation.



Stationary is the only type of motion present in the motions manager when a simulation is created, and cannot be deleted. It is not necessary to create more motions of this type.

What Is a Solid Displacement Motion?

Solid displacement motion, also referred to as the large displacement model, is used with the solid stress model to allow the solid mesh to be geometrically updated in response to solid stress displacements. It is useful for cases involving fluid structure interaction. The Solid Displacement motion node represents this motion and is subject to common menu items.



To use solid displacement in a simulation, activate the Solid Stress continuum model, add this motion, and then select it for a region in the Motion Specification region value.

Pop-Up Menu of the Motions Manager

The following are the pop-up menu items available when you right-click the Motions manager node.

New	Creates a motion.	
	Rotation	Adds a Rotation motion to the Motions node.
	Translation	Adds a Translation motion to the Motions node.
	Rotation And Translation	Adds a Rotation and Translation motion to the Motions node.

Morphing	Adds a new Morphing motion to the Motions node.
DFBI Rotation and Translation	Adds a DFBI Rotation and Translation motion to the Motions node.
DFBI Embedded Rotation	Adds a DFBI Embedded Rotation motion to the Motions node.
DFBI Morphing	Adds a DFBI Morphing motion to the Motions node.
DFBI Superposed Rotation	Adds a DFBI Superposed Rotation motion to the Motions node.
Harmonic Balance Flutter	Adds a Harmonic Balance Flutter motion to the Motions node.
Solid Displacement	Adds a new solid stress Solid Displacement motion to the Motions node.
User-Defined Vertex	Adds a User-Defined Vertex motion to the Motions node.

See also: common menu items.

Working with Motions

Motion objects are applied to regions using the Motion Specification node that can be found in the Physics Values manager node for a region. By default this node is set to use the Stationary motion with the Lab reference frame. To add more motions and set them for use on a region, use the following procedure:

- Open the Tools node. Add a motion to the Motions manager by right-clicking the node and choosing New and then the appropriate type of motion.
- Depending on the type of motion you create, STAR-CCM+ automatically adds a read-only reference frame to the simulation tree. If you require an alternative reference frame, or to add a reference frame without first adding a motion:
  - right-click the Tools > Reference Frames manager node
  - choose New and then the appropriate type of reference frame.
- Having created the necessary motion and/or reference frame, open the Physics Values node of the region, click the Motion Specification node, and select the desired motion and reference frame in the *Properties* window.
- Depending on the motion that is specified, set the Axis or control point list region values.

Currently it is not possible to use non-trivial combinations of a motion and a reference frame. If a motion is selected, the reference frame is automatically set to Lab reference frame. On the other hand, if a reference frame other than Lab is selected, the motion is automatically set to Stationary in that reference frame.

Setting the Motion Specification for the Region

Motion is applied on a region basis in STAR-CCM+. Each region has a Motion Specification node that defines the motion and reference frame that is applied to the region. The default motion is the stationary motion with the Lab Reference Frame. The governing equations for that region are therefore solved within the Lab Reference Frame. To access the properties of the motion specification, select the Motion Specification node.



Motion Specification Properties

Motion	Selects one of the existing motions in the simulation from a drop-down list.
Reference Frame	Selects one of the existing reference frames in the simulation from a drop-down list.

Adding Motions

To create a motion (other than Stationary), right-click the Motions node and select the motion type that you wish to create from the New pop-up menu.



Copying Motions

A motion can be copied. When a motion is copied, all properties are copied, resulting in an exact duplicate of the original.

You can also copy the properties of one motion onto another, provided they are of the same type.

Copying Properties of a Motion

It is possible to copy properties among similar motions of all categories except for motions that do not have properties.

Among motions of the same type, this operation applies all of the properties of the source motion.

It is not possible to copy properties between different types of motions.

Renaming Motions

Any of the motions can be renamed using the standard mechanisms, but be careful that the name you apply still reflects the type of motion. If it does not (as in the example below), identify the type of the motion from the description area at the bottom of the *Properties* window.



Deleting Motions

Motions other than Stationary can be deleted using the standard methods, unless a region uses the motion. In this case the Delete pop-up menu item is grayed and the `<Delete>` key has no effect (no warning is issued).

Simulating Constant Rigid Motion

Examples of systems involving rigid motion are mixers, turbines, compressors, pumps, and other devices where an internal or external component rotates about a fixed axis, or translates along a fixed path. It is assumed that the moving component does not deform.

For an approximate analysis of a constant rigid motion, it is possible to apply a moving reference frame to the region that contains the moving part, and to solve the problem in steady-state. This approach gives a solution that represents the time-averaged behavior of the flow, rather than the time-accurate behavior. The section, Using Moving Reference Frames in Steady-State discusses this matter more fully.

When the time-accurate behavior is required, STAR-CCM+ provides two motion options for rigid motion. Applying either of these motions results in mesh vertices being displaced in each time step. Details are given in the section, Using Rigid Motion.

Using Moving Reference Frames in Steady-State

Applying a moving reference frame to a region generates a constant grid flux in the appropriate conservation equations. This grid flux is calculated based on the properties of the reference frame rather than the local motion of cell vertices. A moving reference frame does not change the position of cell vertices.

For the analysis of turbomachinery flows, moving reference frames can be used for two types of steady-state simulation: frozen rotor and mixing plane. In a frozen rotor simulation, a rotating reference frame is applied to the region of mesh containing the rotor blades. The remainder of the mesh is either static or has a different rotating reference frame applied (as multiple regions can have different reference frames). As the rotor blades are frozen in space with respect to the stator blades, the flow field that the stator blades see depends on the relative position of the stages. Different relative positions result in different solutions, which is a major disadvantage of this approach.

In a mixing plane analysis, the problem of relative positioning is overcome by using an indirect mixing plane interface to couple the rotor and stator regions. Flow conditions on either side of the mixing plane interface are averaged circumferentially to provide a single radial profile that is applied to all cell faces at the interface. The basis for this treatment can be appreciated by considering what would be observed when looking at a fast rotating fan: that is, a single rotating surface emitting a constant stream of air independent of blade position. Increasing the number of blades on the fan, increasing the rotational speed, or simply moving farther away from the fan, would all tend to increase the accuracy of this interpretation. The following figure illustrates a rotor and stator assembly for which a mixing

plane simulation would be appropriate. If periodic interfaces are applied mid-way between the blades, it is only necessary to include a single blade passage for each stage.



Place the mixing plane interface at a point where the time-averaged properties rather than instantaneous flow details are likely to influence the flow.

The mixing plane approach has the advantage that the relative positioning of the stator and rotor is irrelevant. However, it suffers from the fact that wake effects are lost due to the circumferential averaging. If it becomes necessary to model the interaction between stages in greater detail, STAR-CCM+ provides two further options: rigid motion of mesh vertices in a transient analysis, or a harmonic balance analysis of a single stage with a prescribed wake.

Using Rigid Motion of Mesh Vertices

For cases where the motion of an entire region is considered, STAR-CCM+ provides the option to move the mesh vertices of a region during a transient analysis. This mesh motion is specified using either a Rotation motion or a Rotation and Translation motion. For the Rotation motion, it is possible to superpose additional rigid rotations on top of the original rigid rotation. This superposition permits rotations around multiple axes to take place at the same time.

What Is a Rotation Motion?

A Rotation motion is used to move the mesh vertices of a region around a specified axis. The Rotation node represents this motion and has properties and is subject to common menu items. The Superposing Motions child node has its own pop-up menu that is used to add more rigid rotations on top of the original motion.



Adding this motion results in a read-only Rotating reference frame node being added to the Reference Frames manager. This node reflects the properties that are entered for the Rotation motion.

Rotation Motion Properties

The Axis Direction, Axis Origin, and the Rotation Rate scalar property, can be specified in the *Properties* window. In addition to constants, you can enter field function syntax directly. You can select a local coordinate system to define the axis and origin. Also, the Managed Coordinate System property lets you select the coordinate system that the motion moves.

Axis Direction	The Cartesian direction vector specifying the axis of rotation.
Axis Origin	The Cartesian position vector specifying the origin of the axis of rotation.
Rotation Rate	The rotation rate around the axis that with the specified Direction and Origin.

Managed Coordinate Systems	Specifies coordinate systems that the motion is to move. The *Select Objects* dialog of this property includes any local coordinate systems that exist in the simulation.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.

See also: define value action.

Superposing Additional Rigid Rotations

The Superposing Motions node allows you to add more rigid rotations to Translation and Rotation motions.

To add a motion, right-click the node and select New from the pop-up menu.



A Superposing Rotation node is added beneath the Superposing Motions node. This Superposing Rotation node requires the same set of properties as the Rotation node, and contains its own Superposing Motions child node for further rotations.



By successively adding superposing rotations in this way, it is possible to construct complex rotating motions that apply to the whole mesh in the regions for which the original rotation motion is active. For example, pitch, yaw and roll for a flying object can be modeled.



The local coordinate system in which the superposing rotation is defined must be a nested coordinate system whose parent is the coordinate system in which the parent rotation is defined. An example of such coordinate systems is shown below. The rotor motion is defined in the Rotor-csys coordinate system. Coordinate systems blade1-csys and blade2-csys, define

the superposed motions blade1-rotation and blade2-rotation, respectively. These two latter coordinate systems are children of the Rotor-csys coordinate system.



What Is a Translation Motion?

A Translation motion is used to move the mesh vertices of a region with a specified velocity. The Translation node represents this motion. This node has properties and is subject to common menu items. The Superposing Motions child node has its own pop-up menu that is used to add more rigid rotations on top of the original motion.



Translation Motion Properties

Translation Velocity	The translation velocity of the region.
Managed Coordinate Systems	Specifies the coordinate system that uses the motion. The *Select Objects* dialog of this property includes any local coordinate systems that exist in the simulation.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the velocity. The drop-down list of this property includes any local coordinate systems that exist in the simulation.

What Is a Rotation and Translation Motion?

This motion is used to rotate and translate the mesh vertices of a region. The Rotation And Translation node represents this motion. This node has properties and is subject to common menu items.



Adding this motion results in a Rotating and Translating reference frame node being added to the Reference Frames manager.

Rotation and Translation Motion Properties

In addition to Rotation motion properties, this motion lets you specify translation velocity.

Axis Direction	The Cartesian direction vector specifying the axis of rotation.
Axis Origin	The Cartesian position vector specifying the origin of the axis of rotation.
Rotation Rate	The rotation rate around the axis that with the specified Direction and Origin. This option is additive with Translation Velocity.
Translation Velocity	The translation velocity of the region. This property is additive with Rotation Rate.
Managed Coordinate Systems	Specifies the coordinate system that uses the motion. The *Select Objects* dialog of this property includes any local coordinate systems that exist in the simulation.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.

See also: define value action.

Working with Morphing

The Morphing motion in STAR-CCM+ redistributes mesh vertices in response to the movement of control points. The mesh morpher uses control points and their associated displacements to generate an interpolation field

throughout the region which can then be used to displace the actual vertices of the mesh. Each control point has an associated distance vector which specifies the displacement of the point within a single time-step.

Specified Displacement, $\mathbf{d}_i$



Control points can be initialized using existing mesh vertices on a boundary, or can be specified independently of the mesh using a Control Point List in the appropriate region. For control points initialized using mesh vertices on a boundary, the displacement can be specified directly, or by using:

- a grid velocity,
- the displacement field of a finite volume stress calculation,
- a rigid body motion calculation,
- a rigid body motion with a superimposed incremental displacement, or
- the results of an external calculation using the co-simulation feature.

When control points are collected from a boundary, it is not always necessary to create one control point per mesh vertex. A lower ratio can work equally well, and reduce the calculation time required. An automatic thinning algorithm is provided. This algorithm uses an estimated deformation of the mesh to reduce (or thin-out) the number of control vertices. Manual control is also provided through Thin Factor values on morphed boundaries. A Thin Factor of 1 means that all boundary mesh vertices are used, and a factor less than one results in a proportional reduction.

The node, Morphing, is subject to common menu items and has one property. When it is selected in the Motion Specification node for a region, this motion activates a solver and a Morpher boundary condition. In turn the Morpher boundary condition accesses a selection of boundary values that control the morphing process.



Morphing Model Properties

Morphing Order

The morpher can operate in two different modes: one region at a time, or one motion at a time. This property selects the mode to use.

Regionwise

Each region that is associated with this motion is morphed independently, one region at a time. Use this option, for example, in a fluid-solid case, where the solid region has a solid displacement motion and the fluid region has a morphing motion.

Motionwise

All regions that are associated with this motion are morphed at the same time. This method is recommended for use in cases where internal interfaces are present between regions, if:

- the floating morpher boundary method is applied to the boundaries on either side of the interface, and;
- both regions are set to use the same morphing motion.

Linear Fitter

This option extracts a linear transformation from the prescribed motion, and uses this transformation to move vertices. The main multiquadric procedure is then used to treat any remaining motion. If the linear transformation accounts for most of the motion, activate the Morph From Zero property on the Morpher solver.

Activated

Linear motion is treated separately from multiquadric interpolation.

Deactivated

All motion is treated in multiquadric interpolation.

Morphing Model Expert Properties

Automatic Thin-out

Applies an additional stage of control vertex thinning after the user-defined thin factors have been applied. Tick the Morph from Zero property on the Morpher solver when using this option.

Activated

Automated thinning is applied.

Deactivated	No automated thinning is applied. Only user-defined thin factors are accounted for.
Automatic Thin-out Cl	This property controls the target spacing between control points. If the value is smaller, the spacing is smaller resulting in denser control point distribution. If the value is larger, it coarser control point distribution.

In cases with multiple regions, if a given region has no specified control points, the morphing solver is skipped in that region.

Note: In some cases involving interfaces, STAR-CCM+ must intersect the interfaces at the beginning of each time step. This reintersection adds to the computational cost. Further information is given in the section, Working with Interfaces and Morphing.

There is nothing to prevent morphing from being used in 2D and 2D axisymmetric cases. The axis boundary type is treated in the same way as a symmetry boundary type; the in-plane condition is applied to them both.

Alternative frames of reference can be used with the morphing motion. Choose alternative frames of reference in the Reference Frame property of the Motion Specification belonging to the region being morphed.

Understanding the Morphing Procedure

Control points and their associated displacements form the underlying data that the morpher requires to generate an *interpolation field* (see Eqn. (11) in the formulation section). This interpolation field can be considered to act as a field function that is called for every mesh vertex in the region to find the new position for the vertex. A high-level view of the morphing procedure is therefore:

1. The morpher collects control points and their known displacements from boundaries and control point tables within the region. User-specified thin factors are applied. If specified, the linear fitter is applied to update the current location of the vertices.
2. Automatic thinning (if activated) is applied to reduce the number of control points. Automatic thinning is activated by default for new simulations that are created using STAR-CCM+ 6.02 onwards. It is deactivated for simulations that are created in earlier releases, and must be turned on using the Automatic Thin-out property on the Morpher node.

3. After some initial preparation (based on Eqn. (12)), these control points and their known displacements are used to generate the interpolation field (by solving equations (7) and (10)).
4. The interpolation field (Eqn. (11)) is applied to all mesh vertices in the region.
5. Some final adjustments are made to mesh vertices on or near boundaries.

Control points can be specified using vertices that lie on mesh boundaries, and can also be specified using points that are read from a table. These two approaches are discussed below.

Control Points Taken from Mesh Boundaries

When the morpher motion is active for a region, every boundary (apart from boundaries of type symmetry plane) within the region receives a Morpher node in its Physics Conditions folder. Additional nodes for morphing appear in the Physics Values folders, depending on the setting that is chosen for the Morpher node. Boundaries of type symmetry plane have a single morphing-related node in their Physics Values folders. The following screenshot shows these nodes for a wall boundary that is set to displace by a specified amount.



The Morpher node in the Physics Conditions node has a Method property that provides several different ways in which boundaries can interact with the morphing process. Deciding which method to use for each boundary is an

important step in setting up a simulation for morphing. The section Understanding Morpher Boundary Methods has been provided to explain what the different methods are.

For the following methods, the morpher collects Control Points from the boundary surface mesh: Fixed, Displacement, Grid Velocity, Solid Stress, Rigid Body, Co-Simulation, Motion, and Slide on Guide Surface. It is not necessary for the morpher to create a Control Point for every vertex on the boundary surface; a representative subset of mesh vertices is often sufficient. You control the proportion of mesh vertices that is used to create Control Points through the Morpher: ThinFactor node in the Physics Values node of those boundaries from which Control Points are taken. The exception is the Slide on Guide Surface method, which selects the thinning factor automatically

Automatic Thin-Out

When automatic thin-out is activated on the morpher motion, the morpher thins out input vertices automatically to reduce the morpher execution time. The automatic thinner attempts to thin out control vertices by estimating the deformation of the grid due to morphing and placing control vertices efficiently according to this information. This option is active by default for new simulations that are created with STAR-CCM+ 6.02 onwards. It is deactivated when restarting from simulations that are created in earlier versions of STAR-CCM+.

An automatic thin-out factor, Cl, is provided to control the strength of the automatic thinning algorithm. Control vertices are thinned more aggressively as the factor increases. The suggested range of values is from 0.5 (mild thinning) to 5.0 (strong thinning).

Studies have shown that automatic thin-out works best when the Morph From Zero property is activated on the Morpher solver. Further improvements are possible when the Linear Fitter is also activated, as the linear fitter separates any rigid body motion of mesh vertices from their deformation motion.

Linear Fitter Option

The basic idea of the linear fitter is to use the control point data to make a least square fit to a linear transformation between new and old coordinates. Motion due to this linear fit is subtracted from the overall motion during multiquadric interpolation, and added back in afterwards to get the final motion.

A linear transformation includes all the rigid body modes:

- 3 degrees of freedom (DOF) for translation,

- 3 DOF for rotation,

and stretching modes:

- 6 DOF for stretching (that is, the strain in the fluid is a constant).

Therefore the linear fitter solves for a total of 12 unknown degrees of freedom.

Having deduced the linear transformation, the multiquadric interpolation is then used only for the portion of the overall motion that the linear transformation does not cover. The complete motion of the control points is the sum of the linear transformation of the points from their linear location, plus the deformation that is propagated from the multiquadric interpolation.

For example, if the control points were undergoing pure rigid body motion, then the fit to the linear transformation would be exact. The difference between actual motion and the linear transformation would be zero. In this case, the morpher would move the vertices exactly as if you had selected a rigid body motion model, and there would be no morphing due to deformation.

As another example, consider a case where the control points are on a body which is slightly flexible but which has predominately rigid body modes. The linear fit for this case is not exact, but the differences between that and the full motion is small. These small variations would be interpolated to the moving vertex and then the overall rigid body motion (the linear fit) would be added back in for the total morph.

This linear fit option is therefore suited for cases where the rigid body motion usually governs the fluid grid motion.

Thinned Vertex Corrector

When the Vertex Thin Factor on a boundary is set to less than 1, the boundary vertices excluded due to thinning are initially moved according to the interpolation field in step 3 above. It is possible, however, that the final position of these thinned vertices is not precisely the position they would have obtained had they been moved according to the specified displacement on the boundary. A small error could result for other reasons even if vertex thinning is not used. To correct this error, a thinned vertex corrector algorithm is executed within step 4 of the morpher process. The corrector algorithm ensures that the thinned vertices receive the same displacement as the control vertices taken from their respective boundaries.

Control Points Taken from a Table

It is possible to define control points and their displacements directly in space for a morphing analysis. These control points are defined on the Control Point List node that are contained within the Physics Values folder of a region.



For further information, see the section on Using a Control Point List.

Preparing for a Morpher Analysis

The computational resource that is required to run the morpher is directly related to the number of control points present in the regions for which the morpher is active. For this reason, keep the number of control points as low as possible.

1. If the morphing only affects a portion of the mesh, consider creating a separate region for the portion that is morphed. Having a separate region means that control points are only required in the region being morphed. In the example below, morphing is restricted to the small

plunger on the left, which means that the large vessel could be placed in its own region. The morphing motion would then only apply to the region containing the plunger.



2. In regions for which morphing is active, identify those boundaries that can be set to use the Fixed Plane or In Plane morpher boundary method. These boundary methods can be dealt with far more efficiently in the morpher than is possible with the fixed boundary method. Note, however, that these planar boundaries define an infinite plane that must not intersect the mesh at any other point. Further information is provided in the section, Understanding the Interaction of Control Planes and Control Points.
3. If a body is moving parallel to a planar surface, then apply the In Plane morpher boundary method to the planar surface.
4. The Vertex Thin Factor is set to 1.0 by default, but can be safely reduced to 0.25 in many cases. For boundaries that are at a distance from surfaces being actively morphed, a value of 0.1 would not be unreasonable. The automatic thin-out also helps to reduce the number of control points in cases where the Vertex Thin Factor has been set too high.

5. For the initial setup, do not reduce any vertex thin factors but rather activate the automatic thin-out algorithm on the Morpher motion node. This algorithm can provide sufficient thinning without user-specified intervention. Run a short trial case to see if the algorithm reduces the number of control vertices sufficiently.

Understanding Morpher Boundary Methods

You set how the morpher deals with boundaries using the Method property of the morpher boundary condition.

Fixed

The Fixed morpher boundary method supplies control points for use in calculating the interpolation field. In this case, however, the displacement that is associated with the control points is zero. For non-planar boundaries close to the zone of influence, keep the Vertex Thin Factor close to 1.0, as the shape of the Fixed boundary could otherwise be distorted. For planar boundaries far from the zone of influence, the Vertex Thin Factor can be reduced. In the example below, the Vertex Thin Factor on the walls of the funnel section has been set to `0.1`.



For a planar boundary where the infinite plane of the boundary surface does not intersect with the model geometry at any other point, use the Fixed Plane morpher boundary method instead.

Displacement

The Displacement method supplies control points whose displacements you set directly. The displacement is set regardless of time step, which means that if you change the time step size in a simulation, a corresponding change in the distance the control points moves is required.

Grid Velocity

The Grid Velocity method provides essentially the same function as the Displacement method, but takes a velocity specification rather than a distance. The resulting displacement is proportional to the time step used for the simulation.

Fixed Plane

Mesh vertices lying on a fixed plane boundary are not moved in response to the interpolation field calculated by the morpher. This boundary differs from the fixed boundary in that no control points are taken from the boundary this method is applied to. Ultimately this means that a more computationally efficient approach can be taken. Note, however, that applying this boundary condition results in the creation of an infinite plane that could potentially intersect other parts of the model geometry.

The outlet of the funnel section below would be a suitable boundary on which to apply the Fixed Plane method.



In Plane

Applying the In Plane method results in the creation of an infinite plane in the same manner as for the Fixed Plane method. For the In Plane method, however, mesh vertices are permitted to move within the plane of the boundary itself. This behavior is appropriate for simulations similar to the one shown below.



Floating

With the Floating method, mesh vertices on the boundary are not constrained. They move in response to the interpolation field created by the morpher. An example of where this morpher boundary type could be used is for the side walls of a piston undergoing one-dimensional expansion or contraction, as shown in the illustration below.



Solid Stress

This boundary method is intended for use with the finite volume stress model. It behaves in a similar manner to the Displacement boundary type, but takes distance vectors in each time step from the finite volume stress solver.

Rigid Body

This boundary method provides inputs that allow you to define a rigid body motion about several axes. This approach differs from the use of rigid motion objects in that it is only the boundaries that move according to the rigid motion, and not the entire mesh.

Co-Simulation

Apply this boundary method when using co-simulation with Abaqus for a case that involves mechanical coupling. Abaqus supplies the displacements on the boundary at each time step. If absolute values of displacement are supplied (that is, displacement relative to the initial position of the mesh at time zero), activate the Morph From Zero option of the Mesh Morpher solver.

Motion

With this boundary method, you construct a displacement that is composed of:

- a rigid motion that is calculated using the properties of a selected rigid motion object, and
- an incremental displacement that is specified as a vector profile.

The displacements permit the specification of some complex motions such as those that would exist for the blades on a helicopter. In the example that is shown below, a rotation motion sets the rigid rotation of a blade, and a field function sets the incremental displacement. All available rigid motions are referenced in the Rigid Motion boundary value.



Internal

This boundary method allows STAR-CCM+ internal models to add to the displacement. Previously the displacement was user-specified. This method is used by the Fluid Film Melting-Solidification model only.

Slide on Guide Surface

This boundary method allows mesh vertices on the boundary to slide on a guide surface. Three different options can be selected as guide surface for the boundary vertices: the initial boundary surface, a surface defined by an external file or a plane that contains the initial boundary.

When the guide surface is set to be a plane containing the initial boundary, it is possible to specify a motion for the guide surface. This option is not available when the guide surface is the initial boundary surface or an external file surface: in these cases the guide surface motion is set to the default Stationary.

As an example, the following images show how the Slide On Guide Surface boundary method can be used when simulating the fluid motion within a Wankel engine.

The first image shows a comparison between the original and the morphed geometry. The morpher method selected for boundary 1 is Slide on Guide Surface and the guide surface is imported through an external file. For boundary 2, the In Plane morpher method is selected.



A close up of the circled area is displayed in the following image. The morpher method set for boundary 3 is Slide on Guide Surface and the guide surface is selected to be a plane containing boundary 3. The motion specified for this guide surface is a superposing rotation.



Working with Interfaces and Morphing

It is important to apply the correct conditions at interfaces between the regions as failing to do so could result in the mesh being badly distorted. Two different scenarios can exist: internal interfaces and sliding interfaces.

Internal Interfaces Between Regions

For internal interfaces between regions, set the Morphing Order property of the Morpher motion node to Motionwise. Set the morpher boundary conditions on both sides of the interface to Floating if the interface is expected to morph in tandem with vertices surrounding the interface. This condition does not apply to fluid-solid cases where the solid region is using the solid displacements motion and the fluid region uses the morphing motion.

Deactivate the Recalculate Interfaces expert property on the Morpher Solver for internal interfaces between regions.

Sliding Interfaces

At such interfaces, the intention is for the mesh to move differently on the two sides, either to accommodate large motion, or to model a physical boundary. It is still necessary to ensure that the faces on the two sides lie close enough to be intersectable. Appropriate boundary conditions would

be the In Plane condition, or the Displacement condition where the normal component on both sides is identical. The other components could be specified differently on the two sides.

Tick the Recalculate Interfaces expert property on the Morpher Solver for sliding interfaces.

Understanding the Interaction of Control Planes and Control Points

Control planes (in the Fixed Plane and In Plane morpher boundary methods) work by damping out the displacements smoothly to zero over a specified zone of influence. For the Fixed Plane type, the damping applies to displacements in all directions; for the In Plane type, the damping only applies to normal displacements.

Control planes extend infinitely in each direction tangential to the plane. The infinite extension means that it is possible to prevent morphing in other parts by setting a plane that intersects the model. The morpher only considers control points above the plane (the outward direction of the surface normal). Any control points on the opposite side of the plane are ignored. The following diagram illustrates such a case.



(This behavior can be avoided by using the Slide on Guide Surface boundary method. Although computationally more costly, the constraint imposed by this condition does not extend beyond the boundary geometry).

When a mesh vertex is located near or on a control plane, its final displacement due to morphing is calculated by multiplying the predicted displacement from Eqn. (11) by a damping function. The damping function gives zero displacements (or zero normal displacements for an In Plane type) at the plane, and smooth displacements for points away from the plane. The Damping Factor sets the strength of the damping function. It has units of length, and defines the distance normal to the plane within which damping is applied. As a guide, set the Damping Factor to a length that is about a quarter of the distance from the control plane to other boundaries.

Two damping functions are provided to smooth vertices within the region the Damping Factor defines. These functions are identified in the formulation as C_0 *continuous*, Eqn. (13), and C_2 *continuous*, Eqn. (14). A plot is given of both functions.

Potential conflicts can occur at the intersection of boundaries that use control plane methods and boundaries that use methods which supply control points. An example would be the intersection of a boundary using the fixed plane method and a boundary using the displacement method. The prescribed displacement could potentially conflict with the fixed plane. For such cases, the morpher can be set to ignore control points that lie within a certain distance from the control plane. The distance is calculated as the Ignore Factor * Damping Factor. Both of these values are given as properties within the Morpher Plane Options boundary value.

Troubleshooting Negative Volume Cells

It is not uncommon to have reports of negative volume cells when setting up a morpher analysis. When negative volume cells are reported, the first step is to create a threshold derived part using the `Volume` field function for a maximum value of `0`.

- Right-click the Derived Parts node and choose New Part > Threshold...



- In the *Create Threshold* edit panel, do the following:
  - Ensure that the morphed region is selected in the *Input Parts* menu.
  - Choose `Volume` from the *Scalar* menu.
  - Set the *Extraction Mode* to All Below.
  - Leave the Threshold Value at its default setting of `0`.

An example of a panel that is set with these properties is shown below. Click Create to generate the threshold part.



The part is added to a scalar scene to show the location of negative volume cells.



Suggestions for Correcting the Model

1. If the negative volume cells are close to a moving body, increase the Vertex Thin Factor for the boundary on the body. If possible, split the moving body boundary and apply a larger Vertex Thin Factor just for the area near the negative volume cells. This arrangement allows you to use a lower value (0.25 or below) for the remainder of the boundary.
2. If the negative volume cells are close to a fixed boundary, follow the same procedure as in (1).
3. If the negative volume cells are close to a fixed plane boundary, check to see if compression normal to the boundary is causing the negative volumes. If so, increase the damping factor so that a greater number of cell layers are accounted for in the damping zone. If the negative volume cells are being stretched parallel to the boundary, change the boundary to the in-plane type instead.
4. If the negative volume cells are close to an in-plane boundary, increase the damping factor.

Morphing with Multiple Regions

To reduce computational costs, the domain can be divided into regions. In any given region, if there are no control points, the morpher is automatically turned off. Normally, the Fixed morpher boundary method creates control points. If, however, there are no other control points, these Fixed points are discarded and the morpher solution is skipped in that region.

Multiple regions can also be used in cases where grossly different behavior is desired in each region (for example, to preserve some mesh features like boundary extrusion layers). In such cases, the Interface Boundaries must be given identical boundary conditions on the two adjacent regions.

What Is the Mesh Morpher Solver?

The Mesh Morpher solver is designed to control the morphing motion calculations with specialized properties. It becomes available when a region has the Morphing motion selected.



Mesh Morpher Solver Properties

Verbosity

Provides additional output for this value while the simulation is running. This property is useful for debugging problems when they occur.

None	Leaves output at its default.
Low	Adds messages as to which region is morphed.
High	Generates precise equation information about the morphing process as the solver converges.

Morpher Tolerance

Non-dimensional parameter that is used to specify acceptable accuracy in morphing. This parameter is multiplied by the longest edge length of the bounding box enclosing the morphing region.

Initialization needs morphing

Selecting this option causes the morpher to run when the simulation is initialized. Use this option when the control point specifications at time zero define an initial deformation of the mesh.

Activated	The morpher runs at time zero.
Deactivated	The morpher does not run at time zero, but runs during the first time-step to apply deformations at time equal to Δt.

Mesh Morpher Solver Expert Properties

Temporary Storage Retained

Specifies whether to retain temporary storage at the end of the iteration, or not. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

Activated	Retains temporary storage.
Deactivated	Does not retain temporary storage.

Morph At Inner Iterations

Repeats mesh morphing at each inner iteration during the time-step. This action is only meaningful for cases where some boundary condition depends on the solution; it is off by default to reduce computational costs.

Activated	Repeats mesh morphing at each inner iteration.
Deactivated	Leaves off this option.

Morph From Zero

Begins each mesh morphing from the same (zero time) positions. This behavior is helpful for periodic body motions, to ensure that the mesh is also periodic. It is also recommended when automatic thinning is applied.

Activated	Activates this option.
Deactivated	Leaves off this option.

Ignore Gridflux Terms

When this option is activated, the mesh still moves but the "grid" flux is not included in the transport equations. This option provides more stability when marching towards a steady solution with an unsteady solver. It is also useful when you are simply adjusting the shape of the region without solving the transport equations.

Activated	Activates this option.
Deactivated	Leaves off this option.

Recompute Interfaces

When activated, this causes STAR-CCM+ to recompute all interfaces at each time-step. Use this option only if the morphing changes the mesh in different ways on opposite sides of the interfaces.

Activated	Activates this option.
Deactivated	Leaves off this option.

Skip Thinned Vertex Corrector

When this option is activated, it prevents STAR-CCM+ from attempting to correct thinned vertices.

Activated	Activates this option.

	Deactivated	Leaves off this option.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. Do not activate this feature as your first choice, as it is primarily a "debugging" tool and not supported for normal simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Using a Control Point List

The Control Point List is a region value, added by the Morphing motion or the Harmonic Balance Flutter motion. It has properties and is subject to common menu items.



The Control Point List can be used to specify control points and their associated displacements without reference to any boundary. The list is defined as a table containing six columns. The first three columns specify the cartesian components of the initial position of the control points. The second set of three columns specifies the cartesian components of the target position of the control points. An example fragment of such a table is given as follows:

x1	y1	z1	x2	y2	z2
0.55	0.06	0.06	0.56	0.06	0.06
0.55	0.06	0.12	0.56	0.06	0.12
0.55	0.06	0.18	0.56	0.06	0.18

This table is defining a line of three control points (spaced along the Z-axis) whose displacement is 0.01 units in the X direction. The initial position of each control point is set to `[x1,y1,z1]` at the beginning of each time step. The target position is set to `[x2,y2,z2]` in every time step. This behavior is illustrated below. The mesh vertices at the end of the simulation are displaced a considerable proportion, but the control points based on the table remain in their original location.

(1) Initial grid (before morphing)



(2) Grid after several time steps in which morphing was driven by the Control Point List

Control Point List Properties

Table	Selects a table from the Tools > Tables sub-nodes.
Table: Control-Point X0	Selects the column within the table for the initial *x*-coordinate of the control point.
Table: Control-Point Y0	Selects the column within the table for the initial *y*-coordinate of the control point.
Table: Control-Point Z0	Selects the column within the table for the initial *z*-coordinate of the control point.
Table: Control-Point X1	Selects the column within the table for the final *x*-coordinate of the control point.
Table: Control-Point Y1	Selects the column within the table for the final *y*-coordinate of the control point.
Table: Control-Point Z1	Selects the column within the table for the final *z*-coordinate of the control point.

Morpher Boundary Condition

This boundary condition node allows you to determine how the morphing for a boundary is specified. It is activated when the Morpher motion is activated.

Properties

Method Sets the morphing condition.

Fixed	The Fixed option adds the Morpher: ThinFactor node to the Physics Values manager. Points on this boundary are treated as control points with zero displacement.
In Plane	The In Plane option adds a Morpher Plane Options node to the Physics Values manager. Points on this planar boundary are allowed to slide in the plane. Using the in-plane option is much less expensive than using control points.
Displacement	The Displacement option adds Morpher: ThinFactor and Morpher: LinearDisplacement nodes to the Physics Values manager. Points on this boundary have specified displacement.
Grid Velocity	The Grid Velocity option adds Morpher: ThinFactor and Grid Velocity nodes to the Physics Values manager. Points on this boundary are moved using the specified velocity.
Floating	The Floating option does not add any nodes to the Physics Values manager. Points on this boundary have no boundary condition; they move wherever the other control points dictate.
Fixed Plane	The Fixed Plane option adds a Morpher Plane Options manager to the Physics Values node. Points on this planar boundary are fixed. Using this option is much less computationally expensive than using control points.
Solid Stress	The Solid Stress option adds the Morpher: ThinFactor manager to the Physics Values node. Displacements at this boundary are picked up from the internal solid stress solution.

Rigid Body	The Rigid Body option adds the following nodes to the Physics Values manager: • Morpher: AngularDisplacement • Morpher: AngularDisplacementAxis • Morpher: LinearDisplacement • Morpher: ThinFactor This boundary is rigid. A single rigid body displacement is specified for all the control points on this boundary.
Co-Simulation	Use the Co-Simulation option for simulations that involve mechanical coupling with an external CAE tool.
Motion	The Motion option adds the following nodes to the Physics Values manager: • Morpher: Incremental Displacement • Rigid Motion
Internal	The Internal option allows STAR-CCM+ internal models to add to the displacement.
Slide on Guide Surface	The Slide on Guide Surface option adds the following nodes to the Physics Values manager: • Morpher: Guide Surface • Morpher: Guide Surface Motion Points on this boundary are allowed to slide on a specified surface. For stationary and planar guide surfaces, the In-Plane boundary condition is usually more efficient.

Morpher Boundary Values

This section contains descriptions of boundary values that are available for the Morpher motion:

- Grid Velocity
- Morpher: AngularDisplacement
- Morpher: AngularDisplacementAxis
- Morpher: IncrementalDisplacement
- Morpher: LinearDisplacement
- Morpher: ThinFactor
- Morpher Plane Options

- Rigid Motion
- Morpher: Guide Surface
- Morpher: Guide Surface Motion

The particular values that appear depend on the boundary type.

Grid Velocity

This boundary value is activated when the Morpher boundary condition is set to Grid Velocity. It is input as a vector profile. If a user coded profile is used to set this value, the profile must be of type region, rather than type boundary. In this instance, however, only the boundary vertices are used.

Morpher: AngularDisplacement

This boundary value is activated when the Morpher boundary condition is set to Rigid Body. It is input as a scalar profile.

Morpher: AngularDisplacementAxis

This boundary value is activated when the Morpher boundary condition is set to Rigid Body.

Properties

Origin	The Cartesian position vector specifying the origin of the axis of rotation.
Axis	A Cartesian unit vector specifying the direction of the axis of rotation.

Morpher: IncrementalDisplacement

This boundary value is activated when the Morpher boundary condition is set to Motion. It is input as a vector profile, and can be understood as an additional displacement superimposed on the base motion. The base motion is defined using the Rigid Motion boundary value.

Morpher: LinearDisplacement

This boundary value is activated when the Morpher boundary condition is set to Displacement or Rigid Body. It is input as a vector profile.

Properties

Total Displacement — Specifies whether the displacement is an incremental (same amount applied each time step), or a total displacement relative to the initial mesh at time zero. This latter option is used, for example, when the displacements are supplied as absolute values by an external code.

- Activated — The specified displacement is treated as a total value relative to the initial position at time zero.
- Deactivated — The specified displacement is applied to the vertices based on their position at the start of each time step.

Morpher: ThinFactor

This boundary value reduces the number of vertices that are used as control points in order to reduce computational costs. Vertices that are not chosen are given the applied boundary condition (so that the boundary movement is correct).

This object is activated when the Morpher boundary condition is set to one of the following:

- Fixed
- Displacement
- Grid Velocity
- Solid Stress
- Rigid Body

Properties

Vertex Thin Factor — Use only this fraction of boundary vertices as control vertices.

Morpher Plane Options

This boundary value is activated when the Morpher boundary method is set to Fixed Plane or In Plane.

Properties

Damping Factor

Displacements at the boundary are damped smoothly to zero starting at the specified distance. This factor is value L_p in Eqn. (15). The default value of `0.0` means that the distances are calculated automatically.

Continuity Level

Specifies which damping function to use in smoothing vertex displacements approaching the plane. These functions are plotted in the formulation below.

C0 The displacement is damped using Eqn. (13).

C2 The displacement is damped using Eqn. (14).

Expert Properties

Ignore Factor

This factor is multiplied by the Damping Factor to calculate a finite distance. Control points closer than the specified distance to this plane are given special treatment. A value of zero is acceptable.

Rigid Motion

This boundary value is activated when the Morpher boundary method is set to Motion. Its single property is used to select a motion object whose properties are used to define the base motion of the boundary.

Properties

Rigid Motion	Selects a motion object that is used to define the rigid motion of the boundary. The only motion types that can be selected for this boundary value are: • Stationary • Rotation • Rotation and Translation

Morpher: Guide Surface

This boundary value is activated when the Morpher boundary method is set to Slide on Guide Surface. Its single property is used to specify the guide surface the boundary vertices slide on.

Properties

Guide Surface	Allows to select the guide surface the vertices on the boundary slide on. The following options can be selected:	
	Initial surface	Selects the original boundary as the guide surface.
	External file	Specifies the binary STL file named `boundaryName.stl` (boundaryName is to be replaced with the actual name of the boundary), located in the same directory as the `.sim` file, as the guide surface.
	Plane containing this boundary	Selects the plane containing the initial boundary as the guide surface.

Morpher: Guide Surface Motion

This boundary value is activated when the Morpher boundary method is set to Slide on Guide Surface. Its single property is used to specify the motion of the guide surface.

Properties

Guide Surface Motion	Allows to specify the motion of the guide surface. Options other than stationary can be selected only when the Guide surface property for the Morpher: Guide Surface is set to Plane containing this boundary.

Visualizing the Morpher Displacement

You can use the Morpher Displacement field function to analyze how the morpher affects the volume mesh and compare its values with expected results.

Note: Morpher displacement is based upon values at vertices, not cell centers.

To visualize the morpher displacement:

- In a scalar scene, select the Displayers > Scalar 1 > Parts node and set Parts to the boundary on which you want to visualize morpher displacement.
- Select the Scalar Field node and set Function to one of the following:
  - Morpher Displacement > Magnitude: The displacement magnitude of all components.
  - Morpher Displacement > 0: The x-component of displacement.
  - Morpher Displacement > 1: The y-component of displacement.
  - Morpher Displacement > 2: The z-component of displacement.

The following example shows a scalar scene that visualizes the y-component of morpher displacement.



Morphing Motion Formulation

In STAR-CCM+, a set of control vertices defines the initial movement that the morpher imposes on the mesh. Each control vertex is associated with a displacement vector that the morpher uses to move nearby vertices. Using multiquadric interpolation theory [1], these specified displacements are used to generate an interpolation field that applies to the entire region for which the morpher motion applies. When this field has been generated, the morpher uses it to move mesh vertices to their new position.

To generate the interpolation field, a system of equations is created using the control vertices and their specified displacements, where the known displacement $\mathbf{d}'_i$ of each control vertex i is expanded as

$$\mathbf{d}'_i = \sum_{i=1}^{N} f_{b,j}(r_{ij})\boldsymbol{\lambda}_j + \boldsymbol{\alpha} \tag{7}$$

in which $f_{b,j}(r_{ij})$ is a radial basis function of the form

$$f_{b,j}(r_{ij}) = \sqrt{r_{ij}^2 + c_j^2} \tag{8}$$

and r_{ij} is the magnitude of the distance between two vertices:

$$r_{ij} = |\mathbf{x}_i - \mathbf{x}_j| \tag{9}$$

In these equations, $\boldsymbol{\lambda}_j$ is the expansion coefficient, $\mathbf{x}_i$ is the position of vertex i, N is the number of control vertices, and c_j is the basis constant. The basis constant is set to zero in STAR-CCM+. The constant vector $\boldsymbol{\alpha}$ is used to satisfy the additional constraint that

$$\sum_{i=1}^{N} \boldsymbol{\lambda}_j = 0 \tag{10}$$

This additional constraint bounds the expansion for large $\mathbf{x}$. The system that is described in Eqn. (7) is non-singular if $\mathbf{x}_i \neq \mathbf{x}_j$ for all $i \neq j$.

Equations (7) and (10) are solved to give the Cartesian components of all $\boldsymbol{\lambda}_j$ and the components of constant vector $\boldsymbol{\alpha}$. This results in the desired interpolation field:

$$\mathbf{d}'(\mathbf{x}) = \sum_{i=1}^{N} f_{b,j}(r)\boldsymbol{\lambda}_j + \boldsymbol{\alpha} \tag{11}$$

where $f_{b,j}(r)$ applies at all vertices in the mesh, that is,

$$f_{b,j}(r) = \sqrt{r^2 + c_j^2}$$
$$r = |\mathbf{x}\text{-}\mathbf{x}_j|$$

This equation can now be used to move mesh vertices by the calculated displacement $\mathbf{d}'$.

The matrices that are involved in equations (7) and (11) are dense, and a straightforward matrix operation would require $O(N^2)$ operations. Fortunately, an $O(N\log(N))$ method has been developed to solve the linear equation (7) and evaluate equation (11) at many vertex locations. The method involves a conjugate gradient method with a special preconditioner ([2], [3]) and a fast multipole method [4] to accelerate the matrix-vector multiplication.

Plane Boundary Conditions

Two plane boundaries are provided for use with the morpher motion: in plane and fixed plane. For the fixed plane condition, the vertices which lie on the specified plane are fixed, while for the in-plane condition they are permitted to move within the plane itself. An effect similar to the fixed plane condition can be realized by using control vertices but the fixed plane condition itself can be more efficient. The movement of vertices near the

constraint planes is such that they join smoothly with the movement of vertices on the plane. This movement near the constraint planes can be controlled using the damping distance L_p that is described below.

These constraints are imposed before solving Eqn. (7). In preparation, user-specified displacements $\mathbf{d}_i$ are transformed to $\mathbf{d}'_i$ by solving

$$\mathbf{d}_i = \prod_{p=1}^{N_p} [s_p \mathbf{n}_p \mathbf{n}_p^T + t_p (1 - \mathbf{n}_p \mathbf{n}_p^T)] \mathbf{d}'_i \tag{12}$$

where $\mathbf{n}_p$ is the normal of the p-th plane and N_p is the number of planes. The displacements $\mathbf{d}'_i$ are transformed back to $\mathbf{d}_i$ together with other mesh vertices at the end of the morphing process. s_p and t_p are functions that depend on the plane type and the continuity condition you choose. These functions are given as follows:

	C_0 continuous		C_2 continuous	
	s_p	t_p	s_p	t_p
Fixed plane	$f_0(l)$	$f_0(l)$	$f_2(l)$	$f_2(l)$
In-plane	$f_0(l)$	1	$f_2(l)$	1

Here:

$$f_0(l) = \begin{cases} 1 & ; \ l \geq 1 \\ 1 - (1 - l)^3 & ; \ 1 > l \geq 0 \\ 0 & ; \ l < 0 \end{cases} \tag{13}$$

$$f_2(l) = \begin{cases} 1 & ; \ l \geq 1 \\ l^3 (10 - 15l + 6l^2) & ; \ 1 > l \geq 0 \\ 0 & ; \ l < 0 \end{cases} \tag{14}$$

in which

$$l = \frac{\mathbf{n}_p (\mathbf{x} - \mathbf{x}_{0,p})}{L_p} \tag{15}$$

is the ratio of the signed distance from the p-th plane to $\mathbf{x}_i$ and the damping distance L_p. $\mathbf{x}_{0,p}$ is a position vector on the p-th plane and L_p is either entered as a finite distance or calculated automatically as $(L_{min} + L_{max})/2$. L_{min} and L_{max} are respectively the distances from the plane to the nearest and farthest control vertices on the active side of the plane.

Functions f_0 and f_2 in the range from 0 to 1 are plotted below for comparison. A continuity level of C_2 results in a smooth transition.



Solution Steps

The steps that are involved in the morpher stage within a time step are:

1. The control vertex displacements $\mathbf{d}_i$ are transformed to $\mathbf{d}'_i$ according to Eqn. (12).
2. Eqn. (7) is solved with Eqn. (10) to provide an interpolation field.
3. The displacement vectors are transformed from $\mathbf{d}'_i$ to $\mathbf{d}_i$ using Eqn. (12).

Morphing Field Functions

The following primitive field function is made available to the simulation when the morpher motion model is used. This field function has a typical set of base properties.

SavedCoord	This field contains the coordinates before the morphing.

Morpher Displacement — Vector value of the distance that the morpher has moved the mesh vertices (not face centers). As this field is a vector field function, you can plot the x-, y-, z-components, and magnitude.

Motion Bibliography

[1] Hardy, R.L. 1990. "Theory and applications of the multiquadric-biharmonic method", *Comput. Math. Applic.*, **19**, pp. 163-208.

[2] Faul, A.C., Goodsell, G. and Powell, M.J.D. 2005. "A Krylov subspace algorithm for multiquadric interpolation in many dimensions.", *IMA Journal of Numerical Analysis*, **25**, pp. 1-24.

[3] Gumerov, N.A. and Duraiswami, R. 2007. "Fast radial basis function interpolation via preconditioned Krylov iteration.", *SIAM Journal on Scientific Computing*, **29**, pp. 1876-1899.

[4] Beatson, R.K., Powell, M.J.D. and Tan, A.M. 2006. "Fast evaluation of polyharmonic splines in three dimensions", DAMTP 2006/NA03.

Working with User-Defined Vertex Motion

The User-Defined Vertex motion allows you to set a displacement or velocity for every vertex position in the simulation. It is typically intended for use in cases where the other motions cannot give the precise control that is required over the mesh movement. The User-Defined Vertex motion node represents this type of motion. You can add this node using the right-click menu of the Motions manager.



When you select this motion for the region you wish to move, an additional condition, the User-Defined Vertex Option, is added to the Physics Conditions folder of the region. Use this node to set the quantity by which the motion is defined.

Setting the Quantity for User-Defined Motion

User-defined motion can be specified in terms of displacement, grid velocity, new coordinates or coordinate offset. The choice between these quantities is made using the Method property of the User-Defined Vertex Option node.



Method — Specifies the method by which the mesh motion is defined.

- **Fixed** — No vertex motion. Use this method when solving for initial flow fields, particularly in cases for fluid-structure interaction.
- **Displacement** — Move vertices by a specified displacement. Adds the LinearDisplacement node to the Physics Values folder of the region.
- **Grid Velocity** — Calculate vertex movement within a time step using a grid velocity. Adds the Grid Velocity node to the Physics Values folder of the region.

New Coordinate	Specify new coordinates at each time step. Adds the New Coordinate node to the Physics Values folder of the region.
Coordinate Offset	Specify an offset by which vertices move at each time step. Adds the Coordinate Offset node to the Physics Values folder of the region.

Initial Displacement Specifies whether to modify the mesh generation position by adding a displacement.

Activated	The vertices of the mesh are moved by the initial displacement. Adds the Initial Displacement node to the Physics Values folder of the region. For more information, see What Is the Initial Displacement?.
De-activated	The vertices of the mesh remain in the mesh generation position. This position is used as the initial position.

What Is the LinearDisplacement?

Use this method to set the linear displacement of vertices in the parent region for which a user-defined vertex motion is active. This method is valid only with Cartesian coordinate systems. Do not use this method in cylindrical or spherical coordinates. For vertex motion in these coordinate systems, use the New Coordinate or Coordinate Offset methods.

The LinearDisplacement node acts as a vector profile and has a single additional property for the type of displacement.

Properties

Total Displacement — Use this property to specify whether the displacement is taken as incremental (applying the same displacement at each time step), or whether it is taken as a total displacement relative to the initial mesh at time zero. Use this latter option when you supply displacements as an absolute value.

- Activated — The specified displacement is treated as a total value relative to the initial position at time zero.
- Deactivated — The specified displacement is applied to the vertices based on their position at the start of each time step.

Coordinate System — Coordinate system in which the displacements are defined.

Method — The method that is used to define the value for the Linear Displacement.

Dimensions — The dimensionality of the property that is expressed in terms of the basic dimensions, such as length, time (read-only).

What Is the Grid Velocity?

Use this node to set the grid velocity that is applied to vertices in the parent region when a user-defined vertex motion is active. This node acts as a vector profile.

This method is valid only with Cartesian coordinate systems. Do not use this method in cylindrical or spherical coordinates. For vertex motion in these coordinate systems, use the New Coordinate or Coordinate Offset methods.

The coordinate system of the field function must be identical to the coordinate system that is used in the field function definition.

What Is the New Coordinate?

Use this node to specify the new coordinates for the vertices at each time step. The coordinates can be specified in the Cartesian, Cylindrical, or Spherical coordinates system.

Consider a half-section of a cylinder for rotation with a constant angular velocity about its axis:



You can use a field function to rotate this cylindrical section:

- Set the Coordinate System to the cylindrical coordinate system of the mesh for rotation, for example: Laboratory->Cylindrical 1
- Create a field function with a Definition of the form:

```
[ $$Position(@CoordinateSystem("Cylindrical 1"))[0],
  $$Position(@CoordinateSystem("Cylindrical 1"))[1]+0.1,
  $$Position[2] ]
```

- Set the New Coordinate Method to Field Function and select the field function that is defined in the previous step.

The coordinate system of the field function must be identical to the coordinate system that is used in the field function definition.

Properties

Coordinate System	Coordinate system in which the origin is defined.
Method	The method that is used to define the value for the New Coordinate.
Field Function	Defines the value using a field function (typically user defined). Adds the Field Function node as a child to this node.
User Code	Defines the value using a user subroutine. Adds the User Code node as a child to this node.

What Is the Coordinate Offset?

Use this node to specify the change in coordinates of the vertices at each time step. The change in coordinates can be specified in the Cartesian, Cylindrical, or Spherical coordinates system.

Consider the following section of a cylinder for rotation with a constant angular velocity about its axis:



It is not possible to use the Constant method for such a motion: the grid velocities of each vertex are not equal. You can use a field function to rotate this cylindrical section:

- Set the Coordinate System to the cylindrical coordinate system of the mesh for rotation, for example: Laboratory->Cylindrical 1
- Use the Constant method with a value of `[0.0, 0.1, 0.0]` for example.

If the mesh must grow in the radial direction as it is rotated, you cannot use the Constant method. Each radial vertex value increases at a different rate, depending on its radial position. Instead, use a field function such as:

```
[ $$Position(

  @CoordinateSystem("Cylindrical 1"))[0] * 0.05,

  0.1,

  0.0 ]
```



The coordinate system of the field function must be identical to the coordinate system that is used in the field function definition.

Properties

Coordinate System	Coordinate system in which the origin is defined.	
Method	The method that is used to define the value for the Coordinate Offset.	
	Constant	Specifies the value as a constant array. Adds the Constant node as a child to this node.
	Field Function	Defines the value using a field function (typically user defined). Adds the Field Function node as a child to this node.
	User Code	Defines the value using a user subroutine. Adds the User Code node as a child to this node.

What Is the Initial Displacement?

Use this node to adjust the mesh from the mesh generation position to the initial mesh position for the simulation. You specify an initial displacement by which the vertices are moved from the mesh generation position to the initial mesh position of the simulation. The initial displacement is applied only once at the start of the simulation to obtain the initial position for the simulation. The initial position is used to initialize the flow field. Note that selecting *Solution* > Clear Solution... > Reset Mesh resets the mesh back to the mesh generation position.

Properties

Coordinate System	Coordinate system in which the initial displacement vector is defined.
Method	The method that is used to define the value for the Initial Displacement.
Dimensions	The dimensionality of the property that is expressed in terms of the basic dimensions, such as length (read-only).

What Is the User-Defined Vertex Motion Solver?

The user-defined vertex motion solver, which has expert properties, applies your specified displacement or velocity to the mesh.



Expert Properties

Update At Inner Iterations

Repeats the movement at each inner iteration during the time-step. This option is only meaningful for cases where some boundary condition depends on the solution; it is off by default to reduce computational costs.

Activated — Repeats the motion at each inner iteration.

Deactivated — Leaves off this option.

Zero grid flux

When this option is activated, the mesh still moves but the grid flux is not included in the transport equations. This option provides more stability when marching towards a steady solution with an unsteady solver. It is also useful when you are simply adjusting the shape of the region without solving the transport equations.

Activated — Ignores the grid flux.

Deactivated — Includes the grid flux.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. Do not activate this feature, as it is primarily a "debugging" tool and not supported for normal simulation work.

Activated — The solver does not update the solution.

Deactivated — The solver updates the solution.

An Example of User Coding

User coded functions can be used to define the displacements or grid velocity that STAR-CCM+ requires. Such a function requires access to all the vertices in the mesh, obtained as follows:-

```
ucarg((void *)userMotionFunc, "Vertex", "Coord",
sizeof(CoordReal[3]));
```

The argument type here is `"Vertex"`, rather than `"Cell"` or `"Face"`, and the array being requested is `"Coord"`. The argument belongs to the function, `userMotionFunc`, which could have a declaration as follows:-

```
void userMotionFunc(Real (*result)[3], int size, CoordReal
(*vpos)[3])
{
   // user coding here to calculate the displacement
   // or grid velocity, which will be returned
   // in the result array
}
```

In this function, the `result` array would return one of the following quantities:

- incremental displacement;
- total displacement;
- grid velocity.

The particular quantity depends on the choice you have made on the User Defined Vertex Option node and its associated value node.

When a suitable function has been compiled and linked, it appears in the Field Functions list as a vector function. Select the Field Function option for the Method property of the UDV: LinearDisplacement node or Grid Velocity node, and choose the user coded field function from the list.

General Motion Formulation

This section discusses how the motion modeling affects the continuum transport equations. The motion enters these equations in two ways:

- As a grid flux in the convective terms
- As a body force due to system rotation in the momentum equations

Grid Flux

The grid flux at a face, as used in Eqn. (2538) to compute the face mass flux, is calculated as follows.

For a moving reference frame:

$$G_f = \mathbf{v}_g \cdot \mathbf{a}_f \tag{16}$$

where $\mathbf{v}_g$ is the grid velocity and $\mathbf{a}$ is the face area.

For rigid body motion, morphing, and user-defined vertex motion, the first-order grid flux is:

$$G_f = \frac{\delta V_f^n}{\Delta t} \tag{17}$$

The second-order grid flux is:

$$G_f = \frac{(\alpha^2 - 1)\delta V_f^{n+1} - \delta V_f^n}{\alpha(\alpha - 1)\Delta t^{n+1}} \tag{18}$$

where

$$\alpha = 1 + \frac{\Delta t^{n+1}}{\Delta t^n} \tag{19}$$

$$\Delta t^{n+1} = t^{n+1} - t^n$$
$$\Delta t^n = t^n - t^{n-1}$$

and

δV_f^{n+1}	Volume swept by cell face f over time step Δt^{n+1}.
δV_f^n	Volume swept by cell face f over time step Δt^n.

In Eq. (18), the swept volumes are calculated from the grid position at two time levels. n denotes the current time level, $n-1$ the previous time level, and $n+1$ the next time level. Δt is the time step.

Body Force Due to System Rotation

If the reference frame is rotating, the rotational body force ($\mathbf{f}_r$ in Eqn. (100) and Eqn. (135)) is:

$$\mathbf{f}_r = \rho \boldsymbol{\omega} \times \mathbf{v} \tag{20}$$

where $\boldsymbol{\omega}$ is the rotational velocity.

Common Field Functions

The following primitive field function is made available to the simulation with motions. This field function has a typical set of base properties.

Grid Flux	The rate at which a face sweeps out volume due to grid motion.

Modeling Materials

Some simulations that you perform in STAR-CCM+ involve modeling a substance such as air, water, or aluminum. Others involve several substances such as a combusting mixture of gases, or a sloshing mixture of air and water. The Material model handles the definition of these substances.

Additional information is available about:

- Selecting a material model
- Modeling a single-component material
- Modeling a multi-component mixture
- Modeling a multiphase mixture
- Managing single-component materials
- Managing multi-component materials
- Managing multi-part solids
- Modeling solid stress materials—Linear Isotropic Elastic
- Setting material properties
- Using material databases

What Is a Material Model?

The material model is responsible for managing the material, that is, substance or substances, being simulated in the continuum. The material, in turn, is responsible for managing the various thermodynamic and transport properties relevant to that material and to the physical processes being modeled in the continuum.

There are three general types of material models available in STAR-CCM+: single-component, multi-component, and multiphase mixture. Each of these material models manages one specific type of material: a pure substance, a multi-component mixture, or a multiphase mixture.

Selecting a material model early in the model selection process allows STAR-CCM+ to assist you in making subsequent modeling choices appropriately.

A material model is only required if you plan to model fluid flow or heat transfer. If, for example, all you are doing is checking the mesh, a material model is not needed.

What Is a Single-Component Material Model?

The single-component material models are for modeling pure substances such as air, water, or aluminum.

There are three single-component material models in STAR-CCM+: Gas, Liquid, and Solid models. Each manages one single-component material representing the pure substance being modeled. The Gas model manages a gas material (that is, a pure substance in the gaseous phase). The Liquid model manages a liquid material (that is, a pure substance in the liquid phase). The Solid model manages a solid material (that is, a pure substance in the solid phase).

What Is a Multi-Component Material Model?

The multi-component material models are single-phase models for simulating fully miscible mixtures of two or more pure substances in the same phase, such as a combusting mixture of fuel, oxidizer and product.

There are two multi-component material models in STAR-CCM+: Multi-Component Gas and Multi-Component Liquid models. Each manages a multi-component material representing the single-phase, miscible mixture being modeled. The Multi-Component Gas model manages a Gas Mixture material (that is, a miscible mixture of two or more pure substances in the gaseous phase). The Multi-Component Liquid model manages a Liquid Mixture material (that is, a miscible mixture of two or more pure substances in the liquid phase). These models are specified as either reacting or non-reacting.

What Is a Multiphase Mixture Material Model?

The multiphase mixture material model is for simulating two or more immiscible phases, where each phase is composed of a pure gas or liquid substance, such as air and water.

The Multiphase Mixture model in STAR-CCM+ manages objects that control properties for the mixture as a whole. The phase materials and equation of state specifications are managed with the Eulerian Multiphase model.

Selecting a Material Model

To select a Material model,

- Open the Continua node of the simulation tree, then right-click a physics continuum node such as Physics 1. Select the Select models... item from the pop-up menu.



The *Physics Model Selection* dialog opens.

- Select the correct type of Material model for the simulation by clicking the relevant radio button. For example, the following figure shows the single component Gas material model being selected.



The dialog now shows the activated model, with model options other than materials remaining to be selected.



Modeling a Single-Component Material

To model a single-component material, that is, a pure substance, activate either the Gas, Liquid, or Solid model by selecting the appropriate radio button in the Material section of the model selection dialog. Eulerian phases include a Particle model when the Multiphase Segregated Flow model is selected.

What Is a Gas Model?

The Gas model is used to simulate a pure gas.

When the single-component Gas material model is activated, the Gas model node represents the model in the object tree as shown.



Gas Properties

None

See also: common menu items.

What Is a Gas?

The gas material represents a single pure gas substance. The Gas model manages the material as outlined in the section describing single-component material models. By default the Gas model takes the gas material air from the material database.

The following screenshot shows the Gas model node and its child node Air, which represents the pure gas material (air) being modeled.



If you wish to model a gas other than air, you can replace the existing gas material with a different one from the material database. This procedure is outlined in the section describing how to manage single-component materials. Alternatively, you can edit the properties of the gas directly as outlined in the section on setting material properties.

What Is a Liquid Model?

The Liquid model is used to simulate a pure liquid.

When the single-component Liquid material model is activated, the Liquid model node represents the model in the object tree as shown.



Liquid Properties

None

See also: common menu items.

What Is a Liquid?

The liquid material represents a single pure liquid substance. The Liquid model manages the material as outlined in the section describing single-component material models. By default the Liquid model takes the liquid material water from the material database.

The following screenshot shows the Liquid model node and its child node H2O which represents the pure liquid material (water) being modeled.



If you wish to model a liquid other than water, you can replace the existing liquid material with a different one from the material database. This procedure is outlined in the section describing how to manage single-component materials. Alternatively, you can edit the properties of the liquid directly as outlined in the section on setting material properties.

What Is a Solid Model?

The Solid model is used to simulate a pure solid.

When the single-component Solid material model is activated, the Solid model node represents the model in the object tree as shown.



Solid Properties

None

See also: common menu items.

What Is a Solid?

The solid material represents a single pure solid substance. The Solid model manages the material as outlined in the section describing single-component material models. By default the Solid model takes the solid material Aluminum from the material database.

The following screenshot shows the Solid model node and its child node Al which represents the pure solid material (aluminum) being modeled.



If you wish to model a solid other than aluminum, you can replace the existing solid material with a different one from the material database. This procedure is outlined in the section describing how to manage single-component materials. Alternatively, you can edit the properties of the solid directly as outlined in the section on setting material properties.

What Is a Particle Model?

The Particle model is used to represent a dispersed particulate phase in the context of Eulerian multiphase flows. Single particles are not tracked using this model, as in the Lagrangian multiphase framework. Instead, transport equations for the phase model the mean behavior.

Modeling a Multi-Component Mixture

A multi-component mixture is a miscible mixture of two or more pure substances in the same phase. To model a multi-component mixture, activate either the Multi-Component Gas or Multi-Component Liquid model by selecting the appropriate radio button in the *Material* section of the model selection dialog.

What Is a Multi-Component Gas Model?

The Multi-Component Gas model is used to simulate a miscible mixture of two or more pure gases. This model is available for both single phase and multiphase flows.

When the Multi-Component Gas material model is activated, the Multi-Component Gas model node represents the model in the object tree as shown below.



Multi-Component Gas Properties

None

See also: common menu items.

What Is a Gas Mixture?

A Gas Mixture is a material representing a miscible mixture of two or more pure gases. The Multi-Component Gas model manages the material as outlined in the section describing multi-component material models.

The following screenshot shows the Multi-Component Gas model node, the Gas Mixture node it manages, and the child nodes of the mixture: Gas Components and Mixture Properties.



The Gas Components node is a manager node for the gases composing the Gas Mixture. Initially, the Gas Components node is empty. It is here that you add, remove, and replace gases in the mixture as outlined in the section on managing multi-component materials. The material properties of each individual gas component in the mixture are then accessible and can be set as outlined in the section on setting material properties.

The Mixture Properties node is a manager node for the material properties of the entire gas mixture as a whole, distinct from the properties of each individual component. This node is empty initially, but is populated with material property nodes once other physical models are selected. The material properties of the gas mixture can be set as outlined in the section on setting mixture properties.

What Is a Multi-Component Liquid Model?

The Multi-Component Liquid model is used to simulate a miscible mixture of two or more pure liquids. This model is available for single phase and multiphase flows.

When the Multi-Component Liquid material model is activated, the Multi-Component Liquid model node represents the model in the object tree as shown below.



Multi-Component Liquid Properties

None

See also: common menu items.

What Is a Liquid Mixture?

A Liquid Mixture is a material representing a miscible mixture of two or more pure liquids. The Multi-Component Liquid model manages the material as outlined in the section describing multi-component material models.

The following screenshot shows the Multi-Component Liquid model node, the Liquid Mixture node it manages, and the child nodes of the mixture: Liquid Components and Mixture Properties.



The Liquid Components node is a manager node for the liquids composing the Liquid Mixture. Initially, the Liquid Components node is empty. It is here that you add, remove, and replace liquids in the mixture as outlined in the section on managing multi-component materials. The material properties of each individual liquid component in the mixture are then accessible and can be set as outlined in the section on setting material properties.

The Mixture Properties node is a manager node for the material properties of the entire liquid mixture as a whole, distinct from the properties of each individual component. This node is empty initially, but is populated with material property nodes once other physical models are selected. The material properties of the liquid mixture can be set as outlined in the section on setting mixture properties.

What Is a Non-Reacting Flow System?

This system contains a multi-component fluid mixture whose constituents do not react chemically with each other. A node representing such a system (Non-reacting in the following screenshot) appears in the simulation tree when this type of flow problem is specified in the *Physics Model Selection* dialog.



Modeling a Multiphase Mixture

A multiphase mixture is a fluid that is composed of multiple phases: an immiscible mixture of two or more pure substances in either the gas or liquid phase.

To model a multiphase mixture, activate the Multiphase Mixture model.

What Is the Multiphase Mixture Model?

The Multiphase Mixture model is used to simulate an immiscible mixture of two or more pure substances in either the gas or liquid state. It is used with the Eulerian Multiphase model when the Volume of Fluid (VOF) model is active.

The corresponding node (Eulerian Multiphase > Mixture in the following screenshot) is activated as soon as the option Eulerian Multiphase is selected in the *Physics Model Selection* window.



This set of objects only controls properties for the mixture as a whole. They have no properties, but are subject to common menu items. The phase materials and equation of state specifications are managed with the Eulerian Multiphase model.

Managing Single-Component Materials

The single-component material node, which is a child node of the material model, represents the specific type of substance being modeled. For example, as shown in the following screenshot, the Air node exists within the Gas model node, since air is the default gas type when the Gas material model is selected.



The type of material defaults to air, water, or aluminum, depending on the type of the material model (Gas, Liquid, or Solid). A node with a corresponding name appears with properties and a pop-up menu.

While you can rename this node using standard techniques, you can as easily replace the type of material used. It is also possible to set material properties.

Single-Component Material Properties

Database Material	The name of the material as it is listed in the material database (read-only).

Single-Component Material Pop-Up Menu

Replace with	Opens a submenu of all materials of the same phase in the material database.

See also: common menu items.

Replacing Single-Component Materials

To replace a single-component material, right-click the node of the existing material type and select Replace with...



Then select a replacement material from the list in the *Replace Material* dialog and click OK.



The node of that selected material appears in the place of the existing material node.

Setting Material Properties

The Material Properties manager node is a placeholder for the properties of a substance.



This node is empty until either an Equation of State model or a Viscous Regime is selected, after which it becomes populated with material property nodes. A fully populated node is shown below.



Material Properties Manager Properties

None

See common menu items.

Managing Multi-Component Materials

The multi-component mixture node, which is a child node of the material model, is a placeholder for two manager nodes: one for the components of the mixture where you would add substances, the other for the properties of the mixture. For example, in the following screenshot, the Gas Mixture node contains a Gas Components node and a Mixture Properties node.



It is possible to rename the multi-component mixture node in a manner that reflects the specific aspects of the material you are modeling. A mixture can also be copied to another simulation by copying the continuum that contains it.

Multi-Component Mixture Properties

None

See also: common menu items.

Setting Mixture Properties

The information in this section applies equally well to the corresponding nodes inside a multi-component or multiphase material model node, since these nodes function in a similar fashion.



Multi-Component



Multiphase

The Mixture Properties node is a manager node for the material properties of the entire mixture as a whole, distinct from the properties of each individual component. It functions the same as the material properties manager of a single-component material. This node is empty initially, but is populated with material property nodes once other physical models are selected. The material properties of the fluid mixture can be set as outlined in the section describing setting material properties.

Managing Mixture Components

The mixture components manager node (such as Gas Components in the following screenshot) is a placeholder for the nodes of the substances that comprise a mixture.



This node remains unpopulated until you add components.

You can also modify a mixture through the following actions:

- Replacing components
- Reordering components
- Setting component properties

A material can be renamed using the standard techniques.

Materials can be deleted from a mixture by right-clicking on the node and selecting Delete. This action removes the material from the mixture, but not from the material database.

Mixture Components Manager Properties

None

Mixture Components Manager Pop-Up Menu

Gas Mixture

New Gas Components	Opens the *Select Mixture Components* dialog, which displays all gas materials in the material database. Activate the checkboxes in the order in which you want to add components, and click OK.

Liquid Mixture

New Liquid Component	Opens a submenu of all liquid materials in the material database from which you can add a component to the mixture.

Adding Mixture Components

The same method is used to add components to a mixture irrespective of whether a multi-component gas or liquid is being modeled. The component names and material properties are retrieved from the material database, and constant property values are assigned by default. As part of setting up the simulation, review the default material property settings of the mixture as well as the settings of each new component. If necessary, change the settings of the mixture properties, and then the settings of each new component as required (see Managing Multi-component Materials).

The ordering of components can be important for reasons of numerical accuracy. See the section on using the Species model for details.

Adding Components to a Mixture

The following steps explain how to add a component to a mixture:

- Right-click the Gas or Liquid Components node and choose Select Mixture Components.



The *Select Mixture Components* dialog appears. This dialog is essentially the object selection dialog that is applied to material selection.

- In the *Tree View* panel, expand the Material Databases > Standard node (corresponding to the Standard materials database) and its child node (Gases or Liquids).

- Activate the boxes corresponding to the desired components, and click OK.



- The new components appear beneath the Gas or Liquid Components node.



Replacing Mixture Components

To replace a mixture component, right-click the node of the component and select Replace With... The *Replace Material* dialog appears. Expand the nodes as necessary and activate the box corresponding to the replacement component. Click OK to confirm the selection.



Reordering Mixture Components

You can control the order in which the mixture components appear in the object tree. Right-click the Gas or Liquid Components node and choose Reorder Mixture Components... from the menu.



The *Reorder Mixture Components* dialog appears, showing the names of selected components in their order on the tree.



To move the position of a component:

- Click and hold the left mouse button on its name;
- Drag the mouse pointer until it hovers over the destination entry in the list; and,
- Release the mouse button.

The dragged component now occupies the targeted entry, and the component that was there moves up one position. Click OK to save the changes.

Setting Component Properties

The mixture-component material node exists as a child node of the mixture components node.



It contains a sub-node that functions the same as the material properties of a single-component material.



Mixture Component Properties

Database material	The name of the material (read-only).
ID	Unique index for this component within the simulation. If this component is copied to another continuum, the copy within the target continuum is assigned a new, unique, ID.

Mixture Component Pop-Up Menu

Replace with...	Opens the *Replace Material* dialog from which you can replace the selected component.

See also: common menu items.

Managing Multi-Part Solids

Some simulations require you to model many solid parts of different materials. Examples of such simulations are:

- Electronics cooling studies where a motherboard contains tens or hundreds of different components
- Vehicle underhood thermal analysis

The Multi-Part Solid model allows you to define multiple solid materials in a single physics continuum, and to map those materials to individual parts in the region to which the continuum applies. The Multi-Part Solid node represents this model in a physics continuum.

When a physics continuum containing the Multi-Part Solid model is applied to a region, a Material Part Groups node is added to the Physics Values manager for that region. You use the Material Parts Group node to map between solid materials and geometry parts.

The following diagram shows the Multi-Part Solid node in the physics continuum, and the Material Part Groups node in a region to which the physics continuum is applied:



Requirements:

- Multi-part solids can only be used in simulations where geometry parts define the geometry. The identity of these geometry parts must be retained at the region level. Therefore, when you assign geometry parts to regions, include multiple geometry parts in the same region.
  - Select the One region for all parts option in the *Region Mode* menu of the *Assign Parts to Regions* dialog.
- When you have interfaces between parts, the best accuracy can be obtained by setting the interface type to Contact. In this instance, the thermal flux across the interface is calculated in a manner that accounts for the presence of discontinuous properties. Furthermore, Contact interfaces support the application of a Thermal Contact Resistance and/or Heat Source within the interface if necessary. In contrast, if you set the interface between parts to type Internal, the heat flux at the interface is calculated based on averaged properties. This results in a certain amount of smoothing of the temperature field across the interface, which is not always realistic. Therefore, it is recommended that interfaces between parts of dissimilar solid materials be set to type Contact.
- If any material in the simulation has anisotropic thermal conductivity, create a dedicated physics continuum for that material. The Multi-Part Solid model does not currently support materials that require the Anisotropic method for Thermal Conductivity.

To set up a simulation in STAR-CCM+ with multi-part solids:

- Select the Multi-Part Solid model.

- Assign parts to materials.
- Change the interface type between solid parts that are in contact.

This model activates field functions that are useful for post-processing.

Selecting the Multi-Part Solid Model

To use the Multi-Part Solid model:

- In the *Physics Model Selection* dialog, select a model from the *Space* group. This activation also selects the Gradients model automatically from the *Gradient Metrics* group box, provided the Auto-select recommended models checkbox is activated.
- Select the Steady or Implicit Unsteady option from the *Time* group, depending on whether your simulation is transient.
- Select Multi-Part Solid from the *Material* group.

  The *Physics Model Selection* dialog contains your selections similar to the example shown in the following screenshot. In this example, the option Steady was selected.



- Select Coupled Solid Energy or Segregated Solid Energy from the *Optional Models* group.
- Select a model from the *Equation of State* group.

Be sure to assign each region to the proper continuum. Only a solid region is compatible with a continuum that contains the Multi-Part Solid model.

Once you finish activating the physics models, define the solid components.

Defining Solid Components

The Multi-Part Solid model requires you to add the component materials in the same way as you add mixture components for a multi-component liquid or gas. The Solids node manages these component materials. Default properties for each of these components come from the material database.



To add solid components:

- Right-click the Solids node and choose Select Mixture Components...

  The *Select Mixture Components* dialog appears.
- Expand the nodes within the *Select Mixture Components* dialog.

- Activate the checkboxes next to the materials you require; click Apply, then Close, when you are done.



You can manage solid components in the same way as you manage gas or liquid components.

Once these components are in place, work with material properties methods as needed. The choices of available material properties and associated methods are the same as with a Solid model, except that the Anisotropic method for Thermal Conductivity is not available with the Multi-Part Solid model.

Assigning Parts to Materials

When you associate a multi-part solids physics continuum with a region, STAR-CCM+ adds a Material Part Groups node to the Physics Values for the region. Material part groups corresponding to the solid components of the Multi-Part Solid model is listed beneath the Material Part Groups node.



- Select each of the material part groups in turn and assign geometry parts to the Parts property.

If you switch the region to another multi-part solid physics continuum, STAR-CCM+ repopulates the Material Part Groups node with the solids selected in the newly assigned continuum.

Selecting Parts for Materials

In a multi-part solid simulation, each material appears in an associated region as a material part group node (Fe in the following screenshot). To assign parts, select the node of the desired material for access to its properties.



The Parts property of each of these objects opens the object selection dialog.



In this example, the Fe node has `Chip3` and `Chip8` selected, meaning that those chips are made of iron.

Parts can correspond to only one particular material at a time. If you reassign a part to another material, the part disappears automatically from the list for the previous material.

Material Part Group Properties

Parts	List of parts in this material part group. Clicking the property customizer (...) accesses the object selection dialog to add or remove parts.

Post-Processing Multi-Part Solid Data

Various relevant field functions can be post-processed in scenes:

- Thermal Conductivity



- Density



- Specific Heat of Solids



- Solid Material ID -- associating parts with materials visually



For example, suppose that you want to see which parts in the scene are made of zinc. Select the Zn component node in the Solids manager in the continuum model, and get the ID number, which in this case is 1. To associate that number with the color of certain parts in the scalar scene, look at the color bar in the scene -- the parts with the associated color are made of zinc.

Multi-Part Solid Field Functions Reference

The following primitive field functions are made available to the simulation when the Multi-Part Solid model is used.

Cell Part Index	The Index value of the part that is associated with a given cell. Generated by the meshing process. Also used in parts-based reporting.
Solid Material ID	The ID value of the solid material that is associated with a given cell. Identifies which solid component in the continuum defines the material properties for that cell. Also allows visual association of cells defining parts with a particular material.

Modeling Linear Isotropic Elastic

In STAR-CCM+ the available solid stress material model is Linear Isotropic Elastic. When this material model is activated, the Linear Isotropic Elastic node represents the model in the object tree as shown in the following screenshot.



This section explains the following:

- What is the Linear Isotropic Elastic model?
- Selecting the Linear Isotropic Elastic model
- Formulation

What Is the Linear Isotropic Elastic Model?

A material is considered elastic when its stress is linearly related to the strain. If an elastic material is loaded, and then unloaded, it returns to it original undeformed shape. The deformation does not depend on the loading history or the rate of loading and the material exhibits no damping behavior on its own. In other words, it is a reversible process.

The Linear Isotropic Elastic material model works with the Solid Stress model in STAR-CCM+.

Selecting the Linear Isotropic Elastic Model

When the Solid Stress model is selected, the Linear Isotropic Elastic material model is selected automatically, provided the automatic selection of recommended models is activated.



Linear Isotropic Elastic Formulation

The relationship between the stress and strain in an *isotropic* elastic material is:

$$\sigma = 2G\varepsilon + \lambda tr(\varepsilon)\delta \tag{21}$$

where G is the shear modulus and λ is Lame's modulus both with units of stress and δ is the Kronecker delta. Two independent parameters describe the material modulus parameters. Other commonly used parameters are Young's Modulus E, Poisson's ratio ν, and the bulk modulus K. Young's modulus and Poisson's ratio have a simple interpretation for an uniaxial stress loading typical in coupon material tests or a bar loaded only down its axis.

$$\sigma = \begin{bmatrix}\sigma_{xx} & 0 & 0 & 0 & 0 & 0\end{bmatrix}^T \tag{22}$$

In this case, the axial stress is the product of Young's modulus and the axial strain $\sigma_{xx} = E\varepsilon_{xx}$. As the bar is loaded in tension, it elongates down the axis, but it shrinks lateral to the axis. The amount of lateral shrinking is proportional to the axial elongation by the amount of Poisson's ratio:

$$\varepsilon_{yy} = \varepsilon_{zz} = -\nu\varepsilon_{xx} \tag{23}$$

For pure shear loading, the shear stress is product of the shear modulus and the engineering shear strain:

$$\sigma_{xy} = 2G\varepsilon_{xy} = G\gamma_{xy} \tag{24}$$

The bulk modulus relates the mean stress to the volumetric strain:

$$\sigma_m = \frac{(\sigma_{xx} + \sigma_{yy} + \sigma_{zz})}{3} = K(\varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz}) \tag{25}$$

The elastic properties for a solid material are prescribed under continua.

Setting Material Properties Methods

After the selection of an Equation of State model or a Viscous Regime, the Material Properties manager node becomes fully populated with material property nodes, which can be accessed through their properties. Material properties that use polynomials can be plotted instantly.



Material property methods can be set in the following ways, as applicable:

- Using constant values
- Using field function values
- Using polynomials
- Using thermodynamic polynomial data
- Using IAPWS-IF97
- Using Sutherland's law
- Using Table (Reaction Time)
- Using Table(T)
- Using Table(T,P)
- Using the Density Table
- Using the Enthalpy Table
- Using Specific Heat
- Using Latent Heat of Vaporization
- Using Table(B,H)
- Using Tabular Data
- Using the power law
- Using the temperature shift factor
- Using the non-Newtonian generalized models:
  - Using the non-Newtonian generalized Power Law
  - Using the non-Newtonian generalized Cross Fluid
  - Using the non-Newtonian generalized Carreau-Yasuda Fluid
- Using the Schmidt number
- Using the volume-weighted mixture
- Using the mass-weighted mixture
- Using the Lewis number
- Using the PPDF Table
- Using the PCFM method
- Using the Mixture Method for Critical Temperature and Pressure
- Using the Mixture method for molecular weight
- Using the Mixture method for surface tension
- Using the Mathur-Saxena Averaging method for dynamic viscosity
- Using the Mathus-Saxena Averaging method for thermal conductivity

- Using the Elemental Composition Method
- Using the Anisotropic method
- Using the Kinetic Theory (Dars-CFD) Method for Molecular Diffusivity
- Using the Dars-CFD Viscosity Method
- Using the Dars-CFD Thermal Conductivity Method
- Using the Gas Kinetics Method for Specific Heat
- Using the Antoine equation
- Using the Wagner equation
- Using the Clausius-Clapeyron method
- Using the Heat of Formation
- Using the Leidenfrost Temperature
  - Calculating Saturation Temperature
- Using the Standard State Temperature
- Using the Kinetic Theory method for Molecular Diffusivity
- Using the Component Molecular Diffusivities method for Molecular Diffusivity
- Using the linear fraction solid curve method
- Using the Carman-Kozeny Mushy Permeability model
- Using the Metzner model
- Using the Nernst-Einstein method for Charged Species Mobility
- Using the Electrochemical Species method for Electrical Conductivity

To change the method, open the node for the relevant material property. In the *Properties* window, click the drop-down list corresponding to the Method property. This list permits the selection of Constant, Field Function, or one or more material property functions such as Power Law.



Material Property Node Properties

Dimensions

The dimensionality of the property that is expressed in terms of the basic dimensions, such as length, time (read-only). Accessing the customizer opens a *Dimensions* dialog to allow setting the dimensionality of a value or unit by setting the exponent of each of the base dimensions.



However, changing the dimensionality in this dialog has no effect on this property.

Method

The method that is used to define the value for this material property.

Constant

The material property is input as constant. A Constant node is added as a child to this node. Applicable to all material properties.

Field Function	The material property is entered using a field function. A Function node is added as a child to this node. Only applicable to some material properties.
Power Law	Specifies the value using the Power Law. A node by that name is added as a child to this node. Only applicable to some material properties.
Non-Newtonian Generalized Power Law	Specifies the value using the Non-Newtonian Generalized Power Law. A node by that name is added as a child to this node. Only applicable to dynamic viscosity when the models "Liquid" and "Laminar Flow" are selected.
IAPWS-IF97	The material property is calculated using IAPWS-IF97. A node by that name is added as a child to this node. Only applicable to some material properties.
Sutherland's Law	The material property is entered using Sutherland's Law. A node by that name is added as a child to this node. Only applicable to some material properties.
Table (Reaction Time)	The material property is entered using table (reaction time). A node by that name is added as a child to this node. Only applicable to some material properties.
Table(T)	The material property is entered using a table. A node by that name is added as a child to this node. Only applicable to some material properties.
Tabular Data	The material property is entered using a table interpolation method. A Tabular node is added as a child to this node. Only applicable to some material properties.
Polynomial in T	The material property is entered using a Polynomial in T. A node by that name is added as a child to this node. Only applicable to some material properties.
Schmidt Number	The material property is entered using the Schmidt number. A node by that name is added as a child to this node. Only applicable to some material properties.

Volume-Weighted Mixture	The material property is entered using the volume-weighted mixture. A node by that name is added as a child to this node. Only applicable to some material properties.
Mass-Weighted Mixture	The material property is entered using the mass-weighted mixture. A node by that name is added as a child to this node. Only applicable to some material properties.
Lewis Number	The material property is entered using the Lewis number. A node by that name is added as a child to this node. Only applicable to some material properties.
PPDF Table	The material property is entered using the PPDF Table. A node by that name is added as a child to this node. Only applicable to some material properties.
PCFM Table	The material property uses the PCFM Method. A node by that name is added as a child to this node. Only applicable to some material properties.
Mixture	The material property uses a Mixture. A node by that name is added as a child to this node. Only applicable to some material properties.
Anisotropic	The material property uses the Anisotropic method. A node by that name is added as a child to this node. Only applicable to some material properties.
Multi-Component (DarsCFD)	The material property uses the Multi-Component (Dars-CFD) method. A node by that name is added as a child to this node. Only applicable to some material properties.
Dars-CFD Viscosity Method	The material property uses the Dars-CFD Viscosity Method. A node by that name is added as a child to this node. Only applicable to some material properties.
Dars-CFD Thermal Conductivity Method	The material property uses the Dars-CFD Thermal Conductivity Method. A node by that name is added as a child to this node. Only applicable to some material properties.

Antoine Equation	The material property uses the Antoine Equation. A node by that name is added as a child to this node. Only applicable to the Saturation Pressure material property.
Wagner Equation	The material property uses the Wagner Equation. A node by that name is added as a child to this node. Only applicable to the Saturation Pressure material property.
Clausius-Clapeyron	The material property uses the Clausius-Clapeyron method. A node by that name is added as a child to this node. Only applicable to the Latent Heat of Vaporization material property with the Antoine Equation, Wagner Equation, or Polynomial in T methods for the Saturation Pressure.
Linear	The solid volume fraction changes linearly with temperature between the solidus and liquidus temperatures. (Eqn. (2133))
Carman-Kozeny	The material property uses the Carman-Kozeny Mushy Permeability method. A node by that name is added as a child to this node. Only applicable to the Mushy Zone Permeability material property with the Melting-Solidification model.
Metzner	The material property uses the Metzner method. A node by that name is added as a child to this node. Only applicable to the Slurry Viscosity material property with the Melting-Solidification model.
Nernst-Einstein	The material property uses the Nernst-Einstein method. A node by that name is added as a child to this node. Only available for the Charged Species Mobility Material Property.

See also: common menu items.

Plotting Variable Material Properties

When you work with a material property that uses a polynomial function, you can generate an approximate plot preview of the values before you run the simulation. This "pre-processing check" lets you verify that you have entered the values correctly, or otherwise observe the behavior of the polynomial.

To generate the plot preview:

- Select the node of a material property polynomial method. An example of such a node would be Continua > Physics 1 > Models > Gas > Air > Material Properties > Specific Heat > Polynomial in T.
- Click ... (Custom Editor) in the Polynomial property.

A customizer dialog that includes a graph opens.



As you modify the values, the graph changes in response to your entries.



You can also choose units. The Input Units property specifies units for the x-axis, while the Output Units property does so for the y-axis. Unlike constant values, the polynomial coefficients and polynomial ranges (and so x- and y-axes on the polynomial plot) do not rescale when you change the units. These unit selections serve as an indicator of the units that STAR-CCM+ uses when calculating quantities with this polynomial.

Using Constant Values

Generally, most initial conditions are specified as constants. This value is entered as a scalar value in the properties as shown below for Specific Heat.



Constant Value Properties

Constant value	The constant value.

Using Field Functions

Field functions can be used to create non-standard functions for material properties. Suppose that a scalar field function that is called Effective Conductivity is defined as

```
1.5207e-11 * pow($Temperature,3) - 4.857e-8 *
pow($Temperature,2) + 1.0184e-4 * $Temperature - 3.9333e-4
```

This field function would then be selected in the list corresponding to the Scalar Function property in the *Properties* window.



Field Function Value Properties

Scalar Function	The field function that is used to define this material property. A pull down list of available scalar functions is provided.

Using Polynomial in T

This polynomial function is available for certain material properties. The properties for this function are set in the *Properties* window.



Temperature polynomials are specified in piece-wise fashion where each interval i spans the range Tr_i to Tr_{i+1}, has m_i number of terms, and is of the form:

$$X(T)_p = \sum_{k=0}^{m_i} a_{i,k} \cdot T^{bi,k} \; ; \; Tr \le T \le Tr_{i+1} \tag{26}$$

The polynomial has n intervals with interval ranges:

$$Tr_0, Tr_1, Tr_2, \ldots, Tr_{n+1}$$

The number of terms in each interval are:

$$m_0, m_1, m_2, \ldots, m_n$$

The exponents are:

$$\begin{array}{cccc} b_{0,0}, & b_{0,1}, & b_{0,2}, \ldots, & b_{0,m_0}, \\ b_{1,0}, & b_{1,1}, & b_{1,2}, \ldots, & b_{1,m_1}, \\ \vdots & \vdots & \vdots & \vdots \\ b_{n,0}, & b_{n,1}, & b_{n,2}, \ldots, & b_{n,m_n} \end{array}$$

The coefficients for the polynomial are:

$$\begin{array}{cccc} a_{0,0}, & a_{0,1}, & a_{0,2}, \ldots, & a_{0,m_0}, \\ a_{1,0}, & a_{1,1}, & a_{1,2}, \ldots, & a_{1,m_1}, \\ \vdots & \vdots & \vdots & \vdots \\ a_{n,0}, & a_{n,1}, & a_{n,2}, \ldots, & a_{n,m_n} \end{array}$$

Specify a temperature range of the polynomial sufficiently large such that the polynomial is not evaluated outside the range during the analysis. If the polynomial is ever evaluated outside the range, the property is assigned a value of the polynomial at the respective range limit.

Polynomial in T Properties

In the *Properties* window, a single property that is called Polynomial includes a property customizer (...). Click that button to activate a customizer dialog with properties for working with polynomials in STAR-CCM+.



The plot in this dialog gets updated instantly with your changes. You can select units separately for each axis: Input Units for the x-axis, and Output Units for the y-axis. Unlike constant values, the polynomial coefficients and polynomial ranges (and the corresponding x- and y-axes on the polynomial plot) do not rescale when you change the units. These unit selections serve as an indicator of the units that STAR-CCM+ uses when calculating

quantities with this polynomial.

For the Coefficients property, commas separate the entries and the coefficients are in the following order:

$$C_p(T) = C_{p0} + (C_{p1} \cdot T) + (C_{p2} \cdot T^2) + \ldots + (C_{pn} \cdot T^n)$$

A separate section describes an example of working with these properties.

Using Thermodynamic Polynomial Data

Thermodynamic Polynomial Data provides five coefficient polynomial curve fits for C_p (along with the sixth coefficient polynomial curve fit for H and the seventh coefficient polynomial curve fit for S) that are read in from the selected material database (most often from the standard material database, `props.mdb`, provided as a part of STAR-CCM+). A variant of this method is read-only.



Since h_f^0 and T_{ref} are implicit in the curve fit for H, you do not need to supply them as you would if you were to use the Polynomial in T method for C_p.

The coefficients a_1 through a_5 for various temperature ranges define C_p as follows:

$$\frac{C_p}{R} = a_1 + a_2 T + a_3 T^2 + a_4 T^3 + a_5 T^4 \tag{27}$$

These coefficients can be set with a customizer dialog similar to the dialog used for polynomial properties. However, since it can be difficult to make them consistent, it is advisable not to change the default numbers.

The curve fit is:

$$\frac{H_i}{RT} = a_1 + \frac{a_2}{2}T + \frac{a_3}{3}T^2 + \frac{a_4}{4}T^3 + \frac{a_5}{5}T^4 + \frac{a_6}{T} \tag{28}$$

where:

$$\frac{H_i}{R} = \int_{T_{ref}}^{T} \frac{C_p(T)}{R} dT + h_f^0 \Big|_{T_{ref}} \tag{29}$$

Until version 5.04 the values of the polynomial coefficients were scaled with $R_u / M_r(X)$, where R_u is the universal gas constant [J/(kmol K)] and $M_r(X)$ is the molecular weight (kg/kmol) of the substance X. They appeared scaled in this manner in the GUI. Now the values appear unscaled, as they are found in the literature.

The coefficients for various temperature ranges define S as follows:

$$\frac{S}{R} = a_1 \ln T + a_2 T + \frac{a_3 T^2}{2} + \frac{a_4 T^3}{3} + \frac{a_5 T^4}{4} + a_7 \tag{30}$$

In the *Properties* window, a single property that is called NasaPolynomial includes a property customizer (...). Click that button to activate a customizer dialog with properties for working with polynomials in

STAR-CCM+.



The plot in this dialog gets updated instantly with your changes. Separate plots, accessible through tabs, exist for C_p, H and S.

Note: If STAR-CCM+ ever evaluates the polynomial outside the range, it assigns to the property a value of the polynomial at the respective range limit. Depending on the demands of the particular simulation, you can specify a temperature range of the polynomial sufficiently large such that the

polynomial is not evaluated outside the range during the analysis. However, sometimes the polynomials are not valid beyond the range that is given, so manipulate the temperature range with caution.

This section provides the following information about the file `thermdat.dbs`:

- Modifying a copy of the Chemkin database
- Loading a modified copy of the Chemkin database
- The Chemkin database format

Modifying a Copy of the Chemkin Database

The Chemkin database file (`thermdat.dbs`) is essentially a text file, and can be modified with any text editor.

Caution: The `thermdat.dbs` file is the source of the Chemkin database. Before modifying it, make a copy of this file for modification, and keep the original in a safe place.

The file is located as follows (these locations are examples using STAR-CCM+ version 7.04.006):

- On Unix:

  `[INSTALL.DIR]/cd-adapco/starccm+7.04.006/star`

- On Windows:

  `C:\Program Files\CD-adapco\STAR-CCM+ 7.04.006`

Open your copy of the `thermdat.dbs` file in a text editor.




```
THERMO
   300.000  1000.000  5000.000
(CH2O)3            70590C   3H   6O   3     G  0300.00   4000.00  1500.00      1
 0.01913678E+03 0.08578044E-01-0.08882060E-05-0.03574819E-08 0.06605143E-12    2
-0.06560876E+06-0.08432507E+03-0.04662286E+02 0.06091547E+00-0.04710536E-03    3
 0.01968843E-06-0.03563271E-10-0.05665404E+06 0.04525265E+03                   4
(CH3)2SICH2        61991H   8C   3SI  1     G  0300.00   2500.00  1500.00      1
 0.01547852E+03 0.01065700E+00-0.01234345E-05-0.01293352E-07 0.02528715E-11    2
-0.06693076E+04-0.05358884E+03 0.02027522E+02 0.04408673E+00-0.03370024E-03    3
 0.01484466E-06-0.02830898E-10 0.03931454E+05 0.01815821E+03                   4
AL                 62987AL  1               G  0300.00   5000.00  0600.00      1
 0.02559589E+02-0.01063224E-02 0.07202828E-06-0.02121105E-09 0.02289429E-13    2
 0.03890214E+06 0.05234522E+02 0.02736825E+02-0.05912374E-02-0.04033938E-05    3
 0.02322343E-07-0.01705599E-10 0.03886795E+06 0.04363880E+02                   4
AL2H6              62987AL  2H   6          G  0300.00   1500.00  0600.00      1
 0.02634884E+02 0.02135952E+00 0.03154151E-05-0.07684674E-07 0.02335832E-10    2
```

This editing technique lets you add new species as well as modify species properties.

Adding New Species to the Database

Gaseous species and condensed species with only one condensed phase can be in any order. However, the sets for two or more condensed phases of the same species must be adjacent. If there are more than two condensed phases of a species, their sets must be either in increasing or decreasing order according to their temperature intervals.

With the file open in the text editor, scroll to a set of data for any species, preferably one that belongs to the same phase as the species you wish to add.



Select that entire set of data and use the Copy command.

Then paste that data into an appropriate location, and change the entries so that they reflect the data of your species. In order for STAR-CCM+ to read your information correctly, ensure that the spacing is the same as that of other entries.



To complete the addition of a new species, enter data for your species in the proper positions. The details of how various types of data appear in the file are described in the section on the species database format.

Modifying Properties in the Database

Once you have located the value that you want to change, scroll to that value and change it to the new one. In order for STAR-CCM+ to read your information correctly, ensure that the spacing is the same as that of other entries.



Loading a Modified Copy of the Chemkin Database

After you have modified a copy of the Chemkin database, you can get STAR-CCM+ to recognize the changes by loading the `thermdat.dbs` file one of three ways:

- In the File Path property of the Tools > Materials Database > [material database] node, enter the absolute path of your `thermdat.dbs` file. This works in a manner similar to loading material databases.
- If you are using the Dars-CFD Reaction model, you can use a special dialog to import the Chemkin file. This feature also lets you use multiple `.dbs` files.
- In your file management program, move the original `thermdat.dbs` file and transfer the copy in its place, making sure that the copy has the same name, `thermdat.dbs`. Be careful not to overwrite one with the other.

Chemkin Database Format

The information in the `thermdat.dbs` file is organized in a precise structure so that STAR-CCM+ can read it into the simulation. In order for STAR-CCM+ to read the information correctly, any modifications must follow a precise syntax of position and spacing. Never modify the original `thermdat.dbs` file that is supplied with STAR-CCM+. Make a copy of the file for modifications and keep the original somewhere safe.

The following screenshot shows an example of the top of a `thermdat.dbs` file in a text editor.




```
THERMO
   300.000  1000.000  5000.000
(CH2O)3            70590C   3H   6O   3     G  0300.00   4000.00  1500.00      1
 0.01913678E+03 0.08578044E-01-0.08882060E-05-0.03574819E-08 0.06605143E-12    2
-0.06560876E+06-0.08432507E+03-0.04662286E+02 0.06091547E+00-0.04710536E-03    3
 0.01968843E-06-0.03563271E-10-0.05665404E+06 0.04525265E+03                   4
(CH3)2SICH2        61991H   8C   3SI  1     G  0300.00   2500.00  1500.00      1
 0.01547852E+03 0.01065700E+00-0.01234345E-05-0.01293352E-07 0.02528715E-11    2
-0.06693076E+04-0.05358884E+03 0.02027522E+02 0.04408673E+00-0.03370024E-03    3
 0.01484466E-06-0.02830898E-10 0.03931454E+05 0.01815821E+03                   4
AL                 62987AL  1               G  0300.00   5000.00  0600.00      1
 0.02559589E+02-0.01063224E-02 0.07202828E-06-0.02121105E-09 0.02289429E-13    2
 0.03890214E+06 0.05234522E+02 0.02736825E+02-0.05912374E-02-0.04033938E-05    3
 0.02322343E-07-0.01705599E-10 0.03886795E+06 0.04363880E+02                   4
AL2H6              62987AL  2H   6          G  0300.00   1500.00  0600.00      1
 0.02634884E+02 0.02135952E+00 0.03154151E-05-0.07684674E-07 0.02335832E-10    2
```

The file begins as follows:

- The top line starts the file with `THERMO`.
- The second line from the top contains temperature ranges for two sets of coefficients: lowest, common, and highest.

The body of the file contains data for the species. Each set of data for a particular species is composed of four lines, numbered `1` through `4` at the end:

- Line `1` consists of:
  - The name of the species
  - The date
  - The atomic symbols and formula of the species
  - The phase (`G` = gas; `L` = liquid; `S` = solid)
  - The temperature range

- Line 2 contains coefficients a_1 through a_5 for the upper temperature interval.
- Line 3 contains coefficients a_6 and a_7 for the upper temperature interval, and coefficients a_1 through a_3 for the lower temperature interval.
- Line 4 contains coefficients a_4 through a_7 for the lower temperature interval.

After the last set of numbered lines, the bottom line terminates the file with `END`.

Using IAPWS-IF97

The IAPWS-IF97 specification is available for various material properties [8], [9], [10].

This method node has no properties. It acts simply as an indicator that the associated material property follows the IAPWS-IF97 specification. This node is activated automatically when you select one of the IAPWS-IF97 continuum models.



Using Sutherland's Law

The Sutherland's law function is available for various material properties. The coefficient properties for this function are set in the *Properties* window.



Sutherland's law for viscosity is:

$$\frac{\mu}{\mu_0} = \left(\frac{T}{T_0}\right)^{\frac{3}{2}}\left(\frac{T_0 + S}{T + S}\right) \tag{31}$$

where T_0 and μ_0 are the reference temperature and viscosity, respectively, and S is the Sutherland constant.

Sutherland's law for thermal conductivity is:

$$\frac{k}{k_0} = \left(\frac{T}{T_0}\right)^{\frac{3}{2}}\left(\frac{T_0 + S}{T + S}\right) \tag{32}$$

where T_0 and k_0 are the reference temperature and conductivity, respectively, and S is the Sutherland constant.

Sutherland's Law Properties

Reference Temperature	The reference temperature in Eqn. (31) or Eqn. (32) (depending on the material property).
Reference Value	The reference viscosity or conductivity (depending on the material property) in Eqn. (31) or Eqn. (32).
Sutherland Constant	The Sutherland Constant value in Eqn. (31) or Eqn. (32) (depending on the material property).

Using Table (Reaction Time)

A (Reaction Time) table has a column of `Reaction Time` values that specify the physical time of the reaction, in seconds, of each point (row) in the table. The points must be ordered by increasing values of time. By default, the linear-piecewise method is used for interpolating between data points in the table. The properties for this function are set in the *Properties* window.

This type of table is only applicable to unsteady calculations using the Exothermic Solid model, where "Reaction Time" is defined as the accumulated physical time for which the temperature exceeds the autoignition time.



Note: A profile in Reaction Time can also be specified analytically with a user-defined field function that uses the `$ReactionTime` variable.

Table (Reaction Time) Properties

Table: Data	Identifies the selected column in the table file as the `Data` (for example, Position) column.
Table: Reaction Time	Identifies the selected column in the table file as the `Reaction Time` column.
Input Table	Selects a table from among the ones that appear as nodes in the Tables node of the Tools node.
Columns	Displays headings of columns in the table (read-only).

Table (Reaction Time) Expert Properties

Verbose	Provides more output for this value while the simulation is running. This property can be useful for debugging problems when they occur.

Using Table(T)

A (T) table has a column of `temperature` values for each point (row) in the table. By default, linear-piecewise interpolation method is used for interpolating between data points in the table. The properties for this

function are set in the *Properties* window.



Table(T) Properties

Table: Data	Identifies the selected column in the table file as the `Data` (for example, Thermal Conductivity) column.
Table: Temperature	Identifies the selected column in the table file as the `Temperature` column.
Input Table	Selects a table from among the ones that appear as nodes in the Tables node of the Tools node.
Columns	Displays headings of columns in the table (read-only).

Table(T) Expert Properties

Verbose	Provides more output for this value while the simulation is running. This property can be useful for debugging problems when they occur.

Using Table(T,P)

A (T,P) table has a column of `temperature` values and a column of `absolute pressure` values for each point (row) in the table. By default, linear-piecewise interpolation method is used for interpolating between data points in the table. The properties for this function are set in the

Properties window.



Table(T,P) Properties

Table: Data	Identifies the selected column in the table file as the `Data` (for example, density or enthalpy) column.
Table: Temperature	Identifies the table column containing the temperature values.
Table: Absolute Pressure	Identifies the table column containing the absolute pressure values.
Input Table	Selects a table from among the ones that appear as nodes in the Tables node of the Tools node.
Columns	Displays column headings in the table (read-only).

Using the Density Table Method for Density Derivatives

When density is defined by a table of data in temperature (or temperature and pressure for compressible flows), STAR-CCM+ can evaluate the following properties by numerically differentiating the table of data for density:

- The density pressure derivative, derived by differentiating the density table with respect to pressure;
- The density temperature derivative, derived by differentiating the density table with respect to temperature.

The User-Defined Equation of State model requires these properties. For an incompressible equation of state in the presence of an energy model, only the density temperature derivative is required. For a compressible equation of state, both the density pressure and density temperature derivatives are required.



Using the Enthalpy Table Method for Specific Heat

This material property method is available for specifying specific heat when the enthalpy is defined by a table of values in temperature (or in temperature and pressure for a compressible flow). When this option is selected, STAR-CCM+ calculates the specific heat by numerically differentiating the enthalpy table with respect to temperature at constant

pressure. This node has no properties.



Using the Specific Heat Method for Enthalpy

This material property method is available for specifying enthalpy when the User-Defined Equation of State model is active. When selected, STAR-CCM+ calculates the enthalpy by integrating specific heat. It is the default method that is used for calculating enthalpy in STAR-CCM+. The

Specific heat node has no properties.



Using the Enthalpy Difference Method for Heat of Vaporization

This material property method is available for specifying latent heat of vaporization when these models are active:

- Quasi-Steady Single-Component Droplet Evaporation
- User-Defined Single-Component Droplet Evaporation
- Multi-Component Droplet Evaporation

When this option is selected, STAR-CCM+ calculates the latent heat of vaporization directly as a difference in the enthalpy of the gas and liquid states:

$$L(T_{\mathrm{drop}}) = h_{\mathrm{vap}}(T_{\mathrm{drop}}) - h_{\mathrm{liq}}(T_{\mathrm{drop}}) \tag{33}$$

where:

- $L(T_{\mathrm{drop}})$ is the latent heat of the droplet material at the droplet temperature.
- $h_{\mathrm{vap}}(T_{\mathrm{drop}})$ is the static enthalpy of the vapor component corresponding to the droplet material, evaluated at the droplet temperature.

- $h_{liq}(T_{drop})$ is the static enthalpy of the droplet material, evaluated at the droplet temperature.



Using Table(B,H) Method for Permeability

This material property provides a stable iterative scheme for setting magnetic permeability for both convex and concave B-H curves. This method is a preferred method for setting permeability when using a table over the tabular data method.

If you want to select the B-H-mu table that is in the material database, you use this method. The tabular data method does not let you select the permeability table from the material database.

To use Table(B,H), input the table and the columns corresponding to the magnetic field H and the magnetic flux density B.



The first entry in the table must be (0,0). When $\boldsymbol{B}$ becomes greater than the maximum flux density that is provided in the table, the B-H curve is extended at a slope of μ_0 from the last point.

The provided table for this method can look as follows:



	H[A/m]	B[T]
0	0	0
1	0.6848	0.021
2	2.4723	0.042
3	4.962	0.063
4	7.7533	0.084
5	10.4469	0.105
6	12.8003	0.1261
7	14.8842	0.1471
8	16.8062	0.1681
9	18.6742	0.1891
10	20.5934	0.2101
11	22.6048	0.2311

Using Tabular Data

The tabular data method is used to provide values of a material property. The multi-dimensional bilinear interpolation method is used for interpolating between data points in the table. As an example, the tabular data method is described here for the magnetic permeability material property of the Magnetic Vector Potential model.

If you want to select the B-H-mu table that is in the material database, you use the Table(B,H) method. The tabular data method does not let you select the permeability table from the material database.

The tabular data method works as follows:

- Select the Tabular > Table node and, in the *Properties* window, select the table containing the material property data.



- Select the Table > Interpolated node and, in the *Properties* window, set the Column title that matches the material property to be returned by the method.
- Select the input quantity against which the material property is interpolated:
  - Right-click the Table > Variables node and choose New.

    A new Variable child node is added.
  - Select the newly added child node. In the *Properties* window, set the Scalar Function property to the desired quantity. Set the Column title property to the table column that contains values of this quantity.

Setting Magnetic Permeability Property

Following the above procedure, first set the Input table property to the B-H table. The B-H table usually includes three columns: one for the magnetic flux density $\mathbf{B}$, one for the magnetic field $\mathbf{H}$ and one for the calculated

permeability μ that is determined through the relation $\mu = \mathbf{B}/\mathbf{H}$. An example table is shown below. Only two columns are needed for the tabular data method. For this case, either **H**-μ or **B**-μ is used.



	H[A/m]	B[T]	mu[H/m]
0	0.6848349	0.0210084	0.0306765908
1	2.472334	0.04201681	0.0169947952
2	4.961991	0.06302521	0.012701597
3	7.753296	0.08403361	0.010838437
4	10.44691	0.105042	0.0100548392
5	12.80034	0.1260504	0.0098474259
6	14.88418	0.1470588	0.0098802084
7	16.80621	0.1680672	0.0100003035
8	18.67421	0.1890756	0.0101249584
9	20.59343	0.210084	0.010201506
10	22.60479	0.2310924	0.0102231607

Columns that are extracted from this table appear in the Columns property of the Table node.

Having selected the table, choose the mu[H/m] column for the Column title property of the Interpolated child node. This column matches the Magnetic Permeability material property for which the Tabular method has been selected.



Finally, choose either the **H** or **B** scalar field function against which permeability is interpolated. Determining which of these variables to use as the interpolation variable has a direct effect on the convergence of the magnetic vector potential solution (as discussed in the troubleshooting section of electromagnetism modeling).

Here, create a variable and set its Scalar Function property to Magnetic Field: Magnitude, and its Column title property to H[A/m].



Using the Power Law

The Power law function is available for various material properties. The coefficient properties for this function are set in the *Properties* window.



The Power law for viscosity is:

$$\frac{\mu}{\mu_0} = \left(\frac{T}{T_0}\right)^n \tag{34}$$

where T_0 and μ_0 are the reference temperature and viscosity, respectively and n is the Power law exponent.

The Power law for thermal conductivity is:

$$\frac{k}{k_0} = \left(\frac{T}{T_0}\right)^n \tag{35}$$

where T_0 and k_0 are the reference temperature and conductivity, respectively and n is the Power law exponent.

Power Law Properties

Reference Temperature	The reference temperature in Eqn. (34) or Eqn. (35) (depending on the material property).
Reference Value	The reference viscosity or conductivity (depending on the material property) in Eqn. (34) or Eqn. (35).
Temperature Exponent	The temperature exponent in Eqn. (34) or Eqn. (35) (depending on the material property).

Using the Temperature Shift Factor

When temperature effects are present, the non-Newtonian liquid models use a shift factor a_T to adjust viscosity with temperature.



When the simulation is isothermal, $a_T = 1$. For polymers:

$$a_T = \frac{\eta_T}{\eta_{T0}}$$

where η is the viscosity at temperature T and T_0.

The shift factor can be calculated by the Arrhenius model or the Williams-Landel-Ferry (WLF) model.

For the Arrhenius model:

$$\log(a_T) = \frac{E_a}{R}\left(\frac{1}{T} - \frac{1}{T_0}\right)$$

where:

E_a is the activation energy.

R is the universal gas constant.

T_0 is a reference temperature in K.

For the Williams-Landel-Ferry (WLF) model:

$$\log(a_T) = -\frac{C_1(T - T_0)}{C_2 + (T - T_0)}$$

where C_1 and C_2 are positive constants that depend on the material and the reference temperature. If the reference temperature is taken as the glass transition temperature T_g then C_1 and C_2 are universal for a given material. For any polymer, approximate values are $C_1 \approx 15°C$ and $C_1 \approx 50°C$.

Using the Non-Newtonian Generalized Models

To model a non-Newtonian liquid, STAR-CCM+ supports the following methods:

- Non-Newtonian Generalized Power Law
- Non-Newtonian Generalized Cross Fluid
- Non-Newtonian Generalized Carreau-Yasuda Fluid

Using the Non-Newtonian Generalized Models in Turbulent Flows

The Non-Newtonian Generalized models are now available for turbulent flows; previously they were available for laminar flows only.

Note that this approach is not strictly valid, as the turbulence models have been designed for Newtonian fluids; the production at the wall is known to be different for non-Newtonian fluids. Also the turbulence models do not account for any viscoelasticity of the flow. However, in many cases, the results that you obtain from using the (incorrect) turbulence model can be better than the results that you would obtain by assuming a laminar flow.

If you use the Non-Newtonian Generalized models in a turbulent flow, you do so at your own risk.

Using the Non-Newtonian Models in Multiphase Flows

The Non-Newtonian Generalized models are available for multiphase flows. However, the following restrictions apply:

- The Non-Newtonian Generalized Cross Fluid model is not compatible with the Eulerian Multiphase model.
- The dispersed phase cannot be specified as a Non-Newtonian fluid. Consequently, you cannot define any interphase interactions between Non-Newtonian phases (which have to be continuous).

Selecting a Non-Newtonian Model

The graph below shows a comparison of the models from [7]. The three-parameter Cross model and the four-parameter Carreau-Yasuda model both display good agreement with the experimental data for this shear-thinning material.

However, these models require you to provide the values of zero-shear and infinite shear at the Newtonian plateaus. In contrast, the two-parameter Power Law requires fewer inputs and has good agreement in the main-shear thinning (or shear thickening) zone but loses accuracy as the fluid approaches a Newtonian plateau.



Non-Newtonian Generalized Model Formulation

All of the models have the form of the stress tensor:

$$\boldsymbol{\tau}(\dot{\gamma}) = 2\mu(\dot{\gamma})\mathbf{D} \tag{36}$$

where:

$\boldsymbol{\tau}(\dot{\gamma})$ is the dynamic shear viscosity as a function of shear strain-rate. Each method gives $\boldsymbol{\tau}(\dot{\gamma})$.

$\mathbf{D}$ is the rate of deformation tensor.

The rate of deformation tensor is:

$$\mathbf{D} = \frac{1}{2}(\nabla u + (\nabla u)^T) \tag{37}$$

The shear strain rate can be calculated from the deformation tensor:

$$\dot{\gamma} = \sqrt{(2\mathbf{D}:\mathbf{D})} \tag{38}$$

Note: The strain rate in Eqn. (38) is only valid for pure shear flows -- that is, simple shear and pipe flow. In a purely extensional flow field, the strain rate that is calculated from Eqn. (38) over-predicts the extensional strain by a factor of $\sqrt{3}$, which causes under-prediction of the viscosity (for shear-thinning behavior). This situation is a limitation of the generalized models, so treat with caution any flow that includes stretching or accelerating regions.

Using the Non-Newtonian Generalized Power Law

The non-Newtonian Generalized Power Law for viscosity is defined as:

$$\mu_{min} < \left[\begin{array}{c} a_T(\mu_0) \text{when } (\dot{\gamma} \leq (\tau_0/\mu_0)) \\ a_T\left(\dfrac{\tau_0 + k(a_T(\dot{\gamma} - (\tau_0/\mu_0)))^n}{a_T\,\dot{\gamma}}\right) \text{when } (\dot{\gamma} > (\tau_0/\mu_0)) \end{array} \right] < \mu_{max} \tag{39}$$

where:

k is the consistency factor

n is the power law exponent

τ_0 is the yield stress threshold

μ_0 is the yielding viscosity

μ_{min} is the minimum viscosity limit

μ_{max} is the maximum viscosity limit

γ is the strain rate magnitude

a_T is the temperature shift factor.

When the strain rates are smaller than (τ_0/μ_0) the dynamic viscosity is set to μ_0, otherwise the viscosity is given by the power law expression. In both cases, the viscosity is bounded by μ_{min} and μ_{max}.

This law is also known as the Herschel-Bulkley model for Bingham plastics, and is commonly used to model materials like concrete, mud, and toothpaste. This method reduces to the Ostwald de Waele law or the standard power law, when the yield stress threshold is set to zero.

The value of the power law exponent n determines the class of the fluid -

- $n = 1$ -> Newtonian fluid
- $n > 1$ -> Shear-thickening (dilatant) fluid
- $n < 1$ -> Shear-thinning (pseudo-plastic) fluid

For strongly non-linear cases, use a low value of the viscosity under-relaxation factor (between 0 and 1) in order to under-relax the viscosity and prevent solution divergence.

This method is available for the material property Dynamic Viscosity for both laminar and turbulent liquid flow. It can be used with single phase and multiphase flows. The coefficient properties for this function are set in the *Properties* window.



Non-Newtonian Generalized Power Law Properties

Consistency Factor	The value k in Eqn. (39). The default value is 1.0
Power Law Exponent	The value n in Eqn. (39). The default value is 1.0
Yield Stress Threshold	The value τ_0 in Eqn. (39). The default value is 0.0 Pa
Yielding Viscosity	The value μ_0 in Eqn. (39). The default value is 1.0 Pa-s

Non-Newtonian Generalized Power Law Expert Properties

Minimum Viscosity Limit	The value μ_{min} in Eqn. (39). The default value is 0.0 Pa-s
Maximum Viscosity Limit	The value μ_{max} in Eqn. (39). The default value is 10000.0 Pa-s
Viscosity Under-Relaxation Factor	The under-relaxation factor for dynamic viscosity. The default value is 0.7

Using the Non-Newtonian Generalized Cross Fluid

The non-Newtonian Generalized Cross Fluid for viscosity is defined from [6] as:

$$\mu(\dot{\gamma}) = a_T\left(\mu_\infty + \frac{\mu_0 - \mu_\infty}{1 + \left(\frac{a_T\dot{\gamma}}{\dot{\gamma}_c}\right)^m}\right) \tag{40}$$

where:

m is the Cross rate constant.

μ_0 is the zero-shear viscosity, the viscosity at the Newtonian plateau.

μ_∞ is the infinite-shear viscosity.

$\dot{\gamma}_c$ is the critical shear strain-rate at which shear-thinning starts.

a_T is the temperature shift factor.

Newtonian behavior is regained for $m = 0$. As m increases towards unity, the degree of shear-thinning becomes stronger.

This method is available for the material property Dynamic Viscosity for both laminar and turbulent liquid flow. This method is not available for multiphase flows. The coefficient properties for this function are set in the *Properties* window.



Non-Newtonian Generalized Cross Fluid Properties

Cross Rate Constant	The value m in Eqn. (40). The default value is 1.5
Zero Shear Viscosity	The value μ_0 in Eqn. (40). The default value is 0.05 Pa-s
Infinite Shear Viscosity	The value μ_∞ in Eqn. (40). The default value is 0.005 Pa-s
Critical Shear Rate	The value $\dot{\gamma}_c$ in Eqn. (40). The default value is 2.0 /s

Non-Newtonian Generalized Cross Fluid Expert Properties

Viscosity Under-Relaxation Factor	The under-relaxation factor for dynamic viscosity. The default value is 0.7

Using the Non-Newtonian Generalized Carreau-Yasuda Fluid

The non-Newtonian Generalized Carreau-Yasuda Fluid for viscosity comes from [5] and [16]. It is defined as:

$$\mu(\dot{\gamma}) = a_T\left(\mu_\infty + (\mu_0 - \mu_\infty)(1 + (\lambda a_T \dot{\gamma})^a)^{(n-1)/a}\right) \tag{41}$$

where:

n is the power constant

μ_0 is the zero-shear viscosity, the viscosity at the Newtonian plateau

μ_∞ is the infinite-shear viscosity

λ is the relaxation time constant

a is the parameter to control shear-thinning. When $a = 2$, the original Carreau model is regained.

$\dot{\gamma}_c$ is the critical shear strain-rate at which shear-thinning starts.

a_T is the temperature shift factor.

This method is available for the material property Dynamic Viscosity for both laminar and turbulent liquid flow. It can be used with single phase and multiphase flows. The coefficient properties for this function are set in the *Properties* window.



Non-Newtonian Generalized Carreau-Yasuda Fluid Properties

Power Constant	The value n in Eqn. (41). The default value is 0.5
a Parameter	The value a in Eqn. (41). The default value is 2.0
Zero Shear Viscosity	The value μ_0 in Eqn. (41). The default value is 0.05 Pa-s
Infinite Shear Viscosity	The value μ_∞ in Eqn. (41). The default value is 0.005 Pa-s
Relaxation Time	The value λ in Eqn. (41). The default value is 1.5 s

Non-Newtonian Generalized Carreau-Yasuda Fluid Expert Properties

Viscosity Under-Relaxation Factor	The under-relaxation factor for dynamic viscosity. The default value is 0.7

Using the Schmidt Number

The Schmidt number is a non-dimensional parameter that is defined as:

$$S_c = \frac{\nu}{D_{im}} \tag{42}$$

where ν is the kinematic viscosity and D_{im} is the molecular diffusivity of component-i into the mixture.

The Schmidt number function is available for specifying molecular diffusivity of components in a mixture.



The properties for this method are set in the *Properties* window.

Schmidt Number Properties

Schmidt Number	The value of Schmidt number that is used for specifying molecular diffusivity.

Using a Volume-Weighted Mixture

The volume-weighted mixture function is for non-ideal gas mixtures.

Using this function, a given mixture property ϕ_{mix} is defined as:

$$\phi_{mix} = \frac{1}{\sum \left(\frac{y_i}{\phi_i} \right)} \tag{43}$$

where y_i and ϕ_i are the mass fraction and individual material property of component-i in the mixture of N components.

This method is not available for conductivity or viscosity in single-phase multi-component mixtures. It is, however, available for these properties in a multiphase mixture.



Volume-Weighted Mixture Properties

None

Using the Mass-Weighted Mixture

The mass-weighted mixture method calculates a given mixture property by mass-weighting the component property values.



For example, mixture property ϕ_{mix} is:

$$\phi_{mix} = \sum_{i=1}^{N} y_i \phi_i \tag{44}$$

where y_i and ϕ_i are the mass fraction and property values of mixture component-i and N is the total number of components in the mixture.

Mass-Weighted Mixture Properties

None

Using the Lewis Number

The Lewis number L_e is defined as the ratio of thermal diffusion to mass diffusion D_m, expressed as follows:

$$L_e \equiv \frac{k/\rho C_p}{D_m} \tag{45}$$

where ρ, C_p, and k are the density, specific heat, and thermal conductivity of the mixture, respectively.

Given values for L_e, D_m, ρ, and C_p, thermal conductivity k can be calculated as follows:

$$k = L_e D_m \rho C_p \tag{46}$$

The properties for this method are set in the *Properties* window.



Lewis Number Properties

Lewis Number	The value of Lewis number that is used for specifying the ratio of thermal diffusion to mass diffusion.

Using the PPDF Table

The PPDF Table method interpolates property values from the PPDF table.



The Molecular Weight and Specific Heat material properties of the mixture use this method. It is the only method available for these properties when the Presumed Probability Density Function (PPDF) model is active.

PPDF Table Properties

None

Expert Properties

Property Name	Identifies the material property that is being interpolated from the table (read-only).

Using the PCFM Method

The PCFM method becomes active when the Partially-Premixed Coherent Flame Model (PCFM) is used.



The Molecular Weight and Specific Heat material properties of the mixture use this method. It is the only method available for these properties when the PCFM model is active.

PCFM Method Properties

None

Expert Properties

Property Name	Identifies the material property (read-only).

Using the Mixture Method for Critical Temperature and Pressure

The Mixture method calculates the critical temperature $T_{c,m}$ of a mixture as:

$$T_{c,m} = \frac{\sum_i T_{c,i} Y_i / W_i}{\sum Y_i / W_i} \tag{47}$$

where W_i is the molecular weight, Y_i the mass fraction and $T_{c,i}$ the critical temperature of species i. The summation is over all species in the mixture. The critical pressure $P_{c,m}$ of a mixture is similarly calculated as:

$$P_{c,m} = \frac{\sum_i P_{c,i} Y_i / W_i}{\sum Y_i / W_i} \tag{48}$$

where $P_{c,i}$ is the critical pressure of component i.



Properties

None.

Using the Mixture Method for Molecular Weight

When the Mixture method is used for molecular weight, the mixing law is:

$$W_m = \frac{1}{\sum_i Y_i / W_i} \tag{49}$$

where W_i is the molecular weight of species i, Y_i is the mass fraction of species i, and the summation is over all species in the mixture.



Properties

None.

Using the Mixture Method for Surface Tension

When the Mixture method is used for surface tension, the mixture surface tension is:

$$\sigma_m^r = \sum_i X_i \sigma_i^r \tag{50}$$

where σ_i is the surface tension of component i, X_i is its mole fraction, and r is the exponent property of the mixing law.

Mixture Method Properties

Exponent	The exponent property *r* of the mixing law that is used to determine the surface tension of a mixture (Eqn. (50)).

Using the Mathur-Saxena Averaging Method for Dynamic Viscosity

This method becomes available when either the EBU or Homogeneous Reactor combustion model is used. Selecting this method provides the choice of the Chapman-Enskog method for Dynamic Viscosity on the Mixture Component level.

The following properties show up on the mixture component level:

- Dipole momentum -- a measure of polarity of a covalent bond in the molecule (always given in Debye)
- Lennard-Jones Potential -- consists of two parameters:
  - Collision diameter (always given in Angstroms)
  - Potential energy of attraction (J/kmol)

These properties have only one method available for their calculation: Constant. These values are used for calculating the dynamic viscosity of the individual component for the Chapman-Enskog method.

To expose the Mathur-Saxena Averaging method, choose the Mathur-Saxena Averaging option for the Method property of the Dynamic Viscosity node. This node is contained within the Mixture Properties manager for a multi-component gas. The Dynamic Viscosity node contains the Mathur-Saxena Averaging node as shown.



The Mathur-Saxena Averaging methods for Dynamic Viscosity and Thermal Conductivity are specifications of the Multi-Component Mathur-Saxena Averaging Property method $<\Phi>$ where Φ is a material property. This method uses the following formula to compute the property on the mixture level, using values given for the individual mixture components (X_i is the mole fraction of component i):

$$\Phi = \frac{1}{2}\left(\sum_{1}^{n} X_i \Phi_i + \left(\sum_{1}^{n} \frac{X_i}{\Phi_i}\right)^{-1}\right) \tag{51}$$

When the option is set, each component within the Gas Components manager has a Dynamic Viscosity node that contains the Chapman-Enskog node. See Using the Chapman-Enskog Method for Dynamic Viscosity



Using the Chapman-Enskog Method for Dynamic Viscosity

This method becomes available for single- and multi-component gases once an energy model is included. This method is also activated when the Mathur-Saxena Averaging method is selected for dynamic viscosity.



The dynamic viscosity of a simple gas, or of an individual component of a multi-component gas, is calculated with the Chapman-Enskog equation as follows:

$$\mu_i = 2.6693 \times 10^{-6} \frac{\sqrt{M_i T}}{\sigma_i^2 \Omega(T^*)} \tag{52}$$

where:

- $T^* = kT/\varepsilon_i$ is reduced temperature (dimensionless)
- k is the Boltzmann constant
- ε_i is the potential energy of attraction of component i; values given in the standard material database, `props.mdb`, provided as a part of STAR-CCM+
- $\Omega(T^*)$ is the collision integral
- μ_i is the viscosity of component i
- M_i is the molecular weight of component i
- T is temperature (in K)
- σ_i is the collision diameter of component i; values given in the standard material database, `props.mdb`, provided as a part of STAR-CCM+

For a simple gas, there is only one value of i.

Using the Mathur-Saxena Averaging Method for Thermal Conductivity

This method becomes available when either the EBU or Homogeneous Reactor combustion model is used. Selecting this method activates the choice of the Kinetic Theory method for Thermal Conductivity on the Mixture Component level.

The following properties show up on the mixture component level:

- Dipole momentum—a measure of polarity of a covalent bond in the molecule (always given in Debye)
- Lennard-Jones Potential—consists of two parameters:
  - Collision diameter (always given in Angstroms)
  - Potential energy of attraction
- Rotation—rotational relaxation collision number (non-dimensional)
- Molecule Type—describes molecular structure; can be one of the following:
  - 0: atom
  - 1: linear molecule
  - 2: nonlinear molecule

All these properties have only one method available for their calculation: Constant. These values are used for calculating the thermal conductivity of the individual component for the Kinetic Theory method.

To expose the Mathur-Saxena Averaging method, choose the Mathur-Saxena Averaging option for the Method property of the Thermal Conductivity node. This node is contained within the Mixture Properties manager for a multi-component gas. The Thermal Conductivity node contains the Mathur-Saxena Averaging node as shown.



The Mathur-Saxena Averaging methods for Dynamic Viscosity and Thermal Conductivity are specifications of the Multi-Component Mathur-Saxena Averaging Property method $<\Phi>$; see Eqn. (51).

When the option is set, each component within the Gas Components manager has a Thermal Conductivity node that contains the Kinetic Theory node. See Using the Kinetic Theory Method for Thermal Conductivity.



Using the Kinetic Theory Method for Thermal Conductivity

This method becomes available for single- and multi-component gases. This method is also activated when the Mathur-Saxena Averaging method is selected for thermal conductivity.



The thermal conductivity of a simple gas, or of an individual component of a multi-component gas, is calculated as follows:

$$\lambda_i = \frac{\mu_i}{M_i}(f_{trans}C_{v,\,trans} + f_{rot}C_{v,\,rot} + f_{vib}C_{v,\,vib}) \tag{53}$$

where:

$$f_{trans} = \frac{5}{2}\left[1 - \frac{2}{\pi}\frac{C_{v,\,rot}}{C_{v,\,trans}}\frac{A}{B}\right] \tag{54}$$

$$f_{rot} = \frac{\rho D_{ii}}{\mu_i}\left[1 + \frac{2}{\pi}\frac{A}{B}\right] \tag{55}$$

$$f_{vib} = \frac{\rho D_{ii}}{\mu_i} \tag{56}$$

$$A = \frac{5}{2} - \frac{\rho D_{ii}}{\mu_i} \tag{57}$$

$$B = z_{rot} + \frac{2}{\pi}\left(\frac{5}{3}\frac{C_{v,\,rot}}{R_u} + \frac{\rho D_{ii}}{\mu_i}\right) \tag{58}$$

$$C_{v,\,trans} = \frac{3}{2}R_u \tag{59}$$

$$z_{rot} = z_{rot}(298)\frac{F(298)}{F(T)} \tag{60}$$

$$F(T) = 1 + \frac{\pi^{3/2}}{2}\left(\frac{\varepsilon/k}{T}\right)^{1/2} + \left(\frac{\pi^2}{4} + 2\right)\frac{\varepsilon/k}{T} + \pi^{3/2}\left(\frac{\varepsilon/k}{T}\right)^{3/2} \tag{61}$$

and $z_{rot}(298)$ is the value of the rotation collision number at 298 K, defined in the transport properties of the component.

For a linear molecule:

$$C_{v,\,rot} = R_u \tag{62}$$

$$C_{v,\,vib} = C_v - \frac{5}{2}R_u \tag{63}$$

and:

- μ_i is the viscosity of component i, $\mu_i = \mu_i(T)$.
- M_i is the molecular weight of component i.
- λ_i is the thermal conductivity of component i.
- C_v is the molar specific heat at constant volume of the component.
- R_u is the universal gas constant.

For a simple gas, there is only one value of i.

Using the Elemental Composition Method

In order to calculate the molecular weight of a molecule, this method is used to describe its elemental composition. The advantages of this method (as opposed to setting the molecular weight by a constant) are that:

1. The molecular weight is automatically updated when you change the composition of the molecule; and,
2. No round-off is applied to the result. Using the constant method for molecular weight incurs a round-off to the second decimal place.

The constant method is still provided for use with those materials that do not have an elemental composition (such as air), and also to provide backwards compatibility with simulations created in earlier versions of STAR-CCM+.

The Elemental Composition Method node contains a single child node, Atoms, that is used to manage the individual substances that make up the molecule.



In most instances, the Atoms node is populated automatically based on information that is contained in the materials property database. It is, however, possible to populate this node manually.

Populating the Composition

Individual substances are added to the composition by right-clicking on the Atoms node and choosing New Atom from the pop-up menu.



Choosing New Atom opens the *Add atom to elemental composition* dialog, which you can use to set the substance and its quantity in the molecule. A list of available substances is provided in the drop-down menu, and an entry box accepts the number of atoms of the chosen substance



Clicking OK completes the specification and add a corresponding Atom child node to the Atoms manager node.

Removing Substances

Substances can be removed from the composition by right-clicking on their respective Atom nodes and choosing Delete.

What Is an Atom?

The Atom node, which has properties and is subject to common menu items, represents a single substance and its quantity within an elemental composition.



Atom Properties

Value	The number of atoms of this substance in the composition.

Working with Atoms

The Atoms node, which is a child of the Elemental Composition Method node, contains the individual substances (represented by Atom nodes) that make up a particular molecule. It also provides the molecular weight as a read-only property. A pop-up menu is used to populate the node.



The molecular weight W_m of the composition is calculated by summing up the individual molecular weights of each atom multiplied by the atom's quantity, that is:

$$W_m = \sum_i n_i W_i$$

where W_i is the molecular weight of a constituent atom and n_i is the number of such atoms in the molecule.

Atoms Properties

Molecular Weight	The calculated molecular weight of the composition (read-only).

Atoms Pop-Up Menu

New Atom	Opens the *Add atom to elemental composition* dialog for choosing an atom and its quantity.

Using the Anisotropic Method

The Anisotropic method is a placeholder object to indicate that an anisotropic approach to a material property is being used.

This method is available for the following material properties:

- Permittivity of a solid when the Electrostatic Potential model is active.
- Electrical Conductivity of a solid when the Electrodynamic Potential model or Magnetic Vector Potential model is active.
- Thermal Conductivity of a solid when the Coupled Solid Energy or Segregated Solid Energy model is active.

The example below is for the Electrical Conductivity material property of a solid.



The selection of this method adds the corresponding Physics Values node to the solid region of the simulation. The example below shows the Anisotropic Electrical Conductivity node.



Anisotropic Method Properties

Method	The anisotropic value is entered as a tensor profile.

Using the Kinetic Theory (Dars-CFD) Method for Molecular Diffusivity

This method provides the molecular diffusivity of components in a mixture using the kinetic theory calculations in Dars-CFD.



A specific routine in the Dars-CFD package is called to calculate molecular diffusivities of the mixture components. The diffusivity of the species, D_k is related to binary diffusivities D_{ki} as:

$$D_k = \frac{1 - X_k}{\sum_i X_i / D_{ki}} \tag{64}$$

where the mole fraction X_k is related to the mass fraction Y_k as:

$$X_k = \frac{Y_k W}{W_k} \tag{65}$$

W is the molecular weight of the mixture and W_k is the molecular weight of the species k. The binary diffusion coefficients are obtained through a polynomial fit using binary diffusivity values that are calculated for a range of temperatures:

$$D_{ki} = [p \cdot \exp(a_{1,ki} + a_{2,ki}(\log T) + a_{3,ki}(\log T)^2 + a_{4,ki}(\log T)^3)]^{-1} \tag{66}$$

Kinetic Theory (DarsCFD) Properties

None

Using the Dars-CFD Viscosity Method

Dynamic viscosity through Dars-CFD is available for specifying the dynamic viscosity of a mixture.



The dynamic viscosity of individual components is calculated with a variant of the Chapman-Enskog equation, Eqn. (52).

The dynamic viscosity of the mixture is given with the following equation:

$$\mu = \frac{1}{2}\left(\sum_{1}^{n} X_i \mu_i + \left(\sum_{1}^{n} \frac{X_i}{\mu_i}\right)^{-1}\right) \tag{67}$$

where:

- μ is the viscosity of the mixture
- μ_i is the viscosity of component i
- X_i is the mole fraction of component i

Dars-CFD Viscosity Properties

None

Using the Dars-CFD Thermal Conductivity Method

Thermal conductivity through Dars-CFD is available for specifying the thermal conductivity of a mixture.



The thermal conductivity of individual components is calculated from kinetic theory. The thermal conductivity of the mixture is given with the following equation:

$$\lambda = \frac{1}{2}\left(\sum_{1}^{n} X_i\lambda_i + \left(\sum_{1}^{n}\frac{X_i}{\lambda_i}\right)^{-1}\right) \tag{68}$$

where:

- λ is the conductivity of the mixture
- λ_i is the conductivity of component i
- X_i is the mole fraction of component i

Dars-CFD Thermal Conductivity Properties

None

Using the Gas Kinetics Method for Specific Heat

This method becomes available with single- and multi-component gases, together with the Ideal Gas, Constant Density, or Thermal Non-Equilibrium Ideal Gas models, or with any of the Real Gas models: Peng-Robinson, Redlich-Kwong, Soave-Redlich-Kwong, Modified Soave-Redlich-Kwong, and Van der Waals).



The specific heat of a gas consists of contributions from translational, rotational, vibrational and electronic energy [15]:

$$C_{p,\,total} = C_{p,\,trans} + C_{v,\,rot} + C_{v,\,vib} + C_{v,\,elec} \tag{69}$$

where:

- C_p values are specific heats at constant pressure.
- C_v values are specific heats at constant volume.

For each energy type, the contribution to the specific heat is calculated based on the associated partition function, through the following formula:

$$\begin{aligned} \frac{e_i}{R} &= T^2 \frac{\partial}{\partial T} \ln(Q_i) \\ \frac{C_{v,\,i}}{R} &= \frac{T^2}{Q_i}\left[\frac{\partial^2 Q_i}{\partial T^2} + \left(\frac{\partial Q_i}{\partial T}\right)^2 + \frac{2}{T}\frac{\partial Q_i}{\partial T}\right] \end{aligned} \tag{70}$$

where:

- e_i is energy of type i.
- i is $trans$ (translational), rot (rotational), vib (vibrational), or $elec$ (electronic)
- R is the gas constant.
- T is the temperature.
- Q_i is the partition function for energy of type i.

To determine the total specific heat, the partition function for each energy type must be specified. The combined translational and rotational energy $C_{p,\,TR} = C_{p,\,trans} + C_{v,\,rot}$ can be simplified because:

- For the translational energy, the specific heat is effectively constant due to the small spacing between quantum energy levels.
- The rotational energy is assumed fully excited, since the spacing between quantum energy levels is negligible at all but cryogenic temperatures. Rotational energy is zero for monatomic molecules.

The combined translational-rotational specific heat is given as:

$$\begin{aligned} C_{p,\,TR} &= \frac{5}{2}R && \text{for monatomic molecules} \\ C_{p,\,TR} &= \frac{7}{2}R && \text{for linear molecules} \\ C_{p,\,TR} &= \frac{8}{2}R && \text{for non-linear molecules} \end{aligned} \tag{71}$$

The partition function for electronic energy must be specified explicitly for this method. For polyatomic molecules, the partition function for vibrational energy must also be specified. These functions are both specified using a series of modes with characteristic temperature (θ_i) and degeneracy (g_i). The partition function for each of these modes is defined as:

$$Q_{vib} = \prod_i (1 - e^{-\theta_i / T})^{-g_i} \qquad \text{for polyatomic molecules}$$

$$Q_{el} = \sum_i g_i e^{-\theta_i / T} \qquad \text{for all molecules} \tag{72}$$

These degeneracies and characteristic temperatures can be found for various molecules in [11] and [14]. For N_2, O_2, NO, N, and O, these values are in the standard material database provided as part of STAR-CCM+.

Using the Antoine Equation

The Antoine equation is available for the Saturation Pressure material property. The coefficient properties for this equation are set in the *Properties* window.



The Antoine equation for saturation pressure [13] is based on the logarithm of the ratio P_{sat}/P_{atm} to a base of e or 10:

$$\log_n\left(\frac{p_{sat}}{p_{atm}}\right) = A - \frac{B}{T+C} \tag{73}$$

where p_{atm} is 1 bar, T is temperature (K), n is the base of the logarithm, and A, B and C are equation coefficients. Values of the coefficients for many substances are tabulated in [13]. Coefficients must be consistent with the temperature, T, being an absolute temperature in K.

For temperatures above the Maximum Temperature, the saturation pressure is evaluated using the Maximum Temperature; similarly for temperatures below the Minimum Temperature, the Minimum Temperature is used.

Antoine Equation Properties

Coefficient A	The coefficient A in Eqn. (73).
Coefficient B	The coefficient B in Eqn. (73).
Coefficient C	The coefficient C in Eqn. (73).
Minimum Temperature	Specifies the lower limit of the temperature range over which this equation is valid. The default value, 273.15K, is for water. Tables such as those in [13] give minimum temperature values for various substances.
Maximum Temperature	Specifies the upper limit of the temperature range over which this equation is valid. The default value, 647.3K, is for water. Tables such as those in [13] give maximum temperature values for various substances.
Logarithm Base	Specifies the base n in Eqn. (73). e — Coefficients are calibrated for base e, as given in [13]. 10 — Coefficients are calibrated for base 10.

Using the Wagner Equation

The Wagner equation is available for the Saturation Pressure material property. The coefficient properties for this equation are set in the *Properties* window.



The Wagner equation for saturation pressure [13] is:

$$p_{sat} = p_c \exp\left(\frac{A\tau + B\tau^{1.5} + C\tau^3 + D\tau^6}{T_r}\right) \tag{74}$$

in which p_c is the Critical Pressure material property, T_r the reduced temperature T/T_c, T_c the Critical Temperature material property, $\tau = 1 - T_r$ and A, B, C and D are equation coefficients. Values of the coefficients for

many substances are tabulated in [13].

For temperatures above the Maximum Temperature, the saturation pressure is evaluated using the Maximum Temperature; similarly for temperatures below the Minimum Temperature, the Minimum Temperature is used.

Wagner Equation Properties

Coefficient A	The coefficient A in Eqn. (74).
Coefficient B	The coefficient B in Eqn. (74).
Coefficient C	The coefficient C in Eqn. (74).
Coefficient D	The coefficient D in Eqn. (74).
Minimum Temperature	Specifies the lower limit of the temperature range over which this equation is valid. The default value, 275K, is for water. Tables such as those in [13] give minimum temperature values for various substances.
Maximum Temperature	Specifies the upper limit of the temperature range over which this equation is valid. The default value, 647.3K, is for water. Tables such as those in [13] give maximum temperature values for various substances.

Using the Clausius-Clapeyron Method

The Clausius-Clapeyron method is available for the Latent Heat of Vaporization if the Saturation Pressure is given by the Antoine Equation, Wagner Equation, or Polynomial in T methods.



The Clausius-Clapeyron method uses the Clapeyron equation and the assumption of an ideal gas equation of state for the vapor to relate the latent heat of vaporization to the saturation pressure:

$$L = \frac{RT^2}{p_{sat}} \cdot \frac{\mathrm{d}p_{sat}}{\mathrm{d}T} \tag{75}$$

It is the need for $\mathrm{d}p_{sat}/\mathrm{d}T$ which limits the choice of compatible methods for p_{sat}. The ideal gas constant implies that the Molecular Weight material property must be supplied.

Clausius-Clapeyron Properties

None

Using the Heat of Formation

The Heat of Formation is the heat that is evolved when 1 kilogram of the substance is formed from its elements in their respective standard states [J/kg]. This value node (Heat of Formation in the following screenshot) is activated if the Specific Heat property of a mixture component in an EBU or CFM combustion model is to be calculated by a method other than Thermodynamic Polynomial Data. The supplied quantity is always a constant value.



Using the Leidenfrost Temperature

The Leidenfrost temperature is the temperature at which the Leidenfrost effect appears. This effect creates a vapor barrier between liquid and solid phases, reducing heat transfer and adhesion.

Choosing Bai-Gosman Wall Impingement makes the Leidenfrost temperature available in nodes for material properties of liquids under Lagrangian Multiphase > Lagrangian Phases > Phase 1 > Models. You can set the Leidenfrost temperature by selecting one of five methods:

- Constant: Set the temperature to a constant value, or use the default value, 546.03 K, for water at 1 atm.
- Field Function: Define a scalar function determining the temperature.

- Spiegler method:

$$T_L = \frac{27}{32} T_c = 0.84375 T_c \tag{76}$$

Use this method when pressure is much lower than the critical pressure P_c.

- Lienhard method:

$$T_L = T_c \left(0.905 + 0.095 \left(\frac{T_{sat}}{T_c} \right)^8 \right) \tag{77}$$

Use this method when pressure is close to the critical pressure P_c.

- Habchi method:

$$T_L = T_{sat} + \Delta T \tag{78}$$

where

$$\Delta T = \left\{ \begin{array}{ll} T_L{}^1 - T_b & \text{if } (P \le 1) \\ \dfrac{(T_L{}^1 - T_b) - A}{T_c - T_b} (T_c - T_{sat}) + A & \text{if } (P > 1) \end{array} \right\}$$

Use this method when pressure P is close to the critical pressure P_c.

In the preceding equations:

- T_L is the Leidenfrost temperature.
- T_c is the critical temperature. For multi-component drops, T_c is the highest of the critical temperatures of the components.
- T_{sat} is the saturation temperature (boiling point). The saturation temperature is the temperature at which the vapor pressure of the liquid equals the pressure in the cell. It is calculated internally from the saturation pressure p^*. You can define saturation pressure under the Material Properties > Saturation Pressure node of the Liquid model. See Calculating Saturation Temperature.
- $T_L{}^1$ is the "normal Leidenfrost temperature," the value at 1 atm. This is a model parameter, to be determined experimentally or estimated. For multi-component drops, choose $T_L{}^1$ as a function of composition.
- T_b is the normal boiling temperature at 1 atm. For multi-component drops, choose T_b as a function of composition.
- P is pressure.
- A is the larger of 1 or $T_L{}^1 - T_c$.



Calculating Saturation Temperature

The Speigler, Leinhard and Habchi methods for Leidenfrost temperature require the saturation temperature T_{sat} as part of their solution. When the saturation temperature method is set to Iterative, the saturation temperature for the particle is calculated implicitly from the saturation pressure relations of the material. The temperature is calculated using a hybrid bisection/Newton scheme, which finds the temperature at which the vapor pressure of the droplet material is equal to the absolute pressure in the cell.

For multi-component droplet materials, the partial pressure of each component is mass summed to obtain the total vapor pressure. For the numerical method to be successful, the vapor pressure vs. temperature relation must be monotonically increasing.

This temperature is available under the Material Properties > Saturation Temperature node of the Liquid model. Once calculated, it becomes available as a field function. It does not override the Lower Transition Temperature specified in the Bai-Gosman Wall Impingement model.

Using the Standard State Temperature

The Standard State Temperature is the temperature at which the standard state of a substance is defined [*K*]. This value node (Standard State Temperature in the following screenshot) is activated if the Specific Heat property of a mixture component in an EBU or CFM combustion model is to be calculated by a method other than Thermodynamic Polynomial Data. The supplied quantity is always a constant value.



Using the Binary Diffusion Coefficient Method for Molecular Diffusivity

This method uses specified binary diffusion coefficients between the components of the mixture to calculate a molecular diffusivity for each component that is diffusing into the mixture as a whole. For a mixture containing N components, this method requires the specification of $(N(N-1))/2$ diffusion coefficients.

This specification enforces the fact that self-diffusion is zero and the binary diffusion coefficients are symmetric. An array profile specifies these binary diffusion coefficients with all the standard methods available for specifying the individual profiles. The diffusivity of each compound is calculated as:

$$D_m = \frac{1 - x_m}{\sum_{i \neq m} \frac{x_i}{D_{i,m}}} \tag{79}$$

To expose this method, select the Material Properties > Molecular Diffusivity node for a multi-component gas or liquid, then set Method to Binary Diffusion Coefficients. The Molecular Diffusivity > Binary Diffusion Coefficients node appears.



Using the Kinetic Theory Method for Molecular Diffusivity

This method calculates the molecular diffusivity of an individual component using kinetic theory. When this method is used, the following properties show up on the mixture component level:

- Dipole momentum—a measure of polarity of a covalent bond in the molecule (Debye)
- Lennard-Jones Potential—consists of two parameters:
  - Collision diameter (Angstroms)

- Potential energy of attraction
- Polarization—polarizability of a molecule (Cubic Angstroms)

These properties are entered as Constant values only. These values are used for calculating the molecular diffusivity of the individual component for the Kinetic Theory method.

To expose this method, select the Material Properties > Molecular Diffusivity node for a multi-component gas or liquid, then set Method to Kinetic Theory. The Molecular Diffusivity > Kinetic Theory node appears.



The Molecular Diffusivity property on a mixture level is different than other properties in that it actually calculates the molecular diffusivities for each individual component.

The molecular diffusivity D_m of component m is calculated as:

$$D_m = \frac{1 - x_m}{\sum_{i=1,\, i \neq m}^{N} \frac{x_i}{D_{i,m}}} \tag{80}$$

where:

- M_i is the molecular weight of component i
- M_w is the molecular weight of the mixture
- X_i is the molar fraction
- $D_{i,m}$ is the binary diffusivity of components i and m

From gas kinetic theory, the expression of $D_{i,m}$ is based on Chapman-Enskog [12]:

$$D_{i,m} = \frac{2.66 \times 10^{-7} T^{3/2}}{p M_{i,m}^{1/2} \sigma_{i,m}^2 \Omega(T^*)} \tag{81}$$

where:

$$M_{i,m} = \frac{2 M_i M_m}{M_i + M_m} \tag{82}$$

and:

- $\sigma_{i,m}$ is the characteristic Lennard-Jones collision diameter for the pair of component i and component m. $\sigma_{i,m}$ is a function of the Lennard-Jones collision diameter material property, σ_i, that is specified for each component i of the mixture. $\sigma_{i,m}$ is different for polar and non-polar molecules (gases).
- p is the (absolute) static pressure
- $\Omega(T^*)$ is a collision integral that is a function of the reduced temperature, T^*. T^* is defined similarly to the reduced temperature for the viscosity calculation:

$$T^* = \frac{kT}{\varepsilon_{i,m}} \tag{83}$$

where:

- k is the Boltzmann constant = $1.3806503 \times 10^{-23}$ $m^2\ kg\ s^{-2}\ K^{-1}$
- T is temperature

- $\varepsilon_{i,m}$ is the characteristic Lennard-Jones energy for the pair of component i and component m. $\varepsilon_{i,m}$ is a function of the Lennard-Jones potential energy material property, ε_i, that is specified for each component i of the mixture. $\varepsilon_{i,m}$ is different for polar and non-polar molecules (gases).

Using the Component Molecular Diffusivities Method for Molecular Diffusivity

This method allows you to specify the molecular diffusivity of separate components in a multi-component mixture, for cases where the other methods are not sufficient to predict it accurately. Some examples of this situation are:

- Simulations with diffusion dominated flows (that is, where convection is negligible or where there is no turbulence); and,
- Simulations containing mixture components that have significantly different molecular weights (such as H2 with fast diffusion and C10H22 with slow diffusion).

To expose this method, select the Material Properties > Molecular Diffusivity node for a multi-component gas or liquid. Then, set Method to Component Molecular Diffusivity. The Molecular Diffusivity > Component Molecular Diffusivities node appears



When this option is set, you can specify the diffusivities for each component directly using an array profile.

All the usual methods for a scalar profile quantity are available.

See also:

Setting Values Using a Constant Array

Setting Values Using a Composite Array

Using the Nernst-Einstein method for Charged Species Mobility

The default Method for the Charged Species Mobility material property is set to Nernst-Einstein. This method defines the charged species mobility according to Eqn. (2449) to determine the migration of ionic species in electrolytic solutions. You can also define the Charged Species Mobility Method using Constant or Field Function values. See Using Constant Values, or, Using Field Functions.



Note: If the physics continuum contains an energy model, the value of the local temperature is used for T in Eqn. (2449). If no energy model is active, a default value of 293.15 K = 20°C is used.

Using the Electrochemical Species method for Electrical Conductivity

When using the Electrodynamic Potential model, to obtain a consistent setup for applications which require electroneutrality, specify the Electrical Conductivity Method as Electrochemical Species. This option sets the

electrical conductivity according to Eqn. (2448).



See also:

Setting the Electrical Conductivity.

Materials Bibliography

[5] Carreau, P.J. 1972. "Rheological equations from molecular network theories", *J. Rheol.*, **16**(1), pp. 99–127.

[6] Cross, M.M. 1965. "Rheology of non-Newtonian fluids: A new flow equation for pseudoplastic systems", *J. Colloid Sci.*, **20**(5), pp. 417–437.

[7] Escudier, M.P., Poole, R.J., Dales, C., Nouar, C., Desaubry, C., Gresham, L. and Pullum, L. 2005. "Observations of asymmetrical flow behaviour in transitional pipe flow of yield-stress and other shear thinning liquids", *J. Non-Newtonian Fluid Mechanics*, **127**, pp. 143–155.

[8] International Association for the Properties of Water and Steam. 2007. *Revised Release on the IAPWS Industrial Formulation 1997 for the Thermodynamic Properties of Water and Steam*, IAPWS Secretariat, Lucerne, Switzerland (available at www.iapws.org).

[9] International Association for the Properties of Water and Steam. 2008. *Release on the IAPWS Formulation 2008 for the Viscosity of Ordinary Water Substance*, IAPWS Secretariat, Berlin, Germany (www.iapws.org).

[10] International Association for the Properties of Water and Steam. 2011. *Release on the IAPWS Formulation 2011 for the Thermal Conductivity of Ordinary Water Substance*, IAPWS Secretariat, Plzen, Czech Republic (www.iapws.org).

[11] Moore, C. E. 1971. *Atomic Energy Levels as Derived From the Analyses of Optical Spectra*.

[12] Poling, B.E., Prausnitz, J.M., O'Connell, J.P., 2001. *The Properties of Gases and Liquids*, 5th edition, McGraw-Hill, New York.

[13] Reid, R.C., Prausnitz, J.M., and Poling, B.E., 1987. *The properties of gases and liquids*, 4th edition, McGraw-Hill, New York.

[14] Rosen, B., ed. 1970. "Spectroscopic Data Relative to Diatomic Molecules", *International Tables of Selected Constants*, 17.

[15] Vincenti, W. G., and Kruger, C. H. 1965. *Introduction to Physical Gas Dynamics.*

[16] Yasuda, K. 1979. *Investigation of the analogies between viscometric and linear viscoelastic properties of polystyrene fluids*. PhD Thesis, Dept. of Chemical Engineering, Massachusetts Institute of Technology.

Using Material Databases

The material database manager is the method by which STAR-CCM+ provides easy access to information on a wide variety of substances for accurate material modeling in the simulation.

The Tools > Material Databases node, which has its own pop-up menu, is the manager object for working with material databases in the simulation tree.



This section explains the following topics:

- What is a material database?
- How is the Material Data Used in the Simulation File?
- What happens if the database file is missing
- Setting up material databases
- Loading material database files
- Modifying a copy of the material database
- Changing the units of the material properties

What Is a Material Database?

The material database node (Standard in the following screenshot) within the Tools > Material Databases node contains the objects that represent the information from the material database.



When you create a simulation, this node appears with its sub-nodes in the object tree as soon as you select a material model.

Its contents are grouped among three sub-nodes that represent database folders: Gases, Liquids, and Solids. You can also create database folders as well as create database material properties.

The source of the material database is a `props.mdb` file, that is located in your main STAR-CCM+ installation directory (Windows) or in the `star` subdirectory of that directory (Linux). It is possible to load one or more alternate databases. You can also update the database through the pop-up menu.

A second material database is supplied called `propsExt.mdb`, which contains a larger selection of materials that are rarely used.

If you set up a database, the data you enter in the STAR-CCM+ GUI is not saved until you explicitly export it to an external file.

What Is a Database Folder?

A database folder node is a manager containing nodes of database materials.



In the default Standard material database, these folders are used to group database materials as gases, liquids, and solids.

Database Folder Properties

None.

Database Folder Pop-Up Menu

New Material	Creates a database material from a submenu: New Gas, New Liquid, or New Solid.
New Folder	Creates a folder within the database folder. Folders can be useful for organizing many materials.

See common menu items.

What Is a Database Material?

The database material node represents the substance in the material database. It is a manager node for database material properties nodes.



The nodes of these materials have sub-nodes representing database material properties, such as dynamic viscosity and molecular weight. The properties nodes in turn contain objects that represent methods for calculating property values.

Database Material Properties

Symbol	The chemical name of the database material where applicable.
Title	The common name of the database material.

Database Material Pop-Up Menu

New Database Properties...	Activates a dialog that selects database properties from a list.

See common menu items.

What Is a Database Material Property?

The database material property node represents a particular property of the substance in the material database.



Database Material Property Node Properties

Property Tag	A read-only tag denoting the material property.

See common menu items.

What Is a Database Material Property Method?

Each property has one of more of the following property methods that are associated with it:

- Constant
- Polynomial
- Thermodynamic Polynomial Data
- Table(T)

The property Specific Heat also has a Thermodynamic Polynomial Data method available.

Constant

Most properties have a *Constant* method available, which defines the constant value of the property.



To change this property, select the Constant sub-node and make an entry in the Value property.

The value node inside the database material property node represents the value that is set for that property.

Constant Properties

Value	The constant value.

See common menu items.

Polynomial

Polynomial in T is a method for calculating certain quantities as polynomial functions of temperature. The Polynomial node represents polynomial coefficients for density, saturation pressure, and specific heat. This particular type of polynomial method is not available in the supplied material database by default; you must create it. See Creating a Database

Material Property.



Polynomial Properties

Polynomial	Values for the polynomial.

See common menu items.

Thermodynamic Polynomial Data

Thermodynamic Polynomial Data is a method for calculating specific heat with a piecewise, fourth-order (five-coefficient) polynomial in temperature. Two additional coefficients are provided in each temperature range to allow the calculation of enthalpy and entropy.



Thermodynamic Polynomial Data Properties

Polynomial Data	Values for the polynomial.

See common menu items.

Table(T)

A material property that is imported from an XML file uses the Table(T) property method. The data for the material property is contained in a property table that is represented by the Table(T) > Material Database Table node.



To replace the contents of a material database table:

- Right-click the Material Database Table node and select Import Material Table from a file...
- Open the `.csv` file containing the new data.

Table(T) Properties

None

Material Database Table Properties

Interpolated value	Property value that is interpolated from the database table. (Read-only)
Extracted	Properties that are contained in the database table. (Read-only)

Material Database Table Pop-Up Menu

Export...	Export the material table to a `.csv` file.
Import Material Table from a file...	Replace the material data with the contents of a `.csv` file.
Tabulate...	Display the data in tabular form.

Creating a Database Folder

If you are setting up a material database object, part of the process involves creating a folder. Right-click the material database node and select New Folder.



Creating a Database Material

When setting up a material database object, after you create database folders, the next steps involve creating database materials. Right-click on the database folder node and select a new gas, liquid, or solid from the New Material pop-up menu item.



Creating a Database Material Property

To create a database material property:

- Right-click on the database material node and select New Database Properties...

- Select database properties with corresponding methods from the list in the *Select Database Properties* dialog that appears.



The constant values for these new database properties are zero by default.

The Select All button lets you manage the selection of many items efficiently. Simply click it to select the entire group, and then deactivate the individual checkboxes of the ones you wish to exclude. If you want to start over, deactivate all the checkboxes at once by clicking the Deselect All button.

- Click OK when finished.

When a material property option is selected in the dialog, the option disappears from the dialog and a corresponding node appears in the object tree. In the example below, the option Saturation Pressure - Polynomial was selected. The resulting additions to the object tree are nodes representing the Saturation Pressure material property with the Polynomial in T method.



Changing Database Material Properties

Generally material properties can be changed by opening the database material property nodes and entering values in the properties of the sub-nodes. However, your actions in the GUI affect how permanent these changes are.

In principle, changes are made permanently to the material database by modifying a copy of that database. (Take care not to overwrite the original supplied file.) The simulation file itself only retains changes that are made to properties of the database materials that are used to create substances *in the continuum* (within single, multiphase, or multi-component material models). Therefore, if you modify a group of database materials, and then

select only part of them in your continuum, the other modified database materials revert to their original values when you close and reopen the simulation file.

A good approach for modifying properties is as follows:

- For a trial run, change material properties in the continuum.
- Run the simulation.
- If you are satisfied with the results from the new property values of the materials in the continuum:
  - Make a copy of the external database file for use by the simulation.
  - In the database material, enter the custom values that you used in the continuum. Specifically, enter the values in the database material properties values of the database material (for example Tools > Material Databases > Standard > Gases > Air > Density > Constant).
  - Save the values to the database file by right-clicking the material database node and selecting Save.

Material Database Properties

File name	The name of the database file from which the material database object gets its material data.
File path	The path to the database file.

Material Database Pop-Up Menu

New Folder	Creates a database folder.
Import Properties...	Imports material data from a file (`.dbs` format only).
Reload	Reloads the material database.
Save	Saves the material database changes to the external file. Do not use this command with the Standard material database object before you archive a copy of the supplied `props.mdb` file.

See also: common menu items.

How Is the Material Data Used in the Simulation File?

It would be too costly in terms of disk space to duplicate all material data from the `props.mdb` file into every simulation file. Therefore a simulation file only retains the information on materials that are referenced by material model selection. This information includes:

- Names of referenced materials
- The database folder (gas, liquid, or solid) to which each referenced material belongs
- Properties of referenced materials
- Values of the properties

To make data of other materials available, STAR-CCM+ accesses the `props.mdb` file directly. This action occurs when, for example, you want to replace a selected material.



Available replacement materials are listed in the *Replace Material* dialog.



This list is why the `props.mdb` file must be in its proper place when you start a new simulation, or when you open an existing simulation and intend to replace materials. The proper place is your main STAR-CCM+ installation directory (Windows) or the `star` subdirectory of that directory (Linux).

Note: When the materials database does not contain specific properties for a material, STAR-CCM+ inserts default values instead. The default values are those for air (gases), water (liquids), and aluminum (solids).

Additionally, if multiple databases were used in creating a simulation, all of those databases must be in the places they occupied when STAR-CCM+ loaded them into the simulation. Otherwise any unique data that they contain is not accessible within the simulation unless you load the databases again.

What If the Database File Is Missing?

Most of the work you would do with material models depends on the supplied `props.mdb` file residing in its normal place. That place is your main STAR-CCM+ installation directory (Windows) or the `star` subdirectory of that directory (Linux). You can load a material database into an open simulation.

Without a database, you cannot select materials in a new simulation, or change the material selections in an existing simulation. Your only option is to modify properties of material selections that were saved previously in a simulation file.

In a new or existing simulation, STAR-CCM+ lets you know in several ways that the database is missing (in the following example, the standard material database is missing):

- When you attempt to select a new or alternate material model, you get an empty selection dialog, except for any materials selected in a previously saved simulation. (In this example, the simulation file contains only the material Air.)



- After you attempt to select a material model, the material database node appears in the object tree, but with no database folder nodes. (In an existing simulation file, the database folder nodes are also empty except for the referenced materials.)



Setting Up Material Databases

You can add material database objects to the STAR-CCM+ simulation manually. This capability can be helpful for importing legacy (`.dbs`) data.

To create a material database object, right-click on the Material Databases manager node and select New Material Database...



A new, empty node appears with the default name matl.



Subsequent objects have default names that are concatenated with a number, and can be renamed.

To populate this database from within the GUI, create folders. Right-click on the database node and select New Folder. You can then rename the folders as required to organize your database materials.

Once you have prepared the folders, right-click each one to select the database material to create. The database material objects have default names that are based on the type that is selected from the pop-up submenu, concatenated with numbers. Rename these objects according to the database materials you want them to represent. Make entries in both the Symbol and Title properties, though your choice of symbol entry is the more important.

Next, open the database material nodes, select the database material property nodes and enter values for the properties. If you are setting up multiple database materials that have similar properties, you can use the copy-and paste technique as follows:

- Create one of the database materials and set up all its property values.
- Right-click on the node of the database material and select Copy.
- Create another node for the similar database material.
- Right-click that newer node and select Paste. It has all the properties and subobjects of the previous database material.
- Adjust the properties of that newer database material as needed.

Using the File path property of the database object, you can select the file path of your new database file from among the following locations:

- The root directory
- The STAR-CCM+ installation directory
- The same location as the current simulation file



However, the database you create is not saved until you explicitly export it to an external file. STAR-CCM+ reminds you of this fact in the tooltip and *Properties* window description area for every database object, whether the database it represents is saved or not.



To export your new database into an external file, simply right-click on the node and select Save.



The file is saved in `.mdb` format with a default name matching the name of the corresponding node.

Loading Material Database Files

By default, a material database object named Standard is loaded automatically in a new simulation once you select a material model. This object accesses the `props.mdb` file that is located in your main STAR-CCM+ installation directory (Windows) or in the `star` subdirectory of that directory (Linux). To add a material or change the property of one, you would load a separate material database file. Alternatively, you can modify a database file for this purpose and reload it into the simulation. (If you are working with the database file supplied with STAR-CCM+, be sure to make a copy and then move the original to another location.) You can also load legacy material data from `.dbs` files.

Loading a Database File

STAR-CCM+ can accommodate multiple material databases, making them all available in the material model selection process.

The simplest way to prepare to load an additional database file is to place it in the same directory as the ones supplied by default, but with a different name. This location would be your main STAR-CCM+ installation directory (Windows) or the `star` subdirectory of that directory (Linux). Make these preparations before opening your simulation file.

Alternatively, you can set an environment variable. Select Tools > Options and, in the *Options* dialog, select the Environment node. The Material Database Load Path property allows you to enter absolute paths of multiple databases, with each entry separated by a semicolon. By clicking the property customizer (...) you can activate a dialog that provides a complete view of your entries. These paths are *in addition to* the default path. Click OK when finished.

Next, right-click on the Material Databases manager node and select Open Material Database...



A separate *Open Material Database* dialog appears with a list of the `.mdb` files available, including the one(s) you added.



Select one and click OK.

A new node for that database appears in the object tree, and its database file path appears in the *Properties* window.



Reloading a Database File

To reload the material database file, right-click the material database node and select Reload



This action updates all the information in the simulation for that database.

Loading a Legacy Database File

STAR-CCM+ can import the older `.dbs` database format. To import this type of file, start by creating a new (empty) material database object. Then right-click on that object and select Import Properties... The dialog that appears is set to filter the `.dbs` format. Navigate to your file and click Open. Then select an option from the *Import Properties* dialog.



The first option imports the contents of the entire file. The second imports properties only for the materials that are already represented in the database object. This action populates the material database node with three database folders, similar to the structure of the Standard database object.

You can also add more properties to the database object by selecting Import Properties... again, and selecting other files.

Database files that are supplied with STAR-CCM+ include `props.mdb`, `thermdat.dbs` (NASA polynomials), and `moldat.dbs` (transport properties). To combine these files with one such as `props.mdb`, repeat the import action, but select the option Import properties only for materials already in the database in the dialog that appears.

Modifying a Copy of the Material Database

Possible reasons for changing the material database include:

- Adding material data
- Importing a PPDF table (equilibrium or flamelet)
- Importing Chemkin-formatted data

Caution: Before changing the material database, make a copy of the database file for modification, and keep the original in a safe place. An additional helpful measure is to set the original file as read-only.

Located in your main STAR-CCM+ installation directory (Windows) or in the `star` subdirectory of that directory (Linux), the `props.mdb` file is modified through the STAR-CCM+ GUI, rather than opening it with another program.

After you make a copy of the external `props.mdb` file and relocate the original, change properties of database materials and/or add new materials with new properties and enter values for them. Once you are finished, right-click the node of the material database and select Save. (Saving the simulation file alone does not save your database material properties changes to the external database file.)

Changing the Units of the Database Material Properties

By default the database materials have properties that are measured in SI (Systeme Internationale) units. However, you can change the units as they are represented in the database material properties, using the standard technique for changing units in STAR-CCM+.

After making the switch once, you will have the added option of changing the units used for individual database material property values. To access this option for a particular value, click the property customizer (...) in the Value property of the node to activate the Customizer dialog. The properties within that dialog include a drop-down list with a choice of units.



Material Databases Manager Pop-Up Menu

Open Material Database...	Opens an available material database file (`.mdb` format only) from a drop-down list in a separate dialog. This list is updated when the simulation is opened.
New Material Database...	Adds a material database object to the simulation.

Modeling Flow and Energy

STAR-CCM+ offers a wide choice of models for simulating flow and energy when the continuum mechanics formulation of fluid dynamics is valid. The Coupled Flow and Coupled Energy models are best for compressible flows, natural convection problems, and flows with large body forces or energy sources. The Coupled Flow model includes the optional AUSM+ inviscid flux scheme, which is advantageous for flows involving high supersonic and/or hypersonic regimes. The Segregated Flow model is an alternative for incompressible or mildly compressible flows, particularly when computational resources are an issue; it is used with the Segregated Fluid Energy models. For complex time-periodic flows, you can use Harmonic Balance Flow and Energy.

Whichever flow and energy models you choose, modeling options that interact with them include:

- Viscous regime (including inviscid, laminar, and turbulent)
- Equation of state—models that compute the density and the density derivatives with respect to temperature and pressure
- Passive scalar—transport equations for user-defined variables that can be represented using arbitrary scalar values
- Gravity—specifically including the buoyancy source terms in the momentum equations

Additional models accompany the flow and energy ones to create sophisticated simulations for specific applications:

- Thin film —deicing or defogging on a surface such as a cabin windshield
- Porous media—commonly used to show the macroscopic effect of the porous medium on the overall fluid flow

- Non-reflecting boundaries and mixing-plane interfaces—allowing accurate simulations of applications such as multi-stage turbomachinery
- Vorticity confinement—preventing the rapid dissipation of vortices
- Fans—specifically the fan momentum source

In addition to descriptions of the models that are mentioned above, the following sections help clarify how this modeling works:

- Explanation of how pressure is simulated in STAR-CCM+
- Suggestions for initializing the flow
- Formulation of the boundary conditions

When you activate these models, other objects appear in the GUI that are designed to give you added control over the flow modeling. Some of them are common to more than one flow or energy model, or common to flow and energy models. They are listed in the following sections:

- Flow reference
- Energy reference
- Values common to flow and energy models
- Field functions common to flow and energy models

General Navier-Stokes

All of fluid dynamics is based on three physical principles:

- Mass is conserved.
- Newton's second law, $\mathbf{F} = m\mathbf{a}$
- Energy is conserved.

These physical principles are developed into the fundamental governing equations of fluid dynamics: the continuity, momentum, and energy equations. These equations are also known as transport equations or conservation equations.

The form of governing equations that is generally used in Computational Fluid Dynamics (CFD) is known as the Navier-Stokes equations. The Navier-Stokes equations include the effect of viscosity on the fluid flow.

In STAR-CCM+, the finite volume method is used to transform the continuous governing equations into a form that can be solved numerically. Two approaches are taken to solve the governing equations:

- The segregated approach, where the flow equations are solved one after the other and linked using a correction equation.
- The coupled approach, where the coupled system of equations is solved simultaneously.

Each approach has its own merits, which must be considered before proceeding with a flow analysis in STAR-CCM+.

When Is the STAR-CCM+ Continuum Mechanics Formulation Valid?

The Knudsen number Kn determines whether the continuum mechanics formulation of fluid dynamics is valid. The Knudsen number is defined as:

$$Kn = \frac{I}{L} \tag{84}$$

where I is the molecular mean free path and L is the characteristic length of the body. For high Reynolds number, the Knudsen number is:

$$Kn = \frac{M}{Re}\sqrt{\frac{\gamma\pi}{2}} \tag{85}$$

where M is the Mach number, γ is the ratio of specific heats, and the Reynolds number Re is dependent on L.

Flow problems can be divided into four general regimes that are based on the Knudsen number as shown below. These limits are not precise, but rather rules of thumb.

$Kn < 0.01$	Continuum
$0.01 < Kn < 0.1$	Slip flow
$0.1 < Kn < 3$	Transitional
$Kn > 3$	Free molecular flow

STAR-CCM+ can be used in the continuum regime where the Navier-Stokes equations are valid.

In the slip flow regime, the Navier-Stokes equations can produce good results. However, the effects of rarefaction must be taken into account through slip velocity boundary conditions and accounting for temperature jump on the surface.

Solving for transitional flows requires a statistical method, such as a direct simulation Monte Carlo approach. In the free molecular regime, the flow is sufficiently rarefied that molecular collisions can be neglected. The continuum assumption that is built into the Navier-Stokes equations is not applicable to these rarefied gas regimes.

High-Speed Flow Regimes

At high Mach numbers, extreme temperatures occur and dissociation, or even ionization, of the gas must be considered. The Mach number regimes, and the appropriate models to use, vary by gas and upstream temperature and pressure. For air however, the following general guidelines can be used:

Mach Number	Flow Regime	Modeling
$M < 3$	Calorically Perfect	Ideal gas Constant specific heat
$3 < M < 8$	Thermally Perfect	Ideal gas Specific heat using polynomial in T
$8 < M < 30$	Dissociation	Real gas model and/or chemically reacting flow
$M > 30$	Ionization	Requires solving the Navier-Stokes equations together with the Maxwell equations. Not inherently supported in STAR-CCM+

Pressure

STAR-CCM+ simulates various pressure quantities, depending on the model selection:

- Reference Pressure—a device that reduces the numerical roundoff error in the numerical calculations involving pressure.
- Working Pressure—the solution variable that is used in the flow models in STAR-CCM+. In the code and in the documentation this variable is mentioned most of the time as Pressure. Its definition is a matter of computational convenience and depends on the model selection. It is always expressed relative to the Reference Pressure.

  In laminar flows without the influence of gravity, Pressure is equal to the static gauge pressure. If the gravity model is active, Pressure is equal to the piezometric pressure, as shown in equations (257) and (258) for variable and constant density flows, respectively.

  When a turbulence model is used, Pressure includes the contribution of $\frac{2}{3}\rho k$. This contribution can be considered negligible in most flows of engineering interest.
- Static Pressure—the spherical part of the stress tensor acting in fluids (see Eqn. (166)). It happens that the static pressure is the same as the thermodynamic pressure, which is one of the state variables.
- Absolute Pressure—the sum of the static pressure and the reference pressure.
- Relative Total Pressure—the static pressure that results from isentropically bringing the flow to rest in the relative frame of motion.
- Total Pressure—the static pressure that results from isentropically bringing the flow to rest.
- Absolute Total Pressure—the sum of the total pressure and the reference pressure.

Defining the Reference Pressure

The reference pressure is simply a device that is used to reduce the numerical roundoff error in the numerical calculations involving pressure. This reduction is necessary since the differences in pressure are important, and these differences can be small relative to the absolute value of the pressure. For example, in fully incompressible, or very low-Mach number flows. By subtracting a suitable constant reference pressure, a working pressure is obtained that is less prone to roundoff errors.

For simulations that use the Constant Density model or the Polynomial Density model, the actual value of the reference pressure has no relevance to the calculations. However, when using the Ideal Gas model, the reference

pressure is used as in Eqns. (141) and (142). In all cases, the reference pressure is used in the definition of the Absolute Pressure field function according to Eqn. (142).

For simulations using the Ideal Gas model the following guidelines are suggested for setting the reference pressure:

- If the Mach numbers are low (say 0.3 or less), set the reference pressure to some ambient value such that the working pressures are small.
- If the Mach numbers are substantial, you can do one of the following:
  - Set the reference pressure to an ambient value such that the working pressures are small.
  - Set the reference pressure to zero so that working pressures are identical to the absolute pressures.

Specify the pressure location manually when conditions require it.

Reference Pressure Properties

Value	The value of the reference pressure. The customizer triggers the define value action.

Providing a User Specified Pressure Reference Location

It is possible to set the pressure reference location manually when it becomes necessary to do so (as discussed in section, When is the Reference Pressure Location Required?). To set the pressure reference location:

- Select the appropriate solver node and choose User Specified for the Pressure Reference Location property in the *Properties* window.



The Pressure Reference Points manager node, which has its own pop-up menu, is added as a child.

- Right-click this node and choose New. A Pressure Reference Point node is added.



- Select the newly added Pressure Reference Point node and set the point coordinates (along with the reference coordinate system) in the *Properties* window.



When the Pressure Reference Point node is selected, the point is displayed in an active scene.



What Is a Pressure Reference Point?

The properties of a Pressure Reference Point node are used to define a single geometric location that STAR-CCM+ uses in finding the pressure reference cell for a fluid stream. For each fluid stream, the cell closest to one of the defined pressure reference points is used as the pressure reference cell. Multiple pressure reference points are only necessary for problems having multiple disconnected fluid streams; you set one pressure reference point in each fluid stream.

Properties

Point Coordinates	The geometric location of the pressure reference point in the selected coordinate system.
Reference System	A list of local coordinate systems in the simulation for use in defining the point coordinates.

Expert Properties

Enabled	Determines whether the solver users this pressure reference point in finding the reference cell.	
	Activated	The point is included.
	Deactivated	The point is excluded.

Pressure Reference Points Manager Pop-Up Menu

New	Adds a Pressure Reference Point node to the manager.

See also: common menu items.

When Is the Pressure Reference Location Required?

The Pressure Reference Location is an expert property on the Pressure solver for the Segregated Flow model, and on the Coupled Implicit and Coupled Explicit solvers for the Coupled Flow model. It is used to determine whether STAR-CCM+ automatically calculates a location for the reference pressure, or whether you provide the location manually.



When simulating steady flow, or the flow of incompressible fluids, the equation for pressure is singular. If the boundary conditions do not define the pressure level (that is, there are no pressure boundaries), you can anchor the pressure level to a datum value in the solution domain. If the working pressure is not anchored to a known value somewhere in the domain, it can float arbitrarily high and cause numerical problems.

- If absolute pressure levels are not important, set the Pressure Reference Location property to Automatic Selection.

STAR-CCM+ identifies the cell such that the cell centroid has the `min(x), min(y), min(z)` in the domain. This condition can only occur in cells that are next to boundaries.

To determine the appropriate cell, STAR-CCM+ does the following:

a. Checks the centroid of every cell that is next to a boundary, and selects all of the cells that have the smallest X component.

   STAR-CCM+ checks all of the boundaries: wall boundaries as well as flow boundaries.

b. From the cells that are selected in the previous step, STAR-CCM+ selects the cells that have the smallest Y component.

c. From the cells that are selected in the previous step, STAR-CCM+ selects the cell that has the smallest Z component.

STAR-CCM+ uses the selected cell as the pressure reference location.

Note: There is known limitation with this method. Consider a rectangular prism as a flow domain: in the ideal case STAR-CCM+ would select the cell at the lower left front edge of the prism. However, small numerical rounding errors in the centroid calculation and meshing can cause STAR-CCM+ to select a different cell nearer the middle of the left side of the prism.

- If you want to determine absolute pressure levels accurately, set the pressure reference location manually as described in the section, Providing a User Specified Pressure Reference Location. For example, when you want to evaluate the pressure force on wall boundaries. Choose a location where the pressure is known, and where a fixed value is appropriate for the simulation.

The working pressure at the pressure reference location is set using the initial pressure value; the absolute pressure is the sum of the reference pressure plus the initial pressure. If the initial pressure value is changed during a simulation (by using a field function, for example), the pressure at the pressure reference location changes accordingly.

In buoyant flows, the definition of an operating altitude is required. The operating altitude is a position vector defining the location at which the static and piezometric pressures are equal. For more information, see the sections on modeling gravity for variable density and constant density flows.

Modeling the Equation of State

The Equation of State model is used to compute the density and the density derivatives with respect to temperature and pressure. This section describes the various models that are available in STAR-CCM+:

- Constant Density
- IAPWS-IF97
- Polynomial Density
- Ideal Gas
- Ideal Gas (with Combustion)
- Real Gas
- Thermal Non-Equilibrium
- User Defined Equation of State (EOS)

It also includes the following:

- How to select an Equation of State model
- Field function reference

What Is the Constant Density Model?

The Constant Density model makes the assumption that the density is invariant throughout the continuum. It is available for gases, liquids, and solids.

Constant Density Model Properties

None

What Are the IAPWS-IF97 Models?

The models of IAPWS-IF97 (International Association for the Properties of Water and Steam, Industrial Formulation 1997) allow you to run simulations of liquid water or gaseous steam. These models include calculations of density and other thermodynamic properties. The IAPWS-IF97 provides fundamental polynomial equations for the specific Gibbs free energy, $g(p, T)$. Specific volume, internal energy, entropy, enthalpy, heat capacity, and speed of sound are derived from the fundamental equation using appropriate combinations of the dimensionless Gibbs free energy and its derivatives [33].

Validity Regions of IAPWS-IF97

Based on the Gibbs free energy, the IAPWS-IF97 models in STAR-CCM+ are only valid within certain ranges of temperature and pressure. These ranges are represented in a diagram of regions [33] which are essentially the phases of water and steam.



The following table lists the components of the region diagram and how they work with IAPWS-IF97 models in STAR-CCM+.

Component	Ranges of Validity		STAR-CCM+ Models
Region 1 (liquid)	273.15 K ≤ T ≤ 623.15 K	p ≤ 100 MPa	IAPWS-IF97 (Water)
Region 2 (steam)	273.15 K ≤ T ≤ 1073.15 K	p ≤ 100 MPa	IAPWS-IF97 (Steam)
Region 3 (saturation)	This phase region is not supported in STAR-CCM+.		
Boundary 4	This boundary lies between liquid and saturation on one side, and steam on the other.		
Region 5 (steam)	1073.15 K < T ≤ 2273.15 K	p ≤ 50 MPa	IAPWS-IF97 (Steam)

Selecting an IAPWS-IF97 Model

Activate one of the two models depending on whether you are simulating liquid water or gaseous steam:

- For liquid water, select the IAPWS-IF97 (Liquid) model from the *Equation of State* group. It is available once the Liquid material model is activated.
- For gaseous steam, select the IAPWS-IF97 (Steam) model from the *Real Gas Equation of State* group. It is available once the Gas material model and Real Gas model are selected.

When either IAPWS-IF97 model is activated, the following material properties are evaluated using an appropriate IAPWS-IF97 method:

- Critical Pressure
- Critical Temperature
- Dynamic Viscosity
- Molecular Weight
- Saturation Pressure

- Saturation Temperature
- Speed of Sound
- Thermal Conductivity

The IAPWS-IF97 method is the only one available for these properties.

Also when this model is activated, the following fundamental quantities are calculated using IAPWS-IF97:

- Entropy
- Specific Heat

What Is the Polynomial Density Model?

The Polynomial Density model makes the assumption that the density is a function of temperature only, and is available for gases, liquids, and solids.

Polynomial Density Model Properties

None

What Is the Ideal Gas Model?

The Ideal Gas model, available for gases only, uses the ideal gas law to express density as a function of temperature and pressure.

Ideal Gas Model Properties

Incompressible — Controls whether the ideal gas is compressible (dependent on working pressure, p).

- Activated — Removes dependency on working pressure, p.
- Deactivated — Maintains dependency on working pressure, p.

Ideal Gas Model Expert Properties

Density Limiting — Activates or deactivates the Density Limits option that is used to prevent negative densities.

- Activated — Activates Density Limits. A node for this option appears on the simulation tree under the Ideal Gas node.
- Deactivated — Deactivates Density Limits. The default.

Density Limits — Lets you specify the density-limiting parameters.

- Cut-off pressure — The minimum value of absolute pressure that is used to evaluate the density.
- Pressure interval — The interval, centered on the cut-off pressure, over which the blending from minimum to maximum density occurs.

The following figures show the node for the Ideal Gas model, with the Density Limiting expert property activated.



The following figures show the Density Limits node, with the Cut-off pressure and Pressure interval properties exposed.



The following figure shows a comparison of the limited and unlimited density as a function of absolute pressure. As seen from the figure, the limited density smoothly and continuously transitions from its unlimited value to a constant, minimum value, rho_min, over the course of the interval,

P_interval, as specified by the Pressure Interval. The value rho_min is the density that is calculated at the existing temperature. An absolute pressure value as specified by the Cut-off Pressure, P_co.



What Is the Ideal Gas (with Combustion) Model?

The Ideal Gas (with combustion) model allows only the thermodynamic polynomial for calculation of specific heat of individual components of the mixture. By contrast, the standard Ideal Gas model also allows the constant and polynomial in T methods.

The thermodynamic polynomial is the recommended method for getting specific heat (together with the heat of formation and enthalpy) of the mixture. However, in some (rare) instances, this method is not available. In that case you can select the standard Ideal Gas model and pick an alternate method for specific heat calculations.

What Is the Real Gas Model?

Real Gas models are used in situations where the gas properties cannot be explained using the ideal gas law. For example, near the condensation point of gases, near critical points, and at high pressures. Real Gas models take into account non-ideal behavior such as compressibility effects, variable specific heat capacity, van der Waals forces, and non-equilibrium thermodynamic effects.

The following Real Gas models are available in STAR-CCM+:

- Peng-Robinson (available for gases only)
- Redlich-Kwong (available for gases only)
- Soave-Redlich-Kwong (available for gases only)
- Modified Soave-Redlich-Kwong (available for gases only)
- Van der Waals (available for gases only)
- Equilibrium Air (available for air only)
- IAPWS-IF97 (available for steam or liquid water)

Peng-Robinson Model

The Peng-Robinson model registers three more material properties:

- Critical Pressure, p_c
- Critical Temperature, T_c
- Acentric Factor, ω

These properties are constant. Default values for these constants are provided from the material database, though you can change them if desired.

Redlich-Kwong Model

The Redlich-Kwong model registers two more material properties:

- Critical Pressure, p_c
- Critical Temperature, T_c

These properties are constant. Default values for these constants are provided from the material database, though you can change them if desired.

Soave-Redlich-Kwong Model

The model registers three material properties:

- Critical Pressure, p_c
- Temperature, T_c
- Acentric Factor, ω

Modified Soave-Redlich-Kwong Model

Graboski and Daubert [27], [28] modified the Soave-Redlich-Kwong model to improve its behavior for phase equilibrium calculations that involve non-polar molecules. It uses the same material properties as the original Soave-Redlich-Kwong model.

Van der Waals Model

The van der Waals model registers two more material properties:

- Critical Pressure, p_c
- Critical Temperature, T_c

These properties are constant. Default values for these constants are provided from the material database, though you can change them if desired.

Equilibrium Air Model

The equilibrium air equation of state is used for the simulation of high-temperature air, where the effects of molecular dissociation, internal energy excitation, and ionization become significant. For situations where the time scale of these effects is much shorter than the time scale of the flow, these effects can be treated in an equilibrium manner (see Equilibrium Air Formulation). With the equilibrium assumption, this model can capture the effects of dissociation and ionization in air without the computational expense of a multicomponent, reacting simulation.

The curve fits of this model are valid for temperatures less than 30,000 K and for pressures between 10^{-4} and 10^{2} atmospheres. Since the curve fits are based on the behavior of air, the Equilibrium Air equation of state can only be used with air as the working fluid.

This equation of state is useful for the simulation of hypersonic flows, such as nozzle flows or atmospheric re-entry. When Equilibrium Air modeling is activated, the following material properties are evaluated:

- Specific Heat
- Dynamic Viscosity
- Thermal Conductivity
- Speed of Sound

Equilibrium Air Model Expert Properties

Density Limiting — Activates or deactivates the Density Limits option that is used to prevent negative densities.

- Activated — Activates Density Limits. A node for this option appears on the simulation tree under the Equilibrium Air node.
- Deactivated — Deactivates Density Limits. The default.

Thermal Non-Equilibrium

Non-equilibrium occurs at high temperatures and low densities, where the vibrational/electronic energy modes become active yet the density is low enough such that equilibration does not occur. Examples of problems include high-speed atmospheric re-entry flows and nozzles. For problems in this regime, thermal non-equilibrium is required in order to capture temperature profiles and heating accurately.

The Thermal Non-Equilibrium energy model is currently limited to the Coupled Flow model only. The Thermal Non-Equilibrium Ideal Gas model requires the use of the Coupled Thermal Non-Equilibrium model. It is a required model for the Equation of State.

The principal application for this model is hypersonic flows (mainly atmospheric re-entry). Nozzle flows where the effects of vibrational freezing are important can also be simulated. When Thermal Non-Equilibrium modeling is activated, the following material properties are evaluated:

- Vibrational-Electronic Specific Heat, with the methods Thermodynamic Polynomial or Gas Kinetics
- Mode Relaxation Time, with the methods Landau-Teller or Inverse Mole-Fraction Average
- Vibrational Conductivity, with the method Non-Equilibrium Kinetic Theory

Thermal Non-Equilibrium Ideal Gas Model Properties

Incompressible — Controls whether the ideal gas is compressible (dependent on working pressure, p).

- Activated — Removes dependency on working pressure, p.
- Deactivated — Maintains dependency on working pressure, p.

Thermal Non-Equilibrium Ideal Gas Model Expert Properties

Density Limiting — Activates or deactivates the Density Limits option that is used to prevent negative densities.

- Activated — Activates Density Limits. A node for this option appears on the simulation tree under the Thermal Non-Equilibrium Ideal Gas node.
- Deactivated — Deactivates Density Limits. The default.

Thermal Non-Equilibrium Model Properties

Equalize Temperatures — Forces the Vibrational-Electronic temperature to equal the Translational-Rotational temperature. This option is useful for problems with a low degree of thermal non-equilibrium and can be used to overcome start-up problems that are associated with a poor initial condition. Deactivate it when the problem begins to converge towards a solution.

- Activated — Forces the Vibrational-Electronic temperature to equal the Translational-Rotational temperature.
- Deactivated — Allows the Vibrational-Electronic temperature to remain independent of the Translational-Rotational temperature.

Freeze Vibrational Temperature	Controls whether to freeze the Vibrational temperature at its current value. This option is useful for problems with a large degree of thermal non-equilibrium and can be used to overcome start-up problems that are associated with a poor initial condition. Deactivate it when the problem begins to converge towards a solution. **Activated** — Freezes the Vibrational temperature at its current value. **Deactivated** — Allows the Vibrational temperature to change.

What Is the User Defined EOS Model?

The User Defined EOS model allows you to specify the equation of state for a physics continuum using your own approach. Depending on the various model selections you have made for a continuum, some or all of the following material property nodes are added for each material in the continuum:

- Density, ρ
- Density pressure derivative, $(d(\rho))/(d(p))$
- Density temperature derivative, $(d(\rho))/(d(T))$
- Density mass fraction derivative, $(d(\rho))/(d(Y_i))$
- Enthalpy, h
- Speed of sound, c

The User Defined EOS model provides you with various methods for specifying these thermodynamic properties. These methods include constant, field function, temperature polynomial, and the option to input tabular data from a file.

The model also allows you to indicate whether the fluid is compressible, that is, whether density is a function of pressure, by activating the Compressible property of the User Defined EOS model node.



User Defined EOS Properties

Compressible	Controls whether the material is compressible (dependent on working pressure, p).	
	Activated	Material is dependent on working pressure, p.
	Deactivated	Material is not dependent on working pressure.

When the working fluid is declared compressible, the material property Density Pressure Derivative appears with the material. This property defines the partial derivative of the user-defined density with respect to pressure. The Speed of Sound material property is included for compressible materials as well.

If an Energy model is activated in the physics continuum, the material property Density Temperature Derivative is activated. This property defines the partial derivative of the user-defined density with respect to temperature. A node is also added for Enthalpy.

If a Multi-Component Gas or Multi-Component Liquid model is activated in the physics continuum, the material property Density Mass-Fraction Derivative is activated for each mixture component. This property defines the partial derivative of the user-defined density with respect to mass-fraction of the component.

The User Defined Equation of State model is applicable in instances where the density of the working fluid best represented with some user-defined expression or a table of property data. You use this model when the other Equation of State models (Constant Density, Polynomial Density, or Ideal Gas) do not adequately define the density.

As an example, one can choose to model a compressible liquid by specifying density using a field function that is defined as:

$$\rho = \rho_0 + \frac{p}{c^2} \tag{86}$$

where ρ_0 is a constant, p is the pressure, and c is the (constant) speed of sound. In this example, the following function would be specified for Density Pressure Derivative:

$$\frac{d(\rho)}{d(p)} = \frac{1}{c^2} \tag{87}$$

and the constant c would be specified for the speed of sound.

When you specify thermodynamic properties such as density and enthalpy using tabular input data, STAR-CCM+ can evaluate the necessary derivative terms with respect to pressure and temperature by numerically differentiating the tabular data. This calculation eliminates the need to provide more columns of data for the derivative terms.

All other models that interact with the Equation of State model continue to function as usual when the user-defined model is active.

Summary of Available Methods

The following methods are provided when the User Defined EOS model is active:

Available methods for Density:

- Constant
- Field Function

  Additional methods for Density when an energy model is activated:

- Polynomial in T
- Table(T) [EOS Compressible: off]; the density is specified using tabular data as a function of temperature (T).
- Table(T,P) [EOS Compressible: on]; the density is specified using tabular data as a function of temperature (T) and pressure (P).

Available methods for Density Pressure Derivative [EOS Compressible: on]:

- Density Table: the derivative of density with respect to pressure is derived by numerically differentiating the tabular data that is used to specify density. This option is the only available choice when Table(T,P) is selected for Density.
- Constant
- Field Function
- Table(T,P); the derivative of density with respect to P is specified using tabular data as a function of T and P.

Available methods for Density Temperature Derivative (appears with energy model activated):

- Density Table: the derivative of density with respect to temperature is derived by numerically differentiating the tabular data that are used to specify density. This option is the only available choice when Table(T,P) or Table(T) is selected for Density.
- Constant
- Field Function
- Table(T) [EOS Compressible: off]; the derivative of density with respect to T is specified using tabular data as a function of T.
- Table(T,P) [EOS Compressible: on]; the derivative of density with respect to T is specified using tabular data as a function of T and P.

Available methods for Enthalpy:

- Specific Heat: enthalpy is evaluated by integrating specific heat; this method is the default.
- Table, h(T); enthalpy is specified using tabular data as a function of T.
- Table, h(T,P); enthalpy is specified using tabular data as a function of T and P.

Available methods for Specific Heat:

- Constant
- Polynomial in T
- Thermodynamic Polynomial Data

  The above three methods are available when 'Specific Heat' is selected for Enthalpy.
- Enthalpy Table: specific heat is derived by numerically differentiating the tabular data that are used to specify enthalpy with respect to T. This is the only available choice when Table h(T,P) or Table h(T) is selected for Enthalpy.

Available methods for Speed of Sound [EOS Compressible: on]:

- Constant
- Field Function
- Table(T,P); speed of sound is specified using tabular data as a function of T and P.

Using Speed of Sound

The speed of sound is the rate of acoustic propagation in the medium. It is used at certain boundary conditions (pressure outlet, mass flow inlet, and free-stream) when the flow is compressible. (It also appears in the compressibility modification (dilatation dissipation) for the K-Epsilon and Reynolds Stress turbulence models.)

This material property becomes available once you have selected the User Defined EOS model and activated its Compressible property.



You can set the speed of sound for a material using a Constant, a Field Function, or (if the energy model is available) a Table(T,P).

For a multi-component material, you can set the speed of sound for the material as a property of the mixture. For a Eulerian multiphase material, you can set the speed of sound for each phase of the material as a property

of the phase. Eulerian multiphase materials use the speed of sound in the dissipation term in the K-Epsilon and Reynolds Stress turbulence models only.

Note: Eulerian multiphase models do not have a full implementation of this property for high Mach number (M > 0.3).

Selecting an Equation of State Model

Selection of an Equation of State model depends on the selection of a single-phase (single- or multi-component) or multiphase material model.

For a single-phase material model, you select an Equation of State after you select the Coupled Flow model or select the Segregated Flow Model.

- In the *Physics Model Selection* dialog, in the *Equation of State* box, select the relevant model by clicking the corresponding radio button:



The *Physics Model Selection* dialog shows the Equation of State model activated.



For a Real Gas model, seven options become available:

- Peng-Robinson
- Redlich-Kwong
- Soave-Redlich-Kwong
- Modified Soave-Redlich-Kwong
- van der Waals
- IAPWS-IF97 (for steam)
- Equilibrium Air (for air)

Select the appropriate one.



For a multiphase material model, the Eulerian Multiphase model and the Multiphase Equation of State model are automatically activated. The model allows you to specify a Eulerian Phase Equation of State model for each of the Eulerian phases that you define.



Entropy Reference

When the Ideal Gas model is used, entropy is defined as:

$$S = \frac{(T/T_{ref})^{(C_p/R)}}{(P/P_{ref})} \tag{88}$$

where:

- T is temperature K and T_{ref} is the reference temperature.
- P is absolute pressure Pa and P_{ref} is the reference pressure.
- C_p is specific heat J/(kg K), a constant.
- R is a specific gas constant J/(kg K).

This expression can be rewritten as:

$$R\ln(S) = C_p\ln(T/T_{ref}) - R\ln(P/P_{ref}) \tag{89}$$

The more familiar representation of entropy s is:

$$s = C_p\ln\left(\frac{T}{T_{ref}}\right) - R\ln\left(\frac{P}{P_{ref}}\right) \tag{90}$$

Comparison of Eqns. (89) and (90) shows that entropy s can be related to the STAR-CCM+ entropy field function S as follows:

$s = R\ln(S)$ where $T_{ref} = 1\text{K}$ and $P_{ref} = 1\text{Pa}$.

When the IAPWS-IF97 model is activated, entropy is calculated using that method.

Equation of State Field Functions Reference

The following primitive field functions are made available to the simulation when the Ideal Gas model, a Real Gas model (Redlich-Kwong, van der Waals), or the IAPWS-IF97 model is used. These field functions have a typical set of base properties.

Ratio of Specific Heats

Expressed as:

$$\gamma = \frac{C_p}{C_v} = \frac{C_p}{(C_p - R)}$$

where C_p is specific heat J/(kg K), and R is a specific gas constant J/(kg K).

Real Gas/IAPWS-IF97 Field Functions Reference

The primitive field functions in the table below are made available to the simulation when one of the following models is used:

- Peng-Robinson
- Redlich-Kwong
- Soave-Redlich-Kwong
- Modified Soave-Redlich-Kwong
- van der Waals
- IAPWS-IF97

These field functions have a typical set of base properties.

Critical Pressure	Expressed as p_c.
Critical Temperature	Expressed as T_c.
Reduced Pressure	Expressed as $p_r = p/p_c$.
Reduced Temperature	Expressed as $T_r = T/T_c$.

Compressibility Factor	Represents the amount of deviation of the real gas from an ideal gas, expressed as $Z = p/(\rho RT)$. This factor usually takes on values around 1. For the Equilibrium Air model, the compressibility factor represents the degree of dissociation, which is non-ideal behavior, but has values that vary more widely (approximately between 1 and 6). A simplified formula that applies to the equilibrium air model is given as Eqn. (158).

The following primitive field function is available only with the Peng-Robinson, Redlich-Kwong, Soave-Redlich-Kwong, Modified Soave-Redlich-Kwong, or van der Waals model. It has a typical set of base properties.

Real Gas Specific Heat	The value of specific heat at constant pressure of the real gas. This field function accounts for the effects of variations in both temperature and pressure by including the departure from the ideal gas value. The last three terms on the right-hand side of Eqn. (91) give the departure from ideal:

$$C_p = C_p^0 + T\int_{\infty}^{v}\left(\frac{\partial^2 P}{\partial T^2}\right)_v dv - T\left(\frac{\partial P}{\partial T}\right)_v^2\left(\frac{\partial P}{\partial v}\right)_T^{-1} - R \quad (91)$$

The following primitive field function is available only with either IAPWS-IF97 model. It has a typical set of base properties.

IAPWS Region ID	Data that STAR-CCM+ calculates for a particular IAPWS validity region.

The following primitive field functions are available only with the Thermal Non-Equilibrium model. They have a typical set of base properties.

Temperature of Vibrational Mode	Vibrational temperature (characteristic temperature of the vibrational/electronic modes).
Vibrational Enthalpy	Energy associated with the vibrational and electronic energy modes.

Combined Total Enthalpy	Sum of total translational-rotational enthalpy with vibrational enthalpy.
Vibrational Specific Heat	Specific heat for the vibrational/electronic modes.
Mode Relaxation Time	Relaxation time between the vibrational mode and translational mode.
Vibrational Conductivity	Thermal conductivity associated with the vibrational/electronic energy modes.

How Do I Choose Between Coupled and Segregated?

To guide the choice between the Segregated Flow model and the Coupled Flow model, consider their relative strengths and weaknesses:

- The segregated algorithm uses less memory than the coupled.
- The coupled algorithm yields more robust and accurate solutions in compressible flow, particularly in the presence of shocks.
- The coupled algorithm is more robust for high-Rayleigh number natural convection.
- The number of iterations that the coupled algorithm requires to solve a given flow problem is independent of mesh size. However, the number of iterations that the segregated algorithm requires increases with mesh size.
- In some situations the coupled algorithm can be combined with the implicit solver to permit large CFL numbers. This scenario would be analogous to an under-relaxation factor of 1 for all variables in a segregated algorithm. In contrast, the segregated algorithm needs significant under-relaxation for both velocity and pressure and, in compressible flows, energy.

With these strengths and weaknesses in mind, it is suggested that you proceed as follows to select the algorithm:

- Choose the Coupled Flow and Coupled Energy models for compressible flows, natural convection problems, and flows with large body forces or energy sources.
- If computational resources are not an issue, choose the Coupled Flow model for incompressible and/or isothermal flows.
- Choose the Segregated Flow model for incompressible or mildly compressible flows.

Modeling Flow and Energy Using a Coupled Approach

This section explains the following:

- The coupled flow and coupled energy models
- The coupled implicit and coupled explicit solvers
- Coupled flow controls and general setup recommendations
- How to select the coupled flow and coupled energy models
- Running a steady state analysis
- Running a transient analysis
- Troubleshooting

The Coupled Flow Model

The Coupled Flow model solves the conservation equations for mass, momentum, and energy simultaneously using a pseudo-time-marching approach.

One advantage of this formulation is its robustness for solving flows with dominant source terms, such as rotation. Another advantage of the coupled solver is that CPU time scales linearly with cell count; in other words, the convergence rate does not deteriorate as the mesh is refined.

This model can also evaluate inviscid fluxes with the optional AUSM+ scheme, which offers advantages for various cases.

General Preconditioning

The Preconditioning Enabled expert property includes a preconditioning matrix in the governing equations for the Coupled Flow model, in both original and discretized form. Using this matrix improves convergence and accuracy. This improvement applies to solving both compressible and incompressible flows at all speeds.

Unsteady Preconditioning

The Unsteady Low-Mach Preconditioning and Unsteady Preconditioning Max Factor expert properties control treatment of mass-flux pressure dissipation, D_p (described in [31]). When you activate unsteady preconditioning, D_p is a function of the ratio of a local time-step to the user-defined time-step. This local time-step is obtained without preconditioning, assuming CFL = 1.

At high speeds, no unsteady modification is done for D_p. At low speeds (approaching and in the incompressible regime), D_p depends on local flow conditions and also on the local Strouhal number.

At high Strouhal numbers (small time-steps), D_p corresponds to the unpreconditioned mass-flux pressure dissipation. At low Strouhal numbers and towards steady-state (when time-step tends to infinity), D_p moves towards the steady-state low-Mach preconditioned mass-flux pressure dissipation. This move happens because the dissipation depends on the ratio of the local time-step to the user-defined time-step, which decreases as the time-step increases.

Coupled Flow Properties

Integration — Switches between implicit and explicit pseudo-time integration.

- Implicit — The coupled implicit solver is invoked. Adds a Coupled Implicit node to the Solvers node with the option to set a Courant number ramp, grid sequencing, and expert driver for the solver.
- Explicit — The coupled explicit solver is invoked. Adds a Coupled Explicit node to the Solvers node with the option to set a Courant number ramp.

Discretization — Selects the discretization scheme that is used for evaluating face values for convection and diffusion fluxes.

- 1st-order — Selects the first-order upwind discretization scheme.
- 2nd-order — Selects the second-order upwind discretization scheme.
- Central — Selects the central-differencing discretization scheme (available with Large Eddy Simulation (LES) only).
- Bounded-Central — Selects the bounded central-differencing discretization scheme (available with Large Eddy Simulation (LES) only). Adds a Bounded Differencing sub-node to the Coupled Flow node.

Hybrid	Selects the hybrid second-order upwind/central-differencing discretization scheme (available with Detached Eddy Simulation (DES) only).
Hybrid-BCD	Selects the hybrid second-order upwind/bounded central-differencing discretization scheme (available with Detached Eddy Simulation (DES) only). Adds a Bounded Differencing sub-node to the Coupled Flow node.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that STAR-CCM+ uses, do not change properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance.

Unsteady Low-Mach Preconditioning	Activate/deactivate unsteady preconditioning to improve the accuracy for sound propagation in low-speed compressible regimes.
	Activated — Use preconditioning modification for unsteady low-Mach conditions. Only active when Preconditioning Enabled is active.
	Deactivated — Use no preconditioning modification for unsteady low-Mach conditions. The default.
Unsteady Preconditioning Max Factor	The maximum factor for unsteady preconditioning, a value between 0 and 1. The default is 0.95. Larger values can reduce robustness. At zero, the solver uses no high Strouhal number unsteady numerical dissipation.
Flow Boundary Diffusion	Activate/deactivate the flow-boundary diffusion fluxes (or viscous fluxes for flow models)—see Eqn. (2558). This property is activated by default.
	Activated — Activate flow-boundary diffusion.

	Deactivated	Deactivate flow-boundary diffusion.
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Coupled Inviscid Flux	Selects the scheme for the coupled inviscid flux.	
	Roe FDS	Uses the Weiss-Smith preconditioned Roe's flux-difference splitting scheme.
	AUSM+ FVS	Uses Liou's AUSM+ flux-vector splitting scheme.
Positivity Rate Limit	Maximum rate by which temperature is allowed to decrease.	
Explicit Relaxation	A scaling factor that is used to relax all corrections explicitly before they are applied to the variables.	
Minimum Absolute Pressure	Minimum allowable absolute pressure.	
Preconditioning Enabled	Activate/deactivate preconditioning.	
	Activated	Activate preconditioning.
	Deactivated	Deactivate preconditioning.
Minimum Reference Velocity	$U_{r,\,\min}$ from Eqn. (172).	

Maximum Reference Velocity	$U_{r,\max}$ from Eqn. (172).
Pressure difference scale factor	ε from Eqn. (172).

What Is Bounded Differencing?

Bounded differencing lets you specify parameters for:

- The bounded central-differencing discretization scheme when you work with Large Eddy Simulation (LES).
- The hybrid second-order upwind/bounded central-differencing discretization scheme when you work with Detached Eddy Simulation (DES).

The Bounded Differencing node, which has its own properties, is activated when the Discretization property of the Coupled Flow model is set to `Bounded-Central` (with LES) or `Hybrid-BCD` (with DES).



Properties

Upwind blending factor	Specifies the proportion of upwind differencing, ς_{ubf} related to Eqn. (2547). The default maximum of this property, `0.15`, provides the most robustness for the scheme. Reducing it would, in principle, increase accuracy. However, unless you are thoroughly familiar with the theoretical aspects of bounded differencing, do not change this property. The default value reflects optimization for accuracy and performance.

What Is the AUSM+ Scheme?

An optional scheme for evaluating inviscid fluxes is Liou's AUSM+ flux-vector splitting scheme ([35], [17]). Like the Weiss-Smith preconditioned Roe's flux-difference splitting scheme, this scheme is also based on the upwind concept, which is applied consistently on the two physically distinct parts of the inviscid flux: convective and pressure. The AUSM+ inviscid flux scheme is the more appropriate choice for solving flows involving high supersonic and/or hypersonic regimes (Mach number of around 3 and above), when compared with the Roe scheme.

The main advantages of the AUSM+ scheme are:

- Algorithmic simplicity and straightforward extension for complex conservation laws.
- Accurate capture of shock- and contact-discontinuity.
- Numerical solutions that preserve positivity and satisfy entropy.
- Reduced susceptibility to "carbuncle" phenomena.

AUSM+-up is the low-Mach preconditioned formulation that is used in STAR-CCM+. It is included primarily as an option for use with high supersonic and/or hypersonic flows. It also provides uniform accuracy and robust convergence for all Mach number flows.

The AUSM+-up scheme can be employed for solving a wide range of flow problems including incompressible to hypersonic aerodynamics, turbomachinery flows, LES/DES of turbulent flows and aeroacoustics.

The Coupled Energy Model

The Coupled Energy model is an extension of the Coupled Flow model. Together they solve the conservation equations for mass, momentum, and energy simultaneously using a time- (or pseudo-time-) marching approach. This formulation is robust for solving compressible flows and flows with dominant source terms, such as buoyancy.

Due to the preconditioned form of the governing equations that the Coupled Flow and Coupled Energy models use, convergence rate is effectively independent of Mach number, ranging from incompressible through to supersonic regimes.

Enthalpy Formulation

The Enthalpy Formulation option on this model provides a choice between using enthalpy or temperature for the dependent variable when solving the energy conservation equation. This option affects how the energy equation is linearized. It is intended for use with combustion models where changes in enthalpy tend to be much smoother than the steep temperature gradients often found in reaction zones. The option can be considered to provide the same choice for the coupled solver as the Segregated Fluid Enthalpy and Segregated Fluid Temperature models offer for the Segregated Flow solver.

Since the Coupled Energy model is an extension of the Coupled Flow model, it has no properties of its own. Instead, it shares the properties that the Coupled Flow Model node provides.

Coupled Energy Properties

Enthalpy Formulation — Activates or deactivates the enthalpy formulation for the coupled energy model. When activated, enthalpy is the dependent variable instead of temperature.

- Activated — Sets the dependent variable to enthalpy.
- Deactivated — Sets the dependent variable to temperature.

Selecting the Coupled Flow Model

To select the Coupled Flow model, first select a Gas or a Liquid Material model. This selection exposes the Flow models. Select the Coupled Flow radio button.



The *Physics Model Selection* dialog with the Coupled Flow model selected is shown below:



Selecting the Coupled Energy Model

The Coupled Energy model is available for selection only after the Coupled Flow model has been selected. The *Physics Model Selection* dialog looks as shown below, with the Coupled Energy model available in the Optional Models box. Select the Coupled Energy checkbox.



Running a Steady State Analysis

For steady simulations, the coupled solver in STAR-CCM+ employs a time marching scheme to drive the unsteady form of the governing equations to a steady state. In this case, a pseudo-transient term replaces the physical time-derivative. The solution advances in pseudo-time to drive this term to zero. The solution in each cell is advanced independently with an optimal pseudo-time step computed locally according to stability constraints. In this way, convergence to steady state is achieved in the most efficient manner.

Running a Transient Analysis

Transient analysis in STAR-CCM+ is used to solve time-accurate problems, including initial value problems.

Two different time integration schemes are available for transient analysis:

- Implicit Unsteady
- Explicit Unsteady

You choose between these two by activating either the Implicit Unsteady Model or the Explicit Unsteady Model, while also selecting the Coupled Flow model. The Coupled Energy model can be chosen as an option.

What Is Implicit Unsteady?

The Implicit Unsteady approach is appropriate if the time scales of the phenomena of interest are either of the following:

- The same order as the convection and/or diffusion processes (for example, vortex shedding).
- Related to some relatively low frequency external excitation (for example, time-varying boundary conditions or boundary motion).

In the Implicit Unsteady approach, each physical time-step involves some number of inner iterations to converge the solution for that given instant of time. These inner iterations can be accomplished using implicit spatial integration or explicit spatial integration schemes. You specify the physical time-step size that is used in the outer loop. The integration scheme marches inner iterations using optimal pseudo-time steps that are determined from the Courant number.

With the Implicit Unsteady approach, you are required to set the physical time-step size, the Courant number, and the number of inner iterations at each physical time-step.

The transient phenomena being modeled generally governs the physical time-step size. The time step must at least satisfy the Nyquist sampling criterion: more than two time-steps per period are required.

In general, the guidelines for setting the Courant number in steady-state integration schemes also apply to the transient integration schemes.

The number of inner iterations per physical time-step is harder to quantify. Generally, you determine this number by observing the effect that it has on results. Select a number of inner iterations, plot a monitor for one or more specific quantities against iteration, and see whether these monitors are converging within each time-step.

Smaller physical time-steps generally mean that the solution is changing less from one time step to the next; fewer inner iterations are then required. There is an optimal balance of time-step size and number of inner iterations for a given problem and desired transient accuracy.

What Is Explicit Unsteady?

If the unsteady time scales are of the order of the acoustic processes (for example, shock front tracking), the Explicit Unsteady approach is the proper choice.

The Explicit Unsteady approach is effectively the same as the explicit spatial integration scheme, but using the same physical time step for all cells in the domain (rather than local time-stepping). In this case, the solver automatically determines the size of the time-step as Eqn. (199) such that one value satisfies the Courant condition at all points; the minimum allowable time-step is used everywhere. Thus, each iteration becomes a time-accurate advancement of the solution. In this case, the Courant number must be 1 or less, resulting in a physical time-step that varies from one iteration to the next as the flow field changes. Furthermore, preconditioning of the governing equations is omitted; the Explicit Unsteady approach is therefore unsuited for incompressible flow simulations.

What Is the Coupled Implicit Solver?

The Coupled Implicit solver controls the solution update for the Coupled Flow model. If the Coupled Energy Model and the Coupled Species Model are activated, the Coupled Implicit solver controls these models also. The solver is used for implicit spatial integration in both steady and unsteady analyses, using a coupled algebraic multi-grid method.



To set the Courant number for this solver, enter the value in the properties.

The default value of Courant number for the implicit scheme is 50 for unsteady flows, and 5 for steady flows. While these values are fine for many simulations, they are not optimal in all cases.

Generally, there is some optimal value of Courant number that balances the total number of iterations against the time per iteration. This optimum value is a function of many factors: whether the flow is compressible, how close the solution is to convergence, mesh quality and the magnitude of the body forces.

Solving the governing discretized equations implicitly in pseudo-time results in a wider stability margin; this approach permits Courant numbers greater than unity. The large resulting local pseudo-time steps $\Delta\tau$, given by

Eqn. (197), provide relatively fast convergence rates. The penalty for this, however, is a larger storage requirement than the explicit spatial integration scheme.

The following guidelines are recommended:

- If the solution diverges for more than a few iterations after initialization, or if the solver reports that corrections are being limited, it is necessary to reduce the Courant number.
- If the solver reduces the Courant number automatically for more than 1 or 2 iterations, the Courant number must be manually reduced.
- If the Courant number is too high, especially during initial iterations, the convergence can be slowed due to the multigrid solver working too hard. A suggested procedure for monitoring this behavior is as follows:
  - Select the Solvers ⚲ Coupled Implicit ⚲ AMG Linear Solver node.
  - To display the number of inner multigrid, set the Verbosity property to low.
  - There must be about 3–4 inner iterations per global iteration. If there are many more, reduce the Courant number.
- If the Courant number is set very high and the residuals do not decline at some point after many iterations, try reducing the Courant number.
- Consider using ramping to increase the CFL automatically to its final value during the early iterations.

The following sections discuss how to set the following:

- Algebraic multigrid parameters for the Coupled Implicit solver
- Expert initialization for grid-sequencing parameters
- Expert driver parameters
- Convergence accelerators

Coupled Implicit Solver Properties

Courant Number

The value of CFL in Eqn. (197) that controls the local time-step. The default Courant number is 50 for unsteady flows, and 5 for steady flows. In general, and depending on the initial conditions, it is good practice to start at a lower Courant number, and then increase it.

Freeze Flow

Disables the flow solver solution update.

- **Activated** — The flow solver is skipped so that the flow field solution is not updated.
- **Deactivated** — The flow solver is invoked as usual and the flow field solution updated accordingly. (The default.)

Expert Properties

Pressure Reference Location

Provides a choice between using the automatic algorithm of STAR-CCM+ for the reference location or providing it manually, as discussed elsewhere.

- **Automatic Selection** — STAR-CCM+ sets the reference location at the cell that is adjacent to the boundary face that has the smallest X, Y, Z location in the domain.
- **User Specified** — The Pressure Reference Points child node is added. Use this property to add one or more reference locations depending on the number of non-contiguous regions.

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

- **Activated** — Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.
- **Deactivated** — Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

- Activated: Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.
- Deactivated: Do not reset reconstruction gradients. They are recomputed at the next iteration if they had been zeroed.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

- Activated: Retain temporary storage.
- Deactivated: Do not retain temporary storage.

For information about other properties in the *Expert* category, see Expert Properties, p. 6385.

Setting Algebraic Multigrid Parameters

To set the algebraic multigrid parameters for either the pressure solver or the velocity solver:

- Open the node for that solver under the Coupled Implicit node in the Solvers node.
- Open the node for the AMG Linear Solver.

- Enter the algebraic multigrid parameters using the *Properties* window.



Using Grid-Sequencing Initialization

The Expert Initialization node is a sub-node of the Coupled Implicit solver. You can activate or deactivate grid sequencing in its properties.



Expert Initialization Properties

Method	Activates or deactivates grid sequencing for calculation.	
	None	Deactivates grid sequencing.
	Grid Sequencing	Activates grid sequencing. The Grid sequencing node becomes available in the Expert Initialization node.

Setting Grid Sequencing Parameters

With the Coupled Solver, you can select the grid-sequencing (GS) expert initialization, which in essence performs the normal initialization followed by the computation of an approximate inviscid solution to the flow problem. Only flow variables (such as pressure, velocity, and temperature) and species (for mixtures of gases) are initialized.

The Grid Sequencing node allows you to set parameters with its properties. It becomes available when the Method property of the Expert Initialization node is set to `Grid Sequencing`.



When you activate this feature when the flow solution is initialized, the GS initialization performs the following steps (this applies to initializing manually or starting a simulation run):

- A series of coarse meshes is generated.
- Normal initialization of the flow solution is performed on each of the coarse meshes.
- Starting with the coarsest mesh:
  - A number of iterations are run to compute an approximate solution on the current mesh.
  - If either convergence or the maximum number of iterations is reached, the solution is interpolated on the next finer mesh and (unless the next mesh is the finest), the previous step is repeated.
- The process stops when GS has reached the finest mesh level.

GS uses a full implicit incomplete-Newton solution algorithm to compute a first-order inviscid flow solution, which has the following advantages:

- Increased robustness, due to the full implicit algorithm.
- Faster convergence, due to relatively large CFL numbers that can be used.
- Fully automatic coarse-mesh generation and solution.
- An inviscid solution that is a much better approximation of the flow solution than that obtained through normal initialization.
- Faster and more robust convergence of the flow solution, after GS initialization (for both the Coupled and Segregated solvers).

Flexibility

You can interrupt GS at any time by clicking the (Abort) button next to the status bar. (In this event, if a valid coarse mesh solution exists, it is interpolated onto the finest mesh, and it is fully usable.)

You can change GS settings at any time during the initialization. (The solver acknowledges the change through a message in the *Output* window.)

If GS initialization fails (for example due to aggressive settings), the solver attempts to recover from it. At worst, you can simply get a normal (uniform) initialization. In this case, change the GS settings. For example, choose smaller values for the CFL number or the maximum iterations per level, clear the solution and try the initialization again. Also check that the solver is working with a reasonable number of cells per level to converge. The default number of maximum grid levels may not be suitable for problems with a relatively small number of cells in the mesh.

Limitations

GS is not available as an expert initialization method for simulation of heat conduction in solids.

GS, which is regarded as a tool to improve robustness and speed for convergence, is not guaranteed to work in all foreseeable cases. Deactivate it if you encounter situations such as lack of robustness or convergence during the GS process.

Grid Sequencing Properties

Maximum grid levels	Number of coarse grid levels, created automatically, which the grid-sequencing (GS) algorithm uses to compute inviscid solutions, starting from the coarsest mesh.
Maximum iterations per level	Number of solver iterations that are performed during GS, to compute an approximate inviscid solution of the flow problem, on the current coarse mesh.
Convergence tolerance per level	Solution convergence criterion which, together with the maximum number of iterations per level, the solver uses to stop cycling on the current mesh level.
CFL number	Courant-Friedrichs-Levy (CFL) number that the flow solver uses during the grid-sequencing initialization. This number can be different from the Courant number that the solver uses during iterations on the finest mesh (normal run).

Using the Expert Driver

STAR-CCM+ features an automated convergence control tool within the Coupled Implicit solver. This tool, which is known as the Expert Driver, can be activated through the properties of the Solution Driver node that is located beneath the Coupled Implicit solver node.



Solution Driver Properties

Driver	Activates or deactivates the Expert Driver for calculation.	
	None	Deactivates the Expert Driver.
	Expert driver	Activates the Expert Driver. The Expert Driver node becomes available in the Solution Driver node.

This section covers the following:

- Setting Expert Driver parameters
- Plotting solver quantities to assess Expert Driver performance

Setting Expert Driver Parameters

The Expert Driver is in essence a "push-button" convergence control for steady-state simulations with the Coupled Implicit solver. It is a recommended tool for improving the robustness of the Coupled Implicit solver, and increasing the speed of convergence. However, it is not guaranteed to provide the best solution control in all cases. If the solution slows down or becomes unstable, you may need to tune (or deactivate) the feature.

The Expert Driver provides the following levels of control:

1. CFL Control: Automatically modifies the CFL (Courant) number each iteration to ensure that the AMG linear solver converges. The Courant Number property of the Coupled Implicit solver node sets the maximum CFL number. On the Expert Driver node, properties that are connected to this control are:
   - CFL Ramp: Start Iteration, End Iteration, and Start Value. These properties provide an initial linear ramping of the CFL number.
   - CFL Recovery Rate: when the AMG Linear Solver fails to converge, the Coupled solver reduces the CFL number. The Expert Driver can reduce the impact of this reduction by helping to recover the CFL number back to its desired ramp or target value. This property defines this recovery rate.
2. CFL Limiting: Adds an extra layer of solver protection by adjusting the CFL number to achieve a balance between the flow solution coupling and the efficiency of the linear system solution. The Expert Driver achieves this balance by controlling the AMG linear solver convergence behavior so that the number of AMG cycles is as close as possible to the target value given in the CFL Limiting: Target AMG Cycles property.

3. Corrections Control: Computes the under-relaxation factor that gives the quickest rate of decrease in the residual for each coupled flow model. The CFL number control is independent of this property. On the Expert Driver node, properties that are associated with the Corrections Control are:
   - Corrections Control: Minimum Explicit Relaxation, Maximum Explicit Relaxation, URF Recovery Rate, and Relative Residual Cutoff.
   - Corrections Control Update Frequency: Unlike CFL Control and CFL Limiting, which apply to every iteration, the Corrections Control is only updated at the frequency this property sets.
4. Expert-Driver Start Iteration and Cutoff Iteration: Automatically enables and disables the expert driver after the specified number of iterations.
5. Enhanced Dissipation: Significantly improves the convergence and robustness of the Coupled solver during the initial stages of a simulation. The increased numerical dissipation that it provides can lead to inaccuracies in the flow solution. Enhanced Dissipation is available for incompressible flows only. When the simulation reaches the iteration count set by Enhanced Dissipation: Start transition, the solver starts to reduce the amount of extra dissipation. The smooth non-linear reduction continues until the simulation reaches the iteration count set by Enhanced Dissipation: End transition. At this point the extra dissipation is removed completely.

Note: To turn this property off, set both Start and End transition to 1. In all cases, the solver uses extra dissipation for the first 25 iterations, even when Enhanced Dissipation is deactivated.

All relevant controls are provided in the properties of the Expert Driver node.



Expert Driver Properties

Verbose	Provides additional output for this value while the simulation is running. This property can be useful for debugging problems when they occur. Activated — Activates verbose output. Deactivated — Leaves output at its default.
CFL Ramp: Start Iteration	Specifies the iteration at which ramping begins.
CFL Ramp: End Iteration	Specifies the iteration at which ramping ends.
CFL Ramp: Start Value	The CFL value before the ramping begins.
Corrections Control	Optimizes the solution update. The CFL number control is independent of this property. Activated — Activates the corrections control. Deactivated — Deactivates this feature.
Corrections Control: Update Frequency	Specifies the number of iterations to perform before invoking this property.
Corrections Control: URF Recovery Rate	Specifies the rate of increase for the explicit relaxation after a solver-controlled reduction. The default value is `1.1`, which allows a maximum increase of 10%.
CFL Control	Modifies the CFL number appropriately at each iteration to ensure the AMG linear solver convergence, and provides a smooth recovery to the desired value when the CFL number is reduced. Activated — Activates the CFL control. Deactivated — Deactivates this feature.

CFL Limiting	Adjusts the CFL number to strike a balance between the flow solution coupling (as given by the CFL) and the efficiency of the linear system solution.	
	Activated	Activates the CFL limiting.
	Deactivated	Deactivates this feature.
Expert-Driver Start Iteration	Automatically enables the expert driver after the specified number of iterations. The default is 1.	
Expert-Driver Cut-off Iteration	Automatically disables the expert driver after the specified number of iterations. The default is 10000000.	

Expert Driver Expert Properties

Corrections Control: Minimum Explicit Relaxation	Sets the minimum scale to the update with which the Corrections Control property optimizes the solution update.
Corrections Control: Maximum Explicit Relaxation	Sets the maximum scale to the update with which the Corrections Control property optimizes the solution update.
Corrections Control: Relative Residual Cutoff	As the solution progresses towards convergence, this limits the Corrections Control property to the initial phases of the simulation as specified here.
Enhanced Dissipation: Start transition	Sets the iteration at which the solver begins to reduce the amount of extra dissipation. The default value is 1.
Enhanced Dissipation: End transition	By the iteration set in this property, no extra dissipation is included. The default value is 100.
CFL Control: Courant Number Recovery Rate	Specifies the rate of recovery to the desired value, in cases where the CFL Control property reduces the CFL number.
CFL Limiting: Target AMG Cycles	Controls CFL reduction. If the solver reduces the CFL number (when the AMG linear solver did not converge) and the number of AMG cycles (during the CFL recovery period) is larger than the specified “optimal” (or target) number of AMG cycles, the CFL number is additionally reduced (by 10% at a time).

Expert Driver Plots

STAR-CCM+ provides three field functions that allow you to monitor how well the Expert Driver controls the parameters of the Coupled Implicit solver. The functions are: CFL Number, Explicit Relaxation, and Number of AMG Cycles. Plotting these field functions during a simulation help you optimize the problem setup to achieve fastest possible convergence.

CFL Number

A history plot of the CFL Number that shows oscillatory behavior (around a mean value) suggests that you can lower the target CFL Number. The Courant Number property value of the Solvers > Coupled Implicit node defines this target. Conversely, if the CFL number plot shows a monotone increase towards its peak value, you can increase the Courant Number. (In this instance, the increase must be no more than 15–20% at a time).

When the AMG Linear Solver fails to converge, the Coupled solver reduces the CFL Number. The Expert Driver CFL Control helps to recover the CFL Number back to the desired ramp or target value. As recovering from a low CFL Number takes time, it is better to have an appropriate initial CFL Number that does not need to be reduced. Plotting the CFL Number helps in this respect.

Explicit Relaxation

If the history plot of Explicit Relaxation shows general oscillation, it is likely that the fluid problem being solved has a high degree of non-linearity. In this situation, you are recommended to lower the value of Corrections Control: Maximum Explicit Relaxation (and eventually also reduce the value of Corrections Control: Minimum Explicit Relaxation).

Mostly high values of Explicit Relaxation (close to or above `1.0`) indicate rapid and eventless convergence. A higher CFL Number could potentially be used.

Number of AMG Cycles

A history plot of the Number of AMG cycles shows the variations in computational load inside the linear-system solver (AMG). If the Number of AMG Cycles is much larger than the value that is set by the CFL Limiting: Target AMG Cycles property, reduce the target CFL Number for the Implicit Coupled solver.

A low Number of AMG cycles over an extended part of the simulation suggests that you could accelerate convergence by increasing the target CFL Number (but by not more than 15–20% at a time).

Creating the Plots

A plot of the CFL Number can be created from an Expression report:

- Right-click Reports and select New Report > Expression.
- Rename the newly created Expression 1 report to CFL Number.
- In the *Properties* window, set the Definition property to `$CFL`. This definition corresponds to the field function, CFL Number.



- Right-click Reports > CFL Number and select Create Monitor and Plot from Report.

The *CFL Number Monitor Plot* can be visualized when a simulation is running by right-clicking the Plots > CFL Number Monitor Plot node and selecting Open. An example is shown below.



Plots of Explicit Relaxation and Number of AMG cycles are available only when the Expert Driver is activated. The CFL Number can be plotted even though Expert Driver is not activated.

Create a plot of the Explicit Relaxation in the same way as shown for the CFL Number. In this instance, however, rename the Expression report to `Explicit Relaxation`, and set the Definition property to `$UrfED`. An example of this plot is shown below.

Create a plot of the Number of AMG cycles in the same way as shown for the CFL Number. Rename the Expression report to `Number of AMG Cycles`, and set the Definition property to `$NoAMGCyclesED`. An example of this plot is shown below.



Setting Convergence Accelerators

Some flow problems—for example, cases with a wide range of speeds in the domain (including super- or hypersonic flow)—require a lower CFL number for stability. This can happen when using both constant or flow-dependent material properties.

Typically, the mass-flow convergence is slow for these cases and these simulations can take many iterations to converge.

To speed up these computations, STAR-CCM+ features a mass convergence acceleration tool within the Coupled Implicit solver. This tool, the Continuity Convergence Accelerator, is activated through the Convergence Accelerators node that is located beneath the Coupled Implicit solver node.



Convergence Accelerators Properties

Convergence Accelerators	Activates or deactivates the Continuity Convergence Accelerator.	
	None	Deactivates the Continuity Convergence Accelerator. (Default)
	Continuity convergence accelerator	Activates the Continuity Convergence Accelerator. The Continuity Convergence Accelerator node becomes available in the Convergence Accelerators manager.

This section provides documentation on the Continuity Convergence Accelerator as follows:

- Setting Continuity Convergence Accelerator parameters
- Continuity Convergence Accelerator best practices guidelines
- Continuity Convergence Accelerator algorithm

Setting Continuity Convergence Accelerator Parameters

The Continuity Convergence Accelerator (CCA) formulates and solves a pressure-correction equation using the Density-based/Riemann flux discretization. The solution of this additional equation provides updates for pressure, velocity, and other flow field variables in such a way that the overall and individual cell mass imbalances are minimized at each iteration.

The Continuity Convergence Accelerator node represents this tool in the simulation tree.



Continuity Convergence Accelerator Properties

Verbose	Provides more output for this value while the simulation is running. This property is useful for debugging problems when they occur. Activated: Activates verbose output. Deactivated: Leaves output at its default.
Convergence Accelerator Update Frequency	The iteration frequency at which the Continuity Convergence Accelerator is invoked. The default value is 1, meaning that the Continuity Convergence Accelerator is invoked at each iteration.
Under-Relaxation Factor	ω in Eqn. (93). The default value is 0.3.
CCA Cutoff Iteration	Automatically disables the Continuity Convergence Accelerator after the specified number of iterations. The default is 10000000.

Continuity Convergence Accelerator Expert Properties

Pressure Reference Location	Provides a choice between allowing STAR-CCM+ to set the reference pressure location or providing it manually. This property is required for incompressible cases with no pressure boundaries and is ignored otherwise. Automatic Selection: STAR-CCM+ sets the reference location at the cell that is adjacent to the boundary face that has the smallest X, Y, Z location in the domain. (Default) User Specified: The Pressure Reference Points child node is added. Use this method to add one or more reference locations, depending on the number of non-contiguous regions.
Bounded Velocity Corrections	Whether to activate bounded velocity corrections and therefore limit the pressure correction gradient in the formula for velocity correction, Eqn. (95).

	Activated	Makes sure that velocity corrections are "smooth" to provide bounded velocity updates. (Default)
	Deactivated	Does not smooth velocity corrections.
Enhanced Mass-Imbalance Calculations	Whether to compute or correct the mass balances for step 6 of the Continuity Convergence Accelerator Algorithm	
	Activated	The final cell mass balances are computed using the Riemann solver fluxes depending on the flux scheme (see Eqn. (174) or Eqn. (177)). This option is more accurate, at the expense of computational cost.
	Deactivated	The mass flux is corrected. See Eqn. (225). This option is sufficient in most cases. (Default)
Frozen	Prevents or allows the Continuity Convergence Accelerator from updating the solution while iterating.	
	Activated	The Continuity Convergence Accelerator does not update the solution.
	Deactivated	The Continuity Convergence Accelerator updates the solution. (Default)

Continuity Convergence Accelerator Best Practices Guidelines

The Continuity Convergence Accelerator does not modify the final converged solution; it only modifies the path to convergence.

The Continuity Convergence Accelerator has a positive effect on convergence for cases with slow mass balance convergence due to strong non-linearity of the flow and the CFL limitations of the solver. However, the Continuity Convergence Accelerator is not a universal remedy to all convergence or robustness issues in the Coupled Implicit solver. Use the Continuity Convergence Accelerator only when mass-balance convergence is slow.

If convergence issues occur with the Continuity Convergence Accelerator, reduce the under-relaxation factor. For highly rotational and/or high speed flows, lower under-relaxation values and careful under-relaxation ramping (from 0.01 or lower) may be required.

Once activated, the Continuity Convergence Accelerator can be deactivated at any point in the simulation without any negative effect on the Coupled Implicit solver.

Continuity Convergence Accelerator Algorithm

The Continuity Convergence Accelerator pressure-correction discretization uses the Coupled Flow numerical dissipation, not the Rhie-Chow dissipation of the Segregated solver. Choosing the Coupled Flow numerical dissipation formulation over the Rhie-Chow formulation prevents some of the known issues of the Segregated solver, especially in high-speed regimes, such as wiggles at flow discontinuities due to improper numerical dissipation.

When activated, the Continuity Convergence Accelerator algorithm adds a post-update conditioning step to help satisfy better the continuity equation.

The Coupled Implicit solver algorithm with Continuity Convergence Accelerator is:

1. As part of the normal Coupled Implicit iteration update, solve the discretized coupled flow equations, and perform the normal update to obtain an intermediate flow field state p^*, v^* ...
2. Compute the new cell mass fluxes and mass imbalances for this intermediate state, m_f^*.
3. Discretize and solve the Continuity Convergence Accelerator pressure-correction equation with the new mass imbalances in the right-hand side:

$$a_p p'_p + \sum_n a_n p'_n = -\sum_n m_f^* \tag{92}$$

   to produce cell values of the pressure correction p'. The subscript n denotes all neighbors of cell p with faces f.
4. Update cell pressures using the under-relaxation factor ω:

$$p^{n+1} = p^* + \omega p' \tag{93}$$

5. Update the boundary pressure corrections p'_b, and then pressure p_b.
6. Correct the face mass fluxes:

$$m_f^{n+1} = m_f^* + m' \tag{94}$$

   where m' is computed using Eqn. (225).

   If the Enhanced Mass-Imbalance Calculations expert property is activated, the mass imbalances are recomputed using either Eqn. (174) or Eqn. (177).

7. Correct the cell velocities

$$v^{n+1} = v^* - f(\omega)\frac{V\nabla p'}{A_p v} \tag{95}$$

where:

- $f(\omega)$ is a non-linear function of ω
- $\nabla p'$ is the gradient of the pressure corrections
- A_p is the vector of central coefficients for the original coupled equations discretized system representing the velocity equations
- V is the cell volume

8. Update density, total enthalpy, and other variables with the pressure and velocity changes.

The Coupled Explicit Solver

The Coupled Explicit solver controls the solution update for the Coupled Flow model and Coupled Energy model. It is used for explicit integration using a Runge-Kutta multi-stage scheme. The Coupled Explicit node represents this solver in the Solvers node.



To set the Courant number for this solver, enter the value in the properties.

Generally, the Courant Number property is chosen to be 1 for unsteady analysis. For steady analysis, a value of 2 can often be used with two residual smoothing iterations. This restriction is much more severe than the implicit spatial integration scheme, although much less storage is required for the same size problem.

Residual smoothing can be employed along with the explicit spatial integration scheme. This approach helps widen the stability margin and allows slightly larger Courant numbers to be used.

Coupled Explicit Solver Properties

Courant Number

Controls the local time-step at each iteration according to Eqn. (197).

Freeze Flow

Disables the flow solver solution update.

- **Activated** — The flow solver is skipped so that the flow field solution is not updated.
- **Deactivated** — The flow solver is invoked as usual and the flow field solution updated accordingly. (The default.)

Multi-stage Coefficients

Sets the multi-stage coefficient values, see Eqn. (187). The number of values in the vector indicates the number of stages.

Dissipation-stage Flags

Indicates the stages of the multi-stage scheme on which dissipation (and viscous fluxes if appropriate) is evaluated. The number of these flags must be the same as the number of stages in the multi-stage coefficients.

Residual Smoothing

Activates or deactivates residual smoothing.

- **Activated** — Activates smoothing.
- **Deactivated** — Deactivates smoothing.

Residual Smoothing Iterations

Sets the number of residual smoothing iterations using Eqn. (190).

Residual Smoothing Underrelaxation

Sets the residual smoothing under-relaxation factor ε, see Eqn. (189). The default value is 0.5.

Expert Properties

Pressure Reference Location

Provides a choice between using the automatic algorithm of STAR-CCM+ for the reference location or providing it manually, as discussed elsewhere.

- **Automatic Selection** — STAR-CCM+ sets the reference location at the cell that is adjacent to the boundary face that has the smallest X, Y, Z location in the domain.
- **User Specified** — Adds the Pressure Reference Points child node. Use this node to add one or more reference locations depending on the number of non-contiguous regions.

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

- **Activated** — Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.
- **Deactivated** — Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

- **Activated** — Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.
- **Deactivated** — Do not reset reconstruction gradients. They are recomputed at the next iteration if they had been zeroed.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

Activated	Retain temporary storage.
Deactivated	Do not retain temporary storage.

For information about other properties in the *Expert* category, see Expert Properties, p. 6385.

Coupled Flow Controls and General Setup Recommendations

The STAR-CCM+ Coupled Flow model and Coupled Implicit Solver make available several parameters which allow control of stability and convergence for practical CFD cases:

- The Courant Number (CFL) property of the Coupled Implicit and Coupled Explicit solvers controls the size of the local time-steps that are used in the time-marching procedure these two solvers employ. The Courant number plays the same role as under-relaxation parameters in the segregated solver.

  Generally speaking, for a steady-state simulation, a larger CFL number increases the local pseudo-time step size and produces faster convergence. Thus, use the largest CFL number possible while still ensuring that the solver remains within the bounds of stability. (This means that the solver does not start reducing corrections, or hitting temperature and pressure limits, and that the residual has a healthy downward trend.)

- The Explicit Relaxation (an expert property of the Coupled Flow Model) is a scaling factor that is used to relax all corrections explicitly before they are applied to the variables. It is intended to add stability in cases where the solver's first-order linearization of a second-order discretization has trouble. The default value is 1.0, which results in 100% of the corrections being applied. Use a value less than 1.0 to limit the applied corrections explicitly.

- The Positivity Rate Limit (an expert property of the Coupled Flow Model) defines the amount that temperature, a positive definite quantity, is allowed to decrease when corrections are applied at each iteration. It is expressed as a percentage of the current value.

  The default is 0.2, which says temperature is not allowed to decrease by more than 20% (there is no upward limit). If temperature corrections are reduced to meet the Positivity Rate Limit, then all other variable corrections are equally scaled in an attempt to keep everything changing at the same rate.

Note: The last two Coupled Flow model expert properties are linked, in that the Explicit relaxation factor is applied to the corrections first, then the Positivity Rate Limit condition is checked. It is possible that the Explicit relaxation factor could reduce temperature corrections such that the Positivity Rate Limit does not get invoked where it otherwise would.

- Deactivate the Preconditioning Enabled option (an expert property of the Coupled Flow Model) when you are trying to capture unsteady phenomena at small time scales (for instance, in an aeroacoustics analysis). Deactivate this property when the Integration property is set to either Implicit or Explicit.

Setup Recommendations

The setup recommendations are as follows:

- The default CFL number for the Coupled Implicit solver is 50.0.
- You can often increase the CFL number value to 100, 200 or even more, depending on the physical and numerical difficulty of the actual flow problem being solved and grid quality. Laminar flow and good grid quality allow even higher values and very rapid convergence. When using the Coupled Explicit Solver, limit the CFL number to values near unity.
- Turbulent flows require lower CFL numbers and more iterations than laminar flows, since equations for turbulence quantities are solved in a segregated manner outside the coupled solver. However, even for turbulent flows, the CFL number can be increased beyond 1000, especially when the mesh is refined and the simulation is initialized with a coarser mesh solution.
- In many situations, it is helpful to perform some iterations in the segregated solver before switching to the coupled solver. This procedure provides a better initial solution and allows you to set a higher Courant number in the coupled solver.
- If an approximate initial solution is used (uniform values), smaller CFL number values are required at startup.
- You can employ the Linear Ramp capability of the Coupled Implicit solver, to automate desired change of the CFL number during the simulation.
- A low CFL number value of 0.1, or even lower, may be required for startup of some difficult problems (for example, hypersonics or combustion). To ensure that the solver stability is maintained during the entire simulation, use the Linear Ramp to increase the CFL number based on iteration count.

- Depending on the problem (such as very high-speed flow or mesh quality issues) you could use an Explicit relaxation factor as low as 0.2 or 0.25. Very difficult cases may require a value of 0.15 to be able to maintain non-linear stability.
- Typically an Explicit relaxation factor of 0.75–0.85 helps convergence for most problems, without greatly affecting the speed of convergence for those cases that do not require any explicit relaxation.
- Besides the CFL number and the Explicit relaxation factor, a practical way to keep hypersonic cases within physical "bounds" is to use a lower value for the Positivity Rate Limit. The default value is 0.2, although for extreme hypersonic cases it may be necessary to use values as low as 0.05.
- For time accurate simulations, choose the time step to keep the CFL number around 1.0 on average. In this case, it is possible to set under-relaxation factors to 1.0 for all transport equations. The CFL value can be larger where the mesh is fine, and smaller where cells are large.

Mesh Quality Considerations

The mesh quality considerations are as follows:

- Mesh quality does influence the quality of the numerical results, irrespective of the setup: coupled vs. segregated solver, discretization schemes...
- On a mesh of a given quality and sufficient fineness, higher-order schemes yield more accurate results than lower-order ones. For a given cell count, aim to optimize the grid quality because a better mesh gives a more accurate solution regardless of the differencing schemes that are used and other parameters.
- With a poor-quality mesh, the gradient limiter is invoked more often, thus affecting its nominal second-order accuracy. Grid skewing is an important contributing factor (much more than grid stretching) to the loss in nominal accuracy of the solution. Grid design (distribution of cell size, local refinement using feature edges, boundary region, or volume shapes) is important in maximizing the accuracy for a given effort. Two grids with the same number of cells may lead to discretization errors that differ by an order of magnitude. Thus, improving mesh quality is a top priority.
- When solving steady-state problems on fine meshes, start with a much coarser mesh, and then successively refine the mesh by reducing the base size. For example:
    - Design the desired mesh for the highest cell count affordable.

- Increase the base size by a factor of 8 and start the computation with this coarse mesh.

  You may require lower under-relaxation values for the coarsest mesh, but iterations are fast and this does not increase the computing effort much.

- After the convergence criteria are satisfied or maximum specified number of iterations is reached, halve the mesh base size, remesh, and restart the simulation. Repeat until the final mesh size is reached.

In this way, you obtain faster convergence on finer grids by providing a good initial solution. Usually, one needs only a third or a quarter of the number of iterations that would be needed if starting with an initial guess, like constant or zero values.

In addition, you obtain solutions on a series of grids which are systematically refined (same design, only base sizes reduced), which allows for an estimate of discretization errors using Richardson extrapolation. For example, if base size is reduced by a factor of 2 and second-order discretization is used, discretization errors on the finest mesh are equal to one third of the difference between solutions on the finest and next coarser grid. For base size reduction by a factor of 1.5, the errors amount to about 80% of the difference in the two solutions.

Troubleshooting

Flow solution diverges during initial iterations

If the residuals show that the solution is diverging during the first several iterations, and there is no sign that they are decreasing, this indicates a problem.

- Make sure that the overall setup is correct. Some things to look for are:
  - Mesh extents for viscous flow
  - Mesh quality
  - Boundary conditions for both flow and turbulence, if turbulent flow is activated
- Try reducing the Courant number.
  - It is possible to make the Courant number as small as desired to get the solution started. If it seems necessary to run at a Courant number significantly less than 1 for more than just a few iterations, there is most likely something wrong with your solution setup.
  - Once the residuals start coming down again, it is possible to increase the Courant number.

Warnings occur of corrections being limited

Because temperature is a positive definite quantity, STAR-CCM+ tries to limit the corrections so that large changes do not tend to drive this quantity negative.

- It is acceptable for these warnings to occur during startup, as long as they cease after several iterations. If they do not, it is necessary to reduce the Courant number.

Modeling Flow Using a Segregated Approach

This section explains the following:

- The segregated flow model
- The segregated flow solver
- How to select the segregated flow model
- Running a transient analysis
- Troubleshooting

The Segregated Flow Model

The Segregated Flow model solves the flow equations (one for each component of velocity, and one for pressure) in a segregated, or uncoupled, manner. The linkage between the momentum and continuity equations is achieved with a predictor-corrector approach.

The complete formulation can be described as using a colocated variable arrangement (as opposed to staggered) and a Rhie-and-Chow-type pressure-velocity coupling combined with a SIMPLE-type algorithm.

This model has its roots in constant-density flows. Although it can handle mildly compressible flows and low Rayleigh number natural convection, it is not suitable for shock-capturing, high Mach number, and high Rayleigh-number applications.

To set the properties for the Segregated Flow model, select Segregated Flow in the Models node under the required continuum node. The properties are available under the Properties and Expert categories.



Properties

Convection — Selects the convection scheme that is used.

- 1st-order — Selects the first-order upwind convection scheme.
- 2nd-order — Selects the second-order upwind convection scheme.
- Central — Selects the central-differencing convection scheme (available with large eddy simulation (LES) only).
- Bounded-Central — Selects the bounded central-differencing convection scheme (available with large eddy simulation (LES) only). Adds a Bounded Differencing sub-node to the Segregated Flow node.
- Hybrid — Selects the hybrid second-order upwind/central-differencing convection scheme (available with detached eddy simulation (DES) only).
- Hybrid-BCD — Selects the hybrid second-order upwind/bounded central-differencing convection scheme (available with detached eddy simulation (DES) only). Adds a Bounded Differencing sub-node to the Segregated Flow node.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties within the *Expert* category. The values in that category reflect both the design of the model and discretization approaches that have been optimized for accuracy and performance.

Minimum Absolute Pressure — Minimum allowable absolute pressure (for compressible flows).

Secondary Gradients — Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

- On — Include both secondary gradients.
- Off — Exclude both secondary gradients.

Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Limiting Acoustic-CFL

The limiting acoustic CFL, applied for unsteady simulations, to scale the mass flux dissipation. Larger values cause more local dissipation and damp more spurious oscillations, but at the expense of accuracy. Limiting Acoustic-CFL can be specified as a constant, field function, table, or user code. Make sure non-constant values vary smoothly in space.

Per-Model	Specifies Limiting Acoustic-CFL in a node under the Segregated Flow model node. This is the default, with a default value of 1.0.
Per-Region	Specifies Limiting Acoustic-CFL in a node under the Physics Values node for each region.

Delta-V Dissipation

Enable/disable an extra "delta-v" dissipation (in addition to Rhie-Chow) that resembles a term in the coupled flow solver. This dissipation comes from the right-most term in Eqn. (214). This enhancement allows a robust treatment of porous media and similar large body forces by making the segregated solver behave more like the coupled solver, and fixing checker-boarding at a porous/non-porous interface. This property is deactivated by default.

On	Activate delta-V dissipation.
Off	Deactivate delta-V dissipation.

Flow Boundary Diffusion

Enable/disable the flow-boundary diffusion fluxes (or viscous fluxes for flow models) that come from Eqn. (2560). This property is activated by default.

Activated	Activate flow-boundary diffusion.
Deactivated	Deactivate flow-boundary diffusion.

Unsteady Flux Dissipation Corrections	Enable/disable mass flux dissipation corrections for unsteady simulations, modifying the coefficients $\bar{a}_0$ and $\bar{a}_1$ in Eqn. (213), based on CFL.	
	Activated	Activate dissipation corrections.
	Deactivated	Deactivate dissipation corrections, the default.

Bounded Differencing

Bounded differencing lets you specify parameters for the bounded central-differencing convection scheme when you work with large eddy simulation (LES), and for the hybrid second-order upwind/bounded central-differencing convection scheme when you work with detached eddy simulation (DES). The Bounded Differencing node, which has its own properties, is activated when the Convection property of the Segregated Flow model is set to `Bounded-Central` (with LES) or `Hybrid-BCD` (with DES).



Properties

Upwind blending factor	Specifies the proportion of upwind differencing, ς_{ubf} related to Eqn. (2547). The default maximum of this property, `0.15`, provides the most robustness for the scheme. Reducing it would, in principle, increase accuracy. However, unless you are thoroughly familiar with the theoretical aspects of bounded differencing, do not change this property. The default value reflects optimization for accuracy and performance.

Selecting the Segregated Flow Model

Before you can select the Segregated Flow model, first select a Gas or a Liquid Material model.

- In the *Physics Model Selection* dialog, in the *Flow* box, select the Segregated Flow radio button.



The *Physics Model Selection* dialog, with the Segregated Flow model selected, is shown below:



Segregated Flow Solver

The Segregated Flow solver controls the solution update for the Segregated Flow model according to the SIMPLE algorithm. It controls two additional solvers:

- Velocity solver
- Pressure solver

For each of these solvers, you can set under-relaxation factors and algebraic multigrid parameters.

Segregated Flow Solver Properties

Freeze Flow	Disables the flow solver solution update.	
	Activated	The flow solver is skipped so that the flow field solution is not updated.
	Deactivated	The flow solver is invoked as usual and the flow field solution updated accordingly. (The default.)

Segregated Flow Solver Expert Properties

Unless you are thoroughly familiar with the discretization techniques that are used in STAR-CCM+, do not change any properties within the *Expert* category. The values in that category have been optimized for accuracy and performance.

The following expert properties only apply to steady cases.

Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.
	Deactivated	Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

Activated — Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.

Deactivated — Do not reset reconstruction gradients. They are recomputed at the next iteration if they had been zeroed.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

Activated — Retain temporary storage.

Deactivated — Do not retain temporary storage.

Continuity Initialization

Allows the solver to initialize the flow field by optimizing the initial conditions for pressure and velocity through the iterative solution of the continuity equation. The improvement of the initial conditions due to initialization helps to accelerate the convergence of the solution. Adds a node on which to set Continuity Initialization properties.

Activated — The solver uses the optimized values given by the initialization as initial conditions to solve the flow equations according to The Segregated Flow Model.

Deactivated — The solver uses the initial conditions you have specified to solve the flow equations according to The Segregated Flow Model.

For information about other properties within the *Expert* category, see Expert Properties.

Continuity Initialization

The Continuity Initialization node appears when the Continuity Initialization property is activated.



Properties

Continuity Initialization Iterations	Specifies the number of iterations performed during the initialization process. The default value of 3 can be increased for further optimization of the initial conditions.
Continuity Initialization Tolerance	The convergence tolerance during the flow field initialization. Its value overwrites the convergence tolerance for the Pressure Solver during the initialization cycle. When the initialization is complete, the convergence tolerance of the pressure AMG Linear solver is set back to the value you have specified. The default value for the Continuity Initialization Tolerance is 1E-6.

Velocity Solver

The velocity solver controls the under-relaxation factor and algebraic multigrid parameters for the momentum equations. More specifically, it solves the discretized momentum equation to obtain the intermediate velocity field.

Properties

Under-Relaxation Factor	In order to promote convergence, this property is used to under-relax changes of the solution during the iterative process. For steady flows, the default value is 0.7; for unsteady flows, the default value is 0.8. The under-relaxation is implemented in an implicit manner. For more information, see Implicit Iteration.

Expert Properties

Dynamic Local Under-Relaxation	For unsteady flows, the time-derivative term in the momentum equation acts as a form of under-relaxation for the system of equations. Because of this, the user-defined under-relaxation factor can cause a higher level of under-relaxation than the solution requires. Dynamic Local Under-Relaxation causes the solver to compare the relaxation provided by the time derivative term of the momentum equation with the user-defined under-relaxation. The solver then modifies the actual under-relaxation applied based on the result of the comparison.

Activated	Solver compares the two sources of under-relaxation and modifies the final under-relaxation accordingly.
Deactivated	Solver applies the user-defined under-relaxation factor without modifying it. (The default.)

Pressure Solver

The pressure solver controls the under-relaxation factor and algebraic multigrid parameters for the pressure correction equation. More specifically, it solves the discrete equation for pressure correction, and updates the pressure field.

Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. This quantity is ω in Eqn. (238). For steady flows, the default value is 0.3; for unsteady flows, the default value is 0.2.

Expert Properties

Pressure Reference Location	Provides a choice between using the automatic algorithm of STAR-CCM+ for the reference location or providing it manually, as discussed elsewhere.	
	Automatic Selection	STAR-CCM+ sets the reference location at the cell that is next to the boundary face with the smallest X, Y, Z location in the domain.
	User Specified	The Pressure Reference Points child node is added. Use this method to add one or more reference locations depending on the number of non-contiguous regions.

Setting Under-Relaxation Factors

To set the under-relaxation factor for either the pressure solver or the velocity solver, select the node for that solver under the Segregated Flow node in the Solvers node as shown below. The Under-Relaxation Factor property can be entered in the *Properties* window.



The following guidelines are helpful for setting the under-relaxation parameters for the velocity solver and the pressure solver:

Steady State Computations

- The default pressure and velocity under-relaxation factors are conservative. They lead to a converged solution in most cases, including when the grid is poor or the equations are stiff.
- Optimal values of under-relaxation parameters are problem-dependent. The grid aspect also affects the optimum values. In a batch job or any other computation where the progress cannot be monitored, it is safer to use the default values.
- When performing similar calculations, check how the simulation reacts to the variation of under-relaxation factors. Since convergence behavior is similar on different grid densities, it is worthwhile to perform a test on a coarse grid to determine the optimum values for the case. These optimum values can then be applied to a calculation on a denser grid.

- It is an accepted rule of thumb that optimum under-relaxation factors for velocity and pressure add up to 1 for steady-state simulations. If the grid is highly non-orthogonal, keep the under-relaxation factor for pressure lower than this rule suggests. (0.1 is a typical value.)
- If the residuals are converging well, it is acceptable to try increasing the under-relaxation factor for velocity (up to 0.9), and decreasing the pressure under-relaxation factor (to 0.1).
- For compressible flows, you must sometimes decrease the under-relaxation factor for velocity to 0.5 and possibly the pressure under-relaxation factor to 0.1.
- For porous media, you must sometimes to decrease the under-relaxation factor for velocity to 0.5 and possibly the pressure under-relaxation factor to 0.1.
- When solving steady-state problems on fine meshes, start with a much coarser mesh, and then successively refine the mesh by reducing the base size. For example:
  - Design the desired mesh for the highest cell count affordable.
  - Increase the base size by a factor of 8 and start the computation with this coarse mesh.

    You sometimes need lower under-relaxation values for the coarsest mesh, but iterations are fast and this requirement does not increase the computing effort much.
  - After the convergence criteria are satisfied or maximum specified number of iterations is reached, halve the mesh base size, remesh, and restart the simulation. Repeat until the final mesh size is reached.

  In this way, you obtain faster convergence on finer grids by providing a good initial solution. Usually, one needs only a third or a quarter of the number of iterations that would be needed if starting with an initial guess, like constant or zero values.

  In addition, you obtain solutions on a series of grids which are systematically refined (same design, only base sizes reduced), which allows for an estimate of discretization errors using Richardson extrapolation. For example, if you reduce base size by a factor of 2 and use second-order discretization, discretization errors on the finest mesh are equal to one third of the difference between solutions on the finest and next coarser grid. For base size reduction by a factor of 1.5, the errors amount to about 80% of the difference in the two solutions.

Transient Computations

- For transient simulations, provided a reasonably well-resolved time-step is chosen, you can often increase the under-relaxation factor for velocity as high as 1.0 and for pressure up to 0.9 (for example, for LES-type simulations).
- Under-relaxation factors of 0.9 (velocities and scalars) and 0.5 (pressure) work well in most cases.
- These values are higher than in steady-state cases since the contribution of the transient term to the discretized equations has the same effect as under-relaxation.

Setting Algebraic Multigrid Parameters

To set the algebraic multigrid parameters for either the pressure solver or the velocity solver, open the node for that solver under the Segregated Flow Solver node in the Solvers node. Then also open the node for the AMG Linear Solver as shown below. The algebraic multigrid parameters can be entered using the *Properties* window.



In most situations, it is not necessary to change these parameters.

Running a Transient Analysis

Transient analysis uses the unsteady SIMPLE algorithm. The regular segregated flow solver is used to perform inner iterations within each time step.

Transient analysis requires the following steps:

- Activate the Implicit Unsteady model together with the Segregated Flow model.
- Set the under-relaxation factors.
- Set the physical time-step. The transient phenomenon being simulated generally governs the physical time-step.

  For periodic flow, an order of 50–100 time steps per period is appropriate.

  The convective Courant number is a helpful indication for selecting the time step size: for time-accurate simulations, the convective Courant number should be 1.0 on average in the zone of the interest. This value implies that the fluid moves by about one cell per time step.

  For flows with free surfaces, if a second-order scheme for time integration is used, the Courant number must be less than 0.5 in all cells. If a free surface moves more than half a cell per time step, the High-Resolution Interface-Capturing scheme can lead to overshoots or undershoots, and eventually to divergence.

  When computing wave propagation, use more than twice as many time steps per wave period than cells per wavelength.
- Set the maximum inner iterations stopping criterion.

  Smaller physical time-steps generally mean that the solution is changing less from one time-step to the next, so that fewer inner iterations are required. There is clearly an optimal balance of time-step size, under-relaxation factors, and the number of inner iterations for a given problem and desired transient accuracy.

  If you are performing a time-accurate simulation with adequate time step and under-relaxation parameters, 5 iterations per time step are usually sufficient. More iterations per time step (an order of 10) and lower under-relaxation factors (around 0.8) are sometimes required for tough coupled problems like flow with free surface and motion of a floating body. VOF simulations with melting and solidification also require many inner iterations.

  If the convergence of iterations is slow, it is a sign that the time step is too large and will result in significant temporal discretization errors. It is better to reduce the time step than to do many iterations within a large time step. For example, halving the time step allows the number of inner iterations to be reduced, so that the computing effort is not much higher while the accuracy of solution is increased.

Troubleshooting

Flow solution diverges during initial iterations

If the momentum and continuity residuals show solution divergence during the first several iterations and the residuals show no signs of decreasing, this indicates a problem.

- Make sure that the overall setup is correct. Some things to look for are:
  - Mesh extents for viscous flow
  - Mesh quality
  - Boundary conditions for both flow and turbulence, if turbulent flow is activated
- If the flow is incompressible, a slight reduction of the under-relaxation factors for pressure and velocity can stabilize the solution. It is also appropriate to try setting the model properties to use a first-order convection scheme.
- If the flow is compressible, reduce the under-relaxation factor for energy (say to 0.3). Also, the under-relaxation factor for pressure can be set to a small number (say 0.01), for just a few iterations before resuming iterating with the default value of 0.3.

Warnings occur of pressure corrections being limited for compressible flow

For compressible flow, the absolute pressure must be a positive definite quantity. STAR-CCM+ tries to limit the pressure corrections so that large pressure changes do not drive the absolute pressure negative.

- It is acceptable for these warnings to occur during the startup for compressible flow, as long as they cease after several iterations. If they do not, it is necessary to reduce the under-relaxation factor for pressure. This factor can be set to a small number (say 0.01), for just a few iterations, before resuming iterating with the default value of 0.3.
- Ensure that the under-relaxation factor for energy is set to an appropriate value (say 0.3).

Warnings occur about reverse flow at a flow split boundary

In STAR-CCM+, flow split outlet boundaries are not designed to permit recirculation.

- If warnings occur during initial convergence but cease after several iterations, there is no cause for concern.
- If warnings persist, there are two options:
  - Rebuild the mesh but extend the outlet region in such a way as to minimize the risk of recirculation.

- Change the boundary type to a pressure outlet.

Warnings occur about reverse flow at a pressure outlet boundary

In STAR-CCM+, pressure outlet boundaries are designed to permit recirculation, although a unique solution can not exist resulting in convergence problems.

- Recirculation causes a non-uniform pressure distribution at the boundary.
- If a uniform pressure distribution is desired, rebuild the mesh with the outlet region extended in such a way as to prevent recirculation.
- Convergence problems due to unsteady recirculation are indicated by the number of faces at which reversed flow occurs changing from one iteration to the next. (The number of faces is printed in the warning messages in the output window.) If unsteady recirculation occurs, extend the outlet region in such a way as to prevent it.

Modeling Fluid Energy Using a Segregated Approach

This section describes the Segregated Fluid Energy models and the Energy solver. Instructions are provided for selecting a Segregated Fluid Energy model and setting model properties.

A troubleshooting section is provided and the formulation of the equations and discretization approach is also presented.

There are three Segregated Fluid Energy models, which are companion models to the Segregated Flow model:

- Segregated Fluid Enthalpy
- Segregated Fluid Temperature
- Segregated Fluid Isothermal

The first two solve the total energy equation in a continuum using a segregated formulation, while the third uses a constant setting for temperature.

Segregated Fluid Enthalpy

The Segregated Fluid Enthalpy model solves the total energy equation with chemical thermal enthalpy as the solved variable. Temperature is then computed from enthalpy according to the equation of state.

This model is recommended for any simulation involving combustion. It is required with the non-adiabatic PPDF equilibrium combustion models.

Properties

Convection	For guidance on selecting a convection, see Diffusion Term, p. 6349.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.

Expert Properties

Flow Boundary Diffusion — Activate or deactivate the flow-boundary diffusion fluxes that are given in Eqn. (2560). This option is activated by default. A realistic temperature value must be set on flow boundaries when this option is activated.

- Activated — Activate flow-boundary diffusion.
- Deactivated — Deactivate flow-boundary diffusion.

Secondary Gradients — Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

- On — Include both secondary gradients.
- Off — Exclude both secondary gradients.
- Interior Only — Include the interior secondary gradients only.
- Boundaries Only — Include the boundary secondary gradients only.

Segregated Fluid Temperature

The Segregated Fluid Temperature model solves the total energy equation with temperature as the solved variable. Enthalpy is then computed from temperature according to the equation of state.

This model is appropriate for simulations that do not involve combustion.

Properties

Convection — For guidance on selecting a convection, see Diffusion Term, p. 6349.

- 1st-order — Selects the first-order convection.
- 2nd-order — Selects the second-order convection.

Expert Properties

Flow Boundary Diffusion	Activate or deactivate the flow-boundary diffusion fluxes (or viscous fluxes for flow models) given by Eqn. (2558). This option is activated by default. A realistic temperature value must be set on flow boundaries when this option is activated.	
	Activated	Activate flow-boundary diffusion.
	Deactivated	Deactivate flow-boundary diffusion.
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Segregated Fluid Isothermal

The Segregated Fluid Isothermal model keeps the temperature in the continuum constant. For problems where the temperature variations are small and negligible, it would be computationally expensive to solve an ordinary energy transport equation when it would yield a nearly constant field of temperature. Hence, STAR-CCM+ includes the isothermal energy option to provide a constant temperature field for all models that require temperature.

Properties

Continuum Temperature	Specifies the temperature for the continuum with a text entry.

Selecting a Segregated Fluid Energy Model

To select a Segregated Fluid Energy model, first select the Segregated Flow model. The Segregated Fluid Enthalpy, Segregated Fluid Temperature, and Segregated Fluid Isothermal radio buttons become available for selection in the *Optional Models* box:



The *Physics Model Selection* dialog with one of the Segregated Fluid Energy models activated is shown below:



If a non-adiabatic PPDF combustion model is active, Segregated Fluid Enthalpy is selected automatically provided the Auto-select recommended physics models checkbox is activated. This model is the only appropriate choice for an energy model.

Setting Model Properties

To set the properties for the Segregated Fluid Enthalpy model or the Segregated Fluid Temperature model, select its node in the Models node under the required continuum node. The properties are available under the Properties and Expert categories.



The Convection property switches between first-order and second-order upwind convection schemes.

The Segregated Fluid Isothermal model simply has a constant value to set for its property, Continuum Temperature.

Energy Solver

The Energy solver controls the solution update for the Segregated Fluid Energy model. It is used to set the under-relaxation factor and algebraic multigrid parameters for the energy equation.

Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted with the newly computed solution. For the theoretical background, see Implicit Iteration.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

Setting the Under-Relaxation Factor

To set the under-relaxation factor for energy, select the Energy Solver node in the Solvers node as shown below. The Under-Relaxation Factor property for either a fluid or solid material, depending on what you model, can be entered in the *Properties* window.



Setting Algebraic Multigrid Parameters

To set the algebraic multigrid parameters for the energy solver, open the Energy Solver node in the Solvers node. Then also open the node for the AMG Linear Solver as shown below. The algebraic multigrid parameters can be entered using the *Properties* window.



Troubleshooting

Energy solution diverges during initial iterations

If the energy residuals show solution divergence during the first several iterations and the residuals show no signs of decreasing, this scenario indicates a problem.

- Make sure that the overall setup is correct. Some things to look for are:
  - whether the flow field is diverging; if so, refer to the Segregated Flow model troubleshooting
  - mesh quality
  - boundary conditions for both flow and energy.

- A reduction of the under-relaxation factor for energy can stabilize the solution.

Energy solution converges slowly

The under-relaxation factor can be adjusted to control the convergence rate.

- The default Under-Relaxation Factor value of 0.9 is aggressive. It is not advisable to use a larger value than 0.9 for flow simulations.
- For compressible flow problems, the under-relaxation factor is normally kept in the range 0.1-0.5.

Initializing Flow

The process of solution initialization requires that suitable values are specified for the initial conditions. This section provides suggestions for setting initial conditions in both incompressible flows and compressible flows.

STAR-CCM+ also provides the following initialization tools to simplify the process:

- Continuity Initialization for the Segregated Flow Solver
- Grid Sequencing Initialization for the Coupled Flow Solver

Initializing Incompressible Flow

This section presents some suggestions for initializing velocity and for initializing pressure in incompressible flows.

Initializing Velocity

- If the solution domain has a prevalent flow direction, an approximate velocity can be specified in that direction.
- For ducts of complex shapes that do not have a prevalent flow direction, zero is an acceptable guess.

Initializing Pressure

- If the continuum has regions with pressure boundaries, the pressure can generally be set to the value specified on one of the pressure boundaries.
- If there are no pressure boundaries and the flow is incompressible, the initial pressure is arbitrary.

Initializing Compressible Flow

This section presents some suggestions for initializing velocity and for initializing pressure in compressible flows. In addition, a reference is provided with useful isentropic flow relations.

Initializing Velocity

- For uni-directional ducts and free-stream flows over slender bodies, it is best to initialize the flow to the inlet or free-stream values.
- For bent ducts or bluff bodies, it can help to reduce the velocity significantly in these cases. Furthermore, techniques such as changing boundary conditions or ramping boundary values may be required to aid solution convergence.

Initializing Pressure

- If there are no pressure boundaries, the specified initial pressure must not result in a non-physical absolute pressure.
- For free-stream flows, choose the initial pressure such that it is equal to the free-stream pressure.
- For duct flows, choose the initial pressure such that it is equal to or higher than the outlet pressure. Choosing a value higher than the outlet pressure can inhibit temporary reversed flow occurring at the outlet.

Modeling Passive Scalars

Passive scalars are user-defined variables of arbitrary value, assigned to fluid phases or individual particles. They are passive because the do not affect the physical properties of the simulation. An intuitive way to think of passive scalars is as tracer dye in a fluid, but with numerical values instead of colors, and with no appreciable mass or volume.

More elaborate physical effects can be modeled with passive scalars, by careful definition of the source terms, for example, change of composition through radioactive decay.

You can use more than one passive scalar in a simulation. The following situations are examples where this capability can be useful:

- Choosing the best position for a sensor while designing a part with multiple inlet pipes and an exhaust pipe.

  You can use the passive scalars to check the scalar values at discrete points in the computational domain. Then you can determine which of the scalars sends the strongest signal to the sensor.
- Analyzing the mixing of two fluid streams that have the same properties.

  The fluid is represented as single-phase, but you can use multiple passive scalars to examine the effect of mixing.
- Tracing how smoke or any other vapor would convect and diffuse in a room or any fluid domain.

  More specifically, you can model the dissolution of a gas in a liquid.

The Passive Scalar Model

The passive scalar models solve one or more transport equations for the scalar values. The Passive Scalar model is available in STAR-CCM+ when you use the Segregated Flow model, the Coupled Flow model, or the Harmonic Balance model. It is available for a physics continuum and for individual Eulerian phases.

To use passive scalars with Lagrangian phases, see Working with Passive Scalars under Using the Lagrangian Multiphase Model.

You can set the initial numerical value held in the scalar field, and can specify the required value at boundaries. For wall boundaries, set a zero flux instead. You can also define a passive scalar source term and its derivative at the region level.

Selecting the Passive Scalar Model

To select the Passive Scalar model in a physics continuum:

- Select one of the following:
  - Segregated Flow model
  - Coupled Flow model
  - Harmonic Balance model.

  The Passive Scalar radio button becomes available for selection in the *Optional Models* box.

To select the Passive Scalar model for a phase of multiphase flow, select a Multiphase Flow Model for the physics continuum.

The Passive Scalar radio button becomes available for selection under *Optional Models* for each phase.



The Passive Scalar Node

The Passive Scalar model node has a child node, the Passive Scalars manager, which allows you to organize multiple passive scalar material objects and their properties.

For single phase flows, the node appears as follows:



For Eulerian multiphase flows, the node appears as follows:



Properties

None.

Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Convection	The convection scheme to use. For guidance on selecting a convection scheme, see Convection Term, p. 6343.	
	1st-order	First-order convection scheme.
	2nd-order	Second-order convection scheme.

Managing Passive Scalars

STAR-CCM+ provides for easy management of passive scalar materials. The Passive Scalars node, which has a pop-up menu, is the manager object for passive scalar materials in the simulation tree.

Passive Scalars Manager Pop-Up Menu

Use this pop-up menu to add passive scalar material objects to the simulation.

New	Adds a passive scalar material to the simulation.

A Passive Scalar node represents a passive scalar material in the simulation tree. Its child objects function the same as standard material properties managers. They use the Schmidt number function by default; this is useful for when you want the passive scalar to diffuse at a different rate from the fluid. Alternatively, the Method property of the Molecular Diffusivity node can be set for a field function. If you use turbulence, then an additional material property node is activated: Turbulent Schmidt Number. This can be set for a constant or a field function.

This node has properties and is subject to common menu items.



Passive scalar materials can be renamed or deleted.

Passive Scalar Properties

Number	A read-only property that indicates the number of passive scalar materials preceding the current one. For example, for the first passive scalar material that you create, this property would equal zero. These values remain unchanged even when you delete one or more of these objects.
Maximum clip value	Sets the maximum scalar value used for clipping.
Minimum clip value	Sets the minimum scalar value used for clipping.
Clip	Clipping prevents scalar values from exceeding limits you set. Off — No clipping applied.

Below Min, Above Max	Prevents scalar values from rising above the Maximum clip value, and falling below the Minimum clip value.
Below Min	Prevents scalar values from falling below the Minimum clip value.
Above Max	Prevents scalar values from rising above the Maximum clip value.

Transport

Provides the option to deactivate the convection and/or diffusion terms in Eqn. (164) and Eqn. (165). The source term is always active. The transient term is only active for transient simulations.

Convection and Diffusion	Both diffusion and convection are activated.
Convection only	The diffusion term is deactivated.
Diffusion only	The convection term is deactivated.
No Convection or Diffusion	Both convection and diffusion terms are deactivated.

Use GGDH

Uses the Generalized Gradient Diffusion Hypothesis. A Reynolds Stress Turbulence model must be activated first.

Activated	Use GDDH in calculating passive scalar transport.
Deactivated	Do not use GDDH. The default.

The Passive Scalar Solver

The Passive Scalar Solver node is activated with the Passive Scalar model.



Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the new solution replaces the old one. The default value is 0.8. For the theoretical background, see Eqn. (2592) in Implicit Iteration.

Expert Properties

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

- Activated: Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
- Deactivated: Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

- Activated: Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed.
- Deactivated: Do not reset reconstruction gradients. The gradients are recomputed at the next iteration if they have been zeroed.

Temporary Storage Retained

Retains or deletes temporary storage at the end of the iteration. These quantities become available as field functions during subsequent iterations.

- Activated: Retain temporary storage.
- Deactivated: Do not retain temporary storage.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. In general, to freeze flow or turbulence, activate Temporary Storage Retained for both flow and turbulence for one iteration after freezing any solver.

- Activated: The solver does not update the solution.
- Deactivated: The solver does update the solution.

Example of Using a Passive Scalar for Residence Time

You can use the Passive Scalar model to simulate residence time (the mean age of air) for a single phase. Consider a chamber with air flowing in at one opening and flowing out at another, modeled as a steady-state Eulerian flow. A passive scalar is defined on the air, increasing steadily over time in value. In effect, defining the passive scalar attaches a clock to each volume element of air. These virtual clocks can be used to measure the residence time of the air in different parts of the chamber.

Example procedure:

- Set up all your necessary models, boundary conditions, and initial conditions.
- Activate the Passive Scalar model.
- Create a passive scalar and rename it ResidenceTimeAir.
- For all wall boundaries, set the Physics Conditions > Wall Passive Scalar option to Zero Flux.
- Create a field function and rename it ResidenceTimeAirSource.
- In the definition of the field function, supply this syntax:

  ```
  ($ResidenceTimeAir >1000)?0:$Density
  ```

  The value of `1000` represents a maximum time, which is necessary if the flow field has a vortex or recirculation. Otherwise, time would grow to infinity.

  The `$Density` value is present to scale the source term in the passive scalar transport equation.

- Select the Regions > Fluid > Physics Conditions > Passive Scalar Source Option node and activate Passive Scalar Source.



- Open the Regions > [Region] > Physics Values node and select the Passive Scalar Source node.



- In the Method property, select `Composite`.
- Open the Composite node and select the ResidenceTimeAir node.
- In the Method property of the ResidenceTimeAir node, select `FieldFunction`.

- Select the Field Function node and set its Scalar Function property to `ResidenceTimeAirSource`, which is the field function that you defined.



- Create a scalar scene to display `ResidenceTimeAir`.

The simulation is ready to run.

A similar approach can be used for unsteady simulations.

Setting Passive Scalar at Boundaries

For flow boundaries, you can set the amount of passive scalar crossing the boundary as a simple value. For wall boundaries, a choice is presented between a simple value, zero flux, or a specified flux. This choice is made on

the Wall Passive Scalar Option node located within the Physics Conditions folder for a wall boundary. Select this node to access the properties of this option.



Wall Passive Scalar Option Properties

Wall Passive Scalar	Defines whether you want to enter a scalar value or a flux for this condition.	
	Specified Value	Specify a scalar value term. A Passive Scalar Value node is added to the Physics Values node.
	Specified Flux	Specify a flux. A Passive Scalar Flux node is added to the Physics Values node.
	Zero Flux	Do not specify a value.

Setting Passive Scalar as a Value

Passive scalar values can be set at inflow and outflow boundaries using the Passive Scalar node. This node is added to their Physics Values folders when the Passive Scalar model is selected. It is entered as a scalar array profile. For Eulerian and Lagrangian phases, the node is supplied for each phase containing passive scalars.



This node can also be added to a wall boundary depending on the setting of the Wall Passive Scalar Option condition of that boundary.

Setting Passive Scalar as a Flux

Passive scalar can be set as a flux at a wall boundary using the Passive Scalar Flux node. This node can be added to a wall boundary depending on the setting of the Wall Passive Scalar Option condition of that boundary. It is entered as a scalar array profile.



Note: The passive scalar flux value is the mass flux, not the volumetric flux.

Setting a Passive Scalar Region Source

To define a passive scalar source within a region with a single phase flow:

1. Select the Physics Conditions > Passive Scalar Source Option node for the region and activate Passive Scalar Source.



   Two nodes are added to the Physics Values folder for the region: the Passive Scalar Source and Passive Scalar Source Derivative nodes.

2. Use these nodes to set the value and derivative of the source that you wish to set.

To define a passive scalar source within a region with a multiphase flow:

1. Select the Physics Conditions > Passive Scalar Source Option node for the phase and activate Passive Scalar Source.



Two nodes are added to the Physics Values folder for the phase: the Passive Scalar Source and Passive Scalar Source Derivative nodes.

2. Use these nodes to set the value and derivative of the source that you wish to set.

The source term S_{ϕ_i} in Eqn. (164) and Eqn. (165) is not multiplied internally by the volume fraction. When you set the scalar source for an Eulerian or Lagrangian phase in a multiphase simulation, include the phase volume fraction explicitly in the implementation of the source. For example, a passive scalar source term for air in a two-phase Eulerian multiphase simulation containing air and water is:

```
$VolumeFractionAir*$DensityAir
```

Passive Scalar Source

Use this node to set the source term S_ϕ in Eqn. (164). It is entered as a scalar array profile. Include the phase volume fraction explicitly for an Eulerian or Lagrangian phase.

Passive Scalar Source Derivative

This node represents the derivative of each component of the Passive Scalar Source regarding the corresponding passive scalar component, S_{ϕ_j} given in Eqn. (164). The derivative of the user source term is used to linearize Eqn. (164). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution, assuming that the source is set by a differentiable equation. It is entered as a scalar array profile.

Setting an Initial Value of Passive Scalar

Initial values of passive scalars can be set using the Passive Scalar node that appears in the Initial Conditions folder of physics continua. This value can also be set on a region when custom initial conditions are selected.

Passive Scalar Properties

Value	The value and units of the passive scalar initial condition. Clicking the customizer executes the define value action.

Passive Scalar Field Functions Reference

The following primitive field functions are made available to the simulation when the Passive Scalar model is used. These field functions have a typical set of base properties.

Diffusion Coefficient of Passive Scalar x	Diffusion coefficient of a particular passive scalar material, where x equals the total number of passive scalar materials in the simulation when there are two or more.
Diffusivity of Passive Scalar x	Diffusivity of a particular passive scalar material, where x equals the total number of passive scalar materials in the simulation when there are two or more. This field function is activated only when you select the laminar viscous option. This can use the Schmidt number method or a field function.

Passive Scalar x	The field function object that corresponds to a particular passive scalar material, where x equals the total number of passive scalar materials in the simulation when there are two or more. This object makes the passive scalar available as a field function.
Turbulent Schmidt Number of Passive Scalar x	Turbulent Schmidt Number of a particular passive scalar material, where x equals the total number of passive scalar materials in the simulation when there are two or more. This field function is activated only when you select the turbulent viscous option. This can be a constant or a field function.

Modeling the Viscous Regime

STAR-CCM+ lets you specify the *viscous regime* to use in your simulation. The following types of flows can be modeled:

Inviscid Flow

Inviscid flows are an idealization resulting from neglecting the viscous effects in simulating the equations of motion.

The solution of the resulting Euler equations (as opposed to the Navier-Stokes equations) generally results in significant savings of computer resources. Boundary layers and other viscous effects are not resolved. This approximation is only valid for certain physical situations, such as high-Reynolds number compressible aerodynamics.

Viscous Flow

Viscous flows can be classified as either laminar or turbulent. Laminar and turbulent flows occur in nature. Both types are described using the Navier-Stokes equations which include the effects of viscosity, thermal conductivity, and mass diffusion.

Laminar Flow

The term laminar refers to a well-ordered flow, free of macroscopic, non-repeating fluctuations. Laminar flows occur in nature when the Reynolds number (the ratio of viscous to inertial forces) is low enough that transition to turbulence does not occur.

In computational simulations, numerical instabilities can arise from simulating laminar flows at Reynolds numbers that are too large. If you are seeking a steady solution, these instabilities can impede convergence. Therefore, laminar flow simulation is appropriate if you already know that the Reynolds number of the problem is sufficiently low.

Transitional Flow

Transition is a term that refers to the breakdown of laminar flow, through amplification of infinitesimal disturbances, to turbulence. A transitional flow can be defined as one that encompasses this process.

The occurrence of physical instabilities in laminar simulations cannot be relied on as an accurate indication of transition. Furthermore, when using turbulence models in low-Reynolds number simulations, the onset of turbulence in viscous layers cannot be predicted with any reliability by the turbulence model itself.

Strictly speaking, STAR-CCM+ has no means of predicting transition. However, a Transition model is available that allows you to mimic the effect of transition by suppressing the turbulence in a certain pre-defined region.

Turbulent Flow

A flow that is in a state of continuous instability, exhibiting irregular, small-scale, high-frequency fluctuations in both space and time is termed turbulent.

It is strictly possible to simulate turbulent flow directly by resolving all the scales of the flow (termed direct numerical simulation). However, the computer resources that are required are too large for practical flow simulations. Therefore, a suitable turbulence modeling approach must be selected.

Selecting the Viscous Regime

To select a Viscous Regime in STAR-CCM+:

1. In the *Physics Model Selection* dialog, select one of the following:
   - Coupled Flow model
   - Segregated Flow model.

2. In the *Viscous Regime* box, select the appropriate option by clicking the corresponding radio button.



The *Physics Model Selection* dialog shows the selected Viscous Regime model activated.



Modeling Gravity

The Gravity model accounts for the action of gravitational acceleration in STAR-CCM+ simulations. For fluids, it provides two effects:

- The working pressure becomes the piezometric pressure.
- The body force due to gravity can be included in the momentum equations.

The Gravity model also affects other types of simulation, such as material particle flows under the Lagrangian Multiphase framework.

Working Pressure

When the Gravity model is active, the STAR-CCM+ working pressure becomes the piezometric pressure, in which the reference altitude (which you set) is taken into account. The piezometric pressure is given as:

$$p_{piezo} = p_{static} - \rho_{ref}\mathbf{g}(\mathbf{x}\text{-}\mathbf{x}_0)$$

For both variable and constant density flows it is necessary to set the reference altitude $\mathbf{x}_0$, which is used in the calculation of piezometric pressure as shown above. The reference altitude is the point at which the piezometric and static pressure are considered equal.

When setting initial or boundary pressures in simulations for which the Gravity model is active, it is the piezometric pressure value that is being set. STAR-CCM+ calculates the corresponding static pressure.

Body Force Source Term

The Gravity model permits the inclusion of the body force due to gravity in the momentum equations when using the Segregated Flow model or the Coupled Flow model.

For variable-density flows (for example, when using the Ideal Gas model or the Polynomial Density model), the buoyancy effects due to gravity and density variation are modeled directly and are always present. For problems that use the Constant Density model, the buoyancy effects due to small temperature and, so, small density variations, are only included when the optional Boussinesq Model is selected in addition to the Gravity model. The formulation of these source terms is presented below.

For Solids, the Gravity model permits the inclusion of a body force source term, $\mathbf{b}$ in the solid momentum equation Eqn. (2303). This body force is calculated and applied to each cell of the solid continuum.

Selecting the Gravity Model

To select the Gravity model, first select the Segregated Flow model or select the Coupled Flow model.

The Gravity radio button becomes available for selection under *Optional Models*.



After you click the Gravity radio button, the *Physics Model Selection* dialog shows that the Gravity model has been selected.



What Is the Boussinesq Model?

The Boussinesq Model becomes available for selection when you use a single-component gas or liquid and the gravity model is invoked with an energy model and the Constant Density model. This model provides a buoyancy source term that only applies when there are small variations of density due to temperature variations. The error introduced by the Boussinesq model can be estimated by comparing solutions where density depends on temperature with solutions where the density is constant and the Boussinesq model is used.

For the Boussinesq model, specify a reference temperature and coefficient of bulk expansion within the reference conditions and material properties nodes of the physics continuum, respectively.

Selecting the Boussinesq Model

The Boussinesq Model becomes available after you have selected the Gravity model for a constant density flow.



Setting the Gravity Model Parameters

This section deals with the parameters that are required for solving flows in the presence of buoyancy forces. Additional inputs are required for:

- Gravity ($\mathbf{g}$ in Eqn. (256) and Eqn. (259))
- Reference altitude ($\mathbf{x}_0$ in Eqn. (257) and Eqn. (258)
- Reference density for variable density flows (ρ_{ref} in Eqn. (256))
- Reference temperature for constant density flows (T_{ref} in Eqn. (259))
- The coefficient of bulk expansion for use with the Boussinesq model (β in Eqn. (259)).

Setting Gravity Vector

To specify the gravity vector ($\mathbf{g}$ in Eqn. (256) and Eqn. (259)), open the Reference Values node in the relevant continuum node, then select Gravity. A constant vector value can be entered in the *Properties* window. You can also enter field function syntax directly.



Setting Reference Altitude

To specify the reference altitude (h_{ref} in Eqn. (257) and Eqn. (258)), open the Reference Values node in the required continuum node, then select Reference Altitude. A vector can be entered in the *Properties* window.



Although its exact position is arbitrary (similar to the values of reference pressure and reference density), it is reasonable to locate the reference altitude at a central point within the solution domain. This means that the expected range of values for working pressure is balanced around zero, and so provide an optimum condition for the STAR-CCM+ numerical schemes.

The working pressure at pressure boundaries has to be consistent with the selected reference altitude, reference density and reference pressure. If any of these values are changed, it is necessary to ensure that the working pressures set on pressure boundaries are still consistent.

Setting Reference Density for Variable Density Flows

To specify the reference density (ρ_{ref} in Eqn. (256)), open the Reference Values node in the required continuum node, then select Reference Density. A constant value can be entered in the *Properties* window.



Setting Reference Temperature for Constant Density Flows

To specify the reference temperature (T_{ref} in Eqn. (259)), open the Reference Values node in the required continuum node, then select Reference Temperature. A constant value can be entered in the *Properties* window.



Setting the Coefficient of Bulk Thermal Expansion for the Boussinesq Model

The coefficient of bulk thermal expansion (β in Eqn. (259)) is a material property that can be set with one of the following techniques:

- Using constant values
- Using field function values

This property is accessed through the Thermal Expansion node under the specific Gas material node for the required continuum.



Specific Flow Phenomena

Modeling Axisymmetric Swirl

Modeling Porous Media

Modeling Fans

Virtual Disk Model

Confining Vorticity

Modeling Axisymmetric Swirl

The Axisymmetric Swirl model models swirling or rotating flow about a central axis. You can use this model when your simulation includes a swirling flow from an inlet such as a fan interface.

The Axisymmetric Swirl model accounts for the rotational motion in the momentum equations when using the Segregated Flow model or the Coupled Flow model.

For more information, refer to the following:

- Selecting the Axisymmetric Swirl Model
- Using the Axisymmetric Swirl Model
- Axisymmetric Swirl Model Formulation

Selecting the Axisymmetric Swirl Model

The Axisymmetric Swirl model is available only for gas or liquid axisymmetric continua. This model can be used with the Segregated Flow model or the Coupled Flow model.

To select the Axisymmetric Swirl model in the *Physics Model Selection* dialog:

- In the *Space* box, select Axisymmetric.
- In the *Material* box, select Gas or Liquid.
- In the *Flow* box, select Coupled Flow or Segregated Flow.

The Axisymmetric Swirl model is available for selection under *Optional Models*.



After you select the Axisymmetric Swirl model, you must also select an *Equation of State* model, a *Time* model, and a *Viscous Regime* model.

When you close the *Physics Model Selection* dialog, the Axisymmetric Swirl model is shown in the list under the Continua > Physics > Models node. The model has no user-defined properties.

Using the Axisymmetric Swirl Model

In an axisymmetric space, all three velocity components are functions of only the axial and radial coordinates, z and r. You define the swirling flow in your simulation by specifying the appropriate z and r components of the velocity field.

The fluid rotational motion is typically from a swirling inlet such as a fan interface. You model this swirling flow by setting the appropriate boundary conditions at the relevant inlets and walls.

To specify the swirl in your simulation, do the following for each inlet that contributes to the swirling flow:

- If necessary, create the field functions, tables, or user code that you want to use to define the velocity field at the inlet boundary.
- Under Regions > [Region Name] > Boundaries > Inlet > Physics Conditions, select the Velocity Specification node.
- In the *Properties* window, in the Method property drop-down menu, select Components.
- Under Regions > [Region Name] > Boundaries > Inlet > Physics Values, select the Velocity node.
- In the *Properties* window, in the Method property drop-down menu, select the appropriate method.

  The velocity is set using a vector profile. You can define the swirling flow by using field functions, a table, or user code.
- Under the Velocity node, select the sub-node and then, in the *Properties* window, specify the appropriate function or table.

At high rotation velocities, the strong coupling between the momentum equations can make the flow difficult to solve. The high rotation velocity creates a large radial pressure gradient which, in turn, affects the swirl in the fluid. You may find it necessary to first solve the flow for a low rotational velocity, and then gradually increase the swirl velocity until you reach the appropriate values.

Modeling Axisymmetric Swirl

The Axisymmetric Swirl model models swirling or rotating flow about a central axis. You can use this model when your simulation includes a swirling flow from an inlet such as a fan interface.

The Axisymmetric Swirl model accounts for the rotational motion in the momentum equations when using the Segregated Flow model or the Coupled Flow model.

For more information, refer to the following:

- Selecting the Axisymmetric Swirl Model
- Using the Axisymmetric Swirl Model
- Axisymmetric Swirl Model Formulation

Selecting the Axisymmetric Swirl Model

The Axisymmetric Swirl model is available only for gas or liquid axisymmetric continua. This model can be used with the Segregated Flow model or the Coupled Flow model.

To select the Axisymmetric Swirl model in the *Physics Model Selection* dialog:

- In the *Space* box, select Axisymmetric.
- In the *Material* box, select Gas or Liquid.
- In the *Flow* box, select Coupled Flow or Segregated Flow.

The Axisymmetric Swirl model is available for selection under *Optional Models*.



After you select the Axisymmetric Swirl model, you must also select an *Equation of State* model, a *Time* model, and a *Viscous Regime* model.

When you close the *Physics Model Selection* dialog, the Axisymmetric Swirl model is shown in the list under the Continua > Physics > Models node. The model has no user-defined properties.

Using the Axisymmetric Swirl Model

In an axisymmetric space, all three velocity components are functions of only the axial and radial coordinates, z and r. You define the swirling flow in your simulation by specifying the appropriate z and r components of the velocity field.

The fluid rotational motion is typically from a swirling inlet such as a fan interface. You model this swirling flow by setting the appropriate boundary conditions at the relevant inlets and walls.

To specify the swirl in your simulation, do the following for each inlet that contributes to the swirling flow:

- If necessary, create the field functions, tables, or user code that you want to use to define the velocity field at the inlet boundary.
- Under Regions > [Region Name] > Boundaries > Inlet > Physics Conditions, select the Velocity Specification node.
- In the *Properties* window, in the Method property drop-down menu, select Components.
- Under Regions > [Region Name] > Boundaries > Inlet > Physics Values, select the Velocity node.
- In the *Properties* window, in the Method property drop-down menu, select the appropriate method.

  The velocity is set using a vector profile. You can define the swirling flow by using field functions, a table, or user code.
- Under the Velocity node, select the sub-node and then, in the *Properties* window, specify the appropriate function or table.

At high rotation velocities, the strong coupling between the momentum equations can make the flow difficult to solve. The high rotation velocity creates a large radial pressure gradient which, in turn, affects the swirl in the fluid. You may find it necessary to first solve the flow for a low rotational velocity, and then gradually increase the swirl velocity until you reach the appropriate values.

Modeling Porous Media

This section provides the following information about porous media:

- A definition of what is meant by a porous medium in STAR-CCM+
- How to activate a porous region
- How to set the necessary region values of porosity, inertial resistance and viscous resistance
- The treatment of turbulence and energy in porous regions
- The influence of porous resistance on temperature for incompressible flow
- Using porous media to model flow in tubes
- An example of implementing the Ergun equation with a dependence on density

The section Accelerating the Linear Solver provides suggestions on solver settings that can help with convergence of porous cases.

What Is a Porous Medium?

A porous medium is a material whose internal structure consists of solid elements permeated by fine-scale voids that permit the passage of fluids. In many CFD simulations involving this kind of material, it is not the details of the internal flow that are of interest. The macroscopic effect of the porous medium on the overall fluid flow is of interest. In this case, the effect of the porous medium on the flow is defined using lumped parameters. These parameters are typically taken to be resistance coefficients for a source term in the momentum equation. You set these parameters to reflect the effect of the porous media on the flow downstream of the porous section. The flow within the porous section is not relevant as the detailed geometric structure is not represented.

Examples of porous media include filters, packed-bed chemical reactors, honeycomb structures, sponges, and fibrous materials. The methodology can also be applied to other structures such as heat exchangers containing large numbers of tubes or plates. The simulation is assumed to consider a more substantial flow domain of which the heat exchanger is only a part.

In STAR-CCM+, a porous medium is represented at the region level by setting the region type to Porous. Additional conditions and values are added to a porous region to specify the coefficients of flow resistance. These coefficients are used to calculate the porous source term in the momentum equation, as detailed in the formulation. Porous regions are also assigned the porosity value, which is an additional input that is required for

unsteady flow calculations. The porosity value defines the ratio of open volume to total volume within the porous medium. This value is used in the time derivative terms of the continuity, energy, and species equations.

Activating Porous Regions

To activate a porous region, select the region node under the Regions node. Select `Porous Region` from the drop-down list corresponding to the Type property in the *Properties* window. This property allows you to specify the porous inertial resistance and porous viscous resistance.



Setting Porous Region Conditions and Values

The main values that you can set for a porous region are its porosity (for unsteady flows) and the inertial and viscous coefficients that are required for the porous source term in the momentum equation. It is also necessary to specify turbulence quantities directly when a turbulence model is active. For more information, see the section, Working with turbulence in porous media. The section, Working with energy in porous media, presents the quantities available to use in setting thermal quantities for the solid material.

An option is also available to control the use of porous media flux.

Specifying Porosity

This node is an additional input χ required for unsteady flow calculations, and for calculation of the effective thermal conductivity in the porous region. It is defined as the ratio of open volume to total volume of the porous medium, and is entered as a scalar profile.

For unsteady flow calculations, the porosity is used to replace the fluid density ρ in the time derivative terms of the continuity, energy, and species equations by the product $\rho\chi$. This substitution reflects the fact that the fluid occupies only part of the volume. For effective thermal conductivity, the porosity is used to blend the thermal conductivities of the fluid and solid materials that are represented within the porous region, as shown in Eqn. (250).

To specify the porosity of a region:

1. Open the node corresponding to the porous region and then select Physics Values > Porosity.



2. Set Method to the appropriate value.

Specifying Porous Resistance for Single Phase Fluids

The porous inertial resistance tensor ($\mathbf{P}_i$ in Eqn. (263)) and the porous viscous resistance ($\mathbf{P}_v$ in Eqn. (263)) are defined as tensor profiles.

For single phase fluids, you specify this property for the region. For multiphase fluids, you specify this property for each phase within the region. Examples of specifying this property are provided in a later section.

To specify the porous inertial resistance or the porous viscous resistance for a single phase fluid:

1. Open the node corresponding to the porous region and then open the Physics Values node.
2. Select the appropriate node: Porous Inertial Resistance or Porous Viscous Resistance.

   The example shown below is the Porous Inertial Resistance node.



3. Set Method to the appropriate value.

Specifying Porous Resistance for Multiphase Fluids

The porous inertial resistance tensor ($\mathbf{P}_i$ in Eqn. (263)) and the porous viscous resistance ($\mathbf{P}_v$ in Eqn. (263)) are defined as tensor profiles.

For multiphase fluids, as each phase may have a different porous resistance, you specify this property for each phase within the region. For single phase fluids, you specify this property for the region. Examples of specifying this property are provided in a later section.

To specify the porous inertial resistance or the porous viscous resistance for a multiphase fluid:

1. Open the node corresponding to the porous region and then open the Phase Conditions > [Phase] > Physics Values node for the appropriate phase.

2. Select the appropriate node: Porous Inertial Resistance or Porous Viscous Resistance.

   The example shown below is the Porous Inertial Resistance node.



3. Set Method to the appropriate value.

Specifying Porous Media Flux Option

To omit explicitly the convection and viscous fluxes (as in Eqn. (104)), open the node corresponding to the porous region and select the Porous Media Flux Option node in the Physics Conditions node. This option is available only with the Segregated Flow model.



This option is deactivated by default. However, if the checkbox of the Discount Momentum Fluxes property is activated, the fluxes are omitted. As a result, the flow solver is less robust, since approximations are made in the linearization of the discretized equations. Therefore it is recommended that the default setting is used for truly porous media.

Working with Turbulence in Porous Media

The effect of a porous region on turbulent flow depends on the internal structure of the porous medium. Where turbulence is present, the turbulence scales are determined from the geometric structure of the porous medium. As it is not possible for STAR-CCM+ to predict the turbulence scales directly, you specify the appropriate values on the porous region. Turbulence quantities in fluid leaving the porous region are constructed from the user-defined values; they are not transported from the upstream side of the porous region.

A porous region in a turbulent flow has a Turbulence Specification node in its Physics Conditions manager. Depending on the particular turbulence model that is selected, this node contains a Method property that provides several options for how turbulence scales within the porous region can be specified. Value nodes corresponding to the particular method selected are added to the Physics Values manager of the porous region. The appropriate turbulent quantities can be specified on these nodes.

Setting Turbulence Scales on a Porous Region



Working with Energy in Porous Media

Energy transport across porous regions is dealt with by specifying the thermal properties of the solid represented by the porous media in addition to the thermal properties of the fluid passing through. The additional values that are provided in the Physics Values folder of a porous region (when an appropriate fluid energy model is active) are:

- Solid Density
- Solid Specific Heat
- Solid Thermal Conductivity

Examples of Setting Porous Resistance Properties

Porous resistance properties can be specified in several different ways depending on the particular media being modeled and the expected flow conditions through the media. By making it possible to specify porous resistances using field functions and expressions (as well as constant values), STAR-CCM+ provides a framework whereby you can define your own resistance models. This section provides some examples of how to do so:

- Orthotropic viscous resistance
- Isotropic resistance using the Ergun equation (field function example)

Orthotropic Viscous Resistance

Consider a block of honeycomb material, much like a bunch of drinking straws that are glued together. This porous medium is situated in a section of the duct that is oriented at 45 degrees to the x-axis of the Laboratory Coordinate System as shown below.



Measurements of the honeycomb material in a wind tunnel (air is the working fluid) have shown that it results in a pressure drop of 70 Pa per meter of the material when the velocity is 1 m/s. In addition, for the range of velocities of interest, the pressure drop varies linearly with velocity.

Since the porous resistance is a linear function of flow velocity, only the viscous resistance tensor $\mathbf{P}_v$ is specified. In the principal coordinates, the primary coefficient is 70. To restrict the flow in the directions perpendicular to the primary direction, choose a resistance coefficient that is 2-3 orders of magnitude larger than primary coefficient, for example 10000. Choosing a larger resistance does not affect the flow but can adversely affect the convergence characteristics. Therefore the viscous resistance tensor in principal coordinates is chosen to be:

$$\begin{bmatrix} 70 & 0 & 0 \\ 0 & 10000 & 0 \\ 0 & 0 & 10000 \end{bmatrix}$$

The ways that this tensor can be specified in STAR-CCM+ include:

- using the principal tensor profile method.

  The XX, YY and ZZ components are 70, 10000 and 10000 respectively. The XX Axis is the vector (1, 1, 0) relative to the Laboratory Coordinate System, representing a vector that is tilted at 45 degrees to the Laboratory x-axis. The YY Axis is then (-1, 1, 0), also relative to the Laboratory Coordinate System.

- using the composite symmetric tensor method.

  This method is simplified by creating a local Cartesian Coordinate System with (1, 1, 0) as the X Axis Direction corresponding to x′ and (-1, 1, 0) as the Y Axis Direction corresponding to y′. Using this coordinate system, the tensor components are all zero except for the XX, YY, and ZZ components, which are 70, 10000, and 10000 respectively.

Isotropic Resistance Using the Ergun Equation (Forchheimer for Packed Beds)

One of the classic equations describing flow through a porous medium is Darcy's law, which relates the flow velocity to the pressure gradient based on a measure of permeability. This law, which applies to creeping flow, is expressed using [32] as:

$$-\nabla p = \frac{\mu}{k_p} v \tag{96}$$

where μ is the fluid molecular viscosity ($Pa\,s$), k_p is the permeability (considered to be an intrinsic property of the porous medium, m^2) and v is the superficial velocity through the medium ($m\,s^{-1}$).

As the flow velocity increases, the relationship between velocity and pressure gradient becomes nonlinear. Dupuit and Forchheimer (as reported by [32]) proposed the addition of a quadratic term as follows:

$$-\nabla p = \frac{\mu}{k_p} v + \beta \rho v^2 \tag{97}$$

Eqn. (97) is commonly known as the Forchheimer equation [32]. The beta factor β must be deduced depending on the particular flow and medium of interest, and is commonly determined through experiment.

One example of the Forchheimer equation for a particular class of flow is the Ergun equation (as discussed in [17]). This equation is an empirical model for the pressure drop, dp, over a length, L, of fluid flowing through a packed bed:

$$-\frac{dp}{L} = \frac{150\mu(1-\chi)^2 v}{\chi^3 D_p^2} + \frac{1.75\rho(1-\chi)v^2}{\chi^3 D_p} \tag{98}$$

where ρ is the fluid density ($kg\,m^{-3}$), D_p is the mean diameter of particles in the porous medium (m) and χ is the volume porosity (dimensionless). Comparing Eqn. (98) with Eqn. (97) shows that the permeability and beta factor are:

$$\frac{1}{k_p} = \frac{150(1-\chi)^2}{\chi^3 D_p^2}$$

and

$$\beta = \frac{1.75(1-\chi)}{\chi^3 D_p}$$

Comparing Eqn. (97) with Eqn. (263) gives the final viscous and inertial terms as:

$$P_v = \frac{150\mu(1-\chi)^2}{\chi^3 D_p^2} \tag{99}$$

and

$$P_i = \frac{1.75\rho(1-\chi)}{\chi^3 D_p} \tag{100}$$

To implement this model in STAR-CCM+, it is necessary to create user field functions for each of these terms. Field functions are first defined for the volume porosity and mean particle diameter (which are set to constant values in this example). These values are then used in the field functions for the two terms.

Parameter Field Functions

The volume porosity is defined as a scalar field function with no dimensions, and whose Definition is set to the required value (0.39 in this example).



The particle diameter (0.01 m in this example) is defined in a similar way, except that the Dimensions property is set to `Length`.



Viscous Term Field Function

The scalar field function for the viscous term is dimensioned as $Mass \cdot Length^{-3} \cdot Time^{-1}$, and the full Definition is:

```
150.0*$DynamicViscosity*pow((1-$BedPorosity),2)/
(pow($BedPorosity,3)*pow($BedSphereDiameter,2))
```

The function uses the Bed Porosity and Bed Sphere Diameter functions that were defined previously.



This field function is applied to the Porous Viscous Resistance node as an isotropic tensor.

- Select the Physics Values > Porous Viscous Resistance node on the porous region.



- In the *Properties* window, set the Method to Isotropic Tensor using the drop-down menu. The Isotropic Tensor child node is shown underneath the Porous Viscous Resistance node.

- Select the Isotropic Tensor > Isotropic Component node and, in the *Properties* window, set the Method to Field Function.



- Finally, select the Isotropic Component > Field Function node and choose Ergun Viscous Term from the Scalar Function drop-down menu in the *Properties* window.



Inertial Term Field Function

The scalar field function for the inertial term is dimensioned as $Mass \cdot Length^{-4}$, and the Definition is:

```
(1.75*$Density*(1-$BedPorosity))/
(pow($BedPorosity,3)*$BedSphereDiameter)
```


Servers
ergun
simulation
scene/plot
Tools
Annotations
Colormaps
Coordinate Systems
Data Set Functions
Field Functions
Bed Porosity
Bed Sphere Diameter
Ergun Inertial Term
Ergun Viscous Term
Absolute Pressure
Absolute Total Pressure
Ergun Inertial Term - Properties
Properties
Type
Scalar
Dimensions
Mass/Length^4
Function Name
ErgunInertialTerm
Definition
(1.75*$Density*(1-$BedPoros...
Expert
Ignore Boundary Values
Assembly code
(/ (* (* 1.75 ${Density}) (- 1 $...

The fluid density is included in this field function. Any variations in the fluid density therefore affect the inertial resistance. This function is applied to the Porous Inertial Resistance using the isotropic tensor method as before:



Region Porosity

For an unsteady calculation, the volume porosity χ can be set on the region using the Bed Porosity field function. To set the volume porosity:

- Select the Physics Values > Porosity node of the porous region (PackedBed in this example).
- In the *Properties* window, set the Method property to Field Function using the drop-down menu.

- Select the Porosity > Field Function node and, in the *Properties* window, set the Scalar Function property to Bed Porosity.



Influence of Porous Resistance on Temperature for Incompressible Flow

When simulating incompressible flow through a porous medium with the energy model activated, a change in temperature proportional to the user-specified porous resistance takes place. This section explains the reason for this temperature change. Consider a one-dimensional steady incompressible flow of a fluid with constant c_v (specific heat at constant volume) through a porous medium with a constant cross-section. Mass conservation requires that the velocity u is constant. Momentum conservation reduces to a balance between the pressure gradient dp/dx and the porous resistance R:

$$R = \frac{dp}{dx} \tag{101}$$

The energy equation reduces to:

$$\frac{dH}{dx} = 0 \tag{102}$$

and since total enthalpy is $H = c_v T + p/\rho + u^2/2$, the equation becomes:

$$\frac{dT}{dx} = \frac{-1}{\rho c_v}\frac{dp}{dx} = \frac{-R}{\rho c_v} \tag{103}$$

Hence, it is normal to expect a temperature increase or decrease for an incompressible flow through a porous medium.

Using Porous Media to Model Flow in Tubes

In most porous media the additional source terms in the momentum equation, which are used to create the streamwise pressure gradient, dominate the convection and viscous fluxes. For this reason, the convection and viscous fluxes are not explicitly neglected in STAR-CCM+. In some situations, you can model the flow in a tube without properly resolving the flow. In this case, if the porous media tensor coefficients are used to specify the pressure drop, they are unlikely to dominate the convection and viscous fluxes. The resulting pressure drop does not precisely balance the specified body force. By explicitly neglecting the convection and viscous fluxes, the momentum equation reduces to:

$$\frac{dp}{dx} = -(P_i v^2 + P_v v) \tag{104}$$

where P_i and P_v are the equivalent scalar coefficients of the inertial and viscous porous resistance tensors, respectively.

This relation is often erroneously termed “Darcy’s Law” which in fact expresses a linear relationship between the flow rate and the pressure drop through porous media such as groundwater flow. Similar linear relations exist for fully developed, incompressible, laminar pipe flow, where the pressure drop per unit length can be expressed as a function of the mean velocity u, dynamic viscosity μ, and pipe radius r_0:

$$\frac{\Delta P}{\Delta L} = -\frac{8\mu u}{r_0^2} \tag{105}$$

($u = Q/A$, where Q is the volumetric flow rate and A is the cross-sectional area of the pipe.)

In this type of flow, the viscous fluxes dominate. These fluxes could be neglected by using a slip-wall boundary condition. The convective fluxes cancel each other out, unless there are significant entrance effects, in which case the pipe-flow relation can become non-linear in their presence. The convective fluxes also become more significant in the presence of varying fluid properties, such as in compressible flow.

To omit explicitly the convection and viscous fluxes in porous media, the Porous Media Flux Option is available under the Physics Conditions node of porous regions. This option is deactivated by default. When it is activated, the fluxes are omitted. The flow solver is less robust with this option activated, since approximations are made in the linearization of the discretized equations. Therefore it is recommended that the default setting is used for truly porous media.

Modeling Fans

These methods are intended for axial fans where classical fan laws are applicable. Some applications include radiators and chillers, ventilation, and electronics cooling.

STAR-CCM+ provides the following methods for modeling fans:

- The Fan Interface

  The fan interface is a simple model that represents the fan by a zero-thickness interface. This model adds momentum sources, including swirl, to the downstream flow. The model is an implicit implementation that provides quicker convergence than the fan momentum source.

- The Fan Momentum Source

  The fan momentum source is a more complex model than the fan interface. This model can add swirl to the downstream flow, and is applied to a region that represents the three-dimensional volume of the fan. The fan momentum source model uses an explicit implementation, and as such, is slower to converge.

The fan interface and the fan momentum source are simple models that trade accuracy for speed. To obtain more accurate results, use more complex methods that require more time to obtain results.

For more information on complex fan modeling methods, see Advanced Methods for Modeling Fans.

For the theory that STAR-CCM+ uses in computing the fan interface and fan momentum sources, refer to Fan Interface Formulation and Fan Momentum Source Formulation.

If you are interested in modeling the acoustic effects of a fan, refer to recommended practices for modeling acoustic sources in axial fans.

The Fan Interface Model

The fan interface is a simple model that provides quicker convergence than the fan momentum source. It models the pressure rise through an axial fan as a function of local flow rates or velocity. The fan interface models the fan as a circular interface of zero thickness, although there is no restriction on the topology that you use. This interface type is used only between regions belonging to the same fluid continuum.

See Using the Fan Interface, Creating a Fan Interface, and Fan Interface Properties.

When you set up a fan interface, you supply the fan performance curve that corresponds to the physical fan that you are modeling. The fan performance curve specifies the pressure drop across the fan interface as a function of one of the following variables:

- Local velocity
- Mass average velocity
- Volumetric flow rate
- Mass flow rate.

STAR-CCM+ adds momentum sources such that the operating point of the fan lies on the specified fan curve. The momentum sources can include a swirl component, as would occur downstream of a real fan. However, the fan interface converges much quicker than the fan momentum source.

Fan Momentum Source Model

The fan momentum source employs an "actuator disk" type of methodology. With this methodology, the detailed geometry of the fan blades is not modeled. The circular fan disk is meshed as a separate region, and STAR-CCM+ identifies the cells that are in the volume swept out by the fan blades. The appropriate momentum sources, including the swirl component, are applied to these cells. Analysis is steady, so blade passing effects are not included.

The fan momentum source model can simulate the most commonly designed axial fans. The model works well for adverse pressure gradients, where the fan is pushing against a back pressure. You can have more than one fan in the same simulation.

Since the momentum sources are based on theory, there is no guarantee that the resulting flow rate and pressure rise lies in the fan curve of a particular fan. You specify the appropriate fan performance curve. STAR-CCM+ automatically adjusts the momentum sources so that the operating point of the simulated fan falls on the fan curve.

The fan momentum source model uses an explicit implementation, and as such, can be slow to converge.

Advanced Methods for Modeling Fans

To model fans more accurately, you can use the following methods to model a region that contains the geometry of the fan blades:

- Rotating Reference Frames

  A rotating reference frame is a rotating frame of reference that can be applied to regions to generate a constant grid flux.

  This approach gives a solution that represents the time-averaged behavior of the flow, rather than the time-accurate behavior. For more information, see Using Moving Reference Frames in Steady-State.

- Unsteady Sliding Mesh

  When the time-accurate behavior is required, STAR-CCM+ provides the option to move the mesh vertices of a region during a transient analysis. This mesh motion is specified using either a Rotation motion or a Rotation and Translation motion.

  Applying either of these motions results in mesh vertices being displaced in each time step. For more information, see Using Rigid Motion.

Both the moving reference frame and unsteady sliding mesh approaches to fan modeling require that you create a separate region for the moving fan blades. The joins of the regions are interfaces. Conformal meshes are preferred but are not required.

You can use either trimmed cell or polyhedral meshes:

- For polyhedral meshes, all regions can use the same mesh continuum, and the mesh is conformal across the interface.
- For trimmed cell meshes, each region must have its own mesh continuum, and the mesh is not conformal across the interface.

Overview of Fan Performance Curves

A fan performance curve describes the performance of a fan under specific operating conditions. The fan performance curve typically specifies the pressure rise across the fan as a function of volumetric flow rate, for a standard fluid temperature and fan rotation rate. Fan manufacturers provide this information for each fan.

When you model a fan in STAR-CCM+, you specify the appropriate fan performance curve. To obtain the fan performance curve for a particular fan rotation rate and fluid temperature, you can make your own measurements of the fan under the appropriate conditions. Alternatively, STAR-CCM+ can derive it from the manufacturer-supplied fan performance curve: STAR-CCM+ adjusts the fan performance curve using classical fan laws to

compensate for the different fan rotation rate and fluid temperature.

An example of a manufacturer-supplied fan performance curve and the same fan performance curve that is adjusted for particular operating conditions is shown below.



Fan Pressure Components

The fan pressure components that are used in deriving fan performance curves are defined below.

Fan Total Pressure ($p_{f,T}$)	The increase in total pressure (p_t) across the fan: $p_{f,T} = p_{t,2} - p_{t,1}$ The total pressure is the sum of the dynamic pressure (p_d) and the static pressure (p_s): $p_t = p_d + p_s$
Fan Dynamic Pressure ($p_{f,D}$)	The average dynamic pressure at the fan outlet: $p_{f,D} = p_{d,2} = p_{t,2} - p_{s,2}$
Fan Static Pressure ($p_{f,S}$)	The difference between the fan total pressure and the fan dynamic pressure: $p_{f,S} = p_{f,T} - p_{f,D}$ $= p_{t,2} - p_{t,1} - (p_{t,2} - p_{s,2})$ $= p_{s,2} - p_{t,1}$

Fan Configurations

The three essential fan configurations are shown below:

Ducted Fan

The fan has both inlet and outlet ducts.



The dynamic pressure is non-zero at both the inlet and outlet.

In this configuration, the fan static pressure is:

$$
\begin{aligned}
p_{f,S} &= p_{f,T} - p_{f,D} \\
&= p_{t,2} - p_{t,1} - (p_{t,2} - p_{s,2}) \\
&= p_{s,2} - p_{t,1}
\end{aligned}
$$

Forcing Fan

The fan upstream is a plenum and the downstream is a duct.



The dynamic pressure is zero at the upstream, but non-zero at the downstream.

In this configuration, the fan static pressure is:

$$
\begin{aligned}
p_{f,S} &= p_{f,T} - p_{f,D} \\
&= p_{t,2} - p_{t,1} - (p_{t,2} - p_{s,2}) \\
&= p_{s,2} - p_{t,1} \\
&= p_{s,2}
\end{aligned}
$$

Exhausting Fan

The fan upstream is a duct and the downstream is a plenum.



The dynamic pressure is non-zero at the upstream, but zero at the downstream.

In this configuration, the fan static pressure is:

$$
\begin{aligned}
p_{f,S} &= p_{f,T} - p_{f,D} \\
&= p_{t,2} - p_{t,1} - (p_{t,2} - p_{s,2}) \\
&= p_{s,2} - p_{t,1} \\
&= -p_{t,1}
\end{aligned}
$$

Modeling Fan Performance Curves

The fan performance curves that are supplied by fan manufacturers plot the pressure rise across that fan against the flow rate. The pressure rise is defined as the static pressure measured downstream of the fan minus the static pressure that is measured upstream of the fan ($p_{s,2} - p_{s,1}$), or as the static pressure measured downstream of the fan minus the total pressure that is measured upstream of the fan ($p_{s,2} - p_{t,1}$).

For more information on how fan manufacturers produce fan performance curves, refer to the appropriate industry standards:

- ANSI/AMCA Standard 210-07 | ANSI/ASHRAE 51-07
  Laboratory Methods of Testing Fans for Certified Aerodynamic Performance Rating
- BS EN ISO 5801:2008, BS 848-1:2007
  Industrial fans: Performance testing using standardized airways

STAR-CCM+ provides two options to define the pressure rise across the fan:

- Standard

  Defines the fan pressure rise as downstream static pressure minus upstream total pressure.

- Static to Static

  Defines the pressure rise as downstream static pressure minus upstream static pressure.

STAR-CCM+ allows for scaling the fan performance curves to fit new fan rotaton rates and fluid temperatures. See:

- Scaling the Fan Performance Curve
- Specifying the Fan Curve Temperature Scaling

STAR-CCM+ assumes that the dynamic pressure is the same upstream and downstream so the Static to Static option can also be used as total pressure downstream minus total pressure upstream.

When you create a fan interface, choose an appropriate option as specified by the fan manufacturer.

See Specifying the Fan Performance Curve Type.

Using the Fan Interface

To implement a fan interface in your simulation, complete the following tasks:

1. Create the fan interface.

   See Creating a Fan Interface and Fan Interface Properties.

2. Specify the appropriate fan performance curve:

   See Specifying the Fan Performance Curve Type.

   - Import the fan performance curve table into the simulation.

     See Specifying the Fan Performance Curve as a Table.

   - Specify the polynomial function that represents the fan performance curve.

     See Specifying the Fan Performance Curve as a Polynomial Function.

3. If necessary, change the fan swirl, fan verbosity, and fan interface intersection tolerance settings.

   See Modeling the Fan Swirl.

See Setting the Fan Interface Verbosity.

See Adjusting the Fan Interface Intersection Tolerance.

Consider the following when you use a fan interface in your simulation:

- When analyzing results it is appropriate to report mass flow rates and pressures on the interface boundaries. This method is more representative than creating derived parts nearby.
- Radiation models require temperature, and an interface model of a fan has no temperature. Therefore, a Radiation model treats the space that the fan occupies as an ordinary opening.

Note: You can also use a fan performance curve to specify a pressure jump across a boundary. This option is available for Pressure Outlet and Stagnation Inlet boundaries only.

Creating a Fan Interface

Each interface has a direction from side 0 to side 1. When you create the interface, select side 0 first and then select side 1. For a fan interface, side 0 is upstream to side 1.

For a fan curve using a specified mass flow rate $\dot{m}$, the mass flow rate is taken as the maximum of $\dot{m}$ or 0. If you select the sides in the wrong order, the fan model computes the mass flow rate $\dot{m}$ as negative. Therefore, you get an erroneous pressure drop for the fan curve at the limit of $\dot{m} = 0$.

To create a fan interface:

1. Expand the Regions > *Fluid Name* > Boundaries node, select the upstream boundary (side 0), and then hold down the `<Ctrl>` key and select the downstream boundary (side 1).
2. Right-click either of the selected boundaries and select Create Interface.

   The new interface node is added under the Interfaces node.
3. Under the Interfaces node, select the new interface node, and set Type to Fan Interface.

   The appropriate fan nodes are added to the Physics Conditions and Physics Values nodes under the Fan Interface node.

   The Pressure Jump Under-Relaxation Factor node is also added under the Physics Values node.

Fan Interface Properties

The Fan Interface node has properties, conditions, values, and a pop-up menu.



Fan Interface Properties

The properties of a Fan Interface node are described below.

Boundary-0	Specifies the upstream boundary of the fan interface. (Read Only)

Boundary-1	Specifies the downstream boundary of the fan interface. (Read Only)
Topology	Defines the connection type between the interfaces. In-place — Uses the in-place topology. Periodic — Uses the periodic topology. Repeating — Uses the repeating topology.
Type	Specifies the type of interface: Fan Interface.
Connectivity	Shows how the boundaries on two sides of the interface are connected. Only Imprinted connectivity is possible.

Specifying the Fan Performance Curve Type

The fan performance curve specifies the pressure rise across the fan interface. By default, you specify the fan performance curve as a table. Alternatively you can specify the fan performance curve as a polynomial function.

To specify the fan performance curve type:

1. Expand the Interfaces > *Interface Name* > Physics Conditions node, and click Fan Curve Type.
2. In the *Properties* window, set the appropriate Curve Type.

Table	The fan performance curve is supplied as a table. See Specifying the Fan Performance Curve as a Table.
Polynomial	The fan performance curve is supplied as a polynomial function. See Specifying the Fan Performance Curve as a Polynomial Function.

For a linear fan curve, both the table and the polynomial are equivalent if the parameters are set correctly.

Specifying the Fan Performance Curve as a Table

When you set up a fan interface, you can specify the fan performance curve as a table.

Use the following steps:

1. Import the appropriate fan performance curve table into the simulation.
2. Expand the Interfaces > [Interface Name] > Physics Values node, and select Fan Curve Table.



3. In the *Properties* window, set the following properties:

Table	Select the fan performance curve table that you imported into your simulation.
Table: X	Select the column that contains the input variable of the pressure rise function for the selected fan performance table.
Units: X	Select the appropriate input variable units.

Table: P

Select the column that contains the pressure data for the selected fan performance table.

Units: P

Select the appropriate pressure units.

Method

Select the input variable of the pressure rise function. When you choose the Static to Static option, ensure that the specified fan curve is a decreasing function of the input variable. When you choose the Standard option, ensure that [specified fan curve + $\frac{1}{2}pv^2$] is a decreasing function of input variable.

Out of the four fan curve method options, only the Mass Average Velocity option is not linearized and is considered to be less robust. The other options apply a different pressure rise to each face at the fan interface. However, for some problems, it can be more meaningful to apply the same pressure rise to all of the faces on the fan interface. In that scenario, you can use the mass averaged velocity options.

Local Velocity	Sets the curve as a function of local velocity. This option is based on the velocity normal to the face. This option is the default.
Mass Average Velocity	Sets the curve as a function of mass average velocity.
Volumetric Flow Rate	Sets the curve as a function of volumetric flow rate.
Mass Flow Rate	Sets the curve as a function of mass flow rate.

Pressure Rise

Specifies how the pressure rise across the fan interface is defined.

See Modeling Fan Performance Curves.

Standard	The difference between the static pressure downstream and the total pressure upstream of the interface. This is the default option.

Static to Static	The difference between the static pressure downstream and the static pressure upstream of the interface. This definition is the same as the difference between the total pressure downstream and the total pressure upstream of the interface.
Operating Rotation Rate	Specifies the rotation rate used by the simulation. STAR-CCM+ uses this value to scale the fan performance curve to match this rotation rate. If this value is equal to the value of Data Rotation Rate, no scaling is done. See Scaling the Fan Performance Curve. STAR-CCM+ always requires this value regardless of whether swirl velocity is added or not. See Fan Swirl Option.
Data Rotation Rate	Specifies the rotation rate at which the fan curve table or polynomial was measured. See Scaling the Fan Performance Curve.

The fan curve table computes the pressure rise by linear interpolation between the data points in the table. Outside the intervals, the nearest limit is used.

You can scale the fan curve in temperature as well as in rotation rate. See Specifying the Fan Curve Temperature Scaling.

Note: For Pressure Outlet and Stagnation Inlet boundaries, the fan does not add swirl velocity, and the operating and data rotation rates are only used for scaling the performance curve.

Specifying the Fan Performance Curve as a Polynomial Function

When you set up a fan interface, you can specify the fan performance curve as a polynomial function.

The fan curve polynomial computes the pressure rise as:

- Given N intervals $[x_0, x_1], [x_1, x_2], ..., [x_{N-1}, x_N]$ the polynomial in the ith interval is $y_i(x_i)$, evaluated using the coefficients for that interval.
- Outside the intervals, the nearest limit is used:

$$y(x) = \begin{cases} y(x_0) & x \le x_0 \\ y(x_N) & x > x_N \end{cases} \tag{106}$$

To specify the fan performance curve polynomial function:

1. Expand the Interfaces > *Interface Name* > Physics Values node, and then click Fan Curve Polynomial.
2. In the *Properties* window, set the appropriate values.

Polynomial	Specifies the fan performance curve polynomial function. See Specifying the Fan Performance Curve Polynomial.
Method	Specifies the input variable of the pressure rise function. For a full description of this property, see Specifying the Fan Performance Curve as a Table.
Pressure Rise	Specifies how the pressure rise across the fan interface is defined. For a full description of this property, see Specifying the Fan Performance Curve as a Table. For more information on fan performance curves, see Modeling Fan Performance Curves.
Operating Rotation Rate	Specifies the rotation rate used by the simulation. Swirl velocity is not added at the interface unless the Fan Swirl Option is activated. If this value is equal to the value of Data Rotation Rate, no scaling is done. See Scaling the Fan Performance Curve.
Data Rotation Rate	Specifies the rotation rate at which the fan curve table or polynomial was measured. If this value is equal to the value of Operating Rotation Rate, no scaling is done. See Scaling the Fan Performance Curve.

You can scale the fan curve in temperature as well as in rotation rate. See Specifying the Fan Curve Temperature Scaling.

Note: For Pressure Outlet and Stagnation Inlet boundaries, the fan does not add swirl velocity, and the operating and data rotation rates are only used for scaling the performance curve.

Specifying the Fan Performance Curve Polynomial

The *Fan Curve Polynomial - Polynomial* dialog lets you specify the fan performance curve polynomial function.



In the *Properties* panel, set the properties to specify the polynomial function that represents the fan performance curve.

Number of Intervals	Set the number of intervals that define the fan performance curve. You can specify different equations for each interval of the fan curve.

Interval Ranges	Specify the minimum and maximum values of each interval. For example, to specify three intervals [A, B], [B, C], [C, D], enter `[A, B, C, D]`.
Number of Coefficients	Specify the number of coefficients per interval. For example, to specify the coefficients for the intervals [A, B], [B, C], [C, D], enter `[a, b, c]`. In this example, `a` is the number of coefficients for interval [A, B], `b` is the number for interval [B, C], and `c` is the number for interval [C, D].
Coefficients	Specify the coefficients for each of the terms in the polynomial equation. Enter the values for each interval, in order. For example, if a=3, b=2, and c=4, enter the values in the format: `[a1, a2, a3, b1, b2, c1, c2, c3, c4]`.
Exponents	Specify the exponents for each of the terms in the polynomial equation. Enter the values for each interval in order, in the same format as for the coefficients.
Input Units	Specify the units for the input variable (the X-axis) of the polynomial equation. The input variable is the flow rate that you set in the Method property of the Fan Curve Specification.
Output Units	Specify the units for the output variable (the Y-axis) of the polynomial equation. The output variable is the pressure drop across the fan interface.

Specifying the Fan Curve Temperature Scaling

You can modify the fan performance curve in response to temperature using Fan Curve Temperature Scaling under Physics Values. See Scaling the Fan Performance Curve.



Fan Curve Temperature Scaling	Adjusts the fan performance curve according to the fluid temperature used in the simulation. Activated — Scale in response to temperature. Deactivated — Do not scale. This is the default.
Data Temperature	Specifies the temperature at which the fan curve table or polynomial was measured. This value is not used unless Fan Curve Temperature Scaling is activated.

Modeling the Fan Swirl

You can model the fan swirl in your simulation. When the fan swirl is included, Eqn. (275) is used to calculate the swirl downstream of the fan. By default, fan swirl is not included.

For more information, refer to Fan Interface Formulation.

To include the fan swirl in your simulation:

1. Expand the Interfaces > *Interface Name* > Physics Conditions node, and click Fan Swirl Option.

2. In the *Properties* window, activate Fan Swirl Option.

   The Fan Swirl node is added to the Physics Values.



3. In the *Properties* window, set the appropriate values.

Rotation Rate	Specify the fan rotation rate, which is the angular velocity ω in Eqn. (272). The available units are radians per second, revolutions per second, and revolutions per minute.
Fan Efficiency	Specify the fan efficiency, η in Eqn. (271).
Coordinate System	Select the appropriate coordinate system. Types other than Laboratory are not available for selection until you create them. You can use any of the coordinate system types for fan modeling, but the following restrictions apply: • The origin of the coordinate axes must coincide with the center of rotation of the fan. • The z-axis must run along the fan axis and be oriented in the fan-flow direction.

Setting the Fan Interface Verbosity

When fan verbosity is activated, a report is printed to the *Output* window for every iteration. The report includes the calculated operating point of the fan: the pressure rise and flow rate value pair that lies on the fan curve at that iteration.

To activate or deactivate the fan verbosity:

1. Expand the Interfaces > *Interface Name* > Physics Values node, and then click Fan Verbosity.
2. In the *Properties* window, set the Verbosity property to the appropriate value: On or Off.

Adjusting the Fan Interface Intersection Tolerance

The intersection tolerance applies to all direct interfaces. The tolerance is the fraction of the length of the smallest edge that is attached to the vertex.

The default tolerance is 0.05. If the intersection is not 100% successful, you can increase the tolerance value to allow the unmatched faces to be matched.

To adjust the fan interface intersection tolerance:

1. Expand the Interfaces > *Interface Name* > Physics Values node, and then click Intersection.
2. In the *Properties* window, set the Tolerance property to the appropriate value.

Using the Fan Momentum Source

The fan momentum source applies the effect of a rotating fan to the fluid cells within a region that you select. Using a fan momentum source implies that you already have a separate region that is located at the position of the fan.

For more information on the fan momentum source model formulation, see Fan Momentum Source Formulation.

To implement a fan momentum source in your simulation, complete the following tasks:

1. Ensure that you have selected a flow model.
2. Obtain the appropriate fan performance curve and import the data table into the simulation.

   For more information, see Importing the Fan Performance Data Table.

3. In the appropriate region, activate the fan momentum source.

   For more information, see Activating the Fan Momentum Source.

4. Configure the fan momentum source by setting the appropriate properties:
   - Select the fan curve table that you want to use, and specify the Voldot and Pressure columns.
   - Set the data rotation rate and the data temperature.

     For more information, see Specifying the Fan Performance Curve.
   - Specify the upstream and downstream boundaries of the fan momentum source.
   - Specify the operating rotation rate and the blade angle of the fan.

     For more information, see Fan Momentum Source Properties.

5. Set the appropriate fan solver properties.

   For more information, see Working with the Fan Solver.

To plot and monitor the flow rate through the fan, use the Surface Integral report with the velocity component normal to the fan face.

Importing the Fan Performance Data Table

STAR-CCM+ calibrates the fan momentum source in your simulation with performance data from a real fan. To provide this performance data, import the appropriate fan curve table into the simulation.

Shown below is a simplified example of a fan performance data table. The two columns are volumetric flow rate, `Vol`, in cubic meters per second (m^3/s), and pressure rise, `Press`, in Pascals (pa):

```
Vol  Press
0.01 458.58
0.20 286.26
0.40 267.68
0.60 239.05
0.80 220.56
1.00 190.52
1.20 146.84
1.40  98.86
1.60  46.45
1.80 -27.10
2.00 -111.15
2.20 -204.01
```

Ensure that the pressure decreases with increasing flow rate across the fan for all data points in the table. If you use a fan curve that does not meet this requirement, the results are unstable.

Activating the Fan Momentum Source

You activate the fan momentum source in the appropriate fluid region.

To activate the fan momentum source:

1. Expand the Regions > [Region] > Boundaries > Physics Conditions node and select Momentum Source Option.



2. In the *Properties* window, set Momentum Source Option to Fan.



The Fan Momentum Source node, and its child, the Fan Solver node, become accessible.



The Fan Momentum Source node represents fan properties and device data for an axial fan. It is one of the physics values of a region.

Specifying the Fan Performance Curve

After you have imported the fan performance curve table, you specify the data columns to use in your simulation:

1. Expand the Regions > [Region] > Boundaries > Physics Values node and select Fan Momentum Source.



2. In the *Properties* window, set the following Fan Momentum Source properties:

Table	Select the fan performance curve table that you imported into your simulation.
Table: VolDot	Select the column that contains the volumetric flow data for the selected fan performance table.
Table: P	Select the column that contains the pressure data for the selected fan performance table.

Data Rotation Rate

Specify the fan rotation rate at which the fan performance curve was measured. The available units are radians per second, revolutions per second, and revolutions per minute.

If you want to use a different rotation rate in your simulation, you specify the appropriate rate in the Fan Momentum Source Operating Rotation Rate property. STAR-CCM+ automatically adjusts the fan curve data to the new rate by using classical fan laws.

Data Temperature

Specify the temperature at which the fan performance curve was measured. The temperature is taken as the mass flow average of the temperature coming into the fan.

If you use a different temperature in your simulation, STAR-CCM+ automatically adjusts the fan curve data to the new temperature by using classical fan laws.

Fan Momentum Source Properties

You specify the upstream and downstream boundaries of the fan momentum source, and set the operating rotation rate and the blade angle of the fan.

The example below shows the upstream and downstream boundaries of the fan region. These boundaries must belong to interfaces where the fan region joins the fluid region.



If your simulation uses an operating rotation rate which is different to the rate at which the fan performance curve was measured, STAR-CCM+ automatically adjusts the fan curve data to the new rate. For more information, see Scaling the Fan Performance Curve.

You set the fan momentum source properties on the Regions > [Region] > Boundaries > Physics Values > Fan Momentum Source node.

Fan Momentum Source Properties

Table	Select the fan performance curve table that you imported into your simulation.
Table: VolDot	Select the column that contains the volumetric flow data for the selected fan performance table.
Table: P	Select the column that contains the pressure data for the selected fan performance table.
Operating Rotation Rate	By default, the fan speed for the simulation is the Data Rotation Rate, the fan rotation rate at which the fan performance curve was measured. If you want to use a different rotation rate in your simulation, specify the appropriate rate here. The available units are radians per second, revolutions per second, and revolutions per minute. STAR-CCM+ automatically adjusts the fan curve data to the new rate by using classical fan laws.
Data Rotation Rate	Specify the fan rotation rate at which the fan performance curve was measured. The available units are radians per second, revolutions per second, and revolutions per minute. If you want to use a different rotation rate in your simulation, you specify the appropriate rate in the Fan Momentum Source Operating Rotation Rate property. STAR-CCM+ automatically adjusts the fan curve data to the new rate by using classical fan laws.
Data Temperature	Specify the temperature at which the fan performance curve was measured. The temperature is taken as the mass flow average of the temperature coming into the fan. If you use a different temperature in your simulation, STAR-CCM+ automatically adjusts the fan curve data to the new temperature by using classical fan laws.

Blade Angle	Specify the angle that the fan blade makes with respect to the flow. The available units are degrees, radians, and revolutions.
Coordinate System	Select the appropriate coordinate system. Types other than Laboratory are not available for selection until you create them. You can use any of the coordinate system types for fan modeling, but the following restrictions apply: • The origin of the coordinate axes must coincide with the center of rotation of the fan. • The z-axis must run along the fan axis and be oriented in the fan-flow direction.
Upstream Boundary	Select the boundary that defines the upstream side of the fan. Select a boundary of the type in-place interface. If none appears in the object selector, create an in-place interface.
Downstream Boundary	Select the boundary that defines the downstream side of the fan. Select a boundary of the type in-place interface. If none appears in the object selector, create an in-place interface.

Scaling the Fan Performance Curve

A fan performance curve is for a specific fan rotation rate and fluid temperature. If your simulation uses a different fan rotation rate or a different fluid temperature from the fan performance curve that you specified in your simulation, the fan performance curve must be modified to suit those conditions. STAR-CCM+ uses established fan laws to scale fan performance curves from one set of conditions to another, as per Eqns. (107) and (108).

Fan Laws

$$\Delta P_2 = \Delta P_1 \times ((\omega_2)/(\omega_1))^2 \times ((T_1)/(T_2)) \tag{107}$$

$$\psi_2 = \psi_1((\omega_2)/(\omega_1)) \tag{108}$$

where ω is the fan speed (rpm), T is the fluid temperature (K), and ψ is the volumetric flow rate.

Working with the Fan Solver

The Fan Solver node contains properties that control how the fan momentum source is solved. You can specify the fan solver property values to suit your simulation.

You set the fan solver properties on the Regions > [Region] > Boundaries > Physics Values > Fan Momentum Source > Fan Solver node.



Fan Solver Properties

Verbosity

This setting controls the printing of diagnostic data to the *Output* window. The settings are:

0 Prints no information about the fan model.

1 Prints deviation of the swirl velocity from the target, the pressure rise target for the current non-reversed volume flow rate through the fan, and the current pressure rise.

IStart

Specify the iteration number at which to start solving for the fan simulation. This setting lets you delay the start of the fan simulation while your simulation solves for other components. The default is 10, as the solution can be unstable for the first few iterations.

Sigma

Specify the factor that is used to match the target for the swirl. This value is σ in Eqn. (281).

Virtual Disk Model

The virtual disk model is based upon the principle of representing propellers, turbines, rotors, fans, and so on, as an actuator disk. The actuator disk treatment is practical when information about the rotor/propeller behavior is available and its effect on the surrounding is desired.

The action of the actuator disk on the flow field enters the momentum equations in the form of a source term that is distributed over the virtual disk. Various distributions of different fidelity are possible to model the action of the actuator depending on the application area.

For the virtual disk model, there are currently three methods implemented:

- Body Force Propeller Method
- Blade Element Method
- 1D Momentum Method

Body Force Propeller Method

As part of the virtual disk model, the body force propeller method is meant to model the effects of a marine propeller. The method models the flow field interaction of the hull of a ship and the propeller. The flow that is induced by the propeller depends on the flow around the ship hull. Similarly, the hull flow is influenced by the propeller.

The body force propeller method can be used to model the propeller propulsion of a ship as part of a DFBI (Dynamic Fluid Body Interaction) simulation. It can be a clear advantage to use this method regarding mesh size and so reduced computational cost over performing a simulation employing the fully resolved geometry of the propeller. It is useful if you are not interested in analyzing the flow field around the propeller but more importantly need the correct propulsion specification.

The body force propeller method uses the following approach:

- Definition of a virtual disk regarding position and direction in which thrust is produced.
- Specification of a propeller performance curve.
- Specification of an operating point, for example, the rotational speed.
- Specification of an inflow velocity plane that yields an average velocity and density for propeller inflow.

As a result, the radially varying distribution of the axial and tangential forces of the modeled propeller acting on the flow is calculated. The integration of these forces over the disk gives the thrust and torque of the propeller, which are then available for coupling with a DFBI simulation.

See also:

- Setting up and using the Body Force Propeller Method
- Observing Virtual Disk Meshing Requirements for the Body Force Propeller Method
- Selecting the Virtual Disk Model
- Creating a Virtual Disk for the Body Force Propeller Method
- Setting the Propeller Performance Curve for the Body Force Propeller Method
- Specifying the Disk Geometry for the Body Force Propeller and the 1D Momentum Methods
- Specifying the Disk Orientation
- Specifying the Inflow Velocity Plane
- Specifying the Propeller Handedness for the Body Force Propeller Method
- Specifying the Operation Point for the Body Force Propeller Method
- Visualizing the Virtual Disk and the Inflow Velocity Plane
- Reporting Virtual Disk Results
- Virtual Disk Field Functions

Setting up and using the Body Force Propeller Method

The setup of the virtual disk of the body force propeller method does not require the creation of a separate region to which the model is then applied. You create the shape of the virtual disk by specifying the relevant parameters such as radius, thickness, and orientation to an existing mesh. To verify the dimensions and orientation of the virtual disk, you can visualize it by using virtual disk specific field functions.

To set up and run the body force propeller method:

1. Generate a mesh that is suitable for use with the virtual disk model.

   See Observing Virtual Disk Meshing Requirements for the Body Force Propeller Method

2. Select the virtual disk model.

   See Selecting the Virtual Disk Model

3. Create a virtual disk.

   See Creating a Virtual Disk for the Body Force Propeller Method

4. Visualize the virtual disk and the inflow velocity plane.

   See Visualizing the Virtual Disk and the Inflow Velocity Plane

Observing Virtual Disk Meshing Requirements for the Body Force Propeller Method

It is recommended that you refine the mesh in the zone where the virtual disk is located, particularly in the radial direction. This way, you obtain a good distribution of the axial and tangential body force components of the virtual disk.

For details on the body force computation, see Body Force Propeller Method.

You can achieve a suitable mesh refinement by use of a volumetric control that is based on a cylindrical part shape.

For more information, see Mesh Refinement in a Specified Volume.

The diameter of the cylindrical part shape should correspond to the diameter of the inflow velocity plane.

For more information, see Specifying the Inflow Velocity Plane.

If you want to analyze the flow field in the wake of the propeller in more detail, refine the wake zone to a length of approximately fives times the virtual disk diameter.

Note: If you set up a case from scratch, that is, the mesh is not yet available, it is recommended to first activate the virtual disk model. Then specify the geometry and location of the virtual disk. Create the required volumetric controls for mesh refinement and then proceed with the meshing.

Selecting the Virtual Disk Model

Activate the virtual disk model as an optional physical model from the *Physics Model Selection* dialog of the relevant physics continuum.

To select the virtual disk model:

1. Right-click the Continua > Physics 1 > Models node and select Select models...

2. Select the appropriate models:
   - Make sure that Three Dimensional is selected from the *Space* group box. The virtual disk model is only available for this option.
   - Select any option from the *Time* group box except Harmonic Balance, the *Flow* group box, the *Equation of State* group box, and the *Viscous Regime* group box.
   - Select any material model from the *Material* group box except Solid and Multi-Part Solid. The virtual disk model is not available for solid materials.

   The Virtual Disk appears in the *Optional Models* group box.
3. In the *Physics 1 Model Selection* dialog, select Virtual Disk from the *Optional Models* group box.



The Virtual Disk node appears under the Models node in the simulation tree. It contains the Virtual Disks sub node that you use to actually create the virtual disk.



Creating a Virtual Disk for the Body Force Propeller Method

After having selected the virtual disk model, specify the details of the virtual disk.

To create a virtual disk:

1. Right-click the Virtual Disk > Virtual Disks node and select New.

A new child node that is called Virtual Disk 1 is added to the simulation tree.



2. Select the Virtual Disk 1 node and, in the *Properties* window, select the Body Force Propeller Method from the drop-down list.

The additional Propeller Curve, Disk Geometry, Inflow Velocity Plane, Propeller Handedness Option, Operation Point Input Option, and RotationRate(n) sub nodes are added to the simulation tree. The RotationRate(n) sub node can change depending on what you select as Operation Point Input Option



A new local coordinate system for the virtual disk is created automatically. You can find it under the Tools > Coordinate Systems > Laboratory > Local Coordinate Systems > Virtual Disk 1-CSys 1 node. This coordinate system is oriented such that the z-axis is always normal to the plane of the disk and that it points in the direction of thrust that is produced by the virtual disk.

Do not modify the local coordinate system of the disk here in the Local Coordinate Systems node. Modify the virtual disk origin and orientation only under the Virtual Disk node in the relevant physics continuum.



For more information, see Specifying the Disk Orientation.

3. Define the Virtual Disk by configuring the following nodes:

- Propeller Curve

  Set the propeller performance curve as a function of advance ratio.

  See Setting the Propeller Performance Curve for the Body Force Propeller Method

- Disk Geometry

  Define the disk geometry in terms of geometrical dimensions, location, and orientation.

  See Specifying the Disk Geometry for the Body Force Propeller and the 1D Momentum Methods

- Inflow Velocity Plane

  The inflow velocity plane is located in front of the virtual disk. It is always defined in the direction of the disk normal.

  See Specifying the Inflow Velocity Plane

- Propeller Handedness Option

See Specifying the Propeller Handedness for the Body Force Propeller Method

- Operation Point Input Option

  See Specifying the Operation Point for the Body Force Propeller Method

Setting the Propeller Performance Curve for the Body Force Propeller Method

Generally, the propeller performance is available from experiment in terms of thrust coefficient K_T, torque coefficient K_Q, and propeller efficiency η_0 as a function of advance ratio J (see Eqns. (323), (324), (325) for a definition of these values).

To specify the propeller performance curve:

1. Select the Virtual Disk 1 > Propeller Curve node.
2. In the *Properties* window, select the Method that you would like to use to specify the propeller performance curve. There are two options available:
   - Polynomial(J)



   i. Select the Polynomial(J) node and then click the ... (custom editor) next to the Kt = f(J), Kq = f(J) and Eta = f(J).

   ii. In the *Polynomial(J)* dialog, enter the property values of the respective polynomial (see Setting Values Using a Polynomial).

- Table(J)



i. Import the appropriate propeller performance curve table into the simulation (see Reading a Table Data File).

ii. Select the Table(J) node. In the *Table (J) - Properties* window, set the following properties:

Table	Select the propeller performance curve table that you imported into your simulation.
Table: J	Select the column that contains the advance ratio for the selected propeller performance table.
Table: Kt	Select the column that contains the thrust coefficient for the selected propeller performance table.
Table: Kq	Select the column that contains the torque coefficient for the selected propeller performance table.
Table: Eta	Select the column that contains the propeller efficiency for the selected propeller performance table.

Specifying the Disk Geometry for the Body Force Propeller and the 1D Momentum Methods

The virtual disk is a simplified representation of a complex rotating machine such as a propeller or wind turbine. Only the geometrical dimensions of the disk are required. Details such as the number of blades, the chord, and twist variation are not necessary for this method.

To specify the virtual disk geometry:

1. Select the Virtual Disk 1 > Disk Geometry node.



2. In the *Disk Geometry - Properties* window, set the following properties:

Inner Radius	Specify the inner radius of the disk. This radius corresponds to the hub radius of the propeller or the wind turbine.
Outer Radius	Specify the outer radius of the disk. This radius corresponds to the tip radius of the propeller or the wind turbine.
Thickness	Specify the thickness of the propeller or the wind turbine.

Origin	Specify the coordinates of the origin of the disk. The origin is defined in terms of the coordinates and the coordinate system with respect to which the coordinates are specified. Here, the choice of coordinate system is dictated by the selection you make for the Orientation Specification property: • For the option Disk Normal and Coordinate System, the coordinates of the origin are specified with respect to the coordinate system that is selected in that node. • For the option Pitch Angle and Lateral Angle, the coordinates of the origin are always with respect to the laboratory coordinate system.
Local Coordinate System	Reference to the local coordinate system of the disk (read-only).
Orientation Specification	You specify the disk orientation by selecting this option. The disk orientation defines in which direction the thrust is produced. For more information, see Specifying the Disk Orientation. Disk Normal and Coordinate System: Specify the disk orientation in terms of disk normal and coordinate system (default). When you select this option, a Normal and Coordinate System sub node appears in the simulation tree. Pitch Angle and Lateral Angle: Specify the disk orientation in terms of lateral angle and pitch angle. When you select this option, a Lateral And Pitch Angle sub node appears in the simulation tree.

Specifying the Disk Orientation

In addition to the geometric dimensions and the origin of the virtual disk, specify its orientation. The way that the virtual disk is oriented, that is, in which direction the z-axis of the local disk coordinate system points, is the direction in which the virtual disk produces thrust. The positive z-axis points in the direction in which the disk would move if let free.

Before you begin to set the orientation of the virtual disk, it can be helpful to visualize the local disk coordinate system.

To visualize the local virtual disk coordinate system:

1. Open a scene, for example a geometry scene.

   For more details, see Creating a Scene.

2. Click the scene/plot button.

3. In the scene node, for example the Geometry Scene 1 node, click the (custom editor) and set the Coordinate Systems property to [Laboratory->Virtual Disk 1-CSys 1].

4. In the object selection dialog, select Virtual Disk 1-CSys 1 and click OK.



You can now see the current orientation of the local coordinate system of the virtual disk in your currently open scene in the *Graphics* window. The orientation of this local coordinate system adapts in accordance with the input to the disk orientation.

You have two options available to specify the disk orientation:

- Disk orientation in terms of disk normal and coordinate system
- Disk orientation in terms of lateral and pitch angle

To specify the disk orientation in terms of disk normal and coordinate system:

1. Select the Disk Geometry node and set the Orientation Specification to Disk Normal and Coordinate.
2. Select the Disk Geometry > Normal and Coordinate System node.
3. In the *Normal And Coordinate System - Properties* window, set the following properties:

Coordinate System	Specify the coordinate system in regard to which the normal of the virtual disk is defined. You can also use a user-defined local cartesian coordinate system to define the disk normal.
Disk Normal	Specify the direction that is normal to the plane of the disk.

To specify the disk orientation in terms of disk lateral and pitch angle:

1. Select the Disk Geometry node and set the Orientation Specification to Pitch Angle and Lateral Angle.
2. Select the Disk Geometry > Lateral And Pitch Angle node.
3. In the *Lateral and Pitch Angle - Properties* window, set the following properties:

Pitch Angle	To rotate the local disk coordinate system, you specify a pitch angle. The pitch angle rotates the x-z plane around the y-axis.
Lateral Angle	To rotate the local disk coordinate system, you specify a lateral angle. The lateral angle rotates the y-z plane around the x-axis.

Specifying the Inflow Velocity Plane

The inflow velocity plane is a plane that is located upstream of the virtual disk and is always oriented with respect to the direction of the virtual disk normal. The velocity components and the density of the fluid approaching the virtual disk are volume-averaged over this plane. The averaged velocity components are then projected onto the normal plane of the disk to yield one average velocity value.

This averaged value is used by the body force propeller method to compute the advance ratio, the thrust, and the torque (see Eqns. (326), (327), (328)).

The 1D Momentum method uses this averaged value to calculate Eqs. (343) to (360).

To specify the inflow velocity plane:

1. Select the Virtual Disk 1 > Inflow Velocity Plane node.



2. In the *Inflow Velocity Plane - Properties* window, set the following properties:

Local Coordinate System	Reference to the local virtual disk coordinate system (read-only).
Velocity Plane Radius	Specify the radius of the inflow velocity plane. It is recommended that you use a value that is approximately 10% larger than the virtual disk radius.

Velocity Plane Offset

Specify the offset in the normal (z) direction of the virtual disk. You specify the offset with respect to the virtual disk origin.

- For the **body force propeller method**, it is recommended that you use an offset of 10% of the virtual disk diameter.

- For the **1D momentum method**, there is no recommendation for the offset. Note that the offset is in the negative z direction, because generally, the thrust of a wind turbine, and therefore, the positive z-axis is directed downstream of the virtual disk.

Specifying the Propeller Handedness for the Body Force Propeller Method

A propeller is characteristically classified in terms of handedness. You specify the handedness of the propeller to relate the direction of the rotation of the propeller to the direction of thrust of the vessel.

For more detailed information on how the propeller handedness is defined and how it interacts with the local coordinate system of the disk and the operation point, see Propeller Handedness.

To specify the propeller handedness:

1. Select the Virtual Disk 1 > Disk Geometry > Propeller Handedness Option node.



2. In the *Propeller Handedness Option - Properties* window, select the Propeller Handedness Properties.

Propeller Handedness Option Properties

Option	Choose between two different propeller handedness options. For a definition of the propeller handedness and how it affects the directions of rotation, see Propeller Handedness.
	Right Handed — Specifies the propeller is right-handed. A right-handed propeller **always** rotates in the counter-clockwise direction with respect to the positive z-axis of the local coordinate system of the virtual disk.
	Left Handed — Specifies the propeller is left-handed. A left-handed propeller **always** rotates in the clockwise direction with respect to the positive z-axis of the local coordinate system of the virtual disk.

Propeller Handedness

A propeller is - among other parameters - characterized by its handedness: it can be right-handed or left-handed. The handedness relates to the shape of the propeller, which in turn relates the propeller rotation direction to the direction of translation of the vessel.

In STAR-CCM+, the definition of the propeller handedness is modified compared to the traditional definition that is typically used by the marine community.

Propeller Handedness in the Marine Industry

The definition of the handedness requires a reference direction. Traditionally, the marine community defines the propeller handedness with respect to the stern of the vessel:

- A **right-handed** propeller rotates clockwise to produce forward thrust when viewed from the stern of the vessel.
- A **left-handed** propeller rotates counter-clockwise to produce forward thrust when viewed from the stern of the vessel.

The handedness of a propeller affects its shape. When you view a propeller from astern, the leading edges are farther away from you than the trailing edges. For example, the leading edges of a right-handed propeller are on the right and the propeller rotates clockwise when viewed from astern.

Propeller Handedness in STAR-CCM+

The selection of propeller handedness is available as an option as part of the Body Force Propeller method.

Generally, the design of a propeller is optimized to be used for forward motion. However in practice, a propeller can also be used for the reverse motion of a ship. Considering this fact, STAR-CCM+ does not use the stern of the vessel as the reference direction for the definition of left-/right-handedness as is the case for the traditional definition (see Propeller Handedness in the Marine Industry).

In STAR-CCM+, the local coordinate system of the virtual disk is used as the reference for defining the propeller handedness:

- A **right-handed** propeller rotates counter-clockwise when viewed from the positive direction of the z-axis of the local coordinate system of the virtual disk.
- A **left-handed** propeller rotates clockwise when viewed from the positive direction of the z-axis of the local coordinate system of the virtual disk.

The use of the local coordinate system of the virtual disk does not change the definition of handedness, only the reference axis for the direction of rotation is changed. This approach is consistent because the local coordinate

system of the virtual disk changes according to the direction of propulsion of the vessel. The propulsion direction is specified in terms of the orientation specification of the virtual disk.

Rotation Direction

The Propeller Handedness Option is useful to prescribe or change the direction of rotation of the propeller - clockwise or counter-clockwise - in combination with the Operation Point Input Option (see Specifying the Operation Point for the Body Force Propeller Method) and the Orientation Specification of the virtual disk (see Specifying the Disk Orientation).

To specify the operation point, there are three options available:

- Rotation rate with $\omega \geq 0$.
- Thrust with $T \geq 0$.
- Torque with $Q \geq 0$.

The direction of propulsion is specified in terms of the orientation specification of the virtual disk. Forward and reverse propulsion direction is defined in terms of the selected coordinate system and its orientation with respect to the nose and the stern of the vessel.

In the following table, it is assumed that the Disk Normal and Coordinate System option is used as Orientation Specification. The combination of propeller handedness and orientation of the local coordinate system of the virtual disk and its effect on the directions of propulsion and propeller rotation is summarized as follows:

Handedness	Normal	Propulsion Direction	Rotation Direction
Right	(1,0,0)	Forward	CCW
Right	(-1,0,0)	Backward	CCW
Left	(1,0,0)	Forward	CW
Left	(-1,0,0)	Backward	CW

In the following figures, the orientation of the virtual disk is specified in terms of the Body coordinate system.The orientation of the local coordinate system of the virtual disk automatically adapts to the normal direction specified with respect to the Body coordinate system. The direction of propeller rotation for each type of handedness is shown:



Right-handed propeller in forward configuration



Right-handed propeller in backward configuration



Left-handed propeller in forward configuration



Left-handed propeller in backward configuration

Specifying the Operation Point for the Body Force Propeller Method

The operation point defines where on the specified propeller performance curve you want to perform the simulation.

To specify the operation point:

1. Select the Virtual Disk 1 > Operation Point Input Option node.



2. In the *Operation Point Input Option - Properties* window, select the following property:

Option	Choose between three different operation point options.	
	Rotation Rate (n)	Select the rotation rate as the operation point.
	Thrust (T)	Select the thrust as the operation point.
	Torque (Q)	Select the torque as the operation point.

Depending on your previous selection under the Operation Point Input Option node, either the RotationRate (n), the Thrust, or the Torque sub node appears.

3. Specify the respective operating point value and units in this sub node.

Blade Element Method

The blade element method is a method within the framework of the virtual disk model. The blade element method is primarily designed to analyze the complex flow field that is associated with helicopter rotors. The complexity of the flow arises due to mutual aerodynamic interactions among the main rotor, the fuselage, and the tail rotor. The asymmetrical loading over the rotor disk complicates these interactions. A correct solution of the flow field around a rotor plays a crucial role in determining the aerodynamic characteristics and the performance of the helicopter.

The blade element method models the spinning rotor as a distribution of momentum sources. The strength of the source terms and their variations are interactively determined from the rotor geometry and the local velocity field. The number and the geometry of the blades in terms of chord and twist variation along the rotor radius are specified. The blade geometry is not explicitly resolved. The influence of the blade geometry on the flow field is described by the aerodynamic behavior of the blade in terms of lift and drag coefficients that are obtained for the corresponding two-dimensional cross-section of the blade. The aerodynamic characteristic of a blade can vary along the radius for an unlimited number of airfoil cross-sections. In forward flight, the rotor has an asymmetrical loading, which leads to undesired moments with respect to the rotor hub. Therefore, rotor trimming is necessary. The implemented method allows you to trim the rotor to the specified thrust and moments by using a non-linear iterative approach.

The blade element method has numerous advantages over a transient rigid body motion approach for this application. It eliminates the need for a body-fitted mesh along the blades, because the blades are not present physically. The wake structure is well captured and requires no extra model. The virtual disk approach allows the intersecting or overlapping of rotors. There is practically no restriction on the number of rotors. It is possible to study the main rotor and the tail rotor interaction of a single helicopter or numerous helicopters in a flying configuration. The method is not restricted to helicopter applications and can also be used to study automotive radiator fans.

The blade element method provides a computationally cost-effective and reliable method to obtain a flow field, which has enough details to make design decisions leading to an improved overall performance.

See also:

Setting up and Using the Blade Element Method

Creating a Virtual Disk for the Blade Element Method

Specifying the Airfoil Sections

Specifying the Chord Distribution

Specifying the Twist Distribution

Specifying the Disk Geometry for the Blade Element Method

Setting the Rotation Rate for the Blade Element Method

Ramping the Rotation Rate

Specifying the Disk Stick Specification

Specifying the Virtual Disk Trimming

Specifying the Virtual Disk Resolution
Setting the Source Convergence Control
Visualizing the Virtual Disk and the Inflow Velocity Plane
Reporting Virtual Disk Results
Virtual Disk Field Functions

Setting up and Using the Blade Element Method

Similar to the body force propeller method, the setup of the virtual disk for the blade element method does not require the creation of a separate region to which the method is then applied. Instead, you create the shape of the virtual disk by specifying the relevant parameters such as radius, thickness, and orientation to an existing mesh. To verify the dimensions and orientation of the virtual disk, you can visualize it by using virtual disk-specific field functions.

To set up and run the blade element method:

1. Select the virtual disk model.

   See Selecting the Virtual Disk Model
2. Create a virtual disk.

   See Creating a Virtual Disk for the Blade Element Method
3. Visualize the virtual disk.

   See Visualizing the Virtual Disk and the Inflow Velocity Plane

Creating a Virtual Disk for the Blade Element Method

After having selected the blade element method, specify the details of the virtual disk using the following components:

- Geometric definition of the blades
  - Origin and orientation
  - Outer radius and root cut-out radius
  - Number of blades
- Blade information
  - Chord
  - Twist
  - Airfoil data (Drag and lift coefficients as a function of angle of attack and Mach/Reynolds number)
- Rotor conditions

- Rotation rate
- Collective pitch
- Lateral and longitudinal cyclic pitch
- Trim Condition

To create a virtual disk:

1. Right-click the Virtual Disk > Virtual Disks node and select New.

   A new child node that is called Virtual Disk 1 is added to the simulation tree.



2. Select the Virtual Disk 1 node and, in the *Properties* window, select the Blade Element Method from the drop-down list.

The additional Airfoil Sections, Chord Distribution, Twist Distribution, Disk Geometry, Rotation Rate, Disk Stick Specification, Virtual Disk Resolution, and Source Convergence Control sub nodes are added to the simulation tree.



A new local coordinate system for the virtual disk is created automatically. You can find it under the Tools > Coordinate Systems > Laboratory > Local Coordinate Systems > Virtual Disk 1-CSys 1 node. This coordinate system is oriented such that the z-axis is always normal to the plane of the disk.

Do not modify the local coordinate system of the disk here in the Local Coordinate Systems node. Modify the virtual disk origin and orientation only under the Virtual Disk node in the relevant physics continuum.



For more information, see Specifying the Disk Orientation.

3. Define the Virtual Disk by configuring the following nodes:
   - Airfoil Sections

     See Specifying the Airfoil Sections
   - Chord Distribution

     See Specifying the Chord Distribution
   - Sweep Angle Distribution

     See Specifying the Sweep Angle Distribution
   - Twist Distribution

     See Specifying the Twist Distribution
   - Disk Geometry

     See Specifying the Disk Geometry for the Blade Element Method
   - Rotation Rate

     See Setting the Rotation Rate for the Blade Element Method
   - Disk Stick Specification

See Specifying the Disk Stick Specification

- Virtual Disk Resolution

  See Specifying the Virtual Disk Resolution

- Source Convergence Control

  See Setting the Source Convergence Control

Specifying the Airfoil Sections

You can use the Airfoil Sections node to specify drag and lift coefficients as functions of Reynolds or Mach number and angle of attack. Providing this data brings the effect of local blade aerodynamics into the virtual disk simulation. An airfoil section is a cross-section of the blade at a certain radial position. You can either specify one airfoil section for the entire blade or you can subdivide the blade into several cross-sections and provide local drag and lift coefficient data.

To specify an airfoil section:

1. Right-click Virtual Disk 1 > Airfoil Sections and select New.

   An Airfoil Section 1 node with two sub-nodes is added to the simulation tree.



2. In the *Airfoil Section 1 - Properties* window, set the appropriate properties.

   See Airfoil Section Properties

3. Import the airfoil data, that is the drag and lift coefficients, each as a table. The following procedure describes the setup for the lift coefficient as a function of Mach number. The procedure is the same for the drag coefficient and for a dependence on the Reynolds number.

There are two table format options available:

- File table
  i. Import the appropriate lift coefficient table into the simulation.

     See Reading a Table Data File

  ii. Select the Cl (AoA, Mach) node.



  iii. In the *Cl (AoA, Mach) - Properties* window, set the appropriate properties.

     See Lift Coefficient Properties.

- C81 table
  i. Right-click the Virtual Disk 1 > Airfoil Sections node and select Import C81 Airfoil Section Data...

  ii. In the *Open* dialog, set the Airfoil Data Function to (alpha, Mach) and navigate to the C81 airfoil data file. The file contains the data for both the lift and the drag coefficient.

  iii. Click Open.

     Two nodes for the lift and drag coefficients, respectively, are added to the Tools > Tables node.

  iv. Select the Virtual Disk 1 > Airfoil Sections > Airfoil Section 1 > Cl (AoA, Mach) node and set the appropriate properties in the *Cl (AoA, Mach) - Properties* window.

See Lift Coefficient Properties.

Airfoil Section Properties

Airfoil Section Function Specification	You specify the local aerodynamic coefficients of the rotor blade for this airfoil section at a certain radial position. Depending on the selection here, the Cl and Cd sub-nodes change. **Mach Number**: Specify the aerodynamic coefficients as a function of angle of attack and the Mach number. **Reynolds Number**: Specify the aerodynamic coefficients as a function of angle of attack and the Reynolds number.
Normalized Disk span	Dimensionless radial coordinate along the blade with values from 0 through 1. It is defined as $\frac{r}{R}$ where r is the radial coordinate and R is the maximum blade radius. Here, you specify the radial position of the airfoil section.

Lift Coefficient Properties

Table: Mach Number	Select the column that contains the Mach number for the selected input table.
Table: Angle of Attack	Select the column that contains the angle of attack for the selected input table.
Table: Cl	Select the column that contains the lift coefficient for the selected input table.
Input table	Select the table that contains the airfoil section lift coefficient data.

Columns	Lists the imported columns of the selected input table (read-only).
Verbose	Activates additional output.

Specifying the Chord Distribution

The chord is a straight line connecting the leading and the trailing edge of a blade or airfoil. It describes the width of the blade. For the blade element method, you specify the chord length either constant or as a function of the normalized disk span r/R.

To specify the chord distribution:

1. Select the Virtual Disk 1 > Chord Distribution node.



2. In the *Chord Distribution - Properties* window, set the appropriate properties.

   See Chord Distribution Properties.

3. Corresponding to the Method that you specified, perform one of the following steps:
   - Constant
     i. Select the Chord Distribution > Constant node and specify a chord length Value.
   - Polynomial in (r/R)

i. Select the Chord Distribution > Polynomial in (r/R) node and then click the ... (custom editor) next to Polynomial.

ii. In the *Polynomial in (r/R)* dialog, enter the property values of the respective polynomial (see Setting Values Using a Polynomial).

- Table (r/R)

i. Import the appropriate chord distribution table into the simulation (see Reading a Table Data File).

ii. Select the Table (r/R) node. In the *Table (r/R) - Properties* window, set the appropriate properties.

See Table (r/R) Properties.

Chord Distribution Properties

Dimensions	Gives the dimensions of the specified chord (read-only).	
Method	You specify the method how to provide the chord length data.	
	Constant	You set a constant value for the chord.
	Polynomial in (r/R)	You specify the chord distribution as a function of dimensionless disk span in the form of a polynomial.
	Table (r/R)	You specify the chord distribution as a function of dimensionless disk span in the form of a table.

Table (r/R) Properties

Table: Data	Select the column that contains the chord or twist distribution data.
Table: Normalized Span	Select the column that contains the position in the form of normalized span (r/R) for the selected chord or twist distribution.
Input table	Select the chord or twist distribution table that you imported into your simulation

Columns	Displays the extracted columns from the table (read-only).
Verbose	Activates additional output information.

Specifying the Sweep Angle Distribution

Similar to the wings of a subsonic aircraft, the rotor blades of a helicopter can experience the effects of wave drag at the tip of the blades. Wave drag occurs when the air flow reaches the speed of sound at the leading edge at the tip of the blade. To avoid the occurrence of wave drag, the local velocity at the blade tip must be kept below the speed of sound. The local tangential velocity is given by the rotation rate times the blade radius ($v_\varphi = r\omega$). Therefore, these wave drag effects restrict the achievable maximum rotation rates in combination with the blade radius to a certain limit.

To remedy this limitation, the wave drag effects can be reduced by introducing a sweep-back of the rotor blades. Sweeping the blade reduces the local velocity at the wing section by a factor of cosine of the angle of sweep. The sweeping thus increases the critical Mach number.

To specify the sweep angle distribution:

1. Select the Virtual Disk 1 > Sweep Angle Distribution node.



2. In the *Sweep Angle Distribution - Properties* window, set the appropriate properties.

   See Sweep Angle Distribution Properties.

3. Corresponding to the Method that you specified in the previous step, perform one of the following steps:
   - Constant
     - Select the Sweep Angle Distribution > Constant node and specify a sweep angle Value.
   - Polynomial in (r/R)
     i. Select the Sweep Angle Distribution > Polynomial in (r/R) node and then click ... (custom editor) next to Polynomial.
     ii. In the *Polynomial in (r/R)* dialog, enter the property values of the respective polynomial (see Setting Values Using a Polynomial).
   - Table (r/R)
     i. Import the appropriate sweep angle distribution table into the simulation (see Reading a Table Data File).
     ii. Select the Table (r/R) node. In the *Table (r/R) - Properties* window, set the appropriate properties.

     See Table (r/R) Properties.

Sweep Angle Distribution Properties

The Sweep Angle Distribution node requires the setting of several properties. You find a description of these properties in the following table:

Dimensions	Specifies the dimensions of the sweep angle data (read-only). It is set to Angle.	
Method	Specifies how the sweep angle data is provided.	
	Constant	Sets a constant value for the sweep angle.
	Polynomial in (r/R)	Specifies the sweep angle distribution as a function of dimensionless disk span in the form of a polynomial.
	Table (r/R)	Specifies the sweep angle distribution as a function of dimensionless disk span in the form of a table.

Specifying the Twist Distribution

A blade can be built with a twist to achieve different pitch angles between the hub and the tip. You can specify the twist distribution as constant or as a function of normalized disk span r/R.

To specify the twist distribution:

1. Select the Virtual Disk 1 > Twist Distribution node.



2. In the *Twist Distribution - Properties* window, set the appropriate properties.

   See Twist Distribution Properties.

3. Corresponding to the Method that you specified, perform one of the following steps:

   - Constant

     i. Select the Twist Distribution > Constant node and specify a twist length Value.

   - Polynomial in (r/R)

     i. Select the Twist Distribution > Polynomial in (r/R) node and then click the ... (custom editor) next to Polynomial.

     ii. In the *Polynomial in (r/R)* dialog, enter the property values of the respective polynomial (see Setting Values Using a Polynomial).

   - Table (r/R)

     i. Import the appropriate twist distribution table into the simulation (see Reading a Table Data File).

     ii. Select the Table (r/R) node. In the *Table (r/R) - Properties* window, set the appropriate properties.

     See Table (r/R) Properties.

Twist Distribution Properties

Dimensions	Gives the dimensions of the specified twist (read-only). It is set to Angle.	
Method	You specify the method how to provide the twist data.	
	Constant	You set a constant value for the twist.
	Polynomial in (r/R)	You specify the twist distribution as a function of dimensionless disk span in the form of a polynomial.
	Table (r/R)	You specify the twist distribution as a function of dimensionless disk span in the form of a table.

Specifying the Disk Geometry for the Blade Element Method

Here, you define the geometrical dimensions of the disk together with its location and orientation within the computational domain. The disk for the blade element method is of cylindrical shape. Specify the thickness of the disk at least equal to or 10-15% greater than the cell size of the volume mesh where the disk is located. This way you ensure that the cells comprising the disk are continuous (that is, there is no hole in the disk) and that the disk is at least one cell layer thick. It is not recommended that you specify a larger thickness for the disk, because this larger thickness leads to a smearing of the local effect of the blades.

To describe the disk geometry:

1. Select the Virtual Disks > Virtual Disk 1 > Disk Geometry node.



2. In the *Disk Geometry - Properties* window, set the appropriate properties. See Disk Geometry Properties for Blade Element Method

Disk Geometry Properties for Blade Element Method

Number of Blades	You specify the number of blades of the rotor geometry that you want to model
Inner Radius	The inner radius corresponds to the root cut-out radius. The cut-out radius is typically 5% of the outer radius of the rotor.
Outer Radius	The outer radius corresponds to the tip radius of the rotor.
Thickness	You specify the thickness of the virtual disk for the blade element method. The thickness must be chosen such that the virtual disk is approximately one cell layer thick. Therefore, it is recommended that you set the thickness to a value that corresponds to the mesh cell size in the area of the virtual disk.
Origin	Specify the coordinates of the origin of the disk. The origin is defined in terms of the coordinates and the coordinate system with respect to which the coordinates are specified. Here, the choice of coordinate system is dictated by the selection you make for the Orientation Specification property: • For the option Disk Normal and Coordinate System, the coordinates of the origin are specified with respect to the coordinate system that is selected in that node. • For the option Pitch Angle and Lateral Angle, the coordinates of the origin are always with respect to the laboratory coordinate system.
Local Coordinate System	Reference to the local coordinate system of the disk (read-only).

Orientation Specification	You specify the disk orientation by selecting this option. The disk orientation defines in which direction the thrust is produced. For more information, see Specifying the Disk Orientation.
Tip Effect Correction	You specify the fraction of the rotor blade that contributes to the lift generation. The value is specified in terms of the dimensionless normalized span r/R within a range of [0,1]. For more information, see Tip Correction for the Blade Element Method.

Tip Correction for the Blade Element Method

A trailing vortex forms at the tip of each rotor blade. This trailing vortex leads to a higher local inflow near the tip, which effectively reduces the lift generation capability of the blade in this region. This phenomenon is known as **Tip Loss**. Tip loss is the reduction in lift that is due to the influence of tip vortices in the flow.

The blade element method accounts for this effect by including the Tip Effect Correction.

For example, a value of 0.96 for the Tip Effect Correction means that only 96% of the blade generates lift and therefore contributes to the generated thrust. The remaining 4% does not contribute to the lift, but only to the drag.

Setting the Rotation Rate for the Blade Element Method

To set the rotation rate for the blade element method:

1. Select the Virtual Disks > Virtual Disk 1 > Rotation Rate node.



2. In the *Rotation Rate - Properties* window, set the following properties:

Rotation Rate	The rotation rate of the virtual disk.
Ramp Method	Specify the ramping of the rotation rate. For high rotation rates, this ramping can improve convergence.

No Ramp	Ramping of the rotation rate is deactivated.
Linear Ramp	The rotation rate is linearly ramped from a specified initial value to the final specified Rotation Rate. When you select this option, an additional Linear Ramp sub-node appears in the simulation tree. For more information, see Ramping the Rotation Rate.

Ramping the Rotation Rate

If you observe convergence problems that are related to a high rotation rate, you can linearly ramp the rotation rate.

To ramp the rotation rate:

1. Select the Virtual Disks > Virtual Disk 1 > Rotation Rate node and set the Ramp Method to Linear Ramp.

2. Select the Rotation Rate > Linear Ramp node.



3. In the *Linear Ramp - Properties* window, set the following properties:

Start Iteration	You specify the iteration on which to begin the ramping.
End Iteration	You specify the iteration on which to end the ramping.
Initial Value	You specify the low rotation rate value with which to begin the ramping.

Disk Stick Specification

Particularly for aerospace applications, it is important to take the pitch variation of the blades into account. For a helicopter, the pitch angle of the blade θ is important in establishing the angle of attack of the blade with respect to the rotor disk. Due to the asymmetrical nature of the rotor loading in forward flight, the pitch angle needs special treatment.

If the pitch angle of all blades were kept constant, the advancing side would have a higher angle of attack than the receding side.

To avoid such loading, the pitch angle consists of a collective part θ_{r0}, which is the same for each blade and the cyclic pitch component, which varies in the azimuthal direction. The complete blade pitch angle is then defined by the first harmonic term of a Fourier series:

$$\theta_r(\psi) = \theta_{r0} - \theta_{rc}\cos\psi - \theta_{rs}\sin\psi \tag{109}$$

where θ_{r0} is the collective pitch, θ_{rc} is the lateral cyclic pitch, θ_{rs} is the longitudinal cyclic pitch, and ψ is the azimuthal angle of the blade.

You can either specify all three components of the pitch angle (fixed rotor) or you can select the trim algorithm to calculate them. When using the trim routine to calculate these angles, provide an initial guess for the angles. The values of the desired thrust and moments are used to calculate iteratively the new values of the pitch angles.

Specifying the Disk Stick Specification

To specify the disk stick specification:

1. Select the Virtual Disks > Virtual Disk 1 > Disk Stick Specification node.



2. In the *Disk Stick Specification - Properties* window, set the following properties:

Collective Pitch
: The collective part of the pitch angle θ_0 (see Eqn. (109)). The angle that you specify here is the same for all blades. When either Trim Thrust Only or Trim Thrust and Moments are selected for the Virtual Disk Trim Option, this property acts as the initial value for the iterative trim algorithm.

Longitudinal Cyclic Pitch
: The longitudinal cyclic part of the pitch angle θ_{1s} (see Eqn. (109)). When Trim Thrust and Moments are selected for the Virtual Disk Trim Option, this property acts as the initial value for the iterative trim algorithm.

Lateral Cyclic Pitch
: The lateral cyclic part of the pitch angle θ_{1c} (see Eqn. (109)). When Trim Thrust and Moments are selected for the Virtual Disk Trim Option, this property acts as the initial value for the iterative trim algorithm.

Virtual Disk Trim Option
: You specify the trimming type for the virtual disk trimming by selecting this option. For more information, see Specifying the Virtual Disk Trimming.

  No Trimming
  : No trimming is performed. All parts of the pitch angles are fixed and user-specified.

  Trim Thrust Only
  : The iterative trim algorithm adjusts the collective pitch angle only to achieve the desired thrust value. The three sub-nodes Trim Convergence Control, Target Thrust, and Target Pitch and Roll Moment are added to the simulation tree.

  Trim Thrust and Moments
  : The iterative trim algorithm adjusts all parts of the pitch angle (collective, longitudinal cyclic, and lateral cyclic pitch angle) to achieve the desired thrust and moment values. The three sub-nodes Trim Convergence Control, Target Thrust, and Target Pitch and Roll Moment are added to the simulation tree.

Specifying the Virtual Disk Trimming

The trim algorithm is an iterative procedure that adjusts the various parts of the pitch angle to achieve the user-specified target values of thrust and moments. You can follow the iterative development of the pitch angle components by using the Virtual Disk Angle report. For more information, see Reporting Virtual Disk Results.

You have three options available to specify the virtual disk trimming:

- No trimming at all.

  You do not invoke the trim algorithm. All parts of the pitch angle remain fixed as you specified.

- Trim the thrust only.

  Only the collective pitch angle is adjusted to achieve a specified target thrust value.

- Trim the thrust and the moments.

  All parts of the pitch angle (collective, longitudinal cyclic, and lateral cyclic pitch angle) are adjusted to achieve the desired thrust and moment values.

To specify that no trimming is performed:

- Select the Virtual Disks > Virtual Disk 1 > Disk Stick Specification node and set the Virtual Disk Trim Option to No Trimming.

To specify that only the thrust is trimmed:

- Select the Virtual Disks > Virtual Disk 1 > Disk Stick Specification node and set the Virtual Disk Trim Option to Trim Thrust Only.

To specify that the thrust and the moments are trimmed:

- Select the Virtual Disks > Virtual Disk 1 > Disk Stick Specification node and set the Virtual Disk Trim Option to Trim Thrust And Moments.

For the two trimming options, set the details for the trimming:

1. Select the Trim Convergence Control node.
2. In the *Trim Convergence Control - Properties* window, set the following properties:

Trim Under Relaxation Factor	The trim under-relaxation factor for the trim algorithm. The default is 0.5.

Trim Tolerance	The convergence tolerance for the trim algorithm. The default is 0.01.

3. Select the Target Thrust node.
4. In the *Target Thrust - Properties* window, set the following properties:

Target Rotor Thrust	The target value for the rotor thrust.

5. Select the Target Pitch and Roll Moment node.
6. In the *Target Pitch and Roll Moment - Properties* window, set the following properties:

Target Rotor Pitch Moment	The target value for the rotor pitch moment. If you selected the Trim Thrust Only option, this property has no influence on the trim algorithm. You can leave the default value.
Target Rotor Roll Moment	The target value for the rotor roll moment. If you selected the Trim Thrust Only option, this property has no influence on the trim algorithm. You can leave the default value.

Specifying the Virtual Disk Resolution

The cylindrical shape of the virtual disk for the blade element method must be associated with the cells of the underlying volume mesh. To perform this association, you specify the resolution for a two-dimensional interpolation grid in the radial and the azimuthal direction. One element of this virtual disk interpolation grid is called a “bucket”.

Set the virtual disk resolution such that each bucket is associated with at least one cell of the volume mesh. The virtual disk resolution must not be too fine. The resulting size of a bucket must approximately correspond to the cell size of the underlying volume mesh.

For a volume mesh with trimmed cells, you can approximate the required virtual disk resolution. Calculate the length of the virtual disk in radial direction $(r_{out} - r_{in})$ and in azimuthal direction along the inner radius $(2\pi r_{in})$. Divide the lengths by the specified resolution in the respective direction. The resulting value must approximately correspond to the volume mesh cell size in that area.

To specify the virtual disk resolution:

1. Select the Virtual Disks > Virtual Disk 1 > Virtual Disk Resolution node.



2. In the *Virtual Disk Resolution - Properties* window, set the following properties:

Radial Resolution	The resolution of the two-dimensional virtual disk interpolation grid in radial direction.
Azimuthal Resolution	The resolution of the two-dimensional virtual disk interpolation grid in azimuthal direction.

If you specified a virtual disk resolution that is too high, STAR-CCM+ issues an error message. Reduce the grid resolution until the error message disappears. For more details, see Step 1: Cell marking and the allocation to the blade elements/buckets based on the interpolation grid.

Setting the Source Convergence Control

To set the source convergence control:

1. Select the Virtual Disks > Virtual Disk 1 > Source Convergence Control node.



2. In the *Source Convergence Control - Properties* window, set the following properties:

Source Under Relaxation Factor	The under-relaxation factor for the source term.
Source Tolerance	The convergence tolerance for the source term.

1D Momentum Method

As part of the virtual disk model, the 1D Momentum method can be used to model the effects of a wind turbine. Wind turbines convert the wind energy into mechanical energy of rotation, which is then used to generate power. Among the wide variety of types of wind turbine, the 1D Momentum method is restricted to the so-called Horizontal Axis Wind Turbine (HAWT). HAWT uses aerodynamic lift to drive a rotor mounted on a horizontal axis. The rotational energy drives an electrical generator.

The 1D Momentum method is based upon one-dimensional momentum theory with wake rotation. The method imitates the effect of the presence of a wind turbine by adding a source term to the momentum equations. The source term accounts for both axial and tangential effects. For more

information on the formulation, see 1D Momentum Method.

The method requires the performance curves of a wind turbine, that is, the thrust coefficient and the power as functions of wind speed. An inflow velocity plane is used to compute the volume-averaged values for the velocity and the density of air. The inflow velocity plane is placed upstream of the wind turbine. This upstream velocity is used to extract the thrust coefficient and the power values from the performance curves and to calculate the source term.

The method can be used to study wake effects of a single turbine or of a complete wind park. A typical wind park design objective is to avoid placing a wind turbine in the wake of another. The 1D Momentum method is particularly useful for examining various wind park layouts in to order maximize energy production.

Setting Up and Using the 1D Momentum Method

Similar to the other virtual disk methods available, the setup of the virtual disk for the 1D momentum method does not require the creation of a separate region to which the method is then applied. Instead, you create the shape of the virtual disk by specifying the relevant parameters such as radius, thickness, and orientation to an existing mesh. To verify the dimensions and orientation of the virtual disk, you can visualize it by using virtual disk-specific field functions.

To set up and run the 1D momentum method:

1. Select the virtual disk model.

   See Selecting the Virtual Disk Model

2. Create a virtual disk.

   See Creating a Virtual Disk for the 1D Momentum Method

3. Visualize the virtual disk.

   Visualizing the Virtual Disk and the Inflow Velocity Plane

Creating a Virtual Disk for the 1D Momentum Method

To create a virtual disk:

1. Right-click the Virtual Disk > Virtual Disks node and select New.

A new child node called Virtual Disk 1 is added to the simulation tree.



2. Select the Virtual Disk 1 node and, in the *Properties* window, select the 1D Momentum Method from the drop-down list.

   The additional Power Curve, Disk Geometry, Inflow Velocity Plane, and RotationRate(n) nodes are added to the simulation tree.



A new local coordinate system for the virtual disk is created automatically. You can find it under the Tools > Coordinate Systems > Laboratory > Local Coordinate Systems > Virtual Disk 1-CSys 1 node. This coordinate system is oriented such that the positive z-axis is always normal to the plane of the disk.

Do not modify the local coordinate system of the disk here in the Local Coordinate Systems node. Modify the virtual disk origin and orientation only under the Virtual Disk node in the relevant physics continuum.



For more information, see Specifying the Disk Orientation.

3. Define the Virtual Disk by configuring the following nodes:

- Power Curve

  Set the wind turbine power curve as a function of wind speed.

  See Setting the Power Curve for the 1D Momentum Method

- Disk Geometry

  Define the disk geometry in terms of geometrical dimensions, location, and orientation.

  See Specifying the Disk Geometry for the Body Force Propeller and the 1D Momentum Methods

- Inflow Velocity Plane

The inflow velocity plane is located upstream of the virtual disk. For the **body force propeller method**, it is always defined in the positive direction of the disk normal. For the **1D momentum method**, it is always defined in the negative direction of the disk normal.

See Specifying the Inflow Velocity Plane

- RotationRate(n)

  See Specifying the Rotation Rate for the 1D Momentum Method

Setting the Power Curve for the 1D Momentum Method

To specify the propeller performance curve:

1. Select the Virtual Disk 1 > Power Curve node.
2. In the *Properties* window, select the Method that you would like to use to specify the propeller power curve. There are two options available:
   - Polynomial (Wind Speed)



   i. Select the Polynomial (Wind Speed) node and then click the ... (custom editor) next to the Power = f(Wind Speed) and Ct = f(Wind Speed).

   ii. In the *Polynomial(Windspeed)* dialog, enter the property values of the respective polynomial (see Setting Values Using a Polynomial).

- Table (Wind Speed)



i. Import the appropriate propeller performance curve table into the simulation (see Reading a Table Data File).

ii. Select the Table (Wind Speed) node. In the *Table (Wind Speed) - Properties* window, set the following properties:

Table	Select the power performance curve table that you imported into your simulation.
Table: Wind Speed	Select the column that contains the wind speed in m/s for the selected wind turbine power table.
Table: Power	Select the column that contains the power in Watts for the selected wind turbine power table.
Table: Ct	Select the column that contains the thrust coefficient (dimensionless) for the selected wind turbine power table.

Wind Turbine Cut-in and Cut-out Speed

At very low wind speeds, the wind exerts insufficient torque on the wind turbine blades to make them rotate. With increasing wind speed, the turbine begins to rotate and to generate power. The speed at which the turbine first starts to rotate and to generate power is called the **cut-in speed**.

If the wind speed increases beyond a critical velocity, the forces acting on the turbine blades get too large and the turbine can be damaged. Therefore, when this happens, a braking system enforces the wind turbine to stand still. The speed limit where this braking happens is called the **cut-out speed**.

The 1D momentum method takes the effects of the cut-in speed and the cut-out speed into account in the following conditions:

- The averaged air velocity on the inflow velocity plane is lower than the first wind speed entry of the power curve table.
- The averaged air velocity on the inflow velocity plane is greater than the last wind speed entry of the power curve table.

When either of the two cases occur, the 1D momentum source term is taken out of the computation of the momentum equations. This removal of the source term mimics the effect of the cut-in and the cut-out speed.

Specifying the Rotation Rate for the 1D Momentum Method

The rotation rate defines the operation point of the wind turbine.

To specify the rotation rate:

1. Select the Virtual Disk 1 > RotationRate (n) node.
2. In the *RotationRate (n) - Properties* window, set the Rotation Rate.



Visualizing the Virtual Disk and the Inflow Velocity Plane

You can visualize the previously defined virtual disk by use of a field function that become available when the virtual disk model is activated:

- Indicator of the virtual disk.

If you activated the Body Force Propeller method, you can visualize the inflow velocity plane in a similar fashion using the following field function:

- Inflow plane marker of virtual disk.

Each field function is then input into a derived part of the type threshold.

The following procedure is for visualizing the virtual disk, but the procedure for visualizing the inflow velocity plane is identical.

To visualize the virtual disk:

1. Create a scene of any type.

   For more details, see Working with Scenes.

2. To create a new derived part of the type threshold, right-click the Derived Parts node and select New Part > Threshold...
3. In the *Create Threshold* dialog:
   a. In the *Input Parts* box, select the region, in which the virtual disk resides.
   b. In the *Scalar* box, select Indicator of Virtual Disk 1.
   c. Set the Extraction Mode to All Between.
   d. In the *Scalar Range* box, set the Min and Max values both to `1`.
   e. In the *Display* box, select any display option.

f. Click Create and Close.



4. Add the newly created derived threshold part to the currently open scene in the *Graphics* window.
5. To view the cells that are part of the virtual disk, you can either initialize the solution manually or run the simulation.

For more information, see Initializing Manually and see Running the Simulation.

An example of a geometry scene showing the virtual disk with mesh lines displayed on it is shown below.



Reporting Virtual Disk Results

The following reports are available when the virtual disk model is activated:

- Virtual disk force report
- Virtual disk moment report
- Virtual disk angle report
- Virtual disk rotation rate report
- Virtual disk advance ratio report

Some virtual disk reports are not available for all methods. The availability of the reports is shown below, where (x) indicates that the report is available and reports a valid value. (-) indicates that the report is not applicable for the corresponding method and reports a value of 0:

Virtual Disk Report	Body Force Propeller Method	Blade Element Method	1D-Momentum
Force	x	x	x
Moment	x	x	x
Angle	-	x	-
Rotation Rate	x	-	-
Advance Ratio	x	-	-

The following procedure is for the virtual disk force report. The same procedure is used for all of the virtual disk reports.

To create a virtual disk force report:

1. Right-click the Reports node and select New Report > Virtual Disk Force
2. Select Reports > Virtual Disk Force 1 node



3. In the *Properties* window, set the Virtual Disk Reports properties.
4. To display the result, create a monitor and a plot.

   For more information, see Using Report-Based Monitors and see Working with Monitor Plots.

Virtual Disk Reports Properties

Virtual Disk Force Properties

Force Component Option	Specifies the force component to report. The force components are reported in the local coordinate system of the virtual disk.	
	Thrust (Z-Axis)	Specifies the force component in the direction of the z-axis. This direction is the direction in which thrust is produced.
	Pitch Force (X-Axis)	Specifies the force component in direction of the x-axis.
	Lateral Force (Y-Axis)	Specifies the force component in direction of the y-axis.
Virtual Disk	Specifies the virtual disk for which to report the force component.	
Units	Select the units of the force component.	

Virtual Disk Force Expert Properties

Force Type Option	Specifies the active or reactive force type to report.	
	Action of the Fluid on the Virtual Disk	The fluid dynamic forces acting on the virtual disk due to the action of the fluid. Usually selected to report the thrust produced by the virtual disk.
	Reaction of the Virtual Disk on the Fluid	The reaction of the virtual disk corresponds to the force added to the momentum equations in order to model the presence of the virtual disk. It is equal in magnitude, but opposite in sign to the Action of the Fluid on the Virtual Disk.

Virtual Disk Moment Properties

Moment Component Option	Specifies the moment component to report. The moment components are reported in the local coordinate system of the virtual disk.	
	Shaft Torque (Z-Axis)	Specifies the moment component in direction of the z-axis.
	Roll Moment (X-Axis)	Specifies the moment component in direction of the x-axis.
	Pitch Moment (Y-Axis)	Specifies the moment component in direction of the y-axis.
Virtual Disk	Specifies the virtual disk for which to report the moment component.	
Units	Select the units of the moment component.	

Virtual Disk Moment Expert Properties

Moment Type Option	Specifies the active or reactive moment type to report.	
	Action of the Fluid on the Virtual Disk	The moment acting on the virtual disk due to the action of the fluid.
	Reaction of the Virtual Disk on the Fluid	Is equal in magnitude, but opposite in sign to the Action of the Fluid on the Virtual Disk.

Virtual Disk Angle Properties

Angle Option	Specifies which angle to report.	
	Collective Pitch Angle	Specifies the collective pitch angle.
	Longitudinal Cyclic Pitch Angle	Specifies the longitudinal cyclic pitch angle.
	Lateral Cyclic Pitch Angle	Specifies the lateral cyclic pitch angle.
Virtual Disk	Specifies the virtual disk for which to report the angle.	
Units	Select the units of the angle.	

Virtual Disk Rotation Rate Properties

Use this report to obtain information about the rotation rate of the propeller when you select the Thrust or Torque option for Rotation Rate.

Virtual Disk	Select the virtual disk for which you want to report the rotation rate.
Units	Select the units of the rotation rate.

Virtual Disk Advance Ratio Properties

Use this report to obtain information about the advance ratio J of the propeller (see Eqn. (325)).

Virtual Disk	Select the virtual disk for which you want to report the advance ratio.
Units	The advance ratio is dimensionless.

Virtual Disk Field Functions

The following field functions are made available when the virtual disk model is activated.

Inflow Plane Marker of [Virtual Disk]	Cell marker to indicate which volume cell is identified as part of the inflow velocity plane. A value of 1 indicates that the cell is part of the inflow velocity plane. A value of 0 indicates that the cell is not part of the inflow velocity plane. This field function can be used as input to a derived part of the type threshold.
Indicator of [Virtual Disk]	Cell marker to indicate which volume cell is identified as part of the virtual disk. A value of 1 indicates that the cell is part of the virtual disk. A value of 0 indicates that the cell is not part of the virtual disk. This field function can be used as input to a derived part of the type threshold.
Normalized Span of [Virtual Disk]	Normalized radius (values from 0 through 1 where 1 is the outer radius of the virtual disk) of cells within the virtual disk with respect to the local coordinate system of the virtual disk.
Cell Azimuth of [Virtual Disk]	Gives the angular position of a cell within the virtual disk with respect to the local coordinate system of the virtual disk.
Virtual Disk Source Term of Virtual Disk 1	Shows the distribution of the virtual disk source term per unit volume over the disk. This function is only available when Display Source Term is activated on the virtual disk.

Confining Vorticity

Many areas of engineering today require tracking isolated vortices accurately over long distances. Such areas include:

- Aerospace—vortices of aircraft wing-tips, helicopter blades, and other components
- Submarines—tip vortices of various control surfaces
- Wind turbines—tip vortices and their environmental impact.

Unfortunately, due to the built-in dissipation of state-of-the-art numerical schemes that are used in modern CFD codes, any such simulation rapidly dissipates these vortices, unless a special treatment is applied.

One such treatment is the vorticity confinement concept that John Steinhoff *et al.* introduced ([49], [39]). Others have refined this concept [38]. It is implemented in STAR-CCM+ as the Vorticity Confinement model. This model is represented in the object tree by the Vorticity Confinement Model node, which has properties and is subject to common menu items.



An isolated vortex can be defined as the region in space with high values—and sometimes much higher values—of vorticity ($\mathrm{curl}(\mathbf{v})$), also referred to as "vortex tubes". Tracking isolated vortices over long distances is important because of the induced velocity they produce. This velocity in turn influences the flow-field and aerodynamic processes—sometimes at great distances from the place where these strong isolated vortices have

been produced. For example, helicopter blades generate strong blade-tip vortices which have to be either resolved or modeled (with a model such as the Vorticity Confinement model). Otherwise, the simulation of such a flow situation (especially of the rotor in ground effect) can be seriously inaccurate.

The current implementation of the Vorticity Confinement model in STAR-CCM+ is valid for both the Coupled Flow and Segregated Flow models in a low-speed flow regime ($\mathrm{Mach}_\infty \leq 0.3$). However, this model must be chosen with care. The $\mathbf{f}(d^*)$ factor (Eqn. (269)) is used to deactivate effectively the vorticity confinement force next to solid walls (for example inside a boundary layer) for viscous (laminar and turbulent) simulations. The model is deactivated next to the wall, as suggested by the authors [38], [39], [49]. Deactivating the model prevents the "vorticity confinement force" from affecting the boundary layer velocity distribution (and with it, factors such as skin friction).

This section includes instructions on selecting the Vorticity Confinement model, choosing the values for the Vorticity Confinement parameters, as well as its formulation and associated field functions.

Vorticity Confinement Model Properties

The Vorticity Confinement Model has the following properties:

Verbosity	Provides more output for this value while the simulation is running. This option can be useful for debugging problems when they occur.	
	Activated	Activates verbose output.
	Deactivated	Leaves output at its default.
Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration.	
	Activated	Retains temporary storage.
	Deactivated	Does not retain temporary storage.

Expert Properties

Non-Dimensional Wall Distance Coefficient	Prevents the vorticity confinement term from acting inside boundary layers (d^*_{cutoff} in Eqn. (269)).
Model Coefficient	Tunes the strength of vorticity confinement (ε_0 in Eqn. (265)).

Selecting the Vorticity Confinement Model

To select the Vorticity Confinement model, first select the Segregated Flow model or select the Coupled Flow model.

The Vorticity Confinement Model radio button becomes available for selection under *Optional Models*.



Choosing the Vorticity Confinement Model Parameter Values

To choose the values for the non-dimensional wall-distance coefficient d^*_{cutoff} and model coefficient ε_0 that are appropriate for your simulation:

- Create a scalar scene and plot the field function Vorticity Confinement Non-Dimensional Wall-Distance on the parts of interest.

  In the following screenshot, this field function is plotted on the region of interest: the wing surface.



- Set the non-dimensional wall-distance coefficient d^*_{cutoff} to the value that is appropriate for your case.

  In the above example, the computed vorticity confinement force is switched off near the wing surface. Therefore, d^*_{cutoff} is chosen to be 65.

- Run the case.
- Plot the Vorticity Confinement Specific Force field function and verify that:
  - The Vorticity Confinement specific force is zero inside the boundary layer.
  - The Vorticity Confinement specific force is sufficient to preserve the vortex for a sufficiently long distance. The actual value of this specific force depends on the simulation.

The Vorticity Confinement specific force plot for the above example is shown below.



Note: To increase the Vorticity Confinement specific force, you can increase the model coefficient ε_0 .

Solving the Flow Adjoint

The adjoint method is an efficient means to predict the influence of many design parameters and physical inputs on some engineering quantity of interest, that is, on the engineering objective of the simulation. In other words, it provides the sensitivity of the objective (output) with respect to the design variables (input).

Examples of the types of problems to which the adjoint method is applicable are:

- What effect does the shape of a duct (input variable) have on the pressure drop (objective)?
- What is the influence of inlet conditions (input variable) on flow uniformity at the outlet (objective)?
- What areas of the airfoil surface (input variable) have the biggest impact on lift and drag (objectives)?

The advantage of the adjoint method is that the computational cost for obtaining the sensitivities of an objective does not increase with an increasing number of design variables. The computational cost is essentially independent of the number of design variables because the adjoint method requires only a single flow solution and a single adjoint solution for any number of design variables.

The flow adjoint equations form a linear system that is typically solved by means of an iterative defect-correction algorithm. The cost of solving the linear system of equations is similar to solving the primal flow solution in terms of iterations and computational time.

The workflow for performing an adjoint flow analysis is typically the following:

- Set up, run and visualize the flow problem as usual.
- Choose the engineering objective for the problem at hand.
- Set up and run the flow adjoint analysis with regard to the chosen objective.
- Visualize the adjoint solution to obtain sensitivity information with regard to the chosen objective.

There are currently five different types of engineering objectives (cost functions) available: pressure drop, uniformity deviation, force, moment, and tumble/swirl.

See also:

Setting up an Adjoint Flow Analysis
Setting up a Coupled Flow Analysis Compatible with Adjoint
Enhancing the Coupled Flow Solver Convergence Behavior
Selecting the Adjoint Flow Model
Adjoint Cost Functions
Creating the Adjoint Cost Functions
Running the Adjoint Flow Solver
Troubleshooting the Adjoint Flow Solver
Computing the Shape Sensitivity
Setting up a Shape Optimization
Visualizing the Adjoint Flow Solution
Visualizing the Mesh Sensitivity
Adjoint Field Functions

Setting up an Adjoint Flow Analysis

An adjoint flow analysis is performed in two major steps. First, you simulate the steady flow field using the coupled flow solver. This is referred to as the "primal flow field solution". In a following step, you set up and run the adjoint analysis on the previously obtained flow solution. The robustness of the adjoint solution relies heavily on the fidelity of the primal flow field. This means that the primal solution should be well converged, ideally down to machine precision. To enhance the convergence of the primal solution, create a mesh of good quality before you start the simulation.

To set up an adjoint flow analysis:

1. Create a suitable, high-quality volume mesh.

   To verify that you have created a mesh of good quality, it is good practice that you check the mesh quality and improve it if necessary.

   For more information, see Checking the Volume Mesh

   One measure that is particularly useful for improving the quality of polyhedral meshes, is the mesh optimizer. To improve the quality of a polyhedral mesh, increase the number of optimization cycles and increase the quality threshold. The mesh optimization leads to a longer meshing time.

   For more information, see Volume Mesher Controls and Values

2. Set up and run a steady-state flow analysis using models and features compatible with the adjoint solver.

   Ideally, converge the simulation down to machine precision.

   For more information, see Setting up a Coupled Flow Analysis Compatible with Adjoint

   For general considerations regarding a coupled flow analysis, see Modeling Flow and Energy Using a Coupled Approach

   If necessary, enhance the convergence of the simulation by modifying the coupled flow solver numerical parameters.

   For more information, see Enhancing the Coupled Flow Solver Convergence Behavior
3. Select the adjoint flow model.

   For more information, see Selecting the Adjoint Flow Model
4. Create one or several adjoint cost functions.

   For more information, see Creating the Adjoint Cost Functions.
5. Run the adjoint flow analysis using the adjoint flow solver.

   Ideally, converge the adjoint simulation down to machine precision.

   For more information, see Running the Adjoint Flow Solver
6. Set up a shape optimization.

   For more information, see Setting up a Shape Optimization.
7. Visualize the results using the adjoint field functions

   For more information, see Visualizing the Adjoint Flow Solution and Visualizing the Mesh Sensitivity.

For a complete setup of such cases, see the tutorials Adjoint Flow Solver: Flow through an S-Bend and Adjoint Flow Solver: External Flow over a Dual Element Wing.

Setting up a Coupled Flow Analysis Compatible with Adjoint

The computation of the adjoint is based on the solution of the flow. Before you can actually perform the adjoint flow analysis, you need to set up and run a steady-state flow analysis using the coupled flow solver. Note, however, the adjoint solver is compatible with a specific set of the available physical models and features.

It is strongly recommended that you use a double precision version of STAR-CCM+ to perform the simulation.

To set up a coupled flow analysis compatible with adjoint:

1. Right-click the Continua > Physics 1 > Models node and select Select models...
2. Select either the Two Dimensional or the Three Dimensional option from the *Space* group box.

   Adjoint analysis is only compatible with these space models.
3. Select Steady from the *Time* group box.

   Adjoint analysis is only compatible with this time model.
4. Select Gas or Liquid from the *Material* group box.

   Adjoint analysis is only compatible with these material models.
5. Select Coupled Flow from the *Flow* group box.

   Adjoint analysis is only compatible with this flow model.
6. Select Ideal Gas or Constant Density from the *Equation of State* group box.

   Adjoint analysis is only valid with these equation of state models.
7. Select any option from the *Viscous Regime* group box.

   Adjoint analysis is only valid for an inviscid or laminar viscous regime. If the turbulent viscous regime is selected, the turbulent viscosity is treated as a passive quantity.

Note: Adjoint analysis is only valid in the Lab Reference Frame. For porous regions, porosity and porous resistances are treated as passive quantities.

Enhancing the Coupled Flow Solver Convergence Behavior

For a standard flow analysis, the residuals of the flow variables are not required to converge to machine precision. However, to ensure that the adjoint solver converges in a robust manner, it is important to obtain a certain level of convergence in the primal solution.

With the default STAR-CCM+ numerical parameter settings, a satisfactory level of convergence can be reached for a standard flow analysis. For an adjoint flow analysis some additional effort is usually required.

There are some numerical measures that can be taken to enhance the convergence behavior of the coupled flow solver. These measures are not new and can be employed to improve the behavior of the primal flow analysis independent of any subsequent adjoint analysis.

- Grid Sequencing Initialization

A reasonable way of improving the convergence behavior is to start the simulation with a good initialization of the flow field.

For details on how to employ the Grid Sequencing Initialization, see Using Grid-Sequencing Initialization and Setting Grid Sequencing Parameters.

- Continuity Convergence Accelerator (CCA)

  This feature can improve convergence significantly for both constant density and compressible flows.

  For details on how to employ the CCA, see Setting Convergence Accelerators

- Expert Driver

  The Expert Driver is an automated convergence control for steady-state simulations. It can improve the robustness of the coupled flow solver and increase the speed of convergence.

  For more information, see Using the Expert Driver

- Freezing the solution of the turbulence equations.

Freezing the Solution of the Turbulence Equations for Adjoint Flow Analysis

The solution of the turbulence equations can be frozen after a good level of convergence has been reached. This is a legitimate measure because sensitivity with respect to the turbulence is not being considered in the subsequent adjoint analysis.

To freeze the solution of the turbulence equations:

1. Click the Solvers node of the turbulence model that you have selected.

2. In the *Properties* window, activate Solver Frozen.



3. Continue the run to converge the residuals of the remaining variables as much as possible.

Selecting the Adjoint Flow Model

After you have run the steady-state flow analysis, activate the adjoint flow model as an optional physical model from the *Physics Model Selection* dialog of the relevant physics continuum.

To select the adjoint flow model:

1. Right-click the Continua > Physics 1 > Models node and select Select models...
2. In the *Physics Model Selection* dialog, select Adjoint Flow from the *Optional Models* group box.

The Mesh Deformation model is automatically selected when the Adjoint Flow model is selected.



3. Click Close.

The Adjoint Flow node and the Mesh Deformation node appear under the Models node in the simulation tree.



4. In the *Adjoint Flow - Properties* window, set the discretization scheme to use for the adjoint flow model.

Adjoint Discretization	The discretization scheme that is used for the adjoint flow model.	
	2nd-order	This is the default setting. Use the second order discretization scheme for best accuracy of the adjoint solution.
	1st-order	Use the first order discretization scheme in the event the adjoint solver does not converge with the default second order setting. This makes the adjoint solver less accurate but more robust.

Adjoint Cost Functions

An adjoint cost function corresponds to the engineering objective that is examined during the adjoint flow analysis. One engineering objective might be to reduce the pressure drop in a piping system. Another example of an engineering observable quantity is the reduction of drag of the flow around a car.

The cost functions are differentiated to form the right-hand side of the adjoint flow of equations. The adjoint solver solves the system of differential equations with regard to the cost function specified.

Currently, the following cost functions are available:

- Force Cost Function

  The force cost function is the force on a specified set of boundary surfaces in a specified direction. It is linked to a force report. The report can be executed to give the force value.

- Moment Cost Function

  The moment cost function is the moment on a specified set of boundary surfaces with a specified direction and origin. It is linked to a moment report. The report can be executed to give the force value.

- Pressure Drop Cost Function

  The pressure drop cost function is the mass-flow averaged total pressure between two boundary surfaces or between two interfaces representing the high pressure and low pressure boundaries of the system. It is linked to a pressure drop report. The report can be executed to give the pressure drop value.

$$dp = \sum_{\text{High}} \left[\frac{\sum_{\text{f}} \dot{m}_{\text{f}} P_{t_{\text{f}}}}{\sum_{\text{f}} \dot{m}_{\text{f}}} \right] - \sum_{\text{Low}} \left[\frac{\sum_{\text{f}} \dot{m}_{\text{f}} P_{t_{\text{f}}}}{\sum_{\text{f}} \dot{m}_{\text{f}}} \right] \tag{110}$$

- Uniformity Deviation Cost Function

  The uniformity cost function is the deviation of the local normal velocity from the mass-flow averaged value on a specified set of boundary surfaces or on interfaces. It is linked to a uniformity deviation report. The report can be executed to give the uniformity deviation value. The uniformity deviation of the velocity normal to the surface (at each face $u_{\text{f}} = \mathbf{u}_{\text{f}} \cdot \mathbf{n}_{\text{f}}$) is:

$$\gamma = \frac{\sigma}{\bar{u}} \tag{111}$$

  where the variance of normal speed is:

$$\sigma^2 = \frac{\sum_{\text{f}} (u_{\text{f}} - \bar{u})^2 \|\mathbf{a}_{\text{f}}\|}{\sum_{\text{f}} \|\mathbf{a}_{\text{f}}\|}$$

and the average normal speed is:

$$\bar{u} = \frac{\sum_{\mathrm{f}} u_{\mathrm{f}} \cdot \mathbf{a}_{\mathrm{f}}}{\sum_{\mathrm{f}} \|\mathbf{a}_{\mathrm{f}}\|}$$

- Tumble and Swirl Cost Function

  The tumble and swirl cost function is the swirl/tumble ratio (dimensionless) defined by:

$$S_x = \frac{M_x}{I_x} = \frac{\sum_{Cells} \rho_i V_i [(Y_i - Y_m) w_i - (Z_i - Z_m) v_i]}{2\pi \frac{N}{60} \sum_{Cells} \rho_i V_i [(Y_i - Y_m)^2 + (Z_i - Z_m)^2]} \tag{112}$$

$$S_y = \frac{M_y}{I_y} = \frac{\sum_{Cells} \rho_i V_i [(Z_i - Z_m) u_i - (X_i - X_m) w_i]}{2\pi \frac{N}{60} \sum_{Cells} \rho_i V_i [(X_i - X_m)^2 + (Z_i - Z_m)^2]} \tag{113}$$

$$S_z = \frac{M_z}{I_z} = \frac{\sum_{Cells} \rho_i V_i [(X_i - X_m) v_i - (Y_i - Y_m) u_i]}{2\pi \frac{N}{60} \sum_{Cells} \rho_i V_i [(Y_i - Y_m)^2 + (X_i - X_m)^2]} \tag{114}$$

where:

X_i, Y_i, Z_i	Centroid coordinate of cell i (m)
X_m, Y_i, Z_i	Center of mass/volume/origin of the cylinder (m)
u_i, v_i, w_i	Velocity at cell i (m/s)
N	Rotational speed (rev/min)
M_x, M_y, M_z	Angular momentum about the x-, y-, and z-axes (kg m^2/s)
I_x, I_y, I_z	Mass moment of inertia about the x-, y-, and z-axes (kg m^2/s)

- Expression Cost Function
- Mass Flow Cost Function

The mass flow cost function can be used to estimate the sensitivity of mass flow rate through boundaries. It is linked to a mass flow report.

Creating the Adjoint Cost Functions

You define at least one cost function before you start the adjoint analysis. You can also create additional cost functions within the same run.

All cost functions require prior creation of respective reports. These reports serve as input for the cost functions.

Before you create a cost function:

1. Create a report of the respective type. For example, for the force cost function you must first create a force report.



For more information, see Reporting Results.

2. In the *Properties* window of the cost function report, specify the appropriate cost function report properties.

   For more information on cost function report properties, see:

Force	Force
Moment	Moment

Pressure Drop	Pressure Drop
Tumble/Swirl	Tumble/Swirl Report
Uniformity Deviation	Uniformity Deviation Report
Expression	Expression Report The restrictions for a report expression that is associated with an expression cost function are as follows: • The report expression must only reference reports for which there is a corresponding selectable cost function. • The supported operators and standard mathematical functions are + - * / and pow(a,b), sqrt(a), sin(a), cos(a), tan(a).
Mass Flow	Mass Flow

To create cost functions:

1. Right-click the Adjoint Cost Functions node, then click New and select the cost function that you want to create.



A new node for the cost function is created.

2. Select the newly created cost function node.



3. In the *Properties* window, select the required report that you created previously.

Running the Adjoint Flow Solver

The adjoint flow solver solves the linear system of equations for the flow adjoint. The flow adjoint equations are solved by means of an iterative defect-correction algorithm (see Eqn. (302)). Alternatively, a GMRES Krylov solver is available.

After you have activated the adjoint flow model and set up the adjoint cost functions, you can run the adjoint flow solver. The adjoint flow solver is started and controlled separately from the coupled flow solver. It is not recommend that you run the coupled flow solver while the adjoint flow model is active.

To run the adjoint flow solver:

1. Select the Solvers > Adjoint Flow node.



2. In the *Adjoint Flow - Properties* window, set the following numerical properties of the adjoint solver.

Adjoint Driver Selection	Select the solver to use to obtain the adjoint solution.	
	Defect Correction	This is the default option.
	Restarted GMRES	Use the GMRES Krylov solver for "stubborn" cases that are difficult to converge. It accelerates convergence of the adjoint flow algorithm. For more information, see GMRES Krylov Solver.
Courant Number	Adjust the Courant (CFL) number to under-relax and control convergence similarly to standard coupled flow analysis.	

3. The AMG linear solver controls the solution of the linear system of equations during the adjoint flow iterations. The AMG Linear Solver node is configured as for any standard flow analysis (see Algebraic Multigrid). Usually, no changes to the default settings are required.
4. Right-click the Adjoint Flow node and select Step or Run.

GMRES Krylov Solver

For the solution of the linear system of equations for the flow adjoint, you can use the GMRES Krylov solver as an alternative to the default defect correction scheme. Using this solver adds to computation time and memory overhead.

When you have selected the Restarted GMRES option in the *Adjoint Flow - Properties* window, a new GMRES(m) Driver node is added to the simulation tree. Use this node to configure this solver type.

The GMRES(m) Driver node has the following properties:

Verbose	You get additional information on the GMRES solver in the *Output* window.
Krylov Space Dimension	Specify the number of orthogonal Krylov space vectors to use. Increase this number to enhance GMRES-driver convergence and robustness. However, increasing this number increases the memory requirement.
Relative Tolerance	Relative tolerance is used as an additional condition for the GMRES-driver inner-loop convergence and restart.

Troubleshooting the Adjoint Flow Solver

The convergence rate of the linear adjoint flow solution is usually similar to that of the primal coupled flow solution. It takes roughly as many iterations to converge at a similar computational cost.

Similar to a standard flow simulation, it is possible that the iterative process does not converge or even diverges.

One reason for difficult convergence behavior can be the mesh quality. Even though the coupled primal flow solution converges, the mesh quality might not be sufficient for the adjoint flow solution to converge.

- If you suspect that this could be the case, then perform the mesh quality checks as described in Checking the Volume Mesh. If necessary, improve the cell quality of the volume mesh and re-run the primal flow solution.

There are some numerical measures that you can take to improve the adjoint flow solver convergence.

- Reduce the Courant number of the adjoint flow solver.

- Change the default Adjoint Driver Selection to Restarted GMRES. This solver increases the robustness of the adjoint flow algorithm at the cost of increased computational time and memory overhead. For more information, see GMRES Krylov Solver.
- Reduce the order of the adjoint flow model discretization scheme. This option is available for both the Defect Correction as well as the Restarted GMRES solver.

Setting up a Shape Optimization

Shape sensitivity data is the product of the sensitivity of a specific cost function with respect to the mesh vertex position and the sensitivity of the mesh position with respect to either displacement boundaries or control point positions. This data represents to what degree a change in a boundary displacement vertex or design point position will affect a given cost function.

The calculation of the shape sensitivity is based on the previously obtained adjoint flow solution. By adding or subtracting a small fraction of the shape sensitivity to either the displacement boundary vertex positions or the control point final positions, the mesh can be deformed to increase or reduce the cost function, respectively. With the new mesh, you can rerun the primal and adjoint flow solutions in sequence to analyze the solution again.

To set up a shape optimization within the adjoint framework:

1. Specify the initial design point positions.

   See Creating Initial Control Point Positions.
2. Specify the mesh deformation boundary conditions.

   See Specifying Mesh Deformation Boundary Conditions
3. Compute the shape sensitivity with respect to the initial design points.

   See Computing the Shape Sensitivity.
4. Compute new design point positions that are based on the previously computed shape sensitivities.

   See Obtaining New Control Point Positions
5. Compute the mesh deformation and deform the mesh accordingly.

   See Computing the Mesh Deformation.
6. Rerun the primal and adjoint flow solutions.

   See Setting up an Adjoint Flow Analysis.

Creating Initial Control Point Positions

You create the initial control point positions for the shape sensitivity and mesh deformation calculations required by using the *Create Control Points* dialog. Once you create the control points, they are automatically associated with the morpher. You specify the movement of the control points through the Morpher: Displacement profile under the Control Points node.

Note: To compute the shape sensitivity and the mesh deformation directly for the vertices of the boundary, use the morpher boundary type Displacement. For more information, see Specifying Mesh Deformation Boundary Conditions.

When you position the control points around a boundary to be deformed, locate them close enough to have a sufficiently strong influence on it, but far enough away from the boundary to ensure smooth deformations to it. For an airfoil, for example, a good starting point is to place 10 to 20 control points strategically at a distance of 20% of its chord length from the boundary.

To create the initial control point positions:

1. Right-click the Control Points node and select New Control Points, then one of the following options:

- Table...
- Lattice...
- Line...
- Part...

2. In the respective *Create Control Points* dialogs, make the appropriate settings:
   - See Create Control Points from Table Dialog
   - See Create Control Points from Lattice Dialog
   - See Create Control Points from Line Dialog
   - See Create Control Points from Part Dialog
3. Click Create, then Close.

   A new Control Point Region node is created.
4. Select the Control Point Region node and set the Control Point Region Properties.
5. Select the Control Point Region > Physics Values > Morpher: Displacement node and set the appropriate profile method (see Using Vector Profiles).

Create Control Points from Table Dialog

Source Table	Specifies the table from which to extract the control point coordinates with respect to the laboratory coordinate system.
X Data	The column that contains X coordinates.
Y Data	The column that contains Y coordinates.
Z Data	The column that contains Z coordinates.

Create Control Points from Lattice Dialog

Origin	Specifies the origin of the three-dimensional lattice with respect to the laboratory coordinate system, in X, Y, Z coordinates.
Size	Specifies the width, length, and height of the lattice with respect to the origin. Define using dimensions in coordinate directions X, Y, and Z.
Rotation Axis	Specifies the axis of rotation in the laboratory coordinate system, in X, Y, and Z coordinates.
Rotation Angle	Specifies the angle by which to rotate the lattice around the Rotation Axis.
Number in X	The number of control points for the lattice in the X direction.
Number in Y	The number of control points for the lattice in the Y direction.
Number in Z	The number of control points for the lattice in the Z direction.

Create Control Points from Line Dialog

Point 1	Defines X, Y, and Z coordinates of the line start point within the chosen Coordinate System.
Point 2	Defines X, Y, and Z coordinates of the line end point within the chosen Coordinate System.

Coordinate System	Selects the reference coordinate system in which the line start and end points are defined.
Number	The number of control points to create along the defined line.

Create Control Points from Part Dialog

Input Part	Specifies the input part relative to which a three-dimensional cloud of control points is created. The control points are created based on the vertices of that part surface.
Representation	Specifies from which representation the control points are extracted. Input Part can be represented with different meshes or triangulations, so you must specify the appropriate representation from which to extract the control points. The default representation works for most cases. If you create control points along feature curves, you must specify a geometry representation to extract meaningful control points.
Target Control Point Spacing	The approximate distance by which the control points are separated from each other. Adjusts the density of the control point cloud.
Maximum Number of Points	Maximum number of control points to create in the control point region. If you specify a Target Control Point Spacing that yields too many points, a random mask is applied to the control points so that only the Maximum Number of Points is added to the newly created control point region.
Offset from Part Surface	When activated, control points are displaced from the part surface by the given distance.
	Surface Offset: The distance by which the control points are displaced along the normal of the input part surface.

Control Point Region Properties

Index	(Read-only) Unique number to identify the control point region.	
Active	Activated	The control point region is used for the shape sensitivity computation. Associates the control points automatically with the morpher.
	Deactivated	The control point region is not used for the shape sensitivity computation.
Points	(Read-only) The number of control points of this control point region.	

Specifying Mesh Deformation Boundary Conditions

The morpher is used to update and deform the mesh according to the new design point positions (design point displacement) or the vertex positions at a displacement boundary that are computed based on the mesh adjoint.

For more details on how the morpher works, see Working with Morphing.

At each boundary of your computational domain, specify the conditions of how they are morphed. Consider, for example, the simulation of an airfoil within a wind tunnel. Only the airfoil is subject to adjoint shape optimization with the vertices at the boundary being morphed. The surrounding wind tunnel boundaries remain fixed.

To specify mesh deformation boundary conditions:

1. Select the Regions > [Region] > Boundaries > [Boundary] > Physics Conditions > Morpher node.



2. In the *Morpher - Properties* window, set the appropriate properties.

   See Mesh Deformation Morpher Properties

Mesh Deformation Morpher Properties

Method

Defines how the vertices at the boundary are morphed. For more details on the morpher boundary methods, see Understanding Morpher Boundary Methods.

- **Fixed** — The boundary remains fixed. The points on this boundary are treated as design points with zero displacement. Selecting this option adds the Morpher: Thin Factor node to the Physics Values node of the boundary. Specify a value less than 1 for the Vertex Thin Factor to override auto-thinning from the mesh deformation solver.
- **In Plane** — Points on this boundary are allowed to slide in the plane of the boundary. Selecting this option adds the Morpher Plane Options node to the Physics Values node of the boundary.

Floating	Points on this boundary move according to the computed displacement of the design points.
Displacement	Vertices on this boundary are collectively viewed as design points. They move according to the computed displacement of the mesh deformation solver. Selecting this option adds the Morpher: Thin Factor node to the Physics Values node of the boundary. Specify a value less than 1 for the Vertex Thin Factor to override auto-thinning from the mesh deformation solver.

Mesh Deformation Morpher: Thin Factor Properties

Damping Factor	Displacements at the boundary are damped smoothly to zero starting at the specified distance. This factor is value L_p in Eqn. (15). The default value of `0.0` means that the distances are calculated automatically.
Continuity Level	Specifies which damping function to use in smoothing vertex displacements approaching the plane. C0 — The displacement is damped using Eqn. (13). Only this option is valid when evaluating the shape sensitivity. C2 — The displacement is damped using Eqn. (14).
Ignore Factor	This factor is multiplied by the Damping Factor to calculate a finite distance. Control Points closer than the specified distance to this plane are given special treatment. A value of zero is acceptable.

Computing the Shape Sensitivity

The shape sensitivity gives the sensitivity of the cost function with respect to the control points or displacement boundaries. The evaluation of the shape sensitivity is a non-iterative two-step process. The mesh sensitivity is first computed and then, if the Mesh Deformation model is available, the shape sensitivity is solved.

To compute the shape sensitivity:

- Right-click the Solvers > Adjoint Mesh node and select Compute Mesh Sensitivity.



To visualize the results of the computation of the shape sensitivity:

- Visualize the mesh adjoint at displacement boundaries. For more information, see Visualizing the Mesh Adjoint.
- Visualize the mesh sensitivity at all the vertices. For more information, see Visualizing the Mesh Sensitivity. To make this option available, select the Adjoint Mesh node and activate Temporary Storage Retained.

If auto-thinning is disabled and the number of design points exceeds the maximum number of recommended design points as specified in the Mesh Deformation node, you receive a warning message.

To remedy this warning, you have the following options:

- Enable auto-thinning.

  See Computing the Mesh Deformation.
- Decrease the Thin Factor at the Fixed Morpher boundaries.

  See Specifying Mesh Deformation Boundary Conditions.

- Decrease the number of design points in the specified design point table.

  See Creating Initial Control Point Positions.

- Increase the maximum number of recommended design points and ignore the warning.

  See Computing the Mesh Deformation.

Obtaining New Control Point Positions

Based on the computed shape sensitivity at the control points, you compute the new positions of the control points. This computation of the displacement of the control points from their initial position is typically performed by using an external (third-party) optimization program such as LBFGS. You provide tables containing the control point positions and the shape sensitivity at these control points to the external optimization program. You create an internal table that uses the previously created control points as inputs. For more information, see Creating Initial Control Point Positions.

To compute the new control point positions:

1. Right-click the Tools > Tables node and select New Table > XYZ Internal Table.
2. Select the XYZ Internal Table node:
   a. Set Parts to control points that you created previously.
   b. Set Scalars to Adjoint of [Cost Function] w.r.t. Position.
3. Right-click the XYZ Internal Table node and select Extract.
4. Right-click the XYZ Internal Table node and select Export...

   You export the sensitivity data associated with the control point locations to a .csv file.
5. Import the shape sensitivity table into the optimization program of your choice.
6. With the information from the imported tables, compute the control point displacement using the optimization program.
7. From the optimization program, save the control point displacements to a `[Name of control point table].csv` file.
8. Right-click the Tools > Tables node and select New Table > File Table > [Name of control point table].csv.

9. Right-click the Control Points > [Control Region] > Physics Values > Morpher: Displacement node and set Method to Table(x,y,z).
10. Select the Morpher: Displacement > Table (x,y,z) node and set the table properties for the [Name of control point table].

    The morpher uses the new control point positions to deform the mesh.

Computing the Mesh Deformation

To compute the mesh deformation:

1. Select the Solvers > Mesh Deformation node.



2. In the *Mesh Deformation - Properties* window, set the appropriate properties.

   See Mesh Deformation Properties.
3. Right-click the Solvers > Mesh Deformation node and select Deform Mesh.

If the number of design points exceeds the maximum number of recommended design points as specified in the Mesh Deformation node, you receive a warning message. To remedy this warning, you have the following options.

You can visualize the mesh deformation in any suitable scene showing the volume mesh of the morphed geometry.

To undo the mesh deformation:

1. Select *Solution* > Clear Solution...
2. In the *Clear Solutions* dialog, select Reset Mesh.



3. Click OK.

Mesh Deformation Properties

Verbosity	Provides more information on the mesh deformation operation to the *Output* window. None — No additional output. Low — Additional information about deformation point positions. High — Additional information about numerical parameters of mesh deformation morphing operation.
Morpher Tolerance	Non-dimensional parameter that is used to specify acceptable accuracy in morphing. This parameter is multiplied by the longest edge length of the bounding box enclosing the morphing region.
Maximum Recommended Design Points	The recommended upper limit for the number of fixed and displacement design points. The default is 1000. A recommended maximum number of design points is specified, because an excessive high number of design points can become expensive computationally. When the specified limit is exceeded, a warning is issued before the execution of a Compute Mesh Sensitivity and Deform Mesh operation if auto-thinning is disabled. For more information, see Computing the Shape Sensitivity and Computing the Mesh Deformation.

Enable Automatic Thinning	Activated by default. Automatically thins out fixed and displacement boundary vertices for mesh deformation morphing.

Obtaining Boundary Condition Value Sensitivity

You can obtain information about the sensitivity of the cost functions, for example pressure drop or drag, with respect to boundary input parameters such as free-stream Mach number, stagnation temperature, outflow pressure, and other imposed boundary conditions. This information is available when you execute an adjoint-specific report where you specify the boundary parameter with respect to the cost functions for which the sensitivity is computed.

Before you begin to compute the boundary condition value sensitivity:

- Set up at least one cost function.

  See Creating the Adjoint Cost Functions.
- Complete at least one adjoint iteration step.

  See Running the Adjoint Flow Solver.

To obtain boundary condition value sensitivity:

1. Right-click the Reports node and select New Report > Boundary Parameter Sensitivity.
2. Select the Boundary Parameter Sensitivity 1 node.
3. In the *Boundary Parameter Sensitivity 1 - Properties* window, set the following parameters:

Adjoint Cost Function	Select the adjoint cost function for which you want to obtain the boundary condition value sensitivity.
Boundary Parameter	Select the boundary parameter with respect to which you want to obtain the sensitivity.
Parts	Select the input parts onto which you want to evaluate the boundary parameter sensitivity.
Units	Select the units of the previously selected Adjoint Cost Function.

Representation Select a representation. Only the Volume Mesh representation is a valid choice.

Smooth Values When you activate this property, vertex interpolated values are used.

Visualizing the Adjoint Flow Solution

Specific adjoint field functions are available to visualize the adjoint flow solution.

To visualize the adjoint flow solution:

1. Create a scalar scene.

   see Creating a Scalar Scene

2. Click the *scene/plot* button.
3. In the scalar scene, select the Scalar 1 > Scalar Field node.
4. In the *Properties* window, click on the field next to Function. In the drop-down list, select Adjoint. For each of the previously specified cost functions, you can select a menu item.



The adjoint <cost function> field functions are available with regard to continuity, energy, X-, Y- and Z-momentum.

The adjoint solution is displayed in the currently open scalar scene in the *Graphics* window.

An example is shown below:



You can see the adjoint of the pressure drop with regard to continuity. This means that the plot shows you the sensitivity of the cost function, in this case the pressure drop, with regard to mass. (The continuity equation corresponds to the mass conservation equation.) According to Eqn. (301), the flow adjoint represents the sensitivity of the cost function with regard to the flow residual. Consequently, the adjoint with regard to continuity residual indicates that directing mass away from areas with positive adjoint values decreases the cost function, in this case, the pressure drop. While directing mass towards areas with positive adjoint values, increases the cost function. Red areas in the plot indicate that directing mass away from these areas has the most effect on lessening the pressure drop. Blue areas in the adjoint plot also indicate a high sensitivity of the cost function with regard to mass. Directing mass away from areas with negative adjoint values increases the cost function. You can apply this interpretation of results to the other residuals, that is to the momentum and energy residuals, as well.

Visualizing the Mesh Adjoint

Specific sensitivity field functions are available to visualize the mesh adjoint.

To visualize the mesh adjoint:

1. Create a vector scene.

   For more information, see Creating a Vector Scene.

2. Click the *scene/plot* button.
3. In the vector scene, select the Vector 1 > Vector Field node.
4. In the *Properties* window, set Function to the appropriate Adjoint cost function with respect to position.



The mesh adjoint is displayed in the currently open vector scene in the *Graphics* window.

Visualizing the Mesh Sensitivity

Specific sensitivity field functions are available to visualize the mesh sensitivity. To make the field functions available, select the Adjoint Mesh node and activate Temporary Storage Retained.

To visualize the mesh sensitivity:

1. Create a vector scene.

   For more information, see Creating a Vector Scene.

2. Click the *scene/plot* button.
3. In the vector scene, select the Vector 1 > Vector Field node.
4. In the *Properties* window, click on the field next to Function. In the drop-down list, select Sensitivity. For each of the previously specified cost functions, you can select a menu item. Select the adjoint <cost function> field functions with regard to position.



The mesh sensitivity is displayed in the currently open vector scene in the *Graphics* window.

An example is shown below.



You can see the sensitivity of the cost function, in this case, the pressure drop with regard to position. Here, position denotes the coordinates of the vertices of the volume mesh. Red or blue areas in the plot represent large positive and negative sensitivity values. Both areas indicate that the respective cost function is particularly sensitive to that given vertex coordinate position.

Adjoint Field Functions

The following field functions are made available when computing the adjoint flow solution and the shape sensitivity.

Adjoint of [cost function] w.r.t.	Flow adjoint of the respective cost function.	
	Continuity	The adjoint with respect to continuity. This value tells you the sensitivity of the cost function with respect to local changes in mass within the domain.

	X-, Y-, Z-momentum	The adjoint with respect to momentum. This value tells you the sensitivity of the cost function with respect to local changes in X-, Y-, and Z-momentum within the domain.
	Energy	The adjoint with regard to energy. This value tells you the sensitivity of the cost function with regard to local changes in energy within the domain.
	Position	The adjoint with respect to position. This value tells you the sensitivity of the cost function with respect to displacement boundary position.
Sensitivity of [cost function] w.r.t. Position	Sensitivity of the respective cost function with respect to position. The position refers to the coordinates of the vertices of the volume mesh. A high value (positive or negative) tells you that the respective cost function is particularly sensitive to that given vertex coordinate position.	

Boundary Conditions Formulation

This section describes the boundary condition formulation for the flow and energy models, considering each of the following boundary types in turn:

- Velocity inlet
- Symmetry plane
- Wall
- Stagnation inlet
- Pressure outlet
- Mass flow inlet
- Free-stream
- Flow-split outlet

Velocity Inlet

Velocity

The inlet face velocity vector is specified directly.

Pressure

The boundary face pressure is extrapolated from the adjacent cell using reconstruction gradients.

Temperature

The inlet static temperature is specified.

If applicable to the case, it is possible to specify inlet boundary conditions on a reference frame.

Symmetry Plane

Velocity

The shear stress at a symmetry boundary is zero. The face value of velocity is computed by extrapolating the parallel component of velocity in the adjacent cell using reconstruction gradients.

Pressure

The boundary face pressure is extrapolated from the adjacent cell using reconstruction gradients.

Temperature

The heat flux at a symmetry face is zero. The face value of temperature is extrapolated from the adjacent cell using reconstruction gradients.

Wall

Velocity

For *slip walls*, the face value of velocity is computed by extrapolating the parallel component of velocity in the adjacent cell using reconstruction gradients.

For *no-slip walls*, the tangential velocity is explicitly set either to zero (for the case of no wall motion) or to a specified value.

Pressure

The boundary face pressure is extrapolated from the adjacent cell using reconstruction gradients.

Temperature

For inviscid flow, the face value of temperature is extrapolated from the adjacent cell using reconstruction gradients.

For viscous flows, either the temperature or the heat flux is specified. If the heat flux is specified, the wall boundary temperature is computed from the appropriate laminar or turbulent heat flux relationship.

At walls, the wall heat flux is:

$$q_w = -k\frac{dT}{dn} \tag{115}$$

This flux is linearized in terms of the cell and wall temperatures using Eqn. (386). The coefficients of this linearization are functions of the wall model (for example, laminar, turbulent, with or without radiation, with or without boiling). When the cell temperatures are updated, the wall temperature is computed using Eqn. (386).

Stagnation Inlet

Velocity

The boundary face velocity magnitude $|\mathbf{v}_f|$ is obtained from:

$$|\mathbf{v}_f| = \sqrt{2C_p(T_{tf} - T_f)} \tag{116}$$

where C_p is the specific heat, T_{tf} is the total temperature and T_f is the static temperature at boundary face f.

The direction vector is explicitly specified.

Pressure

The boundary face total pressure P_{tf} is specified. For a supersonic flow, the boundary face pressure P_f is specified, while for subsonic flow it is extrapolated from the adjacent cell using reconstruction gradients.

You can specify a pressure jump across the boundary.

This option also applies to pressure outlet boundaries.

Temperature

The boundary face total temperature T_{tf} is specified, and the static temperature is obtained as follows:

$$T_f = \frac{T_{tf}}{(P_{tf}/P_f)^{R/C_p}} \tag{117}$$

If applicable to the case, it is possible to specify inlet boundary conditions on a reference frame.

Pressure Outlet

Velocity

The boundary face velocity is extrapolated from the interior using reconstruction gradients. In the case of the Segregated Flow model, an option exists to use only the normal component of the velocity for inflow. This option can stabilize the solution in certain situations.

Pressure

Inflow

When *inflow* occurs, the boundary pressure p_f is assumed to be given by:

$$p_f = p_{specified} - \frac{1}{2}\rho_f |v_n|^2 \tag{118}$$

where v_n is the normal component of the boundary inflow velocity. This approach discourages backflow from occurring.

Outflow

For *subsonic outflow,* the boundary pressure p_f is specified.

For *supersonic outflow*, the boundary face pressure is extrapolated from the adjacent cell value using reconstruction gradients.

For *radial equilibrium outflow*, the radial pressure gradient exactly balances the centrifugal forces due to rotation. To apply radial equilibrium to a specified pressure boundary condition, one must impose the condition:

$$p_f = p_i + \int_{r_i}^{r} \frac{\rho v_\theta^2}{r} dr \tag{119}$$

where the subscript i refers to the inner location (that is, the hub in a turbomachine) and v_θ is the tangential component of the velocity.

To obtain the function $(\rho v_\theta^2)/r$ as a function of r only, the data is sorted into bins. An arbitrary number of 100 bins is chosen. If any bins are empty, linear interpolation fills them. Since the interpolated data is going to be numerically integrated, the interpolation order is not very important.

Target Mass Flow Rate and Pressure Correction

When the Target Mass Flow Rate Option is enabled, the boundary condition automatically applies a correction to the pressure value specified at the boundary in an attempt to yield the specified mass flow rate.

At each iteration, the value of pressure that is applied at the boundary is given by the specified pressure value plus an adjustment:

$$p_f = p_{specified} + \delta p \tag{120}$$

A pressure correction of $\Delta p \equiv p_f^{n+1} - p_f^n$ is calculated as follows:

$$\Delta p = \frac{\omega \cdot \left|\dot{m}^n\right| \cdot (\dot{m}^n - \dot{m}_{target})}{(\rho_{ave}^n \cdot A^2)} \tag{121}$$

where $\dot{m}^n$ and ρ_{ave}^n are the mass flow rate and average density at the current iteration, and $\dot{m}_{target}$ and A are the target mass flow rate and total area of the boundary, respectively. ω is an under-relaxation factor.

Combining Eqns. (120) and (121), the new pressure adjustment value becomes:

$$\delta p^{n+1} = \frac{\omega \cdot \left|\dot{m}^n\right| \cdot (\dot{m}^n - \dot{m}_{target})}{(\rho_{ave}^n \cdot A^2)} + \delta p^n \tag{122}$$

While the adjustment is applied at every iteration, a new pressure adjustment value is computed according to a specified update frequency.

Dynamic Head and Pressure Specification

Problems arise when flow recirculates at the pressure outlet boundary, in which case that boundary is no longer a real outlet. To prevent this, the model automatically subtracts the *dynamic head* from pressure only on the faces where recirculation occurs. (This is in addition to the pressure terms given in Eqn. (118).) For some applications, however, it is important to maintain the specified pressure at the outlets. In such cases, the user can disable the use of the dynamic head by setting the Pressure specification option to Static for the pressure outlet boundary concerned.

The pressure specification method is present for Pressure Outlet boundary conditions. When this option is set to Static, Eqn. (118) no longer applies and the solver disables the subtraction of the dynamic head from the specified pressure in case of reversed flow or inflow at the pressure outlet boundary.



Pressure Specification Properties

Environmental	Include dynamic head at pressure outlet boundary.
Static	Do not use dynamic head in the case of inflow. Pressure is maintained at the pressure specified by the user.

Setting the Pressure Jump Option

You can specify a pressure jump across the boundary.

This option also applies to Stagnation Inlet boundaries.

Resolving Flow Convergence Failure

When you use the Pressure Jump Option with the Segregated Flow Solver, there can be times when the flow solution does not converge.

- To resolve this problem, activate Delta-V Dissipation in the Segregated Flow model:



Temperature

For outflow, the boundary face temperature is extrapolated from the adjacent cell using reconstruction gradients. For inflow, the specified static temperature is used.

Mass Flow Inlet

Velocity

The velocity is specified as:

$$\mathbf{v}_f = \frac{\dot{m}_{\text{specified}}\vec{\theta}}{\rho(\vec{\theta}\cdot\mathbf{a})} \tag{123}$$

where $\dot{m}_{specified}$ is the specified mass flow rate and $\grave{\theta}$ is the specified flow angle. The mass flow rate $\dot{m}_{specified}$ may be specified either directly or as a flux (that is, mass flow rate per unit area).

Pressure

For *supersonic inflow*, the specified pressure is used. For *subsonic flow*, the boundary face pressure is extrapolated from the adjacent cell using reconstruction gradients.

Temperature

The static temperature is obtained from the specified total temperature using the isentropic definition of total enthalpy.

If applicable to the case, it is possible to specify inlet boundary conditions on a reference frame.

Free-Stream

For a freestream boundary, there are three ways that you can specify the freestream pressure, temperature, and Mach number.

You select the appropriate option as a property of the Region Name > Boundary > Boundary Name > Physics Conditions > Free Stream Option node.



Free Stream Option Properties

Mach Number + Pressure + Temperature	For this option, STAR-CCM+ adds Mach Number, Pressure, and Static Temperature nodes to the Physics Values for the boundary. This option is the default freestream option.
Altitude + Length Scale + Reynolds Number	For this option, STAR-CCM+ adds the Altitude node to the Physics Values for the boundary. The freestream pressure and temperature are computed internally using the specified altitude. Specify the appropriate Reynolds number and length scale values in the Mach Number node under the Physics Values node for the boundary. The freestream Mach number is computed internally using the specified altitude, Reynolds number, and length scale.
Altitude + Mach Number	For this option, STAR-CCM+ adds the Altitude and Mach Number nodes to the Physics Values for the boundary. The freestream pressure and temperature are computed internally using the specified altitude.

When you use the Altitude + Length Scale + Reynolds Number or Altitude + Mach Number options, you also specify the appropriate atmosphere type as a

property of the Physics Conditions > Atmosphere Type Option node.



Atmosphere Type Option Properties

Standard	The US 1976 Standard Atmosphere is used to compute the freestream pressure and temperature from the altitude.
User Table	A user-supplied table is used to compute the freestream pressure and temperature from the altitude.

Using an Atmosphere Table

If you select the User Table option, you must first import the appropriate atmosphere table into STAR-CCM+.

You then set up the atmosphere table properties in the Atmosphere Table node under the Region Name > Boundary > Boundary Name > Physics Values node.



Note: When you use the Altitude + Length Scale + Reynolds Number or Altitude + Mach Number options, ensure that you set initial values that are consistent with the freestream values. You set the initial values in the physics continuum, under the Initial Conditions node, in the Pressure and Static Temperature nodes.

Velocity

For *supersonic inflow*, the velocity is taken to be the free-stream velocity, computed from the specified Mach number and temperature.

For *supersonic outflow*, the boundary face velocity is extrapolated from the cell value using reconstruction gradients.

For *subsonic inflow*, the tangential component of the free-stream velocity is combined with the boundary normal velocity computed from characteristics as follows:

$$\mathbf{v}_f = \mathbf{v}_\infty + (\mathrm{v}_{fn} - \mathrm{v}_{\infty n})\frac{\mathbf{a}}{|\mathbf{a}|} \tag{124}$$

where $\mathrm{v}_{\infty n}$ is the component of free-stream velocity normal to the boundary and v_{fn} is the boundary normal velocity that is obtained from characteristics:

$$\mathrm{v}_{fn} = \frac{1}{2}(\mathrm{v}_{0n}^r + \mathrm{v}_{\infty n}) + \frac{c_0 - c_\infty}{\gamma - 1} \tag{125}$$

In this expression, v_{0n}^r is the normal component of the boundary velocity that is extrapolated from the cell. c_∞ is the speed of sound obtained from the free-stream temperature. c_0 is the speed of sound obtained from T_0^r, the boundary temperature that is extrapolated from the cell, that is:

$$c_0 = \sqrt{\gamma R T_0^r} \tag{126}$$

For *subsonic outflow*, the tangential component of the boundary velocity, that is extrapolated from the cell using the reconstruction gradient $\mathbf{v}_{0n}^r$, is combined with the boundary normal velocity computed from characteristics. The result is:

$$\mathbf{v}_f = \mathbf{v}_0^r + (\mathbf{v}_{fn} - \mathbf{v}_{0n}^r)\frac{\mathbf{a}}{|\mathbf{a}|} \tag{127}$$

where $\mathbf{v}_{fn}$ is again computed using Eqn. (125).

Pressure

For *supersonic inflow*, the pressure is taken to be the free-stream pressure.

For *supersonic outflow*, the boundary face pressure is extrapolated from the adjacent cell value using reconstruction gradients.

For *subsonic inflow*, the pressure is computed from the free-stream pressure using the following isentropic relation:

$$p_f = p_\infty \left(\frac{T_f}{T_\infty}\right)^{C_p/R} \tag{128}$$

For *subsonic outflow*, the pressure is computed from the adjacent cell value using the following isentropic relation:

$$p_f = p_0^r \left(\frac{T_f}{T_0^r} \right)^{C_p/R} \tag{129}$$

Temperature

For *supersonic inflow*, the boundary face temperature is taken to be the free-stream temperature.

For *supersonic outflow*, the boundary face temperature is extrapolated from the adjacent cell value using reconstruction gradients.

For *subsonic flow*, the boundary face temperature is computed as:

$$T_f = \frac{c_f^2}{\gamma R} \tag{130}$$

where:

$$c_f = \frac{1}{2}(c_\infty + c_0) + \frac{1}{4}(\gamma - 1)(\mathrm{v}_{0n}^r - \mathrm{v}_{\infty n}) \tag{131}$$

Flow-Split Outlet

Velocity

The flow velocity at a flow-split outlet boundary i is given by:

$$\mathbf{v}_f = \mathbf{v}_0^r + \frac{x_i}{\rho_f} \frac{\mathbf{a}}{|\mathbf{a}|} \tag{132}$$

where $\mathbf{v}_0^r$ is the velocity that is extrapolated from the adjacent cell value using reconstruction gradients and the vector $\mathbf{a}/|\mathbf{a}|$ is the outward-pointing normal. x_i is a scale factor that is computed for outlet boundary i as follows:

$$x_i = -\frac{f_i}{\sum_{n \text{ outlets}} f_j} \frac{(\dot{m}_{\text{in}} - \dot{m}_i^*)}{\sum_{\text{outlet } i \text{ faces}} |\mathbf{a}|} \tag{133}$$

where f_i is the specified fraction of the flow leaving outlet i, and $\dot{m}_{\text{in}}$ is the total inlet flow, that is defined as:

$$\dot{m}_{\text{in}} = \sum_{\text{non-outlet faces}} \rho_f(\mathbf{v}_f \cdot \mathbf{a} - G) \tag{134}$$

and $\dot{m}_i^*$ is the (unscaled) outlet mass flow rate through outlet i:

$$\dot{m}_i^* = \sum_{\text{outlet } i \text{ faces}} \rho_f \max(\mathbf{v}_0 \cdot \mathbf{a} - G, 0) \tag{135}$$

Pressure

The boundary face pressure is extrapolated from the adjacent cell using reconstruction gradients.

Temperature

The boundary face temperature is extrapolated from the adjacent cell using reconstruction gradients.

Setting the Mass Flux Specification

The flow split can be specified by flow ratio or by mass flow rate. However, all of the flow split boundaries in a continuum must use the same type of mass flux specification. That is, all of the flow split boundaries specify the flow rate using split ratios, or they all specify the flow rate using mass flow rates.

You specify the flow split type in the Mass Flux Specification node.



Mass Flux Specification Properties

Specified Split Ratios	Lets you specify the fraction (f_i) of the total flow that passes through the outlet. Set the Split Ratio property in the Physics Values > Split Ratio node. This boundary type cannot be used in the same continuum as pressure or stagnation boundaries.
Specified Mass Fluxes	Lets you specify the mass flow rate through the outlet. Set the appropriate value under the Physics Values > Mass Flow Rate node. The available methods are: Constant, Field Function, Table (iteration), Table (time), User Code. This option is available only with the Segregated Flow model.

Using a Specified Mass Flow Rate

When you specify a mass flow rate, at every iteration the solver makes a correction to the flow rate at this boundary to ensure that the total flow matches the specified flow rate. If the compressible solver is used, the extrapolated density is used for computing the mass flow rate.

This boundary type can be used in the same continuum as a pressure outlet. However, you need to consider the following:

- In scenarios where the pressure outlet may act as an inlet, you need to be careful how you specify the pressure condition. In the event of reversed flow, the dynamic head is taken out of the specified pressure, which can cause an incorrect pressure profile at the pressure outlet.
- If you specify the mass flow rate using a field function, an average value on that boundary is used as the specified mass flow rate. The mass flow rate must be a single value, not a field function with a value that varies along the boundary.

Modeling Non-Reflecting Boundaries and Mixing-Plane Interfaces

STAR-CCM+ can work effectively with non-reflecting boundaries and mixing-plane interfaces, allowing accurate simulations of applications such as multi-stage turbomachinery.

The implementation of non-reflecting boundaries and explicit mixing-plane interfaces is based on the work of Giles [23], [24], [25] and Saxer [47]. The theory is based on a Fourier analysis of the linearized Euler equations. The approach assumes that the solution at the boundary is decomposed, circumferentially in three-dimensional space, into Fourier modes.

However, these entities differ in the zeroth mode:

- For non-reflecting boundaries, the zeroth mode corresponds to the average solution and receives the standard boundary condition treatment.
- For explicit mixing-plane interfaces, conditions that are used for the zeroth mode solution are reached through an averaging of conserved quantities across the mixing-plane interface.

Additional information is available on these features as follows:

- The non-reflecting boundary condition—its theoretical background and purpose. This information includes setting average total pressure, temperature, and flow angle at the inlet, and pressure at the outlet.
- Details of setting up the non-reflecting option
- The mixing-plane interface—its uses and options.

Flow Reference

This section provides reference information that is common to the flow models in STAR-CCM+:

- Conditions that are shared among the models
- Initial conditions
- Reference values

Conditions

The boundary conditions that are available to all flow models in STAR-CCM+ are:

- Flow direction specification
- Reference frame specification
- Backflow direction specification
- Mass flow option
- Pressure jump option
- Shear stress specification
- Tangential velocity specification
- Target mass flow option
- Turbulence suppression option
- Velocity specification

The non-reflecting option boundary condition is available to the Coupled Flow, and Harmonic Balance Flow and Energy models.

The region conditions that are available to all flow models in STAR-CCM+ are:

- Mass source option
- Momentum source option

This section provides a detailed property reference for these conditions. Region, boundary, or interface conditions are elected by the region, boundary, or interface types, which in turn are limited by the models that have been turned on in the continuum.

Flow Direction Specification

Defines how the flow direction at the boundary is specified.

Properties

Method	Specifies how the direction of flow is input.	
	Boundary Normal	Flow direction is normal to the boundary. No nodes are added to the Physics Values node.
	Components	Direction is specified as a vector—a flow direction node is added to the Physics Values node.
	Angles	Direction is specified as three angles—a flow angles node is added to the Physics Values node.

Reference Frame Specification

You can specify boundary conditions in rotating reference frames for the following types of boundaries:

- Velocity Inlet
- Stagnation Inlet
- Mass Flow Inlet
- Pressure Outlet
- Wall

You can use a rotating reference frame to apply upstream data to a turbomachinery blade.

Each inlet boundary has a Reference Frame Specification node that defines the reference frame in its properties.



Properties

Option — Provides options for the reference frame applied to the boundary.

- **Lab Frame** — Uses the Lab reference frame (default). When the region is not rotating, it has no rotating reference frame attached and uses the Lab reference frame. In such a situation, `Lab Frame` for the boundary has the same effect as the `Region Reference Frame`.
- **Region Reference Frame** — Uses the reference frame of the parent region. When the region has a rotating reference frame attached, a physically valid setup using this condition puts the boundary in the reference frame of the region. This implies that the boundary is rotating with the region at the same RPM, and has the same rotation axis and origin as the region.
- **Local Reference Frame** — Selects a reference frame separately for the boundary. Activates the Boundary Reference Frame Specification node.
  This option allows the boundary to have its own frame (only rotating frame) attributes. For example, this can be used when both the boundary and the region are rotating at a different RPM but have the same axis and origin. A setup where the region and boundary have different values for the RPM, axis, and origin is also possible.

Backflow Direction Specification

The Backflow Direction Specification sets the method by which STAR-CCM+ computes the direction of any flow that enters the simulation through an outflow boundary. For some methods, you must also set the associated Reference Frame Specification.



The continuity equation imposes an internal constraint on the velocity normal to the boundary surface. The velocity normal, U_n, is given as:

$$U_n = \mathbf{U} \cdot \mathbf{a} / |\mathbf{a}|$$

where $\mathbf{a}$ is the area vector.

This constraint means that U_n, the normal component of the velocity $\mathbf{U}$ in the absolute (Laboratory) reference frame must be left unchanged regardless of the overall direction defined for the backflow. If the overall direction is defined by $\mathbf{k}$, then the following formula applies to all Backflow Direction Specification options, apart from the Extrapolated option (where the velocity at the boundary is equal to the internal velocity):

$$\mathbf{U}' = \frac{\mathbf{W} \cdot \mathbf{n}}{\mathbf{k} \cdot \mathbf{n}} \mathbf{k} + \mathbf{V}_{frame} \tag{136}$$

where $\mathbf{U}'$ is the new inflow velocity in the absolute frame and $\mathbf{W} = \mathbf{U} - \mathbf{V}_{frame}$ is the relative velocity.

Eqn. (136) preserves U_n as $\mathbf{U}' \cdot \mathbf{n} = \mathbf{U} \cdot \mathbf{n}$. Eqn. (136) is effectively scaling the desired velocity in order to match U_n.

For the Backflow Direction Specification options, $\mathbf{k}$ is defined as follows:

- Boundary Normal

$$\mathbf{k} = \mathbf{n} \tag{137}$$

- Components

$$\mathbf{k} = \mathbf{V}_{user} \tag{138}$$

  where the vector inflow $\mathbf{V}_{user}$ is explicitly defined in the specified reference frame.

- Angles

$$\mathbf{k} = \mathbf{n}^{\alpha}_{user} \tag{139}$$

  where the direction angles of the inflow velocity are explicitly defined in the specified reference frame. The specified direction angles are internally converted into the components of a unit vector, where $\mathbf{n}^{\alpha}_{user}$ are the unit vector cosine angles.

Backflow Direction Specification Properties

Method — Specifies how the reverse flow direction is computed.

Extrapolated — Extrapolates the flow direction from the interior of the domain. In most situations, this is less stable than the Boundary Normal option. However, it is recommended for situations where the flow is known to be parallel to the pressure boundary, such as a co-flowing jet. This option is set by default for the Coupled Flow model.



Boundary Normal — Assumes that the incoming flow enters the simulation along a vector that is orthogonal to the boundary surface, in the specified reference frame for the boundary. This option is set by default for the Segregated Flow model and generally provides the most robust convergence. If the flow is largely parallel to the boundary, the Extrapolated option should be used instead.

Components — The inflow direction vector $\mathbf{v}_{user}$ is explicitly defined in the specified reference frame. See Eqn. (138).

When the Reference Frame Specification is set to Lab Frame, this value is defined in the Flow Direction node under Physics Values. When the Reference Frame Specification is set to an option other than Lab Frame, this value is defined in the Relative Flow Direction node.

Angles — The direction angles of the inflow velocity are explicitly defined in the specified reference frame. See Eqn. (139).

This value is defined in the Flow Angles node under Physics Values.

Mass Flow Option

Mass flow can be specified either as a flux or as a rate.

Properties

Specification Option	Specifies how the mass flow is input.	
	Mass Flux	Mass flux, or flow rate per unit area, through the boundary. This option adds a mass flux node to the Physics Values node.
	Mass Flow Rate	Mass flow rate through the boundary. This option adds a mass flow rate node to the Physics Values node.

Pressure Jump Option

The pressure jump option lets you specify a pressure change across a boundary.

You specify the pressure jump type as a property of the Pressure Jump Option node.



When you select the pressure jump type, the appropriate nodes are added to the Physics Conditions node and Physics Values node.

The Backflow Direction Specification node under the Physics Conditions node and the Pressure Jump Under-Relaxation Factor node under the Physics Values node apply to all pressure jump types.

Pressure Jump Option Properties

None	No pressure jump exists across the boundary. This option is the default.
Fan	The pressure rise is obtained from a fan curve. At a Stagnation Inlet boundary, the pressure rise is applied in the direction going into the domain. At a Pressure Outlet boundary, the pressure rise is applied in the direction going out of the domain. You specify the appropriate fan performance curve in the Fan Curve Specification node under the Physics Conditions node. You specify the fan curve polynomial or table, and the fan verbosity, in the corresponding nodes under the Physics Values node. For more information, see Specifying the Fan Performance Curve Type.
Porous	The pressure loss is computed as: $$\alpha\rho\|v\|^2 + \beta\|v\|$$ where α is the Porous Inertial Resistance and β is the Porous Viscous Resistance. The pressure loss is always applied in the direction of flow. You set these values in the corresponding nodes under the Physics Values node. For more information, see Porous Baffle Properties.

Loss Coefficient

The pressure loss is computed as:

$$0.5K\rho|v|^2$$

where K is the Pressure Loss Coefficient. The pressure loss is always applied in the direction of flow.

You set this value in the corresponding node under the Physics Values node.

Pressure Jump Under-Relaxation Factor

To make the simulation more stable and improve convergence, you can set the appropriate Pressure Jump Under-Relaxation Factor.

The default setting is 0.1 for the Segregated solver and 1.0 for the Coupled solver. You can modify these values independently, or in conjugation with the Velocity Under-Relaxation Factor (for the Segregated solver).



Shear Stress Specification

Defines whether the boundary is a no-slip wall or not.

Properties

Method Defines the method of entering the value

No-Slip Adds a tangential velocity specification node to the Conditions node.

Slip No shear stress at the wall. No nodes are added to the Physics Values node.

Tangential Velocity Specification

This node is added when the no-slip method is specified in the shear stress specification.



When a Tangential Velocity Specification is used at a wall f, only the tangential component v_τ of the specified velocity $\boldsymbol{v}_{\text{spec}}^{\text{lab}}$ is used. If you specify a velocity with a component normal to the wall, the normal component of $\boldsymbol{v}_{\text{spec}}^{\text{lab}}$ is ignored, since the velocity $\boldsymbol{v}_f$ at the face is computed as:

$$\boldsymbol{v}_f = \boldsymbol{v}_{\text{spec}}^{\text{lab}} - (\boldsymbol{v}_{\text{spec}}^{\text{lab}} \cdot \boldsymbol{a} - G_f)\frac{\boldsymbol{a}}{a^2}$$

where $\boldsymbol{a}$ is the face area vector and $\boldsymbol{v}_{\text{spec}}^{\text{lab}}$ is the velocity in the laboratory frame. The normal component contribution at the face f can only come from the grid flux G_f that is given by Eqn. (16). That is, if the wall itself is moving at a velocity with a non-zero component in the wall normal direction.

Note that:

$$\boldsymbol{v}_{\text{spec}}^{\text{lab}} = \boldsymbol{v}_{\text{spec}}^{\text{mesh}} + \boldsymbol{v}_{\text{mesh}}^{\text{ref}} + \boldsymbol{v}_{\text{ref}}^{\text{lab}}$$

where:

- $\boldsymbol{v}_{\text{spec}}^{\text{lab}}$ is the specified velocity measured in the laboratory frame.
- $\boldsymbol{v}_{\text{spec}}^{\text{mesh}}$ is the specified velocity measured relative to the mesh.
- $\boldsymbol{v}_{\text{mesh}}^{\text{ref}}$ is the velocity of the mesh relative to the reference frame. This quantity is in Region > Physical Values > Motion Specification.
- $\boldsymbol{v}_{\text{ref}}^{\text{lab}}$ is the velocity of the reference frame relative to the laboratory frame.

Properties

The net tangential velocity in the laboratory frame is the vector sum of the tangential velocity with respect to the reference frame selected under the condition Reference Frame Specification and the velocity of the selected frame with respect to the laboratory frame.

Unless specified in an Axis node, axis of rotation is defined by information in the Motion Specification node under the Physics Values node.

Method — Defines the method of entering the tangential velocity at the wall boundary.

- Fixed — Tangential velocity is zero with respect to the reference frame selected under the condition Reference Frame Specification. The following nodes are added to the Physics Values node:
  - No nodes if Lab Frame or Region Reference Frame is selected for reference frame.
  - A Boundary Reference Frame Specification node if Local Reference Frame is selected for reference frame.

Rotation Rate	Tangential velocity is specified by rotation rate with respect to the reference frame selected under the condition Reference Frame Specification. The following nodes are added to the Physics Values node: • A Wall Rotation node if Lab Frame is selected for reference frame. • A Wall Relative Rotation node if Region Reference Frame is selected for reference frame. • A Wall Relative Rotation node and a Boundary Reference Frame Specification node if Local Reference Frame is selected for reference frame.
Vector	Tangential velocity is specified as a vector with respect to the reference frame selected under the condition Reference Frame Specification. The following nodes are added to the Physics Values node: • A Velocity profile if Lab Frame is selected for reference frame. • A Relative Velocity node if Region Reference Frame is selected for reference frame. • A Relative Velocity node and a Boundary Reference Frame Specification node if Local Reference Frame is selected for reference frame.
Local Rotation Rate	Tangential velocity is defined by the rotation rate around the axis specified in the Axis node of the Boundary and with respect to the reference frame selected under the Reference Frame Specification. Besides the Axis node, the following nodes are added to the Physics Values node: • A Wall Rotation node if Lab Frame is selected for reference frame. • A Wall Relative Rotation node if Region Reference Frame is selected for reference frame. • A Wall Relative Rotation node and a Boundary Reference Frame Specification node if Local Reference Frame is selected for reference frame.

Target Mass Flow Option

The Target Mass Flow Option is available at Pressure Outlet boundaries. When this option is used, the Pressure value that is specified at the boundary is automatically corrected in order to achieve a mass flow rate as specified in the Target Mass Flow Rate.

Target Mass Flow Option Properties

Target Mass Flow Option	Specifies whether to use the target mass flow option.	
	Activated	Use the target mass flow option. A Target Mass Flow Rate node is added to the Physics Values node.
	Deactivated	Do not use the target mass flow option.

Turbulence Suppression Option

Defines whether turbulence is to be suppressed near the boundary.

Properties

Turbulence suppression option	Defines whether you want turbulence to be suppressed near the wall.	
	Activated	Suppresses turbulence.
	Deactivated	Does not suppress turbulence.

Velocity Specification

A velocity specification can be entered as a vector, or as a magnitude and direction, in which case the direction would be specified as a direction vector.

Properties

Method — Specifies the method to use for entering velocity.

- Components — Enter the velocity as vector with magnitude. This option adds a Velocity node to the Physics Values node.
- Magnitude Direction — Enter a direction vector and a magnitude. Two nodes are added to the Physics Values node—Flow Direction and Velocity Magnitude.
- Background — This method is available in the Dispersed Multiphase model only. The velocity of the dispersed phase is taken from the background flow.
  For more information, see Velocity Specification.

Wall Surface Specification

Defines whether the boundary is a rough or smooth wall. Available for high-y^+ and all-y^+ wall treatments, but not low-y^+ wall treatments.

Properties

Method — Defines the method of entering the value.

- Smooth — Used regular wall treatment without modification for wall roughness.
- Rough — Modifies the wall treatment to incorporate roughness. Adds wall roughness and roughness height nodes to the Physics Values node.

Wall Roughness Parameters

The wall roughness parameters allow the roughness function to be specified when the wall surface specification is Rough.

Properties

B	Coefficient in roughness function.
C	Coefficient in roughness function.
RplusRough	Roughness parameter R^+ in roughness function corresponding to a fully rough surface.
RplusSmooth	Roughness parameter R^+ in roughness function corresponding to a fully smooth surface.

Expert Properties

Roughness Limiter	Whether to limit the roughness height locally, such that $R^+ = y^+$ when y^+ becomes less than R^+.	
	Activated	Limits the roughness height. (Default)
	Deactivated	Does not limit the roughness height.

Non-Reflecting Option

This boundary condition is on stagnation inlet and pressure outlet boundaries.

In three-dimensional simulations, the boundary must be a surface of revolution. Furthermore, the boundary mesh must be a structured, mapped mesh of quadrilateral faces. (This requirement does not apply to steady-state calculations using the Coupled Flow model.) If these conditions are not satisfied, an error is generated during pre-processing of the boundary mesh before the application of the non-reflecting boundary condition procedure.

This option is available only when the Coupled Flow Ideal Gas, Steady, and Two-Dimensional or Three-Dimensional models are enabled, and only on boundaries of linearly (two-dimensional) or rotationally (three-dimensional) periodic problems. It is appropriate in instances where the flow normal to the non-reflecting boundary remains subsonic.

Properties

Non-Reflecting Option	Specifies whether to use the non-reflecting option.	
	Activated	Use the non-reflecting option.
	Deactivated	Do not use the non-reflecting option.

Note:

No additional nodes appear in the Values node when this option is activated. However, there are some restrictions as to which methods are appropriate for specifying other boundary values. The table below indicates which methods can be used with which boundary values when this option is activated.

	Method			
Boundary Value	Constant	Radial Equilib.	Tabular Data	Field Function
Total pressure	Yes	No	3D[1]	No
Pressure	Yes	Yes[2]	3D[1]	No
Total temperature	Yes	No	3D[1]	No
Flow direction	Yes	No	3D[1]	No

Notes:

[1] External radial tabular data files only, not internal or xyz formatted.
[2] At a pressure outflow boundary only.

All existing Methods for turbulence quantities and other transported scalars are acceptable.

Mass Source Option

A mass source option can be allowed or not.

Properties

Mass source option — Defines whether you want to enter a mass source term.

- Activated — Allows a mass source term to be defined. The following nodes are added to the Physics Values node:
  - Mass Source
  - Mass Source Pressure Derivative
  - Mass Source Temperature Derivative
  - Mass Source Velocity Derivative
- Deactivated — Does not allow a mass source term to be defined.

Momentum Source Option

This condition allows you to select one of the following options:

- No momentum source defined
- A momentum source term defined
- Fan properties and device data defined

Properties

Momentum source option	Defines whether you want to enter a momentum source term, and which type.	
	None	Does not allow a momentum source term to be defined.
	Specified	Allows a momentum source term to be defined. The following nodes are added to the Physics Values node: • Momentum Source • Momentum Source Pressure Derivative • Momentum Source Temperature Derivative • Momentum Source Velocity Derivative
	Fan	Allows fan properties and device data to be defined. A Fan Momentum Source node is added to the Physics Values node.

Initial Conditions

This section lists all the initial conditions that are available for flow models in STAR-CCM+ and it provides a detailed reference of the methods of setting them:

- (Static) Pressure
- Static temperature
- Velocity

(Static) Pressure

The outlet pressure is reasonable to use for the initial value of pressure. If there are no pressure boundaries that are attached to the given continuum, the initial value is arbitrary. It is entered as a scalar profile.

Static Temperature

Energy models require this initial condition. The Coupled Energy model also adds this condition.

It is entered as a scalar profile.

Velocity

For inviscid problems, the initial velocity must be non-zero. Zero initial velocity is also inappropriate if the Turbulence Intensity method is being used to specify initial turbulence values. It is entered as a vector profile.

Reference Values

This section lists all the reference values that must be set for flow models in STAR-CCM+ and it provides a detailed reference of the methods of setting them:

- Gravity
- Maximum allowable temperature
- Minimum allowable temperature
- Minimum allowable wall distance
- Maximum number of eddies
- Reference altitude
- Reference density
- Reference pressure

Gravity

Gravitational acceleration.

Properties

Value	The value and units of the gravity. You can also enter field function syntax directly. Clicking the customizer executes the define value action.

Maximum Allowable Temperature

The largest temperature value that is permitted anywhere in the continuum.

The Energy models (Coupled Energy, Coupled Solid Energy, Segregated Solid Energy, Segregated Fluid Enthalpy, Segregated Fluid Temperature) limit temperature corrections such that the corrected value does not exceed this maximum. If this occurs, a message is printed to the *Output* window.

Properties

Value	The value and units of the maximum allowable temperature. Clicking the customizer executes the define value action. (Default value is 5000 K.)

Minimum Allowable Temperature

The smallest temperature value that is permitted anywhere in the continuum.

The Energy models (Coupled Energy, Coupled Solid Energy, Segregated Solid Energy, Segregated Fluid Enthalpy, Segregated Fluid Temperature) limit temperature corrections such that the corrected value does not go below this minimum. If this occurs, a message is printed to the *Output* window.

Properties

Value	The value and units of the minimum allowable temperature. Clicking the customizer executes the define value action. (Default value is 100 K.)

Minimum Allowable Wall Distance

The wall distance algorithm has been developed to return a positive definite value for almost all cell shapes. However it is possible that some degenerate cells, such as boundary sliver cells, could return a value that is close to zero. Since turbulence models and wall functions require a non-zero value, this could cause numerical problems. For this reason, the minimum allowable wall distance limits the value of the wall distance in the continuum.

If the wall distance is limited in one or more cells, a warning message is displayed. In this case, investigate the cause and attempt to remedy the situation.

Any specific variant of any turbulence model activates this value.

Properties

Value	The value and units of the minimum allowable wall distance. Clicking the customizer executes the define value action. (Default value is 1.0E-6 m.)

Maximum Number of Eddies

The maximum number of eddies parameter sets an upper limit on the number of initial eddies that are generated within the flow domain by the Synthetic Eddy Method. These initial eddies prepare the flow for an analysis that uses the Large Eddy Simulation approach for turbulence. If the maximum limit is reached, a warning is shown in the *Output* window, as this could potentially destroy the desired effect of the initial turbulence. It should not be set too low.

Reference Altitude

A position vector that defines the point at which static and piezometric pressures are equal. This quantity is used in the Gravity model formulation for both variable and constant density flows. Although its exact position is arbitrary (similar to the values of reference pressure and reference density), it is reasonable to locate it at a central point within the solution domain. This means that the expected range of values for working pressure is balanced around zero, and so provide an optimum condition for the STAR-CCM+ numerical schemes.

The working pressure at pressure boundaries has to be consistent with the selected reference altitude, reference density, and reference pressure. If any of these values are changed, it is necessary to make sure that the working pressures set on pressure boundaries are still consistent.

Properties

Value	The value and units of the reference altitude. Clicking the customizer executes the define value action.

Reference Density

The reference density node.

Properties

Value	The value and units of the reference density. Clicking the customizer executes the define value action.

Reference Pressure

The pressures referred to by the flow models in STAR-CCM+ are implicitly gauge pressures. This is important for minimizing roundoff errors, particularly for incompressible problems where pressure changes might be small.

If the Constant Density model is being used, the operating pressure has no bearing on the calculation. However, when the Ideal Gas model is chosen, the absolute pressure is used to obtain the density so that the operating pressure must be specified appropriately.

Properties

Value	The value and units of the reference pressure. Clicking the customizer executes the define value action.

Values Reference

This section contains the following topics:

- Boundary Values

  This section lists all the boundary values that are available for flow and energy models in STAR-CCM+.

- Region Values

  This section provides information about the region values that are available for flow and energy models in STAR-CCM+.

This section also provides a detailed reference of the methods of setting the values. Region, boundary, and interface values are elected by the region, boundary, or interface types, and/or by methods that are selected in corresponding nodes in the boundary conditions node.

Boundary Values

This section lists all the boundary values that are available for flow and energy models in STAR-CCM+.

Flow angle	Allows the flow direction to be defined as the composition of rotations of a unit vector around certain axes of a reference coordinate system. For more information, see Flow angle.
Flow direction	This boundary value is added to the tree when the flow direction specification is set to `Components`. It is input as a vector profile.
Boundary reference frame specification	Allows you to select a reference frame for the boundary. For more information, see Boundary reference frame specification.
Mach number	The Mach number is a dimensionless scalar value, which must be non-negative. It is entered as a scalar profile.

Mass flow rate	Added to the tree when mass flow rate is chosen as the Mass Flow Option. It is entered as a scalar profile. The velocity is determined from the mass flow rate using $\dot{m} = \rho \mathbf{a} \cdot \mathbf{v}$. For more information, see Mass flow rate.
Target mass flow rate	Automatically applies a correction (Eqn. (120)) to the specified Pressure boundary value in an attempt to yield the specified mass flow rate. For more information, see Target mass flow rate.
Mass flux	Added to the tree when mass flux is chosen as the Mass Flow Option. It is entered as a scalar profile. The velocity is determined from the mass flux using $\dot{m}'' = (\rho \mathbf{a} \cdot \mathbf{v})/\lvert\mathbf{a}\rvert$.
Roughness height	Specifies the equivalent sand-grain roughness height r in Eqn. (470), used to determine R^+ in roughness function. It appears when the wall surface specification is Rough, and is entered as a scalar profile.
(Static) pressure	Specify the working pressure in *Pa*. There are no nodes in the Conditions node as there is only one method of entering the value. It is entered as a pressure profile, therefore it has an additional option in the Method property—`Radial Equilibrium`.
Supersonic static pressure	This value, in *Pa*, only appears when the Ideal Gas or Real Gas model is enabled, and is only active when the flow becomes supersonic at the inlet. Further information is given in the documentation for the stagnation inlet boundary. It is entered as a pressure profile.
Total pressure	Specify the total pressure in *Pa*. There are no nodes in the Conditions node as there is only one method of entering the value. It is entered as a scalar profile.
Split ratio	When there are multiple outlet boundaries on one continuum, the fraction of the mass flow passing through each of the boundaries must be specified. The specified fraction is ignored if there is only one outlet boundary in the continuum. This is a simple value with only one method of entering the data—type in a value from `0` to `1`.

Velocity	Specifies the velocity as a vector. It is entered as a vector profile.
Velocity magnitude	Specifies the velocity in *m/s*. It is entered as a scalar profile.
Wall rotation	This node is added when rotation rate is used as the Tangential Velocity Specification. It is entered as a scalar profile.

Flow Angles

Added to the tree when the flow direction specification is set to angles. The properties of this node allow the flow direction to be defined as the composition of rotations of a unit vector around certain axes of a reference coordinate system. These rotations are known as Euler angles. Three different rotation conventions are supported to define the coordinate axes that are used, and the order in which the rotations apply.

Flow Angles Properties

Coordinate System

The coordinate system to use.

- **Cartesian** — Specifies the Cartesian coordinate system.
- **Cylindrical** — Specifies the cylindrical coordinate system.
- **Laboratory** — Specifies the laboratory coordinate system.
- **Spherical** — Specifies the spherical coordinate system.

Rotation Convention

The order in which to apply the rotations that are specified by the angles.

- **X Convention (Z-X-Z)** — For Euler angles (ϕ, θ, ψ), the first rotation is by angle ϕ about the z-axis, the second is by angle θ about the x-axis, and the third is by angle ψ about the z-axis.
- **Y Convention (Z-Y-Z)** — For Euler angles (ϕ, θ, ψ), the first rotation is by angle ϕ about the z-axis, the second is by angle θ about the y-axis, and the third is by angle ψ about the z-axis.
- **Yaw, Pitch, Roll Convention (Z-Y-X)** — For Euler angles (ϕ, θ, ψ), the first rotation is by angle ϕ about the z-axis, the second is by angle θ about the y-axis, and the third is by angle ψ about the x-axis.

Axes Convention

How to treat the axes during rotations.

- **Fixed Axes** — Makes all rotations relative to the original fixed axes.
- **Moving Axes** — Makes second and third rotations relative to the rotated axes.

Reference Vector

A unit vector that is rotated to convert angles and rotation conventions to a direction. Specified as a single three-part, comma-separated number.

Method

Selects the method to use for specifying the angle data.

Constant	Specifies the angle as a single three-part, comma-separated number. A Constant node is added as a child to this node. For two-dimensional cases, only your entries for the *x*- and *y*- directions are relevant to the calculations.
Field Function	Defines the angle using a field function (typically user-defined). A Function node is added as a child to this node.
Table (iteration)	Defines the angle as a function of iteration number. A Table (iteration) node is added as a child to this node.
Table (time)	Defines the angle as a function of physical time. A Table (time) node is added as a child to this node.

Dimensions The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.

Boundary Reference Frame Specification

When activated, the Boundary Reference Frame Specification node selects a reference frame in its properties.



Boundary Reference Frame Specification Properties

Reference Frame	Selects one of the existing reference frames in the simulation from a drop-down list.

Mass Flow Rate

Added to the tree when mass flow rate is chosen as the Mass Flow Option. It is entered as a scalar profile. The velocity is determined from the mass flow rate using $\dot{m} = \rho \mathbf{a} \cdot \mathbf{v}$.

It is possible to specify a negative mass flow rate, negative mass flux, or a negative flow direction at a mass flow inlet boundary so that it behaves like an outlet. This feature is compatible with both segregated and coupled solvers. However, do not use it if any of the following conditions apply (STAR-CCM+ has no automated checks against them):

- Outflow exceeds inflow (for example, +1 kg/s at one inlet and -2 kg/s at the other); this condition would cause the mass flow boundary to attempt to pull more mass than is available.
- The flow near the boundary is much more than Mach 0.2.
- The flow is choked.

Target Mass Flow Rate

Automatically applies a correction (Eqn. (120)) to the specified Pressure boundary value in an attempt to yield the specified mass flow rate. This object can be adjusted with its properties.

Target Mass Flow Rate Properties

Massflow Rate	The mass flow rate that is targeted for the boundary, $\dot{m}_{target}$.
Pressure Adjustment	The calculated pressure adjustment value, δp (read-only).
Minimum Allowable Avg Pressure	Minimum allowable average value of p_f (Eqn. (120)) allowed on the boundary. Used to keep the correction process bounded.
Maximum Allowable Avg Pressure	Maximum allowable average value of p_f (Eqn. (120)) allowed on the boundary. Used to keep the correction process bounded.
Under-Relaxation Factor	Under-relaxation value ω in Eqn. (122).
Update Frequency	Sets the number of iterations or time-steps that must occur before the pressure adjustment value (Eqn. (122)) is updated.

Target Mass Flow Rate Expert Properties

Verbose	Provides additional output for this value while the simulation is running. This property is useful for debugging problems when they occur.	
	Activated	Activates verbose output.
	Deactivated	Leaves output at its default.

Region Values

This section provides information about the region values that are available for flow and energy models in STAR-CCM+.

This section also provides a detailed reference of the methods of setting the values. Region, boundary, and interface values are elected by the region, boundary, or interface types, and/or by methods that are selected in corresponding nodes in the boundary conditions node.

Mass source	This node is added when the Mass Source Option region condition is activated. It is entered as a scalar profile.
Mass source pressure derivative	This node represents linearization of the Mass Source with respect to pressure, S_u in Eqn. (166) (coupled) and $S_u dV$ in Eqn. (200) (segregated). For more information, see Mass source pressure derivative.
Mass source temperature derivative	This node represents linearization of the Mass Source with respect to temperature, S_u in Eqn. (166). For more information, see Mass source temperature derivative.
Mass source velocity derivative	This node represents linearization of the Mass Source with respect to x-, y-, [z-] components of velocity, S_u in Eqn. (166). For more information, see Mass source velocity derivative.
Momentum source	This node is added when the Momentum Source Option region condition is set to `Specified`. It is input as a vector profile.
Momentum source pressure derivative	This node represents linearization of the Momentum Source with respect to pressure, S_u in Eqn. (166). For more information, see Momentum source pressure derivative.

Momentum source temperature derivative	This node represents linearization of the Momentum Source with respect to temperature, S_u in Eqn. (166). For more information, see Momentum source temperature derivative.
Momentum source velocity derivative	This node represents linearization of x-, y-, [z-] components of the Momentum Source with respect to x-, y-, [z-] components of velocity, respectively, S_u in Eqn. (166) (coupled) and $S_u dV$ in Eqn. (200) (segregated). For more information, see Momentum source velocity derivative.

Mass Source Pressure Derivative

This node represents linearization of the Mass Source with respect to pressure, S_u in Eqn. (166) (coupled) and $S_u dV$ in Eqn. (200) (segregated). The derivative of the user source term is used to linearize these equations. Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Mass Source Pressure Derivative is available when either coupled or segregated flow is used. This object is added when the Mass Source Option region condition is activated, and is entered as a scalar profile.

Mass Source Temperature Derivative

This node represents linearization of the Mass Source with respect to temperature, S_u in Eqn. (166). The derivative of the user source term is used to linearize Eqn. (166). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Mass Source Temperature Derivative is available when coupled flow is used. This object is added when the Mass Source Option region condition is activated, and is entered as a scalar profile.

Mass Source Velocity Derivative

This node represents linearization of the Mass Source with respect to x-, y-, [z-] components of velocity, S_u in Eqn. (166). The derivative of the user source term is used to linearize Eqn. (166). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Mass Source Velocity Derivative is available when coupled flow is used. This object is added when the Mass Source Option region condition is activated, and is entered as a vector profile.

Momentum Source Pressure Derivative

This node represents linearization of the Momentum Source with respect to pressure, S_u in Eqn. (166). The derivative of the user source term is used to linearize Eqn. (166). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Momentum Source Pressure Derivative is available when coupled flow is used. This object is added when the Momentum Source Option region condition is set to `Specified`, and is input as a vector profile.

Momentum Source Temperature Derivative

This node represents linearization of the Momentum Source with respect to temperature, S_u in Eqn. (166). The derivative of the user source term is used to linearize Eqn. (166). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Momentum Source Temperature Derivative is available when coupled flow is used. This object is added when the Momentum Source Option region condition is set to `Specified`, and is input as a vector profile.

Momentum Source Velocity Derivative

This node represents linearization of x-, y-, [z-] components of the Momentum Source with respect to x-, y-, [z-] components of velocity, respectively, S_u in Eqn. (166) (coupled) and $S_u dV$ in Eqn. (200) (segregated). The derivative of the user source term is used to linearize these equations. Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Momentum Source Velocity Derivative is added when the Momentum Source Option region condition is set to `Specified`.

When coupled flow is used, this object is input as a tensor profile. With segregated flow, it is input as a vector profile.

Fluid Flow and Energy Field Functions Reference

This section contains the following sub-sections:

- Common Flow and Energy Field Functions
- Porosity Field Functions including:
  - Porosity Parameter Field Functions
  - Viscous Term Field Function
  - Inertial Term Field Function
  - Region Porosity

Common Flow and Energy Field Functions

These field functions are available to both the flow and energy models. These field functions are also available to the Eulerian Multiphase models, except where indicated.

These objects all have the same base properties, with some field functions also having reference value properties (indicated by a `ref` subscript in the following descriptions).

Absolute Pressure	The sum of the static pressure and the reference pressure, which is defined in Eqn. (142). It is used to obtain density by the Ideal Gas Law, Eqn. (141).
Absolute Total Pressure	The sum of the total pressure P_t and the reference pressure p_{ref}: $P_{t,\,\mathrm{abs}} = (P_t + p_{\mathrm{ref}})$
Axial Velocity	The axial component of velocity relative to the region axis. This field function is not available in Eulerian Multiphase models for 3-D cases.
Cell Relative Velocity	The fluid flow velocity at the cell centroid relative to the cell centroid due to motion.

Convective Courant Number	The local convective Courant number, which is the ratio of the physical time-step to the mesh convection time scale, $\Delta t/(dx/V)$ or $V*\Delta t/dx$. For moving mesh cases, V is the velocity relative to the mesh. When you use this field function in a Eulerian Multiphase model, it plots the value for the phase that has the maximum convective Courant number in a cell.
Flow Angle	The boundary flow angle, $\grave{\theta}$.
Grid Relative Velocity	The velocity vector field relative to the grid. This field function is not available in Eulerian Multiphase models.
Helicity	The scalar defined by $H = V \cdot \Omega$ where V is velocity and Ω is vorticity.
Lamb Vector	The cross product of velocity and vorticity, which is defined as $\mathbf{L} \equiv \mathbf{v} \times \nabla \times \mathbf{v}$.
Lambda 2 Criterion	The scalar defined by the second eigenvalue of $(\mathbf{S}^2 + \Omega^2)$ when sorted from minimum to maximum, where $\mathbf{S}$ is the strain-rate tensor and Ω is the spin tensor. Values of $\lambda_2 < 0$ can be interpreted as vortex regions, while values equal to or greater than 0 have no physical interpretation.
Mass Flux	The face mass flux, $\dot{m}_f$.
Mass Imbalance	The imbalance of mass in any cell, which is defined as $\sum_f \dot{m}_f$ for each cell.
Pressure	The working pressure that is the solution variable in the flow models in STAR-CCM+. It is the quantity p in Eqn. (166).

Pressure Coefficient

The value that is defined as $C_p = (p - p_{\mathrm{ref}}) / \left(\frac{1}{2}\rho_{\mathrm{ref}} v_{\mathrm{ref}}^2\right)$ where p_{ref}, ρ_{ref} and v_{ref} are the reference pressure, density, and velocity, respectively.
This field function is not available in Eulerian Multiphase models.

PressureGrad

The pressure gradient ∇P times the cell volume. This becomes available when you activate the Coupled Solver and the Temporary Storage property.

Q Criterion

The scalar value defined by:

$$Q := \frac{1}{2}(\|\Omega\|^2 - \|\mathbf{S}\|^2)$$

where Ω is the spin tensor and $\mathbf{S}$ is the strain-rate tensor. Where positive, this is vorticity-dominated flow; a negative value implies strain-dominated.

Radial Velocity

The radial component of velocity relative to the region axis.
This field function is not available in Eulerian Multiphase models for axisymmetric cases.

Relative Total Pressure

The pressure that is obtained from isentropically bringing the flow to rest in the relative frame of motion.
This field function is not available in Eulerian Multiphase models.

Skin Friction Coefficient

Defined as $C_f = |\vec{\tau}_w| / \left(\frac{1}{2}\rho_{\mathrm{ref}} v_{\mathrm{ref}}^2\right)$ where ρ_{ref} and v_{ref} are the reference density and velocity, respectively.
This field function is not available in Eulerian Multiphase models.

Static Pressure

The spherical part of the stress tensor acting in fluids, which is the same as the actual thermodynamic pressure of the fluid. When the gravity model is active, it is related to the working pressure, (which becomes the piezometric pressure), by Eqn. (257) for variable density flows and by Eqn. (258) for constant density flows.

Tangential Velocity

The tangential component of velocity relative to the region axis.
This field function is not available in Eulerian Multiphase models for axisymmetric cases.

Total Pressure

The total pressure (gauge) is the pressure that is obtained from isentropically bringing the flow to rest. For an ideal gas with constant specific heat, it is defined as

$$P_t = P_{\text{abs}}\left[\left(1 + \frac{(\gamma - 1)}{2}M^2\right)^{\gamma/(\gamma-1)}\right] - P_{\text{ref}}$$

where P_{abs}, P_{ref}, M and γ are the absolute pressure, reference pressure, Mach number, and ratio of specific heats, respectively.
For alternate equations of state, or a non-constant specific heat, P_t is obtained by integrating $dh = dP/\rho$ from static conditions to total conditions.

Total Pressure Coefficient

The value that is defined as $C_p = (P_t - P_{\text{ref}})/\left(\frac{1}{2}\rho_{\text{ref}}\mathrm{v}_{\text{ref}}^2\right)$ where P_{ref}, ρ_{ref} and v_{ref} are the reference pressure, density, and velocity, respectively.
This field function is not available in Eulerian Multiphase models.

Turbulent Charge

The value that is defined as $n = \nabla \cdot \mathbf{L}$, where $\mathbf{L}$ is the lamb vector.
This field function is not available in Eulerian Multiphase models.

Velocity	The velocity vector field.
Vorticity	The vector variable with components ζ_x, ζ_y, ζ_z defined as: $\zeta_x = \frac{\partial w}{\partial y} - \frac{\partial v}{\partial z} \quad \zeta_y = \frac{\partial u}{\partial z} - \frac{\partial w}{\partial x} \quad \zeta_z = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}$ where u, v and w are velocity components in the x, y and z directions. ∇xv is the curl of the velocity field.
Vorticity Confinement Non-Dimensional Wall-Distance	d^* in Eqn. (268).
Vorticity Confinement Specific Force	$\mathbf{f}_\omega$ in Eqn. (264).
Wall Shear Stress	The wall shear stress vector, which is defined at a wall face as $\vec{\tau}_w = (-\mathbf{T} \cdot \mathbf{a}/\lvert\mathbf{a}\rvert)$, where $\mathbf{T}$ is the stress tensor (see coupled viscous fluxes or segregated viscous fluxes). For a multiphase fluid, the wall shear stress vector is the sum of the wall shear stress vectors for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted.
Wall Shear Stress of <phase>	The wall shear stress vector for the specified phase. A version of this field function is created for each phase in a multiphase continuum. This field function is volume fraction weighted for each phase.

Porosity Field Functions

This section includes the following:

- Porosity Parameter Field Functions
- Viscous Term Field Function
- Inertial Term Field Function
- Region Porosity

Porosity Parameter Field Functions

The volume porosity is defined as a scalar field function with no dimensions, and whose Definition is set to the required value (0.39 in this example).



The particle diameter (0.01 m in this example) is defined in a similar way, except that the Dimensions property is set to `Length`.



Viscous Term Field Function

The scalar field function for the viscous term is dimensioned as $Mass \cdot Length^{-3} \cdot Time^{-1}$, and the full Definition is:

```
150.0*$DynamicViscosity*pow((1-$BedPorosity),2)/
(pow($BedPorosity,3)*pow($BedSphereDiameter,2))
```

The function uses the Bed Porosity and Bed Sphere Diameter functions that were defined previously.



This field function is applied to the Porous Viscous Resistance node as an isotropic tensor.

- Select the Physics Values > Porous Viscous Resistance node on the porous region.



- In the *Properties* window, set the Method to Isotropic Tensor using the drop-down menu. The Isotropic Tensor child node is shown underneath the Porous Viscous Resistance node.

- Select the Isotropic Tensor > Isotropic Component node and, in the *Properties* window, set the Method to Field Function.



- Finally, select the Isotropic Component > Field Function node and choose Ergun Viscous Term from the Scalar Function drop-down menu in the *Properties* window.



Inertial Term Field Function

The scalar field function for the inertial term is dimensioned as $Mass \cdot Length^{-4}$, and the Definition is:

```
(1.75*$Density*(1-$BedPorosity))/
(pow($BedPorosity,3)*$BedSphereDiameter)
```


Servers
ergun
simulation
scene/plot
Tools
Annotations
Colormaps
Coordinate Systems
Data Set Functions
Field Functions
Bed Porosity
Bed Sphere Diameter
Ergun Inertial Term
Ergun Viscous Term
Absolute Pressure
Absolute Total Pressure
Ergun Inertial Term - Properties
Properties
Type
Scalar
Dimensions
Mass/Length^4
Function Name
ErgunInertialTerm
Definition
(1.75*$Density*(1-$BedPoros...
Expert
Ignore Boundary Values
Assembly code
(/ (* (* 1.75 ${Density}) (- 1 $...

The fluid density is included in this field function. Any variations in the fluid density therefore affect the inertial resistance. This function is applied to the Porous Inertial Resistance using the isotropic tensor method as before:



Region Porosity

For an unsteady calculation, the volume porosity χ can be set on the region using the Bed Porosity field function. To set the volume porosity:

- Select the Physics Values > Porosity node of the porous region (PackedBed in this example).
- In the *Properties* window, set the Method property to Field Function using the drop-down menu.

- Select the Porosity > Field Function node and, in the *Properties* window, set the Scalar Function property to Bed Porosity.



Fluid Flow Formulation

This section contains the following sub-sections on formulations used in modeling fluid flow:

- Polynomial Density Formulation
- Ideal Gas Formulation
- Real Gas Formulation
- Thermal Equilibrium Formulation
- Passive Scalar Formulation
- Coupled Flow and Energy Formulation
- Segregated Flow Formulation
- Segregated Fluid Energy Formulation
- Gravity Formulation
- Axisymmetric Swirl Model Formulation
- Porous Media Formulation
- Vorticity Confinement Formulation
- Fan Interface Formulation
- Fan Momentum Source Formulation
- Isentropic Flow Formulation
- Adjoint Formulation
- Virtual Disk Formulation

Polynomial Density Formulation

The Polynomial Density model makes the assumption that the density is a function of temperature only, and can be expressed as the following polynomial function of degree n:

$$\rho = \sum_{k=0}^{n_i} a_{i,k} \cdot T^{bi,k} \quad ; \quad (Tr \le T \le Tr_{i+1}) \tag{140}$$

The Polynomial Density model is available for gases, liquids, and solids.

Ideal Gas Formulation

The Ideal Gas model, available for gases only, uses the ideal gas law to express density as a function of temperature and pressure:

$$\rho = \frac{p_{\text{abs}}}{RT} \tag{141}$$

where:

- ρ is the density.
- p_{abs} is absolute pressure:

$$p_{\text{abs}} = p_{\text{static}} + p_{\text{ref}} \tag{142}$$

  where p_{ref} is the reference pressure. For ideal gas flows with gravity, the static pressure p_{static} is related to the working pressure p by Eqn. (257). Otherwise they are the same.
- R is the specific gas constant:

$$R = R_u / M \tag{143}$$

  where R_u is the universal gas constant [8314.4621 J/kmol K] and M is the molecular weight.
- T is the temperature.

"Incompressible" Ideal Gas

The dependency of density on the working pressure p can be removed by activating the Incompressible option. In this case, density is expressed as a function of temperature and reference pressure p_{ref}:

$$\rho = \frac{p_{\text{ref}}}{RT} \tag{144}$$

Real Gas Formulation

This section describes six variations on formulation for real gas models:

- Peng-Robinson Formulation
- Redlich-Kwong Formulation
- Soave-Redlich-Kwong Formulation
- Modified Soave-Redlich-Kwong Formulation
- Van der Waals Formulation
- Equilibrium Air Formulation

Peng-Robinson Formulation

The Peng-Robinson model attempts to simulate the *p-v-T* behavior of real gases. At high pressure and low temperature, this behavior deviates from that predicted by the Ideal Gas equation. This deviation is because as the gas density increases, the molecules of the gas begin to take up a significant portion of the total volume. In addition, under these conditions the intermolecular attractive forces become important as well.

The Peng-Robinson equation is:

$$p = \frac{R_u T}{(v-b)} - \frac{a\alpha(T_r)}{(v^2 + 2bv - b^2)} \tag{145}$$

where:

- p is the pressure.
- v the specific volume.
- R_u the universal gas constant [8314.4621 J/kmol K].
- T the temperature.
- The function $\alpha(T_r)$ is:

$$\alpha(T_r) = [1 + (0.37464 + 1.54226\omega - 0.2699\omega^2)(1 - T_r^{0.5})]^2 \tag{146}$$

  where:

  - T_r is the reduced temperature. $T_r = T/T_c$.
  - ω is the acentric factor of the gas.

- As with the van der Waals equation, constants a and b are a measure of the attractive forces and the co-volume of the particles, respectively. These constants are likewise evaluated from experimentally measured values of critical pressure and temperature, p_c and T_c, and found to be:

$$a = \frac{0.4572 R_u^2 T_c^{2.5}}{p_c} \qquad b = \frac{0.0778 R_u T_c}{p_c} \tag{147}$$

Redlich-Kwong Formulation

The Redlich-Kwong model attempts to simulate the *p-v-T* behavior of real gases. At high pressure and low temperature, this behavior deviates from that predicted by the Ideal Gas equation. This deviation is because, as the gas density increases, the molecules of the gas begin to take up a significant portion of the total volume. In addition, under these conditions the intermolecular attractive forces become important as well.

The Redlich-Kwong equation is:

$$p = \frac{R_u T}{(v-b)} - \frac{a}{[T^{0.5} v(v+b)]} \tag{148}$$

where:

- p is the pressure.
- v the specific volume.
- R_u the universal gas constant [8314.4621 J/kmol K].
- T the temperature.
- As with the van der Waals equation, constants a and b are a measure of the attractive forces and the co-volume of the particles, respectively. These constants are likewise evaluated from experimentally measured values of critical pressure and temperature, p_c and T_c, and found to be:

$$a = \frac{0.4275 R_u^2 T_c^{2.5}}{p_c} \qquad b = \frac{0.0867 R_u T_c}{p_c} \tag{149}$$

Though it contains only two constants and is empirical in nature, the Redlich-Kwong equation of state appears to have considerable accuracy over a wide range of *p-v-T* conditions. It has been shown to be as accurate as the eight-constant Benedict-Webb-Rubin equation of state in several instances. Hence the Redlich-Kwong equation of state has the advantages of simplicity and reasonable accuracy, particularly when compared to other two-constant equations of state.

Soave-Redlich-Kwong Formulation

Soave [48] modified the Redlich-Kwong equation, replacing the $T^{0.5}$ power in the denominator (Eqn. (148)) with a more complicated temperature-dependent expression. This modification improves predictions for liquid densities and vapor-liquid equilibrium, and is:

$$p = \frac{R_u T}{v-b} - \frac{a\alpha(T_r)}{v(v+b)} \tag{150}$$

where:

- p is the pressure.
- v the specific volume.
- R_u the universal gas constant [8314.4621 J/kmol K].
- T the temperature.
- The function $\alpha(T_r)$ is:

$$\alpha(T_r) = [1 + (0.48 + 1.574\omega - 0.176\omega^2)(1 - T_r^{0.5})]^2 \tag{151}$$

where:

- T_r is the reduced temperature.
- ω is the acentric factor of the gas.
- Constants a and b are measures of the attractive forces and the volume of the particles, respectively, given by:

$$a = \frac{0.42748 R_u^2 T_c^{2.5}}{p_c} \qquad b = \frac{0.08664 R_u T_c}{p_c} \tag{152}$$

Modified Soave-Redlich-Kwong Formulation

Graboski and Daubert [27], [28] modified the Soave-Redlich-Kwong model to improve its behavior for phase equilibrium calculations that involve non-polar molecules. All equations remain identical to the Soave-Redlich-Kwong model, except the function $\alpha(T_r)$ which changes to:

$$\alpha(T_r) = [1 + (0.48503 + 1.55171\omega - 0.15613\omega^2)(1 - T_r^{0.5})]^2 \tag{153}$$

Van der Waals Formulation

The van der Waals model attempts to simulate the *p-v-T* behavior of real gases. At high pressure and low temperature, this behavior deviates from that predicted by the Ideal Gas equation. This deviation is because as the gas density increases, the molecules of the gas begin to take up a significant portion of the total volume. In addition, under these conditions the intermolecular attractive forces become important as well.

The van der Waals equation is:

$$(p + a/v^2)(v - b) = R_u T \tag{154}$$

where:

p is the pressure.

v the specific volume.

R_u the universal gas constant [8314.4621 J/kmol K].

T the temperature.

Van der Waals replaced specific volume in the ideal gas relation, $pv = R_u T$, with $(v - b)$ to account for the volume that the particles of the gas occupy, while replacing the pressure with the term $(p + a/v^2)$. The constant b is the co-volume of the particles, and the constant a is a measure of the attractive forces.

The constants a and b are evaluated based on the experimental observation that, in p-v space, the critical isotherm $T = T_c = \text{constant}$ has zero slope and an inflection point at the critical state, $p = p_c$ and $v = v_c$. Thus the following can be written:

$$\frac{dp}{dv}|T_c = 0 \tag{155}$$

and:

$$\frac{d^2p}{dv^2}|T_c = 0 \tag{156}$$

which leads to the following expressions for a and b:

$$a = \frac{27}{64}\frac{{R_u}^2 T_c^2}{p_c} \qquad b = \frac{1}{8}\frac{R_u T_c}{p_c} \tag{157}$$

These equations allow a and b to be determined from experimentally measured values of critical pressure and temperature, p_c and T_c.

Equilibrium Air Formulation

The equilibrium air equation of state is used for the simulation of high-temperature air, where the effects of molecular dissociation, internal energy excitation, and ionization become significant. For situations where the time scale of these effects is much shorter than the time scale of the flow, these effects can be treated in an equilibrium manner. In this case, the properties of the flow, such as the density, specific heat, and transport properties, become expressions of two thermodynamic variables. These expressions typically take the forms of curve fits that account for the chemical reactions and energy modes present at high temperature. In particular, the curve fits of Gupta and others [29] were used. These curve fits specify the compressibility factor, specific heat, enthalpy, viscosity, and thermal conductivity as a function of temperature and pressure. The curve fits are valid for temperatures less than 30,000 K and for pressures between 10^{-4} and 10^2 atmospheres.

The equilibrium air equation of state provides the compressiblity factor as a field function. In the context of this model, the compressiblity factor represents the ratio of the undissociated molecular weight to the mean molecular weight of air [29]:

$$Z(T, P) = \frac{M_o}{\overline{M}} \tag{158}$$

where M_o represents the molecular weight of undissociated air and $\overline{M}$ represents the mean molecular weight of air at the given temperature and pressure. Because the mean molecular weight is a function of the composition, the compressibility factor can be used as a measure of the degree of dissociation in the air, with higher values signifying a large degree of dissociation and ionization.

Thermal Non-Equilibrium Formulation

Thermal non-equilibrium in this context refers to non-equilibrium between the energy modes of a gas. For a gas, the internal energy is composed of translational, rotational, vibrational, and electronic modes. In equilibrium, these energy modes can be adequately described by a single characteristic temperature. For thermal non-equilibrium, the translational and rotational energy modes are described by one temperature, while the vibrational and electronic energy modes are described by an additional vibrational-electronic temperature. The vibrational-electronic temperature is determined by solving an additional energy equation describing the conservation of vibrational-electronic energy, given by Gnoffo and others [26] (simplified to non-ionizing flows by Lockwood [37]):

$$\frac{\partial \rho e_{ve}}{\partial t} + \nabla \cdot (\rho e_{ve} \mathbf{U}) = \nabla \cdot (k_{ve} \nabla T) + \nabla \cdot \left(\rho \sum_s h_{ve,s} D_s \nabla y_s \right) + \sum_{s = \mathrm{mol}} \rho_s \frac{(e^*_{v,s} - e_{v,s})}{<\tau_s>} + \sum_{s = \mathrm{mol}} \dot{\omega}_s e_{ve,s} \tag{159}$$

The exchange of vibrational-electronic energy and translational-rotational energy is modeled through a relaxation term. Gnoffo and others [26] present this relaxation, and provide a simplification which allows the calculation to include fewer species-dependent parameters:

$$e^*_{v,s} - e_{v,s} \approx C^s_{v,v}(T - T_v) \tag{160}$$

$$\sum_{s = \mathrm{mol}} \frac{\rho_s C^s_{v,v}}{<\tau_s>} \approx \frac{\rho C^s_{v,v}}{\bar{\tau}_v} \tag{161}$$

To describe the relationship between vibrational-electronic enthalpy and temperature, the vibrational specific heat must be specified. This vibrational specific heat is specified by simultaneously setting the total specific heat and the translational-rotational specific heat. For this model, the translational-rotational energy modes are assumed fully excited $(F - E)$, giving a constant translational-rotational specific heat $(C_{p,tr})$. With these two values specified, the vibrational specific heat and enthalpy are given by [26]:

$$C_{p,\,ve}(T_{ve}) = C_{p,\,\text{total}}(T_{ve}) - C_{p,\,tr}^{F-E} \tag{162}$$

$$h_{ve}(T_{ve}) = h_{\text{total}}(T_{ve}) - C_{p,\,tr}^{F-E}(T_{ve} - T_{\text{ref}}) - h_o \tag{163}$$

Passive Scalar Formulation

This section describes the source term formulation when passive scalar is included in the simulation. This formulation includes:

- The transport equation for passive scalar
- Wall boundary conditions

Transport Equation

The transport equation for the passive scalar component, ϕ_j, is:

$$\frac{\partial}{\partial t}\int_{\tilde{V}} \rho\phi_j d\tilde{V} + \oint_A \rho\phi_j(\mathbf{v} - \mathbf{v}_g)\cdot d\tilde{\mathbf{a}} = \oint_A \mathbf{D}\cdot d\tilde{\mathbf{a}} + \int_{\tilde{V}} S_{\phi_j} d\tilde{V} \tag{164}$$

where:

- j is the component index.
- $d\tilde{V} = \alpha_i \chi dV$, where α_i is the volume fraction of phase i (set to 1 for single phase flows) and χ is the void fraction.
- $\mathbf{D} = \left(\frac{\mu}{\sigma} + \frac{\mu_t}{\sigma_t}\right)\nabla\phi_j$ is the linear eddy-diffusivity model, the default. Alternatively, if used with the Reynolds Stress model, $\mathbf{D}$ can be modeled using the Generalized Gradient Diffusion Hypothesis (GGDH).
  - μ is the viscosity.
  - μ_t is the turbulent viscosity.
  - σ is molecular Schmidt number—a material property.
  - σ_t is the turbulent Schmidt number—assumed to have a value of 0.9, consistent with the turbulent Prandtl number used for energy.
- $d\tilde{\mathbf{a}} = \alpha_i \chi d\mathbf{a}$
- S_{ϕ_j} is a source term for passive scalar component j. The source term is not multiplied internally by the volume fraction. ϕ_j is assumed to be positive-definite.

Wall Boundary Conditions

At walls, use either a specified flux or a specified value condition. See Setting Passive Scalar at Boundaries.

Coupled Flow and Energy Formulation

This section presents the basic flow and energy equations in their coupled form. The methods that are described in finite volume discretization are extended to the coupled system. Specifically described are:

- Governing equations in continuous integral form
- Preconditioned equations
- Preconditioned equations in discrete form
- Inviscid fluxes
- Viscous fluxes
- Viscous flux discretization
- Explicit time-stepping
- Implicit time-stepping
  - Dual time-stepping
- Pseudo-time step calculation
- Time-step calculation
- Non-reflecting boundary condition

Also included in this section are:

- Nomenclature
- Bibliography

Governing Equations in Continuous Integral Form

The Navier-Stokes equations in Cartesian integral form for an arbitrary control volume V with differential surface area $d\mathbf{a}$ may be written:

$$\frac{\partial}{\partial t}\int_{\tilde{V}} \mathbf{W}\, d\tilde{V} + \oint [\mathbf{F} - \mathbf{G}] \cdot d\tilde{\mathbf{a}} = \int_{\tilde{V}} \mathbf{H}\, d\tilde{V} \tag{165}$$

where:

$$\mathbf{W} = \begin{bmatrix} \rho \\ \rho\mathbf{v} \\ \rho E \end{bmatrix} \quad , \quad \mathbf{F} = \begin{bmatrix} \rho(\mathbf{v}-\mathbf{v}_g) \\ \rho(\mathbf{v}-\mathbf{v}_g)\otimes\mathbf{v}+p\mathbf{I} \\ \rho(\mathbf{v}-\mathbf{v}_g)H+p\mathbf{v}_g \end{bmatrix}$$
$$\mathbf{G} = \begin{bmatrix} 0 \\ \mathbf{T} \\ \mathbf{T}\cdot\mathbf{v}+\dot{\mathbf{q}}'' \end{bmatrix} \quad , \quad \mathbf{H} = \begin{bmatrix} S_u \\ \mathbf{f}_r+\mathbf{f}_g+\mathbf{f}_p+\mathbf{f}_u+\mathbf{f}_\omega+\mathbf{f}_L \\ S_u \end{bmatrix} \tag{166}$$

and

- ρ is the density of the fluid.
- $\mathbf{v}$ is the velocity of the fluid
- E is total energy per unit mass
- p is pressure
- $d\tilde{V} = \alpha_i\chi dV$, where α_i is the volume fraction of phase i (set to 1 for single phase flows) and χ is the void fraction.
- $d\tilde{\mathbf{a}} = \alpha_i\chi d\mathbf{a}$
- $\mathbf{T}$ is the viscous stress tensor.
- $\dot{\mathbf{q}}''$ is the heat flux vector.
- $\mathbf{v}_g$ is the grid velocity vector.

Total energy is related to the total enthalpy H by:

$$E = H - p/\rho \tag{167}$$

where:

$$H = h + |\mathbf{v}|^2/2 \tag{168}$$

and $h = C_pT$. $\mathbf{H}$ is the vector of body forces.

For a full list of terms and their definitions, see Coupled Flow and Energy Nomenclature.

Preconditioned Equations

To provide efficient solution of both compressible and incompressible flows at all speeds, a preconditioning matrix Γ is incorporated into Eqn. (165) as follows [52]:

$$\Gamma\frac{\partial}{\partial t}\int_V \mathbf{Q}dV + \oint(\mathbf{F}-\mathbf{G})\cdot d\mathbf{a} = \int_V \mathbf{H}dV \tag{169}$$

where:

$$\Gamma = \begin{bmatrix} \theta & 0 & \rho_{,T} \\ \theta\mathbf{v} & \rho\mathbf{I} & \rho_{,T}\mathbf{v} \\ \theta H - \delta\,\rho\mathbf{v} & & \rho_{,T}H + \rho C_p \end{bmatrix} \tag{170}$$

and:

$$\mathbf{Q} = \left[p\ \mathbf{v}\ T\right]^{\mathrm{T}}$$

is the dependent vector of primary variables.

In Eqn. (170), $\rho_{,T}$ is the derivative of density with respect to temperature at constant pressure and $\delta = 0$ or 1. For an ideal gas $\rho_{,T} = -p/RT$, $\delta = 1$ and this matrix becomes a member of Turkel's family of preconditioners [49]. For an incompressible fluid, both $\rho_{,T}$ and δ are zero. The parameter θ is defined as:

$$\theta = (1/U_r^2) - (\rho_{,T}/\rho C_p) \tag{171}$$

The reference velocity U_r appearing in Eqn. (171) is chosen such that the eigenvalues of the system remain well conditioned with respect to the convective and diffusive time scales [52]. This is accomplished by limiting U_r such that it does not go below either the local convection or diffusion velocities. An additional limitation on U_r considers local pressure differences and is designed to increase numerical stability in stagnation regions by prohibiting amplification of pressure perturbations there. Thus, the restriction on U_r becomes:

$$U_r = \min\left[\max\left(|\mathbf{v}|, \nu/\Delta x, \varepsilon\sqrt{\delta p/\rho}, U_{r,\,\min}\right), U_{r,\,\max}\right] \tag{172}$$

where Δx is the inter-cell length scale over which the diffusion occurs and δp is the pressure difference between adjacent cells. For compressible flows U_r is further limited such that it does not exceed the local acoustic speed c. The scaling parameter ε appearing in Eqn. (172) is set to 2.

The non-preconditioned Navier-Stokes equations are recovered exactly from Eqn. (165) by setting $1/U_r^2$ to $\rho_{,p}$, the derivative of density with respect to pressure. In this case, Γ reduces to the Jacobian $\partial\mathbf{w}/\partial\mathbf{q}$.

Although Eqn. (169) is conservative for steady-state flows, it is not strictly conservative for time-dependent flows [49]. A dual time-stepping method is used to recover a time-accurate solution of the preconditioned equations for unsteady flows with implicit time stepping.

Preconditioned Equations in Discrete Form

Applying Eqn. (169) to a cell-centered control volume for cell-0, one obtains the following discretized system:

$$V_0 \Gamma_0 \frac{\partial \mathbf{Q}_0}{\partial t} + \sum_f (\mathbf{f}_f + \mathbf{g}_f) \cdot \mathbf{a} = \mathbf{h} V_0 \tag{173}$$

where the summation is over the total number of faces defining cell-0, and $\mathbf{f}_f$ and $\mathbf{g}_f$ are the inviscid and viscous fluxes through face-*f*. V_0 is the volume of cell-0 and $\mathbf{\Gamma}_0$ is the preconditioning matrix that is evaluated in cell-0.

Inviscid Fluxes

This section describes two schemes for evaluating inviscid fluxes:

- The Weiss-Smith preconditioned Roe's flux-difference splitting scheme
- Liou's AUSM+ flux-vector splitting scheme

Mass-Flux Pressure Dissipation

As suggested by Sankaran in [31], the mass-flux pressure dissipation term used in the following sections has been modified to provide improved numerical dissipation for unsteady flows where both sound and advected quantities (such as total enthalpy and entropy) must be simulated accurately. Thus, the numerical dissipation is tuned for the entire range of Strouhal numbers, from zero (steady state) to high values.

This formulation does not require user input for a reference length-scale (as [31] does), but uses local mesh and flow information.

Unlike [31], the STAR-CCM+ formulation uses mass-flux dissipation corrections based on local mesh and flow conditions and does not require you to input a reference length-scale. These corrections are expressed in the preconditioning matrix Γ, used in Eqn. (174).

Weiss-Smith Preconditioned Roe's Scheme

By default, the inviscid fluxes appearing in Eqn. (173) are evaluated by using the Weiss-Smith preconditioned Roe's flux-difference splitting scheme as described in [52] and [51]. This scheme acknowledges that $\mathbf{f}_f$ contains characteristic information propagating through the domain with speed and direction according to the eigenvalues of the system. By splitting $\mathbf{f}_f$ into parts, where each part contains information traveling in a particular direction (that is, characteristic information), and upwind differencing the split fluxes in a manner consistent with their corresponding eigenvalues, the following expression for the value of the flux at each face is obtained:

$$\mathbf{f}_f = \frac{1}{2}(\mathbf{f}_0 + \mathbf{f}_1) - \frac{1}{2}\Gamma|\mathbf{A}|\,\Delta\mathbf{Q} \tag{174}$$

where the subscripts "0" and "1" refer to the cells on either side of face-*f*, Γ is preconditioning matrix, and $\Delta\mathbf{Q}$ is:

$$\Delta\mathbf{Q} = \mathbf{Q}_1^r - \mathbf{Q}_0^r \tag{175}$$

where $\mathbf{Q}_0^r$ and $\mathbf{Q}_1^r$ are the solution vectors from cell-0 and cell-1 interpolated to the face using reconstruction gradients.

The matrix $|\mathbf{A}|$ is defined as:

$$|\mathbf{A}| = \mathbf{M}\,|\Lambda|\,\mathbf{M}^{-1} \tag{176}$$

where Λ is the diagonal matrix of eigenvalues and $\mathbf{M}$ is the modal matrix that diagonalizes $\Gamma^{-1}(\partial\mathbf{f}/\partial\mathbf{Q})$.

By using these reconstructed solution vectors, the discretization scheme becomes formally second-order accurate. In this form, Eqn. (174) can be viewed as a second-order central-difference plus an added matrix dissipation, combining the pressure and velocity differential dissipations. This added matrix dissipation term is not only responsible for producing an upwinding of the convected variables, and of pressure and flux velocity in supersonic flow. It also provides the pressure-velocity coupling that is required for stability and efficient convergence of low-speed and incompressible flows.

The mass-flux pressure dissipation term is modified to provide improved numerical dissipation for the entire range of Strouhal numbers, from Strouhal number zero (steady state) to high Strouhal numbers.

This modification affects the pressure differences term of the mass-flux dissipation. It is inspired by the method proposed by Hosangadi *et al.* [31], but with the following differences:

- It improves the accuracy of simulation for unsteady flows that require both sound and advected quantities (like total enthalpy and entropy).
- It allows you to specify a limit to the allowed amount of dissipation correction, for more robust simulation. The Unsteady Preconditioning Max Factor sets this limit. This factor has a range from 0 (no correction) to 1 (full correction).
- It bases the mass-flux dissipation corrections on local mesh and flow conditions.
- It does not require you to provide a reference length scale.

Liou's AUSM+ Scheme

Inviscid fluxes can be evaluated using Liou's AUSM+ flux-vector splitting scheme, as described in [35], [17] and other references. This scheme is also based on the upwind concept, which is applied consistently on the two physically distinct parts of the inviscid flux: convective and pressure. Thus, the numerical flux is written as:

$$\mathbf{f_f} = \mathbf{f}^{(\mathbf{c})} + \mathbf{P} = \mathbf{m_i^+}(1, u, v, \ldots, H)_0^T + \mathbf{m_i^-}(1, u, v, \ldots, H)_1^T + \mathbf{P}_i \tag{177}$$

Here $\mathbf{m_i}$ is the mass flux, $\mathbf{m_i^+}$ is defined as $0.5(\mathbf{m_i} + |\mathbf{m_i}|)$, $\mathbf{m_i^-}$ is defined as $0.5(\mathbf{m_i} - |\mathbf{m_i}|)$, and $\mathbf{P}_i$ is the pressure flux. The mass flux and pressure flux are computed based on local flow characteristics, emphasizing on correct information propagation inside the fluid for convective and acoustic processes.

The mass-flux pressure dissipation term is modified to provide improved numerical dissipation for the entire range of Strouhal numbers, from Strouhal number zero (steady state) to high Strouhal numbers.

This modification follows the method proposed by Hosangadi *et al.* [31], but with the following differences:

- It bases the mass-flux dissipation corrections on local mesh and flow conditions.
- It does not require you to provide a reference length scale.

Viscous Fluxes

The viscous fluxes are written in terms of the stress tensor $\mathbf{T}$.

In *laminar flow*, the stress tensor $\mathbf{T}$ is:

$$\mathbf{T} = \mathbf{T}_l = \mu\left[\nabla\mathbf{v} + \nabla\mathbf{v}^T - \frac{2}{3}(\nabla \cdot \mathbf{v})\mathbf{I}\right] \tag{178}$$

In *turbulent flow*, the complete stress tensor is:

$$\mathbf{T} = \mathbf{T}_l + \mathbf{T}_t \tag{179}$$

where the turbulent stress tensor $\mathbf{T}_t$ is also known as the Reynolds stress tensor. Unless a Reynolds stress transport model is being used, the Boussinesq approximation is invoked such that:

$$\mathbf{T} = \mu_{eff}\left[\nabla\mathbf{v} + \nabla\mathbf{v}^T - \frac{2}{3}(\nabla \cdot \mathbf{v})\mathbf{I}\right] \tag{180}$$

where the effective viscosity is $\mu_{eff} = \mu + \mu_t$, the sum of the laminar and turbulent viscosities.

The heat flux $\dot{\mathbf{q}}''$ is computed as:

$$\dot{\mathbf{q}}'' = -k_{\text{eff}} \nabla T \tag{181}$$

Viscous Flux Discretization

Boundary Faces

For *slip walls*, the viscous shear force at the wall boundary face, $(\mathbf{T} \cdot \mathbf{a})_f$, set to zero.

For *no-slip walls* in *laminar* flow, Eqn. (2557) is used to represent the velocity gradient tensor at the face:

$$\nabla \mathbf{v}_f = \Delta \mathbf{v} \otimes \vec{\alpha} + \nabla \mathbf{v}_0 - (\nabla \mathbf{v}_0 \cdot \mathbf{ds}) \otimes \vec{\alpha} \tag{182}$$

where:

$$\Delta \mathbf{v} = \mathbf{v}_f - \mathbf{v}_0 \tag{183}$$

and $\mathbf{v}_f$ is the specified wall tangential motion.

The stress tensor at the face is therefore:

$$\mathbf{T}_f = \mu \left[\nabla \mathbf{v} + \nabla \mathbf{v}^T - \frac{2}{3}(\nabla \cdot \mathbf{v})\mathbf{I} \right]_f \tag{184}$$

For *no-slip walls* in *turbulent flow*, it is assumed that only the component of velocity parallel to the wall is of interest. A linear relationship between the wall shear force and the wall-parallel component of the difference in velocity between the wall and the cell is hypothesized:

$$(\mathbf{T} \cdot \mathbf{a})_f = -\gamma(\mathbf{v}^t - \mathbf{v}_f^t) \tag{185}$$

Using the definitions of wall shear stress magnitude ($\tau_w = \rho u^{*2}$) and reference velocity ($u^+ = |\mathbf{v}_p| / u^*$), the coefficient of proportionality is:

$$\gamma = |\mathbf{a}| \frac{\rho u^*}{u^+} \tag{186}$$

The reference velocity u^* is computed according to the specific turbulence model. The value of u^+ is obtained (as a function of y^+) from the appropriate wall law.

Explicit Time-Stepping

An explicit multi-stage time-stepping scheme [32] is used to discretize the time-derivative in Eqn. (169). The solution is advanced from time t to time $t + \Delta t$ with an m-stage Runge-Kutta scheme:

$$\begin{aligned} \mathbf{Q}^{(0)} &= \mathbf{Q}_t \\ \mathbf{Q}^{(i)} &= \mathbf{Q}^{(0)} - \alpha_i \Delta t \Gamma^{-1} \mathbf{R}^{(i-1)} \\ \mathbf{Q}_{t+\Delta t} &= \mathbf{Q}^{(m)} \end{aligned} \tag{187}$$

where $i = 1, 2, \ldots, m$ is the stage counter for the m-stage scheme and α_i is the multi-stage coefficient for the ith stage. The residual $\mathbf{R}^{(i)}$ is computed from the intermediate solution $\mathbf{Q}^{(i)}$ and, for Eqn. (173), is:

$$\mathbf{R}^{(i)} = \frac{1}{V} \sum_f \{ \mathbf{f}(\mathbf{Q}^{(i)}) - \mathbf{g}(\mathbf{Q}^{(i)}) \} \tag{188}$$

Residual Smoothing

Residual smoothing is a mechanism for increasing the explicit time-step size by removing the high-wave number oscillations from the residuals. The residuals for cell i are filtered through a Laplacian smoothing operator:

$$\bar{\mathbf{r}}_i = \mathbf{r}_i + \varepsilon \sum_{\text{neighbors}} (\bar{\mathbf{r}}_j - \bar{\mathbf{r}}_i) \tag{189}$$

where the subscript j refers to the neighboring cells.

This equation is implemented using a Jacobi iteration as follows:

$$\bar{\mathbf{r}}_i^m = \frac{\mathbf{r}_i^0 + \varepsilon \sum_{\text{neighbors}} \bar{\mathbf{r}}_j^{m-1}}{1 + \varepsilon \sum_{\text{neighbors}} 1} \tag{190}$$

where $\mathbf{r}_i^0$ represents the original (unsmoothed) residuals and $\bar{\mathbf{r}}_i^m$ represents the smoothed residuals after iteration m. The default number of smoothing operations is 2 with the under-relaxation factor $\varepsilon = 0.5$. This is typically sufficient to allow the Courant number to be doubled.

Implicit Time-Stepping

An Euler implicit discretization in time of the governing equations Eqn. (165) is combined with a Newton-type linearization of all fluxes to produce the following linearized system in delta form:

$$\left[\mathbf{D}_i + \sum_{j}^{\text{Nfaces}} \mathbf{S}_{j,k}\right] \Delta\mathbf{Q} = -\mathbf{R}_i^n \tag{191}$$

where $\Delta\mathbf{Q} \equiv \mathbf{Q}^{n+1} - \mathbf{Q}^n$. The center and off-diagonal coefficient matrices, $\mathbf{D}_i$ and $\mathbf{S}_{j,k}$, and the residual vector $\mathbf{r}_i^n$ are:

$$\mathbf{D}_i = \frac{V}{\Delta t}\Gamma + \sum_{i}^{\text{Nfaces}} \mathbf{S}_{j,k} \tag{192}$$

$$\mathbf{S}_{j,k} = \left(\frac{\partial \mathbf{F}_j}{\partial \mathbf{Q}_k} - \frac{\partial \mathbf{G}_j}{\partial \mathbf{Q}_k}\right) \tag{193}$$

$$\mathbf{R}_i^n = \sum_{i}^{\text{Nfaces}} (\mathbf{F}_j - \mathbf{G}_j)^n \tag{194}$$

Dual Time-Stepping

Since time-derivative preconditioning destroys the time accuracy of the governing equations, solving for unsteady flows requires a dual time-stepping, with inner iterations in pseudo-time. A preconditioned pseudo-time derivative term is introduced into conservation Eqn. (165) as follows:

$$\frac{\partial}{\partial t}\int_V \mathbf{W}\,dV + \Gamma\frac{\partial}{\partial \tau}\int_V \mathbf{Q}\,dV + \oint [\mathbf{F} - \mathbf{G}] \cdot d\mathbf{a} = \int_V \mathbf{H}\,dV \tag{195}$$

where t denotes the physical time-step which is applied uniformly to all cells in the domain. τ is a locally varying pseudo-time used in the time-marching procedure. As inner iterations progress to convergence, the second term of Eqn. (195) vanishes and Eqn. (165) is recovered.

The pseudo-time derivative is driven to zero with the following iterative algorithm:

$$\mathbf{Q}^{(0)} = \mathbf{Q}_{\tau}$$

$$\left[\Gamma + \frac{3}{2}\frac{\Delta\tau}{\Delta t}\frac{\partial \mathbf{W}}{\partial \mathbf{Q}}\right]\Delta\mathbf{Q} = -\alpha_i \Delta\tau\left\{\mathbf{R}^{(i-1)} + \frac{1}{2\Delta t}(3\mathbf{W}^{(i-1)} - 4\mathbf{W}^{(n)} + \mathbf{W}^{(n-1)})\right\} \tag{196}$$

$$\mathbf{Q}_{\tau+\Delta\tau} = \mathbf{Q}^{(m)}$$

where i is the inner iteration counter, n represents any given physical time level, and $\Delta\mathbf{Q} \equiv \mathbf{Q}^{(i)} - \mathbf{Q}^{(0)}$. Throughout the pseudo-time iterations, $\mathbf{W}^{(n)}$ and $\mathbf{W}^{(n-1)}$ are held constant and $\mathbf{W}^{(i-1)}$ is computed from $\mathbf{Q}^{(i-1)}$. When inner iterations converge, the solution at the next physical time level $\mathbf{W}^{(i+1)}$ is given by $\mathbf{W}(\mathbf{Q}_{\tau+\Delta\tau})$. For each inner iteration, the linear system (block coupled for the explicit scheme or fully coupled for the implicit scheme) is solved to compute $\Delta\mathbf{Q}$.

Pseudo-Time Step Calculation

The local pseudo time-step $\Delta\tau$ associated with the iterations of all steady simulations, or within a time-step of a transient simulation, is computed considering the Courant number and Von Neumann stability conditions:

$$\Delta\tau = \min\left(\frac{\mathrm{CFL}\ V(\mathbf{x})}{\lambda_{\max}(\mathbf{x})}, \frac{\sigma\Delta x^2(\mathbf{x})}{\nu(\mathbf{x})}\right) \tag{197}$$

where CFL is the user-specified Courant number (dimensionless), V is the cell volume (m^3), σ is the Von Neumann number ($\sigma \approx 1$), Δx is a characteristic cell length scale (m) and ν is kinematic viscosity (m^2/s). $\lambda_{\max}$ is the maximum eigenvalue of the system as follows:

$$\lambda_{\max} = |\mathbf{u}\cdot\mathbf{a}| + c|\mathbf{a}| \tag{198}$$

where $\mathbf{u}$ is velocity and c the speed of sound.

The maximum eigenvalue is determined from the pre-conditioned system, Eqn. (173). The field function Convective Courant Number provides the local convective Courant number for a given time-step.

Time-Step Calculation

The physical time-step Δt (seconds) for unsteady simulations is applied uniformly to all cells in the domain.

The time-step Δt, appearing in Eqn. (187) for explicit, unsteady calculations, is also computed by consideration of the Courant number and Von Neumann stability conditions:

$$\Delta t = \underbrace{\min}_{\forall x}\left(\frac{\mathrm{CFL}\ V(\mathbf{x})}{\lambda_{\max}(\mathbf{x})}, \frac{\sigma \Delta x^2(\mathbf{x})}{\nu(\mathbf{x})}\right) \tag{199}$$

The chosen time-step is the minimum local time-step over all the fluid cells in the model, and is equal to the minimum value of the pseudo-time step $\Delta\tau$ (Eqn. (197)) in the domain.

For implicit unsteady computations, Eqn. (192), you specify the time-step.

Non-Reflecting Boundary Condition

The non-reflecting boundary condition approach assumes that the solution at the boundary is decomposed, circumferentially in three-dimensional space, into Fourier modes. The zeroth mode corresponds to the average solution and receives the standard boundary condition treatment. The values you specify at the inlet and outlet correspond to this zeroth mode. At the inlet, you supply flow angle and average values of total pressure and temperature. At the outlet, supply static pressure.

In order to prevent reflections at the boundaries, the sum of the harmonics forms the remaining part of the solution and is treated according to exact two-dimensional theory. This is accomplished by making sure that changes in dependent variables at the boundaries obey particular linear relationships derived by considering characteristic propagation of solutions to the linearized Euler equations [23][24][25][47]. Enforcing these relationships imposes steady-state non-reflecting boundary conditions such that no incoming modes are generated at inflow or outflow boundaries.

Non-reflecting boundary conditions allow you to model a smaller domain. Consider the following overlay of two pressure solutions for a stator-rotor interaction: one on an elongated domain and the other on a shortened domain. The vertical line that is upstream of the stator is the inlet boundary of the shortened domain. Similarly, the vertical line that is downstream of the rotor is the outlet boundary of the shortened domain. The shortened domain uses non-reflecting boundary conditions at both inlet and outlet.

The predicted pressure contours of both domains match, indicating that non-reflecting boundaries can be placed close to airfoils without distortion of the pressure solution.



Coupled Flow and Energy Nomenclature

The following is a guide to the symbols used in this section.

Scalars

c	Speed of sound
C_p	Specific heat
E	Total energy per unit mass
H	Total enthalpy
U_r	Reference velocity
V	Cell volume
h	Enthalpy
p	Pressure
t	Time
α_i	Runge-Kutta multistage coefficient
ε	Residual smoothing under-relaxation factor
τ_w	Wall shear stress magnitude
u^*	Reference velocity that is used in wall functions
u^+	Wall-parallel velocity non-dimensionalized with u^*
γ	Coefficient that is used in turbulent stress calculation
ρ	Density
χ	Porosity

μ	Dynamic viscosity
ν	Kinematic viscosity
Δt	Time-step size
Δx	Cell length scale

Tensors or Column Vectors

$\mathbf{A}$	Inviscid flux Jacobian
$\mathbf{f}_f$	Inviscid face flux
$\mathbf{r}$	Residual
$\mathbf{g}_f$	Viscous face flux
$\mathbf{W}$	Conserved quantities in Navier-Stokes equations
$\mathbf{F}$	Inviscid terms in Euler/Navier-Stokes equations
$\mathbf{G}$	Viscous terms in Navier-Stokes equations
$\mathbf{H}$	Body force
$\mathbf{I}$	Identity matrix
$\mathbf{M}$	Modal matrix
$\mathbf{T}$	Viscous stress tensor
$\mathbf{Q}$	Primitive variables in Euler/Navier-Stokes equations
Λ	Matrix of eigenvalues
Γ	Preconditioning matrix

Cartesian Vectors

$\mathbf{a}$	Face area vector
$\mathbf{ds}$	Vector between cell centroids
$\mathbf{f}_r$	Body force due to rotation
$\mathbf{f}_g$	Body force due to gravity
$\mathbf{f}_p$	Porous media body force
$\mathbf{f}_u$	User-defined body force
$\mathbf{f}_\omega$	Vorticity confinement specific force
$\mathbf{f}_L$	Lorentz force or Laplace force
$\mathbf{v}$	Velocity
$\mathbf{v}_g$	Grid velocity
$\dot{\mathbf{q}}''$	Heat flux vector
$\mathbf{s}$	Vector between cell face and cell centroids
$\mathbf{x}$	Position vector
$\vec{\alpha}$	Face metric quantity

Subscripts

0	cell-0 quantity
1	cell-1 quantity
f	Face quantity
l	Laminar
t	Turbulent

$,p$ Partial derivative with respect to pressure
$,T$ Partial derivative with respect to temperature

Superscripts

r Reconstructed value

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

- Tensor (dyadic) product of two vectors, whose results is a tensor:

$$\mathbf{b}\otimes\mathbf{c} = \begin{pmatrix} b_1c_1 & b_1c_2 & b_1c_3 \\ b_2c_1 & b_2c_2 & b_2c_3 \\ b_3c_1 & b_3c_2 & b_3c_3 \end{pmatrix}$$

Segregated Flow Formulation

This section presents the basic flow equations, together with the extensions to the basic finite volume discretization methods that allow the application of these methods to the equations. These methods include the coupling between the continuity and momentum equations.

The methods that are used are based on a large body of work by Peric, Demirdzic, and colleagues (see, for example, [17], [21] and [44], and the book by Ferziger and Peric [22]) and on two papers by Mathur and Murthy (see [23] and [43]).

First, the governing equations in continuous integral form and viscous fluxes are presented. Then the discretization practices are described, including:

- The discrete momentum equation
- The viscous flux discretization
- The face pressure evaluation
- The continuity equation in discrete form

The SIMPLE algorithm is also summarized, and finally, the nomenclature is presented.

Governing Equations in Continuous Integral Form

The following integral equations represent the Navier-Stokes equations for continuity and momentum:

$$\frac{\partial}{\partial t}\int_V \rho\chi \; dV + \oint_A \rho(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} \; = \; \int_V S_u dV \tag{200}$$

$$\begin{aligned}\frac{\partial}{\partial t}\int_V \rho\chi\mathbf{v}dV + \oint_A \rho\mathbf{v}\otimes(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} \; = \\ -\oint_A p\mathbf{I}\cdot d\mathbf{a} + \oint_A \mathbf{T}\cdot d\mathbf{a} + \int_V (\mathbf{f}_r+\mathbf{f}_g+\mathbf{f}_p+\mathbf{f}_u+\mathbf{f}_\omega+\mathbf{f}_L)dV\end{aligned} \tag{201}$$

The terms on the left-hand side of Eqn. (201) are the transient term and the convective flux. On the right-hand side are the pressure gradient term, the viscous flux and the body force terms. $\mathbf{T}$ is the viscous stress tensor. The body force terms represent the effects of system rotation, buoyancy, porous media, user-defined body forces, vorticity confinement specific force, and (in conducting media) the force due to electromagnetic fields, respectively.

For a full description of these terms, see Segregated Flow Nomenclature.

Viscous Fluxes

In *laminar flow*, the stress tensor $\mathbf{T}$ is:

$$\mathbf{T} \; = \; \mathbf{T}_l \; = \; \mu\left[\; \nabla\mathbf{v} + \nabla\mathbf{v}^T - \frac{2}{3}(\nabla\cdot\mathbf{v})\mathbf{I}\right] \tag{202}$$

In *turbulent flow*, the complete stress tensor is:

$$\mathbf{T} \; = \; \mathbf{T}_l + \mathbf{T}_t \tag{203}$$

where the turbulent stress tensor $\mathbf{T}_t$ is also known as the Reynolds stress tensor. Unless a Reynolds stress transport model is being used, the Boussinesq approximation is invoked such that:

$$\mathbf{T} \; = \; \mu_{eff}\left[\; \nabla\mathbf{v} + \nabla\mathbf{v}^T - \frac{2}{3}(\nabla\cdot\mathbf{v})\mathbf{I}\right] \tag{204}$$

where the effective viscosity is $\mu_{eff} \; = \; \mu + \mu_t$, the sum of the laminar and turbulent viscosities.

For a full description of these terms, see Segregated Flow Nomenclature.

Momentum Equation in Discrete Form

Applying Eqn. (201) to a cell-centered control volume for cell-0, one obtains the following discrete equation for the transport of velocity:

$$\frac{\partial}{\partial t}(\rho\chi\mathbf{v}V)_0 + \sum_f [\mathbf{v}\rho(\mathbf{v}-\mathbf{v}_g)\cdot\mathbf{a}]_f = -\sum_f (p\mathbf{I}\cdot\mathbf{a})_f + \sum_f \mathbf{T}\cdot\mathbf{a} \tag{205}$$

The discrete equation for each velocity component can be expressed implicitly as a linear system as described in Finite Volume Discretization, p. 6341. The transient terms, body forces, and convective flux for each velocity component is discretized in the manner as the scalar quantity described therein. The treatment of the viscous flux is discussed in a later section.

For a full description of these terms, see Segregated Flow Nomenclature.

Viscous Flux Discretization

Interior Faces

To evaluate the stress tensor $\mathbf{T}$, the velocity gradient tensor at the face $\nabla\mathbf{v}_f$ must be written in terms of the cell velocities for purposes of linearization.

Using Eqn. (2555), the velocity gradient tensor at a face is written as:

$$\nabla\mathbf{v}_f = \Delta\mathbf{v}\otimes\vec{\alpha} + \overline{\nabla\mathbf{v}}_f - (\overline{\nabla\mathbf{v}}_f\cdot\mathbf{ds})\otimes\vec{\alpha} \tag{206}$$

where:

$$\Delta\mathbf{v} = \mathbf{v}_1 - \mathbf{v}_0 \qquad \text{and} \qquad \overline{\nabla\mathbf{v}}_f = (\nabla\mathbf{v}_0 + \nabla\mathbf{v}_1)/2 \tag{207}$$

where $\nabla\mathbf{v}_0$ and $\nabla\mathbf{v}_1$ are the explicitly computed velocity gradient tensors in the cells.

For a full description of these terms, see Segregated Flow Nomenclature.

Boundary Faces

For *slip walls*, the viscous shear force at the wall boundary face, $(\mathbf{T}\cdot\mathbf{a})_f$, is set to zero.

For *no-slip walls* in *laminar* flow, Eqn. (2557) is used to represent the velocity gradient tensor at the face:

$$\nabla \mathbf{v}_f = \Delta \mathbf{v} \otimes \vec{\alpha} + \nabla \mathbf{v}_0 - (\nabla \mathbf{v}_0 \cdot \mathbf{ds}) \otimes \vec{\alpha} \tag{208}$$

where:

$$\Delta \mathbf{v} = \mathbf{v}_f - \mathbf{v}_0 \tag{209}$$

and $\mathbf{v}_f$ is the specified wall tangential motion.

The stress tensor at the face is therefore:

$$\mathbf{T}_f = \mu \left[\nabla \mathbf{v} + \nabla \mathbf{v}^T - \frac{2}{3}(\nabla \cdot \mathbf{v})\mathbf{I} \right]_f \tag{210}$$

For *no-slip walls* in *turbulent flow*, it is assumed that only the component of velocity parallel to the wall is of interest. A linear relationship between the wall shear force and the wall-parallel component of the difference in velocity between the wall and the cell is hypothesized:

$$(\mathbf{T} \cdot \mathbf{a})_f = -\gamma(\mathbf{v}^t - \mathbf{v}_f^t) \tag{211}$$

Using the definitions of wall shear stress magnitude ($\tau_w = \rho u^{*2}$) and reference velocity ($u^+ = |\mathbf{v}_p| / u^*$), the coefficient of proportionality is:

$$\gamma = |\mathbf{a}| \frac{\rho u^*}{u^+} \tag{212}$$

The reference velocity u^* is computed according to the specific turbulence model. The value of u^+ is obtained (as a function of y^+) from the appropriate wall law.

For a full description of these terms, see Segregated Flow Nomenclature.

Face Pressure Evaluation

In order to compute the pressure gradient term in Eqn. (205), the pressure must be evaluated at each face:

$$p_f = \frac{\bar{a}_0 p_{f0} + \bar{a}_1 p_{f1}}{\bar{a}_0 + \bar{a}_1} \tag{213}$$

where $\bar{a}_0$ and $\bar{a}_1$ are the average of the momentum coefficients for all components of momentum for cells 0 and 1, respectively. p_{f0} and p_{f1} are interpolated from cell values and reconstruction gradients according to Eqn. (2575).

When the delta-V dissipation expert property is turned on, a dissipation term is incorporated into the face pressure evaluation, as follows:

$$p_f = \frac{\bar{a}_0 p_{f0} + \bar{a}_1 p_{f1}}{\bar{a}_0 + \bar{a}_1} + \frac{2\bar{a}_0\bar{a}_1}{\bar{a}_0 + \bar{a}_1} \cdot \frac{(\mathbf{v}_{f0} - \mathbf{v}_{f1}) \cdot \mathbf{a}}{4(\mathbf{a} \cdot \mathbf{a})} \tag{214}$$

where $\mathbf{v}_{f0}$ and $\mathbf{v}_{f1}$ are also evaluated according to Eqn. (2575).

At boundaries, the cell 1 contribution comes from a fictitious "ghost cell" in which the velocity has been reflected about the boundary face. Therefore Eqn. (213) becomes:

$$p_f = p_{f0} \tag{215}$$

With delta-V dissipation on, Eqn. (214) becomes:

$$p_f = p_{f0} + \bar{a}_0 \cdot \frac{\mathbf{v}_{f0} \cdot \mathbf{a}}{2(\mathbf{a} \cdot \mathbf{a})} \tag{216}$$

at a boundary.

For a full description of these terms, see Segregated Flow Nomenclature.

Continuity Equation in Discrete Form

The discrete continuity equation is written:

$$\sum_f \dot{m}_f = \sum_f (\dot{m}_f^* + \dot{m}_f') = 0 \tag{217}$$

The uncorrected face mass flow rate $\dot{m}_f^*$ is computed after the discrete momentum equations have been solved. The mass flow correction $\dot{m}_f'$ is required to satisfy continuity.

In case of unsteady flow, face fluxes are now completely recomputed at the first iteration of the time-step. This can sometimes result in slightly different residual trends when compared with previous versions of STAR-CCM+.

The following sections describe how the discretized pressure correction equation is obtained, both at interior faces and at boundaries.

Interior Faces

The uncorrected mass flow rate at an interior face can be written in terms of the cell variables as follows:

$$\dot{m}_f^* = \rho_f \left[\mathbf{a} \cdot \left(\frac{\mathbf{v}_0^* + \mathbf{v}_1^*}{2} \right) - G_f \right] - \Upsilon_f \tag{218}$$

where $\mathbf{v}_0^*$ and $\mathbf{v}_1^*$ are the cell velocities after the discrete momentum equations have been solved. $G_f = (\mathbf{a} \cdot \mathbf{v}_g)_f$ is the grid flux. Υ_f is the Rhie-and-Chow dissipation at the face:

$$\Upsilon_f = Q_f(p_R - p_L - \overline{\nabla p_f^*} \cdot d\mathbf{s}) \tag{219}$$

where:

$$Q_f = \rho_f \left(\frac{V_0 + V_1}{\bar{a}_0 + \bar{a}_1} \right) \boldsymbol{\alpha} \cdot \mathbf{a} \tag{220}$$

V_0 and V_1 are the volumes for cell 0 and cell 1, respectively.

$\bar{a}_0$ and $\bar{a}_1$ are the average of the momentum coefficients for all components of momentum for cells 0 and 1, respectively.

p_L is the pressure for cell 0 and p_R is the pressure for cell 1.

The vector $\boldsymbol{\alpha}$ is given by Eqn. (2561).

$\overline{\nabla p^*_f}$ is the volume-weighted average of the cell gradients of pressure, computed using a volume-based interpolation between the gradient values of the two cells

Unsteady Dissipation Corrections

For either coefficient $\bar{a} = \bar{a}_0$ or $\bar{a}_1$, when unsteady dissipation corrections are enabled, the coefficients $\bar{a}$ are modified based on Acoustic CFL.

- Let Acoustic CFL be $CFL_a = c\frac{\Delta t}{\Delta x}$ where:
  - c is the speed of sound.
  - Δt is the simulation time-step.
  - Δx is the local grid size.
- Let limCFL be the Limit Acoustic CFL value specified by the user.

Then:

- If $\mathrm{CFL} > \mathrm{limCFL}$, then $\bar{a}$ is used without modification.
- Otherwise, $\bar{a}$ is replaced with $\bar{a}'$, as follows:

$$\bar{a}' = \bar{a}_u \left(1 + \frac{\mathrm{CFL}_r \, \bar{a}_{uns}}{\max(\bar{a}_u, 1\times10^{-20})} \right)$$

where:

- $\bar{a}_u = \bar{a} - \bar{a}_{uns}$ and $\bar{a}_{uns} = \frac{3}{2}\rho\frac{\mathrm{Vol}}{\Delta t}$
- $\mathrm{CFL}_r = \min\left(1.0, \frac{\mathrm{CFL}}{\mathrm{limCFL}}\right)$

p_L and p_R are the pressures values computed by linear reconstruction at two symmetric points $(\mathbf{x}_L, \mathbf{x}_R)$ with regard to the face-centroid, in the direction of face normal.

$\mathbf{x}_L = \mathbf{x}_f - d_s\mathbf{n}$ and $\mathbf{x}_R = \mathbf{x}_f + d_s\mathbf{n}$

where:

- $\mathbf{n}$ is the unit vector associated with area vector $\mathbf{a}_f$.
- $\mathbf{x}_f$ is the position vector of the face centroid.
- d_s is a scalar distance, where:
  - If $\max\left(\frac{V_1}{V_0}, \frac{V_0}{V_1}\right) > 10$, then $d_s = \frac{0.5}{|\alpha|}$.
  - Otherwise, $d_s = \min(d_{s0}, d_{s1})$, where:
    - $d_{s0} = d\mathbf{x}_0 \cdot \mathbf{n}$ (the projection of $d\mathbf{x}_0 = \mathbf{x}_f - \mathbf{x}_0$ onto the face area normal)
    - $d_{s1} = d\mathbf{x}_1 \cdot \mathbf{n}$ (the projection of $d\mathbf{x}_1 = \mathbf{x}_f - \mathbf{x}_1$ onto the face area normal).

Density Correction

For compressible flow, the density correction ρ' is introduced. Noting that:

$$\begin{aligned} \dot{m}_f &= (\rho + \rho')_f(v_{fn}^* + v'_{fn})|\mathbf{a}| \\ &= (\rho_f v_{fn}^* + \rho'_f v_{fn}^* + \rho_f v'_{fn} + \rho'_f v'_{fn})|\mathbf{a}| \end{aligned} \tag{221}$$

where the subscript "*fn*" denotes the face-normal component, define:

$$\rho_f v'_{fn}|\mathbf{a}| \equiv -Q_f(p'_1 - p'_0) \tag{222}$$

where p'_1 and p'_0 are the cell pressure corrections, and:

$$\rho'_f v_{fn}^*|\mathbf{a}| = \frac{\dot{m}_f^*}{\rho_f}\left(\frac{\partial\rho}{\partial p}\right)_T p'_{\mathrm{upwind}} \tag{223}$$

where:

$$p'_{\mathrm{upwind}} = \begin{cases} p'_0 & \text{for } \dot{m}_f^* > 0 \\ p'_1 & \text{for } \dot{m}_f^* < 0 \end{cases} \tag{224}$$

Since you can choose to neglect the second-order term $\rho'_f v_{fn}'$, these equations can be combined to define the mass flow correction as:

$$\dot{m}'_f = Q_f(p'_0 - p'_1) + \frac{\dot{m}^*_f}{\rho_f}\left(\frac{\partial \rho}{\partial p}\right)_T p'_{\text{upwind}} \tag{225}$$

Discrete Pressure Correction Equation

The discrete pressure correction equation is obtained from Eqns. (217) and (225) and is written in coefficient form as:

$$\bar{a}\, p'_p + \sum_n a_n p'_n = r \tag{226}$$

The residual r is simply the net mass flow into the cell:

$$r = -\sum_f \dot{m}^*_f \tag{227}$$

Boundary Faces with Specified Velocity

In the cases where the boundary velocity is specified, such as wall, symmetry, and inlet boundaries, the value of $\dot{m}^*_f$ is calculated directly from the known velocity $\mathbf{v}^*_f$ on boundaries as:

$$\dot{m}^*_f = \rho_f(\mathbf{a} \cdot \mathbf{v}^*_f - G_f) \tag{228}$$

For these boundaries, a Neumann condition is used for the pressure correction:

$$p'_f = p'_0 \tag{229}$$

and the mass flux corrections are zero.

For a full description of these terms, see Segregated Flow Nomenclature.

Pressure Boundaries

On a specified-pressure boundary (stagnation inlet and pressure outlet), the pressure corrections are not zero. The uncorrected boundary mass flux is:

$$\dot{m}^*_f = \rho_f(\mathbf{v}_f \cdot \mathbf{a} - G_f) - \Upsilon_f \tag{230}$$

where $\mathbf{v}_f$ is the boundary velocity. The dissipation Υ_f is defined as:

$$\Upsilon_f = Q_f(p^*_f - p^*_0 - \overline{\nabla p^*_0} \cdot \vec{ds}) \tag{231}$$

where:

$$Q_f = \frac{\rho_f V_0}{\bar{a}_0}(\vec{\alpha} \cdot \mathbf{a}) \tag{232}$$

Similar to Eqn. (225), you can postulate a mass flux correction of the form:

$$\dot{m}'_f = -Q_f(p'_0 - p'_f) + \mathbf{v}_f \cdot \mathbf{a}\left(\frac{\partial \rho}{\partial p}\right)_T p'_{\text{upwind}} \tag{233}$$

For *supersonic inflow*, $p'_f = 0$ and $p'_{\text{upwind}} = 0$, and $\dot{m}'_f = 0$.

For *supersonic outflow*, $p'_0 = p'_f$ and $p'_{\text{upwind}} = p'_0$, and Eqn. (233) becomes:

$$\dot{m}'_f = \mathbf{v}_f \cdot \mathbf{a}\left(\frac{\partial \rho}{\partial p}\right)_T p'_0 \tag{234}$$

For *subsonic outflow*, $p'_f = 0$ and $p'_{\text{upwind}} = p'_0$, and Eqn. (233) becomes:

$$\dot{m}'_f = \left[Q_f + \mathbf{v}_f \cdot \mathbf{a}\left(\frac{\partial \rho}{\partial p}\right)_T\right] p'_0 \tag{235}$$

For *subsonic inflow*, $p'_{\text{upwind}} = 0$. However, even though this value simplifies Eqn. (233) somewhat, it is still necessary to eliminate p'_f from the equation.

Noting that:

$$\dot{m}'_f = \frac{\partial \dot{m}^*_f}{\partial p_f} p'_f = -\frac{\mathbf{v}_f \cdot \mathbf{a}}{|\mathbf{v}_f|^2} p'_f \tag{236}$$

this equation is combined with Eqn. (233) to obtain:

$$\dot{m}'_f = \frac{\mathbf{v}_f \cdot \mathbf{a}\, Q_f}{\mathbf{v}_f \cdot \mathbf{a} - Q_f |\mathbf{v}_f|^2} \quad \text{and} \quad p'_f = \frac{|\mathbf{v}_f|^2 Q_f}{Q_f |\mathbf{v}_f|^2 - \mathbf{v}_f \cdot \mathbf{a}} p'_0 \tag{237}$$

A similar approach is followed for the free-stream boundary type, with slightly different results.

For a full description of these terms, see Segregated Flow Nomenclature.

SIMPLE Solver Algorithm

The SIMPLE algorithm is used to control the overall solution. This algorithm is summarized as follows:

1. Set the boundary conditions.
2. Computed the reconstruction gradients of velocity and pressure.

3. Compute the velocity and pressure gradients.
4. Solve the discretized momentum equation. This creates the intermediate velocity field $\mathbf{v}^*$.
5. Compute the uncorrected mass fluxes at faces $\dot{m}_f^*$.
6. Solve the pressure correction equation. This produces cell values for the pressure correction p'.
7. Update the pressure field:

$$p^{n+1} = p^n + \omega p' \tag{238}$$

   where ω is the under-relaxation factor for pressure.
8. Update the boundary pressure corrections p'_b.
9. Correct the face mass fluxes:

$$\dot{m}_f^{n+1} = \dot{m}_f^* + \dot{m}'_f \tag{239}$$

10. Correct the cell velocities:

$$\mathbf{v}^{n+1} = \mathbf{v}^* - \frac{V\nabla p'}{\mathbf{a}'^{\mathrm{v}}_p} \tag{240}$$

   where $\nabla p'$ is the gradient of the pressure corrections, $\mathbf{a}'^{\mathrm{v}}_p$ is the vector of central coefficients for the discretized linear system representing the velocity equation, and V is the cell volume.
11. Update density due to pressure changes.
12. Free all temporary storage.

For a full description of these terms, see Segregated Flow Nomenclature.

PISO Solver Algorithm

The PISO algorithm is suited for transient cases where the time-step is smaller than the flow requires, that is, the Convective Courant number is less than 1. In such cases, coupling with a third-party code or other external factors such as regulatory requirements typically determine the time-step used. For all other cases, a larger time-step and the SIMPLE algorithm are preferred.

In STAR-CCM+, PISO is currently only compatible with cases using the standard K-Epsilon model, a single phase, and no combustion.

The PISO algorithm shares many features with the SIMPLE algorithm. The algorithm is summarized as follows:

1. Set the boundary conditions.

2. Compute the reconstruction gradients of velocity and pressure.
3. Compute the velocity and pressure gradients.
4. Solve the discretized momentum equation. This creates the intermediate velocity field v^*.
5. Compute the uncorrected mass fluxes at faces m_f^*.
6. Construct the pressure correction equation and test for convergence of the PISO algorithm. If PISO has converged, then set the last iteration flag to true.
7. Solve the pressure correction equation. This produces cell values for the pressure correction p'.
8. Update the pressure field:

$$p^{n+1} = p^n + \omega p' \tag{241}$$

   where ω is the under-relaxation factor for pressure. If this is the last iteration, ω is equal to one.

9. Update the boundary pressure corrections p'_b.
10. Correct the face mass fluxes:

$$m_f^{n+1} = m_f^* + m'_f \tag{242}$$

11. Correct the cell velocities:

$$v^{n+1} = v^* - \frac{V\omega\nabla p'}{a_p^v} \tag{243}$$

where $\nabla p'$ is the gradient of the pressure corrections, a_p^v is the vector of central coefficients for the discretized linear system representing the velocity equation, and V is the cell volume.

12. Update density due to pressure changes.
13. Solve any PISO-coupled equations such as for Species or Energy.
14. Recalculate density from the equation of state (except for the last corrector).
15. If this is the last iteration, then go to step 18.
16. Explicitly update the momentum equation with corrections due to the neighbor cells.
17. Go to step 5 and continue until convergence is reached.
18. Free all temporary storage.

For a full description of these terms, see Segregated Flow Nomenclature.

Segregated Flow Nomenclature

The following is a guide to the symbols used in this section.

Scalars

a	linear coefficient
$\dot{m}_f$	face mass flow
G_f	grid flux
Q_f	coefficient in dissipation
V	cell volume
p	pressure
r	residual
T	temperature
t	time
τ_w	wall shear stress magnitude
u^*	reference velocity that is used in wall functions
u^+	wall-parallel velocity non-dimensionalized with u^*
Υ_f	Rhie-Chow-type dissipation
γ	coefficient that is used in turbulent stress calculation
ρ	density
χ	porosity, the fraction of volume not occupied by porous media
μ	dynamic viscosityporosity is
ν	kinematic viscosity

Tensors or Column Vectors

$\mathbf{I}$	identity matrix
$\mathbf{T}$	viscous stress tensor

Cartesian Vectors

$\mathbf{a}$	face area vector
$\mathbf{ds}$	vector between cell centroids
$\mathbf{f}_r$	body force due to rotation
$\mathbf{f}_g$	body force due to gravity
$\mathbf{f}_p$	porous media body force
$\mathbf{f}_u$	user-defined body force
$\mathbf{f}_\omega$	vorticity confinement specific force
$\mathbf{f}_L$	Lorentz force or Laplace force
$\mathbf{v}$	velocity
$\mathbf{v}_g$	grid velocity
$\mathbf{x}$	position vector
$\vec{\alpha}$	face metric quantity

Subscripts

0	cell-0 quantity
1	cell-1 quantity
f	face value
fn	face normal component
eff	effective (as in effective viscosity)
p	center coefficient
n	neighbor coefficient
f	face quantity
l	laminar
t	turbulent

Superscripts

r	reconstructed value
*	predictor value
′	correction

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

- Tensor (dyadic) product of two vectors, whose results is a tensor:

$$\mathbf{b} \otimes \mathbf{c} = \begin{pmatrix} b_1c_1 & b_1c_2 & b_1c_3 \\ b_2c_1 & b_2c_2 & b_2c_3 \\ b_3c_1 & b_3c_2 & b_3c_3 \end{pmatrix}$$

Segregated Fluid Energy Formulation

In this section, the total energy equation in integral and discrete form is presented together with the individual terms in the equations:

- The convection term
- The diffusion term
- The unsteady term
- The viscous work term

Finally, the nomenclature is presented.

Energy Equation

The energy equation written in integral form is:

$$\frac{d}{dt}\int_V \rho E dV + \oint_A [\rho H(\mathbf{v} - \mathbf{v}_g) + \mathbf{v}_g p] \cdot d\mathbf{a} = \tag{244}$$
$$-\oint_A \dot{\mathbf{q}}'' \cdot d\mathbf{a} + \oint_A \mathbf{T} \cdot \mathbf{v} d\mathbf{a} + \int_V \mathbf{f} \cdot \mathbf{v} dV + \int_V s_u dV$$

where E is the total energy, H is the total enthalpy, $\dot{\mathbf{q}}''$ is the heat flux vector, $\mathbf{T}$ is the viscous stress tensor, $\mathbf{v}$ is the velocity vector, and $\mathbf{v}_g$ is the grid velocity vector. $\mathbf{f}$ is the body force vector representing the combined body forces of rotation, gravity, and the fan model in Eqn. (201) (Segregated) and Eqn. (166) (Coupled); body forces due to porous media are not included. s contributes more energy source terms, such as those that you specify.

Total energy is related to the total enthalpy H by:

$$E = H - p/\rho \tag{245}$$

where:

$$H = h + |\mathbf{v}|^2/2 \tag{246}$$

and $h = C_p T$.

For a full description of these terms, see Segregated Fluid Energy Nomenclature.

Energy Equation in Discrete Form

Applying Eqn. (244) to a cell-centered control volume (V_0), one obtains:

$$\frac{d}{dt}(\rho E V_0) + \sum_f \{[\rho H(\mathbf{v} - \mathbf{v}_g) + \dot{\mathbf{q}}'' \cdot \mathbf{a} - (\mathbf{T} \cdot \mathbf{v})] \cdot \mathbf{a}\}_f = (\mathbf{f} \cdot \mathbf{v} + s)V_0 \tag{247}$$

For a full description of these terms, see Segregated Fluid Energy Nomenclature.

Convection Term

The convection term is evaluated as described in Finite Volume Discretization, p. 6341. In this case, the convected quantity is the total enthalpy, H.

Diffusion Term

Interior Faces

The heat flux vector is:

$$\dot{\mathbf{q}}'' = -k_{\text{eff}} \nabla T \tag{248}$$

where k_{eff} is the effective thermal conductivity:

$$k_{\text{eff}} = k + \frac{\mu_t C_p}{\sigma_t} \tag{249}$$

and for porous regions:

$$k_{\text{eff}} = \chi k_{fluid} + (1 - \chi) k_{solid} \tag{250}$$

with σ_t being the turbulent Prandtl number. The diffusive flux at an interior face is then discretized as described in Finite Volume Discretization, p. 6341 as follows:

$$\dot{\mathbf{q}}'' \cdot \mathbf{a} = -k_{\text{eff}} \nabla T \cdot \mathbf{a} = -k_{\text{eff}} [(T_1 - T_0)\vec{\alpha} \cdot \mathbf{a} + \overline{\nabla T} \cdot \mathbf{a} - (\overline{\nabla T} \cdot \mathbf{ds})\vec{\alpha} \cdot \mathbf{a}] \tag{251}$$

Boundary Faces

For inviscid flows, the heat flux is zero.

For viscous flows, the heat flux is required. This value is specified directly in the case of adiabatic or specified heat flux walls. It can also be inferred as a convective heat transfer condition, such that:

$$\dot{\mathbf{q}}'' \cdot \mathbf{a} = h(T_\infty - T_f) \tag{252}$$

where h is a convective heat transfer coefficient and T_∞ is a reference temperature.

For *laminar flow* in which the wall temperature is specified, the heat flux is obtained as follows:

$$\dot{\mathbf{q}}'' \cdot \mathbf{a} = -k_{\text{eff}} \nabla T \cdot \mathbf{a} = -k_{\text{eff}} [(T_f - T_0)\vec{\alpha} \cdot \mathbf{a} + \nabla T \cdot \mathbf{a} - (\nabla T \cdot \mathbf{ds})\vec{\alpha} \cdot \mathbf{a}] \tag{253}$$

For *turbulent flow* in which the wall temperature is specified, thermal wall laws are employed as follows:

$$\dot{\mathbf{q}}'' \cdot \mathbf{a} = -\frac{\mathbf{a}\rho C_p u^*}{t^+}(T_b - T_0) \tag{254}$$

where t^+ is defined in terms of the appropriate thermal wall law.

For a full description of these terms, see Segregated Fluid Energy Nomenclature.

Unsteady Term

For all non-porous regions, the unsteady term is computed as $\frac{d}{dt}\int_V \rho E dV$.

For porous regions, it is computed as:

$$\frac{d}{dt}\int_V \rho E dV = \frac{d}{dt}\int_V ((\chi(\rho E)_{fluid} + (1-\chi)(\rho C_p)_{solid} T) dV) \tag{255}$$

For a full description of these terms, see Segregated Fluid Energy Nomenclature.

Viscous Work Term

The viscous work term at a face, $[(\mathbf{T} \cdot \mathbf{v}) \cdot \mathbf{a}]$, is evaluated in a straightforward fashion using the stress tensor that is calculated in the Segregated Flow model formulation.

For a full description of these terms, see Segregated Fluid Energy Nomenclature.

Segregated Fluid Energy Nomenclature

The following is a guide to the symbols used in this section.

Scalars

E	total energy
C_p	specific heat
H	total enthalpy
V	cell volume
T	temperature
t^+	non-dimensional temperature in wall units
u^*	reference velocity that is used in wall functions
k	thermal conductivity
t	time
γ	porous media volume porosity
ρ	density
χ	porosity,
μ	dynamic viscosity
ν	kinematic viscosity

Tensors or Column Vectors

$\mathbf{T}$	viscous stress tensor
$\mathbf{R}$	porous resistance tensor

Cartesian Vectors

$\mathbf{a}$	face area vector
$\mathbf{f}$	body force vector
$\mathbf{v}$	velocity
$\mathbf{v}_g$	grid velocity
$\dot{\mathbf{q}}''$	heat flux vector
$\vec{\alpha}$	face metric quantity

Subscripts

0	cell-0 quantity
t	turbulent
eff	effective (as in effective conductivity)
∞	reference value

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Scalar (inner, double-dot) product of two vectors, whose result is a scalar:

$$\mathbf{a} \cdot \mathbf{b} = a_i b_i$$

Gravity Formulation

This section describes the source term formulation when gravity is included in the simulation. These source terms would appear in Eqn. (166) for the Coupled Flow model or in Eqn. (201) for the Segregated Flow model. The two alternative formulations are for the following:

- Variable density flows
- Constant density flows

Variable Density Flows

For variable density flows, the buoyancy source term is:

$$\mathbf{f}_g = (\rho - \rho_{\mathrm{ref}})\mathbf{g} \tag{256}$$

where:

- $\mathbf{g}$ is the gravitational vector.
- ρ_{ref} is the reference density.

The working pressure becomes the piezometric pressure, and the static pressure is given by:

$$p_{\mathrm{static}} = p + \rho_{\mathrm{ref}}\,\mathbf{g}(\mathbf{x} - \mathbf{x}_0) \tag{257}$$

where $\mathbf{x}_0$ is a position vector, termed the operating altitude, at which the piezometric and static pressures are equal.

For VOF multiphase flows the density variation could be very significant (several orders of magnitude), and result in numerical instabilities. To obtain a robust and stable numerical solution in such cases, set the reference density to zero or to the density of the lightest phase.

Constant Density Flows

The working pressure becomes the piezometric pressure, and the static pressure is given by:

$$p_{\mathrm{static}} = p + \rho\mathbf{g}(\mathbf{x} - \mathbf{x}_0) \tag{258}$$

where $\mathbf{x}_0$ is a position vector, termed the operating altitude, at which the piezometric and static pressures are equal.

If the Boussinesq model is enabled, the buoyancy source term is approximated as:

$$\mathbf{f}_g = \rho \mathbf{g} \beta (T_{\text{ref}} - T) \tag{259}$$

where:

- $\mathbf{g}$ is the gravitational vector.
- β is the coefficient of bulk expansion.
- T_{ref} is the operating temperature.

Axisymmetric Swirl Model Formulation

The equations being modeled by the Axisymmetric Swirl model are the 2D conservation equations for mass and momentum. In a cylindrical coordinate system, the 2D conservation equation is:

$$\frac{\partial}{\partial t}\int_V \rho \boldsymbol{u} r dV + \oint (\rho \boldsymbol{u} \otimes \boldsymbol{u} + p\boldsymbol{I}) \cdot r d\mathbf{a} \tag{260}$$

$$= \oint \boldsymbol{T} \cdot r d\mathbf{a} + \int \frac{1}{r} \begin{bmatrix} 0 \\ p + \rho u_\theta{}^2 - \tau_{\theta\theta} \\ -\rho u_r u_\theta + \tau_{\theta r} \end{bmatrix} r dV$$

where

$$\boldsymbol{u} = \begin{bmatrix} u_z & u_r & u_\theta \end{bmatrix}^T$$

and

$$\frac{\partial}{\partial \theta}(\ ) = 0$$

and

$$A_\theta = 0$$

and $\boldsymbol{T}$ is the stress tensor

$$\boldsymbol{T} = \mu \begin{bmatrix} 2\frac{\partial u_z}{\partial z} - \frac{2}{3}\nabla \bullet \boldsymbol{u} & \frac{\partial u_z}{\partial r} + \frac{\partial u_r}{\partial z} & \frac{\partial u_\theta}{\partial z} \\ \frac{\partial u_z}{\partial r} + \frac{\partial u_r}{\partial z} & 2\frac{\partial u_r}{\partial r} - \frac{2}{3}\nabla \bullet \boldsymbol{u} & \frac{\partial u_\theta}{\partial r} - \frac{u_\theta}{r} \\ \frac{\partial u_\theta}{\partial z} & \frac{\partial u_\theta}{\partial r} - \frac{u_\theta}{r} & 2\frac{u_r}{r} - \frac{2}{3}\nabla \bullet \boldsymbol{u} \end{bmatrix}$$

and $\boldsymbol{I}$ is the identity matrix.

Porous Media Formulation

The porous source term appears in the momentum equations of the coupled and segregated flow solvers (see equations (166) and (201) respectively), and is:

$$\mathbf{f}_p = -\mathbf{P} \cdot \mathbf{v} \tag{261}$$

and

$$\mathbf{f}_p = -\alpha \mathbf{P} \cdot \mathbf{v} \tag{262}$$

for single phase and multiphase flows respectively, where $\mathbf{P}$ is the porous resistance tensor and α is the volume fraction of a particular phase.

The porous resistance tensor is given by:

$$\mathbf{P} = \mathbf{P}_\mathrm{v} + \mathbf{P}_\mathrm{i}|\mathbf{v}| \tag{263}$$

where $\mathbf{P}_\mathrm{v}$ and $\mathbf{P}_\mathrm{i}$ are the viscous (linear) and inertial (quadratic) resistance tensors, respectively.

Vorticity Confinement Formulation

The basic approach of the Vorticity Confinement model consists of adding a force term $\mathbf{f}_\omega$ to momentum equations (Eqn. (166) for coupled flow and Eqn. (201) for segregated) to help "confine" the vorticity.

Vorticity confinement is a force per unit volume $[N/m^3]$. Its purpose is to keep the vorticity from spreading. The source and its derivation can be found in the references: [38], [39], [49].

It is computed as:

$$\mathbf{f}_{\omega} = -\varepsilon\rho(\mathbf{n} \times \omega) \tag{264}$$

where $\omega = \mathrm{curl}(\mathbf{v})$, ρ is fluid density and $\mathbf{n}$ is the unit vector that is aligned with $(\nabla|\omega|)/(|\nabla|\omega||)$. The following diagram depicts the basic spatial relationship among ω and the vectors $\mathbf{n}$ and $\mathbf{v}$.



$$\varepsilon = \varepsilon_0 V_{\omega_1} \mathbf{f}(d^*) \tag{265}$$

where the model coefficient ε_0 (non-dimensional) is set to 0.04 in three-dimensional cases and 0.02 in two-dimensional cases by default, and:

$$V_{\omega_1} = (|\nabla|\omega||)h^2 \tag{266}$$

where h is the local grid size, computed as:

$$h = \sqrt[ND]{\mathrm{Vol}_{\mathrm{cell}}} \tag{267}$$

d^* is the non-dimensional wall distance:

$$d^* = W_d\left(\frac{V_{\mathrm{scale}}}{\nu}\right) \tag{268}$$

where W_d is wall distance, ν is kinematic viscosity of fluid, and V_{scale} is the mass average of velocity magnitude.

$$\mathbf{f}(d^*) \triangleq \min\left(\max\left(\left(\frac{d^* - d^*_{\mathrm{cutoff}}}{0.3d^*_{\mathrm{cutoff}}}\right), 0\right), 1\right) \tag{269}$$

where d^*_{cutoff} is user-specified non-dimensional wall distance.

Both the model coefficient ε_0 and non-dimensional cutoff wall distance can be retuned for different simulation. Default values have been found to be appropriate for external aerodynamics.

Fan Interface Formulation

This section describes how STAR-CCM+ computes the fan interface swirl velocity, based on fan efficiency, angular velocity, and pressure rise across the fan.

Note: This formulation derives the swirl from the Euler equations, and assumes that the fan consists of a single rotor. This approach is not valid for a rotor-stator pair, or for more blade rows.

Assuming that the axial velocity, radial velocity, and density are unchanged across the fan, the total pressure jump ΔP_{total} across the fan is related to the static pressure jump ΔP_{static}, derived from the performance curve, as:

$$\Delta P_{total} = \Delta P_{static} + \frac{1}{2}\rho[V_{\theta\,dn}^2 - V_{\theta\,up}^2] \tag{270}$$

The subscripts dn and up represent downstream and upstream of the fan respectively. V_θ represents the swirl velocity. The upstream swirl velocity $V_{\theta\,up}$ is the tangential component of the velocity at the upstream boundary of the fan interface.

The efficiency is given as:

$$\eta = \frac{(\text{Power Imparted to Air})}{(\text{Power Consumed})} \tag{271}$$

$$\eta = \frac{\dot{Q}\Delta P_{total}}{\dot{m}\Delta(\omega r V_\theta)} \tag{272}$$

Where $\dot{Q}$ and $\dot{m}$ are volume flow rate and mass flow rate respectively, ω is the fan angular velocity, and r is the distance from the center of the fan.

Rearranging these equations gives:

$$V_{\theta\,dn} - V_{\theta\,up} = \frac{\Delta P_{total}}{\eta\omega r\rho} \tag{273}$$

Eqn. (270) and Eqn. (273) combine to give the following quadratic equation:

$$\frac{1}{2}V_{\theta\,dn}^2 + \eta\omega r V_{\theta\,dn} + \left[\frac{\Delta P_{static}}{\rho} + \eta\omega r V_{\theta\,up} - \frac{1}{2}V_{\theta\,up}^2\right] \tag{274}$$

The solution of this equation is:

$$V_{\theta\,dn} = \eta\omega r \pm \sqrt{(\eta\omega r)^2 - 2\left[\frac{\Delta P_{static}}{\rho} + \eta\omega r V_{\theta\,up} - \frac{1}{2}V_{\theta\,up}^2\right]} \tag{275}$$

STAR-CCM+ uses only the negative square root. In the case of a complex root, STAR-CCM+ uses only the real part. This equation is used to calculate the swirl velocity which is added due to the fan.

Fan Momentum Source Formulation

This section describes the theory that STAR-CCM+ uses in computing the fan momentum sources.

To understand the forces that the fan blades impart to the fluid, consider the velocity vector diagram below. The fan axis is aligned in the z-direction. A single fan blade is shown in the fan volume.



Fluid approaches the fan with velocity V_{in}, which is assumed to be uniform and to have no tangential component $(V_{in} = V_{axial})$. The fan rotates in the θ-direction with the angular velocity ω. The fan blades are at an angle β to the fluid flow and are assumed to have no curvature.

The fluid velocity leaving the fan, $V_{out,\,absolute}$, is the vector sum of the relative velocity $V_{out,\,relative}$ (assumed to be parallel to the blade angle) and the rotational velocity ωr, imparted to the fluid by the fan blades.

Continuity implies that $V_{axial,\,in} = V_{axial,\,out}$. Conservation of momentum states that the sum of the forces acting on a volume is equal to the net flow of momentum across the surfaces of that volume. Applying this principle and assuming the momentum flux contributions apply only to the volume force (surface forces are not affected), the equation for the θ-direction becomes:

$$F_\theta = \rho V_{axial} A(V_{\theta,out} - V_{\theta,in}) \tag{276}$$

where F_θ is the tangential force, ρ is the fluid density, and A is the surface area of the fan blade in the θ-direction. From the previous vector diagram, it can be shown that:

$$V_{\theta,out} = \omega r - V_{axial} \tan\beta \tag{277}$$

Recall that $V_{\theta,in}$ is assumed to be zero so that Eqn. (276) can be written as:

$$F_\theta = \rho V_{axial} \cdot A \cdot (\omega r - V_{axial} \tan\beta) \tag{278}$$

Eqn. (278) is an expression for the tangential force of the fan blades.

Tangential Component

The target tangential velocity (see Eqn. (277)) is obtained as:

$$V_{\theta_TARGET} = (\omega \times \text{radius_from_fan_center} - V_{AXIAL} \times \tan\beta) \tag{279}$$

Therefore $V_{\theta_TARGET_AVG}$ is calculated as the average of the V_{θ_TARGET} at the downstream fan boundary. The average of the actual tangential velocity is calculated as:

$$V_{\theta_ACTUAL_AVG} = (V \bullet \theta)_{AVERAGE_DOWNSTREAM_BOUNDARY} \tag{280}$$

where θ is a unit vector in the tangential direction.

A factor σ is defined as:

$$\sigma = \sigma \times V_{\theta_TARGET_Avg} / V_{\theta_ACTUAL_AVG} \tag{281}$$

Finally, the tangential part of the source term addition to the momentum equation is:

$$v^2 \times \sigma \times V_{\theta_TARGET} \times \theta \tag{282}$$

Where v is the distance of the current cell from the fan upstream boundary, calculated along the fan axis. The term v ensures that the swirl increases gradually in the fan region.

Axial Component

The axial component of the source term addition to the momentum equation controls the pressure rise through the fan region. At each time step, the non-reversed volume flow rate is calculated through the fan region, and the corresponding pressure rise is obtained from the experimental table. This pressure rise is the target pressure rise that the fan

must achieve for the current mass flow rate. As the pressure rise is over the entire fan region, the target pressure rise per unit length is $\Delta P_{TARGET} / (Length_{FAN_REGION})$.

The actual pressure rise in the current time step is calculated as:

$$\Delta P_{ACTUAL} = P_{AVERAGE_DOWNSTREAM_BOUNDARY} - P_{AVERAGE_UPSTREAM_BOUNDARY} \quad (283)$$

Therefore the addition to the momentum equation source term is:

$$\frac{\Delta P_{TARGET}}{\Delta P_{ACTUAL}} \times \frac{\Delta P_{TARGET}}{Length_{FAN_REGION}} \vec{X} \quad (284)$$

where $\vec{X}$ is a unit vector in the axial direction of the fan.

Isentropic Flow Formulation

The following isentropic relations are useful when initializing compressible flow cases:

$$\frac{T_t}{T} = 1 + \frac{(\gamma - 1)}{2} M^2 \quad (285)$$

$$\frac{P_t}{P} = \left(1 + \frac{(\gamma - 1)}{2} M^2\right)^{\gamma/(\gamma - 1)} \quad (286)$$

$$V = M\sqrt{\gamma R T} \quad (287)$$

$$\frac{\dot{m}}{A} = \sqrt{\frac{\gamma}{R}} \frac{P_t}{\sqrt{T_t}} M \left(1 + \frac{(\gamma - 1)}{2} M^2\right)^{1/2 - \gamma/(\gamma - 1)} \quad (288)$$

$$\gamma = \frac{C_p}{C_p - R} \quad (289)$$

$$R = \frac{R_u}{M_r(X)} \quad (290)$$

where:

- A is area
- $\dot{m}$ is mass flow rate
- M is Mach number
- $M_r(X)$ is molecular weight (relative molecular mass) of molecular entity X
- P is static pressure

- P_t is total pressure
- R is specific gas constant
- R_u is universal gas constant [8314.4621 J/kmol K]
- T is static temperature
- T_t is total temperature
- V is velocity magnitude
- γ is ratio of specific heats

Supersonic Static Pressure Formulation

Supersonic static pressure, which can be given at a stagnation inlet or at a mass flow inlet, is used based on the speed of flow at the inflow boundary. If the velocity at the inlet is subsonic (Mach number $M < 1.0$), then supersonic static pressure is ignored. If $M \geq 1.0$, then the supersonic static pressure is employed in the application of the boundary condition. See for example Eqn. (286). See also Stagnation Inlet.

Adjoint Formulation

The adjoint formulation is divided into several subsections:

Simulation Inputs and Outputs

Sensitivity of the Objective with respect to the Design Inputs

Adjoint Method

Flow Adjoint

Adjoint Flow Solver

Mesh Sensitivity

Mesh Deformation

Cost Function Sensitivity with Respect to Boundary Parameters

Simulation Inputs and Outputs

Consider a use-case where the primary inputs (design variables) $\mathbf{d}$ are control points parametrizing a body surface mesh $\mathbf{x}_{\text{surf}}$

$$\mathbf{x}_{\text{surf}} = \mathbf{x}_{\text{surf}}(\mathbf{d})$$

The surface mesh is then input to the mesh generator to yield the volume mesh $\mathbf{x}$

$$\mathbf{x} = \mathbf{x}(\mathbf{x}_{\text{surf}}(\mathbf{d}))$$

The volume mesh, together with the boundary conditions for a given set of physics are input to the flow solver. With a given volume mesh and a set of physics to model, the flow solver provides the flow solution **w**

$$\mathbf{w} = \mathbf{w}(\mathbf{x}(\mathbf{x}_{\text{surf}}(\mathbf{d})))$$

The objective (cost function) L is further derived from the flow solution on the current mesh:

$$L = L(\mathbf{w}(\mathbf{x}(\mathbf{x}_{\text{surf}}(\mathbf{d}))),\mathbf{x}(\mathbf{x}_{\text{surf}}(\mathbf{d}))) \tag{291}$$

Sensitivity of the Objective with respect to the Design Inputs

To compute the sensitivity of the cost function L with respect to the design variables **d**, Eqn. (291) is differentiated in the forward direction. The chain rule is used starting from inputs through to the objective (outputs) to give the total sensitivity:

$$\begin{aligned}\frac{dL}{d\mathbf{d}} &= \frac{\partial L}{\partial \mathbf{x}} \cdot \frac{\partial}{\partial \mathbf{x}_{\text{surf}}}\mathbf{x} \cdot \frac{\partial \mathbf{x}_{\text{surf}}}{\partial \mathbf{d}} + \frac{\partial L}{\partial \mathbf{w}} \cdot \frac{\partial}{\partial \boldsymbol{x}}\mathbf{w} \cdot \frac{\partial}{\partial \mathbf{x}_{\text{surf}}}\mathbf{x} \cdot \frac{\partial \mathbf{x}_{\text{surf}}}{\partial \boldsymbol{D}} \\ &= \left[\frac{\partial L}{\partial \mathbf{x}} + \frac{\partial L}{\partial \mathbf{w}} \cdot \frac{\partial}{\partial \mathbf{x}}\mathbf{w}\right] \frac{\partial}{\partial \mathbf{x}_{\text{surf}}}\mathbf{x} \cdot \frac{\partial \mathbf{x}_{\text{surf}}}{\partial \mathbf{d}}\end{aligned} \tag{292}$$

where:

$[\partial L/\partial \mathbf{x} + \partial L/\partial \mathbf{w} \cdot \partial \mathbf{w}/\partial \mathbf{x}]$ is the sensitivity of the objective (and the flow solution) with respect to the volume mesh,

$\partial \mathbf{x}/\partial \mathbf{x}_{\text{surf}}$ is the sensitivity of the volume mesh with respect to the surface mesh and

$\partial \mathbf{x}_{\text{surf}}/\partial \mathbf{d}$ is the sensitivity of the surface mesh with respect to the design variables.

Adjoint Method

Typically, there are a only a few cost functions L and many design variables **d.** In this case, the adjoint method provides the most economical approach for computing sensitivities. Formation of the adjoint begins by transposing Eqn. (292) as follows:

$$\frac{dL}{d\mathbf{d}}^{\text{T}} = \frac{\partial \mathbf{x}_{\text{surf}}^{\text{T}}}{\partial \boldsymbol{d}} \cdot \left[\frac{\partial \mathbf{x}}{\partial \mathbf{x}_{\text{surf}}}^{\text{T}} \cdot \left(\frac{\partial L}{\partial \mathbf{x}}^{\text{T}} + \frac{\partial \mathbf{w}}{\partial \mathbf{x}}^{\text{T}} \cdot \frac{\partial L}{\partial \mathbf{w}}^{\text{T}}\right)\right] \tag{293}$$

Next, Eqn. (293) is evaluated by a series of matrix-vector multiplications from right to left. The right-most terms not involving **d** are precomputed and stored as a single vector. Consequently, the sensitivities with regard to the many design variables are obtained in a single matrix-vector multiplication.

Flow Adjoint

To compute the sensitivity of the flow solution **w** with respect to the volume mesh coordinates **x**, consider the residual of the discretized flow equations that has been driven to zero by means of the primal flow solver:

$$\mathbf{R}[(\mathbf{w}(\mathbf{x}),\mathbf{x})] = 0 \tag{294}$$

Linearizing Eqn. (294) about this converged state gives:

$$\left[\frac{\partial}{\partial \mathbf{w}}\mathbf{R}\right]\frac{\partial}{\partial \mathbf{x}}\mathbf{w} + \frac{\partial}{\partial \mathbf{x}}\mathbf{R} = 0 \tag{295}$$

which can be rewritten as:

$$\frac{\partial \mathbf{w}}{\partial \mathbf{x}}^{\mathrm{T}} = \frac{\partial \mathbf{R}}{\partial \mathbf{x}}^{\mathrm{T}}\left[\frac{\partial}{\partial \mathbf{w}}\mathbf{R}\right]^{-\mathrm{T}} \tag{296}$$

Inserting Eqn. (296) into Eqn. (293) yields:

$$\frac{dL}{d\mathbf{d}}^{\mathrm{T}} = \frac{\partial \mathbf{x}_{\mathrm{surf}}}{\partial \mathbf{d}}^{\mathrm{T}} \cdot \left[\frac{\partial \mathbf{x}}{\partial \mathbf{x}_{\mathrm{surf}}}^{\mathrm{T}} \cdot \left(\frac{\partial L}{\partial \mathbf{x}}^{\mathrm{T}} - \left(\frac{\partial \mathbf{R}}{\partial \mathbf{x}}^{\mathrm{T}} \cdot \frac{\partial \mathbf{R}}{\partial \mathbf{w}}^{-\mathrm{T}} \cdot \frac{\partial L}{\partial \mathbf{w}}^{\mathrm{T}}\right)\right)\right] \tag{297}$$

Defining the flow adjoint $\mathbf{\Lambda}$ in terms of the cost function L, the flow solution **w**, and the residual of the discretized flow equations **R** as follows:

$$\mathbf{\Lambda} = \left[\frac{\partial}{\partial \mathbf{w}}\mathbf{R}\right]^{-\mathrm{T}} \cdot \frac{\partial L}{\partial \mathbf{w}}^{\mathrm{T}} \tag{298}$$

and substituting this into Eqn. (297) gives:

$$\frac{dL}{d\mathbf{d}}^{\mathrm{T}} = \frac{\partial \mathbf{x}_{\mathrm{surf}}}{\partial \mathbf{d}}^{\mathrm{T}} \cdot \left[\frac{\partial \mathbf{x}}{\partial \mathbf{x}_{\mathrm{surf}}}^{\mathrm{T}} \cdot \left(\frac{\partial L}{\partial \mathbf{x}}^{\mathrm{T}} - \frac{\partial \mathbf{R}}{\partial \mathbf{x}}^{\mathrm{T}} \cdot \mathbf{\Lambda}\right)\right] \tag{299}$$

The flow adjoint is the inner-most (first) quantity necessary to obtain the total sensitivity.

Adjoint Flow Solver

The solution of the flow adjoint $\mathbf{\Lambda}$ requires solving the linear system of equations, given by (see Eqn. (298)):

$$\left[\frac{\partial}{\partial \mathbf{w}}\mathbf{R}\right]^{\mathrm{T}} \cdot \mathbf{\Lambda} = \frac{\partial L}{\partial \mathbf{w}}^{\mathrm{T}} \tag{300}$$

where

$[\partial \mathbf{R}/\partial \mathbf{w}]$: Sensitivity of the flow residual with respect to the flow variables. This represents the exact (2nd order) Jacobian of the flow equations.

$\partial L/\partial \mathbf{w}$: Sensitivity of the cost function with respect to the flow variables.

Note that the flow adjoint is an expression of the sensitivity of the cost function L with respect to the flow residual $\mathbf{R}$:

$$\mathbf{\Lambda}^{\mathrm{T}} \sim \frac{\partial L}{\partial \mathbf{R}} \tag{301}$$

The flow adjoint equations are solved by means of the following iterative defect-correction algorithm:

$$[\mathbf{P}]^{\mathrm{T}}(\mathbf{\Lambda}^{\mathrm{n+1}} - \mathbf{\Lambda}^{\mathrm{n}}) = \frac{\partial L}{\partial \mathbf{w}}^{\mathrm{T}} - \left(\left[\frac{\partial}{\partial \mathbf{w}}\mathbf{R}\right]^{\mathrm{T}} \cdot \mathbf{\Lambda}^{\mathrm{n}}\right) \tag{302}$$

where:

$[\mathbf{P}]$: Matrix preconditioner. Given by the 1st order Jacobian already implemented with in the coupled implicit solver.

$(\partial \mathbf{R}/\partial \mathbf{w})^{\mathrm{T}} \cdot \mathbf{\Lambda}^{\mathrm{n}}$: Matrix-vector product of flow adjoint solution at current iteration and exact Jacobian. Evaluated by two-pass accumulation due to extended stencil of full, 2nd order exact Jacobian.

The solution of the linear flow adjoint problem typically exhibits similar convergence rate as the primal flow solution and takes roughly as many iterations at similar computational cost.

Mesh Sensitivity

In Eqn. (299), the flow adjoint is obtained first by iterative solution. The remaining sensitivities are obtained by a straightforward matrix-vector multiplication running backwards up the chain. Next in the sequence is the mesh sensitivity:

$$\frac{dL}{d\mathbf{x}} = \left[\frac{\partial L}{\partial \mathbf{x}}^{\mathrm{T}} - \left(\frac{\partial \mathbf{R}}{\partial \mathrm{x}}^{\mathrm{T}} \cdot \mathbf{\Lambda}\right)\right] \tag{303}$$

where:

$\partial \mathbf{R} / \partial \mathbf{x}$: Sensitivity of flow residual with respect to the volume mesh coordinates.

$\partial L / \partial \mathbf{x}$: Sensitivity of cost function with respect to the volume mesh coordinates.

This gives the sensitivity of the cost function L with respect to the vertices of the volume mesh $\mathbf{x}$. The evaluation of the mesh sensitivity is a one-step, non-iterative process.

Mesh Deformation

Based upon the computed design point sensitivity, the new design point positions are calculated. The mesh is then deformed by use of the morpher to account for the new design point positions.

Evaluating the Interpolation Function

Using radial basis functions, the interpolated position $\mathbf{s}$ about the reference position $\mathbf{x}$ is given by:

$$\mathbf{s}(\mathbf{x}) = \boldsymbol{\alpha} + \sum_{i=1}^{N} \boldsymbol{\beta}_j \phi_j(r_j(\mathbf{x})) \tag{304}$$

where:

$$r_j(\mathbf{x}) = \left\| \mathbf{x} - \mathbf{x}_j \right\|$$

There are two types of radial basis functions: those with compact support and those with global support. Here, only one global support function is considered: the multiquadratic biharmonics function $\phi(r)$ given as:

$$\phi_j(r_j) = \sqrt{r_j^2 + c_j^2}$$

and the constant $c_j = 0, \forall j$. For a discrete set of M reference positions in the domain Ω, evaluation of all interpolation positions in that domain can be evaluated by casting Eqn. (304) as:

$$\begin{bmatrix} \mathbf{s}_1^{\Omega} \\ \vdots \\ \mathbf{s}_M^{\Omega} \end{bmatrix} = \begin{bmatrix} 1_1 \\ \vdots \\ 1_M \end{bmatrix} \boldsymbol{\alpha} + \begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} \\ \vdots & \ddots & \vdots \\ \phi_{M,1} & \cdots & \phi_{M,N} \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta}_1 \\ \vdots \\ \boldsymbol{\beta}_N \end{bmatrix} \tag{305}$$

where the components of the $M \times N$ matrix are computed via the radial basis function $\phi_{i,j} = \phi_j(r_j(\mathbf{x}_i))$.

Evaluating the Interpolation Function Coefficients

To make use of the continuous interpolation function Eqn. (304) or its discrete form Eqn. (305), the coefficients $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are required. This is achieved by making use of N known interpolation positions, or control positions $\mathbf{s}^{\Gamma}$, and the additional constraint:

$$\sum_{i=1}^{N} \boldsymbol{\beta}_j = \mathbf{0} \tag{306}$$

From Eqn. (304) and Eqn. (306), a system of equations is constructed to solve the unknown coefficients:

$$\begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} & 1_1 \\ \vdots & \ddots & \vdots & \vdots \\ \phi_{N,1} & \cdots & \phi_{N,N} & 1_N \\ 1_1 & \cdots & 1_N & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta}_1 \\ \vdots \\ \boldsymbol{\beta}_N \\ \boldsymbol{\alpha} \end{bmatrix} = \begin{bmatrix} \mathbf{s}_1^{\Gamma} \\ \vdots \\ \mathbf{s}_N^{\Gamma} \\ \mathbf{0} \end{bmatrix} \tag{307}$$

State, Objective, and Derivative

From a functional point of view, the independent variables of the morpher are $\mathbf{s}_j^{\Gamma}$, for $j = 1, \ldots, N$. Its dependent variable is $\mathbf{s}$ in the continuous form and $\mathbf{s}_i^{\Omega}$, for $i = 1, \ldots, M$ in the discrete. The positions $\mathbf{x}$ are constants of the system. Since only the discrete form of the morpher is of interest, the required Jacobian is then:

$$\mathbf{J} = \frac{d\mathbf{s}^{\Omega}}{d\mathbf{s}^{\Gamma}} \tag{308}$$

Derivative

Tangent

Dropping function arguments and adopting algorithmic differentiation notation for derivative terms (that is $\dot{\mu}$ and $\bar{\mu}$ for tangents and adjoints, respectively), a tangent of the Jacobian Eqn. (308) is:

$$\begin{bmatrix} \dot{\mathbf{s}}_1^{\Omega} \\ \vdots \\ \dot{\mathbf{s}}_M^{\Omega} \end{bmatrix} = \begin{bmatrix} 1_1 \\ \vdots \\ 1_M \end{bmatrix} \boldsymbol{\alpha} + \begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} \\ \vdots & \ddots & \vdots \\ \phi_{M,1} & \cdots & \phi_{M,N} \end{bmatrix} \begin{bmatrix} \dot{\boldsymbol{\beta}}_1 \\ \vdots \\ \dot{\boldsymbol{\beta}}_N \end{bmatrix} \tag{309}$$

whose coefficients $\dot{\boldsymbol{\alpha}}$ and $\dot{\boldsymbol{\beta}}$ are found from solving the following system:

$$\begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} & 1_1 \\ \vdots & \ddots & & \vdots \\ \phi_{N,1} & \cdots & \phi_{N,N} & 1_N \\ 1_1 & \cdots & 1_N & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta}_1 \\ \vdots \\ \boldsymbol{\beta}_N \\ \boldsymbol{\alpha} \end{bmatrix} = \begin{bmatrix} \dot{\mathbf{s}}_1^{\Gamma} \\ \vdots \\ \dot{\mathbf{s}}_N^{\Gamma} \\ \mathbf{0} \end{bmatrix} \tag{310}$$

Therefore, the requirement for solving the tangent derivative is the capacity to evaluate Eqn. (305) and Eqn. (307) for specific inputs.

Adjoint

Formulation of the adjoint begins with transposing each term of Eqn. (309)to get:

$$\bar{\boldsymbol{\alpha}} = \begin{bmatrix} 1_1 \\ \vdots \\ 1_M \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \bar{\mathbf{s}}_1^{\Omega} \\ \vdots \\ \bar{\mathbf{s}}_M^{\Omega} \end{bmatrix}$$

$$\begin{bmatrix} \bar{\boldsymbol{\beta}}_1 \\ \vdots \\ \bar{\boldsymbol{\beta}}_N \end{bmatrix} = \begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} \\ \vdots & \ddots & \vdots \\ \phi_{M,1} & \cdots & \phi_{M,N} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \bar{\mathbf{s}}_1^{\Omega} \\ \vdots \\ \bar{\mathbf{s}}_M^{\Omega} \end{bmatrix} \tag{311}$$

Upon evaluating Eqn. (311), the adjoint derivative of Eqn. (310) is solved:

$$\begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} & 1_1 \\ \vdots & \ddots & & \vdots \\ \phi_{N,1} & \cdots & \phi_{N,N} & 1_N \\ 1_1 & \cdots & 1_N & 0 \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \bar{\mathbf{s}}_1^{\Gamma} \\ \vdots \\ \bar{\mathbf{s}}_N^{\Gamma} \\ \boldsymbol{t} \end{bmatrix} = \begin{bmatrix} \bar{\boldsymbol{\beta}}_1 \\ \vdots \\ \bar{\boldsymbol{\beta}}_N \\ \bar{\boldsymbol{\alpha}} \end{bmatrix} \tag{312}$$

where the coefficient $\mathbf{t}$ is discarded. Since the matrix in Eqn. (312) is symmetric, it is in the same form as Eqn. (310). Eqn. (312) has a non-zero entry in the last element of the right-hand side. This difference is handled by solving an auxiliary equation:

$$\begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} & 1_1 \\ \vdots & \ddots & & \vdots \\ \phi_{N,1} & \cdots & \phi_{N,N} & 1_N \\ 1_1 & \cdots & 1_N & 0 \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \mathbf{s}'^{\Gamma}_1 \\ \vdots \\ \mathbf{s}'^{\Gamma}_N \\ \mathbf{t} \end{bmatrix} = \begin{bmatrix} \bar{\boldsymbol{\beta}}_1 - c_1 \\ \vdots \\ \bar{\boldsymbol{\beta}}_N - c_N \\ \mathbf{0} \end{bmatrix} \tag{313}$$

where c_j, for $j = 1, \ldots, N$ are computed from

$$\begin{bmatrix} \phi_{1,1} & \cdots & \phi_{1,N} & 1_1 \\ \vdots & \ddots & & \vdots \\ \phi_{N,1} & \cdots & \phi_{N,N} & 1_N \\ 1_1 & \cdots & 1_N & 0 \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \bar{\boldsymbol{\alpha}}/N \\ \vdots \\ \bar{\boldsymbol{\alpha}}/N \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} c_1 \\ \vdots \\ c_N \\ \bar{\boldsymbol{\alpha}} \end{bmatrix} \tag{314}$$

The solution for Eqn. (312) can be obtained by $\bar{s}_i^{\Gamma} = s'^{\Gamma}_i + \bar{\boldsymbol{\alpha}}/N$, for $i = 1, \ldots, N$.

Cost Function Sensitivity with Respect to Boundary Parameters

Consider a cost function of the form:

$$L = L(\mathbf{w}, \mathbf{w}^{\infty}) \tag{315}$$

where $\mathbf{w}$ is the solution vector of the following set of nonlinear, discrete equations:

$$\mathbf{R}(\mathbf{w}) = \mathbf{0}$$

and $\mathbf{R}$ is the residual vector of the flow equations. The solution of the flow adjoint $\boldsymbol{\Lambda}$ is obtained from Eqn. (300). Since the solution $\mathbf{w}$ to this set of equations depends on the imposed boundary conditions $\mathbf{w}^{\infty}$, there is an implicit dependence of $L = L(\mathbf{w}, \mathbf{w}^{\infty})$ on $\mathbf{w}^{\infty}$ in addition to the explicit dependence. The objective is to quantify this dependency. Quantifying the dependency is achieved in a straight-forward manner by applying the chain rule:

$$\frac{dL}{d\mathbf{w}^{\infty}} = \frac{\partial L}{\partial \mathbf{w}^{\infty}} + \left(\frac{\partial L}{\partial \mathbf{w}}\right)^{\mathrm{T}} \frac{\partial}{\partial \mathbf{R}}\mathbf{w} \frac{\partial}{\partial \mathbf{w}^{\infty}}\mathbf{R} \tag{316}$$

The above equation can be simplified by using:

$$\left(\frac{\partial L}{\partial \mathbf{w}}\right)^{\mathrm{T}} \frac{\partial}{\partial \mathbf{R}}\mathbf{w} = \boldsymbol{\Lambda}$$

This simplification leads to:

$$\frac{dL}{d\mathbf{w}^{\infty}} = \frac{\partial L}{\partial \mathbf{w}^{\infty}} + \boldsymbol{\Lambda} \frac{\partial}{\partial \mathbf{w}^{\infty}}\mathbf{R} \tag{317}$$

From Eqn. (317) it can be seen that the computation of the cost function sensitivity with respect to boundary conditions involves two steps:

- The differentiation of the explicit dependency of the cost function on the imposed boundary conditions.

- The accounting of its implicit dependence that is given by the second term in Eqn. (317), which requires obtaining the global adjoint solution.

Virtual Disk Formulation

The virtual disk models the effect of a rotating turbo machine on the surrounding flow field. The geometry of the turbo machine is not resolved. Instead, the body forces created by the rotating turbo machine are distributed over a cylindrical disk volume.

See also:

Body Force Propeller Method

Blade Element Method

1D Momentum Method

Body Force Propeller Method

The body force propeller method is used to model the effects of a propeller such as thrust and torque and thereby creating propulsion without actually resolving the geometry of the propeller. The method employs a uniform volume force $\mathbf{f}_b$ distribution over the cylindrical virtual disk. The volume force varies in radial direction.

The radial distribution of the force components follows the Goldstein optimum and is given by:

$$f_{bx} = A_x r^* \sqrt{1 - r^*} \tag{318}$$

$$f_{b\theta} = A_\theta \cdot \frac{r^* \sqrt{1 - r^*}}{r^*(1 - r_h') + r_h'} \tag{319}$$

$$r^* = \frac{r' - r_h'}{1 - r_h'} \tag{320}$$

$$r_h' = \frac{R_H}{R_P} \quad \text{and} \quad r' = \frac{r}{R_P}$$

where f_{bx} is the body force component in axial direction, $f_{b\theta}$ is the body force component in tangential direction, r is the radial coordinate, R_H is the hub radius and R_P is the propeller tip radius.

The constants A_x and A_θ are computed as:

$$A_x = \frac{105}{8} \cdot \frac{T}{\pi\Delta(3R_H + 4R_P)(R_P - R_H)} \tag{321}$$

$$A_\theta = \frac{105}{8} \cdot \frac{Q}{\pi\Delta R_P(3R_H + 4R_P)(R_P - R_H)} \tag{322}$$

where T is the thrust, Q is the torque, and R_P is the hub radius of the propeller.

The computation of the body force components requires several user inputs. A propeller performance curve needs to be specified, which gives the dimensionless thrust coefficient K_T, the torque coefficient K_Q, and the propeller efficiency η_0 as functions of the advance ratio J.

$$K_T = \frac{T}{\rho n^2 d^4} \tag{323}$$

$$K_Q = \frac{Q}{\rho n^2 d^5} \tag{324}$$

$$J = \frac{V_A}{nD_P} \tag{325}$$

where V_A is the speed of advance of the propeller, n the rotation rate, and D_P the propeller diameter.

Further inputs are the position of the propeller within the computational domain, the direction of the propeller rotational axis, and the direction of rotation.

The simulation is performed for a certain operating point. The operating point can be specified by either of these quantities:

- Rotation rate n
- Thrust T
- Torque Q

If the operation point is given by rotation rate n, the procedure for obtaining the body force component distribution over the virtual disk is the following:

1. The advance ratio J is calculated as:

$$J = \frac{V_{\text{Inflow Plane}}}{nD_P} \tag{326}$$

2. The thrust coefficient K_T and the torque coefficient K_Q are interpolated from the specified propeller performance curve:

$$K_T, K_Q = f(J)$$

3. With K_T and K_Q available, thrust T and torque Q are calculated for the propeller:

$$T = \frac{K_T \rho V_{\text{Inflow Plane}}^2 \; D^2{}_P}{J^2} \tag{327}$$

$$Q = \frac{K_Q \rho V_{\text{Inflow Plane}}^2 \; D_P^3}{J^2} \tag{328}$$

4. With T and Q available, the axial and the tangential body force components f_{bx} and $f_{b\theta}$ are calculated according to Eqns. (318) to (322).

If the operation point is given by thrust T or by torque Q, the first step of above procedure is replaced by:

1. With a given T or Q, the advance ratio is calculated by solving the following equation numerically:

$$f(J) = K_T - K_T' \text{ or } f(J) = K_Q - K_Q' \tag{329}$$

where K_T and K_Q are evaluated from the propeller performance curve. K_T' and K_Q' are evaluated as:

$$K_T' = \frac{J^2 T_{\text{Operating Point}}}{\rho_{\text{Inflow Plane}} V_{\text{Inflow Plane}}^2 \; D_P^2} \tag{330}$$

$$K_Q' = \frac{J^2 Q_{\text{Operating Point}}}{\rho_{\text{Inflow Plane}} V_{\text{Inflow Plane}}^2 \; D_P^3} \tag{331}$$

With the advance ratio J available, the calculation of the axial and tangential body force components follows the same procedure as for the rotation rate operating point continuing with the second step.

Blade Element Method

The blade element method formulation is based on the methodology proposed by Rajagoplan and others. [53]. It is primarily designed to model the effects of a helicopter rotor on the flow field without actually prescribing the exact geometry of rotor blades. This approach consists of two parts. The first step is to identify the location of cells within the computational domain that belong to the virtual disk and to which the source term is applied. The second step involves the calculation of the momentum source term vector for each cell, which was identified and marked in the first step.

Step 1: Cell marking and the allocation to the blade elements/buckets based on the interpolation grid

The information for this step is user-specified and provided through the virtual disk nodes:

- Disk Geometry

  This describes the geometrical dimensions and the orientation of the disk

- Virtual Disk Resolution

  This specifies the size of the two-dimensional interpolation grid in polar coordinates, representing the actual blade elements. The blade properties such as chord, twist, and aerodynamic characteristics prevail throughout the whole element.

The source term is calculated per blade element on the interpolation grid and then transferred to the finite volume mesh. Each element is associated with a number of cells on the finite volume mesh. For the coupling between the local flow conditions and the modeled rotor, the information is taken from the finite volume cell which is the closest to the center of a blade element on the interpolation grid. Thus, the specified size of the interpolation grid is the actual computational effort of the influence of the rotor on the flow field. The sketch below explains the flow field coupling between the finite volume mesh (red) and the virtual disk interpolation grid (black).



The marking procedure is itself a two-step procedure. First, all the cells that fall within the circular disk area are marked to receive the source term. Secondly, each of the marked cells is assigned a bucket/element number corresponding to its location with respect to the interpolation grid. So each of the cells that was marked during the first step belongs to one of the elements/buckets of the interpolation grid.

This approach requires that each element/bucket has at least one cell in its domain. If the finite volume mesh is very coarse and the interpolation grid is too fine then you are asked to reduce the resolution. The following two figures show the interpolation grid for two different resolutions.





For the fine interpolation grid in the second figure, it can be seen that no finite volume cell comes under the domain of some for some of the elements/buckets located at the innermost radius.

Step 2: Calculation of Momentum Source Term Strength

The second part of the blade element method calculates the strength of momentum source term for the elements on the interpolation grid. Based on the geometry of the rotor (that is orientation and origin), a disk local cartesian coordinate system (X', Y', Z') is created. This cartesian coordinate system has its origin at the center of the rotor and the Z' -axis in the direction corresponding to the thrusting direction of the disk. The Z' -axis is

perpendicular to the plane of rotation, while the X' and Y' axes lie in the plane of the disk. The following sketch demonstrates the disk local coordinate system:



The rotational speed ω is always oriented along the Z'-direction. $+\omega$ corresponds to a counter-clockwise wise (CCW) rotation and $-\omega$ corresponds to a clock-wise (CW) rotation when viewed from the $+Z'$ direction.

The local flow field velocity in the laboratory coordinate system is:

$$\mathbf{v} = v_x\hat{\imath} + v_y\hat{\jmath} + v_z\hat{k} \tag{332}$$

The local flow velocity is transferred to the disk local coordinate system using the transformation matrix M:

$$\begin{bmatrix} v_{x'} \\ v_{y'} \\ v_{z'} \end{bmatrix} = M \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} \tag{333}$$

The blade has a velocity due to rotation, which can be written in the disk local coordinate system as:

$$\mathbf{v}^{bl} = \boldsymbol{\omega}' \times \mathbf{R}'$$

where $\boldsymbol{\omega}' = \omega_{x'}\hat{\imath} + \omega_{y'}\hat{\jmath} + \omega_{z'}\hat{k}$ and $\mathbf{R}' = R_{x'}\hat{\imath} + R_{y'}\hat{\jmath} + R_{z'}\hat{k}$.

$\mathbf{v}^{bl}$ can be rewritten as:

$$\mathbf{v}^{bl} = (\omega_{y'}R_{z'} - \omega_{z'}R_{y'})\hat{\imath} + (\omega_{z'}R_{x'} - \omega_{x'}R_{z'})\hat{\jmath} + (\omega_{x'}R_{y'} - \omega_{y'}R_{x'})\hat{k} \tag{334}$$

As $\omega_{x'} = \omega_{y'} = 0$, thus $\boldsymbol{\omega}' = \omega_{z'}$. With this, $\mathbf{v}^{bl}$ can be written as:

$$\mathbf{v}^{bl} = -\omega_{z'}R_{y'}\hat{\imath} + \omega_{z'}R_{x'}\hat{\jmath} \tag{335}$$

The velocity relative to the blade is thus given by:

$$\begin{aligned}\mathbf{v}^{rel} &= \mathbf{v}' - \mathbf{v}^{bl} \\ &= (v_{x'} + \omega_{z'}R_{y'})\hat{\imath} + (v_{y'} - \omega_{z'}R_{x'})\hat{\jmath} + v_{z'}\hat{k}\end{aligned}$$

To calculate the effective angle of attack α, the induced angle of attack β, the blade twist angle θ, and the pitch angle θ_r, which varies in the azimuthal direction, must be taken into account. The induced angle of attack is calculated as:

$$\beta = \operatorname{atan}\left[\frac{-v_{z'}^{rel}}{\sqrt{(v_{x'}^{rel})^2 + (v_{y'}^{rel})^2}}\right] \tag{336}$$

The pitch angle θ_r is defined as:

$$\theta_r = \theta_{r0} - \theta_{rc}\cos\psi - \theta_{rs}\sin\psi \tag{337}$$

where ψ is the azimuth angle, θ_{r0} is the collective pitch angle, θ_{rc} is the lateral cyclic pitch angle, and θ_{rs} is the longitudinal cyclic pitch angle.

The final angle of attack is then calculated as:

$$\alpha = \theta_r + \theta - \beta \tag{338}$$

Knowing the angle of attack α and the blade relative velocity $\mathbf{v}^{rel}$ experienced by the airfoil, the Mach number or Reynolds number can be calculated. Which of the two numbers is calculated depends on in which format the aerodynamic characteristics of the airfoil section are available. The aerodynamic force coefficients C_l, $C_d = f(\alpha, Ma)$ or $f(\alpha, Re)$ are obtained from the user-specified tables. The lift and drag forces are then computed as:

$$\begin{aligned}L' &= \frac{1}{2}\cdot\rho\cdot\left|\mathbf{v}^{rel}\right|^2\cdot C_l\cdot c\cdot dr \\ D' &= \frac{1}{2}\cdot\rho\cdot\left|\mathbf{v}^{rel}\right|^2\cdot C_d\cdot c\cdot dr\end{aligned} \tag{339}$$

where c is the blade chord. These forces are then converted into the normal and tangential force (transformation from the wind axis to the rotor axis):

$$
\begin{aligned}
F_n' &= L'\cos\beta - D'\sin\beta \\
F_t' &= L'\sin\beta - D'\cos\beta
\end{aligned}
$$

Next, the forces are converted into the disk local coordinate system:

$$
\begin{aligned}
F_x' &= F_t' \cdot \sin\psi \\
F_y' &= -F_t' \cdot \cos\psi \\
F_z' &= F_n'
\end{aligned}
\tag{340}
$$

Finally, the force vector is transformed into the laboratory coordinate system:

$$
\begin{bmatrix} F_x \\ F_y \\ F_z \end{bmatrix} = M^{\mathrm{T}} \begin{bmatrix} F_x' \\ F_y' \\ F_z' \end{bmatrix}
\tag{341}
$$

The instantaneous force acting on the fluid element is then $-\mathbf{F}$.

The time-averaged source terms $\mathbf{S} = (S_x, S_y, S_z)$ are added to the discretized momentum conservation equations. Since the blade only spends a finite fraction of its total revolution time passing through the control volume, the time-averaged source terms are scaled by the time fraction:

$$
\mathbf{S} = \frac{b\Delta\phi}{2\pi}(-\mathbf{F})
\tag{342}
$$

where b is the number of blades and $\Delta\phi$ is the angular distance through which the blade traverses in passing through a blade element.

1D Momentum Method

A wind turbine extracts mechanical energy from the kinetic energy of the wind. It acts as an actuator disk, slowing down the wind speed from U_0 to U at the rotor plane, and further slowing it down to U_1 in the wake. Moving upstream to downstream, the velocity decreases and the streamlines diverge to comply with continuity.

The 1D momentum method adds a source term to the momentum equations to model the effects of a wind turbine. The source term accounts for both the axial and the tangential effects as given by the Ideal Horizontal Axis Wind Turbine (HAWT) with wake rotation. When wake rotation is included in the analysis, the induced velocity at the rotor consists of an axial component a and a tangential component a' acting in the rotor plane.

The formulation is based on the use of an annular stream tube with a radius r and a thickness dr. (Note that this is not the disk thickness.) This definition results in a cross-sectional area dA of $2\pi r dr$. The following figure shows the schematics and the notation that is used in the formulation:



Thrust Distribution

Total thrust of a wind turbine is given by:

$$T = \frac{1}{2}\rho_0 U_0^2 \; C_t \pi (R_0^2 - R_i^2) \tag{343}$$

Thrust on an annular ring of width dr at a radius r is:

$$\Delta T = \frac{1}{2}\rho_0 U_0^2 \; C_t 2\pi r dr \tag{344}$$

The theory assumes an infinitesimally thin disk with zero thickness, whereas the source terms are added for a virtual disk of a finite thickness δ. Therefore, the elemental thrust ΔT is to be understood as thrust per Δr annulus per unit disk thickness δ:

$$\Delta T = \frac{1}{2}\rho_0 U_0^2 \; C_t 2\pi r \Delta r \tag{345}$$

Using Eqn. (343), this can be rewritten as:

$$\Delta T = T \frac{2r\Delta r}{(R_0^2 - R_i^2)} \tag{346}$$

Thrust per unit volume within the annulus thus reads:

$$\frac{\Delta T}{\Delta V} = \frac{1}{2} \cdot \frac{\rho_0 U^2{}_0 C_t}{\delta} = \frac{T}{\delta\pi(R_0^2 - R_i^2)} \tag{347}$$

where $\delta\pi(R_0^2 - R_i^2)$ is the geometrical volume of the disk.

Thrust per unit volume is independent of r. Thrust depends on r quadratically, but the volume of each ring also increases quadratically with r, so thrust per unit volume is constant. For each cell within a ring:

$$T_{cell} = \frac{\Delta T}{\Delta V} V_{cell} \tag{348}$$

Eqn. (348) needs to be scaled, because not all cells are fully within the disk. The scaling factor is:

$$\frac{V_{geom}}{\sum_{\cdot\cdot} V_{cell}}$$

where

$$V_{geom} = \pi(R_0^2 - R_i^2)$$

The thrust per cell is given by:

$$T_{cell} = T \cdot \frac{V_{cell}}{\sum_{\cdot\cdot} V_{cell}} \tag{349}$$

T_{cell} corresponds to the axial force F_{axial} for a given cell, whose cell center lies within the virtual disk.

Torque Distribution

The torque distribution corresponds to the calculation of the tangential force. With a given velocity at the inflow plane U_0, the power P delivered by the wind turbine can be extracted from the user specified power curve table or polynomial. From this value, the torque is evaluated as:

$$Q = \frac{P}{\Omega}$$

where Ω is the rotation rate. This torque is distributed as tangential force over the disk cells.

The elemental torque dQ for an annular ring of width dr and thickness δ are obtained by applying the conservation of angular momentum. On an annular ring element, this gives:

$$dQ = d\dot{m}(\omega r)(r) = (\rho U 2\pi r dr)(\omega r)(r) \tag{350}$$

Since $U = U_0(1-a)$ and $a' = \omega/(2\Omega)$, the expression for the elemental torque reduces to:

$$dQ = 4a'(1-a)\frac{1}{2}\rho U_0 \Omega r^2 2\pi r dr \tag{351}$$

Integrating Eqn. (351):

$$\int dQ = 4a'(1-a)\frac{1}{2}\rho U_0 \Omega 2\pi r \int_{R_i}^{R_0} r^3 dr \tag{352}$$

The torque now is:

$$Q = a'(1-a)\frac{1}{2}\rho U_0 \Omega \pi ({R_0}^4 - {R_i}^4) \tag{353}$$

To derive an expression for torque per unit of annular ring volume, Eqn. (351) is rewritten by use of Eqn. (353):

$$\Delta Q = \frac{4Qr^3 \Delta r}{{R_0}^4 - {R_i}^4} \tag{354}$$

Torque per unit volume of annular ring $\Delta V = 2\pi\delta r \Delta r$ is then:

$$\frac{\Delta Q}{\Delta V} = \frac{2Qr^2}{\pi\delta \cdot ({R_0}^4 - {R_i}^4)} \tag{355}$$

Torque for each cell within the annular ring is:

$$Q_{cell} = \frac{\Delta Q}{\Delta V} \cdot V_{cell} = \frac{2Qr_{cell}^2 V_{cell}}{\pi\delta \cdot ({R_0}^4 - {R_i}^4)} \tag{356}$$

Scaling needs to be applied, because $\sum V_{cell} \neq V_{geom}$. The scaling factor reads:

$$scf = \frac{V_{geom} \cdot \frac{1}{2}(R_0^2 + R_i^2)}{\sum_{\cdot\cdot}^{\cdot\cdot} r_{cell}^2 \, V_{cell}} \tag{357}$$

With scaling applied:

$$Q_{cell} = \frac{2Qr_{cell}^2 V_{cell}}{\pi\delta \cdot (R_0^{\ 4} - R_i^{\ 4})} \cdot scf \tag{358}$$

With $V_{geom} = \pi(R_0^2 - R_i^2)\delta$, Eqn. (358) can be further simplified:

$$Q_{cell} = \frac{Qr_{cell}^2 V_{cell}}{\sum_{\cdot\cdot\ \cdot\cdot} r_{cell}^2 \ V_{cell}} \tag{359}$$

The tangential force $F_{\mathrm{tan}gential}$ for a given cell whose cell center lies inside the virtual disk is given by:

$$F_{\mathrm{tan}gential} = \frac{Qr_{cell}^2 V_{cell}}{\sum_{\cdot\cdot\ \cdot\cdot} r_{cell}^2 \ V_{cell}} \tag{360}$$

The axial and the tangential forces are used to assemble a source term vector with respect to the laboratory coordinate system. This source term is then added to the momentum equations.

Fluid Flow Bibliography

[17] Bird, R.B., Stewart, W.E., and Lightfoot, E.N. 1966. “Transport Phenomena”. John Wiley & Sons, New York.

[18] Burton,T., Jenkins, N., Sharpe, D., Bossanyi, E. 2011. “Wind Energy Handbook”, 2nd ed., Wiley.

[19] Chima, R.V., and Liou, M.-S. 2003. “Comparison of the AUSM+ and H-CUSP Schemes for Turbomachinery Applications”, NASA/TM-2003-212457.

[20] Demirdzic, I., Lilek, Z., and Peric, M. 1993. “A collocated finite volume method for predicting flows at all speeds”, *Int. J. for Numerical Methods in Fluids*, **16**, pp. 1029-1050.

[21] Demirdzic, I. and Muzaferija, S. 1995. “Numerical method for coupled fluid flow, heat transfer and stress analysis using unstructured moving meshes with cells of arbitrary topology”, *Comput. Methods Appl. Mech. Eng.*, pp. 1-21.

[22] Ferziger, J.H. and Peric, M. 2002. “Computational Methods for Fluid Dynamics”, 3rd rev. ed., Springer-Verlag, Berlin.

[23] Giles, M.B. 1988. “Non-reflecting boundary conditions for the Euler equations”, CFDL Report 88-1, MIT Dept. of Aero. and Astro.

[24] Giles, M.B. 1990. "Non-reflecting boundary conditions for Euler equation calculations", *AIAA Journal*, **28**(12):2050-2058.

[25] Giles, M.B. 1991 "UNSFLO: A numerical method for unsteady flow in turbomachinery", Gas Turbine Laboratory Report GTL 205, MIT Dept. of Aero. and Astro.

[26] Gnoffo, P. A., Gupta, R. N., and Shinn, J. L. 1989. "Conservation Equations and Physical Models for Hypersonic Air Flows in Thermal and Chemical Nonequilibrium", NASA TP-2867.

[27] Graboski, M. S. and Daubert, T. E. 1978. "A Modified Soave Equation of State for Phase Equilibrium Calculations: Hydrocarbon Systems", *Ind. Eng. Chem. Process Des. Dev.*, **17**(4), pp. 443-448.

[28] Graboski, M. S. and Daubert, T. E. 1978. "A Modified Soave Equation of State for Phase Equilibrium Calculations: Systems Containing CO_2, H_2S, N_2, and CO", *Ind. Eng. Chem. Process Des. Dev.*, **17**(4), pp. 448-454.

[29] Gupta, R. N., Lee, K. P., Thompson, R. A., and Yos, J. M. 1991. "Calculations and Curve Fits of Thermodynamic and Transport Properties for Equilibrium Air to 30000 K", NASA STI/Recon Technical Report N, 92, 11285.

[30] Hansen, M. 2007. "Aerodynamics of Wind Turbines, 2nd ed. Routledge.

[31] Hosangadi, A., Sachdev, J. , Sankaran, V. 2012. "Improved Flux Formulations for Unsteady Low Mach Number Flows". ICCFD7 proceedings.

[32] Huang, H., and Ayoub, J. 2006. "Applicability of the Forchheimer Equation for Non-Darcy Flow in Porous Media". SPE Journal, March 2008, pp. 112-122.

[33] International Association for the Properties of Water and Steam. 2007. *Revised Release on the IAPWS Industrial Formulation 1997 for the Thermodynamic Properties of Water and Steam*, IAPWS Secretariat, Lucerne, Switzerland (available at www.iapws.org).

[34] Jameson, A., Schmidt, W. and Turkel, E. 1981. "Numerical solution of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes", AIAA Paper 81-1259.

[35] Liou, M.-S. 1996. "A Sequel to AUSM: AUSM+", *J. Comput. Phys.*, **129**: 364-382.

[36] Liou, M.-S. 2001. "Ten Years in the Making--AUSM-Family", AIAA 2001-2521.

[37] Lockwood, B. A. 2012. "Gradient-based Approaches for Sensitivity Analysis and Uncertainty Quantification within Hypersonic Flows", Ph.D. Dissertation, University of Wyoming.

[38] Loehmer, R. 2009. "On Limiters for Minimal Vorticity Dissipation", AIAA 2009-135.

[39] Lynn, N. and Steinhoff, J. 2007. "Large Reynolds Number Turbulence Modeling with Vorticity Confinement", AIAA 2007-3965.

[40] Makkonen, L. 2000. "Models for the growth of rime, glaze, icicles and wet snow on structures." *Phil. Trans. R. Soc. Lond. A*, **358**, pp. 2913-2939.

[41] Manwell, J.F., McGowan, J., and Rogers, A. 2009. "Wind Energy Explained: Theory, Design and Application", 2nd ed. Wiley.

[42] Mathur, S.R. and Murthy, J.Y. 1997. "Pressure-based method for unstructured meshes", *Numerical Heat Transfer*, *Part B: Fundamentals*, **31**(2), pp. 195-214.

[43] Mathur, S.R. and Murthy, J.Y. 1997. "Pressure boundary conditions for incompressible flow using unstructured meshes", *Numerical Heat Transfer*, *Part B: Fundamentals*, **32**(3), pp. 283-298.

[44] Peric, M., Kressler, R., and Scheuerer, G. 1988. "Comparison of finite-volume numerical methods with staggered and colocated grids", *Computers & Fluids*, **16**(4), pp. 389-403.

[45] Poling, B. E., Prausnitz, J. M., and O'Connell, J. P., "The Properties of Gases & Liquids", 5th Ed., McGraw-Hill, 2001.

[46] Reid, R. C., Prausnitz, J. M., and Poling, B. E., "The Properties of Gases & Liquids", 4th Ed., McGraw-Hill, 1987.

[47] Saxer, A.P. 1992. "A numerical analysis of three-dimensional inviscid rotor/stator interactions using non-reflecting boundary conditions", Ph.D. Thesis. MIT Dept. of Aero. and Astro.

[48] Soave, G. 1972. "Equilibrium Constants from a Modified Redlich–Kwong Equation of State", *Chem. Eng. Sci.*, **27**, pp. 1197-1203.

[49] Steinhoff, J., Lynn, N. and Wang, L. 2005. "Computation of High Reynolds Number Flows Using Vorticity Confinement: I. Formulation", UTSI Preprint.

[50] Turkel, E. 1987. "Preconditioned methods for solving the incompressible and low speed compressible equations", *Journal of Computational Physics*, **72**, pp. 277-298.

[51] Weiss, J.M., Maruszewski, J.P., and Smith, W.A. 1999. "Implicit solution of preconditioned Navier-Stokes equations using algebraic multigrid", *AIAA Journal*, **37**(1), pp. 29-36.

[52] Weiss, J.M., and Smith, W.A. 1995. "Preconditioning applied to variable and constant density flows", *AIAA Journal*, **33**(11), pp. 2050-2057.

[53] Rajagopalan, R.G. 1993. "Three Dimensional Analysis of a Rotor in Forward Flight", *Journal of the American Helicopter Society*, Vol. 38.

[54] Yang, Z. et al. 2000. "Recent Improvements to a Hybrid Method for Rotors in Forward Flight", AIAA Conference Reno.

Modeling Heat Transfer

Heat transfer is the study of energy in transit due to a temperature difference in a medium or between media. Heat transfer extends thermodynamic analysis through the study of the modes of energy transfer and through development of relations to calculate energy transfer rates.

This transfer of heat takes place through three distinct modes of heat transfer:

- Heat Conduction
- Convective Heat Transfer
- Thermal Radiation

In most problems, heat transfer is coupled to fluid flow through convection, and coupled to thermodynamics through fluid properties (for example, density, material conductivity, and fluid viscosity) and phase change (for example, condensation, evaporation, melting, and freezing). Therefore, from a modeling perspective, heat transfer must be seen as a multi-physics activity.

STAR-CCM+ provides additional models that let you create sophisticated simulations for specific applications:

- Boiling — where you are interested in modeling the improved heat transfer due to boiling, and the creation of the vapor bubbles has little effect on the flow
- Moisture Boiling in Solids — where you are interested in modeling the boiling of moisture within a solid material
- Thin Film — where you are interested in simulating deicing or defogging on a surface
- Heat exchangers — heat transfer between two fluid streams: a hot and a cold stream (such as in radiators and air coolers)

In addition to descriptions of the models mentioned above, the following sections help clarify how this modeling works:

- Information on solid energy models
- Explanation of modeling fully developed energy
- Instructions for setting an environment boundary condition for heat transfer

When you activate these models, other objects appear in the GUI that are designed to give you added control over the flow modeling. Some of them are common to more than one flow or energy model, or common to flow and energy models. They are described in more detail in the following sections:

- Flow reference
- Energy reference
- Values common to flow and energy models
- Field functions common to flow and energy models
- Energy formulation

Heat Conduction

Heat conduction is the flow of internal energy from a region of higher temperature to a region of lower temperature by the interaction of the adjacent particles (for example, atoms, molecules, ions, or electrons) in the intervening space. This flow can be viewed as the transfer of energy from the more energetic particles to the less energetic particles of a substance due to interactions between the particles.

Conduction takes place in all forms of matter: solids, liquids, gases, and plasmas. In solids, heat conduction is due to the combination of vibrations of the molecules in a lattice and diffusion of free electrons. In gases and liquids, heat conduction is due to collisions between molecules and the diffusion of molecules during their random motion. Conduction is greater in solids because of the relatively close fixed spatial relationships between atoms, whereas fluids and gases are less conductive due to the large distance between atoms.

The law of heat conduction, also known as Fourier's law, provides a means of calculating the local heat flux. This law can be stated in two equivalent forms:

- The integral form, which considers the amount of energy flowing into or out of a body as a whole.
- The differential form, which considers the flow rates or fluxes of energy locally:

$$\mathbf{q}'' = -k\nabla T \tag{361}$$

  In this expression, $\mathbf{q}''$ $[W/m^2]$ is the local heat flux vector, k $[W/mK]$ is the thermal conductivity of the material, and ∇T $[K/m]$ is the temperature gradient. The thermal conductivity of the material can be a strong function of temperature, pressure (for gases), spatial location, and can exhibit anisotropic behavior (that is, it varies with orientation and is represented by a second-order tensor). The negative sign in the expression is because $k > 0$ and heat always moves from a high temperature region to a low temperature (as defined in the Second Law of Thermodynamics).

Thermal conductivities vary from material to material, being the largest for metallic solids, smaller for nonmetallic solids, very small for liquids, and extremely small for gases.

The major objective of performing a simulation that includes conductive heat transfer is to determine the temperature field in a medium, taking into account the conditions that are imposed on the boundaries. The First Law of Thermodynamics (that is, Conservation of Energy) and Fourier's law can be used to derive the Heat Diffusion Equation (a second-order partial

differential equation in temperature, or enthalpy, or energy) from which the temperature field can be calculated. This temperature solution can be steady (a function of space only and not changing with time), or unsteady (the temperature field varies in both space and time as, for example, in metal quenching applications).

Convective Heat Transfer

Convective heat transfer is the transfer of thermal energy by the movement of a fluid. In engineering practice, it is used to provide specific temperature changes. Examples include heat exchangers, electronic device cooling, and cooling of turbine blades.

Convective heat transfer comprises two mechanisms: energy transfer due to random molecular motion (diffusion), and energy transfer due to the larger-scale bulk motion of the fluid (advection - the collective movement of groups of molecules).

Convective heat transfer to or from a solid requires the transfer of heat by the bulk motion of the fluid and the transfer of heat by diffusion. The contribution due to random molecular motion dominates near the surface where the fluid velocity is low. At the interface between the surface and the fluid, the fluid velocity is zero (that is, there is no relative motion between the fluid and surface) and heat transfer is transferred by diffusion only.

Convection is usually described as being either *natural* ("free") or *forced*, although other mechanisms also exist (for example, gravitational, granular, thermomagnetic, capillary action-driven, or Marangoni and Weissenberg driven effects).

See also: Natural Convection and Forced Convection.

Convective heat transfer at a surface is governed by Newton's law of cooling:

$$q_s'' = h(T_s - T_{ref}) \tag{362}$$

In this expression, q_s'' $[W/m^2]$ is the local surface heat flux, h $[W/m^2K]$ is the local convective heat transfer coefficient, T_s is the surface temperature, and T_{ref} is a characteristic temperature of the fluid moving over the surface.

Newton's law of cooling expresses a linear relationship between the local surface heat flux and the difference between the local surface and fluid temperatures. This linear relationship is only an approximation: in reality

the relationship can be strongly nonlinear. Because flow conditions can vary from point to point on the surface, both q_s'' and h can vary as a function of space and time.

Natural Convection

In a gravitational field, natural convection occurs due to temperature differences which affect the density, and thus the relative buoyancy, of the fluid. Denser components fall, while lighter (less dense) components rise, leading to bulk fluid movement.

Natural convection is more likely and more rapid when any of the following conditions apply:

- a greater variation in density between two mixing fluids
- a larger acceleration due to gravity that drives the convection
- a larger distance through the convecting medium.

Natural convection is less likely and/or less rapid when any of the following conditions apply:

- a more rapid diffusion (which acts to diffuse away the thermal gradient that causes the convection)
- a more viscous (sticky) fluid.

To quantify when natural convection is possible, the dimensionless Rayleigh number is defined and can be viewed as the ratio of buoyant and viscous forces times the ratio of momentum and thermal diffusivities. Natural convection becomes possible when the Rayleigh number exceeds 10^5.

Forced Convection

In forced convection, fluid movement results from external sources (for example, a fan, pump, or the action of a propeller), and is typically used to increase the rate of heat transfer at a surface in cooling or heating applications.

In cases of mixed convection (that is, natural convection and forced convection occurring together) you can determine how much is due to forced convection and how much is due to natural convection. The relative magnitudes of the Rayleigh, Prandtl, and Reynolds numbers indicate which form of convection dominates:

If $\frac{Re}{PrRa^2} \gg 1$, forced convection can be neglected.

If $\frac{Re}{PrRa^2} \ll 1$, natural convection can be neglected.

If the ratio is approximately one, then both forced and natural convection need to be taken into account.

Thermal Radiation

Thermal radiation is the emission of electromagnetic waves from all matter that has a temperature greater than absolute zero, and represents a conversion of thermal energy into electromagnetic energy. It is generated by the thermal motion of charged particles in matter, which results in charge-acceleration and dipole oscillation. This drives the electrodynamic generation of coupled electric and magnetic fields, which cause the emission of thermal radiation.

Electromagnetic radiation, or light, does not require the presence of matter to propagate, and travels the fastest in a vacuum. All forms of matter emit radiation. For gases and some semi-transparent solids (for example, glass and salt crystals at elevated temperatures) emission is a volumetric phenomenon (that is, emission is an integrated effect of local emission throughout the volume). For most solids and liquids, radiation that is emitted from the interior molecules is strongly absorbed by adjoining molecules. Therefore, because all of the emitted radiation originates from molecules near the surface, it can be considered as a surface phenomenon.

Electromagnetic energy propagation can be described at various levels of complexity, depending on the features of the radiation field to be described. For example, if polarization is important, the Tokes parameters are used to describe the intensity field. If wave effects are important, the vector wave quantities from electromagnetic theory such as electric vector, magnetic vector, and the Poynting vector are used. Wave effects become significant when the characteristic length scale for transport, such as spacing between the surfaces, is similar to the wavelength.

For more information, see Modeling Radiation.

For calculating heat transfer in most situations, however, this level of description is overly complicated. For most heat transfer applications, thermal radiation can be treated as unpolarized (multiple reflections and scattering usually nullify polarization effects) and incoherent (waves or photons are usually out of phase). The length scale for transport is usually

much larger than the wavelength of radiation, so the limiting description of geometric optics (that is, the wavelength approaches zero) can be applied: the waves are described as bundles of rays carrying energy in a small volume that is associated with the solid angle in the direction of the rays.

The maximum flux at which radiation can be emitted from a surface is given by the Stefan-Boltzmann law:

$$q''_{bb} = \sigma T_s^{\,4} \tag{363}$$

In this expression, q''_{bb} $[W/m^2]$ is the local surface heat flux, T_s is the temperature of the surface, and σ is the Stefan-Boltzmann constant $\sigma = 5.67 \times 10^{-8}$ $[W/m^2K^4]$. Such a surface is called an ideal radiator or blackbody.

The heat flux that is emitted by a real surface is less than the heat flux of an ideal radiator, and is given by:

$$q''_e = \varepsilon q''_{bb} = \varepsilon\sigma T_s^{\,4} \tag{364}$$

where $0 \le \varepsilon \le 1$ is the surface emissivity (that is, the relative ability of a surface to emit energy by radiation).

If radiation is incident upon a surface, a portion is absorbed, a portion is reflected, and a portion is transmitted through the material:

$$q''_{inc} = q''_{abs} + q''_{ref} + q''_{trans} = \alpha q''_{inc} + \rho q''_{inc} + \tau q''_{inc} \tag{365}$$

The constants in this expression are:

$0 \le \alpha \le 1$	The surface absorptivity: the fraction of incident radiation that gets absorbed.
$0 \le \rho \le 1$	The reflectivity: the fraction of incident radiation that gets reflected away from the surface.
$0 \le \tau \le 1$	The transmissivity: the fraction of incident radiation that gets transmitted through the material.

Absorptivity, reflectivity, and emissivity depend on the surface temperature, surface roughness, emission angle, and wavelength of the radiation (the radiation is not monochromatic, but comprises a continuous dispersion of photon energies).

Verification and Validation of Results

Verification and validation are independent procedures that are used together to ensure that a code meets requirements and specifications, and that it fulfills its intended purpose. In simple terms, validation can be expressed by the query "*Are you building the right thing?*" and verification by the query "*Are you building it right?*"

In other words, verification determines whether the programming and computational implementation of the conceptual model (that is, the partial differential equation system) is correct by comparison to exact analytical results. For example, does the temporal and spatial rate of convergence match the theory? Verification is more concerned with mathematics than engineering. By contrast, validation determines whether the computational simulation agrees with physical reality by comparison to experimental results. For example, in quenching, does the temporal variation of the surface temperature match experimental results?

Because experimental data is not always available for complex multi-physics applications, you must use a building-block approach consisting of phases that involve successively more complex flow physics, geometry, and their interactions. When validating multi-physics problems combined with complex geometry and forcing, the initial validation problem should be as simple as possible and only involve single physics and simple geometry. For example, if you are interested in forced flow over a cooled turbine blade, you could start with isothermal flow over a flat plate. Once the simulation code has been validated using single physics and simple geometry, you can add more physics and increased geometrical complexity until the actual complete system is being simulated (that is, all of the relevant physics and geometry are present).

Designing a Heat Transfer Optimization Study

In mathematics, optimization is the selection of the best element (with respect to some criteria) from a set of available alternatives. In the simplest case, an optimization problem consists of maximizing or minimizing a real function (that is, the objective function) by systematically choosing input values from within an allowed set and computing the value of the function.

More generally, optimization includes finding the "best available" values of some objective function given a defined domain, including a variety of different types of objective functions and constraints. Multi-objective optimization (that is, an optimization problem that has more than one

objective function) adds significant complexity to optimization problems. It is not always possible to satisfy all of the objective functions, so a trade-off can be required. For example, you cannot minimize both the surface temperature and the surface heat flux. Optimization problems are also often multi-modal: they possess multiple good solutions that can all be globally good (that is, they have the same objective function value) or the solutions could be a mix of globally and locally good solutions. When classical optimization techniques are used to obtain multiple solutions, they do not perform satisfactorily due to their iterative approach. That is, it is not guaranteed that different solutions are obtained even with different starting points in multiple runs of the algorithm. However, evolutionary algorithms are able to obtain multiple solutions and have become a popular technique for cases in which more than one solution is desired.

As far as the engineer is concerned, there are two important objectives that optimization can help with:

- Code tuning and validation

  For example, to reproduce previous computations or experiments of interest.

- Meeting design specifications and constraints by allowing you to choose the ”best” set of models, parameters, and geometry for the current application.

  For example, what shape maximizes heat flux but is constrained to have a fixed volume.

When you use optimization in flow and heat transfer applications, address the following points:

- Decide on the objective function (also known as a cost function or utility function) that you are interested in minimizing or maximizing.

  For example, minimize the maximum surface temperature, minimize the downstream vorticity after a pipe bend or junction, or maximize the surface heat flux.

- Decide on the set of models and parameters that are important and need to be considered in the search for objective function extrema.

  For example, material properties, length of a boundary, surface shape, domain size, and boundary conditions.

  It is useful to perform a one-parameter-at-a-time sensitivity study over the input parameter space, and identify any parameters that have a significant impact on the solution (these parameters are sometimes referred to as ”big knobs”). You can measure sensitivity by monitoring changes in the objective function value, looking at its partial derivatives, or by using linear regression. However, this approach does not fully explore the input space, since it does not take into account the

simultaneous variation of input variables, and therefore cannot detect input variable interactions. Modelers frequently prefer this technique because of its simplicity and because when the model fails the modeler knows immediately which input parameter caused the failure.

- The mathematical techniques that are used to solve optimization problems depend on the form of the objective function and constraints, as well as the desired use of the results (for example, code tuning or design).

  You must decide which optimization technique best fits the current problem: is finding a local solution good enough, or must you find the global solution, or do you require multiple solutions?

The simplest situations are unconstrained problems, or problems in which all the constraints can be expressed as equality relationships. For these types of problems, the technique of Lagrangian multipliers can be used efficiently to find the optimal solution.

Another technique that is frequently used is response surface analysis. In this technique, you define a finite set of plausible parameter value collections (for example, based on uniform latin hypercube sampling) using the truncated input parameter space (obtained from, for example, the one-parameter-at-a-time sensitivity analysis), and then fit a response surface (either local or global) to the objective function results. Once the response surface has been created, it is easy to use calculus to define a system whose solutions represent the minimum or maximum of the objective function over the truncated input parameter space considered. However, solving this system is not always easy. This technique works well when the objective function is locally or globally smooth with respect to input parameter space. This smoothness makes it possible to get a good representation of the solution hyper-surface with the least number of objective function values.

For problems that involve inequality relationships, the Lagrangian multiplier method is not suited to solving problems of this type efficiently. Modern mathematical programming techniques, such as linear- and nonlinear-programming, in which both the objective and constraint can be linear or non-linear functions, must be used to find optimal solutions.

Convective Heat Transfer Coefficients

Convective heat transfer is the transfer of energy between a surface and a fluid moving over the surface. It is primarily used as a boundary condition when solving solid/fluid heat transfer problems, such as quenching, electronic device cooling, and heat exchangers.

Convective heat transfer is governed by Newton's law of cooling:

$$q_s'' = h(T_s - T_{ref}) \tag{366}$$

where q_s'' is the local surface heat flux (that is, power per unit area), h is the local convective heat transfer coefficient, T_s is the surface temperature, and T_{ref} is a characteristic temperature of the fluid moving over the surface.

Newton's law of cooling expresses a linear relationship between the local surface heat flux and the difference between the local surface and fluid temperatures, and is only an approximation. Because flow conditions can vary from point to point on the surface, both q_s'' and h can vary as a function of space and time.

In Eqn. (366), q_s'', T_s, and T_{ref} are fundamental in nature, while the heat transfer coefficient is a constant of proportionality that relates the three fundamental parameters. While q_s'' and T_s are unambiguous, there is some latitude in the choice of the fluid temperature T_{ref}. Depending upon the choice of T_{ref}, the heat transfer coefficient is different in order to satisfy Eqn. (366). Therefore, heat transfer coefficients cannot be defined without also defining a reference temperature: there is an infinite number of heat transfer coefficient and reference temperature combinations that give the same surface heat flux.

In STAR-CCM+, the conceptual centerpiece for modeling convective heat transfer at the wall for turbulent flows stems from the standard wall functions (SWF). Consider the expression for the local surface heat flux:

$$q_s'' = \frac{\rho_f(y_c)C_{p,f}(y_c)u_\tau}{T^+(y^+(y_c))}(T_s - T_c) \tag{367}$$

You can then equate the terms from Eqn. (366) and Eqn. (367) and define an h in terms of the local flow conditions.

In this equation, ρ_f is the fluid density, $C_{p,f}$ is the fluid-specific heat capacity, u_τ is a velocity scale that is based on the wall shear stress, T^+ is the dimensionless temperature, $y^+ = (u_\tau y_c)/\upsilon_f$ is a Reynolds number, and y_c and T_c are the normal distance and temperature of the near-wall cell, respectively.

The Standard Wall Functions (SWFs) are a set of semi-empirical functions that are used to satisfy the flow physics in the near-wall region (the boundary layer). They give relationships for T^+ and u_τ in terms of the laminar and turbulent Prandtl numbers, the dimensionless near-wall flow velocity, and the turbulent kinetic energy.

Four heat transfer coefficients are defined in STAR-CCM+:

- Local Heat Transfer Coefficient
- Specified y+ Heat Transfer Coefficient
- Heat Transfer Coefficient
- Virtual Local Heat Transfer Coefficient

These heat transfer coefficients are all grounded to varying degrees in the standard wall function treatment. The local heat transfer coefficient is the one that is used internally within the STAR-CCM+ solver. The other three heat transfer coefficients are computed during post-processing.

The differences between these three heat transfer coefficient post-processing options are discussed in Guidelines for Heat Transfer Coefficients.

Local Heat Transfer Coefficient

The heat transfer coefficient option that is alluded to by Eqn. (367) is the one that the STAR-CCM+ code uses internally to calculate the local surface heat flux.

If you equate the Local Heat Transfer Coefficient and Local Reference Temperature terms from Eqn. (366) and Eqn. (367):

$$h_1 = \frac{\rho_f(y_c)C_{p,f}(y_c)u_\tau}{T^+(y^+(y_c))} \text{ and } T_{ref,1} = T_c \tag{368}$$

you have the pair (h_1 and $T_{ref,1}$) which can be viewed or exported.

In STAR-CCM+, the local surface heat transfer coefficient is always calculated (and used internally) in the code using the Standard Wall Functions and Eqn. (368). However, there can be times when you want to view or export heat transfer coefficients that are defined in other ways. For example, to compare the results to previous work or to use with other applications. In general, you can accomplish this by post-processing the simulation results and defining the one-to-one linear mapping (based on equating the surface heat fluxes):

$$h_{user} = h\left(\frac{T_s - T_c}{T_s - T_{ref}}\right) \tag{369}$$

In this way, you can define your own h and T_{ref} pair, where h, T_c, and T_s are known from the simulation, and you pick T_{ref} to suit your requirements.

Specified y+ Heat Transfer Coefficient

The first built-in post-processing option in STAR-CCM+ calculates the heat transfer coefficient as in Eqn. (368), but at a user-specified y^+ value instead of the value that is associated with the near-wall cell:

$$h_2 = \frac{\rho_f(y_c)C_{p,f}(y_c)u_\tau}{T^+(y^+_{user})} \tag{370}$$

Then, using the surface heat flux as calculated in Eqn. (367), you solve Eqn. (366) for the reference temperature:

$$T_{ref,2} = T_s - \frac{q_s''}{h_2} \tag{371}$$

The result is the pair (h_2 and $T_{ref,2}$) which can be viewed or exported.

This definition of the reference temperature does not guarantee that $T_{ref,2}$ lies within the minimum and maximum temperature range of the local fluid region. If $T_{ref,2}$ lies outside this temperature range, then either the minimum or maximum temperature is used in its place, and the heat transfer coefficient is modified accordingly.

Heat Transfer Coefficient

The second built-in post-processing option in STAR-CCM+ calculates the surface heat flux using Eqn. (367), and then recasts Eqn. (366) as:

$$h_3 = \frac{q_s''}{T_s - T_{ref,3}} \tag{372}$$

The heat transfer coefficient is calculated using a user-defined reference temperature ($T_{ref,3}$).

The result is the pair (h_3 and $T_{ref,3}$) which can be viewed or exported.

Virtual Local Heat Transfer Coefficient

The third built-in post-processing option in STAR-CCM+ does not need to solve the energy equation in the fluid. Instead, this option uses the Local Heat Transfer Coefficient (Eqn. (368)), which is calculated using only flow field information. You specify the specific heat, the molecular Prandtl number, and the turbulent Prandtl number as reference values for this heat transfer coefficient.

The result is the pair (h_1 and $T_{ref,4}$) which can be viewed or exported.

Guidelines for Heat Transfer Coefficients

This section discusses the recommendations for accurate use of the Standard Wall Functions and the built-in post-processing heat transfer coefficients.

STAR-CCM+ always uses Eqn. (367) to obtain the local surface heat flux when solving for fluid temperatures (and/or solid temperatures in CHT applications). This expression embodies the important boundary layer physics, and it is important that you follow the recommendations to ensure its proper application.

The heat transfer coefficients are STAR-CCM+ post-processing results which can be used for comparing to other solutions, visualization, or exporting to other applications such as Abaqus or Nastran. The particular choice of heat transfer coefficient does not affect the simulation results, except in the case when STAR-CCM+ is two-way-coupled to another application that solves for the temperature in the solid. The choice of heat transfer coefficient that is applied in an external application does affect the results in that application. In this case, it is important that you choose the correct option, as the choice of heat transfer coefficient that is exported to the external application does affect the STAR-CCM+ results.

However, if you are solving for both the fluid and solid temperatures within STAR-CCM+ (that is, Conjugate Heat Transfer), your choice of post-processing heat transfer coefficients option has no effect on the simulation results.

In general, standard wall functions give reasonable accuracy for most high-Reynolds-number, wall-bounded flows. However, standard wall functions reach their limitation when the flow conditions differ too much from the ideal conditions that are used to define the functions. This limit can be reached in the following cases:

- Pervasive low-Reynolds-number or near-wall effects (for example, flow through a small gap or highly viscous low-velocity fluid flow).

- Massive transpiration through the wall (such as blowing or suction).
- Severe pressure gradients leading to boundary layer separation.
- Strong body forces (for example, flow near rotating disks or buoyancy-driven flows).
- High three-dimensionality in the near-wall region (for example, Ekman spiral flow or strongly skewed 3D boundary layers).

Thus to get good results when using the standard wall functions, the problem should not contain any of the physics conditions that are outlined above. The y^+ of the near-wall cell should be in the inertial sublayer ($30 < y^+ < 150$), and temperature-dependent fluid properties should be used to ensure maximum accuracy. If the near-wall cell lies in the buffer layer ($5 < y^+ < 30$) between the viscous and inertial sublayers, the "All y^+" option (default) should be used which smoothly connects the viscous and inertial sublayers.

The differences between the heat transfer coefficient post-processing options are discussed in the following sections:

- Heat Transfer Coefficient
- Local Heat Transfer Coefficient
- Specified y^+ Heat Transfer Coefficient
- Virtual Local Heat Transfer Coefficient

Heat Transfer Coefficient

This option is primarily used for comparing heat transfer coefficients to experimental data or other previous work.

When you use this option, note the following:

- It is recommended that you follow all recommendations for getting accurate results using the standard wall functions.
- For external flows where there is a well-defined free-stream temperature outside the thermal boundary layer, the usual practice is to use this temperature as the reference temperature.

  For internal flows without axial recirculation, it is often practical to define the reference temperature as the local bulk temperature that is obtained by solving the one-dimensional (stream-wise) energy equation. If the bulk temperature varies little relative to the overall fluid-wall temperature difference, you can set the reference temperature to the inlet temperature.
- It is recommended that temperature-dependent fluid properties be used to ensure maximum accuracy.

If you use temperature-dependent fluid properties, all fluid properties should be temperature-dependent, rather than just some of the properties.

- A poor choice of reference temperature can lead to a negative heat transfer coefficient.
- This option is not recommended for coupled use with other applications.

Local Heat Transfer Coefficient

When you use this option, note the following:

- To ensure maximum accuracy, use temperature-dependent fluid properties.

  If you use temperature-dependent fluid properties, all fluid properties should be temperature-dependent, rather than just some of the properties. For example, to get the Prandtl Number correct you must account for temperature dependence in both the momentum and thermal diffusivity.
- Because y^+ of the near-wall cell can lie in the viscous sublayer, the inertial sublayer, or anywhere in between, this option can give heat transfer coefficients that are different from the values that are found in classic textbooks.
- This option is not recommended for coupled use with other applications unless there is a reason for not using Specified y^+ Heat Transfer Coefficient.

  For more information, see Specified y^+ Heat Transfer Coefficient.

Specified y^+ Heat Transfer Coefficient

When you use this option, note the following:

- This heat transfer coefficient and reference temperature pair is recommended to be used when coupling STAR-CCM+ to other applications.

  This heat transfer coefficient and reference temperature pair is less dependent on the mesh than the Local Heat Transfer Coefficient, gives better conditioned systems at solve time, and gives the fastest convergence for a given number of coupling-iterations. One coupling-iteration involves solving for the flow field given the initial solid wall temperature and then solving for the solid temperature field given h and T_c (that is, the near wall cell temperature) from the flow field computation. A single coupling-iteration can give good results, although at least two coupling-iterations are recommended.

  For accuracy and numerical stability reasons, it is not recommended that you use alternative techniques of exchanging boundary conditions between STAR-CCM+ and other applications. For example, when the fluid code passes the local heat flux to the solid code and the solid code passes the local wall temperature to the fluid code.
- This option can have implications for cycle averaging and mapping.
- It is recommended that you follow all recommendations for getting accurate results using the standard wall functions.
- The user-specified value of y^+ should be less than approximately 150. A value of 100 has been shown to work well.
- It is recommended that temperature-dependent fluid properties be used to ensure maximum accuracy.

  If you use temperature-dependent fluid properties, all fluid properties should be temperature-dependent, rather than just some of the properties.

Virtual Local Heat Transfer Coefficient

When you use this option, note the following:

- It is recommended that you follow all recommendations for getting accurate results using the SWFs.
- This option is useful when you are not interested in the temperature distribution in the fluid and it is easy to define a reasonable global reference temperature. For example, in an external/internal flow where the temperature of the fluid is not expected to change much due to surface heat flux.

For external flows where there is a well-defined free-stream temperature outside the thermal boundary layer, the usual practice is to use this temperature as the reference temperature. For internal flows without axial recirculation, it is often practical to define the reference temperature as the local bulk temperature that is obtained by solving the one-dimensional (stream-wise) energy equation. If the bulk temperature varies little relative to the overall fluid-wall temperature difference, you can set the reference temperature to the inlet temperature.

- The y^+ of the near-wall cell should be in the inertial sublayer $30 < y^+ < 150$ because it is more likely that the user specified reference temperature is close to what is used in the SWFs to calculate h.
- This option should not be used for laminar flows because calculating the surface heat flux accurately requires the near-wall cell to lie in the viscous sublayer. However, as mentioned above, having the near-wall cell in the viscous sublayer is not recommended.
- This option is not recommended for coupled use with other applications.

Setting an Environment Boundary Condition for Heat Transfer

The main purpose of this feature is to permit the modeling of both convection and radiation heat transfer between the external sides of domain boundaries and the environment. This modeling is achieved through a certain type of thermal boundary condition (or more specifically, an option for Thermal Specification) called the Environment condition (in contrast to a Temperature, Flux, or Convection condition).

This condition is analogous to a Baffle Interface where the physics continuum on one side of the baffle is replaced with the environment. Similar to a baffle interface, the Environment condition can have radiative transfer to/from both surface sides and distinct surface temperatures for each side. In contrast to the baffle interface, the Environment condition is applied using a single boundary (using the Thermal Specification option). This single boundary behaves like a dual-sided boundary when the condition is activated. In contrast, the baffle interface is applied using two single-sided boundaries.

Modeling external radiation transfer with the environment condition is only available with the Surface-to-Surface radiation transfer model.

Understanding the Physics and Modeling

The physics that is being considered is shown in the following diagram.



Specifically, the physics of interest consists of:

- Convection and radiation from the internal domain to the internal side of the boundary. The internal surface can absorb, emit, reflect, and transmit radiation.

- Conduction through the boundary. The boundary can have any finite thermal resistance, including zero.
- Convection and radiation from the environment to the external side of the boundary. The external surface can absorb, emit, reflect, and transmit radiation.

The modeling of this physics is depicted in the following diagram. In contrast to the existing Convection thermal condition, the Environment thermal condition additionally considers the following:

- Absorption and emission of radiation on the external side.
- A non-zero wall thermal resistance.
- An external surface temperature distinct from the internal surface temperature. This distinction is the result of allowing from a non-zero wall thermal resistance.

The last two items are not necessarily unique to the Environment condition (that is, they could potentially be applied for the other thermal conditions). Currently, however, the non-zero thermal resistance is only available with the Environment condition. Thermal heat capacity is not considered with the Environment condition. If the wall has sufficient thermal capacity and an unsteady analysis is being considered, then the wall can be modeled using a thin solid. An Environment condition can optionally be applied to the external boundary of the solid.



$$q_{inc,ext} = q_{inc,sol} + q_{inc,env} + q_{inc,other\ bodies}$$

Setting the Environment Condition

The feature is activated on a per-boundary basis by changing the Thermal Specification option for the given boundaries.



Specifically, the steps are:

- Activate the Environment condition:
  - Select the Physics Conditions > Thermal Specification node for the boundary and set the Condition property to Environment.
- Specify boundary values under the Physics Values node:
  - Specify the External Emissivity and External Reflectivity for the outward facing side.

These settings are in addition to the standard Surface Emissivity and Surface Reflectivity for the inward-facing side. The existing Surface Transmissivity serves as the transmissivity for both sides (it must be symmetrical to satisfy the second law, just like at interfaces).

- Specify the Heat Transfer Coefficient and Ambient Temperature for convection from the external side to the environment.

  These settings are the same profiles that are currently used for the Convection condition.

- Specify the Thermal Resistance for the wall.

  The resistance can be specified as a total wall resistance or as a multi-layer resistance. The multi-layer resistance is specified using the `Multi-Layer Resistance` method on the Thermal Resistance profile. This setting is similar to the multi-layer thermal resistance for baffles.

- Specify radiation environment attributes in the radiation-specific models as usual.

  The attributes include the environment radiation temperature and the solar loads. Unlike the per-boundary inputs, these attributes are global (for example, the specified solar loads are applied to all relevant boundaries). This step has not changed from previous usage, but the specified radiative environment conditions now interact with the external sides of boundaries that have the Environment condition activated.

- Run the simulation.

- Post-process the results.

  The field functions for the inward facing side have the standard names/titles (such as Temperature or Boundary Irradiation).

  The field functions for the outward facing side have names/titles that reflect the “External” namespace. The field function name is prefixed with “External” (such as ExternalTemperature). The field function titles are suffixed with “on External Side” (for example Temperature on

External Side). Typical results are shown in the following screenshot, an example of a field function display for internal and external sides of boundaries.



Using the Environment Condition with S2S External Ray Tracing

The external sides of the boundaries with the Environment condition activated are patched, and view factors are computed for these patches as with any other patches. Often the external sides only see the environment and thus have a single unity view factor to the environment. However, if the external side of the boundary "sees" any other boundaries that also have an environment condition (or have an outward S2S surface orientation), then the boundary has view factors to these other boundaries.

Limitations and Restrictions

- External radiation modeling is only available for S2S (not currently available for DOM).
- External surfaces with the Environment condition activated must only see other boundaries that have either of the following:
  - have the Environment condition activated
  - are attached to an opaque solid region that has its S2S surface orientation that is set to `Outward`.

  This restriction is similar to that governing the outward S2S surface orientation.

- The Environment condition must not be applied to boundaries of regions that have an outward S2S surface orientation. The Environment condition provides an outward facing surface (the primary role of the outward orientation).
- The Environment condition is not currently available for flow openings such as inlets and outlets.

Example: Net Radiative Transmission Calculation on External Transmissive Boundaries

With the environment condition, there is an easy way to calculate the net radiative transmission on external transmissive boundaries. The following information illustrates the complete energy balance for a simulation with transparent external boundaries. It involves a window with convection from the outside: `hinf=10W/m2-K, Tinf=350`.

For this case with the environment condition, the net transmission in the window balances the heat transfer at all the boundaries (to about five digits relative to rate at the inlet):

```
        Surface Integral of Net Transmission on Volume Mesh
Part                                    Value (W)
------------------------------          -------------
Body 2: Window                           5.895503e-01
                                        -------------
                          Total:         5.895503e-01

        Heat Transfer of Boundary Heat Transfer on Volume
Mesh
Part                                    Value (W)
------------------------------          -------------
Body 2: Box                              0.000000e+00
Body 2: Inlet                           -4.314370e+02
Body 2: Outlet                           4.323347e+02
Body 2: Side                             0.000000e+00
Body 2: Tubes                            0.000000e+00
Body 2: Window                          -3.105699e-01
                                        -------------
                          Total:         5.871046e-01
```

The transmission was computed using a user field function `[$BoundaryTransmissivity*($ExternalBoundaryIrradiation - $BoundaryIrradiation)]` that has been surface integrated over the transmissive boundary (that is, the window).

Example: Results for an Auto Cabin

The feature is applied to a thermal analysis of an auto cabin. Specifically, the Environment condition is applied to the boundaries representing the windows of the car. The windows are partially transparent, allowing both the solar radiation and the environment radiation from the outside to be absorbed on the external side and to be transmitted into the cabin interior. In addition, the windows have finite thermal resistance to conduction.

The following four illustrations display various field functions for the internal and external sides of the driver side window—the internal side being displayed on the left and the external on the right. The following screenshot displays the direct solar irradiation on internal and external sides of the driver window. As expected, the flux on the internal side is zero, and the flux on the external size is finite. The highest flux is near the top as the window bends further from vertical.



The following screenshot shows the net radiative heat flux on both sides.



The flux on the external side is highly influenced by the solar flux (both direct and diffuse). The flux is relatively high over the entire glass. In contrast, the net radiative flux on the internal side is much less. The internal side largely interacts with the interior of the cabin which has temperatures similar in scale to the window temperature. The higher fluxes on the bottom of the internal side are due to solar radiation coming through the window. This radiation hits the horizontally oriented part of the door just below the window, and is then diffusely reflected back into the internal side of the window.

The following screenshot shows the temperature on the internal and external sides.



The window has a thermal resistance so the temperatures can be different. However, since the resistance is relatively small, the temperatures end up being similar, with the external temperatures slightly higher than the internal ones.

Lastly, this next screenshot displays the net boundary heat fluxes for the internal and external sides.



In order to conserve energy at the window, these fluxes must be equal and of opposite sign. As seen in the figure, this behavior is realized. The external side generally has a net heat flux to the window from the environment. The internal side generally has a net heat flux from the window to the cabin interior. The fluxes on the opposite sides balance each other to machine precision.

Modeling Fully Developed Energy

This section explains what fully developed flows are, presents formulation of the conditions for fully developed energy and the algorithm, and describes situations where errors are incurred. Information is provided about variable heating situations, validation, and recirculation.

What Are Fully Developed Flows?

"Fully developed flows" occur in long tubes or tube banks. For simple flows without heat transfer or transported scalars, all flow variables except pressure can be treated as periodic. Pressure is typically modeled with a "jump" condition at the interface that compensates for the streamwise pressure drop through the system.

When energy is present, a fully developed situation can be deemed to occur if the profiles of temperature and/or enthalpy become self-similar when using an appropriate scaling. In these cases, the self-similarity relations can be used to rescale the thermal field leaving the domain in order to reintroduce it at the inflow periodic boundary. This rescaling is done in STAR-CCM+ by using the Fully Developed Flow Option. The treatment of the thermal field is more complex than pressure, since both convection and diffusion are taken into account. In addition, the limitations on which physical models are present in order for the fully developed condition to occur must also be considered.

Under Which Conditions Is the Fully Developed Assumption Valid for Energy?

Consider the flow in a duct in which the streamwise coordinate is x, and the cross-stream coordinate is y. By definition, fully developed heat transfer exists when the heat transfer coefficient h is independent of x:

$$h = \frac{\dot{q}_w}{T_w - T_b} = -\frac{k}{T_w - T_b}\left.\frac{\partial T}{\partial y}\right|_w \neq f(x) \tag{373}$$

For a constant-property fluid, the bulk temperature T_b is defined:

$$T_b = \frac{1}{AU_m}\int_A uTdA \tag{374}$$

and the mean velocity U_m is defined as:

$$U_m = \frac{1}{A}\int_A udA \tag{375}$$

It is assumed that u and U_m are independent of x, but T_b is a function of x. Since T_w and T_b are independent of y, the following is obtained:

$$Nu = \frac{hl}{k} = -\frac{\partial}{\partial(y/l)}\left(\frac{T_w - T}{T_w - T_b}\right)\bigg|_w \neq f(x) \tag{376}$$

implying the existence of a dimensionless temperature profile:

$$\theta(y/l) = \frac{T_w - T}{T_w - T_b} \neq f(x) \tag{377}$$

or equivalently, since $\partial\theta/\partial x = 0$:

$$\frac{\partial T_w}{\partial x} - \frac{\partial T}{\partial x} = \frac{T_w - T}{T_w - T_b}\left(\frac{\partial T_w}{\partial x} - \frac{\partial T_b}{\partial x}\right) \tag{378}$$

In the *constant wall heat flux* case, since both h and $\dot{q}''$ are constant, $T_w - T_b$ equals a constant, so that:

$$\frac{\partial T}{\partial x} = \frac{\partial T_w}{\partial x} = \frac{\partial T_b}{\partial x} \tag{379}$$

In other words, for a constant wall heat flux, both the bulk temperature and the wall temperature increase linearly with streamwise distance (see the figure of constant heat flux). This linear increase means that the temperature profiles at various streamwise locations differ only by a constant offset as follows:

$$T(r, x) - T_b(x) = T(r, x + L) - T_b(x + L) \tag{380}$$

This means that the problem can be solved *without knowledge of the wall temperature*, if one knows the increase in bulk temperature that is due to the (constant) wall heat flux.



Behavior of bulk and wall temperatures for fully developed thermal fields with a constant heat flux

In the *constant wall temperature* case, the partial derivative of temperature in the streamwise direction in the energy equation, which results from the convective term, is replaced with:

$$\frac{\partial T}{\partial x} = \theta \frac{\partial T_b}{\partial x} \tag{381}$$

such that the rest of the energy equation can be written as a function purely of y/l. However, the temperature profiles must be scaled using both the bulk temperature and the wall temperature. Furthermore, the bulk

temperature increases non-linearly with streamwise distance x in the duct, as shown in the following figure, resulting in a corresponding non-linear increase in local temperature.



Behavior of bulk and wall temperatures for fully developed thermal fields with a wall temperature

Therefore this section defines two and only two situations that permit fully developed heat transfer in a duct:

- Wall with constant heat flux
- Walls with a constant temperature

The resulting temperature profiles can be different, depending on the geometry. The following figure shows the profile shapes for laminar constant-property flow through circular ducts.



Temperature profiles for fully developed thermal fields in laminar pipe flows

The constant heat flux boundary condition applies to a wall subject to electrical resistance heating or to a radiant flux. The constant temperature wall is approached when the outer wall is subject to phase change such as condensation or boiling. The two heating situations are also applicable to heat exchangers. The constant heat flux situation can correspond to a counterflow heat exchanger when the fluid capacity rates are the same. The constant temperature solution is applicable to any heat exchanger where one material has a very much larger capacity rate than the other. In addition to evaporators and condensers, the other material can be a heated solid of large thermal capacity.

How Does the Fully Developed Energy Algorithm Work?

The basic mechanism for the fully developed energy algorithm is the rescaling of the temperature profile from the interface boundary at the outflow to use at the inflow interface boundary. This mechanism is different according to whether the constant temperature wall assumption or constant heat flux assumption is being used. For the constant temperature wall assumption, the scaling is:

$$\frac{T_w - T}{T_w - T_b} = \text{const} \tag{382}$$

so that the wall temperature and the bulk temperature must be specified.

For the constant heat flux assumption, the scaling law is:

$$T - T_b = \text{const} \tag{383}$$

so that bulk temperature must be specified.

In Which Situations Will Errors Be Incurred?

Errors occur with mesh issues, compressibility effects, property variations, and non-contiguous regions.

Mesh Issues

In addition to regular mesh refinement considerations, particularly in the wall-normal direction where large gradients occur, care must be taken to resolve the flow in the streamwise direction. Unlike simple fully developed duct flows without energy, which can often use as few as 3 - 5 cells in the streamwise direction, a larger number of cells is required when energy is considered. This requirement is true of the constant wall temperature heating condition, since the fluid heats up rapidly as the flow enters the domain. Grid refinement studies are recommended to confirm that the grid resolution is adequate.

Compressibility Effects

In truly compressible flows, viscous dissipation causes the fluid to heat up. With adiabatic walls, this heating effect causes the flow to accelerate. This scenario is called “fanno flow.” The fully developed energy treatment in STAR-CCM+ is not designed for fanno flows or compressible duct flows with other heating situations. However, use of the ideal gas law is not precluded and if employed, the Mach number must be sufficiently low that these effects do not dominate and cause stability problems.

Property Variations

The existence of a similarity solution for energy in a fully developed duct flow is most easily demonstrated for constant property flows. However, the use of temperature-dependent fluid properties is not precluded, and small property variations do not have a large effect on the solution. However, caution must be exercised when using variable properties in case they affect the accuracy or stability of the solution.

Non-Contiguous Regions

STAR-CCM+ permits regions to be non-contiguous, in general. However, the fully developed energy algorithm must search through all regions that are connected to the interface to evaluate contributions from associated cells

and boundaries. For this reason, it is recommended that non-contiguous regions are not used. If in any doubt, use the Split Non-Contiguous function on the region.

What About Variable Heating Situations?

The similarity solutions that cause a fully developed energy situation in a duct are limited to:

- Constant wall heat flux
- Constant wall temperature

However, given the flexibility of the boundary conditions in STAR-CCM+, no attempt is made to ensure that the values specified for wall heat flux or temperature are indeed constant. If a boundary condition is specified that deviates from the profile scaling that is used at the fully developed interface, an error is incurred. The significance of the error clearly depends on how incorrect the assumption is.

The specified wall heat flux situation is likely to be more forgiving than the specified wall temperature situation because the linear behavior of the energy equation permits superposition of solutions. Therefore a wall heat flux that varies periodically (such as a sinusoidal wave) most likely produces a well-behaved solution.

How Can One Verify That Fully Developed Energy Is Valid for a Specific Heating Situation?

The similarity solutions that cause a fully developed energy situation in a duct are limited to:

- Constant wall heat flux
- Constant wall temperature in **constant property flows**

Any deviations from these situations cast doubt on the validity of the approach. One way to check that the fully developed approach is valid is to create a simplified version of the domain and manually replicate it 20 or more times. Then, verify that the velocity and thermal fields are indeed periodic (using the appropriate scaling). This check requires a significant effort, but is necessary to remove residual doubt.

What Happens If Recirculation Occurs at the Fully Developed Interface?

Care has been taken to test the fully developed interface for an interface at which recirculating flow occurs. However, problems can arise if the portion of flow entering approaches the portion of the flow leaving. Therefore, if possible, it is preferable to place the interface at a location where recirculation does not occur.

Modeling Boiling

Some flows are characterized by subcooled nucleate boiling on solid surfaces, where the creation of the vapor bubbles is insignificant for the flow. In such situations, if you are interested in modeling the improved heat transfer due to boiling, you can add boiling to a STAR-CCM+ simulation with the flow of a single-component liquid.

This section describes the following:

- Selecting a boiling model
- Setting boiling model parameters

For information about the formulation that is implemented in STAR-CCM+, see Single-Phase Boiling Formulation.

Selecting a Boiling Model

To select the boiling model that is based on Rohsenow's expression for the nucleate boiling:

- In the *Physics Model Selection* dialog, specify the material model as a single-component liquid by selecting Liquid in the *Material* group.

- Make appropriate selections of flow and energy models. The Boiling model becomes available in the *Optional Models* box.



- When you select Boiling the *Physics Model Selection* dialog shows the model as activated, with other model options remaining to be selected.



Setting Boiling Model Parameters

When a simulation with boiling is selected, STAR-CCM+ creates a node that is called Boiling in the simulation tree, as shown in the following screenshot.



In the *Properties* window, specify the boiling model parameters:

Boiling Properties

Latent heat	The latent heat of evaporation.
Vapor density	The density of vapor.
Surface tension	The surface tension coefficient between the liquid and its vapor.
C_qw	The empirical coefficient C_{qw} in the Rohsenow expression for the wall heat flux, Eqn. (390). This value varies with the liquid-surface combination.
n_p	The Prandtl number exponent n_p (1.73 by default) in Eqn. (390).

Modeling Moisture Boiling in Solids

When a wet solid substance is heated above the boiling temperature of the moisture, it can become a significant heat sink due to the latent heat of vaporization. The Moisture Boiling model in STAR-CCM+ simulates this heat transfer effect.



This model computes the boiling of a liquid (moisture) within a solid material. The model considers only energy transport -- it does not simulate the evaporated gas phase and distribution. However, this model does calculate distribution and volume fraction of moisture within the solid.

The moisture in the solid does not influence the material properties of the solid such as conductivity or density. The boiling process influences the temperature distribution within the solid. (Boiling is an isothermal process, so the temperature remains constant.)

Moisture Boiling Properties

Density	The density of the moisture.
Latent heat	The latent heat of vaporization of the moisture.
Boiling Temperature	The boiling temperature of the moisture.

See also:

Setting up a Moisture Boiling Simulation

Moisture Boiling Field Functions Reference

Setting up a Moisture Boiling Simulation

To use moisture boiling in a STAR-CCM+ simulation:

1. Select the Moisture Boiling model:
   a. Include the following models with your selections:
      - The Solid material model.
      - The Segregated Solid Energy model.

b. Select the Moisture Boiling model from the *Optional Models* group box.



When you close the model selection dialog, the Moisture Boiling model is shown in the list under the Continua > Physics > Models node.



2. Specify the initial moisture volume fraction as an initial condition.



3. Set the static temperature which is also an initial condition.
4. Set the boiling temperature in the Boiling Temperature property of the Moisture Boiling model node.

   The temperature values have the following effects:

   - When the initial temperature is greater than the boiling temperature, the phase change (boiling) of moisture results in a temperature decrease down to boiling temperature. This process requires the volume fraction of moisture to exceed zero.

- When the initial temperature is less than boiling temperature, and if the cell temperature increases due to a net heat flux into the cell, temperature first increases up to the boiling temperature. Heat is consumed to vaporize moisture isothermally at the boiling temperature. When the moisture volume fraction reaches zero, the temperature increases further.

Moisture Boiling Field Functions Reference

The following primitive field function is made available to the simulation when the Moisture Boiling model is used. It has a typical set of base properties.

Moisture Volume Fraction	The fraction of volume that the moisture represents.

Modeling Thin Film

This section covers the following:

- What is the Thin Film model?
- What is the Deicing model?
- What is the Defogging model?
- Selecting the Thin Film model
- What is the Deicing solver?
- What is the Defogging solver?

Also included is the following reference information:

- Boundary conditions
- Boundary values and initial conditions
- Field functions that this model activates

What Is the Thin Film Model?

The Thin Film model provides a simplified approach for simulating deicing or defogging inside a cabin windshield or other surface. The object has no properties of its own, but is subject to common menu items.



The limitations on this model are as follows:

Thin Film Model

- Once a boundary is switched from deactivated to having a fog or ice layer, the simulation must be run in transient mode.
- It is possible to have Deicing and Defogging in the same region, but not on the same boundary.

Deicing Model

- You cannot use a contact boundary. Deicing can only be applied to wall boundaries and the type must be convection or heat flux.

Defogging Model

- Defogging can only be applied to wall and contact interfaces (solid/fluid).
- Currently the saturation vapor pressure curve is based on water vapor.

What Is the Deicing Model?

The Deicing model, which is made available by the Thin Film model, is a simplified approach to obtain deicing on a windshield or other surface. The model object has an expert property and is subject to common menu items.



This object also contains the Ice material, which lets you work with material properties for this film.

Major hypotheses are:

- The ice layer is not represented as a phase in the volume mesh, but is modeled using functions, thus the ice layer does not flow.
- The model formulation implicitly assumes that the temperature surrounding the ice is above the freezing temperature. If this assumption is not correct, the model could produce some non-physical values (for example, the ice could start melting at temperatures which are below the freezing temperature).
- A steady-state calculation is performed to initialize the convection flow field.
- The best method is convection, which you would select for the Thermal Specification boundary condition.
- During the transient run, it is recommended that you leave all equations on. However, you can use a large time-step because the melting of the ice layer does not affect the flow field.



Deicing Model Properties

None

Deicing Model Expert Properties

Cemp	Specifies the C_{emp} empirical coefficient in Eqn. (401). This coefficient specifies the proportion of heat flux that is absorbed at the ice-gas interface.

What Is the Ice Material?

This object has a sub-node that works like the material properties manager: it contains material properties objects that are specific to this type of thin film.



Ice Material Properties

None

See also: common menu items.

What Is the Defogging Model?

The Defogging model, which is made available by the Thin Film model, is a simplified approach to obtaining a defogging simulation inside a cabin windshield or surface. The model object can be controlled with properties and is subject to common menu items.



This object also contains the film material, which lets you work with material properties for this film.

Major hypotheses are:

- The fog layer is not represented as a phase in the inner volume mesh, but is modeled using functions, thus the fog layer does not flow.
- A steady state calculation is performed to initialize the convection flow field.

- During the transient run, it is recommended that you leave all equations on. However, you can use a large time-step because the condensation/evaporation of the fog layer does not affect the flow field.



Defogging Model Properties

LimitSuperSaturation	Controls whether supersaturation (humidity exceeding 100%) is prevented.	
	Activated	Prevents supersaturation.
	Deactivated	Allows supersaturation.

Defogging Model Expert Properties

The settings within the *Expert* category can be useful if you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+. The values in that category reflect both the design of the model and the discretization approaches that have been optimized for accuracy and performance. If you are not familiar with these details, do not change any settings.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.

Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Convection — Selects the convection scheme that is used. For guidance as to convection scheme selection, see Diffusion Term, p. 6349.

1st-order	Selects the first-order upwind convection scheme.
2nd-order	Selects the second-order upwind convection scheme.

Cemp — Specifies the C empirical coefficient.

What Is the Film Material?

This object has a sub-node that works like the material properties manager: it contains material properties objects that are specific to this type of thin film.



Film Material Properties

None

See also: common menu items.

Selecting the Thin Film Model

To select the Thin Film model:

- In the *Physics Model Selection* dialog, specify a material model in the *Material* group.
- Make appropriate selections of flow or solid models, and an energy model, depending on whether you are using the Deicing or Defogging model. If you are using Deicing, you can apply it to a fluid or a solid region. However, defogging requires a flow model for a fluid region.

The Thin Film model becomes available in the *Optional Models* group.



Once you activate the Thin Film model, the Deicing and Defogging models become available in the *Thin Film* group.



Select one of these models, depending on the type of film you are simulating.

What Is the Deicing Solver?

This solver, which is subject to common menu items, controls the solution of the Deicing model in all the continua for which the model is activated.



Deicing Solver Properties

None

Deicing Solver Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

What Is the Defogging Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the Defogging model in all the continua for which the model is activated.



Defogging Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the old solution is supplanted with the newly computed solution. The default value is 0.9.

Defogging Solver Expert Properties

The properties in the *Expert* category are designed for use in a temporary "debug" situation, at the expense of simulation accuracy and extra memory usage. It is not recommended that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Temporary Storage Retained

Retains or deletes temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

- Activated: Retain temporary storage.
- Deactivated: Do not retain temporary storage.

Reconstruction Frozen

Specifies whether reconstruction gradients are updated with each iteration.

- Activated: Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
- Deactivated: Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

- Activated: Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you deactivate the check-box again.
- Deactivated: Do not reset reconstruction gradients. If the gradients have been zeroed, they are recomputed at the next iteration.

Boundary Conditions Reference

This section details boundary conditions that are available for the Thin Film model:

- Thin Film Option
- Vapor Mass Fraction Option
- Vapor Mass Fraction Specification

Thin Film Option

This node allows you to determine how the thin film across a wall boundary is specified. It is activated when the Thin Film model is activated.

Properties

Option	Activates defogging or deicing for the boundary.	
	Disable	The option is turned off for the boundary. No nodes are added to the Physics Values node.
	Fog Layer	This option, which is available when the Defogging model is activated, adds a Fog Layer Thickness node to the Values node, which must be entered as a scalar profile.
	Ice Layer	This option, which is available when the Deicing model is activated, adds an Ice Layer Thickness node to the Values node, which must be entered as a scalar profile.

Vapor Mass Fraction Option

This node, which is activated when the Defogging model is activated, allows you to determine how the defogging across a wall boundary is specified.

Properties

Wall Vapor Mass Fraction Scalar	Activates a fog layer at the boundary.	
	Zero Flux	No fog layer is present. No nodes are added to the Physics Values node.
	Film	Specifies saturation conditions are specified for the scalar when a fog layer is present.

Vapor Mass Fraction Specification

This node, which is activated when the Defogging model is activated, allows you to determine how the defogging is specified across the following types of boundaries:

- Free stream
- Mass flow inlet

- Velocity inlet
- Pressure outlet
- Stagnation inlet

This object is also an initial condition.

Properties

Method	The method that is used to define the defogging at the boundary.	
	Vapor Mass Fraction	This option adds a Vapor Mass Fraction node to the Values node, which must be entered as a scalar profile.
	Vapor Humidity	This option adds a Vapor Humidity node to the Values node, which must be entered as a scalar profile.

Boundary Values and Initial Conditions Reference

This section contains reference descriptions of boundary and initial condition values that are available for the Thin Film model:

- Fog Layer Thickness
- Ice Layer Thickness
- Vapor Humidity
- Vapor Mass Fraction

The boundary values are set according to the boundary type, and by methods that are selected in corresponding nodes in the Conditions node.

Fog Layer Thickness

This value node is available when the Defogging model is selected. It is visible if `Fog Layer` is selected as the Option property for the Thin Film Option node of the Conditions node. It is entered as a scalar profile.

Ice Layer Thickness

This value node is available when the Deicing model is selected. It is visible if `Ice Layer` is selected as the Option property for the Thin Film Option node of the Conditions node. It is entered as a scalar profile.

Vapor Humidity

This value node is available when the Defogging model is selected. It is visible if `Vapor Humidity` is selected as the Option property for the Vapor Mass Fraction Specification node of the Conditions node. It is entered as a scalar profile.

Vapor Mass Fraction

This value node is available when the Defogging model is selected. It is visible if `Vapor Mass Fraction` is selected as the Option property for the Vapor Mass Fraction Specification node of the Conditions node. It is entered as a scalar profile.

Thin Film Field Functions Reference

The following field functions are made available to the simulation when the Thin Film model is used with the Defogging model.

Film Density	Density of liquid film [kg/m^3].
Film Diffusivity	The diffusion of vapor in air when the liquid film evaporates [m^2/s].
FilmLatentHeat	Latent heat due to condensation/evaporation of the film vapor [J/kg].
Fog Layer Thickness Value	Thickness of the film [m].
Vapor Humidity Value	Partial pressure of water vapor in a mixture to the saturated vapor pressure of the water at the mixture temperature [%].
Vapor Mass Fraction Diffusion Coefficient	Diffusion of vapor mass fraction in air [m^2/s].
Vapor Mass Fraction Value	Ratio of the vapor mass fraction to the total mass fraction for each cell [mv/mt].

The following field functions are made available to the simulation when the Thin Film model is used with the Deicing model.

Ice Density	Density of the ice layer [kg/m^3].
Ice Latent Heat	Latent heat due to melting/solidification of the ice layer $[J/kg]$.
Ice Layer Thickness Value	Thickness of the ice layer [m].
Ice Specific Heat	Specific heat of the ice layer $[J/kgK]$.

Modeling Heat Exchangers

Heat Exchangers in STAR-CCM+ model the heat transfer between two fluid streams—a hot and a cold stream (such as in radiators and air coolers). This section describes how to use the single stream and dual stream heat exchanger options in STAR-CCM+.

In the single stream option, only one of the fluid streams (usually the cold stream) is modeled explicitly, while the second stream is assumed to have a uniform temperature that you specify. In the dual stream option, both the streams are modeled explicitly by defining two overlapping regions and creating a heat exchanger interface between them. In both the options, heat is added or removed from the participating cells as a source or sink term in the corresponding fluid energy equation. It is assumed in this section that both the streams exist only in a single phase (liquid or gas) and no phase change occurs inside the heat exchangers.

Before choosing any heat exchanger option, an energy model must have been selected as part of the physics continuum of the heat exchanger region.

Using the Single Stream Heat Exchanger Option

In a single stream heat exchanger, one stream is assumed to have a uniform temperature and the other stream is modeled by specifying the heat exchanger enthalpy source. This section defines what is meant by the heat exchanger enthalpy source in STAR-CCM+, and describes how to activate this type of heat exchanger and specify its properties. The formulation of this heat exchanger is also described.

What Is the Heat Exchanger Enthalpy Source?

The heat exchanger enthalpy source is a device which introduces or removes a specified amount of heat flow in a region. Local temperature with respect to a user-specified reference temperature establishes the distribution of this set amount of heat source or sink through the space of the region. In this way, the distribution of heat within the fluid is similar to what would be seen with a heat exchanger. It is associated with one fluid stream only and is a net source or sink with the model instead of a transfer from one continuum to another.

This device is a simple way to incorporate a known heat load into (or from) a model with a minimum of effort and a minimal computational burden. Sample applications include air conditioner evaporators, condensers, charge air coolers, radiators, electric heaters, and electronic devices.

The method is relatively insensitive to mesh size and utilizes resolution corresponding the grid fineness.

The Energy Source Option activates the heat exchanger enthalpy source.

Activating the Heat Exchanger Enthalpy Source

Before you can follow this procedure, you must have selected an energy model.

To activate the heat exchanger enthalpy source, open the Regions node, the individual region node, and then the Physics Conditions node. Select the Energy Source Option node.



Select `Heat Exchanger` in the drop-down list corresponding to the Energy Source Option property in the *Properties* window. This selection allows you to set heat exchanger properties.



Setting Heat Exchanger Enthalpy Source Properties

You control the heat exchanger enthalpy source by setting physics values along with a physics condition of a region. These objects are activated when you select `Heat Exchanger` as the Energy Source Option of the region.



The values include:

- Heat exchanger minimum temperature difference
- Heat exchanger boundaries
- Heat exchanger first iteration
- Heat exchanger temperature

The region condition that gets added, Heat Exchanger Exit Temperature Specification, lets you choose between specifying the heat transfer rate Q and predicting the outlet temperature T_{out}, or vice versa. When this condition is left at `Inferred`, you specify the heat transfer rate by accessing the Heat Exchanger Energy Transfer physics value node.

However, you can change the setting of the Outlet Temperature property of the Heat Exchanger Exit Temperature Specification to `Specified`.



This selection activates the Target Exit Temperature node.



By accessing this object, you can specify the temperature at the outlet while letting STAR-CCM+ compute the overall heat transfer rate.

Working with Heat Exchanger Energy Transfer

This region value node represents the total energy transfer for the heat exchanger enthalpy source.



Properties

Value	The fixed rate of heat transfer, Q_{total}, from and/or to the device, typically measured in Watts.

Working with Heat Exchanger Minimum Temperature Difference

This node represents the minimum allowable temperature difference between the following:

- The specified heat exchanger temperature and the maximum fluid temperature for a heater (that is, when the specified Heat Exchanger Energy Transfer is positive)

- The specified heat exchanger temperature and the minimum fluid temperature for a cooler (that is, when the specified Heat Exchanger Energy Transfer is negative).



For a heater, the hot stream is assumed to have an infinite thermal capacitance, and the specified heat exchanger temperature T_h refers to the constant temperature of the hot stream. Now, given the heat exchanger energy transfer rate Q, the inlet cold stream temperature T_{ci}, and the cold stream specific heat c_{pc}, an energy balance could be used to predict the cold stream outlet temperature T_{co} as follows:

$$T_{co} = T_{ci} + \frac{Q}{\dot{m}_c c_{pc}} \tag{384}$$

where $\dot{m}_c$ is the mass flow rate of the cold stream. Physically, it is impossible to have $T_h < T_{co}$. However, if the heat exchanger temperature value is erroneously specified below the maximum cold stream temperature in the heat exchanger, the specified heat exchanger minimum temperature difference is used to correct the heat exchanger temperature. The heat exchanger temperature stays above the predicted maximum cold stream temperature. When this error occurs, a warning message appears on the console, which is an indication that the heat exchanger temperature specification is determined to be unphysical.

Similarly, for a cooler, the cold stream is assumed to have an infinite thermal capacitance, and the specified heat exchanger temperature T_c refers to the constant temperature of the cold stream. Now, given the heat exchanger energy transfer rate Q, the inlet hot stream temperature T_{hi}, and the hot stream specific heat c_{ph}, an energy balance could be used to predict the hot stream outlet temperature T_{ho} as follows:

$$T_{ho} = T_{hi} - \frac{abs(Q)}{\dot{m}_h c_{ph}} \tag{385}$$

where $\dot{m}_h$ is the mass flow rate of the hot stream. Physically, it is impossible to have $T_c > T_{ho}$. However, if the heat exchanger temperature value is erroneously specified above the minimum hot stream temperature in the heat exchanger, the specified heat exchanger minimum temperature difference is used to correct the heat exchanger temperature. The heat exchanger temperature stays below the predicted minimum hot stream temperature. When this error occurs, a warning message appears on the console, which is an indication that the heat exchanger temperature specification is determined to be unphysical.

Properties

Value	The minimum temperature difference between the reference heat exchanger temperature and the maximum/minimum fluid temperature for a heater/cooler respectively.

Working with Heat Exchanger Boundaries

The properties of this region value node are used to specify the upstream and downstream interface boundaries that the Heat Exchanger (Single Stream) report uses.



This node assumes that the heat exchanger region interfaces with two other regions as shown in the example below. You specify the boundaries belonging to these interfaces in the property page of this node.



Properties

Upstream Interface Boundary	The interface boundary through which fluid enters the heat exchanger region. Selected in the object selector from the list of interface boundaries for the region.
Downstream Interface Boundary	The interface boundary through which fluid leaves the heat exchanger region. Selected in the object selector from the list of interface boundaries for the region.

Working with Heat Exchanger First Iteration

This node specifies the iteration from which the heat exchanger enthalpy source is active and begins its contribution to the source term of the energy equation.



Ideally, activate the heat exchanger source term when the flow solution is beginning to conform to its final pattern (at least in terms of direction). Conformity could typically be achieved within the first 50 iterations of the flow solver for a moderately complex geometry.

Properties

First Heat Exchanger Iteration	The iteration number at which the heat exchanger enthalpy source is activated.

Working with Heat Exchanger Temperature

This region value node represents the reference temperature T_{ref} of the heat exchanger enthalpy source. It is entered as a scalar profile.



The local fluid temperature at each cell is subtracted from the reference temperature T_{ref} to determine the relative amount of heat that is transferred from and/or to the cell. See Eqn. (411).

For pop-up menu information, see common menu items.

Working with Target Exit Temperature

The Target Exit Temperature node, which has properties, lets you specify the temperature at the outlet while allowing STAR-CCM+ to compute the overall heat transfer rate. This node becomes available when the Outlet Temperature property of the Heat Exchanger Exit Temperature Specification region condition node is set to `Specified`.



Properties

Target Outlet Temperature	Target outlet temperature of the fluid leaving the heat exchanger.
Heat Transfer Rate	The overall heat transfer rate that STAR-CCM+ computes (read-only).

Expert Properties

Under-Relaxation Factor	This parameter is specifically for predicting the overall heat transfer rate. At each iteration, it governs the extent to which the new calculation supplants the old one. The default value is 0.1.

Using the Dual Stream Heat Exchanger Option

Two versions of the dual stream heat exchanger option exist in STAR-CCM+: basic and actual.

This section describes the heat exchanger topology necessary to simulate the dual stream heat exchanger in STAR-CCM+, how to create the heat exchanger topology, and define a heat exchanger interface between two regions.

The dual stream heat exchanger model supports the following:

- A fluid-solid interface, where the heat exchange occurs between a solid region and a fluid region.
- A fluid-fluid interface, where the heat exchange takes place between two fluid regions.

For a fluid-fluid type heat exchanger, the following options are available for specifying the heat transfer rate in the heat exchanger:

- UAL Polynomial
- UAL Table
- UAG Table
- Q Table
- Q Map
- UAL Map

An option is included to specify the target heat rejection rate to enable predicting the hot flow inlet boundary temperature.

For a fluid-solid type heat exchanger, the following option is made available:

- Heat Exchanger UAL

The formulation for this option describes the above quantities and discusses the implementation aspects.

What Are the Basic and Actual Dual Stream Heat Exchanger Options?

The basic and actual dual stream heat exchanger options use different approaches for the solution:

- The basic option solves for a simpler version of the energy equation, Eqn. (244). The simpler version only uses the convective term and the source term to yield a solution independent of the solution that is generated for the overall simulation. This separate solution consists of pseudo-temperatures for the hot and cold regions of the heat exchanger, which are used to compute the mean temperature difference and hence the local heat transfer coefficient, UAL.

- In order to compute the mean temperature difference and hence the UAL, the actual option uses the flow and energy values that are calculated for the simulation itself.

The basic option uses a wider range of properties settings than the actual. Specifics are given in the descriptions of affected objects.

When using the basic option, ensure that the core mesh is constructed of uniform cells. That is, all the cells in the heat exchanger region must be of comparable cell volume, as the cell volumes do not enter the calculation of the UAL table in the basic model. This restriction does not apply when using the actual option.

Note that UA refers to the product of the Overall heat transfer coefficient (U) and the Area of the heat exchanger (A), specified in W/K. UAL refers to the local (or cell value of) heat transfer coefficient (in W/K), with L referring to local.

Selecting the Basic or Actual Option

This step is part of defining a heat exchanger interface. To select one of these options, open the node of the heat exchanger interface and select the Physics Conditions > Heat Exchanger Method node. Select `Basic Dual Stream` or `Actual Dual Stream` from the drop-down list of the Option property.



What Is the Heat Exchanger Topology?

Consider the heat exchange between a cold fluid stream and a hot fluid stream. These streams are modeled as two different physics continua (each having different material properties), as shown below:



The figure below describes the topology that is required to simulate the dual stream heat exchanger in STAR-CCM+ (A 2D example is shown for clarity):



The heat exchanger topology that is shown above comprises six regions - cold-in, cold-core, cold-out, hot-in, hot-core, and hot-out. The regions cold-in, cold-core and cold-out represent the cold stream. Their physics continuum is selected as the Cold Continuum.



The regions hot-in, hot-core and hot-out represent the hot stream. Their physics continuum that is set as the Hot Continuum.



Define internal (or in-place) boundary interfaces between the common boundaries of cold-in and cold-core, cold-core and cold-out, hot-in and hot-core, and hot-core and hot-out. The actual heat exchange takes place between the

two regions cold-core and hot-core. Both regions must have an identical mesh (that is, they must be created from the same mesh continuum and must have the same number of cells and cell topology between them). Define a heat exchanger interface between them.

For maximum accuracy when using the basic dual stream heat exchanger, ensure that the mesh used for the core region is constructed of uniform cells. The accuracy of results that the actual dual stream heat exchanger predicts does not depend on the uniformity of the mesh.

Creating the Heat Exchanger Topology

This section explains how to create the heat exchanger topology, starting with the following set of regions:



The figure above shows five regions - cold-in, cold-core, cold-out, hot-in and hot-out. Ensure that the regions cold-in, cold-out and cold-core have their physics continuum set as Cold Continuum and the regions hot-in and hot-out have their physics continuum set as Hot Continuum. Copy the cold-core region as shown in the following screenshot:



To ensure that both the regions are created from the same mesh continuum, copy the region after the volume mesh has been generated.

Paste the copied region into the Regions node as shown below:



Rename the new region as hot-core and change its physics continuum to Hot Continuum. Now, the simulation has six regions - cold-in, cold-core, cold-out, hot-in, hot-core, hot-out. Define internal (or in-place) boundary interfaces between the common boundaries of cold-in and cold-core, cold-core and cold-out, hot-in and hot-core, and hot-core and hot-out. The creation of the heat exchanger topology is now complete.

You can also specify multiple inlets and outlets for a heat exchanger. The figure below shows a heat exchanger with a single inlet and a single outlet.



The next figure shows the heat exchanger with three inlets and two outlets.



Defining a Heat Exchanger Interface

To define a heat exchanger interface, a heat exchanger topology must exist. The topology must contain two identical regions with a different continuum associated with each region.

Once the two overlapping regions exist, the heat exchanger interface can then be created by selecting the two regions (holding down the CTRL key) and choosing Create Interface from the right-click menu. While selecting the two regions, select the region representing the cold stream before the region representing the hot stream.



By default, the Heat Exchanger Method stays Inactive, which indicates that the source terms of the energy equation are not affected. When the Basic Dual Stream option is chosen as the Heat Exchanger Method, a new physics condition, Heat Exchanger Data Specification appears. This node allows you to choose an input option.



In order to identify the one-to-one correspondence between the cells belonging to the hot and cold regions, the “Direct Region Interface Intersection” allows you to change the tolerance between the cell centroids.

What Is the UAL Polynomial?

This polynomial option assumes that UAL is known as a function of the cold flow velocity v.



The known polynomial can be specified by selecting the UAL Polynomial node. The following example illustrates the specification of the function $UAL = 0.5 + 0.03v^2 + 0.005v^3$ defined within the range $v = 2\frac{m}{s}$ to $v = 50\frac{m}{s}$



When you use the actual dual stream heat exchanger option, two additional (expert) properties become available:

- Hot Mass Flow Rate Specified—When activated, this property causes the hot stream to use the mass flow rate that the table specifies.
- UAL Frozen—When activated, this property freezes the UAL table in its current state.

What Is the UAL Table?

This option assumes that the local heat transfer coefficient, UAL, is known as a table of the cold flow velocity. For a known local cell velocity of the cold stream, the UAL is determined by interpolating from this table. The UAL table must contain two columns of data representing the cold stream velocity and UAL. The table must be arranged in an ascending order of the velocities.

The UAL Table option is selected as shown below.



The UAL table can be imported into the simulation through the Tools node.



Once imported, the table can be selected into the UAL Table option.



When you use the actual dual stream heat exchanger option, two additional (expert) properties become available:

- Hot Mass Flow Rate Specified—When activated, this property causes the hot stream to use the mass flow rate that the table specifies.
- UAL Frozen—When activated, this property freezes the UAL table in its current state.

What Is the UAG Table?

This option assumes that the overall heat transfer coefficient (times the heat exchanger area) across the entire heat exchanger is known as a table of the cold stream mass flow rate (say, from an experimental source). You also specify the inlet temperatures of the cold and hot streams under test conditions, from which the overall heat transfer across the heat exchanger can be determined.

For a basic dual stream exchanger, the local UAL is computed in the following manner: a linear system of algebraic equations is solved iteratively by considering the heat balance between the cold and the hot

streams, preserving the target heat exchange rate, for each row of the UAG table (that is, for each value of the cold stream mass flow rate). For each row of the UAG table:

$$UAL = \frac{UAG \cdot (T_{actual,\,hot} - T_{actual,\,cold})}{net\Delta T}$$

where $T_{actual,\,hot}$ and $T_{actual,\,cold}$ are the area averaged hot and cold inlet temperatures that are taken from the actual simulation, and:

$$net\Delta T = \frac{\sum_i (T_{hot,\,i} - T_{cold,\,i}) \cdot Vol_i}{\frac{1}{NC} \cdot \sum_i Vol_i}$$

where Vol_i is the cell volume and NC is the number of cells in the heat exchanger region. The summations are across all cells of the heat exchanger.

For an actual dual stream heat exchanger, the local hot and cold temperatures are extracted from the simulation, obviating the need to solve any linear system of equations.

This UAL table is computed internally only once initially (or whenever any input for the UAG table is changed). The table is then used to interpolate the UAL for a given cold flow local cell velocity in the energy equation.

The `UAG Table` option is selected as shown below.



UAG refers to global (or) overall heat transfer coefficient for the entire heat exchanger, with the G referring to global.

Properties

UAG Table: Heat Transfer Units	Default heat transfer units that the UAG table uses.
UAG Table	The table containing the columns cold mass flow rate and UAG, which must be arranged in increasing order of the mass flow rate.
UAG Table: Cold Mass Flow Rate	The column within the UAG table representing the cold stream mass flow rate.
UAG Table: UAG	The column within the UAG table representing the overall heat transfer rate.
Cold Fluid Density	Test condition density of the cold fluid (basic dual stream heat exchanger option only).
Hot Fluid Mass Flow Rate	Test condition mass flow rate of the hot stream (basic dual stream heat exchanger option only).
Cold Fluid Inlet Temperature	Test condition inlet temperature of the cold stream (basic dual stream heat exchanger option only).
Hot Fluid Inlet Temperature	Test condition inlet temperature of the hot stream (basic dual stream heat exchanger option only).
Cold Fluid Specific Heat	Test condition specific heat of the cold stream (basic dual stream heat exchanger option only).
Hot Fluid Specific Heat	Test condition specific heat of the hot stream (basic dual stream heat exchanger option only).
Cold Upstream Boundaries	Upstream boundary interfaces of the cold stream.
Hot Upstream Boundaries	Upstream boundary interfaces of the hot stream.

Expert Properties

Maximum Solver Iterations	The maximum number of iterations permissible for solving the linear system of equations to compute UAL (basic dual stream heat exchanger option only).
Convergence Tolerance	Tolerance on converging the solver to the target heat transfer rate (basic dual stream heat exchanger option only).

Hot Mass Flow Rate Specified	When activated, causes the hot stream to use the mass flow rate that the table specifies (actual dual stream heat exchanger option only).
UAL Frozen	When activated, freezes the UAL table in its current state (actual dual stream heat exchanger option only).

What Is the Q Table?

This option assumes that the overall heat transfer Q across the heat exchanger is known as a table of the cold stream mass flow rate (say, from an experimental source) for one value of hot stream mass flow rate.

In order to compute the local UAL, a UAG table is first computed from Q using the expression:

$$UAG = \frac{Q}{T_{h,\,in} - T_{c,\,in}}$$

where $T_{hot,\,in}$ and $T_{cold,\,in}$ are the test rig hot and cold temperatures that are specified as properties of the Q Table.

For a basic dual stream heat exchanger, a linear system of algebraic equations is solved iteratively by considering the heat balance between the cold and the hot streams, preserving the target heat exchange rate, for each row of the UAG table (that is, for each value of the cold stream mass flow rate) as described in the UAG table section.

For an actual dual stream heat exchanger, the local hot and cold temperatures are extracted from the simulation, obviating the need to solve any linear system of equations.

This UAL table is computed internally only once initially (or whenever any input for the Q table is changed). It is then used to interpolate the UAL for a given cold flow local cell velocity in the energy equation.

The Q Table is selected as shown below.



Properties

Q Table	The table containing the columns cold mass flow rate and overall heat transfer, which must be arranged in increasing order of the mass flow rate.
Q Table: Q	The column within the Q table representing the overall heat transfer rate.
Q Table: Heat Transfer Units	Default heat transfer units used that the Q table uses.
Q Table: Cold Mass Flow Rate	The column within the Q table representing the cold stream mass flow rate.
Hot Fluid Mass Flow Rate	Test condition mass flow rate of the hot stream (basic dual stream heat exchanger option only).
Cold Fluid Inlet Temperature	Test condition inlet temperature of the cold stream (basic dual stream heat exchanger option only).
Hot Fluid Inlet Temperature	Test condition inlet temperature of the hot stream (basic dual stream heat exchanger option only).
Cold Fluid Specific Heat	Test condition specific heat of the cold stream (basic dual stream heat exchanger option only).
Hot Fluid Specific Heat	Test condition specific heat of the hot stream (basic dual stream heat exchanger option only).
Cold Flow Density	Test condition density of the cold fluid (basic dual stream heat exchanger option only).
Cold Upstream Boundaries	Upstream boundary interfaces of the cold stream.
Hot Upstream Boundaries	Upstream boundary interfaces of the hot stream.

Expert Properties

Maximum Solver Iterations	The maximum number of iterations permissible for solving the linear system of equations to compute UAL (basic dual stream heat exchanger option only).
Convergence Tolerance	Tolerance on converging the solver to the target heat transfer rate (basic dual stream heat exchanger option only).

Hot Mass Flow Rate Specified	When activated, causes the hot stream to use the mass flow rate that is specified in the table (actual dual stream heat exchanger option only).
UAL Frozen	When activated, freezes the UAL table in its current state (actual dual stream heat exchanger option only).

What Is the Q Map?

A Q Map provides STAR-CCM+ with a series of curves from which it calculates the local heat transfer coefficient, UAL. One set of curves on the map provides UAL values for constant cold stream mass flow rates m_{cold} and varying hot stream mass flow rates m_{hot}. These curves are intersected by other curves that provide UAL values for varying m_{cold} with constant m_{hot}.

The Q Map option may be selected as shown below:



The interpolation method for computing the local UAL is different for the Basic Dual Stream and the Actual Flow Dual Stream heat exchanger methods.

Q Map (Basic Dual Stream) Properties

Q Table: Heat Transfer Units	Default heat transfer units that the Q table uses.
Q Table	The table containing the columns cold mass flow rate, hot mass flow rate, and overall heat transfer.
Q Table: Q	The column within the Q map representing the overall heat transfer rate.
Q Table: Cold Mass Flow Rate	The column within the Q map representing the cold stream mass flow rate.
Q Table: Hot Mass Flow Rate	The column within the Q map representing the hot stream mass flow rate.
Cold Fluid Density	Test condition density of the cold fluid.
Hot Fluid Density	Test condition density of the hot fluid.
Cold Fluid Inlet Temperature	Test condition inlet temperature of the cold stream.
Hot Fluid Inlet Temperature	Test condition inlet temperature of the hot stream.
Cold Fluid Specific Heat	Test condition specific heat of the cold stream.
Hot Fluid Specific Heat	Test condition specific heat of the hot stream.
Cold Upstream Boundaries	Upstream boundary interfaces of the cold stream.
Hot Upstream Boundaries	Upstream boundary interfaces of the hot stream.

Q Map (Basic Dual Stream) Expert Properties

Cold Flow Exponent	Exponent ec in Eqn. (419).
Hot Flow Exponent	Exponent eh in Eqn. (419).
Goodness of fit	A read-only value ranging between 0 and 1 that indicates how well the fitted curve represents the data. Typically, this value can be expected to be within 0.96 and 0.98 for a reasonable fit.
Coefficient C0	Coefficient C_0 in Eqn. (419). This computed value is read-only.
Coefficient C1	Coefficient C_1 in Eqn. (419). This computed value is read-only.

Coefficient C2	Coefficient C_2 in Eqn. (419). This computed value is read-only.
Maximum Solver Iterations	The maximum number of iterations permissible for solving the linear system of equations to compute UAL.
Convergence Tolerance	Tolerance on converging the solver to the target heat transfer rate.

Q Map (Actual Flow Dual Stream) Properties

Q Table: Heat Transfer Units	Default heat transfer units that the Q Table uses.
Q Map	The table containing the columns cold mass flow rate, hot mass flow rate, and overall heat transfer.
Q Map: Q	The column within the Q map representing the overall heat transfer rate.
Q Map: Cold Mass Flow Rate	The column within the Q map representing the cold stream mass flow rate.
Q Map: Hot Mass Flow Rate	The column within the Q map representing the hot stream mass flow rate.
Cold Fluid Inlet Temperature	Test condition inlet temperature of the cold stream.
Hot Fluid Inlet Temperature	Test condition inlet temperature of the hot stream.
Cold Upstream Boundaries	Upstream boundary interfaces of the cold stream.
Hot Upstream Boundaries	Upstream boundary interfaces of the hot stream.

What Is the UAL Map?

This method is available with the actual dual stream heat exchanger option only.

The UAL Map allows a map of curves (V_{cold}, V_{hot} and UA_l) to specify the UAL directly when multiple heat exchanger curves are available. This option requires you to specify a table containing m_{cold}, m_{hot} and UA_l. That is, specify the mass flow rates of the cold and hot streams rather than the velocities. If the velocities of the two streams are known, they can be multiplied by the product of the average density and inlet flow area, thus creating the UAL map table.

The UAL Map is selected as shown below.



Properties

UAL Map	The table containing the columns cold mass flow rate, hot mass flow rate, and local heat transfer coefficient multiplied by the area.
UAL Map: Cold Mass Flow Rate	The column within the UAL map representing the cold stream mass flow rate.
UAL Map: Hot Mass Flow Rate	The column within the UAL map representing the hot stream mass flow rate.
UAL Map: UAL	The column within the UAL map representing the local heat transfer coefficient multiplied by the area.
UAL Map: Heat Transfer Units	Default heat transfer units that the UAL map uses.
Cold Upstream Boundaries	Upstream boundary interfaces of the cold stream.
Hot Upstream Boundaries	Upstream boundary interfaces of the hot stream.

Expert Properties

Cold Flow Exponent	Exponent ec in Eqn. (419).
Hot Flow Exponent	Exponent eh in Eqn. (419).
Goodness of fit	A read-only value ranging between 0 and 1 that indicates how well the fitted curve represents the data. Typically, this value can be expected to be within 0.96 and 0.98 for a reasonable fit.
Coefficient C0	Coefficient C_0 in Eqn. (419). This computed value is read-only.
Coefficient C1	Coefficient C_1 in Eqn. (419). This computed value is read-only.
Coefficient C2	Coefficient C_2 in Eqn. (419). This computed value is read-only.
UAL Frozen	When activated, freezes the UAL table in its current state (actual dual stream heat exchanger option only).

What Is Target Heat Rejection?

STAR-CCM+ allows you to set a target heat rejection rate for the heat exchanger without setting the temperature at the hot inlet. STAR-CCM+ then predicts the temperature that is required at the inlet to sustain the heat rejection rate. This Target Heat Rejection option is selected by setting the Hot Inlet Temperature property of the Hot Stream Inlet Temperature Specification node to Inferred.



When you select this option, STAR-CCM+ iteratively modifies the hot stream inlet boundary temperature until the specified heat rejection rate is achieved.

The default option for the Hot Stream Inlet Temperature Specification node is Specified. This option indicates that you know and specify the hot stream inlet boundary temperature. The Target Heat Rejection node is added to the interface Physics Values when you choose Inferred.



Target Heat Rejection Properties

Target Heat Rejection Rate	The heat transfer rate that is desired from the heat exchanger.
Hot Inlet Boundaries	The hot stream boundaries at which the temperature must be predicted.
Temperature Jump Solver URF	Under-relaxation factor for the heat exchanger temperature jump solver.

Target Heat Rejection Expert Properties

Temperature Prediction Accuracy	An advanced option which allows you to control the accuracy of temperature prediction.

What Is the Heat Exchanger UAL?

This option is made available for a solid-fluid type dual stream heat exchanger, which is created between a solid region and a fluid region or vice versa. For a solid-fluid type heat exchanger, when the basic dual stream option is chosen as the Heat Exchanger Method, the Heat Exchanger UAL option is added as a physics value node to the fluid region. This option is entered as a scalar profile.



The heat exchanger UAL scalar field represents the local (or cell) heat exchange rate, that may either be specified as a constant or a function of the fluid properties. This user-specified value is added to the source term of the fluid energy equation and subtracted from the source term of the solid energy equation.

Modeling Energy in Solids

This section provides information about the properties of the solid energy models in STAR-CCM+:

- Coupled Solid Energy
- Segregated Solid Energy
- Modeling Heat Conduction in Solid Shells
- Modeling Exothermic Solids
- Solid Energy Formulation

Modeling Solid Energy Using a Coupled Approach

The Coupled Solid Energy model is available for selection after you enable the Solid model during material model selection.



Properties

Implicit	Uses a checkbox to select the integration scheme to use.	
	Activated	The coupled implicit solver is invoked. Adds a Coupled Implicit node to the Solvers node.
	Deactivated	The coupled explicit solver is invoked. Adds a Coupled Explicit node to the Solvers node.

Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Modeling Solid Energy Using a Segregated Approach

The Segregated Solid Energy model is available for selection after you enable the Solid model during material model selection.



Properties

None.

Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Modeling Heat Transfer in Solid Shells

Problems where thin solid shells are involved in heat transfer are common in engineering practice. For thin solid shells, STAR-CCM+ can model heat conduction within the shell, and radiative heat transfer between the surface of the shell and the surfaces of other objects. Modeling of in-plane shell heat conduction is required when the following apply:

- The heat conduction in the lateral direction of the shells cannot be neglected.
- The shells are so thin that their meshing is not practical.

Heat transfer in solid shells happens on the surface of the shell. Therefore, to provide support for the discretization on surfaces, include the Shell Three-Dimensional model in a physics continuum for solid shells.



The rest of the model selection is identical to that of solid materials in conjugate heat transfer problems.

A shell region is a surface domain in space and can be created using region boundaries. Each shell region is assigned a unique name that is based on the name of the boundary which is used for the shell region creation, although it can be renamed.

A shell region is joined to another region, or shell region, using an interface, so that information can be passed between them. An interface between a shell region and the region which owns the boundary that is used to create the shell region is created automatically during the shell region creation.

You can model radiation in solid shells as you do in other volumetric solids. Specifically, you can:

- Set emissivity, reflectivity, specularity, and transmissivity in shell boundaries.
- Model radiative transfer between the external side of the shell and the surfaces of other objects, using the environment boundary condition.
- Model opaque shells by turning off the radiative transfer option within the shell region. (The option is on by default.)

See also:

- Specifying the Solid Shell Thickness for the Region
- Conducting Solid Shells Formulation

Specifying the Solid Shell Thickness for the Region

The Shell Thickness node represents the thickness of a solid shell. This node appears under the Physics Values node of a shell region. The shell thickness is entered as a scalar profile.



Modeling Exothermic Solids

If temperature exceeds the ignition temperature of an exothermic (mold) material for the first time, a certain amount of energy is released over time. In a simulation, the cells that represent the exothermic material each burn for the same amount of time, but at different times. This difference in timing makes it difficult to model an exothermic solid through a user-defined volumetric heat source.

To quantify this energy in STAR-CCM+, use the Exothermic Solid model. This model is a simple combustion model that works like a heat source. When temperature in the solid exceeds ignition temperature, a user-specified heat source simulates combustion.

Two material properties, Autoignition Temperature and Power of Reaction, suffice to describe the entire effect of the exothermic solid material on the energy balance. For example, if a simulation involves two different exothermic solid materials, they would differ only in those material properties. There are no reaction constants or, for instance, no Arrhenius equation from which the energy release over time is computed; the Power of Reaction material property gives this quantity directly.

The main simplification is the Reaction Time field function. It returns the time that the temperature exceeds the Auto-Ignition Temperature.

The application for this model is exothermic mold materials. They are used to enhance feeding systems by extending the solidification time of the reservoirs of liquid metal which are intended to compensate for shrinkage on solidification. The exothermic reaction is initiated when molten metal meets the exothermic mold material.

See also:

Selecting the Exothermic Solid Model

Working with the Autoignition Temperature

Working with the Power of Reaction

Exothermic Solid Field Functions Reference

Selecting the Exothermic Solid Model

To select the Exothermic Solid model:

1. Include the following models with your selections:
   - The Solid material model
   - The Segregated Solid Energy model
   - The Implicit Unsteady time model
   - The Casting optional model

2. Select the Exothermic Solid model from the *Optional Models* group box.



When you close the model selection dialog, the Exothermic Solid model is shown in the list under the Continua > Physics > Models node. The model has no user-defined properties.



Working with the Autoignition Temperature

When you activate the Exothermic Solid model, you can specify the autoignition temperature as a material property. In the following screenshot, the example shows this property available in the Continua > Physics > Models > Solid > Al > Material Properties node.



For a calculation method, choose a constant value (K) or a user-defined field function.

Working with the Power of Reaction

When you activate the Exothermic Solid model, you can specify the power of reaction as a material property. In the following screenshot, the example shows this property available in the Continua > Physics > Models > Solid > Al > Material Properties node.



For a calculation method, choose a constant (W/kg), a user-defined field function, or a Table (Reaction Time) table method. With the Table (Reaction Time) method, the reaction time is used for interpolation of the table and the physical time (as for the regular Table (Time) method).

Exothermic Solid Field Functions Reference

The following primitive field functions are made available to the simulation when the Exothermic Solid model is used. They have a typical set of base properties.

Autoignition Temperature	The temperature at which a material ignites, and at which reaction time begins to accumulate.
Power of Reaction	The energy release during the reaction time of the material.
Reaction Time	Accumulates the time that the temperature exceeds the autoignition Temperature. If the temperature decreases below the autoignition temperature, the reaction time remains constant until the temperature exceeds the autoignition temperature again.

Modeling Average Circumferential Heat Flux

In a system with a rapidly spinning, axially symmetric body, thermal flux characteristics average out around the circumference of the spinning body. A rotating drive shaft near an engine and a spinning motor in a casing give examples of such systems.

Use the Circumferential Flux Averaging model for such systems. This is a heat transfer model for the special case of rapidly rotating bodies, part of the moving reference frame feature. The model maps out equal-width strips of constant temperature along the axis of the spinning body, using the mesh boundaries for averaging.

The model covers two general cases:

- A solid body spinning rapidly in an interfacing fluid medium. Convection, conduction, and radiation operate within in the fluid. Conduction and radiation operate within the solid. The interface is opaque to radiation.
- A solid or fluid body spinning rapidly in a region without mesh, with no medium between the body and its environment. The body and the external environment interact only by thermal radiation between their boundaries, across the empty meshless region. Within the body, convection (if it is fluid), conduction, and radiation can all operate. Convection and radiation heat transfer operate between the boundaries of the body and the environment.

Only use the Circumferential Heat Flux Averaging model under the following conditions:

- The body in question must be axisymmetric, with little or no circumferential variation in any properties or physical conditions. Minor variations from symmetry may still give acceptable results, but use careful judgement.
- The axisymmetric body must be spinning very fast, and the difference between the angular velocity of the spinning body and any other bodies in the system is either very large or very small.

  "Very fast" means the spinning body must complete several rotations within its thermal relaxation time for this model to be applicable.
- If you want to include radiation modeling in the simulation, interface boundaries and wall boundaries between the rotating region and the surroundings must be opaque to radiation. (The model does not allow radiation from a non-rotating body to pass through the boundaries of a rotating body.)

For more information, see:

- Selecting the Circumferential Flux Averaging Model
- Using the Circumferential Heat Flux Averaging Model with an External Boundary
- Using the Circumferential Heat Flux Averaging Model on an Interface
- Circumferential Heat Flux Averaging Formulation

Selecting the Circumferential Flux Averaging Model

To select the Circumferential Heat Flux Averaging model:

1. Select any energy model, but do ***not*** select any of these incompatible models:
   - Explicit or Implicit Unsteady time models
   - Models that require unsteady time
   - Eulerian models
   - Harmonic Balance models
   - The Fluid Film model
2. Select the Circumferential Flux Averaging model from the *Optional Models* group box.



You can select the Circumferential Heat Flux Averaging model for any Physics node in the simulation tree, if the conditions for the model apply.

When you close the *Physics Model Selection* dialog, the Circumferential Heat Flux Averaging model is shown in the list under the Continua > Physics > Models node. The model has no user-defined properties.



Using the Circumferential Heat Flux Averaging Model on an Interface

Use this method in the case of a solid body spinning rapidly in an interfacing fluid medium. Before you begin, set up a rotating reference frame for the spinning region attached to the interface. See A Body Spinning in an Interfacing Medium.

1. Select the Circumferential Heat Flux Averaging model for both continua on each side of the interface.
2. Under the Interface node, open the Physics Conditions node for the interface being averaged.

3. Select Circumferential Heat Flux Averaging Option and, in the *Properties* window, click to activate the option.



4. Open the Physics Values node for the boundary.
5. Select Axial Strip Specification Method and, in the *Properties* window, enter a value for the Number of Axial Strips.



Run the simulation.

Using the Circumferential Heat Flux Averaging Model with an External Boundary

Use this method in the case of a solid or fluid body spinning rapidly in a region without mesh, with no medium between the body and its environment, represented by the boundary of the region beyond the meshless one. Before you begin, set up a rotating reference frame for the spinning region. See A Body Spinning in a Meshless Region.

1. Under the Regions node, open the Physics Conditions node for the rotating region.
2. Select Thermal Specification and, in the *Properties* window, select Convection or Environment from the Condition drop-down menu.



3. Select Circumferential Heat Flux Averaging Option and, in the *Properties* window, click to activate the option.



4. Open the Physics Values node for the boundary.
5. Select Axial Strip Specification Method and, in the *Properties* window, enter a value for the Number of Axial Strips.



Run the simulation.

Energy Reference

This section provides reference information that is common to the energy models in STAR-CCM+. The common information that is covered here includes:

- Conditions (boundary and region)
- Values (boundary and region)

Static temperature is the initial condition available for energy models in STAR-CCM+.

Conditions

This section lists all the boundary conditions available with the energy models in STAR-CCM+:

- Thermal specification
- User wall heat flux coefficient specification

It also lists region conditions:

- Energy source option
- Convective Velocity Option

In addition, detailed property information is provided.

Boundary conditions are elected by the boundary type, which in turn are limited by the models that have been turned on in the continuum. The boundary types that are available for each of the energy models are listed in the Boundary Types section for the relevant model.

Thermal Specification

This node allows you to determine how the energy flow across the boundary is specified.

Properties

Method — The method that is used to define the heat flow at the boundary.

- Heat Flux — Specifies the amount of energy flowing across the boundary in W/m^2. Activating this option adds a Heat Flux node under the Physics Values node, which must be entered as a scalar profile.

Heat Source	Specifies a user-defined total heat source in W. Activating this option adds a Heat Source node under the Physics Values node, which must be entered as a scalar profile.
Temperature	Specifies the boundary temperature in K. Activating this option adds a Temperature node under the Physics Values node, which must be entered as a scalar profile.
Adiabatic	Specifies that there is no energy transfer across the boundary. Activating this option does not require definition of any values.
Convection	Specifies convection flux across the boundary in W/m^2. Activating this option adds Ambient Temperature and Heat Transfer Coefficient nodes to the Physics Values node, which must be entered as scalar profiles.
Environment	Enables the Environment condition across the boundary. Activating this option adds the following nodes, which generally must be entered as scalar profiles, to the Physics Values node: • Ambient Temperature • External Emissivity • External Reflectivity • Heat Transfer Coefficient • Thermal Resistance

User Wall Heat Flux Coefficient Specification

This node allows you to control the heat flux relationship at walls and interfaces.

The wall heat flux $\dot{\mathbf{q}}''_w$ is linearized as:

$$\dot{\mathbf{q}}''_w = A + BT_c + CT_w + D(T_w)^4 \tag{386}$$

where T_c and T_w are the cell and wall temperatures, respectively. A, B, C, and D are the net wall heat flux coefficients.

In the absence of radiation, the heat flux from the wall into the domain is given by $k\nabla T_w$, where ∇T_w is the temperature gradient at the wall.

This is discretized as described by Eqn. (2558):

$$D_f = \Gamma_f \nabla\phi_f \cdot \mathbf{a} = \Gamma_f[(\phi_1 - \phi_0)\vec{\alpha} \cdot \mathbf{a} + \overline{\nabla\phi} \cdot \mathbf{a} - (\overline{\nabla\phi} \cdot \mathbf{ds})\vec{\alpha} \cdot \mathbf{a}]$$

where:

- $\Gamma_f = k$
- $\nabla\phi_f = \nabla T_w$
- $\nabla\phi_0 = \nabla T_c$ (the temperature gradient in the near-wall cell)
- $\phi_f = T_w$
- $\phi_0 = T_c$

Prior to taking the dot product with area ***a***, this expression is:

$$k\nabla T = k[(T_w - T_c)\vec{\alpha} + \nabla T_c - (\nabla T_c \cdot \mathbf{ds})\vec{\alpha}]$$

which can be rearranged to:

$$k\nabla T = k[\nabla T_c - (\nabla T_c \cdot \mathbf{ds})\vec{\alpha}] - k\vec{\alpha}T_c + k\vec{\alpha}T_w + 0T_w{}^4$$

so the internal heat flux coefficients are:

- $A_{\text{internal}} = k[\nabla T_c - (\nabla T_c \cdot \mathbf{ds})\vec{\alpha}]$
- $B_{\text{internal}} = -k\vec{\alpha}$
- $C_{\text{internal}} = k\vec{\alpha}$
- $D_{\text{internal}} = 0$

The net wall heat flux coefficients A, B, C, and D are composed of the sum of internally calculated and user-defined coefficients:

$$\begin{aligned} A &= A_{\text{internal}} + A_{\text{user}} \\ B &= B_{\text{internal}} + B_{\text{user}} \\ C &= C_{\text{internal}} + C_{\text{user}} \\ D &= D_{\text{internal}} + D_{\text{user}} \end{aligned} \tag{387}$$

The internal wall heat flux coefficients are evaluated first, in the usual manner according to the active physical processes (for example conduction, radiation and/or boiling). Then the net coefficients are evaluated by summing the internal and user-defined values.

The internal coefficients can be included as necessary in the user-defined coefficients using the field functions:

A_{internal} = `$InternalWallHeatFluxCoefficientA ("Internal Wall Heat Flux Coefficient, A")`

B_{internal} = `$InternalWallHeatFluxCoefficientB ("Internal Wall Heat Flux Coefficient, B")`

C_{internal} = `$InternalWallHeatFluxCoefficientC ("Internal Wall Heat Flux Coefficient, C")`

D_{internal} = `$InternalWallHeatFluxCoefficientD ("Internal Wall Heat Flux Coefficient, D")`

So, for example, you can completely replace net coefficient A with some function f by defining:

$$A_{\text{user}} = f - A_{\text{internal}} \tag{388}$$

such that the net coefficient becomes:

$$A = A_{\text{internal}} + (f - A_{\text{internal}}) = f \tag{389}$$

The resultant net wall heat flux coefficients are available using the field functions:

A = `$NetWallHeatFluxCoefficientA ("Net Wall Heat Flux Coefficient, A")`

B = `$NetWallHeatFluxCoefficientB ("Net Wall Heat Flux Coefficient, B")`

C = `$NetWallHeatFluxCoefficientC ("Net Wall Heat Flux Coefficient, C")`

D = `$NetWallHeatFluxCoefficientD ("Net Wall Heat Flux Coefficient, D")`

The user-defined heat flux coefficients do not affect the amount of heat that enters the domain. These coefficients determine how much of the heat flux goes towards increasing the wall temperature, and how much goes towards increasing the cell temperature.

User-defined wall heat flux coefficients can be enabled on Wall or Contact Interface boundaries by setting the Method property of this node.

Properties

Method — Controls whether to specify the heat flux relationship.

- None — Leaves off this feature.
- Specified — Activates this feature. This option adds the following nodes to the Physics Values node, which must be entered as scalar profiles:
  - User Wall Heat Flux Coefficient, A
  - User Wall Heat Flux Coefficient, B
  - User Wall Heat Flux Coefficient, C
  - User Wall Heat Flux Coefficient, D

Energy Source Option

Properties

Unless otherwise specified, the conditions below apply to both fluid and solid regions.

Energy Source Option	Provides energy source options for the region.	
	None	Do not specify an energy source.
	Volumetric Heat Source	Specifies a user-defined volumetric heat source in W/m^3. Activating this option when using the Segregated solver adds the following nodes under the Physics Values node: • Volumetric Heat Source • Energy Source Temperature Derivative Activating this option when using the Coupled solver adds the following nodes under the Physics Values node: • Volumetric Heat Source • Energy Source Pressure Derivative • Energy Source Temperature Derivative • Energy Source Velocity Derivative
	Total Heat Source	Specifies a user-defined total heat source in W. Activating this option when using the Segregated solver adds the following nodes under the Physics Values node: • Heat Source • Energy Source Temperature Derivative Activating this option when using the Coupled solver adds the following nodes under the Physics Values node: • Heat Source • Energy Source Pressure Derivative • Energy Source Temperature Derivative • Energy Source Velocity Derivative

Specific Heat Source	Specifies a user-defined specific heat source in W/kg. Activating this option when using the Segregated solver or the Coupled solver adds the following nodes under the Physics Values node: • Specific Heat Source • Energy Source Temperature Derivative
Heat Exchanger	Specifies a heat exchanger source term. Activating this option adds several Heat Exchanger Enthalpy Source nodes under the Physics Conditions and Physics Values nodes. This option is only available with fluid regions.

Convective Velocity Option

This node, which appears in solid regions when either of the Segregated Solid Energy or Coupled Solid Energy models are enabled, allows you to activate a convective velocity ($\mathbf{v}_s$ in Eqn. (424)) for simulating energy convection within the solid region. The specified velocity field is a field of pure rotation or translation, and is intended for use in steady cases only. It is recommended that the solid energy under-relaxation factor be reduced to a value between 0.5 and 0.6 when convection is included.

This option should only be used when the region is a pure body of rotation or translation. This is not the case when holes or grooves are present, as illustrated below. All free surfaces should be tangential to the direction of rotation or translation.

Convective Velocity Option: only for bodies of rotation



Properties

Convective Velocity Option	Provides a convective velocity field for the region.	
	None	Do not specify a convective velocity.
	Specified	Specify a convective velocity. A Convective Velocity node is added to the Values node.

Values

This section lists all the boundary values that are available for energy models in STAR-CCM+ and it provides a detailed reference for the methods of setting them:

- Ambient temperature
- External emissivity
- External reflectivity
- Heat flux
- Heat transfer coefficient
- Static temperature

- Thermal resistance
- Total temperature
- User wall heat flux coefficient, A
- User wall heat flux coefficient, B
- User wall heat flux coefficient, C
- User wall heat flux coefficient, D

It also lists region values:

- Anisotropic thermal conductivity
- Volumetric heat source
- Heat source
- Specific heat source
- Energy source pressure derivative
- Energy source temperature derivative
- Energy source velocity derivative
- Solid density
- Solid specific heat
- Solid thermal conductivity
- Convective Velocity

Boundary values are elected by the boundary type, and/or by methods that are selected in corresponding nodes in the boundary conditions node. The boundary types that are available are listed in the Boundary Type section of each model.

Ambient Temperature

This value node is visible if `Convection` or `Environment` is selected as the Method property for the Thermal Specification node of the Conditions node. It is input as a scalar profile.

External Emissivity

Certain energy and radiation models use this boundary value. It is only visible if the `Environment` method is chosen as the Thermal Specification and is entered as a scalar profile.

External Reflectivity

Certain energy and radiation models use this boundary value. It is only visible if the `Environment` method is chosen as the Thermal Specification. It can be entered as a scalar profile or can be calculated automatically using the Auto-Calculate method, which is available for reflectivity.

Heat Flux

This value node is only visible if `Heat Flux` is selected as the Method property for the Thermal Specification node of the Conditions node. It is input as a scalar profile in W/m^2.

A positive specified heat flux value q means that heat is flowing into the domain. This is opposite to the convention that the boundary heat flux:

$$q_w = \boldsymbol{q} \cdot \boldsymbol{n}$$

is negative for an inflow of heat (due to outward pointing surface normals at boundaries). The sign of q is switched at this point.

Heat Transfer Coefficient

Energy models require this boundary value. It is visible if the `Convection` or `Environment` method is chosen as the Thermal Specification and is entered as a scalar profile.

Static Temperature

The temperature of the boundary. This value node is visible if `Temperature` is selected as the Method property for the Thermal Specification node of the Conditions node. It is entered as a scalar profile.

Thermal Resistance

This boundary value appears when the `Environment` method is chosen as the Thermal Specification. It functions in a manner similar to the Thermal Resistance baffle interface value.

Total Temperature

The total temperature of the boundary. It is entered as a scalar profile. Not available in regions of type solid.

User Wall Heat Flux Coefficient, A

This boundary value is available with energy models. It is only visible if the `Specified` method is chosen as the User Wall Heat Flux Coefficient Specification and is entered as a scalar profile. By default it has a constant value of zero.

User Wall Heat Flux Coefficient, B

This boundary value is available with energy models. It is only visible if the `Specified` method is chosen as the User Wall Heat Flux Coefficient Specification and is entered as a scalar profile. By default it has a constant value of zero.

User Wall Heat Flux Coefficient, C

This boundary value is available with energy models. It is only visible if the `Specified` method is chosen as the User Wall Heat Flux Coefficient Specification and is entered as a scalar profile. By default it has a constant value of zero.

User Wall Heat Flux Coefficient, D

This boundary value is available with energy models. It is only visible if the `Specified` method is chosen as the User Wall Heat Flux Coefficient Specification and is entered as a scalar profile. By default it has a constant value of zero.

Anisotropic Thermal Conductivity

This value node is added to a solid region when the thermal conductivity material property of a solid is set to the Anisotropic method. It is entered as a tensor profile. The method of the tensor profile is set by selecting the Anisotropic Thermal Conductivity node and using the drop-down menu in the *Properties* window, as shown below.



Volumetric Heat Source

This node is added when the Energy Source Option region condition is set to Volumetric Heat Source. It is entered as a scalar profile.

Heat Source

This value specifies the total heat source in W. It is available when you activate either the Heat Source Energy Source Option for a region, or the Heat Source Thermal Specification for a boundary.

Specific Heat Source

This node is added when the Energy Source Option region condition is set to Specific Heat Source. It is entered as a scalar profile.

This value specifies the heat source in W/kg.

Energy Source Pressure Derivative

This node represents the derivative of the Energy Source with respect to pressure, S_u in Eqn. (8). The derivative of the user source term is used to linearize Eqn. (8). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution when the function for the source term is a function of pressure. If the source is constant, or not a function of pressure, leave this value at zero.

The Energy Source Pressure Derivative is available when coupled flow is used. This object is added when the Energy Source Option region condition is set to `Specified`, and is entered as a scalar profile.

Energy Source Temperature Derivative

This node represents linearization of the Energy Source with respect to temperature, S_u in Eqn. (8) (coupled) and $s_u dV$ in Eqn. (85) (segregated). The derivative of the user source term is used to linearize these equations. Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution when the function for the source term is a function of temperature. If the source is constant, or not a function of temperature, leave this value at zero.

The Energy Source Temperature Derivative is available when either coupled or segregated flow is used. This object is added when the Energy Source Option region condition is set to `Specified`, and is entered as a scalar profile.

Energy Source Velocity Derivative

This node represents linearization of the Energy Source with respect to x-, y-, [z-] components of velocity, S_u in Eqn. (8). The derivative of the user source term is used to linearize Eqn. (8). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution when the energy source is a function of velocity. If the source is constant, or not a function of velocity, leave this value at zero.

The Energy Source Velocity Derivative is available when coupled flow is used. This object is added when the Energy Source Option region condition is set to `Specified`, and is entered as a vector profile.

Solid Density

This node is only available in regions of the porous type. The density that is provided here is used to specify the solid density ρ_{solid} in Eqn. (96). The fluid density ρ_{fluid} in Eqn. (96) is the density property that is specified for the fluid continuum. It is entered as a scalar profile.



Solid Specific Heat

This node is only available in regions of the porous type. The specific heat that is provided here is used to specify the solid specific heat $C_{p\text{-}solid}$ in Eqn. (96). It is entered as a scalar profile.



Solid Thermal Conductivity

This node is only available in regions of the porous type. The thermal conductivity that is provided here is used to specify k_{solid} in Eqn. (250). The fluid thermal conductivity k_{fluid} in Eqn. (250) is extracted from the property

that is specified for the fluid continuum. It is entered as a tensor profile, and it is possible to use field functions and other profile methods to define the components of the tensor.



Convective Velocity

This node, which has properties, is added when the Convective Velocity Option for a solid region has been set to Specified. Rotation is specified by providing a rotation axis, an origin for that axis, and a rotation rate. Translation is specified by providing the translation velocity. User-defined coordinate systems can be used when specifying quantities.



Properties

Direction	Provides a convective velocity field for the region.
Origin	The Cartesian position vector specifying the origin of the axis of rotation.

Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.
Rotation Rate	The rotation rate around the axis that is specified by Direction and Origin. This is additive with Translation Velocity.
Translation Velocity	The translation velocity of the region. This is additive with Rotation Rate.

Heat Transfer Field Functions Reference

The following primitive field functions are made available to the simulation when the following models are used:

- Coupled Energy
- Segregated Fluid Energy

These objects all have the same base properties, with some field functions also having reference value properties (indicated by a `ref` subscript in the following descriptions).

Boundary Advection Heat Flux

The magnitude of the advection heat flux vector normal to the boundary. The quantity includes the energy that is carried by the fluid across the domain boundary and is only defined at flow openings. Consistent with the other heat fluxes, the sign convention is positive for transfer out of the domain.
A version of this field function is created for each phase in a multiphase continuum.

Boundary Advection Heat Transfer

The heat transferred by advection at the boundary. The quantity is essentially the rate that energy is advected into and out of the domain by the moving fluid. Consistent with the other heat transfer rates, the sign convention is positive for transfer leaving the domain.
A version of this field function is created for each phase in a multiphase continuum.

Boundary Conduction Heat Flux

The magnitude of the conduction heat flux vector normal to the boundary, defined as $(\dot{\mathbf{q}}''_{\text{conduction}} \cdot \mathbf{a})/|\mathbf{a}|$. This quantity includes molecular and turbulent diffusion effects at the boundary (positive denotes transfer out of the domain). The advection contribution at open flow boundaries is separately included in the Boundary Advection Heat Flux.
For a multiphase fluid, the boundary conduction heat flux is the sum of the boundary conduction heat fluxes for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted.

Boundary Conduction Heat Flux of <phase>	The magnitude of the conduction heat flux vector normal to the boundary for the specified phase. A version of this field function is created for each phase in a multiphase continuum. This field function is volume fraction weighted for each phase.
Boundary Conduction Heat Transfer	The heat transferred by conduction at the boundary, defined as $\dot{\mathbf{q}}''_{\text{conduction}} \cdot \mathbf{a}$. This quantity includes molecular and turbulent diffusion effects at the boundary (positive denotes transfer out of the domain). The advection contribution at open flow boundaries is separately included in the Boundary Advection Heat Transfer. For a multiphase fluid, the boundary conduction heat transfer is the sum of the boundary heat transfer values for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted.
Boundary Conduction Heat Transfer of <phase>	The heat transferred by conduction at the boundary for the specified phase. A version of this field function is created for each phase in a multiphase continuum. This field function is volume fraction weighted for each phase.
Boundary Heat Flux	The magnitude of the heat flux vector normal to the boundary, defined as $(\dot{\mathbf{q}}'' \cdot \mathbf{a})/\lvert\mathbf{a}\rvert$. For a multiphase fluid, the boundary heat flux is the sum of the boundary heat fluxes for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted. When the VOF model is used, this sum includes the wall boiling heat flux.
Boundary Heat Flux of <phase>	The magnitude of the heat flux vector normal to the boundary for the specified phase. A version of this field function is created for each phase in a multiphase continuum. This field function is volume fraction weighted for each phase.

Boundary Heat Flux Radiation Coefficient

The boundary heat flux radiation coefficient, defined by $\sigma\varepsilon$, where σ is the Stefan-Boltzmann constant and ε is the total emissivity. This quantity can be used in external CAE packages to express the radiative emission flux in terms of the boundary temperature:

$$\dot{\mathbf{q}}''_{\text{rad,emission}} = DT_w^4$$

where D is the boundary heat flux radiation coefficient.

Boundary Heat Transfer

The heat transferred at the boundary, defined as $\dot{\mathbf{q}}'' \cdot \mathbf{a}$.
For a multiphase fluid, the boundary heat transfer is the sum of the boundary heat transfer values for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted.

Boundary Heat Transfer of <phase>

The heat transferred at the boundary for the specified phase. A version of this field function is created for each phase in a multiphase continuum.
This field function is volume fraction weighted for each phase.

Effective Conductivity

The effective thermal conductivity given by $k_{\text{eff}} = k + (\mu_t C_p)/\sigma_t$ where σ_t is the turbulent Prandtl number.
For porous regions this field function is given by

$$k_{\text{eff}} = \alpha[k + (\mu_t C_p)/\sigma_t] + (1-\alpha)k_{solid}$$

External Ambient Temperature

The boundary ambient temperature that you specify for convection thermal boundary conditions. This field function is only available on boundaries for which the `Convection` option has been selected as the Thermal Specification.

External Heat Transfer Coefficient	The boundary external heat transfer coefficient that you specify for convection thermal boundary conditions. This field function is only available on boundaries for which the `Convection` option has been selected as the Thermal Specification.
Heat Exchanger Density of <phase, region>	The density of the specified phase in the specified region of a heat exchanger.
Heat Exchanger Dynamic Viscosity of <phase, region>	The viscosity of the specified phase in the specified region of a heat exchanger.
Heat Exchanger Energy Transfer	The amount of energy from the hot to the cold flow of the heat exchanger.
Heat Exchanger Enthalpy of <phase, region>	The enthalpy of the specified phase in the specified region of a heat exchanger.
Heat Exchanger Mass Fraction of <phase, region>	The mass fraction of the specified phase in the specified region of a heat exchanger.
Heat Exchanger Specific Heat of <phase, region>	The specific heat of the specified phase in the specified region of a heat exchanger.
Heat Exchanger Temperature Difference	The local temperature difference between the hot and cold fluid streams at each cell in the heat exchanger region. This function is valid only for a single-phase dual-stream heat exchanger model.

Heat Exchanger Temperature Jump	The temperature jump (positive or negative) that is applied to the specified temperature boundary condition at the hot fluid stream inlet boundary. This function is valid only when using the Target Heat Rejection option for a single-phase dual-stream heat exchanger model.
Heat Exchanger Thermal Conductivity of <phase, region>	The thermal conductivity of the specified phase in the specified region of a heat exchanger.
Heat Exchanger U Velocity of <phase, region>	The X component of the fluid velocity for the specified phase in the specified region of a heat exchanger.
Heat Exchanger V Velocity of <phase, region>	The Y component of the fluid velocity for the specified phase in the specified region of a heat exchanger.
Heat Exchanger Volume Fraction of <phase, region>	The volume fraction of the specified phase in the specified region of a heat exchanger.
Heat Exchanger W Velocity of <phase, region>	The Z component of the fluid velocity for the specified phase in the specified region of a heat exchanger.
Heat Transfer Coefficient	The wall boundary heat transfer coefficient defined by: $$h = (\dot{\mathbf{q}}''_{\text{conduction}} \cdot \mathbf{a})/[\lvert\mathbf{a}\rvert(T_{\text{ref}} - T_f)]$$ where T_f is the wall boundary temperature and T_{ref} is the specified reference temperature in the field function.

Local Heat Transfer Coefficient

The local heat transfer coefficient is the heat transfer coefficient between the near-wall cell and the boundary. It is dependent upon the distance between the boundary and the near-wall cell, the media's thermal properties and, in the case of turbulent flow, the near-wall turbulence state. Because of the dependence on the mesh, this quantity should not be considered a bulk heat transfer coefficient.
For a multiphase fluid, the local heat transfer coefficient is the sum of the local heat transfer coefficients for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted.

Local Heat Transfer Coefficient of <phase>

The local heat transfer coefficient of the specified phase. A version of this field function is created for each phase in a multiphase continuum.
This field function is volume fraction weighted for each phase.

Local Heat Transfer Reference Temperature

The local heat transfer reference temperature is the effective temperature that may be used to linearize the total heat flux to the boundary with respect to the boundary temperature: $\dot{\mathbf{q}}''_{\text{total}} = h_{\text{local}}(T_{\text{ref}} - T_w)$, where h_{local} is the local heat transfer coefficient and T_w is the boundary temperature. The effect of radiation is explicitly included in this quantity. This quantity is useful for export to external CAE packages with the local heat transfer coefficient.
For a multiphase fluid, the local heat transfer reference temperature is given by:

$$\sum(h_i \times Tr_i)/\sum h_i$$

where h_i is the local heat transfer coefficient of phase i and Tr_i is the local heat transfer reference temperature of phase i.

Local Heat Transfer Reference Temperature of <phase>

The local heat transfer reference temperature of the specified phase. A version of this field function is created for each phase in a multiphase continuum.
This field function is not volume fraction weighted.

Nusselt Number	The non-dimensional heat transfer coefficient $Nu = hL_{\mathrm{ref}}/k$, where h is the heat transfer coefficient, k is the thermal conductivity and L_{ref} is the specified reference length in the continuum.
Porous Solid Conductivity	Solid conductivity that is valid only for porous regions. This field function reflects the data that you enter as part of the Solid Thermal Conductivity physics value.
Relative Total Temperature	The temperature obtained by isentropically bringing the flow to rest in the relative frame of motion.
Specified Y+ Heat Transfer Coefficient	The specified Y+ heat transfer coefficient is calculated in the same way as the local heat transfer coefficient (see Eqn. (370)), but at a user-specified y^+ value instead of the value that is associated with the near-wall cell. For a multiphase fluid, the specified Y+ heat transfer coefficient is the sum of the specified Y+ heat transfer coefficients for each phase. This summation is valid because the phase-specific coefficients are already volume fraction weighted. To make this field function available, the Turbulence models must be activated.
Specified Y+ Heat Transfer Coefficient of <phase>	The specified Y+ heat transfer coefficient of the specified phase. A version of this field function is created for each phase in a multiphase continuum. This field function is volume fraction weighted for each phase. To make this field function available, the Turbulence models must be activated.

Specified Y+ Heat Transfer Reference Temperature

The specified Y+ heat transfer reference temperature as given by Eqn. (371).
For a multiphase fluid, the specified Y+ heat transfer reference temperature is given by:

$$\sum(h_i \times Tr_i)/\sum h_i$$

where h_i is the specified Y+ heat transfer coefficient of phase i and Tr_i is the specified Y+ heat transfer reference temperature of phase i.
To make this field function available, the Turbulence models must be activated.

Specified Y+ Heat Transfer Reference Temperature of <phase>

The specified Y+ heat transfer reference temperature of the specified phase. A version of this field function is created for each phase in a multiphase continuum. This field function is not volume fraction weighted.
To make this field function available, the Turbulence models must be activated.

Thermal Conductivity

Reflects the material property data that you enter as part of the physics continuum.

Total Temperature

The total temperature is the temperature that is obtained by isentropically bringing the flow to rest. For an ideal gas with constant specific heat, it is defined as:

$$T_t = T\left(1 + \frac{(\gamma - 1)}{2}M^2\right)$$

where T, M, and γ are the static temperature, Mach number, and ratio of specific heats, respectively.

Energy Formulation

This section contains the following sub-sections on formulations used in modeling energy:

- Single-Phase Boiling Formulation
- Thin Film Formulation
- Single-Stream Heat Exchanger Formulation
- Dual-Stream Heat Exchanger Formulation
- Circumferential Heat Flux Averaging Formulation
- Solid Energy Formulation

Single-Phase Boiling Formulation

There are many models of nucleate-boiling heat transfer, and none of them is accepted as a standard, more-or-less reliable model. The main mechanisms which these models try to mimic are summarized below [380]:

- Evaporation (latent-heat) transport—Heat is supplied near the wall to bubbles which then move away into the bulk liquid. In subcooled boiling, there could be simultaneous evaporation at the base of a bubble and condensation at its tip.
- Micro-convection—Growth and collapse of bubbles in the wall proximity cause random liquid motion.
- Vapor-liquid exchange—A "Reynolds analogy" model in which bubble growth causes an exchange of liquid between the wall and bulk regions.
- Surface quenching—A variation of the vapor-liquid exchange, assuming transient conduction to the cold liquid contacting the wall after bubble departure.
- Wake flow—Liquid motion behind a departing bubble causes convection from the wall.
- Enhanced natural convection—Bubble columns produce a cellular flow pattern similar to natural convection above a large horizontal surface.
- Thermo-capillary flow—Small variations in surface tension due to temperature differences between the base and tip of a bubble cause a jet of hot liquid to flow away from the wall.

The empirical correlation that Rohsenow [382] presented is used to calculate the surface heat flux due to boiling:

$$q_{bw} = \mu_l h_{lat} \sqrt{\frac{g(\rho_l - \rho_v)}{\sigma}} \left(\frac{C_{p_l}(T_w - T_{sat})}{C_{qw} h_{lat} Pr_l^{n_p}} \right)^{3.03} \tag{390}$$

In this equation, μ_l, h_{lat}, C_{p_l}, ρ_l, and Pr_l are the dynamic viscosity, latent heat, specific heat, density, and Prandtl number of the liquid phase, n_p is the Prandtl number exponent (1.73 by default), g is gravity, and ρ_v is the vapor density, σ is the surface tension coefficient at the liquid-vapor interface, T_w is the wall temperature, T_{sat} is the saturation temperature, and C_{qw} is an empirical coefficient varying with the liquid-surface combination.

The vapor mass generation rate $\dot{m}_{ew}$ over the area that nucleation sites cover is:

$$\dot{m}_{ew} = \frac{C_{ew} q_{bw}}{h_{lat}} \tag{391}$$

where C_{ew} is a model constant stating how much of the boiling heat flux is used for creation of vapor bubbles.

If the Rohsenow correlation is applied outside its range of applicability (for example, to the film boiling regime), unrealistically high heat fluxes could result. Fluid temperatures can become higher than the near wall temperature. This behavior stems from the fact that the Rohsenow correlation does not depend on the fluid temperature; heat enters the domain irrespective of the fluid temperature.

To prevent this condition, the heat flux that the correlation calculates is multiplied by:

$$max\left[0, min\left(\left(\frac{T_w - T}{Tw - T_{sat}}\right), 1\right)\right] \tag{392}$$

where T is the fluid temperature near the heated wall. Thus, if $T < T_{sat}$, the Rohsenow correlation is used directly, and if $T > T_w$ the boiling heat flux is zero. For fluid temperatures which are between the wall and the saturation temperature, only a fraction of the heat flux predicted by the correlation is used.

Thin Film Formulation

This section includes the following formulation:

- Defogging
- Deicing

Defogging

The defogging model is based on solving an additional scalar transport equation that represents the mass fraction of water vapor. A source/sink term for the scalar is considered for condensation/evaporation of the fog layer as well as the latent heat required for transition.

When there is a difference between the water vapor content at the fog layer and the cell next to this surface, the model calculates a rate of evaporation or condensation depending on the conditions. The assumptions are:

- The vapor content in the air does not affect the thermal properties of the vapor-air mixture.
- The water vapor mass is neglected with respect to the total mass in a cell.

The rate of mass transfer per unit surface $[\mathrm{kg/s\text{-}m^2}]$ is:

$$\dot{m} = \rho_g \beta_g \cdot CG \ \ln \frac{CG(1 - C_s)}{CS(1 - C_g)} \tag{393}$$

where ρ_g is density $[\mathrm{kg/m^3}]$.

$$\beta_g = C_{\mathrm{emp}} \cdot D_V / L \cdot 0.045 \cdot Re^{0.8} Sc^{0.43} [\mathrm{m/s}] \tag{394}$$

where C_{emp} is the empirical constant useful for calibration (0.05-0.9) for windshield applications, L is a characteristic length (cubic root of cell volume next to fog layer boundary) [m], D_V is diffusion of vapor in air $[\mathrm{m^2/s}]$, Re is the Reynolds number (which is based on the gas state and the characteristic length L), and Sc is the Schmidt number, and:

$$CG = 0.622 + 0.378 C_g \tag{395}$$

$$CS = 0.622 + 0.378 C_s \tag{396}$$

where C_g is the actual concentration of vapor in air and C_s is the saturation concentration of vapor in air.

For a given temperature, the saturation pressure is calculated from the following expression:

$$p_s(T) = 611.85 e^{\frac{17.502(T - 273.15)}{T - 32.25}} \tag{397}$$

The amount of evaporated/condensed vapor in one time-step is:

$$d\dot{m} = \dot{m} \cdot dt \tag{398}$$

The thickness of the liquid film is updated as:

$$s = s_{t-1} + dm/\rho_l \tag{399}$$

Deicing

Deicing is assumed to be a quasi-steady process. An ice layer at temperature T_{ice} is melted with a wall heat flux q, of which a proportion C_{emp} is absorbed in the ice layer. Melted ice at the freezing temperature T_F escapes as run-off which cannot refreeze in the current model. The mass and energy fluxes that are associated with this process are illustrated in the diagram below:



The mass balance for the deicing process is simply

$$\dot{m}_{melt} = \dot{m}_{run-off} \tag{400}$$

As it is assumed that run-off cannot refreeze, there is the additional constraint that $\dot{m}_{run-off} \geq 0$ and hence from Eqn. (400), $\dot{m}_{melt} \geq 0$.

The energy balance for this process, neglecting the kinetic energy that is associated with the mass fluxes, is:

$$\dot{m}_{melt}\, h_s(T_{ice}) + q = \dot{m}_{run-off}\, h_L(T_F) + (1 - C_{emp})q \tag{401}$$

where $h_s(T_{ice})$ is the enthalpy of the solid ice at the temperature of the ice layer, and $h_L(T_F)$ is the enthalpy that is associated with the melting ice at the freezing temperature. These enthalpies are related to temperature by:

$$h_s(T) = C_{p,s}T \tag{402}$$

$$h_L(T) = C_{p,s}T + L \tag{403}$$

Using these relations with the mass balance, and rearranging, gives an explicit expression for the melting rate:

$$\dot{m}_{melt} = \frac{C_{emp} q}{C_{p,s}(T_F - T_{ice}) + L} \tag{404}$$

The constraint, $\dot{m}_{melt} \geq 0$, must be imposed on this equation. If Eqn. (404) predicts a negative $\dot{m}_{melt}$, it is clipped to zero.

Given $\dot{m}_{melt}$, the ice thickness, s, can be updated using:

$$\frac{ds}{dt} = -\frac{\dot{m}_{melt}}{\rho_s} \tag{405}$$

where ρ_s is the density of the ice layer.

Accretion

The preceding analysis is extended to include an additional incident mass flux of liquid droplets, each with mass m_d , temperature T_d and velocity v_d . This expression is intended to represent liquid droplets impinging on the surface. The various contributions to the overall energy balance are depicted below:



In the above diagram, $A_{s,b}$ is the surface area of the boundary face onto which the droplets impinge and Δt is the time-step. The term:

$$\frac{\sum m_d\left[h_L(T_d) + \frac{1}{2}v_d^2\right]}{A_{s,b}\Delta t}$$

indicates the overall energy flux due to the impinging droplets.

The mass balance is therefore given as:

$$\dot{m}_{melt} + \dot{m}_d = \dot{m}_{run-off} \tag{406}$$

where:

$$\dot{m}_d = \frac{\sum m_d}{A_{s,b}\Delta t} \tag{407}$$

In addition to the constraint that $\dot{m}_{run-off} \geq 0$, there is also $\dot{m}_d \geq 0$. The constraint on $\dot{m}_{melt}$ is now $\dot{m}_{melt} \geq -\dot{m}_d$, that is, freezing is allowed but no faster than the incoming liquid rate.

The energy balance for this process is:

$$\dot{m}_{melt}\, h_s(T_{ice}) + \frac{\sum m_d\left[h_L(T_d) + \frac{1}{2}v_d^2\right]}{A_{s,b}\Delta t} + q = (1 - C_{emp})q + \dot{m}_{run-off}\, h_L(T_F) \tag{408}$$

giving:

$$\dot{m}_{melt} = \frac{C_{emp}q + \frac{\sum m_d\left[C_{p,s}(T_d - T_F) + \frac{1}{2}v_d^2\right]}{A_{s,b}\Delta t}}{C_{p,s}(T_F - T_{ice}) + L} \tag{409}$$

A similar analysis can be made for the accretion of solid particles.

Accretion Ratio

Following Makkonen [38], it is useful to define the accretion ratio for non-zero $\dot{m}_d$ as:

$$\alpha = -\frac{\dot{m}_{melt}}{\dot{m}_d} \tag{410}$$

This quantity characterizes the nature of the accretion process. When $\alpha = 1$ there is no run-off and the accretion is "dry", which tends to form rime ice. Otherwise there is some run-off, and the "wet" growth tends to form glaze ice. The accretion ratio is sometimes also known as the "freezing fraction".

Single-Stream Heat Exchanger Formulation

STAR-CCM+ computes the local cell heat exchange, Q_c, from the specified target total heat exchange, Q_{total}. Q_{total} is given as:

$$Q_c = Q_{total} \frac{V_c \upsilon_c (T_{ref} - T_c)}{\sum V_c \upsilon_c (T_{ref} - T_c)} \tag{411}$$

where:

- Q_{total} is the total heat source (positive Q_{total}) or sink (negative Q_{total}) transferred to/from the region, in Watts.
- T_{ref} is the reference temperature, which is the constant temperature of the fluid stream that is not modeled in the simulation. You specify this value in the simulation, in kelvins.

  You also specify ΔT_{min}, which is the minimum temperature difference that must exist between the two streams.

  However, if the user-specified value is physically incorrect, the T_{ref} value is changed in the code as follows:

  - For a heater (that is, Q_{total} is positive), T_{ref} (which is the constant hot stream temperature in this case) is set to:

    $$T_{ref} = T_{max} + \Delta T_{min} \quad \text{if} \quad T_{ref} < T_{max}$$

    In other words, if you specify a (hot) reference temperature that is less than the maximum cold stream temperature T_{max}, the reference temperature is reset. The cold stream is the one that is being modeled in the simulation. Therefore the reference temperature is set equal to the maximum cold stream temperature plus the minimum temperature difference that must exist between the two streams.

  - For a cooler (that is, Q_{total} is negative), T_{ref} (which is the constant cold stream temperature in this case) is set to:

    $$T_{ref} = T_{min} - \Delta T_{min} \quad \text{if} \quad T_{ref} > T_{min}$$

    In other words, if you specify a (cold) reference temperature that is more than the minimum hot stream temperature T_{min}, the reference temperature is reset. The hot stream is the one that is being modeled

in the simulation. Therefore the reference temperature is set equal to the minimum hot stream temperature minus the minimum temperature difference that must exist between the two streams.

- Q_c is the cell heat source/sink, in Watts.
- V_c is the cell volume, in m^3.
- υ_c is the cell velocity, in m/s.
- T_c is the cell temperature, in kelvins.

Dual-Stream Heat Exchanger Formulation

This section provides formulation for the two methods of heat exchanger modeling:

- Basic
- Actual

Basic Method

In this section, "local" refers to the value stored in each cell.

The goal in the dual stream heat exchanger is to specify or compute a local/cell heat transfer coefficient (multiplied by the local heat exchanger area) UA_l as a function of the local velocity of the cold stream. The local heat exchange between the hot and the cold stream, Q_{cell}, is given by:

$$Q_{cell} = \pm UA_l(T_{hot} - T_{cold}) \tag{412}$$

where the "+" is for the cold stream and "-" is for the hot stream (that is, hot stream loses heat to the cold stream). T_{hot} and T_{cold} refer to the local temperatures of the hot and cold streams respectively. This heat exchange is then added as a source term to the corresponding fluid energy equation after scaling by the cell volume, that is:

$$S_{cell} = Q_{cell}V_{cell}(Nc)/V_{Hx}$$

where S_{cell} is the cell source term, V_{cell} is the cell volume, Nc is the number of cells in the heat exchanger, and V_{Hx} is the volume of all the cells in the heat exchanger.

UAL Polynomial Option

In the UAL Polynomial option, the specified polynomial function is used to compute UA_l as a function of the local velocity of the cold stream. In the UAL Table option, UA_l is interpolated from the specified table of cold flow velocity vs. UA_l.

UAG Table Option

In the UAG Table option, UA_g refers to the overall heat transfer coefficient (multiplied by the heat exchanger area) across the entire heat exchanger, such that,

$$Q = UA_g(T_{\text{hot,in}} - T_{\text{cold,in}}) \tag{413}$$

where Q is the overall heat transfer across the entire heat exchanger, and $T_{\text{hot,in}}$ and $T_{\text{cold,in}}$ refers to the inlet temperatures of the hot and cold streams respectively. The UAG Table is specified as a table of cold flow mass flow rate vs. UA_g, which may be obtained from test/experimental data. You also specify other data pertaining to the test conditions including:

- Hot stream mass flow rate $\dot{m}_h$.
- Inlet temperatures of the cold $T_{\text{cold,in}}$ and hot $T_{\text{hot,in}}$ streams.
- Specific heats of the cold Cp_c and hot Cp_h streams.
- Density of the cold fluid ρ_c.

Using these input data, the specified UAG table is internally converted into a UAL table (as in the UAL Table option). This conversion is performed by considering a heat balance between the cold and hot streams as detailed in the following algorithm.

Repeat the following steps for each entry in the UAG Table:

1. Retrieve cold stream mass flow rate $\dot{m}_c$ and UA_g from the UAG Table.
2. Compute the target overall heat transfer rate:

$$Q_{target} = UA_g(T_{\text{hot,in}} - T_{\text{cold,in}})$$

3. Guess $UA_l = UA_g$ / Nc, where Nc is the total number of cells in the heat exchanger.
4. Build a linear system of equations, by constructing the following set of equations in each cell:

$$\begin{aligned} \dot{m}_c Cp_c(T_c - T_{ci}) &= UA_l(T_h - T_c) \\ \dot{m}_h Cp_h(T_{hi} - T_h) &= UA_l(T_h - T_c) \end{aligned} \tag{414}$$

where T_c, T_h, T_{ci} and T_{hi} are local cell unknown temperatures as shown in the following figure.



Assemble the system for all the cells in the heat exchanger, and solve for the unknown cell temperatures T_c and T_h. Boundary conditions are required only at the inlet boundaries, at which the cell temperatures are set to the user-specified $T_{\text{cold,in}}$ and $T_{\text{hot,in}}$ for the cold and hot streams respectively.

5. Compute the mean temperature difference between the cold and hot streams:

$$\overline{\Delta T} = \frac{\sum'(T_h - T_c)}{(Nc)} \tag{415}$$

6. Compute the estimated overall heat transfer $Q = UA_l\overline{\Delta T}Nc$.
7. Update $UA_l = Q_{target}/(\overline{\Delta T}Nc)$.
8. Repeat steps 4 to 7 until the user-specified Maximum Solver Iterations have reached or the convergence criteria $(Q - Q_{target})/Q_{target} < \varepsilon$ is met, where ε is the user-specified convergence tolerance.
9. Store the converged UA_l and cold flow velocity $V_c = \dot{m}_c/(\rho_c A_c)$ where A_c is the heat exchanger cold stream inlet boundary area into the internal UAL table.

The above algorithm is invoked once before the first energy solver iteration, and the internally stored UAL table is reused in the subsequent iterations for computing the energy source term as in Eqn. (412). Whenever any UAG Table input is changed, the internal UAL table is recomputed. UA_l in Eqn. (412) is then interpolated based on the local cold stream velocity from this internal UAL table.

The maximum allowable heat transfer rate in a heat exchanger occurs when the temperatures of each stream do not change at all from inlet to outlet (maximum temperature difference). Then the heat transfers to or from each stream are:

$$Q_{hot} = \dot{m}_h Cp_h (T_{\text{hot,in}} - T_{\text{cold,in}}) \tag{416}$$

and

$$Q_{cold} = \dot{m}_c Cp_c (T_{\text{hot,in}} - T_{\text{cold,in}}) \tag{417}$$

These heat transfers would have to be equal, and therefore the maximum heat transfer is limited to the minimum of these two values: the one for which $\dot{m}Cp$ is a minimum. Usually, this minimum value is for the hot fluid since coolants have high specific heats.

The target heat transfer rate Q_{target} must be chosen such that it is always smaller than the minimum of Eqn. (416) and Eqn. (417), that is:

$$Q_{transfer} < min(\dot{m}_h Cp_h, \dot{m}_c Cp_c)(T_{\text{hot,in}} - T_{\text{cold,in}}) \tag{418}$$

Q Table Option

The Q Table option is identical in functionality to the UAG Table option, except that you specify the overall heat transfer Q instead of UA_g in the input table. UA_g is then computed using Eqn. (413).

Q Map Option

To compute the local UAL, a linear system of algebraic equations is solved iteratively by considering the heat balance between the cold and the hot streams, preserving the target heat exchange rate, for each row of the Q Map (that is, for each value of the cold stream and hot stream mass flow rates). This UAL table is computed internally only once initially (or whenever any input for the Q Map is changed). The UAL table is then used to determine the coefficients in the following curve using least squares regression.

$$\frac{1}{UAL} = C_0 + \frac{C_1}{V_c^{ec}} + \frac{C_2}{V_h^{eh}} \tag{419}$$

In Eqn. (419), C_0, C_1 and C_2 are the coefficients that are determined by fitting the UAL table data to the above curve using least squares regression, V_c and V_h are the cold flow and hot flow local velocities, and ec and eh are the cold flow and hot flow exponents (which by default are taken to be 0.8). After the coefficients are determined, Eqn. (419) is used to determine the local UAL for a given cold flow and hot flow local cell velocities in the energy equation.

Actual Method

Application of local heat exchanger energy sources into the hot/cold energy equations

At every solver iteration, the following steps are carried out to include the actual flow dual stream heat exchanger sources into the energy equations of the hot and cold fluid streams:

1. The UAL Table is computed.
2. For every cell i of the heat exchanger region, the UAL Table is interpolated to compute the local heat transfer coefficient ual_i, given the cold and hot mass fluxes that are entering cell i
3. The local energy source is computed as:

$$src_i = \frac{ual_i(T_{hi} - T_{ci})V_i}{\frac{1}{NC}\sum_{i=1}^{NC} V_i}$$

   where:

   - T_{hi} and T_{ci} are the temperatures of the hot and cold streams in cell i
   - V_i is the volume of cell i
   - NC is the total number of cells in the heat exchanger region
4. Apply src_i as an energy source to the cold stream energy equation and as an energy sink to the hot stream energy equation.

Computation of the UAL Table

For each entry of the user input table (UAG Table, Q Table, or Q Map), a corresponding UAL value is computed, thus generating the UAL Table.

The UAL value is computed using the following expression:

$$UAL = \frac{\Gamma}{\Delta T_{net}}$$

where

$$\Delta T_{net} = \frac{\sum_{i=1}^{NC} [(T_{hi} - T_{ci})V_i]}{\frac{1}{NC}\sum_{i=1}^{NC} V_i}$$

When a UAG Table is supplied:

$$\Gamma = UAG(\bar{T}_{in}^{hot} - \bar{T}_{in}^{cold})$$

where:

UAG is the user-table value, and

$$\bar{T}_{in} = \frac{\sum_f [T_f A_f]}{\sum_f A_f}$$

is the average inlet temperature computed at the inlet interface boundary of the heat exchanger region. The summation is over all the cell faces. T_f and A_f are the cell face values of temperature and surface area.

When a Q Table or a Q Map is supplied:

$$\Gamma = \frac{Q(\bar{T}_{in}^{hot} - \bar{T}_{in}^{cold})}{T_{user}^{hot} - T_{user}^{cold}}$$

where:

Q is the user-table value,

T_{user}^{hot} and T_{user}^{cold} are the user-supplied test rig inlet temperature values of the hot and cold streams respectively (these values are specified in STAR-CCM+ as properties of Q Table or Q Map), and

$$\bar{T}_{in} = \frac{\sum_f [T_f A_f]}{\sum_f A_f}$$

is the average inlet temperature computed at the inlet interface boundary of the heat exchanger region. The summation is over all the cell faces. T_f and A_f are the cell face values of temperature and surface area.

Q Map Option

The Q Map option of the Actual Flow Dual Stream heat exchanger model uses bilinear interpolation to compute UAL values from an UAL map. Consider the graphical representation of a computed UAL map:



The points of intersection between the constant m_{cold} and m_{hot} represent data in the computed UAL map. The UAL at a general point (m_c, m_h) is given by:

$$UAL = u_{c1} + q_c(u_{c2} - u_{c1})$$

where

$$u_{c1} = u_1 + q_h(u_2 - u_1)$$

$$u_{c2} = u_3 + q_h(u_4 - u_3)$$

and

$$q_c = \frac{m_c - m_c^{low}}{m_c^{high} - m_c^{low}}$$

$$q_h = \frac{m_h - m_h^{low}}{m_h^{high} - m_h^{low}}$$

u_1, u_2, u_3 and u_4 are the UAL values at points (m_c^{low}, m_h^{low}), (m_c^{low}, m_h^{high}), (m_c^{high}, m_h^{low}) and (m_c^{high}, m_h^{high}) respectively.

For example, consider the following UAL table for a given Q Map table:

m_{cold}	m_{hot}	UAL
1.0	4.0	10.0
2.0	4.0	20.0
1.0	8.0	30.0
2.0	8.0	40.0

The UAL for a cold mass flow rate of 1.5 kg/s and a hot mass flow rate of 5.0 kg/s is computed as:

$$UAL = \frac{10 + 20}{2} + \left(\frac{5.0 - 4.0}{8.0 - 4.0}\right) \times \left[\frac{30.0 + 40.0}{2} - \frac{10.0 + 20.0}{2}\right] = 20$$

Circumferential Heat Flux Averaging Formulation

STAR-CCM+ bases the averaging of heat flux around the circumference of a spinning body on either of two cases:

- A body spinning in an interfacing medium.
- A body spinning in a meshless region, interacting across the body boundary.

Both cases use the following conditions:

- The body in question is axisymmetric, with little or no circumferential variation in any properties, physical conditions, or geometry.
- Differences between the angular velocity of the axisymmetric body and any other bodies in the system are either very large or very small.
- Interface boundaries and wall boundaries between the rotating region and the surroundings are opaque. (The model does not allow radiative heat transfer between rotating and non-rotating regions.)

A Body Spinning in an Interfacing Medium

In this case, a solid body spins rapidly in an interfacing fluid medium. Convection, conduction, and radiation operate within in the fluid. Conduction and radiation operate within the solid. The interface is opaque to radiation.



Using averaging, the boundary energy balance for the fluid-side interface is:

$$(\overline{h_f T_f} - \overline{h_f} T_{fb}) + \overline{(T_f / R_f)} - \overline{(1/R_f)} T_{fb} + \overline{q_{af}} - \sigma \varepsilon_E^{total} T_{fb}^{\ 4} = -\frac{q_b}{2} + \frac{(T_{fb} - T_{sb})}{R_b} \tag{420}$$

where:

- The overline bar indicates averaging. The average value of a variable x is weighted by boundary face area A over a given axial strip n:

$$\bar{x} = \frac{\sum_{\text{strip n}} x_n \times A_n}{\sum A_n}$$

- h_f is the convection coefficient of the fluid side.
- T_f is the cell-center temperature on the fluid side.
- T_{fb} is the fluid-boundary temperature.
- T_{sb} is the solid-boundary temperature.
- R_f is the thermal resistance from the fluid cell center to the boundary.

- q_{af} is the radiation heat flux in the fluid due to absorption.
- σ is the Stefan-Boltzmann constant.
- $\varepsilon_E^{\mathrm{total}}$ is the total boundary emissivity.
- q_b is the heat flux in the interface.
- R_b is the thermal resistance of the boundary.

On the left-hand side of Eqn. (420):

- $(\overline{h_f T_f} - \overline{h}_f T_{fb})$ is the convection term.
- $\overline{(T_f / R_f)} - \overline{(1/R_f)} T_{fb}$ is the conduction term.
- $\overline{q_{af}} - \sigma \varepsilon_E^{\mathrm{total}} T_{fb}^{\ 4}$ is the radiation term.

On the right-hand side of Eqn. (420), the term $\frac{(T_{fb} - T_{sb})}{R_b}$ implies conduction through the boundary.

The boundary energy balance for the solid is:

$$\frac{1}{R_s}(T_s - T_{sb}) + q_{as} - \sigma\, \varepsilon_A^{\mathrm{total}} T_{sb}^{\ 4} = -\frac{q_b}{2} + \frac{(T_{fb} - T_{sb})}{R_b} \tag{421}$$

where:

- R_s is the thermal resistance from the solid cell center to the boundary.
- T_s is the cell-center temperature on the solid side.
- q_{as} is the radiation heat flux in the solid due to absorption.
- $\varepsilon_A^{\mathrm{total}}$ is the total boundary absorbtivity.

On the left-hand side of Eqn. (421):

- $\frac{1}{R_s}(T_s - T_{sb})$ is the conduction term.
- $q_{as} - \sigma \varepsilon_A^{\mathrm{total}} T_{sb}^{\ 4}$ is the radiation term.

For both equations, assume $T_f > T_{fb} > T_{sb} > T_s$.

A Body Spinning in a Meshless Region

In this case, a solid or fluid body spins rapidly in a region without mesh, with no medium between the body and its environment, represented in the following figure by an external boundary. The body and the external environment interact only by thermal radiation between their boundaries,

across the empty meshless region. Within the body, convection (if it is fluid), conduction, and radiation can all operate. Convection and radiation heat transfer operate between the boundaries of the body and the environment.



Using averaging, the boundary energy balance for the external side of the body is:

$$(\overline{h_\infty T_\infty} - \overline{h_\infty} T_{fb}) + \overline{q_{af}} - \sigma \varepsilon_E^{\text{total}} T_{fb}^{\ 4} = \frac{(T_{fb} - T_{sb})}{R_b} \tag{422}$$

where:

- h_∞ is the convection coefficient of the external side.
- T_∞ is the temperature on the external side.

Other variables are defined as in the previous section.

On the left-hand side of Eqn. (422):

- $(\overline{h_\infty T_\infty} - \overline{h_\infty} T_{fb})$ is the convection term.
- $\overline{q_{af}} - \sigma \varepsilon_E^{\text{total}} T_{fb}^{\ 4}$ is the radiation term.

The boundary energy balance for the internal side of the body is:

$$h_s(T_s - T_{sb}) + \frac{1}{R_s}(T_s - T_{sb}) + q_{as} - \sigma \varepsilon_A^{\text{total}} T_{sb}^{\ 4} = -\frac{(T_{fb} - T_{sb})}{R_b} \tag{423}$$

Variables are defined as in the previous section.

On the left-hand side of Eqn. (423):

- $h_s(T_s - T_{sb})$ is the convection term.

- $\frac{1}{R_s}(T_s - T_{sb})$ is the conduction term.
- $q_{as} - \sigma\varepsilon_A^{\text{total}} T_{sb}^{\ 4}$ is the radiation term.

Solid Energy Formulation

The governing equation for energy transport in a solid is given as follows:

$$\frac{d}{dt}\int_V \rho C_p T dV + \oint_A \rho C_p T \mathbf{v}_s \cdot d\mathbf{a} = -\oint_A \mathbf{q}'' \cdot d\mathbf{a} + \int_V s dV \tag{424}$$

where:

- ρ is the solid density.
- C_p is the specific heat.
- T is the temperature.
- $\mathbf{q}''$ is the heat flux vector.
- $\mathbf{v}_s$ is the solid convective velocity (only applicable when the region is a pure body of rotation).

Conducting Solid Shells Formulation

The modeling of heat transfer in thin solid structures is based on the modeling assumption that the temperature variation normal to the shell surface is piecewise-linear. The consequence of this assumption is that the relevant heat fluxes can be estimated having only one cell across the shell thickness. The energy transport in solid (in-plane conducting) shells is described with the same transport equation as the energy transport in solids.

Modeling Species

This chapter provides information about the chemical species models in STAR-CCM+. A species model is activated whenever a multi-component liquid or multi-component gas is chosen from the *Material* model section of the *Physics Model Selection* dialog. The specific models that are discussed are the Coupled Species model and the Segregated Species model. Some general advice on working with species is also given.

Reference sections are provided on the following:

- Setting Species Boundary and Initial Conditions
- The Segregated Species Solver
- Setting a species region source
- Formulation
- Field Functions Reference

What Is the Coupled Species Model?

The Coupled Species model is an extension of the Coupled Flow model. Together with the Coupled Energy model, it solves the conservation equations for mass, momentum, and energy of a multi-component fluid mixture simultaneously using a time- (or pseudo-time-) marching approach. For a mixture of `N` components, the Coupled Species model introduces `N-1` transport equations into the system. Together with global mass continuity, these equations provide a means for updating the field of `N` mass fractions defining the mixture composition. Inter-equation coupling is fully maintained between all equations in the system, including species transport with respect to momentum and energy, and vice versa.

For isothermal simulations of multi-component fluid mixtures, it is possible to omit the Coupled Energy model, in which case the Coupled Species model, together with the Coupled Flow model, solves the conservation equations for mass and momentum simultaneously.



See also: Segregated Species/Coupled Species Model Properties

What Is the Segregated Species Model?

The Segregated Species model solves the species continuity equations for a multi-component fluid mixture. For a mixture of N components, the Segregated Species model solves N-1 transport equations sequentially. Together with global mass continuity, these equations provide a means for updating the field of N mass fractions defining the mixture composition.



See also: Segregated Species/Coupled Species Model Properties.

Working with Species

The Species models are often used with combustion models, and so part of the process for setting up a simulation that involves a Species model also involves selecting a combustion model. In general, however, when a Species model is active it is necessary to:

- Specify mixture components (the different species) and associated properties
- Set species initial distribution
- Set species fractions at relevant boundaries

Specifying Mixture Components

A particular species is defined as one component in a multi-component gas mixture or liquid mixture. For a gas, the component collection is held within the Gas Mixture node of the Multi-Component Gas model node in the Models tree of the relevant physics continuum. These nodes are shown below. For liquids, this arrangement becomes the Liquid Mixture node of the Multi-Component Liquid model node.



The components of the gas or liquid mixture are defined by adding new components to the Gas Components or Liquid Components node. This can be done manually as described in the section, Adding Mixture Components, or automatically by importing Chemkin-formatted data.

Each new component's name and material properties are retrieved from the database with default property settings and values. You should review the default material property settings of the mixture as well as those of each new component. If necessary, change the settings of the mixture properties and then those of each new component accordingly (see Managing Multi-component Materials).

Note that:

- In order for a chemical species to be available for any given chemical reaction equation (see Defining Chemical Reactions) it must first be added to the Gas Components node.
- Some care should be taken in deciding the ordering of mixture components, as discussed in the section, Ordering of Mixture Components.
- If the database does not include a component that is needed for your combusting mixture, make a local copy of the STAR-CCM+ material database (file `props.mdb`) and add the required component and its physical properties to it.
- For an EBU or CFM combustion simulation, you most likely want to select Thermodynamic Polynomial Data (if it is available) as the method for the Specific Heat property of each mixture component. If you choose an alternative method for Specific Heat (for example, Constant or Polynomial in T), then you are required to provide Heat of Formation and Standard State Temperature values for that component as well.

Ordering of Mixture Components

In order to maximize numerical accuracy, it is recommended that the component expected to have the greatest overall mass fraction (sometimes referred to as the *background* component) be specified as the `N-th` component of the mixture. The reason for this is that both the Segregated Species model and the Coupled Species model solve species transport equations for all but the last (`Nth`) component of the mixture (so solving `N-1` equations). These `N-1` equations are used, along with the global mass continuity constraint that all mass fractions must sum to one, to update mass fractions of all components. This means that any discretization and/or machine round-off errors in mass fraction are absorbed by the `Nth` component. Usually these errors are small, particularly in relation to the mass fraction of the background component, and their impact is negligible. If, however, the `Nth` component were set to a minor component, such as a contaminant, the significance of these numerical errors could increase.

In the example gas mixture composition shown below, the O2 component is expected to have the largest overall mass fraction, and so has been specified last in the list of components.



Further Notes

- In a combusting mixture involving air, the background component is typically nitrogen (N2) and hence it should be the last component added to the Gas Components node.
- Mixture components can be reordered to get the desired background species (see Reordering Mixture Components).
- In a CFM combustion simulation involving a mixture of N components, the fuel species (also the first species) concentration is computed from the mean regress variable based on Eqn. (1085). Then the concentrations of species 2 to N-1 are computed based on the stoichiometric relationship from the specified reaction scheme. At last the concentration of species N is computed by 1 minus the sum of concentrations of species 1 to N-1. For the same reasons as outlined above, the Nth place should also be reserved for the component with the largest mass fraction in a CFM combustion simulation.

Segregated Species/Coupled Species Model Properties

The Segregated Species and Coupled Species models share certain properties.

Properties

Convection (Segregated Species only)	Selects the convection scheme that is used.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

Expert Properties

Flow Boundary Diffusion	Allows you to activate or deactivate diffusion at flow boundaries. Activated by default.	
Multi-component Diffusion	When activated, diffusive flux for a given component becomes a function of the mass fraction gradient of all other components. Deactivated by default.	
	Activated	Multi-component diffusion is used to calculate the diffusive flux in the species transport equation. When activated, you can only calculate Molecular Diffusivity using the Binary Diffusion Coefficients method or Kinetic Theory Diffusivity method.
	Deactivated	Fick's Law is used to calculate the diffusive flux in the species transport equations instead of multi-component diffusion.
Secondary Gradients (Segregated Species only)	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Species Boundary and Initial Conditions

The distribution of species within a continuum or across flow boundaries can be defined using either mass fraction or mole fraction. In STAR-CCM+, the same set of nodes are provided for flow boundaries and initial conditions.

Use the Species Specification node within the Physics Conditions folder of a boundary or initial condition node to choose between mass fraction or mole fraction as the method for specifying species distribution. The corresponding value node, Species Mass Fraction or Species Mole Fraction, is added to the relevant Physics Values folder.

On wall boundaries, the Wall Species Option allows you to define the wall as impermeable to species transfer, or one having a specified mass fraction or specified flux.

What Is the Species Specification?

This node is added to the Initial Conditions node of physics continua, and to the Physics Conditions node of flow boundaries, after you have selected components for the mixture.



The single Method property of this node is used to switch between specifying the species distribution in terms of mass fraction or mole fraction.

Species Specification Properties

Method — Defines the basis on which initial species distribution is defined.

- Mass fraction — Mixture components are defined in terms of mass fraction. When this option is selected, a Species Mass Fraction sub-node is added under the Initial Conditions node (or the Physics Values node for a flow boundary).
- Mole fraction — Mixture components are defined in terms of mole fraction. When this option is selected, a Species Mole Fraction sub-node is added under the Initial Conditions node (or the Physics Values node for a flow boundary).

Using Species Mass Fraction

The mass fraction of a component is the ratio of its mass over the total mass of the mixture, and is represented by the Species Mass Fraction node. It is a scalar quantity that is used for both initial and boundary conditions, and is defined by one of the following options:

- Constant specifies a fixed value for all mixture components. The required values are entered in the *Properties* window as a scalar array profile in the order in which the components were defined.



Alternatively, you may click the property customizer box, marked with an ellipsis and located to the right of the Value property array. This opens an input editor, shown below, which allows a more convenient way of entering these values as it lists component names side-by-side with the specified values.



- Composite specifies the distribution of each component separately; distributions for each component are entered as scalar profiles.



Both initial and boundary values that you specify are normalized prior to being used so that they sum to 1.0. Before normalizing, any value entered larger than 1.0 is set to 1.0. After this, normalization will be applied across all fractions. For example, [1.0, 2.0, 3.0] becomes [1.0, 1.0, 1.0], which normalizes to [0.333, 0.333, 0.333].

Using Species Mole Fraction

The mole fraction of a species is the ratio of its molar amount over the mixture molar amount, and is represented by the Species Mole Fraction node. It is a scalar quantity that is used for both initial and boundary conditions, and is defined by one of the following options:

- Constant specifies a fixed value for all mixture components. The required values are entered in the *Properties* window as a scalar array profile in the order in which the components were defined.



As for the Species Mass Fraction, the property customizer can be accessed by clicking the ellipsis to the right of the Value property array. This opens an input editor that identifies the different components in the array.

- Composite specifies the distribution of each component separately; distributions for each component are entered as scalar profiles.



What Is the Wall Species Option?

For species transfer, this option defines a wall boundary as impermeable or as having a specified value or flux. The Coupled Species model or the Segregated Species model activates it.

Limitations:

- The bulk velocity of the wall species influx or outflux is negligible.
- The gain or loss of momentum due to wall species influx or outflux is negligible.

Post-Processing

You can monitor the species flux at the wall with the Boundary Species Flux field function. To activate this field function, enable the Temporary Storage Retained expert property of the Segregated Species solver or the Coupled Implicit solver. This feature lets you verify your specification of the boundary conditions.

Properties of the Wall Species Option

Method	Defines the permeability of the boundary.	
	Impermeable	Specifies the boundary as impermeable.
	Specified Flux	Specifies a species flux at the boundary. The energy that is associated with the wall species influx or outflux is appropriately added. Adds a Species Flux node to the Physics Values node.
	Specified Value	Specifies a species mass fraction at the boundary. Adds a Species Mass Fraction node to the Physics Values node.

Using the Species Flux

The species flux is the mass flow rate of species added or removed per unit area. The Species Flux node is added when the Method property of the Wall Species Option boundary condition is set to `Specified Flux`. It is entered as a scalar array profile.

What Is the Segregated Species Solver?

This object controls the solution update for the Segregated Species model.

Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.9.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

Setting a Species Region Source

The Species Source is a region source that is activated by the Coupled Species or the Segregated Species model. It specifies the rate at which a unit mass of the species is added or removed per unit volume of the region

(Mass/Time/Volume). In the Physics Conditions of a region, you can enable user-specified species sources by ticking the checkbox of the Species Source Term property.

To define a species source within a region:

- Select the Species Source Option node within the Physics Conditions folder for the region.
- Tick the Species Source Term property in the *Properties* window.



One additional node appears in the Physics Conditions folder for the region: the Species Source Temperature Option, and two nodes are added to the Physics Values folder for the region: the Species Sources and Species Source Jacobians nodes. Use these nodes to set the value and derivative of the source you wish to set.

Species Source Temperature Option

When you add a species source to a simulation, ensure that an equivalent source of energy is added or removed with the species mass. STAR-CCM+ can provide this energy source automatically by calculating the energy that is associated with the species source at a particular temperature. To have STAR-CCM+ add the equivalent energy automatically:

- Select the Species Source Temperature Option node. In the *Properties* window, set the Method property to Specified Value.



An additional node, the Species Source Temperature, is added to the Physics Values folder for the region, where the temperature of the user species source can be specified.

If you leave the Method property as None, you must supply your own energy source using the Energy Source Option on the region.

Species Sources

This node can be used to specify the species source, S_{Y_i}, of each of the defined species. They can be specified as constant or composite, since this value is entered as a scalar array profile.

Species Source Jacobians

This node can be used to specify the Jacobian, $d(S_{Y_i})/d(Y_i)$, for each species. This value is entered as a scalar array profile.

Species Source Temperature

This node enables you to specify the temperature of the user species sources. STAR-CCM+ adds an amount of energy corresponding to the mass of the user species sources and this temperature. The default option uses a constant temperature value set in Kelvin. The alternative options are described in the scalar profile properties section.

Species Transport Formulation

The transport equation for the mass fraction Y_i of a species is given as:

$$\frac{\partial}{\partial t}\int_{\tilde{V}} \rho Y_i d\tilde{V} + \oint_A \rho Y_i(\mathbf{v} - \mathbf{v}_g) \cdot d\tilde{\mathbf{a}} = \oint_A \left[J_i + \frac{\mu_t}{\sigma_t}\nabla Y_i \right] \cdot d\tilde{\mathbf{a}} + \int_{\tilde{V}} S_{Y_i} d\tilde{V} \tag{425}$$

where:

- i is the component index.
- ρ is the overall density.
- $d\tilde{V} = \alpha_i \chi dV$, where α_i is the volume fraction of phase i (set to 1 for single phase flows) and χ is the void fraction.
- $d\tilde{\mathbf{a}} = \alpha_i \chi d\mathbf{a}$
- $\mathbf{v}$ is the velocity.
- $\mathbf{v}_g$ is the grid velocity.
- μ_t is the turbulent dynamic viscosity.
- σ_t is the turbulent Schmidt number.
- S_{Y_i} is a user specified region source term for species i.
- J_i is the diffusive flux.

The turbulent Schmidt number σ_t is assumed to have a constant default value of 0.9. This value can be modified if necessary.

When multi-component diffusion is deactivated, Fick's Law is used to calculate the diffusive flux J_i, given by:

$$J_i = \rho D_{i,m} \nabla Y_i \tag{426}$$

where D_m is the molecular diffusivity.

When multi-component diffusion is activated, the scalar diffusivity of Fick's Law is replaced by a matrix, and the diffusive flux for component i becomes a function of the mass fractions for all components:

$$J_i = \rho \sum_{j=1}^{N} D_{i,j} \nabla Y_j \tag{427}$$

$D_{i,j}$ represents the multi-component diffusion coefficients which are calculated using the Maxwell-Stefan equations. These equations define the diffusive flux implicitly as a function of mole fraction gradients, which are given as:

$$\nabla x_i = \frac{M_w}{\rho} \sum_{j=1, j \neq i}^{N} \frac{x_i J_j}{M_j \mathfrak{D}_{i,j}} - \frac{x_j J_i}{M_i \mathfrak{D}_{i,j}} \tag{428}$$

$\mathfrak{D}_{i,j}$ represents the binary diffusion coefficient between component i and j. Writing these equations in matrix form gives:

$$[B] \nabla Y = [A] J \tag{429}$$

where $[B]$ represents the mapping from mass fraction gradients to mole fraction gradients and $[A]$ represents the Maxwell-Stefan equations. By inverting $[A]$ and multiplying by the matrix $[B]$, the multi-component diffusion coefficients are calculated.

$$D_{i,j} = \rho \sum_{k=1}^{N-1} A_{i,k}^{-1} B_{k,j} \tag{430}$$

For *N* components, STAR-CCM+ solves *N-1* transport equations. The mass fraction of the *N*th component is found from the global mass continuity constraint that all mass fractions must sum to 1.

Species Transport Field Functions Reference

The following field functions are made available when a Multi-Component Material has been specified.

Effective Mass Diffusivity of *Species*	The effective mass diffusivity, $D_{\text{eff}} = D_m + \mu_t/(\rho\sigma_t)$, where D_m is the molecular diffusivity, μ_t is the turbulent viscosity, and σ_t is the turbulent Schmidt number. For laminar flows, $\mu_t = 0$, and $D_{\text{eff}} = D_m$.
Mass Fraction of *Species*	The mass fraction, $Y_i = m_i/m_m$, of species *i* in mixture *m*.
Molar Concentration of *Species*	The molar concentration, $C_i = \rho_m Y_i/W_i$, of species *i* in a mixture *m*, where ρ_m is the mixture density, Y_i is the species mass fraction, and W_i is the species molecular weight.
Mole Fraction of *Species*	The mole fraction, $X_i = Y_i W_m/W_i$, of species *i* in a mixture *m*, where Y_i is the species mass fraction, W_m is the molecular weight of the mixture, and W_i is the molecular weight of the species.

Modeling Turbulence and Transition

The turbulence models in STAR-CCM+ are responsible for providing closure of the governing equations in turbulent flows.

This section first provides general overviews of:

- Selecting a turbulence modeling approach
- Selecting a transition modeling approach
- Using Reynolds-Averaged Navier-Stokes (RANS) turbulence models
- Using large eddy simulation (LES)
- Using detached eddy simulation (DES)
- Using synthetic eddy method (SEM) for initialization and inflow
- Using wall treatment models

Details of specific turbulence models are also given for:

- Spalart-Allmaras (RANS and DES)
- K-Epsilon
- K-Omega
- Reynolds Stress Transport
- Smagorinsky Subgrid Scale Turbulence
- Dynamic Smagorinsky Subgrid Scale Turbulence
- WALE Subgrid Scale Turbulence

The following approaches to modeling transition are described:

- Turbulence Suppression

- Gamma ReTheta

Information is included about an algebraic formulation for the turbulent heat flux.

Reference information is also provided for the following objects shared by turbulence models:

- Shared conditions and values
- Field functions

Selecting a Turbulence Modeling Approach

Three basic approaches to modeling turbulence are available in STAR-CCM+:

- Models that provide closure of the Reynolds-Averaged Navier-Stokes (RANS) equations.
- Large eddy simulation (LES).
- Detached eddy simulation (DES).

Most simulations will rely on the first approach. The second two approaches (LES and DES) are best used after carefully reviewing the applicable literature to gain confidence that the grid resolution requirements can be met and that the computational costs incurred by resolving the small time and length scales are indeed justified.

Two methods are available to select the turbulence modeling approach:

- Using auto-select
- Without using auto-select

Using Auto-Select

This section describes how to select a turbulence modeling approach using automatic selection of recommended models. This is the default mode of selection, and is indicated by the Auto-select recommended physics models checkbox being ticked.

The selection of the Turbulent radio button in the Viscous Regime group box will automatically select turbulence models that provide closure of the Reynolds-Averaged Navier-Stokes (RANS) equations as a turbulence

modeling approach. These modeling choices will be displayed in the *Reynolds-Averaged Turbulence* group box as shown below, although the choice of *Enabled Models* might differ.



While the end result will depend on the choice made above, a complete selection will be made, including the variant of the model and the wall treatment model. The result for the K-Epsilon model is shown below.



Without Using Auto-Select

This section describes how to select a turbulence model that supports DES without using the automatic selection of recommended models feature. This selection method should be used if large eddy simulation or detached eddy simulation is desired. Note too that these two turbulence modeling approaches are only available if the following prerequisite models are selected:

- the Three-Dimensional model;
- the Implicit Unsteady model;
- the Segregated Flow model.

To select the turbulence modeling approach, first select the viscous regime. In the *Model Selection* dialog, make sure that the Auto-select recommended physics models checkbox is cleared. Then, in the *Turbulence* group box, select the radio button corresponding to the appropriate turbulence method. The availability of methods depends partly on your selection of a time model.



Depending on your selection, you will now be able to select:

- a (RANS) turbulence model;

- a subgrid scale model for large eddy simulation (LES); or
- a turbulence model for detached eddy simulation.

The result for the SST K-Omega model is shown below.



Selecting a Transition Modeling Approach

This section describes the use of the transition models in STAR-CCM+, specifically the following aspects:

- What is transition?
- What is a transition model?
- Selecting a transition model

What Is Transition?

The term "transition" refers to the phenomenon of laminar to turbulence transition in boundary layers. Three primary modes of transition are typically involved:

1. Natural transition, in which a laminar boundary layer subjected to weak disturbances becomes linearly unstable beyond a critical Reynolds number at which Tollmien-Schlichting waves start to grow.
2. Bypass transition, the process of transition in response to large disturbances outside the boundary layer, typically free stream turbulence levels in excess of 1%.
3. Separation induced transition, in which separation of the laminar boundary layer induces transition. The laminar boundary layer often reattaches in response to the enhanced mixing caused by the turbulent flow, forming laminar separation bubble upstream of the transition location.

What Is a Transition Model?

A transition model is a model that can be combined with a turbulence model to predict the onset of transition in a turbulent boundary layer. Transition models are necessary because turbulence models themselves cannot be relied upon to predict the onset of transition.

In STAR-CCM+, two methods are currently available to account for transition:

1. The Turbulence Suppression model, a model that allows the user to mimic the effect of transition by suppressing the turbulence in a certain pre-defined region.
2. The Gamma ReTheta Transition model, a correlation-based model that is coupled with the SST K-Omega model to provide a true predictive capability.

Selecting a Transition Model

The transition model that becomes available depends on your choice of RANS turbulence model:

- If you use the SST K-Omega model, both the Turbulence Suppression model and the Gamma ReTheta Transition model become available.
- With other turbulence models, Turbulence Suppression is the transition option available to select.

Using Reynolds-Averaged Navier-Stokes Turbulence Models

This section deals with models that provide closure for the Reynolds-Averaged Navier-Stokes (RANS) equations. In this context, unless otherwise stated, it should be understood that the term "turbulence model" refers to RANS closure models. The following topics are covered:

- What is a turbulence model?
- How to decide on a turbulence model
- Setting model parameters
- Setting solver parameters
- Troubleshooting
- Turbulent flow initialization

Selection of a turbulence model is covered in a separate section.

What Is a Turbulence Model?

To obtain the Reynolds-Averaged Navier-Stokes (RANS) equations, the Navier-Stokes equations for the instantaneous velocity and pressure fields are decomposed into a mean value and a fluctuating component. The averaging process may be thought of as time averaging for steady-state situations and ensemble averaging for repeatable transient situations. The resulting equations for the mean quantities are essentially identical to the original equations, except that an additional term now appears in the momentum transport equation. This additional term is a tensor quantity, known as the Reynolds stress tensor, which has the following definition:

$$\mathbf{T}_t \equiv -\rho\overline{\mathbf{v}'\mathbf{v}'} = -\rho\begin{bmatrix} \overline{u'u'} & \overline{u'v'} & \overline{u'w} \\ \overline{u'v'} & \overline{v'v'} & \overline{v'w'} \\ \overline{u'w'} & \overline{v'w'} & \overline{w'w} \end{bmatrix} \tag{431}$$

The challenge is thus to model the Reynolds stress tensor $\mathbf{T}_t$ of Eqn. (431) in terms of the mean flow quantities, and hence provide closure of the governing equations. Two basic approaches are used in STAR-CCM+:

- Eddy viscosity models; and
- Reynolds stress transport models.

It is widely acknowledged that turbulence models are inexact representations of the physical phenomena being modeled, and no single turbulence model is the best for every flow simulation. Consequently, it is necessary to provide a suite of models that reflect the current state of the art.

What is an Eddy Viscosity Model?

Eddy viscosity models use the concept of a turbulent viscosity μ_t to model the Reynolds stress tensor as a function of mean flow quantities. The most common model is known as the Boussinesq approximation:

$$\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I} \tag{432}$$

where $\mathbf{S}$ is the strain tensor:

$$\mathbf{S} = \frac{1}{2}(\nabla \mathbf{v} + \nabla \mathbf{v}^T) \tag{433}$$

and k is the turbulent kinetic energy.

While some simpler models rely on the concept of mixing length to model the turbulent viscosity in term so mean flow quantities (similar to the Smagorinsky subgrid scale model used in LES), the eddy viscosity models in STAR-CCM+ solve additional transport equations for scalar quantities that enable the turbulent viscosity μ_t to be derived. These include Spalart-Allmaras, K-Epsilon and K-Omega models.

Since the assumption of Eqn. (432) that the Reynolds stress tensor is linearly proportional to the mean strain rate is known to be flawed, some two-equation models extend the linear approximation to include non-linear constitutive relations.

What is a Reynolds Stress Transport Model?

Reynolds stress transport models, also known as second-moment closure models, solve transport equations are solved for each component of the Reynolds stress tensor.

How Do I Decide on a Turbulence Model?

There are four major classes of turbulence models currently in STAR-CCM+. This section presents broad guidelines as to the applicability of each of these. Further guidance on selecting the specific model variants can be found within the sections that provide details about the models.

- Spalart-Allmaras models are a good choice for applications in which the boundary layers are largely attached and separation is mild if it occurs. Typical examples would be flow over a wing, fuselage or other aerospace external-flow applications.

  The Spalart-Allmaras models for RANS equations are not suited to flows that are dominated by free-shear layers, flows where complex recirculation occurs (particularly with heat transfer), or natural convection. This statement does not apply to the Spalart-Allmaras detached eddy model.

- K-Epsilon models provide a good compromise between robustness, computational cost and accuracy. They are generally well suited to industrial-type applications that contain complex recirculation, with or without heat transfer.

- K-Omega models are similar to K-Epsilon models in that two transport equations are solved, but differ in the choice of the second transported turbulence variable. The performance differences are likely to be a result of the subtle differences in the models, rather than a higher degree of complexity in the physics being captured. These models have seen most application in the aerospace industry. Therefore, they are recommended as an alternative to the Spalart-Allmaras models for similar types of applications.

- Reynolds stress transport models are the most complex and computationally expensive models offered in STAR-CCM+. They are recommended for situations in which the turbulence is strongly anisotropic, such as the swirling flow in a cyclone separator.

Setting Model Parameters

To set the properties for a turbulence model, select the model in the Models node under the required continuum node. The properties are available under the Properties and Expert categories.



Each of the turbulence models will have different properties, but all have the Convection property. This causes the convection scheme to switch between first-order and second-order upwind.

The properties in the *Expert* category should not generally be altered.

Setting Solver Parameters

Every turbulence model has two solvers associated with it: a turbulence solver that controls the under-relaxation and multigrid parameters for the turbulence transport equations, and a turbulent viscosity solver that allows the parameters governing the update of turbulent viscosity to be set.

Setting Turbulence Solver Under-Relaxation

To set the under-relaxation factor for the turbulence solver, select the turbulence solver node in the Solvers node. The example below shows the selection of the K-Epsilon Turbulence solver . The Under-Relaxation Factor property can be entered in the *Properties* window.



The default under-relaxation factor should be adequate for most situations. These are shown below for the various turbulence solvers:

Turbulence Solver		Under-Relaxation Factor
Spalart-Allmaras		0.7
K-Epsilon		0.8
K-Omega		0.8
Reynolds-Stress Transport		0.6

Setting Turbulence Solver Algebraic Multigrid Parameters

To set the algebraic multigrid parameters for the turbulence solver, open the turbulence solver node in the Solvers node. Then also open the node for the AMG Linear Solver. The example below shows the selection corresponding to the K-Epsilon Turbulence Solver. The algebraic multigrid parameters can be entered using the *Properties* window.



Setting Turbulent Viscosity Update Parameters

To set the parameters governing the update of turbulent viscosity, select the turbulent viscosity solver node in the Solvers node as shown below. Two properties are available in the *Properties* window.



The Maximum Ratio property sets the maximum allowable ratio of turbulent to laminar viscosity.

The Under-Relaxation Factor property governs what portion of the newly computed turbulent viscosity field is used to update the turbulent viscosity. The default value of 1 implies that the turbulent viscosity will be entirely computed from the new field.

Troubleshooting

Solution of the transported turbulence variables diverges

The residual for any of the transported turbulence variables increases significantly after the startup transient, indicating that the solution of the turbulence field is diverging.

- Make sure the velocity field is not diverging and hence causing the turbulence model to diverge.
- Slightly reducing the turbulence solver under-relaxation factor (say to between 0.3 and 0.5) might stabilize the solution.
- In situations where mesh quality is impeding convergence, using the first-order convection scheme might help.

Solution of the transported turbulence variables converges slowly.

The flow residuals are converging well, but the residuals for the transported turbulence variables are not decreasing as fast.

- You can increase the turbulence solver under-relaxation factor if you are confident that the solution is otherwise progressing well, but a value of more than 0.9 is not recommended, since it might result in solution divergence.

Unrealistically high turbulent viscosity impedes flow solution

During the first few iterations, unrealistically high turbulent viscosity values are impeding the flow solution.

- Reducing the turbulence solver under-relaxation factor will retard the update of turbulent viscosity.
- It is possible to make this factor quite small (say 0.01) to cause only a minimal update, but this would typically be done for a few iterations only.

Warnings of limited turbulent viscosity occur

Repeated warnings occur that the turbulent viscosity is being limited in some cells.

- This is both common and acceptable during the first few iterations due to startup transients.
- If the warnings continue to occur as the flow solution nears convergence, it is possible to increase the Maximum Ratio to prevent the warnings.
- Generally one would only increase the Maximum Ratio if the simulation has a very high Reynolds number and no errors seem to be present in the simulation setup. The turbulent viscosity ratio is proportional to a turbulent Reynolds number, where the velocity scale is typically given as $k^{1/2}$ and the length scale as $(k^{3/2})/\varepsilon$ (for a model based on the eddy viscosity concept). If you expect very high turbulent Reynolds numbers (which occur when the mean flow Reynolds number is high, such as in a planetary boundary layer), you could try increasing the Maximum Ratio limit by one or two orders of magnitude. If you do not expect high turbulent Reynolds numbers, it is likely that another problem is present in your simulation, such as an incorrect boundary condition, or a mesh that hasn't been scaled correctly.

Initializing Turbulent Flow

When initializing turbulent flows, three turbulence specification methods are typically presented (depending on the turbulence model):

- Direct specification of the turbulence variables (such as turbulent kinetic energy and turbulent dissipation rate)
- Specification of turbulence intensity and turbulent length scale
- Specification of turbulence intensity and turbulent viscosity ratio

To avoid divergence during the initial iterations, it helps to have reasonable initial values. Admittedly, it might be hard to find a single number that describes the turbulence adequately. However, a good rule of thumb is that the turbulence intensity should be in the range $0.01 < I < 0.1$ and the turbulent viscosity ratio should be in the range $10 < \mu_t/\mu < 100$. These values mean that there is some additional viscosity over and above the laminar viscosity, and the dissipation scales are not unreasonable.

Using Large Eddy Simulation

This section explains the following:

- What is large eddy simulation (LES)?
- Selecting large eddy simulation
- What is a subgrid scale model?
- How to select a subgrid scale model
- Understanding the LES wall treatment

Also included are reference sections on turbulence specification and field functions as they relate to large eddy simulation.

What Is Large Eddy Simulation?

Large eddy simulation (LES) is an inherently transient technique in which the large scales of the turbulence are solved for, and the small-scale motions are modeled. One justification is that by modeling "less" of the turbulence and explicitly solving for more of it, the error in the turbulence modeling assumptions is not as consequential. Furthermore, it is hypothesized that the smaller eddies are self-similar and thus lend themselves to simpler and more universal models. The downside of the approach is the computational expense, which although less than direct numerical simulation, is still excessive.

The equations that are solved for large eddy simulation are obtained by a filtering rather than an averaging process. The filtered equations may be rearranged into a form that looks identical to unsteady RANS equations. However, the turbulent stress tensor, $\mathbf{T}_t$, now represents the subgrid scale stresses. This tensor is modeled using the Boussinesq approximation as follows:

$$\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I} \tag{434}$$

where μ_t and k are the subgrid scale turbulent viscosity and kinetic energy, respectively. $\mathbf{S}$ is the strain rate tensor that is computed from the resolved velocity field:

$$\mathbf{S} = \frac{1}{2}(\nabla \mathbf{v} + \nabla \mathbf{v}^T) \tag{435}$$

A subgrid scale model is required for the subgrid scale viscosity, μ_t.

Mean-flow quantities must be obtained by gathering statistics either over a long physical time, and/or from a homogeneous spatial coordinate. Field function monitors are provided in STAR-CCM+ to monitor the field statistics.

Inflow Boundaries and Initialization

One of the difficulties in large eddy simulations is representing the flow upstream of the flow domain, as an inadequate amount of information introduces sources of error. In STAR-CCM+ the Synthetic Eddy Method is implemented to automatically provide turbulent eddies across inflow boundaries and provide an initial perturbated flow field.

Working with Subgrid Kinetic Energy

Subgrid kinetic energy is a physical by-product of the STAR-CCM+ flow. LES always uses a subgrid scale model—Smagorinsky or WALE—which in turn directly computes turbulent viscosity. Within this computation is an implicit level of subgrid kinetic energy.

You may need to post-process this function only in certain kinds of STAR-CCM+ cases. For example, in Lagrangian simulations that include turbulent dispersion, subgrid kinetic energy is contained in turbulent fluctuations that the calculation has not taken into account. Therefore this field function can provide an approximate analysis of these unresolved fluctuations.

The values for subgrid kinetic energy are affected by the C_t parameter in the Smagorinsky expert properties or WALE expert properties.

Working with the Smagorinsky Subgrid Scale Wall Damping Function

The Smagorinsky subgrid scale model can give inaccurate scaling at walls, so the Van Driest damping function (Eqn. (762)) modifies the length scale closer to the wall. Visualizing this factor as a field function (SGS Wall Damping Factor) allows you to see the extent of the domain that is affected by this damping function.

Selecting Large Eddy Simulation

In order to select the large eddy simulation model in the *Physics Model Selection* dialog for a physics continuum, it is first necessary to deactivate the Auto-select recommended Physics models checkbox before choosing the

Turbulent option from the *Viscous Regime* group box. Some model selections must be made before the *Viscous Regime* group box becomes visible. The following graphic illustrates this operation:



The *Physics Model Selection* dialog displays the *Turbulence* group box on completion. Choose the Large Eddy Simulation model from this group as shown:



The dialog is updated to present the *Subgrid Scale Turbulence* group box.



Activating the Auto-select recommended Physics models checkbox results in an automatic selection of the options marked by an asterisk in the *Subgrid Scale Turbulence* group box and the *LES Wall Treatment* group box (that appears after choosing the subgrid model). If this route is not followed, it is necessary to choose one of the options in the *Subgrid Scale Turbulence*

group box, and one option from the *LES Wall Treatment* group box. In the dialog that is shown below, the WALE Subgrid Scale model has been selected with the All y+ Wall Treatment option:



What Is a Subgrid Scale Model?

Large eddy simulation requires closure of the filtered Navier-Stokes equations by modeling the subgrid scale stress tensor using the Boussinesq approximation (see Eqn. (432)). A subgrid scale model is required to model the subgrid scale viscosity, μ_t. Currently, three subgrid scale models are available in STAR-CCM+:

1. Smagorinsky Subgrid Scale model. This is the original and most basic subgrid scale model, upon which many more advanced models are based. One weakness of this model is that an ad hoc Van Driest damping function is required for proper results in wall-bounded flows. The algorithm that is used to implement the Van Driest damping has the potential to incur a significant parallel overhead (both memory and communication) when the geometry is complex. Therefore, this model should be avoided for large-scale computations on complex geometries. It is, however, useful for reference simulations, since it is the original algebraic subgrid scale model formulation.

2. Dynamic Smagorinsky Subgrid Scale model. This model computes a local time-varying constant. This dynamic variation of the constant allows the model to obtain the correct result for wall-bounded flows without the use of damping functions.
3. WALE Subgrid Scale model. This is a more modern subgrid scale model that uses a novel form of the velocity gradient tensor in its formulation. Its advantage is that it does not require any form of near-wall damping, so it is more suited to large-scale parallel computation on complex geometries. Experience within STAR-CCM+ has shown this model to be reasonably insensitive to the value of the model coefficient. It is, therefore, the recommended subgrid scale model.

Selecting a Subgrid Scale Model

This section describes how to select a subgrid scale model. Two methods are described:

- Using auto-select
- Without using auto-select

Using Auto-Select

This section describes how to select a subgrid scale model using automatic selection of recommended models. When selecting a turbulence modeling approach, ensure that the Auto-select recommended physics models checkbox is ticked before choosing Large Eddy Simulation. This causes the WALE Subgrid Scale model to be selected, since it is the recommended subgrid

scale modeling option in STAR-CCM+.



Without Using Auto-Select

This section describes how to select a subgrid scale model without using automatic selection of recommended models. When selecting a turbulence modeling approach, ensure that the Auto-select recommended physics models checkbox is not ticked before choosing Large Eddy Simulation. You can

now select the radio button corresponding to the Smagorinsky Subgrid Scale model, the Dynamic Smagorinsky Subgrid Scale model, or the WALE Subgrid Scale model.



Understanding the LES Wall Treatment

Resolving the viscous sublayer with LES is potentially expensive, since unlike RANS simulations, it is not sufficient to stretch the mesh only normal to the wall. The streaky structures that develop in the near-zero wall region also require adequate mesh definition in the spanwise and streamwise directions. Therefore, it may be desirable to avoid having to resolve the viscous sublayer.

To facilitate this choice, two wall treatments are available:

- Low y+ wall treatment

- All y+ wall treatment

What Is the Low y+ Wall Treatment?

The Low y^+ Wall Treatment assumes that the viscous sublayer is resolved and thus wall laws are not needed. The wall shear stress is computed as it would be in a laminar flow or in a direct numerical simulation.

What Is the All y+ Wall Treatment?

The All y^+ Wall Treatment makes no assumption about how well the viscous sublayer is resolved. By using a blended wall law to estimate shear stress, the result is similar to the Low y+ wall treatment if the mesh is fine enough. If the mesh is coarse enough ($y^+ > 30$), the wall law is equivalent to a logarithmic profile.

Synthetic Turbulence Specification

Synthetic turbulence specification is available as a boundary condition and as an initial condition. It allows you to specify the method for defining the initial turbulent velocity fluctuations in a continuum, at a velocity inlet, free stream, or mass flow inlet boundaries.

Properties

Method	Choose the method for defining turbulence.	
	No fluctuations	Leaves the specified mean velocity unperturbed.
	Intensity + Length Scale	Calculation of an isotropic field of velocity fluctuations that are superimposed on the specified mean velocity. A Turbulence Intensity node and a Turbulent Length Scale node are added to the Values node.
	Reynolds Stress + Length Scale	Calculation of an anisotropic field of velocity fluctuations that are superimposed on the specified mean velocity. Nodes for Turbulent Length Scale, UU Reynolds stress, VV Reynolds stress, WW Reynolds stress, UV Reynolds stress, UW Reynolds stress, and VW Reynolds stress are added to the Values node.

LES Field Functions Reference

The following primitive field function is made available to the simulation when you activate an aeroacoustics or combustion model along with WALE or Smagorinsky turbulent scales. This object has typical base properties.

SGS Turbulent Kinetic Energy	Subgrid kinetic energy, k_{SGS} in Eqns. (765) and (779).

Using Detached Eddy Simulation

This section explains the following:

- What is detached eddy simulation (DES)?
- How to select a DES model
- Using the hybrid second-order upwind/central scheme

What is Detached Eddy Simulation?

The unsteady RANS equations are applicable to transient situations where the unsteadiness is either imposed, such as by a time-varying boundary condition, or is inherent, such as the vortex shedding in a massively separated flow. In the latter case, transient simulations often yield better results than attempting use a steady-state approach. However, successful unsteady RANS simulations require that the time scales of the turbulence be disparate from the mean-flow unsteadiness. Furthermore, the limitations of the turbulence model may preclude good unsteady results.

Detached eddy simulation (DES) is a hybrid modeling approach that combines features of Reynolds-Averaged (RANS) simulation in some parts of the flow and large eddy simulation (LES) in others. DES turbulence models are set up so that boundary layers and irrotational flow regions are solved using a base RANS closure model. However, the turbulence model is intrinsically modified so that, if the grid is fine enough, it will emulate a basic LES subgrid scale model in detached flow regions. In this way, one gets the best of both worlds: a RANS simulation in the boundary layers and an LES simulation in the unsteady separated regions.

While the technology holds great promise for certain types of simulations, it must be cautioned that DES is not the answer to all turbulence modeling problems. The creation of suitable grids is something of an art, and an unsteady simulation is required, with statistics gathered to obtain a mean field as in a LES. Field function monitors are provided in STAR-CCM+ to monitor the field statistics.

Selecting a DES Model

This section describes how to select a turbulence model. Two methods are described:

- Using auto-select
- Without using auto-select

Using Auto-Select

This section describes how to select a turbulence model that supports DES using automatic selection of recommended models. When selecting a turbulence modeling approach, ensure that the Auto-select recommended physics models checkbox is ticked before choosing Detached Eddy Simulation. This will cause the Spalart-Allmaras DES model to be selected, since it is the recommended detached eddy simulation modeling option.



Without Using Auto-Select

This section describes how to select a turbulence model that supports DES without using automatic selection of recommended models. When selecting a turbulence modeling approach, ensure that the Auto-select recommended physics models checkbox is not ticked before choosing Detached Eddy

Simulation. You can now select the radio button corresponding to the Spalart-Allmaras DES model or the SST (Menter) K-Omega Detached Eddy model.



Using the Hybrid Second-Order Upwind/Central Difference Scheme

As discussed in the section on discretizing convection, the second-order upwind is the most accurate scheme for steady-state calculations, while the kinetic energy-preserving properties of the central differencing scheme make it appropriate for large eddy simulation (LES). When performing detached eddy simulation (DES), the Segregated Flow model makes available a hybrid second-order upwind/central scheme where the blending factor is computed according to the features of the flow field.

It should be noted that, like DES itself, this feature is not a mature development, and some experimentation might be required to tune it. It is suggested that the simulation be initiated using the second-order upwind scheme. Then the coefficient CdesTimeLim can be adjusted so that the field function quantity DES Upwind Blending Factor looks like it is indicating unity for areas such as attached shear layers and irrotational free-stream flow, and zero for areas where LES is required such as massively separated regions. Once the field function looks correct, the hybrid second-order upwind/central scheme can be activated.

Using the Synthetic Eddy Method for Initialization and Inflow

This section explains the following:

- What is the synthetic eddy method (SEM)?
- Initialization and inflow using the Synthetic Eddy Method
- The mathematical formulation of the Synthetic Eddy Method

What Is the Synthetic Eddy Method?

The specification of realistic inflow boundary conditions is of paramount importance in Large Eddy Simulations of spatially developing flows. A number of methods have been proposed, with varying degrees of complexity and computational cost. The simplest approach consists of superimposing Gaussian noise to the mean velocity. However, this approach has the disadvantage of producing a decorrelated velocity field that is incapable of sustaining the growth of coherent turbulent structures.

The Synthetic Eddy Method (SEM) proposed by Jarrin *et al.* [57], is a simple, yet effective, approach for generating the instantaneous velocity field that is required at inflow boundaries by LES computations. The method retains the conceptual basis of the vortex method [55], but is more flexible and virtually mesh-independent.

Within the SEM framework, the turbulent flow field is seen as a superposition of spinning eddies, whose spin (clockwise or counterclockwise) and position are drawn from a normalized uniform distribution. The size of the eddies is the characteristic scale of turbulence. As eddies must span at least two cells in order to produce a correlated signal, the minimum mesh spacing imposes a limit on this turbulence scale.

A user-assigned Reynolds stress tensor provides the correlation function that the method requires. If this stress tensor is not known, the correlation function uses only normal components that are directly derived from a given turbulence intensity (assumption of isotropy).

The synthetic eddies generated at the inflow are convected and recycled in the computational domain with the mean inflow velocity.

The turbulent structures that SEM produces are only an approximation of real turbulence, and must be allowed to develop the proper correlations. The correlations develop naturally as the eddies are convected downstream. To allow this behavior to occur, it is important to leave sufficient distance

between the inflow boundary and the geometry of interest. The necessary distance varies from case to case, but in channel flow simulations, a distance of at least ten half-heights is required.

Initialization and Inflow Using the Synthetic Eddy Method

The Synthetic Eddy Method is activated when any of the LES or DES turbulence models are selected. You can access its settings through the Synthetic Turbulence Specification node which is located under both the Initial Conditions node of a physics continuum and Physics Conditions node of an inflow boundary.



To activate SEM:

- Specify either `Intensity + Length Scale` or `Reynolds Stress + Length Scale` in the Method property of the *Properties* window.



- Provide either a value or distribution for the turbulence intensity or the six components of the (symmetric) Reynolds stress tensor.

The values for turbulence intensity or the Reynolds stress tensor may come from experimental data, from precursor RANS or LES simulations, or from some other approximation. If running a precursor simulation, Reynolds stress data can only be obtained by running an LES simulation, or a RANS simulation using a Reynolds Stress Transport turbulence model. For other turbulence models, the isotropic turbulent intensity can be computed from the turbulent kinetic energy.

In many cases, a constant length scale, representative of eddy size, is sufficient for the Turbulent Length Scale. When using a RANS precursor simulation, the integral length scale can be calculated and used. For wall bounded flows, the wall distance may be a suitable approximation of the length scale when this information is not available from elsewhere. If the given length scale is inappropriate for the flow, the resolved turbulence may take much longer to develop, or disappear altogether.

When computing the SEM velocities, the number of eddies that are used in the calculation is inversely proportional to the cube of the length scale. To prevent the number of eddies becoming too large and affecting run-time and memory requirements, it is good practice to limit the minimum value of the length scale.

Properties

Method	Choose the method for defining turbulence.	
	None	Leaves the specified mean velocity unperturbed.
	Intensity + Length Scale	Calculation of an isotropic field of velocity fluctuations that are superimposed on the specified mean velocity. A Turbulence Intensity node and a Turbulent Length Scale node are added to the Values or Initial Conditions node.
	Reynolds Stress + Length Scale	Calculation of an anisotropic field of velocity fluctuations that are superimposed on the specified mean velocity. Nodes for Turbulent Length Scale, UU Reynolds stress, VV Reynolds stress, WW Reynolds stress, UV Reynolds stress, UW Reynolds stress, and VW Reynolds stress are added to the Values or Initial Conditions node.

Synthetic Turbulence Mass Flow Scaling

If you use the Synthetic Eddy Method at inflow boundaries, there is no guarantee of constant mass flow. This can cause problems for the solver, especially for incompressible cases. In such cases, use the Synthetic Turbulence Mass Flow Scaling Specification to add a correction that maintains constant mass flow by scaling the computed fluctuations.



This specification is available whenever the Synthetic Turbulence Specification is available.

Properties

Mass Flow Scaling	SEM scales fluctuations to maintain constant mass flow across boundaries.	
	Activated	Scaling in effect.
	Deactivated	The default.

SEM Formulation

The fluctuating velocity signal is generated on an interval

$$[L_{x,\,min}, L_{x,\,max}] \times [L_{y,\,min}, L_{y,\,max}] \times [L_{z,\,min}, L_{z,\,max}] \tag{436}$$

The signal uses a shape in:

$$[-\sigma_x, +\sigma_x] \times [-\sigma_y, +\sigma_y] \times [-\sigma_z, +\sigma_z] \tag{437}$$

The interval limits are based on the geometric limits of the inflow section, as shown in the figure below, and σ_i is the mean eddy length scale in coordinate direction i.



The shape function $\mathbf{f}(\mathbf{x})$ of the eddies satisfies the following normalization property:

$$\frac{1}{\Delta_x \Delta_y \Delta_z} \int_{-\Delta_x/2}^{+\Delta_x/2} \int_{-\Delta_y/2}^{+\Delta_y/2} \int_{-\Delta_z/2}^{+\Delta_z/2} \|\mathbf{f}(\mathbf{x})\| d\mathbf{x} = \mathbf{1} \tag{438}$$

where:

$$\Delta_i = L_{i,\,max} - L_{i,\,min} + 2\sigma_i \tag{439}$$

and i is one of the three coordinate directions. Any function satisfying Eqn. (438) is suitable to be taken as a shape function. In STAR-CCM+ a *tent* function is used to represent the influence the k-th eddy exerts at position $\mathbf{x}_k$ on a point $\mathbf{x}$ on the inflow surface:

$$f(x_i) = \max\left[0, 1 - \frac{|x_i - x_{i,k}|}{\sigma_{i,k}}\right] \tag{440}$$

where:

x_i is the coordinate of point $\mathbf{x}$ in the i-th direction.

$x_{i,k}$ is the coordinate of the center of eddy k in the i-th direction.

Each eddy therefore has a position $\mathbf{x}_k$, a length scale $\boldsymbol{\sigma}_k$, and a rotation (*spin*) direction which is given as:

$$\varepsilon_{k,i} = \pm 1 \tag{441}$$

As shown in [56] the sign of the rotation is chosen randomly from a uniform distribution to ensure that the generated velocity signal satisfies:

$$\overline{v_i' v_j'} = \delta_{ij} \tag{442}$$

In Eqn. (442), the overbar denotes ensemble averaging and δ_{ij} is the Kronecker delta.

The eddy turbulent length scale $\boldsymbol{\sigma}_k$ can in principle be a function of space and vary according to the coordinate direction. In the current implementation, however, the turbulent length scale is assumed to be *isotropic*: while it may vary according to position in space, it must have the same value in each coordinate direction at any one point. The reason for this is that preliminary studies have shown that for increasing values of anisotropy in the turbulent length scale, there is a progressively larger departure from the analytical normalization properties of the shape function. This results in distortion of the instantaneous velocity signal, which fails to recover the prescribed moments. This area is currently the subject of ongoing research.

The coordinates of the centers of eddies are chosen randomly in a volume $\Delta_x \times \Delta_y \times \Delta_z$ where Eqn. (439) provides each dimension. The number of eddies, N, is obtained from the ratio of this volume to the average eddy volume, computed as $\bar{\sigma}^3$, where $\bar{\sigma}$ is the average eddy length scale.

Invoking the assumption of frozen turbulence, the position of each eddy k is advanced in each coordinate direction i according to:

$$x_{i,k}^{new} = x_{i,k}^{old} + \Delta t V_i \tag{443}$$

where:

Δt is the time interval.

The convective speed V is obtained as an average over the inflow area of the mean inflow velocity profile. The fluctuating velocity signal is therefore obtained as:

$$v_i'(\mathbf{x}, t) = \frac{1}{\sqrt{N}} \sum_{k=1}^{N} \varepsilon_{i,k} f_k(\boldsymbol{\sigma}_k, \mathbf{x}, t) \tag{444}$$

The instantaneous velocity field is computed as:

$$v_i(\mathbf{x}, t) = V_i(\mathbf{x}, t) + \underbrace{a_{ij} v_i'(\mathbf{x}, t)}_{v_i''} \tag{445}$$

In Eqn. (445), the factors a_{ij} are obtained from the Cholesky decomposition of the Reynolds stress tensor as:

$$\begin{pmatrix} \sqrt{R_{11}} & 0 & 0 \\ R_{21}/a_{11} & \sqrt{R_{22} - a_{21}^2} & 0 \\ R_{31}/a_{11} & (R_{32} - a_{21}a_{31})/a_{22} & \sqrt{R_{33} - a_{31}^2 - a_{32}^2} \end{pmatrix} \tag{446}$$

where:

$$R_{ij} = \overline{v_i'' v_j''} \tag{447}$$

The assumption of isotropy is invoked when you provide turbulence intensity at the inlet rather than the individual components of the stress tensor. In this case, the normal components of the Reynolds stress tensor are calculated from the intensity and mean inflow velocity according to the following relations:

$$I = \sqrt{\frac{\overline{v''^2}}{V^2}}$$

$$k = \frac{3}{2} \overline{v''^2}$$

$$V^2 = V_i V_i$$

which results in:

$$R_{ii} = I^2 V^2 \tag{448}$$

SEM Bibliography

[55] S. Benamadouche, N. Jarrin, Y. Addad, and D. Laurence. 2006. "Synthetic turbulent inflow conditions based on a vortex-method for large eddy simulation", *Progress in Computational Fluid Dynamics*, **6**, pp. 50-57.

[56] N. Jarrin. 2008. "Synthetic inflow boundary conditions for the numerical simulation of turbulence", Ph.D. Thesis, University of Manchester (UK).

[57] N. Jarrin, S. Benamadouche, D. Laurence, and R. Prosser. 2006. "A synthetic-eddy-method for generating inflow conditions for large eddy simulations". *International Journal of Heat and Fluid Flow*, **27**, pp. 585-593.

Using Wall Treatment Models

This section explains the following key concepts:

- Wall treatment
- Wall function
- The Wall Distance solver
- The low-Reynolds number approach
- Wall laws

Guidelines for how to decide on a wall treatment and checking the wall y+ value, instructions for selecting a wall treatment and using wall roughness, and formulation of wall laws are also provided.

What Is a Wall Treatment?

A wall treatment in STAR-CCM+ is the set of near-wall modeling assumptions for each turbulence model. This term is different from the term wall functions, which refers to only one type of wall treatment.

Three types of wall treatment are provided:

- The high-y^+ wall treatment implies the wall-function-type approach: this treatment assumes that the near-wall cell lies within the logarithmic region of the boundary layer.
- The low-y^+ wall treatment is suitable only for low-Reynolds number turbulence models: this treatment assumes that the viscous sublayer is properly resolved.
- The all-y^+ wall treatment is a hybrid treatment that attempts to emulate the high-y^+ wall treatment for coarse meshes, and the low-y^+ wall treatment for fine meshes. It is also formulated with the desirable characteristic of producing reasonable answers for meshes of intermediate resolution (that is, when the wall-cell centroid falls within the buffer region of the boundary layer).

Note: Some types of wall treatment are not available for some turbulence models.

The wall treatments have been specialized according to each turbulence model, since assumptions specific to that model must be made for the wall boundary conditions for the turbulence quantities. Both the high-y^+ and all-y^+ wall treatments share a common need to specify profiles of the mean flow quantities in the near-wall region of turbulent boundary layers. These profiles are termed wall laws.

What Is a Wall Function?

A wall function is the set of mathematical relations that are used to obtain the boundary conditions for the continuum equations whereby:

- a consistent set of assumptions is made regarding the distributions of velocity, turbulence, and other scalar quantities;
- it is assumed that the turbulence model is valid only outside the viscous-dominated region of the boundary layer, and the viscous-affected region of the boundary layer is not resolved; and
- it is assumed that the centroid of the near-wall cell lies within the logarithmic region of the boundary layer.

The main advantage of using wall functions is the significant savings in terms of near-wall mesh resolution.

What Is the Wall Distance Solver?

Wall distance is a parameter that represents the distance from a cell centroid to the nearest wall face with a non-slip boundary condition. All turbulence models require this parameter.

The solver makes an exact projection calculation in real space, based on a triangulation of the surface mesh. The calculation is accelerated by use of search trees.

The wall distance solver controls the wall distance computation for all continua in which wall distance is calculated. This solver becomes available when turbulence models are activated. The wall distance calculation takes place during the initialization step of the simulation.

Information on the progress of this calculation is shown in the progress bar and status bar beneath the STAR-CCM+ explorer window.

Wall Distance Solver Properties

Verbosity	Uses a numeric entry to activate extra information for the Wall Distance solver while the simulation is running. This property is useful for debugging problems when they occur. The only valid entries are `0` (default), `1`, and `2`, and have the following effects: • `0` = None: Leaves output at its default. • `1` = Low: Displays convergence history for wall distance. • `2` = High: Generates precise equation information as the solver converges.
Minimum Tree Size Threshold	Controls the size of the search tree used by the wall distance solver. By default, the search tree can contain 500,000 data points. Increasing this number can speed up the wall distance calculation at the expense of consuming additional memory. For a parallel simulation, the maximum tree size is computed as the higher of the following numbers: • Number of cells per processor • Minimum tree size threshold By forcing the tree size to be at least the value set by the minimum tree size threshold, STAR-CCM+ limits the number of times that the wall distance solver rebuilds the search tree.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. Activated — The solver does not update the solution. Deactivated — The solver updates the solution.

What Is the Low-Reynolds Number Approach?

The low-Reynolds number approach employs a turbulence model that is valid throughout the boundary layer, including the viscous sublayer.

To resolve the viscous sublayer, this approach requires a sufficiently fine mesh. There is little or no need for modeling to obtain the wall boundary conditions. The computational expense that is associated with this approach can be significant, particularly for large Reynolds numbers.

What Is a Wall Law?

A wall law is a mathematical description of mean flow quantities, such as velocity, temperature and species concentration, in turbulent boundary layers.

The most basic wall laws apply only to equilibrium situations, although there have been some attempts to extend wall laws to include non-equilibrium effects.

Two types of wall laws are used in STAR-CCM+:

- standard wall laws, which are slope-discontinuous between the laminar and turbulent profiles;
- blended wall laws, which include a buffer region that smoothly blends the laminar and turbulent profiles together.

You cannot select wall laws; they are chosen automatically based on the behavior of the turbulence model.

Turbulence models that contain damping functions (or behave as if they do) typically use blended wall laws in their all-y^+ wall treatment. The standard wall law is used with high-y^+ wall treatments.

How Do I Decide on a Wall Treatment?

There are many applications where a wall-function-type mesh is adequate, and many applications where it is necessary to resolve the viscous-affected region. However, if the near-wall mesh resolution is not consistent with the modeling assumptions, significant errors can result.

The following guidelines can be used in the selection of the wall treatment:

- The high-y^+ wall treatment might save some computer storage by avoiding the computation of wall distance when used with certain high-Reynolds number models. This treatment is suitable for use with models that do not explicitly damp the turbulence in the near-wall region. (Examples of such models include the Spalart-Allmaras high-Reynolds number model, the standard K-Epsilon model, the realizable K-Epsilon model, the standard K-Omega model, and the SST K-Omega model.)

  A good rule of thumb is that the wall-cell centroid should be situated in the logarithmic region of the boundary layer ($y^+ > 30$). However, it is unlikely that significant error results from y^+ values as low as 12.

  In some cases, a few cells have a small value of y^+, especially near stagnation or separation. This situation is acceptable.

- The low-y^+ wall treatment makes no explicit modeling assumptions at all. Use this treatment only if the entire mesh is fine enough for y^+ to be approximately 1 or less.
- The all-y^+ wall treatment is the most general: use this treatment whenever available.

The above guidelines for both the high-y^+ wall treatment and the low-y^+ wall treatment reflect suitable lower limits on y^+. However, when wall-function-type meshes are used, many factors determine the upper limit on y^+, particularly the Reynolds number and the presence of pressure gradients. Since the potential for errors increases with large values of y^+, it is prudent to aim for values from 30 through 50.

Not all wall treatments are available with every model or model variant. Some high-Reynolds number models do not contain the ability to attenuate the turbulence in the viscous-affected regions and therefore only include the high-y^+ wall treatment. Some low-Reynolds number models are offered only with the low-y^+ wall treatment, and the K-Epsilon two-layer models are only provided with the all-y^+ wall treatment.

Checking the Wall y+ Value

The wall y+ value determines whether the near-wall spacing is appropriate for the wall treatment that is selected for a simulation. There is no automated way to verify that the wall spacing is appropriate. You need to run an initial exploratory simulation, and plot a scalar scene with Wall Y+ as the scalar field function on a wall boundary part. Use this scene to determine whether your mesh is appropriate for the wall treatment that you have chosen.

To use a low-y^+ wall treatment, the wall y^+ value at the walls of your body or region of interest should be lower than 1, as shown below.



In such cases, the boundary layer is resolved accurately. Low y^+ values are achieved when using a fine prism layer.

On the other hand, higher y^+ values can be achieved by using a coarser mesh at the wall. This kind of mesh is required when using the high y^+ wall treatment.



Using Wall Roughness

As detailed in the section describing the formulation of the wall laws, a reduction of the turbulent wall function coefficient E (see Eqn. (456)) incorporates roughness. Reducing E reproduces the effect of the erosion of the inner layer by moving the logarithmic region of the wall law downward.



As shown in the accompanying figure, for large values of roughness parameter R^+, the logarithmic profile might no longer intersect with the assumed linear profile in the sublayer. If this occurs in the fully rough regime $R^+ > R^+_{rough}$, there is no problem, since the sublayer is irrelevant. If these curves do not intersect in the transitional roughness regime $R^+_{smooth} < R^+ < R^+_{rough}$, STAR-CCM+ uses the logarithmic profile, with the proviso that u^+ is never allowed to be less than zero.

Caution: It is not physically meaningful for $y^+ < R^+$, so take care that the distance from each wall-adjacent cell centroid to the wall is larger than the wall roughness height r.

By default, STAR-CCM+ limits the roughness height locally such that $R^+ = y^+$ when this condition is violated. You can remove this limit by deactivating the Roughness Limiter Wall Roughness Parameters expert property.

Selecting a Wall Treatment

If auto-select is used for choosing a turbulence model, the choice of wall treatment is made on your behalf.

If the turbulence model and wall treatment are selected without using auto-select, a group box is presented for the selection of the wall treatment after the model selection. In the example below, the *K-Epsilon Two-Layer All y+ Wall Treatment* group box is shown.



After selection, the *Model Selection* dialog appears similar to below.



See How Do I Decide on a Wall Treatment? See also Eqn. (449) for the definition of y^+ and its relation to u^*.

High y+ Wall Treatment Properties

None.

Low y+ Wall Treatment Properties

None.

All y+ Wall Treatment Properties

Iterative Ustar — Prevents or allows iterative computation of the reference velocity u^* from the wall law.

- Activated — The reference velocity u^* is computed iteratively according to the wall law.
- Deactivated — The reference velocity u^* is computed according to the specific turbulence model.

Two-Layer All y+ Wall Treatment Properties

Iterative Ustar — Prevents or allows iterative computation of the reference velocity u^* from the wall law.

- Activated — The reference velocity u^* is computed iteratively according to the wall law.
- Deactivated — The reference velocity u^* is computed according to the specific turbulence model.

Wall Laws Formulation

In this section, details of the wall laws that are used in STAR-CCM+ are presented. This formulation includes:

- the non-dimensional quantities that are used to express the laws;
- the asymptotic limits of the wall laws for the viscous sublayer and logarithmic regions;
- the standard wall laws; and
- the blended wall laws.

Finally, the nomenclature and bibliography are presented.

Non-Dimensional Quantities

The following non-dimensional quantities are used:

$$y^+ = \frac{yu^*}{\nu} \tag{449}$$

$$u^+ = \frac{u}{u^*} \tag{450}$$

$$t^+ = (T - T_w)\frac{\rho C_p u^*}{\dot{q}''} \tag{451}$$

In these equations, the following nomenclature applies:

y is the normal distance from the wall to the wall-cell centroid.

u is the component of wall-cell velocity parallel to the wall.

T is the wall-cell temperature.

T_w is the wall-boundary temperature.

$\dot{q}''$ is the wall heat flux.

ν is the kinematic viscosity.

ρ is the density.

C_p is the specific heat.

u^* is a reference velocity.

Although the reference velocity is often related to the wall shear stress ($u^* = \sqrt{\tau_w/\rho}$), in actual practice the reference velocity is derived from a turbulence quantity specific to the particular turbulence model. For blended and two-layer models, the reference velocity u^* can be computed iteratively from the wall law.

The wall laws are set up to provide u^+ and t^+ (and their first derivatives) as a function of y^+ and other relevant quantities, such as molecular and turbulent Prandtl numbers.

Asymptotic Limits

The wall laws differ only in their treatment in the buffer region; the viscous-sublayer and log-layer behaviors are identical.

Viscous Sublayer

The velocity distribution is modeled as:

$$u^+_{\text{lam}} = y^+ \tag{452}$$

The temperature distribution (in the absence of viscous dissipation) is modeled as:

$$t^+_{\text{lam}} = \sigma u^+_{\text{lam}} \tag{453}$$

The non-dimensional quantity q^+_{lam} introduces the effect of viscous dissipation:

$$q^+_{\text{lam}} = \frac{\rho u^* u^2}{2\dot{q}''} \tag{454}$$

Logarithmic Layer

The velocity distribution is modeled as:

$$u^+_{\text{trb}} = \frac{1}{\kappa}\ln(E' y^+) \tag{455}$$

where:

$$E' = \frac{E}{f} \tag{456}$$

and the default values of the coefficients are $\kappa = 0.42$ and $E = 9.0$. The value of the roughness function f in Eqn. (456) is unity for smooth walls, and for rough walls is computed according to Eqn. (471).

The temperature distribution (in the absence of viscous dissipation) is modeled as:

$$t^+_{\text{trb}} = \sigma_t (u^+_{\text{trb}} + P) \tag{457}$$

where the function P, which governs the velocity at which the logarithmic and viscous regions of the thermal profiles intersect, is by Jayatilleke [58]:

$$P = 9.24\left[\left(\frac{\sigma}{\sigma_t}\right)^{3/4} - 1\right]\left[1 + 0.28\exp\left(\frac{-0.007\sigma}{\sigma_t}\right)\right] \tag{458}$$

and σ is the molecular Prandtl number, and σ_t is the turbulent Prandtl number.

The non-dimensional quantity q^+_{trb} introduces the effect of viscous dissipation:

$$q^+_{\text{trb}} = \frac{\rho u^*}{2\dot{q}''}\left[u^2 - \left(\frac{\sigma}{\sigma_t} - 1\right)(u^+_c u^*)^2\right] \tag{459}$$

where u^+_c is a fictitious non-dimensional velocity that would occur at the intersection of the laminar and turbulent temperature profiles. It is computed from Eqn. (458) as:

$$u^+_c = \frac{P}{\sigma - \sigma_t} \tag{460}$$

Standard Wall Laws

The standard wall laws have slope discontinuities between the laminar and logarithmic regions. They are given by:

$$u^+ = \begin{cases} u^+_{\text{lam}} & \text{for } y^+ \le y^+_m \\ u^+_{\text{trb}} & \text{for } y^+ > y^+_m \end{cases} \tag{461}$$

$$t^+ = \begin{cases} \sigma(u^+_{\text{lam}} - q^+_{\text{lam}}) & \text{for } y^+ \le y^+_t \\ \sigma_t(u^+_{\text{trb}} + P + q^+_{\text{trb}}) & \text{for } y^+ > y^+_t \end{cases} \tag{462}$$

For the momentum wall law, a Newton iteration determines the intersection of the viscous and fully turbulent regions y^+_m. A similar iterative process determines y^+_t, the intersection of the viscous and fully turbulent regions for the thermal wall law.

Blended Wall Laws

The blended wall laws are intended to represent the buffer layer by appropriately blending the viscous sublayer and logarithmic regions. For momentum, Reichardt's law [61] is used:

$$u^+ = \frac{1}{\kappa}\ln(1 + \kappa y^+) + C\left[1 - \exp\left(-\frac{y^+}{D}\right) - \frac{y^+}{D}\exp(-by^+)\right] \tag{463}$$

where:

$$D = y^+_m \tag{464}$$

$$C = \frac{1}{\kappa}\ln\left(\frac{E'}{\kappa}\right) \tag{465}$$

$$b = \frac{1}{2}\left(\frac{D\kappa}{C} + \frac{1}{D}\right) \tag{466}$$

Reichardt's law is also used to compute the near-wall reference velocity u^* if the appropriate option is selected; see Eqn. (449). For temperature, Kader's law [59] modified slightly for the effects of wall roughness is used:

$$t^+ = \exp(-\Gamma)\ (t^+_{\text{lam}} - q^+_{\text{lam}}) + \exp\left(-\frac{1}{\Gamma}\right)\ (t^+_{\text{trb}} - q^+_{\text{trb}}) \tag{467}$$

where:

$$\Gamma = \frac{0.01c(\sigma y^+)^4}{1 + \frac{5}{c}\sigma^3 y^+} \tag{468}$$

and

$$c = e^{\left(\frac{1}{f} - 1\right)} \tag{469}$$

where f is the roughness function.

Roughness Function

The roughness function f that modifies the log-law coefficient E is a function of the roughness parameter

$$R^+ = \frac{ru^*}{\nu} \tag{470}$$

where r is the equivalent sand-grain roughness height, that you specify. You obtain the value of the equivalent sand-grain roughness height that is appropriate for your model either from the literature or empirically. To determine r from an experiment:

- Measure the friction factor for the Reynolds number of your case in the experiment.
- Plot your measurements on a Moody diagram.
- Deduce the equivalent sand-grain roughness height from the Moody diagram curve.

The function that is used is a slightly expanded version of the expression that was given in [62]:

$$f = \begin{cases} 1 & \text{for } R^+ \le R^+_{\text{smooth}} \\ \left[B\left(\dfrac{R^+ - R^+_{\text{smooth}}}{R^+_{\text{rough}} - R^+_{\text{smooth}}}\right) + CR^+\right]^a & \text{for } R^+_{\text{smooth}} < R^+ < R^+_{\text{rough}} \\ B + CR^+ & \text{for } R^+ > R^+_{\text{rough}} \end{cases} \tag{471}$$

where the exponent a is expressed as:

$$a = \sin\left[\frac{\pi}{2}\frac{\log(R^+/R^+_{\text{smooth}})}{\log(R^+_{\text{rough}}/R^+_{\text{smooth}})}\right] \tag{472}$$

and the default values of the coefficients are from [62]:

$$B = 0 \quad , \quad C = 0.253 \quad , \quad R^{+}_{\text{smooth}} = 2.25 \quad , \quad R^{+}_{\text{rough}} = 90$$

Wall Laws Nomenclature

The following is a guide to the terms used in this section.

Scalars

E	wall function coefficient
f	roughness function
u	component of wall-cell velocity parallel to the wall
u^*	reference velocity that is used in wall functions
u^+	wall-parallel velocity non-dimensionalized with u^*
R^+	roughness parameter
y^+	non-dimensional wall coordinate
y	normal distance from the wall to the wall-cell centroid
P	function of Prandtl numbers
q^+	dimensionless viscous dissipation
t^+	dimensionless temperature
T	wall-cell temperature
T_w	wall-boundary temperature
$\dot{q}''$	wall heat flux magnitude
C_p	specific heat
κ	von Karman constant
ρ	density
μ	dynamic viscosity
μ_t	turbulent viscosity
σ	Prandtl number
σ_t	turbulent Prandtl number
τ	wall shear stress magnitude
Γ	blending function in Kader's wall law

Subscripts

lam	laminar
turb	turbulent

Wall Laws Bibliography

[58] Jayatilleke, C.L. 1969. "The influence of Prandtl number and surface roughness on the resistance of the laminar sub-layer to momentum and heat transfer", *Progress in Heat and Mass Transfer*, **1**, pp. 193-330.

[59] Kader, B.A. 1981. "Temperature and Concentration Profiles in Fully Turbulent Boundary Layers", *Int. J. Heat Mass Transfer*, **24**, pp. 1541-1544.

[60] Launder, B.E. and Spalding, D.B. 1974. "The Numerical Computation of Turbulent Flows", *Computer Methods in Applied Mechanics and Engineering*, **3**, pp. 269-289.

[61] Reichardt, H. 1951. "Vollstaendige Darstellung der turbulenten Geschwindigkeitsverteilung in glatten Leitungen", *Z. Angew. Math. Mech,* **31**(7), pp. 208-219.

[62] Cebeci, T. and Bradshaw, P. 1977. *Momentum Transfer in Boundary Layers,* Hemisphere Publishing/McGraw-Hill, pp. 176-180.

Using Spalart-Allmaras Turbulence

This section explains the following key concepts:

- What are Spalart-Allmaras Turbulence models?
- Understanding the Spalart-Allmaras wall treatment
- How to choose a model and wall treatment
- Working with Spalart-Allmaras solvers
- Formulation of the model equations

Also included are reference sections on the following topics as they relate to the Spalart-Allmaras model:

- Turbulence specification
- Modified diffusivity
- Modified Diffusivity Source
- Field functions

What Are the Spalart-Allmaras Turbulence Models?

The Spalart-Allmaras turbulence models [68] solve a single transport equation that determines the turbulent viscosity. This situation is in contrast to many of the early one-equation models that solve an equation for the transport of turbulent kinetic energy and required an algebraic prescription of a length scale.

The original model was developed primarily for the aerospace industry, and has the advantage of being readily implemented in an unstructured CFD solver, unlike the more traditional aerospace models such as Baldwin-Lomax [63] and Johnson-King [65]. This advantage has resulted in its popularity increasing as the use of unstructured CFD methods has grown more widespread in the aerospace industry.

The authors of the original Spalart-Allmaras turbulence model presented results for attached boundary layers and flows with mild separation (such as flow past a wing). It is reasonable to expect that these cases are the types of flows for which the model yields the best results. Wilcox [73] presents free-shear flow spreading rates for the model. While acceptable results are obtained for wake, mixing layer and radial jet flows, the predicted spreading rates for plane and round jets are inaccurate. Therefore, Wilcox concludes that the model is not suited to applications involving jet-like free-shear regions. It is also likely to be less suited to flows involving

complex recirculation and body forces (such as buoyancy) than two-equation models such as K-Epsilon and K-Omega or Reynolds Stress Transport.

Three variants of the model are available in STAR-CCM+:

- Standard Spalart-Allmaras
- High-Reynolds number Spalart-Allmaras
- Spalart-Allmaras detached eddy model

In addition, separate versions of the first two variants exist for use with harmonic balance.

What Is the Standard Spalart-Allmaras Model?

In its standard form, the Spalart-Allmaras model is a low-Reynolds number model, meaning that it is applied without wall functions. According to the formulation of the model, the entire turbulent boundary layer, including the viscous sublayer, can be accurately resolved and the model can be applied on fine meshes (small values of y^+). This model can be used with the all y^+ wall treatment for all simulations.

Standard Spalart-Allmaras Properties

Deformation Option	Selects the modified or standard production term.	
	Modified	Selects the modified production term. This option is recommended for most simulations. See Eqn. (486).
	Original	Selects the standard production term. This option allows comparison with results from other CFD codes that lack the modified production term. See Eqn. (484).
Convection	For guidance on selecting a convection, see Convection Term, p. 6343.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.

Standard Spalart-Allmaras Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not make any changes within the *Expert* category. The values in that category reflect both the design of the model and discretization approaches that have been optimized for accuracy and performance. Tampering with them can diminish the effectiveness of the model.

Cb1	The coefficient C_{b1}, Eqn. (474).
Cb2	The coefficient C_{b2}, Eqn. (473).
Cprod	The coefficient C_{prod}, Eqn. (486).
Cv1	The coefficient C_{v1}, Eqn. (482).
Cw2	The coefficient C_{w2}, Eqn. (491).
Cw3	The coefficient C_{w3}, Eqn. (490).

Kappa	The von Karman constant κ, Eqns.(479) and (493).	
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Sigma	The coefficient $\sigma_{\tilde{\nu}}$, Eqn. (473).	

What Is the High-Reynolds Number Spalart-Allmaras Model?

In the derivation of the model, the authors present a high-Reynolds number version in which the viscous damping within the buffer layer and viscous sublayer is not included. This form of the model, analogous to a K-Epsilon model with wall functions, has also been presented in STAR-CCM+. It is only suited to coarse, wall-function-type meshes where y^+ values generally exceed about 30. The standard model with All y+ Wall Treatment can give similar results to this model on coarse meshes. Therefore, activate the High-Reynolds number formulation only if you suspect that the viscous damping functions adversely affect results. This problem is likely to occur if y^+ values are in the intermediate range of 10–30.

High-Reynolds Number Spalart-Allmaras Properties

Deformation Option	Selects the modified or standard production term.	
	Modified	Selects the modified production term. This option is recommended for most simulations. See Eqn. (486).
	Original	Selects the standard production term. This option allows comparison with results from other CFD codes that lack the modified production term. See Eqn. (484).
Convection	For guidance on selecting a convection, see Convection Term, p. 6343.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.
Curvature Correction Option	Selects whether curvature correction is applied to eddy viscosity in curved streamlines, or when solving the transport equation in rotating frames of reference.	
	Off	Curvature correction is not applied. (Default)
	On	Curvature correction is applied.

High-Reynolds Number Spalart-Allmaras Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not make any changes within the *Expert* category. The values in that category reflect both the design of the model and discretization approaches that have been optimized for accuracy and performance. Changing these properties can diminish the effectiveness of the model.

Cb1	The coefficient C_{b1}, Eqn. (474).
Cb2	The coefficient C_{b2}, Eqn. (473).
Cprod	The coefficient C_{prod}, Eqn. (486).

Cw2	The coefficient C_{w2}, Eqn. (491).	
Cw3	The coefficient C_{w3}, Eqn. (490).	
Kappa	The von Karman constant κ, Eqns.(479) and (492).	
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Sigma	The coefficient $\sigma_{\tilde{\nu}}$, Eqn. (473).	

What Is the Spalart-Allmaras Detached Eddy Model?

Detached eddy simulation (DES) [69] is a modeling approach that combines features of Reynolds-Averaged (RANS) simulation in part of the flow and large eddy simulation (LES) in unsteady, separated regions. The DES version of the Spalart-Allmaras model contains modifications to the damping functions recommended by Squires and others [71]. These modifications do not affect the predictions of fully turbulent flows, but prevent spurious upstream propagation of the eddy viscosity into attached laminar regions. Recent updates by Spalart and others [70] termed DDES (delayed detached eddy simulation) have been incorporated to enhance the ability of the model to distinguish between LES and RANS regions. In addition, the improved delayed detached eddy simulation (IDDES) formulation of Shur and others [66] is available as an option. This combines DDES with an improved RANS-LES hybrid model aimed at wall modeling in LES when the grid resolution supports it.

Inflow Boundaries and Initialization

The Synthetic Eddy Method is implemented to provide turbulent eddies automatically across inflow boundaries and provide an initial perturbated flow field.

Spalart-Allmaras Detached Eddy Properties

Formulation Option	Selects the DES formulation to use.	
	DDES	Uses DDES formulation.
	IDDES	Uses IDDES formulation. This selection activates the IDDES Coefficients sub-node for this model.
Low Re Correction	Option to include the low-Reynolds number correction factor ψ in Eqn. (494).	
	Activated	ψ is computed as in Eqn. (497).
	Deactivated	ψ is set to 1.
Convection	For guidance on selecting a convection, see Convection Term, p. 6343.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.
Curvature Correction Option	Selects whether curvature correction is applied to eddy viscosity in curved streamlines, or when solving the transport equation in rotating frames of reference.	
	Off	Curvature correction is not applied. (Default)
	On	Curvature correction is applied.

Spalart-Allmaras Detached Eddy Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not make any changes within the *Expert* category. The values in that category reflect both the design of the model and discretization approaches that have been optimized for accuracy and performance. Tampering with them can diminish the effectiveness of the model.

Cb1	The coefficient C_{b1}, Eqn. (474).
Cb2	The coefficient C_{b2}, Eqn. (473).

Cdes	The DES length scale coefficient C_{des}, Eqn. (494).
CdesTimeLim	The DES blended upwind scheme time scale coefficient C_{desT}, Eqn. (513).
Cv1	The coefficient C_{v1}, Eqn. (482).
Cv2	The coefficient C_{v2}, Eqn. (480).
Cw2	The coefficient C_{w2}, Eqn. (491).
Cw3	The coefficient C_{w3}, Eqn. (490).
Kappa	The von Karman constant κ, Eqns.(479) and (492).
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.
	On — Include both secondary gradients.
	Off — Exclude both secondary gradients.
	Interior Only — Include the interior secondary gradients only.
	Boundaries Only — Include the boundary secondary gradients only.
Sigma	The coefficient $\sigma_{\tilde{\nu}}$, Eqn. (473).

Working with the IDDES Coefficients

The IDDES Coefficients node represents the terms C_t in Eqn. (503), C_l in Eqn. (504), and C_{dt} in Eqn. (508) as constant coefficients. The default settings of these coefficients are 1.63, 3.55, and 8.0, respectively, and can be changed

with a text entry in the corresponding expert properties. This object is activated as a sub-node of the Spalart-Allmaras Detached Eddy model node when the Formulation Option property of the latter is set to `IDDES`.



IDDES Coefficients Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not make any changes within the *Expert* category. The values in that category reflect both the design of the model and discretization approaches that have been optimized for accuracy and performance. Tampering with them can diminish the effectiveness of the model.

Ct	The coefficient C_t, see Eqn. (503).
Cl	The coefficient C_l, see Eqn. (504).
Cdt	The coefficient C_{dt}, see Eqn. (508).

Spalart-Allmaras Detached Eddy Field Functions Reference

The following primitive field functions are made available to the simulation when the Spalart-Allmaras Detached Eddy turbulence model is used. These objects all have the same base properties.

DES Low Re Correction Factor	The low-Reynolds number correction function ψ defined in Eqn. (497)
DES Upwind Blending Factor	The blending function for the hybrid upwind scheme
Detached Eddy Length Scale	The cell length scale that is used for detached eddy simulation Δ, see Eqn. (494).
IDDES Length Scale	The cell length scale that is used for detached eddy simulation with the IDDES formulation option Δ_{IDDES}, see Eqn. (509).

Understanding the Spalart-Allmaras Wall Treatment

There are three alternative wall treatments available to compute wall-cell parameters that are related to velocity gradients, wall fluxes and stresses, and source terms in the $\tilde{\nu}$ transport equation (Eqn. (473)). These alternatives are:

- Low y+ wall treatment
- High y+ wall treatment
- All y+ wall treatment

What Is the Low y+ Wall Treatment?

The Low y+ wall treatment implements the model and boundary conditions in low-Reynolds number form as suggested by Spalart and Allmaras. Damping functions are used for the source terms in the transport equation, which is solved without modification all the way to the wall cell. Wall laws are not used to evaluate shear stress. Only use these if the mesh is fine enough to resolve the viscous sublayer. If y^+ values substantially exceed 1, significant errors result. If there is any doubt, use the All y+ wall treatment.

What Is the High y+ Wall Treatment?

The High y+ wall treatment is essentially a wall-function approach for cases where the near-wall cell centroid falls within the logarithmic region of the boundary layer. The model does not use damping functions, so that $f_{v1} = 1$ and $f_{v2} = 0$ in Eqns. (480) and (482) and the source terms of the model in the wall cell are suitably modified using the standard wall laws. This treatment is only available with the High-Reynolds number Spalart-Allmaras model.

What Is the All y+ Wall Treatment?

The All y+ Wall Treatment is a hybrid model. It provides a more realistic modeling than either the low-Re or the high-Re treatments, for when the wall-cell centroid falls in the buffer region. This wall treatment is meant to give results similar to the low-y^+ treatment as $y^+ \rightarrow 0$ and to the high-Re treatment for $y^+ > 30$. This treatment uses damping functions for the source terms in the transport equation, but the source terms of the model in the wall cell are suitably modified using the blended wall laws. This wall treatment is recommended for most simulations.

How Do I Choose a Spalart-Allmaras Model and Wall Treatment?

The available Spalart-Allmaras models and wall treatments are shown below.

Spalart-Allmaras Variant	Wall Treatment
Standard	low-y^+, all-y^+
High-Reynolds Number	high-y^+
DES	low-y^+, all-y^+

Some recommendations regarding which model to use follow:

- The Standard Spalart-Allmaras model, combined with the all-y^+ wall treatment, is recommended for most, if not all, applications.
- The Standard Spalart-Allmaras model, combined with the low-y^+ wall treatment, could be used in situations where the mesh is known to be fine enough to resolve the viscous sublayer everywhere. The model is designed to give substantially the same results as the all-y^+ wall treatment.

- The High-Reynolds number Spalart-Allmaras model is provided as a check on the effect of the damping functions on the performance of the mode when the mesh is not coarse enough to be unaffected by the damping functions in the model ($20 < y^+ < 30$). It is not recommended for general-purpose applications.
- The DES Spalart-Allmaras model is only available with unsteady simulations, since detached eddy simulation is an inherently unsteady phenomenon.

Working with Spalart-Allmaras Solvers

This section describes the solvers for the Spalart-Allmaras family of turbulence models:

- Spalart-Allmaras Turbulence
- Spalart-Allmaras Turbulent Viscosity

It also lists the residual monitors that are added to the simulation by these solvers.

Spalart-Allmaras Turbulence Solver

This solver controls the solution of the Spalart-Allmaras model in all the continua for which the model is activated.

Spalart-Allmaras Turbulence Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the newly computed solution replaces the old solution. The default value is 0.7. For the theoretical background, see Implicit Iteration, p. 6362.

Spalart-Allmaras Turbulence Solver Expert Properties

For information about properties in the *Expert* category, see Expert Properties, p. 6385.

Spalart-Allmaras Turbulent Viscosity Solver

This solver controls the update of the turbulent viscosity.

Spalart-Allmaras Turbulent Viscosity Solver Properties

Maximum Ratio	The maximum ratio of turbulent to laminar viscosity (μ_t/μ) permitted during the update process. This property is used to prevent unphysically high values of turbulent viscosity from occurring on the path to convergence.
Under-Relaxation Factor	An under-relaxation factor ω_μ for the update of turbulent viscosity. The default value is 1.0. Let μ_t^n be the value of the previous iteration, and μ_t^{new} be the value that is computed during the current iteration using Eqn. (493). Then the update is controlled as follows: $\mu_t^{n+1} = \omega_\mu \mu_t^{new} + (1-\omega_\mu)\mu_t^n$

Spalart-Allmaras Turbulent Viscosity Solver Expert Properties

For information about properties in the *Expert* category, see Expert Properties, p. 6385.

Residual Monitors

The Spalart-Allmaras turbulence solvers add the following residual monitors to the top-level Monitors node of a simulation tree.

Sa_nut	Modified diffusivity $\tilde{\nu}$ in Eqn. (473).

The properties of all residual monitors are the same.

Spalart-Allmaras Turbulence Formulation

The formulation of the Spalart-Allmaras models includes the basic transport equation and the individual terms of the model:

- Non-conservative diffusion
- Turbulent production
- Turbulent dissipation

In addition, this section presents the following:

- Relation for turbulent viscosity
- Turbulent length scale
- Model coefficients

- Evaluation of reference velocity u*
- Boundary, region, and initial conditions

Finally, the nomenclature and bibliography are presented.

Basic Transport Equation

The transport equation for the Spalart-Allmaras model is:

$$\frac{d}{dt}\int_V \rho\tilde{\nu}dV + \int_A \rho\tilde{\nu}(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \frac{1}{\sigma_{\tilde{\nu}}}\int_A (\mu+\rho\tilde{\nu})\nabla\tilde{\nu}\cdot d\mathbf{a} + \int_V [C_{b2}\rho(\nabla\tilde{\nu}\cdot\nabla\tilde{\nu}) + G_{\tilde{\nu}} - \Upsilon_{\tilde{\nu}} + S_{\tilde{\nu}}]\ dV \tag{473}$$

where $S_{\tilde{\nu}}$ is the user-specified source term, and the transported variable $\tilde{\nu}$ is a modified diffusivity. The terms on the right-hand side represent diffusion, production, and dissipation.

Non-Conservative Diffusion

The first term inside the last integral of the above equation is a non-conservative diffusion term. If it is discretized as an explicit source term, it can lead to convergence problems. Therefore, following the recommendations of Spalart and Allmaras, it is rewritten and combined with the conservative diffusion term as follows:

$$\frac{1}{\sigma_{\tilde{\nu}}}\int_A (\mu+\rho\tilde{\nu})\nabla\tilde{\nu}\cdot d\mathbf{a} + \int_V C_{b2}\rho(\nabla\tilde{\nu}\cdot\nabla\tilde{\nu})\ dV \approx \frac{1+C_{b2}}{\sigma_{\tilde{\nu}}}\int_A (\mu+\rho\tilde{\nu})\nabla\tilde{\nu}\cdot d\mathbf{a} - \frac{C_{b2}}{\sigma_{\tilde{\nu}}}\int_V (\mu+\rho\tilde{\nu})\nabla\tilde{\nu}^2 dV \tag{474}$$

Turbulent Production

Shur and others [67] proposed a correction term for streamline curvature effects. The model acts directly on the eddy-viscosity by modifying the production term $G_{\tilde{\nu}}$:

$$G_{\tilde{\nu}} = (1-f_{t2})\ C_{b1} f_{r1}\rho\tilde{S}\tilde{\nu} \tag{475}$$

where C_{b1} is a constant and the rotation function f_{r1} is:

$$f_{r1} = (1+c_{r1})\ \frac{2r^*}{1+r^*}[1-c_{r3}\tan^{-1}(c_{r2}r)] - c_{r1} \tag{476}$$

with $r^* = |\mathbf{S}|/|\Omega|$ and:

$$r = 2\Omega_{ik} S_{jk} \left[\frac{DS_{ij}}{Dt} \right] \frac{1}{D^4} \tag{477}$$

where $D^2 = 0.5(|\mathbf{S}|^2 + |\Omega|^2)$. The model constants in Eqn. (476) are $c_{r1} = 1$, $c_{r2} = 12$ and $c_{r3} = 1$.

The function f_{t2} is:

$$f_{t2} = \begin{cases} 0 & \text{for the high-Reynolds number model} \\ C_{t3}\exp(-C_{t4}\chi^2) & \text{for the standard and DES models} \end{cases} \tag{478}$$

The deformation parameter is:

$$\tilde{S} = f_{v3} S + \frac{\tilde{\nu}}{\kappa^2 \tilde{d}^2} f_{v2} \tag{479}$$

where κ is the von Karman constant, $\tilde{d}$ is the length scale, and the damping functions f_{v2} and f_{v3} are defined as:

$$f_{v2} = \begin{cases} 0 & \text{for the high-Reynolds number model} \\ 1 - \dfrac{\chi}{1 + \chi f_{v1}} & \text{for the standard model} \\ \left(1 + \dfrac{\chi}{C_{v2}}\right)^{-3} & \text{for the DES model} \end{cases} \tag{480}$$

$$f_{v3} = \begin{cases} 1 & \text{for the high-Reynolds number model} \\ 1 & \text{for the standard model} \\ \dfrac{(1 + \chi f_{v1})(1 - f_{v2})}{\chi} & \text{for the DES model} \end{cases} \tag{481}$$

The additional damping function f_{v1} is:

$$f_{v1} = \begin{cases} 1 & \text{for the high-Reynolds number model} \\ \dfrac{\chi^3}{\chi^3 + C_{v1}^3} & \text{for the standard and DES models} \end{cases} \tag{482}$$

$$\text{and: } \chi = \frac{\tilde{\nu}}{\nu} \tag{483}$$

Two alternatives are provided for evaluation of the scalar deformation S:

The original model uses the magnitude of the vorticity:

$$S = |\mathbf{W}| = \sqrt{2\mathbf{W}:\mathbf{W}^T} \tag{484}$$

where the vorticity tensor $\mathbf{W}$ is defined:

$$\mathbf{W} = \frac{1}{2}(\nabla\mathbf{v} - \nabla\mathbf{v}^T) \tag{485}$$

An alternative suggested in Dacles-Mariani and others [64] combines the strain rate and vorticity tensor magnitudes as follows:

$$S = |\mathbf{W}| + C_{\text{prod}} \min\,[0, |\mathbf{S}| - |\mathbf{W}|] \tag{486}$$

where:

$$|\mathbf{S}| = \sqrt{2\mathbf{S}:\mathbf{S}^T} = \sqrt{2\mathbf{S}:\mathbf{S}} \tag{487}$$

$$\mathbf{S} = \frac{1}{2}(\nabla\mathbf{v} + \nabla\mathbf{v}^T) \tag{488}$$

and the coefficient $C_{\text{prod}} = 2$. This alternative production method is the default for RANS calculations. However, it is not available for detached eddy simulation since tests have show that it does not produce the correct decay of turbulence when the DES equations are in “LES mode”.

Turbulent Dissipation

The dissipation term (for both the high- and low-Reynolds number versions) is modeled as:

$$\Upsilon_{\tilde{\nu}} = \rho\left(C_{w1}f_w - \frac{C_{b1}}{\kappa^2}f_{t2}\right)\left(\frac{\tilde{\nu}}{\tilde{d}}\right)^2 \tag{489}$$

where:

$$f_w = g\left(\frac{1 + C_{w3}^6}{g^6 + C_{w3}^6}\right)^{1/6} \tag{490}$$

$$g = r + C_{w2}(r^6 - r) \tag{491}$$

$$r \equiv \frac{\tilde{\nu}}{\tilde{S}\kappa^2\tilde{d}^2} \tag{492}$$

and f_{t2} is given in Eqn. (478).

Relation for Turbulent Viscosity

The turbulent viscosity μ_t is computed from:

$$\mu_t = \rho\tilde{\nu}f_{v1} \tag{493}$$

where f_{v1} is defined in Eqn. (482).

Turbulent Length Scale

The length scale, $\tilde{d}$, is a function of the distance to the nearest wall. It is computed as follows:

$$\tilde{d} = \begin{cases} d & \text{for the high-Reynolds number model} \\ d & \text{for the standard model} \\ d - f_d \max(0, d - \psi C_{\text{des}}\Delta) & \text{for the DES model (DDES)} \\ \tilde{f}_d(1 + f_e)d + (1 - \tilde{f}_d)\psi C_{\text{des}}\Delta_{\text{IDDES}} & \text{for the DES model (IDDES)} \end{cases} \tag{494}$$

DDES Formulation of the DES Model

For the DES model, C_{des} is a coefficient and Δ is the largest distance between the cell center under consideration and the cell centers of the neighboring cells. The function f_d is [70]:

$$f_d = 1 - \tanh[(8r_d)^3] \tag{495}$$

where:

$$r_d = \frac{\tilde{\nu}}{\sqrt{\nabla\mathbf{v}:\nabla\mathbf{v}^T}\kappa^2 d^2} \tag{496}$$

The optional low-Reynolds number correction function ψ is designed to prevent activation of the low-Reynolds number terms when in "LES" mode [70]. It is defined as:

$$\psi^2 = \min\left\{100, \frac{1 - C_{b1}([f_{t2} + (1 - f_{t2})f_{v2}]/f_w^* \kappa^2 C_{w1})}{(1 - f_{t2})f_{v1}f_{v3}}\right\} \tag{497}$$

where the damping functions f_{t2}, f_{v2}, f_{v3} and f_{v1} are defined as appropriate for the DES model. The coefficient f_w^* is a computed iteratively as a nonlinear function of the other model coefficients, as described in [70]. It has a value of 0.424 when the standard model coefficients are used.

IDDES Formulation of the DES Model

For the IDDES model, two more functions are introduced to the length scale calculation to add Wall-Modeled LES (WMLES) capability, a blending function f_B and the so-called "elevating" function f_e:

$$f_B = \min[2\exp(-9\alpha^2), 1.0] \tag{498}$$

$$\alpha = 0.25 - \frac{d}{\Delta} \tag{499}$$

$$f_e = \max[(f_{e1} - 1), 0]\psi f_{e2} \tag{500}$$

$$f_{e1} = \begin{cases} 2\exp(-11.09\alpha^2) & \text{if } \alpha \geq 0 \\ 2\exp(-9.0\alpha^2) & \text{if } \alpha < 0 \end{cases} \tag{501}$$

$$f_{e2} = 1.0 - \max(f_t, f_l) \tag{502}$$

$$f_t = \tanh[(C_t^2 r_{dt})^3] \tag{503}$$

$$f_l = \tanh[(C_l^2 r_{dl})^{10}] \tag{504}$$

$$r_{dt} = \frac{\nu_t}{\sqrt{\nabla\mathbf{v}:\nabla\mathbf{v}^T}\kappa^2 d^2} \tag{505}$$

$$r_{dl} = \frac{\nu}{\sqrt{\nabla\mathbf{v}:\nabla\mathbf{v}^T}\kappa^2 d^2} \tag{506}$$

The introduction of the low-Reynolds number correction function ψ in the formulation of f_e is unrelated to the low-Reynolds number correction role of this function in the LES mode of DDES, and is purely empirical.

C_t and C_l are model constants.

The WMLES and DDES branches of the model are combined using a modified version of the DDES f_d function as follows:

$$\tilde{f}_d = \max((1 - f_{dt}), f_B) \tag{507}$$

$$f_{dt} = 1 - \tanh[(C_{dt} r_{dt})^3] \tag{508}$$

C_{dt} is a model constant.

The IDDES model also uses an altered version of the mesh length scale Δ_{IDDES}, computed as:

$$\Delta_{\mathrm{IDDES}} = \min(\max(0.15*d, 0.15*\Delta, \Delta_{\min}), \Delta) \tag{509}$$

where $\Delta_{\min}$ is the smallest distance between the cell center under consideration and the cell centers of the neighboring cells.

Model Coefficients

The model coefficients have the following values:

$$\begin{gathered} C_{b1} = 0.1335 \quad , \quad C_{b2} = 0.622 \quad , \quad \sigma_{\tilde{\nu}} = 2/3 \\ C_{v1} = 7.1 \quad , \quad C_{v2} = 5 \quad , \quad C_{w1} = \frac{C_{b1}}{\kappa^2} + \frac{(1 + C_{b2})}{\sigma_{\tilde{\nu}}} \\ C_{w2} = 0.3 \quad , \quad C_{w3} = 2.0 \quad , \quad \kappa = 0.41 \quad , \quad C_{\mathrm{des}} = 0.65 \\ C_{t3} = 1.1 \quad , \quad C_{t4} = 2.0 \end{gathered} \tag{510}$$

Evaluation of Reference Velocity *u**

Spalart and Allmaras devised the model such that $u^* = \tilde{\nu}/(\kappa \tilde{d})$ would remain constant throughout the viscous sublayer as well as in the log region. However, they concede that "ν_t does not follow its ideal linear log-law behavior $\kappa y u_\tau$ far up..." This means that the assumption of a linear relationship causes inaccuracies in u^* at large y^+ values. In STAR-CCM+, a proprietary quadratic relation is used that yields more realistic values of u^* for coarse meshes.

Function for Hybrid Upwind Scheme

Recognizing the superiority of an upwind scheme for RANS computations and the benefits of central differencing for LES, Travin and others [72] proposed the following blending coefficient for DES:

$$\sigma = \tanh\{[\min(A, 10)]^3\} \tag{511}$$

$$A = \max\left(\frac{l_{\text{ratio}}}{\max(g, 10^{-20})} - 0.5, 0\right) \tag{512}$$

$$l_{\text{ratio}} = C_{\text{des}}\Delta\sqrt{\frac{\rho C_\mu^{1.5}\max(\mathcal{S}, C_{\text{desT}}/\tau)}{\mu + \mu_t}} \tag{513}$$

$$g = \tanh\{[\min(B, 10)]^4\} \tag{514}$$

$$B = \frac{2W\max(S, W)}{\max(\mathcal{S}^2, 10^{-20})} \tag{515}$$

$$\mathcal{S}^2 = \frac{W^2 + S^2}{2} \tag{516}$$

with

$$W = \sqrt{2\mathbf{W}:\mathbf{W}^T} \tag{517}$$

$$S = \sqrt{2\mathbf{S}:\mathbf{S}^T} = \sqrt{2\mathbf{S}:\mathbf{S}} \tag{518}$$

and the vorticity tensor W and strain rate tensor S defined according to Eqns. [483] and [488] respectively.

The local convective scale, τ is defined as:

$$\tau = \frac{\Delta}{|\mathbf{v}|} \tag{519}$$

and the coefficient C_{desT} has a default value of 1.

Boundary, Region, and Initial Conditions

Wall Boundaries

The value of $\tilde{\nu}$ is set to zero at the wall.

Flow, Region, and Initial Conditions

When defining values for flow boundary, region and initial conditions, you have several choices for specifying the modified turbulent diffusivity $\tilde{\nu}$:

- Enter the value directly.
- Have it derived from a specified turbulence intensity (I) and length scale (L) using the relation:

$$\tilde{\nu} \approx C_{\mu}^{0.25} \sqrt{\frac{3}{2}} I \mathrm{v} L \tag{520}$$

where v is the local velocity magnitude and the coefficient C_{μ} has a default value of 0.09.

- Have it derived from a specified turbulent viscosity ratio μ_t/μ, requiring a solution of the following equation:

$$\frac{\mu_t}{\mu} = \frac{\rho \tilde{\nu} f_{v1}}{\mu} \tag{521}$$

Spalart-Allmaras Turbulence Nomenclature

The following is a guide to the terms used in this section.

Scalars

d	distance to the nearest wall
$\tilde{d}$	turbulent length scale
t	time
κ	von Karman constant
$\tilde{\nu}$	modified turbulent diffusivity
ρ	density
μ	dynamic viscosity
ν	kinematic viscosity
$\sigma_{\tilde{\nu}}$	turbulent Schmidt number
χ	ratio of modified turbulent diffusivity to kinematic viscosity
u^*	reference velocity that is used in wall functions
$G_{\tilde{\nu}}$	turbulent production
$\Upsilon_{\tilde{\nu}}$	turbulent dissipation
I	turbulence intensity
L	turbulent length scale
$\tilde{S}$	modified deformation parameter
V	cell volume
v	velocity magnitude

Tensors or Column Vectors

S	strain rate tensor
W	vorticity tensor

Cartesian Vectors

a	face area vector
v	velocity
$\mathbf{v}_g$	grid velocity

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S}:\mathbf{T} = S_{ij}T_{ji}$$

Spalart-Allmaras Turbulence Bibliography

[63] Baldwin, B.S. and Lomax, H. 1978. "Thin-Layer Approximation and Algebraic Model for Separated Turbulent Flows", AIAA-78-0257.

[64] Dacles-Mariani, J., Zilliac, G., Chow, J.S., and Bradshaw, P. 1995. "Numerical/experimental study of a wingtip vortex in the near field", *AIAA Journal*, **33**(9), pp. 1561-1568.

[65] Johnson, D.A. and King, L.S. 1985. "A Mathematically Simple Turbulence Closure Model for Attached and Separated Turbulent Boundary Layers", *AIAA Journal*, **23**(11), pp. 1684-1692

[66] Shur, M.L., Spalart, P.R., Strelets, M.Kh., and Travin, A.K. 2008. "A hybrid RANS-LES approach with delayed-DES and wall-modelled LES capabilities", *International J. Heat and Fluid Flow*, **29**(6), pp. 1638-1649.

[67] Shur, M. L., Strelets, M. K., Travin, A. K., and Spalart, P. R. 2000. "Turbulence Modeling in Rotating and Curved Channels: Assessing the Spalart-Shur Correction", *AIAA Journal*, **38**(5), pp. 784-792.

[68] Spalart, P.R., and Allmaras, S.R. 1992. "A one-equation turbulence model for aerodynamic flows", AIAA-92-0439.

[69] Spalart, P.R., Jou., W.H., Strelets, M. and Allmaras, S.R. 1997. "Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach", 1st AFOSR Int. Conf. on DNS/LES Aug, 4-8 1997. In: Advances in DNS/LES, C. Liu and Z. Liu Eds., Greyden Press, Columbus, OH.

[70] Spalart, P.R., Deck, S., Shur, M.L., Squires, K.D., Strelets. M. and Travin, A. 2006. "A new version of detached eddy simulation, resistant to ambiguous grid densities", *Theor. Comput. Fluid Dynamics*, **20**, pp. 181-195.

[71] Squires, K. D., Forsythe, J. R., Strang, W. Z., Wurtsler, K. E., Tomaro, R.F., Grismer, M. J. and Spalart, P.R. 2002. "Progress on detached-eddy simulation of massively separated flows", AIAA 2002-1021.

[72] Travin, A., Shur, M., Strelets, M. and Spalart, P.R. 2002. "Physical and numerical upgrades in the detached eddy simulation of complex turbulence flows", in *Advances in LES of Complex Flows*, Kluwer Academic Press, pp. 239-254.

[73] Wilcox, D.C. 1998. "Turbulence Modeling for CFD". 2nd edition, DCW Industries, Inc.

Turbulence Specification

Turbulence specification is available as a boundary condition and as an initial condition. It allows you to specify the method for defining the initial turbulent velocity fluctuations in a continuum, at a velocity inlet, free stream, or mass flow inlet boundaries.

Properties

Method	Choose the method for defining turbulence.	
	Modified Turbulent Diffusivity	Direct specification of the value of the modified turbulent diffusivity ($\tilde{\nu}$). A Modified Diffusivity node is added to the Values node.
	Intensity and Length Scale	Calculation of the value of the modified turbulent diffusivity ($\tilde{\nu}$) from a turbulence intensity and length scale (Eqn. (520)). Do not use this method if the velocity field is initialized to zero (that is, V = 0). A Turbulent Intensity node and a Turbulent Length Scale node are added to the Values node. Additionally, a Turbulent Velocity Scale node appears in the Initial Conditions node.
	Turbulent Viscosity Ratio	Calculation of the value of the modified turbulent diffusivity ($\tilde{\nu}$) from the ratio of turbulent to laminar viscosity using Eqn. (493). A Turbulent Viscosity Ratio node is added to the Values node.

Modified Diffusivity

Modified diffusivity is available as a boundary condition, initial condition value, and region value. It allows you to specify the method for defining the initial turbulent velocity fluctuations in a continuum, for a region or at a boundary.

As a boundary value and initial condition value, modified diffusivity is added to the Values node by selecting `Modified Turbulent Diffusivity` as the Method in the Physics Conditions > Turbulence Specification node. It requires direct specification of the value of $\tilde{\nu}$. It is input as a scalar profile.

As a region value, modified diffusivity is added to the Values node by selecting `Modified Turbulent Diffusivity` as the Method in the Physics Conditions > Turbulence Specification node of the region. It requires direct specification of the value of $\tilde{\nu}$.

Properties of the Modified Diffusivity Region Value

Value	The value of the turbulence intensity.

See also: define value action (triggered by property customizer).

Modified Diffusivity Source

This region value is added to the Values node by activating the Turbulence Source Option region condition. It is entered as a scalar profile.

Spalart-Allmaras Turbulence Field Functions Reference

The following primitive field functions are made available to the simulation when the Spalart-Allmaras turbulence model is used. These objects all have the same base properties.

Modified Diffusivity	The transported variable $\tilde{\nu}$.

Using K-Epsilon Turbulence

This section explains the following:

- What are the K-Epsilon Turbulence models?
- What are the Non-Linear Constitutive relations?
- Exposing Reynolds Stresses for Standard Non-Linear K-Epsilon Models
- What is the K-Epsilon C3e coefficient?
- What is the 'Realizability coefficient?
- How to choose a model and wall treatment
- Understanding wall treatment
- Working with K-Epsilon solvers
- Formulation of the model equations

Also included are reference sections on the following topics as they relate to the K-Epsilon model:

- Boundary and region types
- Shared conditions
- Region and Boundary conditions and values
- Field functions

What Are the K-Epsilon Turbulence Models?

A K-Epsilon turbulence model is a two-equation model in which transport equations are solved for the turbulent kinetic energy k and its dissipation rate ε. Various forms of the K-Epsilon model have been in use for several decades, and it has become the most widely used model for industrial applications. Since the inception of the K-Epsilon model, there have been countless attempts to improve it. The most significant of these improvements have been incorporated into STAR-CCM+.

In its original form, the K-Epsilon turbulence model was applied with wall functions, but was later modified to use the following approaches for resolving the viscous sublayer:

- Low-Reynolds number
- Two-layer

STAR-CCM+ has a choice of eight different K-Epsilon turbulence models:

- Standard K-Epsilon

- Standard Two-Layer K-Epsilon
- Standard Low-Reynolds number
- Realizable K-Epsilon
- Realizable Two-Layer K-Epsilon
- Abe-Kondoh-Nagano low-Reynolds number
- Elliptic Blending K-Epsilon
- V2F low-Reynolds number

Non-linear constitutive relations are available with the three standard K-Epsilon turbulence model variants.

What Is a Low-Reynolds Number K-Epsilon Model?

A low-Reynolds number K-Epsilon model is one that can be used to resolve the viscous sublayer. The most common approach is to apply damping functions to some or all of the coefficients in the model (C_μ, $C_{\varepsilon 1}$ and $C_{\varepsilon 2}$). These damping functions modulate the coefficients as functions of a turbulence Reynolds number, often also incorporating the wall distance. Literally dozens of models incorporating damping functions have been proposed in the literature, and two have been incorporated into STAR-CCM+: the Standard low-Reynolds number model and the Abe-Kondoh-Nagano model. The V2F low-Reynolds number model, however, does not use wall distance or damping functions; it is claimed to be valid throughout the flow domain while automatically modeling the region close to solid surfaces.

What Is the Two-Layer Approach?

The two-layer approach, first suggested by Rodi [86], is an alternative to the low-Reynolds number approach that allows the K-Epsilon model to be applied in the viscous sublayer. In this approach, the computation is divided into two layers. In the layer next to the wall, the turbulent dissipation rate ε and the turbulent viscosity μ_t are specified as functions of wall distance. The values of ε specified in the near-wall layer are blended smoothly with the values computed from solving the transport equation far from the wall. The equation for the turbulent kinetic energy is solved in the entire flow. This explicit specification of ε and μ_t is arguably no less empirical than the damping function approach, and the results are often as good or better. In STAR-CCM+, the two-layer formulations work with either low-Reynolds number type meshes $y^+ \sim 1$ or wall-function type meshes $y^+ > 30$.

Several two-layer formulations have been proposed, and three have been implemented in STAR-CCM+, two for shear-driven flows and one for buoyancy-driven flows. Of the two shear-driven models, neither is clearly better. The Wolfstein model is activated by default, but the Norris-Reynolds model is offered as an option since some users find it preferable. Only use the Xu natural-convection model for buoyant flows.

What Is the Standard K-Epsilon Model?

The Standard K-Epsilon Model is a de facto standard version of the two-equation model that involves transport equations for the turbulent kinetic energy k and its dissipation rate ε. The transport equations are of the form suggested by Jones and Launder [80], with coefficients suggested by Launder and Sharma [82]. Some additional terms have been added to the model in STAR-CCM+ to account for effects such as buoyancy and compressibility. An optional non-linear constitutive relation is also provided.

Properties

Constitutive Relation — Selects the type of constitutive relation.

- **Linear** — Selects the linear constitutive relation. Use this selection for most simulations. See Eqn. (592).
- **Quadratic** — Selects the quadratic constitutive relation (see Eqn. (593)) and activates two sub-nodes:
  - Non-Linear Cmu Coefficients
  - Non-Linear Quadratic Coefficients
- **Cubic** — Selects the cubic constitutive relation (see Eqn. (594)) and activates three sub-nodes:
  - Non-Linear Cmu Coefficients
  - Non-Linear Quadratic Coefficients
  - Non-Linear Cubic Coefficients

Convection — Selects the convection scheme.

- **1st-order** — Selects the first-order upwind convection scheme.
- **2nd-order** — Selects the second-order upwind convection scheme.

Realizability Option — Activates realizability constraints on the time scale. See Eqn. (540).

- **None** — No realizability.
- **Durbin Scale Limiter** — Applies a minimum constraint on the time scale. This constraint activates the Realizability Coefficient sub-node for this model.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties with the *Expert* category. The values in that

category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

C1e	The coefficient $C_{\varepsilon 1}$, see Eqn. (541).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (541).
Cmu	The coefficient C_{μ}, see Eqn. (541).
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.

None	Neglects the term $C_{\varepsilon 3} G_b$.
Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (532).
Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.

Normal Stress Term

This property is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same even it this term is not explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k is poorly initialized, activating this option affects the solution convergence.

Activated

The linear stress tensor is modeled with Eqn. (592), and the linear production is modeled using Eqn. (526).

Deactivated

The linear stress tensor is modeled as $\mathbf{T}_{t,\,\mathrm{lin}} = -\rho\overline{\mathbf{v}'\mathbf{v}'} = 2\mu_t\mathbf{S}$, rather than the full expression of Eqn. (592). Furthermore, the linear stress that is used in the turbulent production is assumed to be $G_k = \mu_t S^2$ rather than the full expression of Eqn. (526).

Sarkar

The coefficient C_M, see Eqn. (534).

Secondary Gradients

Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Sigma_e	The coefficient σ_ε, see Eqn. (592).
Sigma_k	The coefficient σ_k, see Eqn. (592).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

What Is the Standard Two-Layer K-Epsilon Model?

The Standard Two-Layer K-Epsilon model combines the Standard K-Epsilon model with the two-layer approach. The coefficients in the models are identical, but model gains the added flexibility of an all y+ wall treatment.

Properties

Constitutive Relation — Selects the type of constitutive relation.

- **Linear** — Selects the linear constitutive relation of Eqn. (592). Use this selection most simulations.
- **Quadratic** — Selects the quadratic constitutive relation (see Eqn. (593)) and activates two sub-nodes:
  - Non-Linear Cmu Coefficients
  - Non-Linear Quadratic Coefficients
- **Cubic** — Selects the cubic constitutive relation (see Eqn. (594)) and activates three sub-nodes:
  - Non-Linear Cmu Coefficients
  - Non-Linear Quadratic Coefficients
  - Non-Linear Cubic Coefficients

Two-Layer Type — Selects the type of two-layer formulation.

- **Shear Driven (Wolfstein)** — Selects the two-layer formulation of Wolfstein [90]. Use this selection for flows where buoyancy forces do not dominate. See Eqns. (606) to (609).
- **Buoyancy Driven (Xu)** — Selects the Xu two-layer formulation. [91]. Use this selection for flows where buoyancy forces dominate. See Eqns. (610) to (619).
- **Shear Driven (Norris-Reynolds)** — Selects the two-layer formulation of Norris and Reynolds [85]. Use this selection for flows where buoyancy forces do not dominate. See Eqns. (613) to (618).

Convection — Selects the convection scheme.

- **1st-order** — Selects the first-order upwind convection scheme.
- **2nd-order** — Selects the second-order upwind convection scheme.

Yap Correction	Defines whether to include the Yap correction of Eqn. (536) in the transport equation for ε.	
	Activated	The Yap correction is included.
	Deactivated	The Yap correction is not included.
Realizability Option	Activates realizability constraints on the time scale. See Eqn. (540).	
	None	No realizability.
	Durbin Scale Limiter	Applies a minimum constraint on the time scale. This constraint activates the Realizability Coefficient sub-node for this model.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties with the *Expert* category. The values in that category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

C1e	The coefficient $C_{\varepsilon 1}$, see Eqn. (592).	
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (592).	
Cmu	The coefficient C_{μ}, see Eqn. (592).	
Two-Layer Delta ReY	The value of $\Delta \mathrm{Re}_y$ used in Eqn. (602).	
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.	
	None	Neglects the term $C_{\varepsilon 3} G_b$.
	Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (532).

Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.

Normal Stress Term

An explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.

This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ are absorbed into the pressure, and cause the pressure to be slightly different. However, the results are the same whether this term is explicitly included or not.

In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k happens to be poorly initialized, activating this option can affect the solution convergence.

Activated	The linear stress tensor is modeled with Eqn. (592), and the linear production is modeled using Eqn. (526).
Deactivated	The linear stress tensor is modeled as $\mathbf{T}_{t,\,\mathrm{lin}} = -\rho\overline{\mathbf{v}'\mathbf{v}'} = 2\mu_t\mathbf{S}$, rather than the full expression of Eqn. (592). Furthermore, the linear stress that is used in the turbulent production is assumed to be $G_k = \mu_t S^2$ rather than the full expression of Eqn. (526).

Two-Layer ReY*	The value of $\mathrm{Re}_y{}^*$ used in Eqn. (601).
Sarkar	The coefficient C_M, see Eqn. (534).
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces. On — Include both secondary gradients. Off — Exclude both secondary gradients. Interior Only — Include the interior secondary gradients only. Boundaries Only — Include the boundary secondary gradients only.
Sigma_e	The coefficient σ_ε, see Eqn. (592).
Sigma_k	The coefficient σ_k, see Eqn. (592).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Yap Cl	The Yap correction for the coefficient C_l
Yap Cw	The Yap correction for the coefficient C_w

What Is the Realizable K-Epsilon Model?

A successful recent development is the realizable K-Epsilon model that Shih and others [88] developed. This model contains a new transport equation for the turbulent dissipation rate ε. Also, a critical coefficient of the model, C_μ, is expressed as a function of mean flow and turbulence properties, rather than assumed to be constant as in the standard model. This procedure lets the model satisfy certain mathematical constraints on the normal stresses consistent with the physics of turbulence (realizability). The concept of a variable C_μ is also consistent with experimental observations in boundary layers.

The realizable K-Epsilon model is substantially better than the standard K-Epsilon model for many applications, and can generally be relied upon to give answers that are at least as accurate. Both the standard and realizable models have been implemented in STAR-CCM+ with a two-layer approach, which activates them to be used with fine meshes that resolve the viscous sublayer.

Properties

Curvature Correction Option

Alters the turbulent energy production term according to local rotation and vorticity rates. It supplies the effects of strong (streamline) curvature and frame-rotation. For more information, see Curvature Correction.

- On — Activates the Curvature Correction Option.
- Off — Leaves the Curvature Correction Option inactive. This is the default setting.

Convection

Selects the convection scheme.

- 1st-order — Selects the first-order upwind convection scheme.
- 2nd-order — Selects the second-order upwind convection scheme.

Expert Properties

Aside from $C_{\varepsilon 1}$, which is computed using Eqn. (550), the expert properties are identical to the expert properties of the standard K-Epsilon model. The coefficients in the model have different values as specified by Eqn. (552). The coefficient C_μ is computed dynamically from Eqn. (547), but the value that is specified is used to compute initial or boundary values.

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties within the *Expert* category. The values in that category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (552).	
Cmu	The coefficient C_{μ}, see Eqn. (547).	
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (543) is calculated.	
	None	Neglects the term $C_{\varepsilon 3} G_b$.
	Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (532).
	Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
	Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.

Normal Stress Term

This property is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same even if this term is not explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (88)). If k is poorly initialized, activating this option affects the solution convergence.

Activated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (526).
Deactivated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.

Sarkar

The coefficient C_M, see Eqn. (534).

Secondary Gradients

Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Sigma_e	The coefficient σ_ε, see Eqn. (552).
Sigma_k	The coefficient σ_k, see Eqn. (552).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

What Is the Realizable Two-Layer K-Epsilon Model?

The Realizable Two-Layer K-Epsilon model combines the Realizable K-Epsilon model with the two-layer approach. The coefficients in the models are identical, but model gains the added flexibility of an all y+ wall treatment.

Properties

Two-layer type	Selects a two-layer formulation.	
	Shear Driven (Wolfstein)	Selects the two-layer formulation of Wolfstein [90]. Use this selection for flows where buoyancy forces do not dominate. See Eqns. (606) to (609).
	Buoyancy Driven (Xu)	Selects the two-layer formulation of Xu and others. [91]. Use this selection for flows where buoyancy forces dominate. See Eqns. (610) to (619).
	Shear Driven (Norris-Reynolds)	Selects the two-layer formulation of Norris and Reynolds [85]. Use this selection for flows where buoyancy forces do not dominate. See Eqns. (613) to (618).
Curvature Correction Option	Alters the turbulent energy production term according to local rotation and vorticity rates. It supplies the effects of strong (streamline) curvature and frame-rotation. For more information, see Curvature Correction.	
	On	Activates the Curvature Correction Option.
	Off	Leaves the Curvature Correction Option inactive. This is the default setting.
Convection	Selects the convection scheme.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

Expert Properties

Aside from $C_{\varepsilon 1}$, which is computed using Eqn. (551), the expert properties are identical to the expert properties of the standard K-Epsilon model. The coefficients in the model have different values, as specified by Eqn. (553).

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties in the *Expert* category. The values in that category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (592).
Cmu	The coefficient C_{μ}, see Eqn. (592).
Two-Layer Delta ReY	The value of $\Delta \mathrm{Re}_y$ used in Eqn. (602).
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.

None	Neglects the term $C_{\varepsilon 3} G_b$.
Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (532).
Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.

Normal Stress Term

This property is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same even if this term is not explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k is poorly initialized, activating this option affects the solution convergence.

Activated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (526).
Deactivated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.

Two-Layer ReY*

The value of $\mathrm{Re}_y{}^*$ used in Eqn. (601).

Sarkar

The coefficient C_M, see Eqn. (534).

Secondary Gradients

Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.

	Boundaries Only	Include the boundary secondary gradients only.
Sigma_e	The coefficient σ_ε, see Eqn. (592).	
Sigma_k	The coefficient σ_k, see Eqn. (592).	
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.	
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.	

What Is the Standard Low-Reynolds Number K-Epsilon Model?

The low-Reynolds number model by Lien and others [83] is dubbed the "Standard Low-Reynolds Number K-Epsilon Model" because it has identical coefficients to the Standard K-Epsilon model, but provides more damping functions. These functions let it be applied in the viscous-affected regions near walls. This model is recommended for natural convection problems, particularly if the Yap correction is invoked.

Properties

Yap Correction	Defines whether to include the Yap correction of Eqn. (536) in the transport equation for ε.	
	Activated	The Yap correction is included.
	Deactivated	The Yap correction is not included.
Constitutive Relation	Selects the type of constitutive relation.	

Linear	Selects the linear constitutive relation. Use this selection for most simulations. See Eqn. (592).
Quadratic	Selects the quadratic constitutive relation (see Eqn. (593)) and activates two sub-nodes: • Non-Linear Cmu Coefficients • Non-Linear Quadratic Coefficients
Cubic	Selects the cubic constitutive relation (see Eqn. (594)) and activates three sub-nodes: • Non-Linear Cmu Coefficients • Non-Linear Quadratic Coefficients • Non-Linear Cubic Coefficients

Realizability Option	Activates realizability constraints on the time scale. See Eqn. (540).	
	None	No realizability.
	Durbin Scale Limiter	Applies a minimum constraint on the time scale. This selection activates the Realizability Coefficient sub-node for this model.

Convection	Selects the convection scheme.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties within the *Expert* category. The values in that

category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

Normal Stress Term

This property is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same even if this term is not explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k is poorly initialized, activating this option affects the solution convergence.

Activated

The linear stress tensor is modeled with Eqn. (592), and the linear production is modeled using Eqn. (526).

Deactivated

The linear stress tensor is modeled as $\mathbf{T}_{t,\,\mathrm{lin}} = -\rho\overline{\mathbf{v}'\mathbf{v}'} = 2\mu_t\mathbf{S}$, rather than the full expression of Eqn. (592). Furthermore, the linear stress that is used in the turbulent production is assumed to be $G_k = \mu_t S^2$ rather than the full expression of Eqn. (526).

Yap Cw

The Yap correction for the coefficient C_w

Yap Cl

The Yap correction for the coefficient C_l

Cd0 coefficient

The coefficient C_{d0}, see Eqn. (561).

Cd1 coefficient	The coefficient C_{d1}, see Eqn. (561).
Cd2 coefficient	The coefficient C_{d2}, see Eqn. (561).
C coefficient	The coefficient C, see Eqn. (556).
D coefficient	The coefficient D, see Eqn. (555).
E coefficient	The coefficient E, see Eqn. (555).
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.
	On: Include both secondary gradients.
	Off: Exclude both secondary gradients.
	Interior Only: Include the interior secondary gradients only.
	Boundaries Only: Include the boundary secondary gradients only.
C1e	The coefficient $C_{\varepsilon 1}$, see Eqn. (592).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (592).
Cmu	The coefficient C_{μ}, see Eqn. (592).
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.
	None: Neglects the term $C_{\varepsilon 3} G_b$.
	Boundary Layer Orientation: Computes $C_{\varepsilon 3}$ according to Eqn. (532).

Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.

Sarkar	The coefficient C_M, see Eqn. (534).
Sigma_e	The coefficient σ_ε, see Eqn. (592).
Sigma_k	The coefficient σ_k, see Eqn. (592).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

What Is the Abe-Kondoh-Nagano Low-Reynolds Number K-Epsilon Model?

The Abe-Kondoh-Nagano model [74] is known to work well for a wide range of complex flows. It has different coefficients than the Standard K-Epsilon model, and uses different damping functions than the Standard Low-Reynolds Number model. This model is a good choice for applications such as compact heat exchangers, where the Reynolds numbers are low but the flow is relatively complex.

Properties

Convection	Selects the convection scheme.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.
Realizability Option	Activates realizability constraints on the time scale. See Eqn. (540).	
	None	No realizability.
	Durbin Scale Limiter	Applies a minimum constraint on the time scale. This constraint activates the Realizability Coefficient sub-node for this model.

Expert Properties

The same expert properties as the expert properties of the standard K-Epsilon model are available.

Note: The coefficients in the model have different values, as specified by Eqn. (599).

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties in the *Expert* category. The values in that category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

C1e	The coefficient $C_{\varepsilon 1}$, see Eqn. (592).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (592).
Cmu	The coefficient C_{μ}, see Eqn. (592).
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.

None	Neglects the term $C_{\varepsilon 3} G_b$.
Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (532).
Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.

Normal Stress Term

This property is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same even if this term is not explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k is poorly initialized, activating this option affects the solution convergence.

Activated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (526).
Deactivated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.

Sarkar

The coefficient C_M, see Eqn. (534).

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.
	On — Include both secondary gradients.
	Off — Exclude both secondary gradients.
	Interior Only — Include the interior secondary gradients only.
	Boundaries Only — Include the boundary secondary gradients only.
Sigma_e	The coefficient σ_ε, see Eqn. (592).
Sigma_k	The coefficient σ_k, see Eqn. (592).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

What Is the Elliptic Blending K-Epsilon Model?

The concept of elliptic relaxation was proposed by Durbin [78] for Reynolds-stress models. The initial model required the solution of six additional transport equations, but this number was later reduced to a single additional equation. The model was later simplified by Manceau and Hanjalic [84] to make it more industry-friendly.

The elliptic relaxation model also led to the development of some two-equation eddy-viscosity models, starting with the $\overline{v^2} - f$ model. One of the most robust models, according to its authors, is the $BL\text{–}\overline{v^2} - k$ model of

Billard and Laurence [75]. Important modifications to this model make it truly robust for complex flow geometries.

The benefits of using the Billard - Laurence model are:

- An improvement on the existing realizable $k-\varepsilon$ model in terms of accuracy, especially in the near-wall region.
- An improvement in terms of stability over the SST $k-\omega$ model.

The formulation of the Elliptic Blending K-Epsilon model is given in Elliptic Blending K-Epsilon Model.

Properties

Convection	Selects the convection scheme.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.
Variable C2e Option	Defines whether to use the variable or the constant value of the coefficient $C_{\varepsilon 2}$ in the transport equation for ε (Eqn. (575)).	
	On	Activates the Variable C2e Option.
	Off	Leaves the Variable C2e Option inactive. This is the default setting ($C_{\varepsilon 2}$ = 1.83).

Expert Properties

Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.	
	None	Neglects the term $C_{\varepsilon 3} G_b$.
	Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (532).
	Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
	Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.
Cmu	The coefficient C_μ, see Eqn. (570).	
C1e	The coefficient $C_{\varepsilon 1}$, see Eqn. (575).	
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (580).	
C_T	The coefficient C_T, see Eqn. (572).	

Sigma_k	The coefficient σ_k, see Eqn. (574).
Sigma_e	The coefficient σ_ε, see Eqn. (575).
Sarkar	The coefficient C_M, see Eqn. (534).
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.
	On — Include both secondary gradients.
	Off — Exclude both secondary gradients.
	Interior Only — Include the interior secondary gradients only.
	Boundaries Only — Include the boundary secondary gradients only.
C1	The coefficient C_1, see Eqn. (582).
C2	The coefficient C_2, see Eqn. (582).
C3e	The coefficient $C_{\varepsilon 3}$, see Eqn. (575).
Cl	The coefficient C_L, see Eqn. (573).
Ceta	The coefficient C_η, see Eqn. (573).

Ck	The coefficient C_k, see Eqn. (573).
Sigma_phi	The coefficient σ_φ, see Eqn. (576).
Phi Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Alpha Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Normal Stress Term	This property is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation. This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same even it this term is not explicitly included. In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k is poorly initialized, activating this option affects the solution convergence.

Activated	The linear stress tensor is modeled with Eqn. (592), and the linear production is modeled using Eqn. (526).
Deactivated	The linear stress tensor is modeled as $\mathbf{T}_{t,\,\mathrm{lin}} = -\rho\overline{\mathbf{v}'\mathbf{v}'} = 2\mu_t\mathbf{S}$, rather than the full expression of Eqn. (592). Furthermore, the linear stress that is used in the turbulent production is assumed to be $G_k = \mu_t S^2$ rather than the full expression of Eqn. (526).

What Is the V2F Low-Reynolds Number K-Epsilon Model?

The V2F K-Epsilon model ([76], [93], [94]) is known to capture the near-wall turbulence effects more accurately, which is crucial for the accurate prediction of heat transfer, skin friction and flow separation. This model solves two more turbulence quantities, namely the normal stress function and the elliptic function, in addition to k and ε. This model is designed to handle wall effects in turbulent boundary layers and to accommodate non-local effects.

Properties

Realizability Option	Activates realizability constraints on the time scale.	
	Durbin Scale Limiter	Activates realizability constraints.
	None	Deactivates realizability constraints.
Convection	Selects the convection scheme.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties in the *Expert* category. The values in that category

reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

Normal Stress Term

This is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same whether or not this term is explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). If k is poorly initialized, activating this option affects the solution convergence.

Activated

The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (526).

Deactivated

The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.

A

The coefficient A, see Eqn. (591).

C1

The coefficient C_1, see Eqn. (591).

C2

The coefficient C_2, see Eqn. (591).

Cl

The coefficient C_L, see Eqn. (591).

CmuV2	The coefficient C_μ, see Eqn. (591).
V2 Minimum	The minimum value that the transported variable ϑ^2 is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
F Minimum	The minimum value that the variable f is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.
	On — Include both secondary gradients.
	Off — Exclude both secondary gradients.
	Interior Only — Include the interior secondary gradients only.
	Boundaries Only — Include the boundary secondary gradients only.
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (523) is calculated.
	None — Neglects the term $C_{\varepsilon 3} G_b$.
	Boundary Layer Orientation — Computes $C_{\varepsilon 3}$ according to Eqn. (532).
	Thermal Stratification — Computes $C_{\varepsilon 3}$ according to Eqn. (533).

Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the properties of the node.
Cmu	The coefficient C_{μ}, see Eqn. (569).
C1e	The coefficient $C_{\varepsilon 1}$, see Eqn. (591).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (591).
Ct	The coefficient C_t, see Eqn. (591).
Ceta	The coefficient C_{η}, see Eqn. (591).
Sigma_k	The coefficient σ_k, see Eqn. (591).
Sigma_e	The coefficient σ_{ε}, see Eqn. (591).
Sarkar	The coefficient C_M, see Eqn. (534).
Tke Minimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

Shared Values

Elliptic Function

When the V2F K-Epsilon model is selected, this boundary value and initial condition value is added to the Values node by selecting `K + Epsilon` as the Method in the Conditions > Turbulence Specification node. The elliptic function f is a redistribution term in the $\overline{\vartheta^2}$ equation.

Properties

This value is input as a scalar profile.

Normal Stress Function

When the V2F K-Epsilon model is selected, this boundary value and initial condition value, $\overline{\vartheta^2}$, is added to the Values node by selecting `K + Epsilon` as the Method in the Conditions > Turbulence Specification node.

Properties

This value is input as a scalar profile.

What Are the Non-Linear Constitutive Relations?

The Non-Linear Constitutive relations [83] account for anisotropy of the turbulence by extending the Boussinesq approximation to include non-linear functions of the strain and rotation tensors. These appear as

sub-nodes of the Standard K-Epsilon, Standard Two-Layer K-Epsilon, or Standard Low-Reynolds Number model node, depending on the Constitutive Relation property setting of that node.



They improve turbulence prediction by including effects such as secondary flows, streamline curvature and rotation. Two relations are available:

- Quadratic (Cmu and Quadratic coefficients)
- Cubic

What Is the Quadratic Relation?

By allowing the turbulence to have an anisotropic effect on the flow, the quadratic non-linear constitutive relation (see Eqn. (593)) improves the performance of the K-Epsilon Model, letting the model predict secondary flows better. It is represented in the object tree by the Cmu and Quadratic coefficients, each has Expert properties.

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties in the *Expert* category. The values in that category reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

Cmu Coefficient Expert Properties

Ca0	The coefficient C_{a0}, see Eqn. (598).
Ca1	The coefficient C_{a1}, see Eqn. (598).
Ca2	The coefficient C_{a2}, see Eqn. (598).
Ca3	The coefficient C_{a3}, see Eqn. (598).

Quadratic Coefficient Expert Properties

Cnl1	The coefficient C_{NL1}, see Eqn. (598).
Cnl2	The coefficient C_{NL2}, see Eqn. (598).
Cnl3	The coefficient C_{NL3}, see Eqn. (598).
Cnl6	The coefficient C_{NL6}, see Eqn. (598).
Cnl7	The coefficient C_{NL7}, see Eqn. (598).

What Is the Cubic Relation?

The Cubic Relation extends the non-linear constitutive relation of the Quadratic Relation by including cubic functions of the velocity gradients (see Eqn. (594)). These additional terms improve the sensitivity of the K-Epsilon Model to streamline curvature and rotation.

Cubic Relation Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change any properties in the *Expert* category. The values in that category

reflect both the design and the discretization approaches that have been optimized for accuracy and performance in the model. Tampering with them can diminish the effectiveness of the model.

Cnl4	The coefficient C_{NL4}, see Eqn. (598).
Cnl5	The coefficient C_{NL5}, see Eqn. (598).

Exposing Reynolds Stresses for Standard Non-Linear K-Epsilon Models

You can extract the Reynolds Stresses from simulations involving non-linear K-Epsilon models. To expose these stresses:

1. Select the relevant Standard K-Epsilon model node in the physics continuum and set Constitutive Relation to a non-linear value:
   - Quadratic

- Cubic


Servers
Star 1
simulation
Continua
Physics 1
Models
Coupled Energy
Coupled Flow
Gas
Gradients
High y+ Wall Treatment
Ideal Gas
K-Epsilon Turbulence
Reynolds-Averaged Navier-Stokes
Standard K-Epsilon
Steady
Three Dimensional
Turbulent
Reference Values
Standard K-Epsilon - Properties
Properties
Constitutive Relation
Linear
Realizability Option
Linear
Convection
Quadratic
Expert
Cubic
Normal Stress Term

2. Navigate to the Solvers > K-Epsilon Turbulence node, and activate the Temporary Storage Retained option.



The Reynolds Stresses field functions are exposed. You can see the relevant field functions in the Tools > Field Function node. See Reynolds Stress Transport Turbulence Field Functions Reference.



Note: If the simulation is two-dimensional, only the UU, VV, and UV stresses are exposed. All other stresses are equal to zero.

What Is the K-Epsilon C3e Coefficient?

The C3e node represents the term $C_{\varepsilon 3}$ in Eqn. (523) as a constant coefficient. The default setting of this object is `1.0`, and can be changed with a text entry in its C3e expert property. It is activated as a sub-node of any K-Epsilon model node when the Buoyancy Production of Dissipation expert property of the latter is set to `Constant Coefficient`.



C3e Coefficient Expert Properties

C3e	The coefficient $C_{\varepsilon 3}$, see Eqn. (523).

What Is the Realizability Coefficient?

Two-equation models predict an unexpectedly large growth of turbulent kinetic energy in stagnation point flows [94]. This growth can have an adverse effect on the rest of the flow solution. To overcome this effect, a lower limit on the turbulence time scale has been proposed that imposes the "realizability" constraint on the eddy-viscosity formula. This time scale

realizability constraint can be optionally applied to several of the K-Epsilon turbulence models in STAR-CCM+. To apply it, set the Realizability Option property of the parent model object to `Durbin Scale Limiter`.



Properties

Realizability Coefficient	Value of the realizability time scale coefficient, C_T, as used, for example, in Eqn. (540). The default value is 0.6.

The effect of this coefficient can be demonstrated by observing the difference between two solutions of flow over an airfoil. In the first solution that is shown below, the Realizability option has not been activated. The build-up of turbulence at the stagnation point is clearly observed. In the second solution, the Realizability option has been activated.





How Do I Choose a K-Epsilon Model and Wall Treatment?

The available K-Epsilon models and wall treatments are shown in the table below.

K-Epsilon Model	Wall Treatment
Standard	high-y^+
Standard Two-Layer	all-y^+
Realizable	high-y^+
Realizable Two-Layer	all-y^+
Standard Low-Reynolds Number	low-y^+, all-y^+
Abe-Kondo-Nagano Low-Reynolds Number	low-y^+, all-y^+
V2F Low-Reynolds Number	low-y^+, all-y^+

Below are some recommendations regarding the choice of model:

- The Standard K-Epsilon model and the Realizable K-Epsilon model are suitable for coarse meshes, where the wall-cell y^+ values are typically 30 and above. The Realizable model generally gives results at least as good as the Standard model, but typically better.
- The Standard Two-Layer K-Epsilon model and the Realizable Two-Layer K-Epsilon model offer the most mesh flexibility. They can be used with the same meshes as the high-Reynolds number versions. They give good results on fine meshes (that is, low-Reynolds number type or low-y^+ meshes), and also produces the least inaccuracies for intermediate meshes (that is, $1 < y^+ < 30$).

The Standard low-Reynolds number model, Abe-Kondoh-Nagano low-Reynolds number model, and V2F low-Reynolds number model are recommended for truly low-Reynolds number applications. The standard low-Reynolds number model is suitable for natural convection flows, for situations when it is desirable to have a low-Reynolds number version of the Standard K-Epsilon model, or when a Non-linear Constitutive Model is needed with a low-Reynolds number treatment. The Abe-Kondo-Nagano model is appropriate for other complex low-Reynolds number flows, where a linear constitutive model and damping-function approach is adequate. The V2F model is recommended for applications where it is crucial to capture the near-wall turbulence effects more accurately.

If there is uncertainty as to which turbulence model to use in a given situation, the Realizable Two-Layer K-Epsilon model would be a reasonable choice. If the mesh is coarse, it provides results that are quite close to the

version without the two-layer formulation. If the mesh is fine enough to resolve the viscous sublayer, it produces results similar to a low-Reynolds number model.

Understanding the K-Epsilon Wall Treatment

In the K-Epsilon models, the wall treatment performs the following functions:

- Specifies the reference velocity u^* to be used in the wall laws
- Computes a special value of turbulent production G_k in the wall cell
- Computes a special value of turbulent dissipation ε in the wall cell

The following wall treatments are available as appropriate to the specific K-Epsilon model:

- High y+ Wall Treatment
- Low y+ Wall Treatment
- All y+ Wall Treatment
- Two-Layer All y+ Wall Treatment

The equations that are used for each are detailed in the K-Epsilon model formulation, and a brief description of each follows.

What Is the High y+ Wall Treatment?

High y^+ Wall Treatment is essentially the classic wall-function approach, where wall shear stress, turbulent production and turbulent dissipation are all derived from equilibrium turbulent boundary layer theory.

For more information on the wall treatment formulation, see Wall Treatment.

What Is the Low y+ Wall Treatment?

Low y^+ Wall Treatment is consistent with low-Reynolds number models. This wall treatment assumes that the viscous sublayer is resolved and thus wall laws are not needed.

For more information on the wall treatment formulation, see Wall Treatment.

What Is the All y+ Wall Treatment?

All y^+ Wall Treatment is a hybrid approach that seeks to recover the behaviors of the other two wall treatments in the limit of fine or coarse meshes. This wall treatment was designed to give results similar to the low-y^+ treatment as $y^+ \to 0$ and to the high-y^+ treatment for $y^+ > 30$. It also gives reasonable results for intermediate meshes where the cell centroid falls in the buffer layer. This treatment is the recommended wall treatment where provided.

For more information on the wall treatment formulation, see Wall Treatment.

What Is the Two-Layer All y+ Wall Treatment?

Two-Layer All y^+ Wall Treatment is a formulation that is identical to the All y+ Wall Treatment, but contains a wall boundary condition for ε that is consistent with the two-layer formulation. This treatment is the recommended wall treatment where provided.

For more information on the wall treatment formulation, see Wall Treatment.

Working with K-Epsilon Solvers

This section describes the various solvers that are associated with the K-Epsilon family of models:

- Turbulence
- Turbulent viscosity

It also lists the residual monitors that are added to the simulation by these solvers.

What Is the K-Epsilon Turbulence Solver?

This solver controls the solution of the K-Epsilon model in all the continua for which the model is activated. For each transported variable k and ε the basic steps that are involved in the solution update are as follows:

1. Update boundary conditions.
2. Compute the reconstruction gradients and cell gradients.
3. Create the linear system of Eqn. (2591) using the discretization methods that are described in the Formulation section.
4. Compute the residual sum $R = \sum_{\text{cells}} |r|$ for monitoring convergence.

5. Solve the linear system.
6. Update the transported variable field.

Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the newly computed solution supplants the old solution. The default value is 0.8. For the theoretical background, see Eqn. (2592) in Implicit Iteration.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties, p. 6385.

What Is the K-Epsilon Turbulent Viscosity Solver?

This solver controls the update of the turbulent viscosity.

Let μ_t^n be the value of the previous iteration, and μ_t^{new} be the value that is computed during the current iteration. The update is controlled as follows:

$$\mu_t^{n+1} = \omega_\mu \mu_t^{new} + (1 - \omega_\mu)\mu_t^n$$

Properties

Maximum Ratio	The maximum ratio of turbulent to laminar viscosity (μ_t/μ) permitted during the update process. This ratio prevents unphysically high values of turbulent viscosity from occurring on the path to convergence.
Under-Relaxation Factor	An under-relaxation factor ω_μ for the update of turbulent viscosity. The default value is 1.0.

Using Residual Monitors

The K-Epsilon turbulence solvers add the following residual monitors to the top-level Monitors node of a simulation tree.

Tdr	Turbulent dissipation rate.
Tke	Turbulent kinetic energy.

The properties of all residual monitors are the same.

K-Epsilon Turbulence Formulation

In this section, the formulations of the K-Epsilon model variants are presented:

- Standard
- Realizable
- Standard Low-Reynolds Number
- Abe-Kondoh-Nagano Low-Reynolds Number
- Elliptic Blending K-Epsilon
- V2F Low-Reynolds Number
- Non-linear Constitutive Models
- Two-layer formulations (Standard and Realizable)
- Wall treatment formulations
- Formulations of boundary, region, and initial conditions

Also included in this section are:

- Nomenclature
- Bibliography

Standard K-Epsilon Model

Basic Transport Equations

The transport equations for the standard K-Epsilon model are:

$$\frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_k}\right)\nabla k \cdot d\mathbf{a} + \int_V [G_k + G_{nl} + G_b - \rho((\varepsilon - \varepsilon_0) + \Upsilon_M) + S_k] dV \tag{522}$$

$$\frac{d}{dt}\int_V \rho \varepsilon dV + \int_A \rho \varepsilon(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right)\nabla \varepsilon \cdot d\mathbf{a} + \int_V \frac{1}{T}[C_{\varepsilon 1}(G_k + G_{nl} + C_{\varepsilon 3}G_b) - C_{\varepsilon 2}\rho(\varepsilon - \varepsilon_0) + \rho\Upsilon_y + S_\varepsilon]\ dV \tag{523}$$

where:

S_k and S_ε are the user-specified source terms.

ε_0 is the ambient turbulence value in the source terms that counteracts turbulence decay [89].

For Eulerian multiphase cases, where more than one phase exists, the equations are:

$$\frac{d}{dt}\int_V \alpha_i\rho_i k_i dV + \int_A \alpha_i\rho_i k_i(\mathbf{v}_i - \mathbf{v}_g)\cdot d\mathbf{a} = \int_A \alpha_i\left(\mu_i + \frac{\mu_i^t}{\sigma_k}\right)\nabla k_i \cdot d\mathbf{a} + \int_V \alpha_i[G_i^k + G_i^b - \rho_i((\varepsilon_i - \varepsilon_0) + Y_i^M) + S_i^k + S_i^{kr}]dV + \sum_{i \to i}(m_{ij}k_j^{(ij)} - m_{ji}k_i) \tag{524}$$

$$\frac{d}{dt}\int_V \alpha_i\rho_i\varepsilon_i dV + \int_A \alpha_i\rho_i\varepsilon_i(\mathbf{v}_i - \mathbf{v}_g)\cdot d\mathbf{a} = \int_A \alpha_i\left(\mu_i + \frac{\mu_i^t}{\sigma_\varepsilon}\right)\nabla\varepsilon_i \cdot d\mathbf{a} + \int_V \alpha_i\frac{1}{T}[(C_{\varepsilon 1}(G_i^k + G_i^{nl} + C_{\varepsilon 3}G_i^b) - C_{\varepsilon 2}\rho_i(\varepsilon_i - \varepsilon_0)) + \rho_i Y_i^y + S_i^\varepsilon + S_i^{\varepsilon r}]\, dV + \sum_{i \to i}(m_{ij}\varepsilon_j^{(ij)} - m_{ji}\varepsilon_i) \tag{525}$$

where:

S_i^{kr} and $S_i^{\varepsilon r}$ are source terms for the continuous phase when modeling Eulerian particle induced turbulence.

α_i is the volume fraction of each phase, ***i***. For single phase flows α_i is automatically set to 1.

Turbulent Production

The production G_k is evaluated as:

$$G_k = \mu_t S^2 - \frac{2}{3}\rho k\nabla\cdot\mathbf{v} - \frac{2}{3}\mu_t(\nabla\cdot\mathbf{v})^2 \tag{526}$$

where $\nabla\cdot\mathbf{v}$ is the velocity divergence and S is the modulus of the mean strain rate tensor:

$$S = |\mathbf{S}| = \sqrt{2\mathbf{S}:\mathbf{S}^T} = \sqrt{2\mathbf{S}:\mathbf{S}} \tag{527}$$

and:

$$\mathbf{S} = \frac{1}{2}(\nabla\mathbf{v} + \nabla\mathbf{v}^T) \tag{528}$$

Non-linear Production

If a non-linear model constitutive relation is used, then the production due to the non-linear parts of the stress must be included. This production is termed the "non-linear" production G_{nl}. It is expressed as:

$$G_{nl} = \begin{cases} \nabla\mathbf{v} : (\mathbf{T}_{t,\,\mathrm{quad}} - \mathbf{T}_{t,\,\mathrm{lin}}) & \text{for the quadratic model} \\ \nabla\mathbf{v} : (\mathbf{T}_{t,\,\mathrm{cubic}} - \mathbf{T}_{t,\,\mathrm{lin}}) & \text{for the cubic model} \end{cases} \tag{529}$$

where $\mathbf{T}_{t,\,\mathrm{lin}}$, $\mathbf{T}_{t,\,\mathrm{quad}}$ and $\mathbf{T}_{t,\,\mathrm{cubic}}$ are given by Eqns. (592), (593) and (594), respectively.

Buoyancy Production

The production due to buoyancy G_b is evaluated as:

$$G_b = \beta \frac{\mu_t}{\sigma_t}(\nabla T \cdot \mathbf{g}) \tag{530}$$

where β is the coefficient of thermal expansion, $\mathbf{g}$ is the gravitational vector, ∇T is the temperature gradient vector, and σ_t is the turbulent Prandtl number.

For constant density flows using the Boussinesq approximation, you specify β. For ideal gases, the following relation is used:

$$\beta = -\frac{1}{\rho}\frac{\partial\rho}{\partial T} \tag{531}$$

The available literature is not clear as to the specification of the coefficient $C_{\varepsilon 3}$. By default, it is computed according to [79] as:

$$C_{\varepsilon 3} = \tanh\frac{|\mathbf{v}_b|}{|\mathbf{u}_b|} \tag{532}$$

where $\mathbf{v}_b$ is the velocity component parallel to $\mathbf{g}$, and $\mathbf{u}_b$ is the velocity component perpendicular to $\mathbf{g}$.

This formulation tends to set the coefficient to zero outside natural convection boundary layers. Alternatively, $C_{\varepsilon 3}$ can be taken as constant everywhere, or specified as follows:

$$C_{\varepsilon 3} = \begin{cases} 1 & \text{for } G_b \geq 0 \\ 0 & \text{for } G_b < 0 \end{cases} \tag{533}$$

Compressibility Modification

The dilatation dissipation Υ_M is modeled according to Sarkar as:

$$\Upsilon_M = \frac{C_M k \varepsilon}{c^2} \tag{534}$$

where c is the speed of sound and $C_M = 2$.

Yap Correction

The Yap correction Υ_y requires the computation of the wall distance. Therefore, it is only available when the standard K-Epsilon model is used together with the two-layer model.

It is defined as:

$$\Upsilon_y = C_w \frac{\varepsilon^2}{k} \max\left[\left(\frac{l}{l_\varepsilon} - 1\right)\left(\frac{l}{l_\varepsilon}\right)^2, 0\right] \tag{535}$$

and the length scales are defined:

$$l = \frac{k^{3/2}}{\varepsilon} \tag{536}$$

$$l_\varepsilon = C_l d \tag{537}$$

where d is the distance to the wall. The coefficients C_w and C_l are defined as:

$$C_l = 2.55 \quad , \quad C_w = 0.83 \tag{538}$$

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho C_\mu k T \tag{539}$$

Turbulent Time Scale

The turbulent time scale is computed as:

$$T = \begin{cases} \max\left(\frac{k}{\varepsilon}, C_t\sqrt{\frac{\nu}{\varepsilon}}\right) & \text{without realizable scale option} \\ \min\left(\max\left(\frac{k}{\varepsilon}, C_t\sqrt{\frac{\nu}{\varepsilon}}\right), \frac{C_T}{\sqrt{3}C_\mu S}\right) & \text{with realizable scale option} \end{cases} \tag{540}$$

In Eqn. (540), C_T is the realizable time scale coefficient and has a default value of 0.6.

Model Coefficients

$$C_{\varepsilon 1} = 1.44 \ , \ C_{\varepsilon 2} = 1.92 \ , \ C_\mu = 0.09 \ , \ \sigma_k = 1.0 \ , \ \sigma_\varepsilon = 1.3 \ , \ C_t = 1 \tag{541}$$

Realizable K-Epsilon Model

Basic Transport Equations

The transport equations for the realizable K-Epsilon model are:

$$\begin{gathered} \frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \\ \int_A \left(\mu + \frac{\mu_t}{\sigma_k}\right)\nabla k \cdot d\mathbf{a} + \int_V [f_c G_k + G_b - \rho((\varepsilon - \varepsilon_0) + \Upsilon_M) + S_k] dV \end{gathered} \tag{542}$$

$$\begin{gathered} \frac{d}{dt}\int_V \rho \varepsilon dV + \int_A \rho \varepsilon(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \\ \int_A \left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right)\nabla \varepsilon \cdot d\mathbf{a} + \int_V \left[f_c C_{\varepsilon_1} S\varepsilon + \frac{\varepsilon}{k}(C_{\varepsilon_1} C_{\varepsilon_3} G_b) - \frac{\varepsilon}{k+\sqrt{\nu\varepsilon}} C_{\varepsilon_2}\rho(\varepsilon - \varepsilon_0) + S_\varepsilon\right] dV \end{gathered} \tag{543}$$

where:

S_k and S_ε are the user-specified source terms.

ε_0 is the ambient turbulence value in the source terms that counteracts turbulence decay.

f_c is the curvature correction factor.

For Eulerian multiphase cases, where more than one phase exists, the equations are:

$$\frac{d}{dt}\int_V \alpha_i\rho_i k_i dV + \int_A \alpha_i\rho_i k_i(\mathbf{v}_i - \mathbf{v}_g)\cdot d\mathbf{a} = \int_A \alpha_i\left(\mu_i + \frac{\mu_i^t}{\sigma_k}\right)\nabla k_i \cdot d\mathbf{a} + \int_V \alpha_i[f_i^c G_i^k + G_i^b - \rho_i((\varepsilon_i - \varepsilon_0) + Y_i^M) + S_i^k + S_i^{kr}]dV + \sum_{j\neq i}(m_{ij}k_j^{(ij)} - m_{ji}k_i) \tag{544}$$

$$\frac{d}{dt}\int_V \alpha_i\rho_i \varepsilon_i dV + \int_A \alpha_i\rho_i \varepsilon_i(\mathbf{v}_i - \mathbf{v}_g)\cdot d\mathbf{a} = \int_A \alpha_i\left(\mu_i + \frac{\mu_i^t}{\sigma_\varepsilon}\right)\nabla \varepsilon_i \cdot d\mathbf{a} + \int_V \alpha_i\left[f_i^c C_{\varepsilon_1} S\varepsilon + \frac{\varepsilon_i}{k_i}(C_{\varepsilon_1}C_{\varepsilon_3}G_i^b) - \frac{\varepsilon_i}{k_i + \sqrt{\nu\varepsilon_i}}C_{\varepsilon_2}\rho_i(\varepsilon_i - \varepsilon_0) + S_i^\varepsilon + S_i^{\varepsilon r}\right]dV + \sum_{j\neq i}(m_{ij}\varepsilon_j^{(ij)} - m_{ji}\varepsilon_i) \tag{545}$$

where:

the subscript i denotes the different phases.

α_i is the volume fraction of each phase i. For single phase flows α_i is automatically set to 1.

S_i^{kr} and $S_i^{\varepsilon r}$ are source terms for the continuous phase when modeling Eulerian particle induced turbulence.

f_i^c is the curvature correction factor for each phase.

Turbulent Production

The turbulent production term G_k is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (526)).

Buoyancy Production

The production term G_b is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (530)).

Compressibility Modification

The dilatation dissipation Υ_M is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (534)).

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho C_\mu \frac{k^2}{\varepsilon} \tag{546}$$

where the coefficient C_μ is no longer constant as with the standard K-Epsilon model, but is instead:

$$C_\mu = \frac{1}{A_0 + A_s U^{(*)} \frac{k}{\varepsilon}} \tag{547}$$

where:

$$U^{(*)} = \sqrt{\mathbf{S}:\mathbf{S} + \mathbf{W}:\mathbf{W}} \tag{548}$$

where $\mathbf{S}$ is the strain rate tensor as defined in Eqn. (528), and $\mathbf{W}$ is the rotation rate tensor that is given by Eqn. (595).

The coefficients are:

$$\begin{aligned} A_s &= \sqrt{6}\cos\phi \\ \phi &= \frac{1}{3}\mathrm{acos}(\sqrt{6}W) \\ W &= \frac{S_{ij}S_{jk}S_{ki}}{\sqrt{S_{ij}S_{ij}}^3} \\ A_0 &= 4.0 \end{aligned} \tag{549}$$

where S is the modulus of the mean strain rate tensor (see Eqn. (527)).

Model Coefficients

The coefficient $C_{\varepsilon 1}$ is defined as follows:

$$C_{\varepsilon 1} = \max(0.43, \frac{\eta}{5+\eta}) \tag{550}$$

where:

$$\eta = \frac{Sk}{\varepsilon} \tag{551}$$

The rest of the coefficients defining the model have the following values:

$$C_{\varepsilon 2} = 1.9 \quad , \quad \sigma_k = 1.0 \quad , \quad \sigma_\varepsilon = 1.2 \tag{552}$$

Standard (Lien) Low-Reynolds Number K-Epsilon Model

Basic Transport Equations

The transport equations for the Standard Low-Reynolds Number K-Epsilon model are:

$$\frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_k}\right)\nabla k \cdot d\mathbf{a} + \int_V [G_k + G_b - \rho((\varepsilon - \varepsilon_0) + \Upsilon_M) + S_k]dV \tag{553}$$

$$\frac{d}{dt}\int_V \rho \varepsilon dV + \int_A \rho \varepsilon(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right)\nabla \varepsilon \cdot d\mathbf{a} + \int_V \frac{1}{T}[C_{\varepsilon 1}(G_k + G_{nl} + G' + C_{\varepsilon 3}G_b) - C_{\varepsilon 2}f_2\rho(\varepsilon - \varepsilon_0) + \rho\Upsilon_y + S_\varepsilon]\ dV \tag{554}$$

where S_k and S_ε are the user-specified source terms, ε_0 is the ambient turbulence value in source terms that counteract turbulence decay, and G' is an additional production term that is given by:

$$G' = D\ f_2\left(G_k + 2\mu\frac{k}{y^2}\right)\exp(-ERe_y^2) \tag{555}$$

f_2 is a damping function that is defined as:

$$f_2 = 1 - C\exp(Re_T^2) \tag{556}$$

where:

$$Re_y = \frac{\sqrt{k}y}{\nu} \tag{557}$$

and:

$$Re_t = \frac{k^2}{\varepsilon\nu} \tag{558}$$

The coefficients C and E are assigned default values of 0.3 and 0.00375, respectively.

In the original model, the value of the coefficient D is $D = C_{\varepsilon 2}/C_{\varepsilon 1} \approx 1.3$. However, DNS data for low-Reynolds number channel flow suggests that better results are obtained with $D = 1$, which is the default value that is used in STAR-CCM+.

Turbulent Production

The turbulent production term G_k is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (526)).

Non-linear Production

The non-linear production term G_{nl} is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (529)).

Buoyancy Production

The production term G_b is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (530)).

Compressibility Modification

The dilatation dissipation Υ_M is defined identically to the corresponding term in the standard K-Epsilon model (see Eqn. (534)).

Yap Correction

The Yap correction Υ_y is defined identically to the corresponding term in the standard k-ε model (see Eqn. (535)).

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho C_\mu f_\mu k T \tag{559}$$

The time scale T is computed as in Eqn. (540) and f_μ is a damping function that is given by:

$$f_\mu = 1 - \exp[-(C_{d0}\sqrt{Re_y} + C_{d1} Re_y + C_{d2} Re_y{}^2)] \tag{560}$$

where:

$$C_{d0} = 0.091 \quad , \quad C_{d1} = 0.0042 \quad , \quad C_{d2} = 0.00011 \tag{561}$$

Eqn. (560) is slightly different in form to the original function by Lien *et al.* [83], which was written

$$f_\mu = [1 - \exp(-0.0198 Re_y)]\left(1 + \frac{5.29}{Re_y}\right) \tag{562}$$

Although these functions have a slightly different asymptotic behavior as $Re_y \to 0$, they behave otherwise almost identically.

Model Coefficients

The model coefficients are identical to the corresponding ones in the standard k-ε model, see Eqn. (541).

Non-linear Constitutive Relations

Since Lien and others [83] proposed the non-linear constitutive relations used in the standard k-ε model, these relations are also available in the Lien k-ε model.

Abe-Kondoh-Nagano Low-Reynolds Number K-Epsilon Model

Basic Transport Equations

The transport equations for the Abe-Kondoh-Nagano k-ε model are:

$$\frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_k}\right)\nabla k \cdot d\mathbf{a} + \int_V [G_k + G_b - \rho((\varepsilon - \varepsilon_0) + \Upsilon_M) + S_k] dV \tag{563}$$

$$\frac{d}{dt}\int_V \rho \varepsilon dV + \int_A \rho \varepsilon(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right)\nabla \varepsilon \cdot d\mathbf{a} + \int_V \frac{1}{T}\frac{\varepsilon}{k}[C_{\varepsilon 1}(G_k + C_{\varepsilon 3}G_b) - C_{\varepsilon 2} f_2 \rho(\varepsilon - \varepsilon_0) + \rho \Upsilon_y + S_\varepsilon]\ dV \tag{564}$$

where S_k and S_ε are the user-specified source terms, ε_0 is the ambient turbulence value in source terms that counteract turbulence decay, and f_2 is a damping function that is given by:

$$f_2 = \left[1 - \exp\left(-\frac{Re_\varepsilon}{3.1}\right)\right]^2 \left\{1 - 0.3\exp\left[-\left(\frac{Re_t}{6.5}\right)^2\right]\right\} \tag{565}$$

Re_t is as defined in Eqn. (558) and:

$$Re_\varepsilon = \frac{(\nu\varepsilon)^{1/4} y}{\nu} \tag{566}$$

Turbulent Production

The turbulent production term G_k is defined identically to the corresponding term in the standard k-ε model (see Eqn. (526)).

Buoyancy Production

The production term G_b is defined identically to the corresponding term in the standard k-ε model (see Eqn. (530)).

Compressibility Modification

The dilatation dissipation Υ_M is defined identically to the corresponding term in the standard k-ε model (see Eqn. (534)).

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho C_\mu f_\mu k\ T \tag{567}$$

The time scale T is computed as in Eqn. (540) and f_μ is a damping function that is given by:

$$f_\mu = \left[1 - \exp\left(-\frac{Re_\varepsilon}{14}\right)\right]^2 \left\{1 + \frac{5}{Re_t^{3/4}} \exp\left[-\left(\frac{Re_t}{200}\right)^2\right]\right\} \tag{568}$$

Model Coefficients

$$C_{\varepsilon 1} = 1.5\ ,\ C_{\varepsilon 2} = 1.9\ ,\ C_\mu = 0.09\ ,\ \sigma_k = 1.4\ ,\ \sigma_\varepsilon = 1.4\ ,\ C_t = 1 \tag{569}$$

Elliptic Blending K-Epsilon Model

The details of the derivation of the baseline model can be found in the publication by Billard and Laurence [75]. This section describes the transport equation and explains the changes that were introduced in the model.

The turbulent eddy-viscosity is computed as:

$$\mu_t = \rho C_\mu k \varphi \tau_{lim} \tag{570}$$

where the turbulent time-scale τ and length-scale L are given by:

$$\tau = \sqrt{\left(\frac{k}{\varepsilon}\right)^2 + C_t\frac{\nu}{\varepsilon}} \tag{571}$$

$$\tau_{lim} = \max\left(\tau, \frac{C_T}{\sqrt{3}C_\mu\varphi S}\right) \tag{572}$$

$$L = C_L\sqrt{\left(\frac{k}{\varepsilon}\right)^3 + C_\eta^2\sqrt{\frac{\nu^3}{\varepsilon}}} \tag{573}$$

Transport Equations

The transport equations for the four variables k, ε, $\varphi = \overline{v^2}/k$, and α are:

$$\frac{d\rho k}{dt} = P_k + P_b - \rho\varepsilon + \frac{\partial}{\partial x_j}\left[\left(\frac{\mu}{2} + \frac{\mu_t}{\sigma_k}\right)\frac{\partial k}{\partial x_j}\right] \tag{574}$$

$$\frac{d\rho\varepsilon}{dt} = \frac{1}{\tau}(C_{\varepsilon 1}P_k + C_{\varepsilon 1}C_{\varepsilon 3}P_b - \rho E - C_{\varepsilon 2}\,\varepsilon) + \frac{\partial}{\partial x_j}\left[\left(\frac{\mu}{2} + \frac{\mu_t}{\sigma_\varepsilon}\right)\frac{\partial \varepsilon}{\partial x_j}\right] \tag{575}$$

$$\frac{d\rho\varphi}{dt} = \rho(1-\alpha^3)f_w + \rho\alpha^3 f_h - (P_k + P_b)\frac{\varphi}{k} + \frac{\partial}{\partial x_j}\left[\left(\frac{\mu}{2} + \frac{\mu_t}{\sigma_\varphi}\right)\frac{\partial \varphi}{\partial x_j}\right] \tag{576}$$

$$L^2\frac{\partial^2\alpha}{\partial x_j\partial x_j} = \alpha - 1 \tag{577}$$

with

$$S = \sqrt{2S_{ij}S_{ij}} \text{ and } S_{ij} = \frac{1}{2}\left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right)$$

The turbulent production P_k and extra-dissipation term E are given by:

$$P_k = \mu_t S^2 - \frac{2}{3}\frac{\partial U_i}{\partial x_i}\left(\rho k + \mu_t\frac{\partial U_j}{\partial x_j}\right) \tag{578}$$

$$E = -2C_k\nu\nu_t(1-\alpha)^3\frac{k}{\varepsilon}\left(\frac{\partial\left\|2S_{ij}n_j\right\|n_k}{\partial x_k}\right)^2 \tag{579}$$

with $n_i = \nabla\alpha/\|\nabla\alpha\|$ representing the wall-normal direction. The buoyancy production P_b uses the same definition as that used in other turbulence models.

One important difference from other two-equation models is the definition of the coefficient of the destruction term, $C_{\varepsilon 2}$. This term is defined as a function of the turbulent kinetic energy gradient, to reduce its value in the defect layer:

$$C_{\varepsilon 2} = 1.83 + \alpha^3(1 - 1.83)\tanh\left[\left|\frac{\partial}{\partial x_j}\left(\frac{\mu_t}{\rho\sigma_k}\frac{\partial k}{\partial x_j}\right)\varepsilon^{-1}\right|^{3/2}\right] \tag{580}$$

The production and destruction terms in Eqn. (576) are:

$$f_w = -\frac{\varphi\varepsilon}{2k} \tag{581}$$

$$f_h = -\frac{1}{\tau}\left(C_1 - 1 + C_2\frac{P_k + P_b}{\rho\varepsilon}\right)\left(\varphi - \frac{2}{3}\right) \tag{582}$$

Model Coefficients

The constants that are used in the Elliptic Blending K-Epsilon model are:

$$\begin{gathered}
\sigma_k = 1 \\
\sigma_\varepsilon = 1.5 \\
\sigma_\varphi = 1 \\
C_\mu = 0.22 \\
C_t = 4 \\
C_{\varepsilon 1} = 1.44 \\
C_{\varepsilon 3} = 1 \\
C_k = 2.3 \\
C_L = 0.164 \\
C_1 = 1.7 \\
C_2 = 0.9 \\
C_\eta = 75 \\
C_T = 1
\end{gathered} \tag{583}$$

Changes From Base Model

The main changes from the original model proposed by Billard and Laurence [75] are:

- Removal of an additional cross-diffusion term in the φ transport equation (Eqn. (576)), for stability reasons

- The constant in the definition of $C_{\varepsilon 2}$, which acts on its decay in the defect layer, has been increased from 0.4 to 1.
- The extra-dissipation term E, which accounts for viscous wall-effects, was moved from the k-equation back to the ε-equation. This change is for stability reasons and to avoid unwanted relaminarization effects in some scenarios.

Boundary Conditions and All-y+ Wall Treatment

The following set of boundary conditions are used near walls:

$$\lim_{d \to 0} k = 0 \ , \quad \lim_{d \to 0} \varepsilon = \frac{\nu k}{d^2} \ , \quad \lim_{d \to 0} \varphi = 0 \ , \quad \lim_{d \to 0} \alpha = 0$$

where d is the distance from the first cell-center to the wall.

The Elliptic Blending K-Epsilon model all-y+ wall-treatment is a combination of the baseline STAR-CCM+ version and the version given in Billard and Laurence [75].

V2F Low-Reynolds Number Model

Basic Transport Equations

The transport equations for the V2F k-ε model are:

$$\frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g) \cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_k}\right) \nabla k \cdot d\mathbf{a} + \int_V [G_k + G_b - \rho(\varepsilon - \varepsilon_0) + \Upsilon_M + S_k] dV \tag{584}$$

$$\frac{d}{dt}\int_V \rho \varepsilon dV + \int_A \rho \varepsilon(\mathbf{v}-\mathbf{v}_g) \cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right) \nabla \varepsilon \cdot d\mathbf{a} + \int_V \left[\rho C_{\varepsilon_1}\left(1 + a\sqrt{\frac{k}{\vartheta^2}}\right)\frac{G_k}{T} - \rho C_{\varepsilon_2}(\varepsilon - \varepsilon_0)\frac{\varepsilon}{k} + S_\varepsilon\right] dV \tag{585}$$

$$\frac{d}{dt}\int_V \rho\overline{\vartheta^2}dV + \int_A \rho\overline{\vartheta^2}(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_k}\right)\nabla\overline{\vartheta^2}\cdot d\mathbf{a} + \int_V \left[\rho\ \min\left\{k\ f, -\frac{1}{T}\left[(C_1-6)\overline{\vartheta^2} - \frac{2k}{3}(C_1-1)\right] + C_2 G_k\right\} - \frac{6\rho\overline{\vartheta^2}\varepsilon}{k} + S_{\overline{\vartheta^2}}\right]dV \tag{586}$$

$$\int_A L^2\nabla f\cdot d\mathbf{a} + \int_V \left[-f + \frac{1}{T}(C_1-1)\left(\frac{2}{3} - \frac{\overline{\vartheta^2}}{k}\right) + C_2\frac{C_{\mu_{\vartheta^2}}\overline{\vartheta^2}\mathbf{S}^2}{\varepsilon} + \frac{5\overline{\vartheta^2}}{kT} + S_f\right]dV = 0 \tag{587}$$

where S_k, S_ε, $S_{\overline{\vartheta^2}}$ and S_f are the user-specified source terms, and ε_0 is the ambient turbulence value in source terms that counteract turbulence decay.

Turbulent Production

The turbulent production term G_k is defined identically to the corresponding term in the standard k-ε model (see Eqn. (526)).

Buoyancy Production

The production term G_b is defined identically to the corresponding term in the standard k-ε model (see Eqn. (530)).

Compressibility Modification

The dilatation dissipation Υ_M is defined identically to the corresponding term in the standard k-ε model (see Eqn. (534)).

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho\ T\ \min\left(C_{\mu_{\text{std}}}k, C_{\mu_{\vartheta^2}}\vartheta^2\right) \tag{588}$$

Turbulent Time and Length Scales

The turbulent time scale is computed as:

$$\Gamma = \begin{cases} \max\left(\frac{k}{\varepsilon}, C_t\sqrt{\frac{\nu}{\varepsilon}}\right) & \text{without realizable scale option} \\ \max\left(\min\left(\frac{k}{\varepsilon}, \frac{\alpha k}{\sqrt{3}C_{\mu_{\vartheta^2}}\overline{\vartheta^2}S}\right), C_t\sqrt{\frac{\nu}{\varepsilon}}\right) & \text{with realizable scale option} \end{cases} \tag{589}$$

The turbulent length scale is defined as:

$$L = C_L \max\left(\min\left(\frac{k^{3/2}}{\varepsilon}, \frac{k^{3/2}}{\sqrt{3}C_{\mu_{\vartheta^2}}\overline{\vartheta^2}S}\right), C_\eta\left(\frac{\nu^3}{\varepsilon}\right)^{1/4}\right) \tag{590}$$

Model Coefficients

$$\begin{gathered} a = 0.045 \quad , \quad \alpha = 0.6 \quad , \quad C_{\varepsilon 1} = 1.4 \\ C_{\varepsilon 2} = 1.9 \quad , \quad C_{\mu_{\text{std}}} = 0.09 \quad , \quad C_{\mu_{\vartheta^2}} = 0.22 \quad , \quad \sigma_k = 1.0 \quad , \quad \sigma_\varepsilon = 1.3 \\ C_1 = 1.4 \quad , \quad C_2 = 0.3 \quad , \quad C_\eta = 70.0 \quad , \quad C_t = 6.0 \quad , \quad C_L = 0.23 \end{gathered} \tag{591}$$

Nonlinear Constitutive Relations

The Boussinesq approximation that is used by default in all K-Epsilon models implies a linear constitutive relation:

$$\mathbf{T}_{t,\,\text{lin}} = -\rho\overline{\mathbf{v}'\mathbf{v}'} = 2\mu_t\mathbf{S} - \frac{2}{3}(\mu_t\nabla\cdot\mathbf{v} + \rho k)\mathbf{I} \tag{592}$$

where $\mathbf{S}$ is the strain rate tensor that is defined in Eqn. (528).

The non-linear constitutive model [83] accounts for anisotropy by adding non-linear functions of the strain and rotation tensors.

The quadratic model is written:

$$\begin{aligned} \mathbf{T}_{t,\,\text{quad}} = \mathbf{T}_{t,\,\text{lin}} &- 4C_1\mu_t\frac{k}{\varepsilon}\left[\mathbf{S}\cdot\mathbf{S} - \frac{1}{3}\mathbf{I}(\mathbf{S}:\mathbf{S})\right] - 4C_2\mu_t\frac{k}{\varepsilon}(\mathbf{W}\cdot\mathbf{S} + \mathbf{S}\cdot\mathbf{W}^T) \\ &- 4C_3\mu_t\frac{k}{\varepsilon}\left[\mathbf{W}\cdot\mathbf{W}^T - \frac{1}{3}\mathbf{I}(\mathbf{W}:\mathbf{W}^T)\right] \end{aligned} \tag{593}$$

and the cubic model is:

$$\mathbf{T}_{t,\,\text{cubic}} = \mathbf{T}_{t,\,\text{quad}} - 8C_4\mu_t\frac{k^2}{\varepsilon^2}[(\mathbf{S}\cdot\mathbf{S})\cdot\mathbf{W} + \mathbf{W}^T\cdot(\mathbf{S}\cdot\mathbf{S})] - 8C_5\mu_t\frac{k^2}{\varepsilon^2}(\mathbf{S}:\mathbf{S} - \mathbf{W}:\mathbf{W}^T)\mathbf{S}$$

(594)

where $\mathbf{S}$ is the strain rate tensor as defined in Eqn. (528) and $\mathbf{W}$ is the vorticity tensor that is given by:

$$\mathbf{W} = \frac{1}{2}(\nabla\mathbf{v} - \nabla\mathbf{v}^T) \tag{595}$$

The coefficients in these equations are:

$$\begin{aligned}
C_\mu &= \frac{C_{a0}}{C_{a1} + C_{a2}S + C_{a3}W} \\
C_1 &= \frac{C_{NL1}}{(C_{NL6} + C_{NL7}S^3)C_\mu} \\
C_2 &= \frac{C_{NL2}}{(C_{NL6} + C_{NL7}S^3)C_\mu} \\
C_3 &= \frac{C_{NL3}}{(C_{NL6} + C_{NL7}S^3)C_\mu} \\
C_4 &= C_{NL4}C_\mu^2 \\
C_5 &= C_{NL5}C_\mu^2
\end{aligned} \tag{596}$$

where:

$$\begin{aligned}
S &= \frac{k}{\varepsilon}\sqrt{2\mathbf{S}:\mathbf{S}} \\
W &= \frac{k}{\varepsilon}\sqrt{2\mathbf{W}:\mathbf{W}}
\end{aligned} \tag{597}$$

and:

$$\begin{gathered}
C_{a0} = 0.667 \quad , \quad C_{a1} = 1.25 \quad , \quad C_{a2} = 1 \quad , \quad C_{a3} = 0.9 \\
C_{NL1} = 0.75 \quad , \quad C_{NL2} = 3.75 \quad , \quad C_{NL3} = 4.75 \quad , \quad C_{NL4} = -10 \\
C_{NL5} = -2 \quad , \quad C_{NL6} = 1000 \quad , \quad C_{NL7} = 1
\end{gathered} \tag{598}$$

The variable coefficient C_μ given by Eqn. (596) replaces the constant value in the relation for turbulent viscosity, Eqn. (539).

Two-Layer Formulation

The two-layer model blends a one-equation model, which solves for k but prescribes ε algebraically with distance from the wall, with the two-equation k-ε model. It is available with both the standard and realizable k-ε models in STAR-CCM+.

The two-layer model is parameterized as a length scale function, $l_\varepsilon = f(y, \mathrm{Re}_y)$, and a turbulent viscosity ratio function, $\mu_t/\mu = f(\mathrm{Re}_y)$, where the turbulent Reynolds number Re_y is given by Eqn. (557).

The dissipation rate that is computed from the two-layer model is simply:

$$\varepsilon = \frac{k^{3/2}}{l_\varepsilon} \tag{599}$$

As per the suggestion of Jongen [81], the following wall-proximity indicator is used to combine the two-layer formulation with the full two-equation model:

$$\lambda = \frac{1}{2}\left[1 + \tanh\left(\frac{\mathrm{Re}_y - \mathrm{Re}_y{}^*}{A}\right)\right] \tag{600}$$

where $\mathrm{Re}_y{}^*$ defines the limit of applicability of the two-layer formulation. A value of $\mathrm{Re}_y{}^* = 60$ is used in STAR-CCM+.

The constant A determines the width of the wall-proximity indicator. By defining a width such that the value of λ is within 1% of its far-field value for a given variation of $\Delta\mathrm{Re}_y$, one obtains the following relation between A and $\Delta\mathrm{Re}_y$:

$$A = \frac{|\Delta\mathrm{Re}_y|}{\mathrm{atanh}\,0.98} \tag{601}$$

In STAR-CCM+, $\Delta\mathrm{Re}_y = 10$.

The turbulent viscosity from the k-ε model is blended with the two-layer value as follows:

$$\mu_t = \lambda\mu_t\big|_{k-\varepsilon} + (1-\lambda)\mu\left(\frac{\mu_t}{\mu}\right)_{2\mathrm{layer}} \tag{602}$$

Since cell values of ε can be specified by setting the coefficients of the solution matrix to satisfy:

$$\frac{a_p}{\omega}\Delta\varepsilon_p = a_p(\varepsilon_p^{n+1}\big|_{\mathrm{specifed}} - \varepsilon_p^n) \tag{603}$$

the discretized transport equation for ε can be combined with this equation to obtain:

$$\frac{a_p}{\omega}\Delta\varepsilon_p + \sum_n a_n \lambda \Delta\varepsilon_n = \lambda\left(b - a_p\varepsilon_p^n - \sum a_n\varepsilon_n^n\right) + (1-\lambda)a_p(\varepsilon_p^{n+1}\Big|_{2\text{layer}} - \varepsilon_p^n) \tag{604}$$

Wolfstein Model

The one-equation model of Wolfstein [90] is as follows:

$$l_\varepsilon = c_l y\left[1 - \exp\left(-\frac{\mathrm{Re}_y}{A_\varepsilon}\right)\right] \tag{605}$$

$$A_\varepsilon = 2c_l \tag{606}$$

$$c_l = \kappa C_\mu^{-3/4} \tag{607}$$

where $C_\mu = 0.09$ and $\kappa = 0.42$.

The turbulent viscosity is:

$$\frac{\mu_t}{\mu} = \mathrm{Re}_y C_\mu^{1/4}\kappa\left[1 - \exp\left(-\frac{\mathrm{Re}_y}{A_\mu}\right)\right] \tag{608}$$

where $A_\mu = 70$.

Xu Model

Xu and others [91] have recently developed a two-layer model specifically for natural convection. This model can be formulated as follows:

$$l_\varepsilon = \frac{8.8y}{1 + 10/y_v{}^* + 5.15\times 10^{-2} y_v{}^*}\frac{1}{\sqrt{\overline{v^2}/k}} \tag{609}$$

$$\frac{\mu_t}{\mu} = \frac{0.544 y_v{}^*}{1 + 5.025\times 10^4 y_v{}^{*1.65}} \tag{610}$$

$$\overline{v^2}/k = 7.19\times 10^{-3}\mathrm{Re}_y - 4.33\times 10^{-5}\mathrm{Re}_y^2 + 8.8\times 10^{-8}\mathrm{Re}_y^3 \tag{611}$$

$$y_v{}^* = \mathrm{Re}_y\sqrt{\overline{v^2}/k} \tag{612}$$

Norris-Reynolds Model

The model by Norris and Reynolds [85] is as follows:

$$l_\varepsilon = \frac{c_l \mathrm{Re}_y y}{\mathrm{Re}_y + c_\varepsilon} \tag{613}$$

where:

$$c_\varepsilon = 5.3 \tag{614}$$

$$c_l = \kappa C_\mu^{-3/4} \tag{615}$$

$$\kappa = 0.42 \tag{616}$$

$$\mathrm{Re}_y = \frac{\sqrt{k}y}{\nu} \tag{617}$$

$$\frac{\mu_t}{\mu} = \mathrm{Re}_y \kappa C_\mu^{1/4}\left[1 - \exp\left(-\frac{\mathrm{Re}_y}{A_\mu}\right)\right] \tag{618}$$

where C_μ is defined by the turbulence model (normally 0.09) and $A_\mu = 50.5$.

Wall Treatment

The following equations present the wall-treatment formulations. In general, the high-y^+ formulation is equivalent to the traditional wall-function approach, and the low-y^+ formulation is equivalent to the traditional low-Reynolds number approach where no modeling beyond the assumption of laminar flow is necessary in the wall cells.

For the all-y^+ formulation, a blending function g is defined in terms of the wall-distance Reynolds number:

$$g = \exp\left(-\frac{Re_y}{11}\right) \tag{619}$$

and $Re_y = \sqrt{k}y/\nu$, a wall-distance-based Reynolds number.

Reference Velocity

$$u^* = \begin{cases} \sqrt{C_\mu^{1/2} k} & ; \quad \text{high-}y^+ \\ \sqrt{g\nu u/y + (1-g)C_\mu{}^{1/2} k} & ; \quad \text{low-}y^+ \\ \sqrt{g\nu u/y + (1-g)C_\mu{}^{1/2} k} & ; \quad \text{all-}y^+ \end{cases} \tag{620}$$

The reference velocity u^* is not needed in the calculation for low-y^+ simulations (using a low-Reynolds number model). However, it is still needed when computing the wall y^+ for post-processing. Since the assumption is that $u^* \equiv \sqrt{\tau_w/\rho} \approx \sqrt{\nu u/y}$ gives an inaccurate result for coarse meshes, the hybrid approach is used for the low-y^+ wall treatment as well.

Wall-Cell Production

$$G_k = \begin{cases} \frac{1}{\mu}\left(\rho u^* \frac{u}{u^+}\right)^2 \frac{\partial u^+}{\partial y^+} & ; \quad \text{high-}y^+ \\ \mu_t S^2 & ; \quad \text{low-}y^+ \\ g\mu_t S^2 + (1-g)\frac{1}{\mu}\left(\rho u^* \frac{u}{u^+}\right)^2 \frac{\partial u^+}{\partial y^+} & ; \quad \text{all-}y^+ \end{cases} \tag{621}$$

Wall-Cell Dissipation

$$\varepsilon = \begin{cases} \frac{u^{*3}}{\kappa y} & ; \quad \text{high-}y^+ \\ \frac{2\nu k}{y^2} & ; \quad \text{low-}y^+ \\ g\frac{2\nu k}{y^2} + (1+g)\frac{u^{*3}}{\kappa y} & ; \quad \text{all-}y^+ \\ \frac{k^{3/2}}{l_\varepsilon} & ; \quad \text{two-layer all-}y^+ \end{cases} \tag{622}$$

Boundary, Region, and Initial Conditions

Wall Boundaries

At walls, a Neumann boundary condition is used for the turbulent kinetic energy k, that is, $\partial k/\partial n|_w = 0$. The turbulent dissipation rate ε is specified in the wall cells according to the appropriate method in the wall treatment.

Flow, Region, and Initial Conditions

When defining values for flow boundary, region and initial conditions, you have several choices for specifying the turbulent kinetic energy and turbulent dissipation rate:

- Enter the values directly.

- Have them derived from a specified turbulence intensity (I) and length scale (L) using the relations:

$$k \approx \frac{3}{2}(I\mathrm{v})^2$$
$$\varepsilon \approx \frac{C_\mu^{3/4} k^{3/2}}{L} \tag{623}$$

  where v is the turbulent velocity scale and C_μ is a coefficient of the model.

- Have them derived from a specified turbulence intensity (I) and turbulent viscosity ratio μ_t/μ using the relations:

$$k \approx \frac{3}{2}(I\mathrm{v})^2$$
$$\varepsilon \approx \frac{\rho C_\mu k^2}{(\mu_t/\mu)\ \mu} \tag{624}$$

  where v is the turbulent velocity scale and C_μ is a coefficient of the model.

K-Epsilon Turbulence Nomenclature

The following is a guide to the terms used in this section.

Scalars

c	speed of sound
d	distance to the nearest wall
k	turbulent kinetic energy
l	turbulent length scale
t	time
u^*	reference velocity that is used in wall functions
u^+	wall-parallel velocity non-dimensionalized with u^*
G_k	turbulent production
G_{nl}	non-linear turbulent production
G_b	turbulent production due to buoyancy
I	turbulence intensity
L	turbulent length scale
Re_y	wall-distance-based turbulent Reynolds number
T	temperature
S	strain rate parameter
V	cell volume
v	velocity magnitude
β	coefficient of thermal expansion

$\underline{\varepsilon}$	turbulent dissipation rate
ϑ^2	normal stress function
f	elliptic function
ρ	density
μ	dynamic viscosity
μ_t	turbulent viscosity
σ_t	turbulent Prandtl number
ν	kinematic viscosity
σ_k, σ_ε	turbulent Schmidt numbers
Υ_M	dilation dissipation
Υ_y	Yap correction term
C_T	realizable time scale coefficient

Tensors or Column Vectors

$\mathbf{I}$	identity matrix
S	strain rate tensor
$\hat{\mathbf{S}}$	modified strain rate tensor
$\mathbf{T}$	viscous stress tensor
$\mathbf{W}$	vorticity tensor
$\hat{\mathbf{W}}$	modified vorticity tensor

Cartesian Vectors

$\mathbf{a}$	face area vector
g	gravitational vector
$\mathbf{v}$	velocity
$\mathbf{u}_b$	velocity perpendicular to the gravitational vector
$\mathbf{v}_b$	velocity parallel to the gravitational vector
$\mathbf{v}_g$	grid velocity

Subscripts

lin	linear
quad	quadratic
cubic	cubic
t	turbulent

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Divergence of vector quantity:

$$\nabla \cdot \mathbf{v} = \frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a} \cdot \mathbf{b} = a_i b_i$$

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S}:\mathbf{T} = S_{ij} T_{ji}$$

K-Epsilon Turbulence Bibliography

[74] Abe, K., Kondoh, T., and Nagano, Y. 1994. "A new turbulence model for predicting fluid flow and heat transfer in separating and reattaching flows --- 1. Flow field calculations", *Int. J. Heat Mass Transfer*, **37**, pp. 139-151.

[75] Billard, F. and Laurence, D. 2012. "A robust $k-\varepsilon$-v^2/k elliptic blending turbulence model applied to near-wall, separated and buoyant flows", *International Journal of Heat and Fluid Flow* **33** (1), pp. 45–58.

[76] Davidson, L., Nielsen, P.V., and Sveningsson, A. 2003. "Modifications of the $\overline{v^2}-f$ Model for Computing the Flow in a 3D Wall Jet", *Turbulence, Heat and Mass Transfer*, **4**, pp. 577-584.

[77] Durbin, P.A. 1991. "Near-wall turbulence closure modeling without damping functions", *Theoretical and Computational Fluid Dynamics*, **3**, pp. 1-13.

[78] Durbin, P.A. 1993. "A Reynolds stress model for near-wall turbulence", *Journal of Fluid Mechanics* **249** (1), pp 465–498.

[79] Henkes, R.A.W.M., van der Flugt, F.F., and Hoohendoorn, C.J. 1991. "Natural convection in a square cavity calculated with low-Reynolds number turbulence models", *Int. J. Heat Mass Transfer*, **34**, pp. 1543-1557.

[80] Jones, W.P., and Launder, B.E. 1972. "The Prediction of Laminarization with a Two-Equation Model of Turbulence", *Int. J. Heat and Mass Transfer*, **15**, pp. 301-314.

[81] Jongen, T. 1998. "Simulation and Modeling of Turbulent Incompressible Flows", Ph.D. Thesis, Lausanne EPFL.

[82] Launder, B.E., and Sharma, B.I. 1974. "Application of the Energy Dissipation Model of Turbulence to the Calculation of Flow Near a Spinning Disc", Letter in Heat and Mass Transfer, vol. 1, no. 2, pp 131-138.

[83] Lien, F.S., Chen, W.L., and Leschziner, M.A. 1996. "Low-Reynolds number eddy-viscosity modelling based on non-linear stress-strain/vorticity relations", Proc. 3rd Symp. on Engineering Turbulence Modelling and Measurements, 27-29 May, Crete, Greece.

[84] Manceau, R. and Hanjalic, K. 2002. "Elliptic blending model: A new near-wall reynolds-stress turbulence closure" *Physics of Fluids* **14**, pp 744.

[85] Norris, L.H., and Reynolds, W.C. 1975. "Turbulent channel flow with a moving wavy boundary", Report No. FM–10, Department of Mechanical Engineering, Stanford University, USA.

[86] Rodi, W. 1991. "Experience with Two-Layer Models Combining the k-e Model with a One-Equation Model Near the Wall", 29th Aerospace Sciences Meeting, January 7-10, Reno, NV, AIAA 91-0216.

[87] Sarkar, S., and Balakrishnan, L. 1990. "Application of a Reynolds-stress turbulence model to the compressible shear layer", ICASE Report 90-18, NASA CR 182002.

[88] Shih, T.-H., Liou, W.W., Shabbir, A., Yang, Z. and Zhu, J. 1994. "A New *k*-ε Eddy Viscosity Model for High Reynolds Number Turbulent Flows -- Model Development and Validation", NASA TM 106721.

[89] Spalart, P.R. and Rumsey, C.L. 2007. "Effective Inflow Conditions for Turbulence Models in Aerodynamic Calculations", *AIAA Journal*, **45**(10), pp. 2544-2553.

[90] Wolfstein, M. 1969. "The velocity and temperature distribution in one-dimensional flow with turbulence augmentation and pressure gradient", *Int. J. Heat Mass Transfer*, **12**, pp. 301-318.

[91] Xu, W., Chen, Q., and Nieuwstadt, F.T.M. 1998. "A new turbulence model for near-wall natural convection", *Int. J. Heat Mass Transfer*, **41**, pp. 3161-3176.

[92] Yap, C. 1987. "Turbulent heat and momentum transfer in recirculating and impinging flows", Ph.D. Thesis, Faculty of Technology, University of Manchester, U.K.

[93] Lien, L. S., Kalitzin, G., Durbin, P. A. 1998 "RANS modeling for compressible and transitional flows", *Center for Turbulence Research - Proceedings of the Summer Program 1998*

[94] Durbin, P. A. 1996, "On the k-e stagnation point anomaly", *Int. J. Heat and Fluid Flow* **17**, 89-90.

Region, Boundary, and Phase Types Reference

This section defines the region, boundary, and phase conditions that are required for each boundary type that is supported by the K-Epsilon turbulence model:

- Free stream
- Mass flow inlet
- Pressure outlet
- Stagnation inlet
- Velocity inlet

In addition, information is provided about the values for each region type:

- Porous

Note: Region and boundary types that do *not* require setting any conditions or values are *not* listed.

Free Stream

The free stream boundary type requires the following boundary conditions when the K-Epsilon turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Turbulent dissipation rate.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulence intensity.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulent length scale.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Mass Flow Inlet

The mass flow inlet boundary type requires the following boundary conditions when the K-Epsilon turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Turbulent dissipation rate.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulence intensity.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.

- Turbulent length scale.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Pressure Outlet

The pressure outlet boundary type requires the following boundary conditions when the K-Epsilon turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Turbulent dissipation rate.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulence intensity.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulent length scale.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Stagnation Inlet

The stagnation inlet boundary type requires the following boundary conditions when the K-Epsilon turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Turbulent dissipation rate.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.

- Turbulence intensity.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulent length scale.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Velocity Inlet

The velocity inlet boundary type requires the following boundary conditions when the K-Epsilon turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Turbulent dissipation rate.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulence intensity.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulent length scale.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Porous

The porous region type requires the following conditions:

- Turbulence specification

The following region values are also required:

- Turbulent dissipation rate.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulence intensity.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `K + Epsilon`.
- Turbulent length scale.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This value is only available if the Method property of the Conditions Turbulence Specification is set to Intensity + Viscosity Ratio.

Shared Conditions Reference

This section lists condition nodes that can be common among regions, boundaries, phases, and initial conditions:

- Turbulence specification

Turbulence Specification

Allows you to specify the method for defining the turbulence profile at a boundary or phase condition.

Properties

Method Choose the method for defining turbulence.

K + Epsilon	Direct specification of the values of k and ε. A Turbulent Kinetic Energy node and a Turbulent Dissipation Rate node are added to the Values node.
Intensity + Length Scale	Calculation of k and ε from a specified turbulence intensity I and length scale L using Eqn. (623). Do not use this method if the velocity field is initialized to zero. A Turbulent Intensity node and a Turbulent Length Scale node are added to the Values node. Additionally, a Turbulent Velocity Scale node appears in the Initial Conditions node.
Intensity + Viscosity Ratio	Calculation of k and ε from the turbulence intensity (I) and the ratio of turbulent to laminar viscosity (μ_t/μ) using Eqn. (624). Do not use this method if the velocity field is initialized to zero. A Turbulent Viscosity Ratio node and a Turbulence Intensity node are added to the Values node. Additionally, a Turbulent Velocity Scale node appears in the Initial Conditions node.

Regions, Boundary, and Phase Conditions and Values Reference

Set the following region, boundary, and phase conditions for the K-Epsilon turbulence models, depending on the selected boundary type:

- Turbulence specification

The boundary type elects the boundary conditions, based on the models that have been activated in the continuum. The boundary conditions that are required for each boundary type when the K-Epsilon turbulence model is activated are listed here.

In addition, set the following values for these models, again depending on the selected boundary type:

- Turbulent dissipation rate
- Turbulence intensity

- Turbulent kinetic energy
- Turbulent length scale
- Turbulent viscosity ratio

Boundary values are elected by the boundary type and/or by methods that are selected in corresponding nodes in the boundary conditions node.

K-Epsilon Turbulence Field Functions Reference

The following primitive field functions are made available to the simulation when the K-Epsilon turbulence model is used. These objects all have the same base properties.

Turbulent Kinetic Energy	The transported variable k.
Turbulent Dissipation Rate	The transported variable ε.

Using K-Omega Turbulence

This section explains the following:

- What are the K-Omega Turbulence models?
- What are the Low Re damping parameters?
- What are the compressibility parameters?
- What are the curvature correction parameters?
- What is the realizability coefficient?
- How to choose a model and wall treatment
- Understanding wall treatment
- Working with K-Omega solvers
- Formulation of the model equations

Also included are reference sections on the following topics as they relate to the K-Omega model:

- Boundary and region types
- Shared conditions and values
- Boundary conditions and values
- Region conditions and values
- Field functions

What Are the K-Omega Turbulence Models?

The K-Omega model is a two-equation model that is an alternative to the K-Epsilon model. The transport equations that are solved are for the turbulent kinetic energy k and a quantity called ω, which is defined as the specific dissipation rate, that is, the dissipation rate per unit turbulent kinetic energy ($\omega \sim \varepsilon / k$).

The book by D.C. Wilcox [95] is the most comprehensive reference on the K-Omega model, discussing the origin of the model, comparing it to other models, and presenting the latest version of the model. As the originator of the K-Omega model, Wilcox touts the superiority of his model over the K-Epsilon model, and the superiority of the Omega transport equation over other scale equations.

One reported advantage of the K-Omega model over the K-Epsilon model is its improved performance for boundary layers under adverse pressure gradients. Perhaps the most significant advantage, however, is that it may be applied throughout the boundary layer, including the

viscous-dominated region, without further modification. Furthermore, the standard K-Omega model can be used in this mode without requiring the computation of wall distance.

The biggest disadvantage of the K-Omega model, in its original form, is that boundary layer computations are sensitive to the values of ω in the free stream. This translates into extreme sensitivity to inlet boundary conditions for internal flows, a problem that does not exist for the K-Epsilon models. The versions of the model included in STAR-CCM+ have been modified in an attempt to address this shortcoming.

There are three versions of the K-Omega model in STAR-CCM+:

- Standard K-Omega model
- SST K-Omega model
- SST K-Omega detached eddy model

What Is the Standard K-Omega Model?

Wilcox revised his original model in 1998, and then in 2006 [102], to account for several perceived deficiencies in the original version (1988). These revisions include:

- A revised set of model coefficients
- Two corrections to account for sensitivity to free-stream/inlet conditions, both based on products of ∇k and $\nabla\omega$. The first model, introduced in 1998, is a modification of the turbulent kinetic energy equation. The more recent revision [102] is the introduction of a cross-diffusion term that is similar to that used in the SST K-Omega model
- A correction to improve the free-shear-flow spreading rates of the model
- A compressibility correction
- Low-Reynolds number corrections that allow the K-Omega model to be better applied in the prediction of low-Reynolds number or transitional flows

However, the validation results published in Wilcox's book are typically for two-dimensional, primarily parabolic, flows. Until further validations for complex flows are widely published, the corrections should be used with caution. Therefore, each correction has been included as an option in STAR-CCM+. To regain the 1988 form of the standard K-Omega model, all these corrections should be omitted and the model coefficients changed to their original values. To use the most recent updates of the Standard K-Omega model [102], select the Cross-diffusion Limiter in the Free-Shear Modification, activate the Vortex-Stretching Modification, and set the

Realizability Option to Durbin Scale Limiter.

Standard K-Omega Model Properties

Free-Shear Modification

Specifies what method to use for the free-shear modification.

None	The modification not included.
Dissipation Limiter	The modifications detailed in Eqns. (637) and (638) are included.
Cross-diffusion Limiter	The modification detailed in Eqns. (639) and (640) are included.

Vortex Stretching Modification

Specifies whether the vortex modification as detailed in Eqns. (634) through (636) is incorporated.

Activated	Include the modification.
Deactivated	Exclude the modification.

Convection

For guidance on selecting a convection scheme, see Convection Term, p. 6343.

1st-order	Selects the first-order convection.
2nd-order	Selects the second-order convection.

Curvature Correction Option

Alters the turbulent energy production term according to local rotation and vorticity rates. It supplies the effects of strong (streamline) curvature and frame-rotation. For more information, see Curvature Correction.

On	Activates the Curvature Correction Option.
Off	Leaves the Curvature Correction Option inactive. This is the default setting.

Realizability Option

Enables realizability constraints on the time scale. See Eqn. (649).

None	No realizability.

Durbin Scale Limiter	Applies a minimum constraint on the time scale. This activates the Realizability Coefficient subnode for this model.

Compressibility Correction	Specifies whether the dilatation dissipation correction that is described in Eqns. (656) through (658) is incorporated. This correction is irrelevant in incompressible flows.	
	Activated	The modification is incorporated. This activates the Compressibility Parameters subnode for this model.
	Deactivated	The modification is not included.

Low Re Damping Modification	Specifies whether the low-Reynolds number modification as detailed in Eqns. (651) through (655) is incorporated.	
	Activated	The modification is incorporated. This activates the Low Re Damping Parameters subnode for this model.
	Deactivated	The modification is not included.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, we recommend that you do not change properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

Alpha	The coefficient α, see Eqn. (633).
Beta	The coefficient β, see Eqn. (629).
BetaStar	The coefficient β^*, see Eqn. (628).

Normal Stress Term

This is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation.
This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same whether this term is explicitly included.
In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). Activating this option may affect the solution convergence if k happens to be poorly initialized.

Activated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (630)
Deactivated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.

Sdr Minimum

The minimum value that the transported variable ω is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

Secondary Gradients

Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Sigma_k

The coefficient σ_k, see Eqn. (628).

Sigma_w	The coefficient σ_ω, see Eqn. (629).
TkeMinimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

What Is the SST K-Omega Model?

The problem of sensitivity to free-stream/inlet conditions was addressed by Menter [96], who recognized that the ε transport equation from the standard K-Epsilon model could be transformed into an ω transport equation by variable substitution. The transformed equation looks similar to the one in the standard K-Omega model, but adds an additional non-conservative cross-diffusion term containing the dot product $\nabla k \cdot \nabla \omega$. Inclusion of this term in the ω transport equation potentially makes the K-Omega model give identical results to the K-Epsilon model. Menter suggested using a blending function (which includes functions of wall distance) that would include the cross-diffusion term far from walls, but not near the wall. This approach effectively blends a K-Epsilon model in the far-field with a K-Omega model near the wall. Purists may object that the blending function crossover location is arbitrary, and could obscure some critical feature of the turbulence. Nevertheless, the fact remains that this approach cures the biggest drawback to applying the K-Omega model to practical flow simulations.

In addition, Menter also introduced a modification to the linear constitutive equation and dubbed the model containing this modification the SST (shear-stress transport) K-Omega model. The SST model has seen fairly wide application in the aerospace industry, where viscous flows are typically resolved and turbulence models are applied throughout the boundary layer.

SST K-Omega Model Properties

Convection

For guidance on selecting a convection, see Convection Term, p. 6343.

1st-order	Selects the first-order convection.
2nd-order	Selects the second-order convection.

Curvature Correction Option

Alters the turbulent energy production term according to local rotation and vorticity rates. It supplies the effects of strong (streamline) curvature and frame-rotation. For more information, see Curvature Correction.

On	Activates the Curvature Correction Option.
Off	Leaves the Curvature Correction Option inactive. This is the default setting.

Realizability Option

Enables realizability constraints on the time scale. See Eqn. (665).

None	No realizability.
Durbin Scale Limiter	Applies a minimum constraint on the time scale. This activates the Realizability Coefficient subnode for this model.

Compressibility Correction

Specifies whether the dilatation dissipation correction that is described in Eqns. (679) through (681) is incorporated. This correction is irrelevant in incompressible flows.

Activated	The modification is incorporated. This activates the Compressibility Parameters subnode for this model.
Deactivated	The modification is not included.

Low Re Damping Modification

Specifies whether the low-Reynolds number modification as detailed in Eqns. (676) through (678) is incorporated.

Activated	The modification is incorporated. This activates the Low Re Damping Parameters subnode for this model.
Deactivated	The modification is not included.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, we recommend that you do not make changes to any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

a1	The coefficient a_1 (SST model only), see Eqn. (668).
Beta1	The coefficient β_1, see Eqn. (670).
Beta2	The coefficient β_2, see Eqn. (671).
BetaStar	The coefficient β^*, see Eqn. (628).
Kappa	The coefficient κ, see Eqns. (670) and (671).
Normal Stress Term	This is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation. This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same whether this term is explicitly included. In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). Activating this option may affect the solution convergence if k happens to be poorly initialized.

Activated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (630)
Deactivated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.

Sdr Minimum

The minimum value that the transported variable ω is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

Secondary Gradients

Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Neglect both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Sigma_k1

The coefficient σ_{k1}, see Eqn. (670).

Sigma_k2

The coefficient σ_{k2}, see Eqn. (671).

Sigma_w1

The coefficient $\sigma_{\omega 1}$, see Eqn. (670)

Sigma_w2

The coefficient $\sigma_{\omega 2}$, see Eqn. (671)

TkeMinimum

The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

What Is the SST K-Omega Detached Eddy Model?

Detached eddy simulation (DES) [69] is a modeling approach that combines features of Reynolds-Averaged (RANS) simulation in part of the flow and large eddy simulation (LES) in unsteady, separated regions. The delay factor introduced by Menter and Kuntz [99] has been adopted for the DES formulation of the SST K-Omega model. This modification enhances the ability of the model to distinguish between LES and RANS regions on computational meshes where spatial refinement could give rise to ambiguous behavior. In addition, the improved delayed detached eddy simulation (IDDES) formulation of Shur *et al.* [66] is available as an option. This combines DDES with an improved RANS-LES hybrid model aimed at wall modeling in LES when the grid resolution supports it.

Inflow Boundaries and Initialization

The Synthetic Eddy Method is implemented to automatically provide turbulent eddies across inflow boundaries and provide an initial perturbated flow field.

SST K-Omega Detached Eddy Model Properties

Formulation Option Selects the DES formulation to use.

- **DDES** Uses DDES formulation.
- **IDDES** Uses IDDES formulation. This activates the IDDES Coefficients subnode for this model.

Convection For guidance on selecting a convection, see Convection Term, p. 6343.

- **1st-order** Selects the first-order convection.
- **2nd-order** Selects the second-order convection.

Curvature Correction Option Alters the turbulent energy production term according to local rotation and vorticity rates. It supplies the effects of strong (streamline) curvature and frame-rotation. For more information, see Curvature Correction.

- **On** Activates the Curvature Correction Option.
- **Off** Leaves the Curvature Correction Option inactive. This is the default setting.

Realizability Option Enables realizability constraints on the time scale. See Eqn. (665).

- **None** No realizability.
- **Durbin Scale Limiter** Applies a minimum constraint on the time scale. This activates the Realizability Coefficient subnode for this model.

Compressibility Correction Specifies whether the dilatation dissipation correction that is described in Eqns. (679) through (681) is incorporated. This correction is irrelevant in incompressible flows.

- **Activated** The modification is incorporated. This activates the Compressibility Parameters subnode for this model.
- **Deactivated** The modification is not included.

Low Re Damping Modification — Specifies whether the low-Reynolds number modification as detailed in Eqns. (676) through (678) is incorporated. This option is incompatible with the Gamma ReTheta Transition model.

- Activated — The modification is incorporated. This activates the Low Re Damping Parameters subnode for this model.
- Deactivated — The modification is not included.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, we recommend that you do not change any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

CdesKw — The DES length scale coefficient for the model's K-Omega branch, see Eqn. (686).

CdesKe — The DES length scale coefficient for the model's K-Epsilon branch, see Eqn. (686).

CdesTimeLim — The DES blended upwind scheme time scale coefficient C_{desT}, calculated as for the Spalart-Allmaras turbulence model in Eqn. (513).

a1 — The coefficient a_1 (SST model only), see Eqn. (668).

Kappa — The coefficient κ, see Eqns. (670) and (671).

BetaStar — The coefficient β^*, see Eqn. (628).

Beta1 — The coefficient β_1, see Eqn. (670).

Sigma_k1 — The coefficient σ_{k1}, see Eqn. (670).

Sigma_w1	The coefficient $\sigma_{\omega 1}$, see Eqn. (670)
Beta2	The coefficient β_2, see Eqn. (671).
Sigma_k2	The coefficient σ_{k2}, see Eqn. (671).
Sigma_w2	The coefficient $\sigma_{\omega 2}$, see Eqn. (671)
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Normal Stress Term	This is an explicit term that directly incorporates divergence and turbulent kinetic energy, $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$, according to the full Boussinesq approximation. This setting is off by default, in which case the quantity $-2/3(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ is absorbed into the pressure, and causes the pressure to be slightly different. However, the results are the same whether this term is explicitly included. In incompressible flow only the gradients of pressure matter, so this setting has no effect on the results. In compressible flow, however, the absolute value of pressure is used in the Ideal Gas Law (Eqn. (99)). Activating this option may affect the solution convergence if k happens to be poorly initialized.

Activated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t \mathbf{S} - \frac{2}{3}(\mu_t \nabla \cdot \mathbf{v} + \rho k)\mathbf{I}$ and production is computed using Eqn. (630)

Deactivated	The stress tensor is modeled as $\mathbf{T}_t = 2\mu_t\mathbf{S}$ and production is modeled using the simplified expression $G_k = \mu_t S^2$.
TkeMinimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
Sdr Minimum	The minimum value that the transported variable ω is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

Working with the IDDES Coefficients

The IDDES Coefficients node represents the terms C_t in Eqn. (694), C_l in Eqn. (695), and C_{dt} in Eqn. (699) as constant coefficients. The default settings of these coefficients are 1.87, 5.0, and 20, respectively, and can be changed with a text entry in the corresponding expert properties. This object is activated as a subnode of the SST K-Omega Detached Eddy model node when the latter's Formulation Option property is set to `IDDES`.



IDDES Coefficients Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, we recommend that you do not change any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

Ct	The coefficient C_t, see Eqn. (694).
Cl	The coefficient C_l, see Eqn. (695).
Cdt	The coefficient C_{dt}, see Eqn. (699).

Menter SST K-Omega Detached Eddy Field Functions Reference

The following primitive field functions are made available to the simulation when the Menter SST K-Omega Detached Eddy model is used. These objects all have the same base properties.

DDES Correction Factor	The correction factor ϕ defined in Eqn. (682). This becomes available when you activate the Temporary Storage Retained property of the K-Omega Turbulence solver.
DES Upwind Blending Factor	The blending function for the hybrid upwind scheme.
Detached Eddy Length Scale	The cell length scale that is used for detached eddy simulation, expressed as Δ in Eqn. (682).
IDDES Length Scale	The cell length scale that is used for detached eddy simulation with the IDDES formulation option Δ_{IDDES}, see Eqn. (700).

What Are the Low Re Damping Parameters?

The Low Re Damping Parameters node provides the properties with which you can control the damping modification. It is activated as a subnode of any K-Omega model node when the latter's Low Re Damping Modification property is enabled.



This feature is disabled by default. It is not needed for low-y^+ wall treatment, only for "transitional flows". Although further development is required to make it totally reliable as a transition model, it can handle laminar bubble transition to some degree. It is *not* compatible with the Gamma ReTheta Transition model.

Low Re Damping Parameters Properties

RBeta	The coefficient Re_β. See Eqn. (651) for the Standard K-Omega model. See Eqn. (676) for the SST K-Omega and SST K-Omega Detached Eddy models.
Rk	The coefficient R_k. See Eqn. (652) for the Standard K-Omega model. See Eqn. (677) for the SST K-Omega and SST K-Omega Detached Eddy models.
Rw	The coefficient R_ω. See Eqn. (653) for the Standard K-Omega model. See Eqn. (678) for the SST K-Omega and SST K-Omega Detached Eddy models.

What Are the Compressibility Parameters?

The Compressibility Parameters node provides the properties with which you can control the compressibility correction. It is activated as a subnode of any K-Omega model node when the latter's Compressibility Correction property is enabled.



Compressibility Parameters Properties

Zeta_Star	The coefficient ξ^* . See Equations (656) and (657) for the Standard K-Omega model. See Equations (679) and (680) for the SST K-Omega and SST K-Omega Detached Eddy models.

What Is Curvature Correction?

Curvature correction is a factor that alters the turbulent energy production term according to local rotation and vorticity rates. It supplies the effects of strong (streamline) curvature and frame-rotation.

The Curvature Correction Parameters node provides the properties with which you can control the curvature correction factor. It is activated as a subnode of any K-Omega model node when the Curvature Correction Option property is enabled.

This feature is disabled by default. It is only necessary for cases that have strong streamline curvature or a rapidly rotating reference frame. In addition, we recommend you activate it only when you have a stable solution, to improve the results.



Curvature Correction Parameter Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this parameter, we recommend that you do not make changes to any properties. Tampering with them may diminish the effectiveness of the parameter.

Cr1	The coefficient C_{r1}. See Equation (642) for Curvature Correction. The default value is 0.04645.
Cr2	The coefficient C_{r2}. See Equation (642) for Curvature Correction. The default value is 0.25.
Cct	The coefficient C_{ct}. See Equation (645) for Curvature Correction. The default value is 2.
Upper limit	The maximum allowed value of the curvature correction factor, f_C, set by C_{max}. See Equation (642) for Curvature Correction. The default value is 1.25.

What Is the Realizability Coefficient?

It has been observed [94] that two-equation models predict an unexpectedly large growth of turbulent kinetic energy in stagnation point flows, which can have an adverse effect on the rest of the flow solution. To overcome this, a lower limit on the turbulence time scale has been proposed that imposes the "realizability" constraint on the eddy-viscosity formula. This time scale

realizability constraint is applied by default to the K-Omega turbulence models in STAR-CCM+. The Realizability Option property of the parent model object is set to `Durbin Scale Limiter`.



Properties

Realizability Coefficient	Value of the realizability time scale coefficient C_T as used, for example, in Eqn. (649). The default value is 0.6.

The effect of this coefficient can be demonstrated by observing the difference between two solutions of flow over an airfoil. In the first solution shown below, the Realizability option has not been activated. The build-up of turbulence at the stagnation point is clearly observed. In the second solution, the Realizability option has been activated.





How Do I Choose a K-Omega Model and Wall Treatment?

The Standard K-Omega model in its original form is sensitive to the levels of turbulence in the free stream. Wilcox has attempted to correct this and other limitations of the model by including several corrections in a recent revision. These corrections are incorporated into STAR-CCM+. However, they need to be used with caution until they are properly validated for complex flow situations.

Menter's SST K-Omega model uses the insensitivity to free-stream conditions of the K-Epsilon model in the far-field, while retaining the advantages of the K-Omega model near walls.

Both K-Omega turbulence models are available with all three wall treatments: high-y^+, low-y^+ and all-y^+. For most applications, the recommended choice of models is the SST K-Omega model combined with the all-y^+ wall treatment. It is recommended that only the default parameters be used with the SST K-Omega model.

Understanding the K-Omega Wall Treatment

In the K-Omega models, the wall treatment performs the following functions:

- Specifies the reference velocity u^* to be used in the wall laws
- Computes a special value of turbulent production G_k in the wall cell
- Computes a special value of specific dissipation ω in the wall cell

Three wall treatments are available as appropriate to the specific K-Omega model:

- High y+ Wall Treatment
- Low y+ Wall Treatment
- All y+ Wall Treatment

The equations that are used for each are detailed in the K-Omega model formulation, and a brief description of each follows.

What Is the High y+ Wall Treatment?

The high y^+ Wall Treatment is essentially the classic wall-function approach, where wall shear stress, turbulent production and turbulent dissipation are all derived from equilibrium turbulent boundary layer theory.

What Is the Low y+ Wall Treatment?

The low-y^+ Wall Treatment assumes that the viscous sublayer is resolved and thus wall laws are not needed.

What Is the All y+ Wall Treatment?

The all y^+ Wall Treatment is a hybrid approach that seeks to recover the behaviors of the other two wall treatments in the limit of very fine or very coarse meshes. It is a design goal that this wall treatment should give results similar to the low-y^+ treatment as $y^+ \to 0$ and to the high-y^+ treatment for $y^+ > 30$. It also gives reasonable results for intermediate meshes where the cell centroid falls in the buffer layer. This is the recommended wall treatment.

What Are Roughness Effects?

Roughness effects are effects such as wall shear stress (skin frition) and near-wall turbulence. They result from surface imperfections too small for the mesh to resolve, but which still affect the flow. Because of these effects, wall models need to be employed.

Following Wilcox [95], the K-Omega model modifies specific dissipation ω as a function of wall roughness in the cases of low Reynolds number and the all y^+ wall treatment.

Omega for Low Reynolds Numbers

In a low-Re treatment, the K-Omega model modifies specific dissipation ω in wall cells to be:

$$\omega = \frac{u_*^2}{\nu} S_r \tag{625}$$

with

$$S_r = \begin{cases} \left(\frac{50}{R^+}\right)^2 & \text{if } 5 \le R^+ \le 25 \\ \frac{100}{R^+} & \text{if } R^+ \ge 25 \end{cases} \tag{626}$$

where u_* is the friction velocity, ν is kinematic viscosity, and R^+ is the roughness parameter.

Omega for the All y+ Wall Treatment

In the all y^+ wall treatment, the K-Omega model modifies specific dissipation ω in wall cells to be:

$$\omega = g\frac{u_*^2}{\nu}S_r + (1-g)\frac{u_*}{\kappa y\sqrt{C_\mu}} \tag{627}$$

where g is the blending function described in Eqn. (701), with the additional requirement that $g = 0$ if $R^+ > R^+_{\text{rough}}$ where R^+_{rough} is the roughness parameter for a fully rough surface (i.e., a surface where roughness effects are so large, they leave no visible viscous flow behavior near the wall). In addition, u_* is the friction velocity, ν is kinematic viscosity, S_r is the function described above in Eqn. (626), κ is the von Karman constant, and C_μ (= 0.09) is a model coefficient.

Working with K-Omega Solvers

This section describes the various solvers that are used by the K-Omega turbulence models:

- Turbulence
- Turbulent viscosity

It also lists the residual monitors that are added to the simulation by these solvers.

What Is the K-Omega Turbulence Solver?

This solver controls the solution of the K-Omega model in all the continua for which the model is activated. For each transported variable—k and ω—the basic steps that are involved in the solution update are as follows:

1. Update boundary conditions.
2. Compute the reconstruction gradients and cell gradients.
3. Create the linear system of Eqn. (2591) using the discretization methods that are described in the K-Omega model formulation.
4. Compute the residual sum for monitoring convergence:

$$R = \sum_{\text{cells}} |r|$$

5. Solve the linear system.
6. Update the transported variable field.

Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.8. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.
Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.
	Deactivated	Update reconstruction gradients with each iteration.

Reconstruction Zeroed — Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

- Activated — Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.
- Deactivated — Do not reset reconstruction gradients. They are recomputed at the next iteration if they had been zeroed.

What Is the K-Omega Turbulent Viscosity Solver?

This solver controls the update of the turbulent viscosity.

Properties

Maximum Ratio	The maximum ratio of turbulent to laminar viscosity (μ_t/μ) permitted during the update process. This is used to prevent unphysically high values of turbulent viscosity from occurring on the path to convergence.
Under-Relaxation Factor	An under-relaxation factor ω_μ for the update of turbulent viscosity. The default value is 1.0. Let μ_t^n be the value of the previous iteration, and μ_t^{new} be the value computed during the current iteration. The default value is 0.8. The update is controlled as follows: $\mu_t^{n+1} = \omega_\mu \mu_t^{new} + (1-\omega_\mu)\mu_t^n$

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Using Residual Monitors

The K-Omega turbulence solvers add the following residual monitors to the top-level Monitors node of a simulation tree.

Sdr	Specific dissipation rate
Tke	Turbulent kinetic energy

The properties of all residual monitors are the same.

K-Omega Turbulence Formulation

In this section, the formulation of the following K-Omega model variants is presented:

- Standard
- SST
- SST detached eddy model

The wall treatments that are used and the boundary, region, and initial conditions are described, followed by the nomenclature and bibliography.

Standard K-Omega Model

Basic Transport Equations

The transport equations for the standard K-Omega model are:

$$\frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A (\mu + \sigma_k \mu_t)\nabla k \cdot d\mathbf{a} + \int_V \left(G_k - \rho\beta^* f_{\beta^*}(\omega k - \omega_0 k_0) + S_k\right) dV \tag{628}$$

$$\frac{d}{dt}\int_V \rho \omega dV + \int_A \rho \omega(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A (\mu + \sigma_\omega \mu_t)\nabla \omega \cdot d\mathbf{a} + \int_V (G_\omega - \rho\beta f_\beta(\omega^2 - \omega_0^2) + S_\omega) dV \tag{629}$$

where S_k and S_ω are the user-specified source terms. k_0 and ω_0 are the ambient turbulence values in source terms that counteract turbulence decay [100].

Turbulent Production

The production G_k is evaluated as:

$$G_k = \mu_t f_c S^2 - \frac{2}{3}\rho k \nabla \cdot \mathbf{v} - \frac{2}{3}\mu_t (\nabla \cdot \mathbf{v})^2 \tag{630}$$

where $\nabla \cdot \mathbf{v}$ is the velocity divergence and S is the modulus of the mean strain rate tensor:

$$S = |\mathbf{S}| = \sqrt{2\mathbf{S}:\mathbf{S}^T} = \sqrt{2\mathbf{S}:\mathbf{S}} \tag{631}$$

and:

$$\mathbf{S} = \frac{1}{2}(\nabla\mathbf{v} + \nabla\mathbf{v}^T) \tag{632}$$

The production of ω is evaluated as:

$$G_\omega = \rho\alpha\left[\left(\alpha^* S^2 - \frac{2}{3}(\nabla\cdot\mathbf{v})^2\right) - \frac{2}{3}\omega\nabla\cdot\mathbf{v}\right] \tag{633}$$

Vortex-Stretching Modification

The function f_β is a "vortex-stretching modification" designed to overcome the round-jet/plane-jet anomaly [95]. It is defined as follows:

$$f_\beta = \frac{1 + 70\chi_\omega}{1 + 80\chi_\omega} \tag{634}$$

where:

$$\chi_\omega = \frac{(\mathbf{W}\cdot\mathbf{W}):\mathbf{S}}{(\beta^*\omega)^3} \tag{635}$$

and $\mathbf{W}$ is the rotation rate tensor that is given by:

$$\mathbf{W} = \frac{1}{2}(\nabla\mathbf{v} - \nabla\mathbf{v}^T) \tag{636}$$

If the vortex-stretching modification is not required, f_β is set to unity.

Free-Shear Modification

The function f_{β^*} is designed to ameliorate the dependence of the model on free-stream boundary conditions. If the dissipation limiter is selected, f_{β^*} is defined as follows:

$$f_{\beta^*} = \begin{cases} 1 & \text{for } \chi_k \le 0 \\ \dfrac{1 + 680\chi_k^2}{1 + 400\chi_k^2} & \text{for } \chi_k > 0 \end{cases} \tag{637}$$

where:

$$\chi_k = \frac{\nabla k\cdot\nabla\omega}{\omega^3} \tag{638}$$

If the cross-diffusion limiter is selected, the following term is added to the specific dissipation rate equation (629):

$$CD_\omega = \sigma_d \frac{\rho}{\omega} \frac{\partial k}{\partial x_j} \frac{\partial \omega}{\partial x_j} \tag{639}$$

where:

$$\sigma_d = \begin{cases} 0, & \frac{\partial k}{\partial x_j} \frac{\partial \omega}{\partial x_j} \le 0 \\ \frac{1}{8}, & \frac{\partial k}{\partial x_j} \frac{\partial \omega}{\partial x_j} > 0 \end{cases} \tag{640}$$

Additionally, the function f_{β^*} is set to unity, and the Vortex-Stretching Modification function f_β is set as:

$$f_\beta = \frac{1 + 85\chi_\omega}{1 + 100\chi_\omega} \tag{641}$$

If the free-shear modification is not required, f_{β^*} is set to unity.

Curvature Correction

The transport equation for the turbulent kinetic energy is insensitive, by construction, to stabilizing and destabilizing effects usually associated with strong (streamline) curvature and frame-rotation. These effects can be incorporated by using a curvature correction factor, which alters the turbulent kinetic energy production term according to the local rotation and vorticity rates.

The curvature correction factor f_c [97] is

$$f_c = \min\left(C_{\max}, \frac{1}{C_{r_1}(|\eta| - \eta) + \sqrt{1 - \min(C_{r_2}, 0.99)}} \right) \tag{642}$$

where

$$\eta = \tau^2 (\mathbf{S} : \mathbf{S} - \mathbf{W} : \mathbf{W}) \tag{643}$$

The time-scale τ is limited in order to have the correct near-wall asymptotic behavior

$$\tau = \max(\tau_1, \tau_3); \qquad \tau_1 = \frac{1}{\beta^* \omega}; \qquad \tau_2 = 6\sqrt{\frac{\nu}{\beta^* k \omega}}; \qquad \tau_3 = (\tau_1{}^n \tau_2)^{\frac{1}{n}+1} \tag{644}$$

where n = 1.625.

The strain-rate tensor **S** is as defined in Eqn. (630), and the absolute rotation-rate tensor **W** is

$$\boldsymbol{W} = \boldsymbol{W}^l + \boldsymbol{W}^f + (C_{\mathrm{ct}} - 1)\boldsymbol{W}^S \tag{645}$$

where

$$\boldsymbol{W}^l = \frac{1}{2}(\nabla v - \nabla v^T), \qquad \boldsymbol{W}^f = \bar{\varepsilon} w^f, \qquad \boldsymbol{W}^S = \boldsymbol{W}^f - \bar{\varepsilon}(\boldsymbol{A}^{-1} \cdot w) \tag{646}$$

$\bar{\varepsilon}$ is the Levi-Civita tensor, and $\boldsymbol{W}^l$, $\boldsymbol{W}^f$, and $\boldsymbol{W}^S$ are the contributions due to the vorticity computed in the local frame of reference, the rotating frame of reference (defined by the rotation vector w), and the streamline curvature, respectively. For the last contribution, two additional terms are needed:

$$\boldsymbol{A} = \boldsymbol{I} - \frac{3\boldsymbol{S}^2}{\mathbf{2S:S}} \qquad \text{and} \qquad w = \bar{\varepsilon}\,\frac{\boldsymbol{S} \cdot (D_t \boldsymbol{S}) - (D_t \boldsymbol{S}) \cdot \boldsymbol{S}}{\mathbf{2S:S}} \tag{647}$$

where $D_t \boldsymbol{S}$ is the strain-rate total derivative tensor. For more details on the derivation of this method, see [97].

The constants used in the model are

$$C_{\mathrm{r1}} = 0.04645, \qquad C_{\mathrm{r2}} = 0.25, \qquad C_{\mathrm{ct}} = 2, \qquad C_{\max} = 1.25$$

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho k T \tag{648}$$

Turbulent Time Scale

The turbulent time scale is computed using Durbin's realizability constraint [98] as:

$$T = \begin{cases} \dfrac{\alpha^*}{\omega} & \text{without realizable scale option} \\ \min\left(\dfrac{\alpha^*}{\omega}, \dfrac{C_T}{\sqrt{3}S}\right) & \text{with realizable scale option} \end{cases} \tag{649}$$

where C_T is the realizable time scale coefficient with a default value of 0.6, and S is the modulus of the mean strain rate tensor defined in Eqn. (631).

Model Coefficients

Without the low-Reynolds number modifications, the coefficients defining the model have the following values:

$$\beta^* = 0.09 \quad , \quad \beta = \frac{9}{125} \quad , \quad \alpha = \frac{13}{25}$$
$$\alpha^* = 1 \quad , \quad \sigma_k = 0.5 \quad , \quad \sigma_\omega = 0.5 \tag{650}$$

Low-Reynolds Number Modifications

If the low-Reynolds number modifications are enabled, the coefficients β^*, α and α^* are replaced by the following:

$$\beta^*_{\text{low-re}} = \beta^* \cdot \frac{4/15 + (Re_T/Re_\beta)^4}{1 + (Re_T/Re_\beta)^4} \tag{651}$$

$$\alpha_{\text{low-re}} = \alpha \cdot \frac{\alpha_0 + Re_T/Re_\omega}{1 + Re_T/Re_\omega} \cdot \frac{1}{\alpha^*} \tag{652}$$

$$\alpha^*_{\text{low-re}} = \frac{\beta/3 + Re_T/Re_k}{1 + Re_T/Re_k} \tag{653}$$

where:

$$Re_\beta = 8 \quad , \quad Re_k = 6 \quad , \quad Re_\omega = 2.95 \quad , \quad \alpha_0 = \frac{1}{9} \tag{654}$$

and the turbulent Reynolds number is defined as:

$$Re_T = \frac{k}{\omega\nu} \tag{655}$$

Compressibility Correction

If the compressibility correction is enabled (to account for dilatation dissipation), the coefficient β^* is replaced by

$$\beta^*_{\text{comp}} = \beta^*[1 + \xi^* F(M_t)] \tag{656}$$

where β^* is given either by the constant coefficient of Eqn. (650) or if the low Reynolds number correction is used, the expression of Eqn. (651). Similarly, the coefficient β is replaced by:

$$\beta_{\text{comp}} = \beta - \beta^* \xi^* F(M_t) \tag{657}$$

The value of the coefficient is $\xi^* = 1.5$ and the compressibility function is defined as

$$F(M_t) = \max[0, M_t^2 - M_{t0}^2] \tag{658}$$

where $M_t^{\ 2} = \dfrac{2k}{a^2}$ and $M_{t0} = \dfrac{1}{4}$.

k is the turbulent kinetic energy and a is the speed of sound.

SST K-Omega Model

Basic Transport Equations

The transport equations for the SST K-Omega model are:

$$\frac{d}{dt}\int_V \rho k dV + \int_A \rho k(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A (\mu + \sigma_k \mu_t)\nabla k \cdot d\mathbf{a} + \int_V \left(\gamma_{eff} G_k - \gamma' \rho \beta^* f_{\beta^*}(\omega k - \omega_0 k_0) + S_k\right) dV \tag{659}$$

$$\frac{d}{dt}\int_V \rho \omega dV + \int_A \rho \omega(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A (\mu + \sigma_\omega \mu_t)\nabla \omega \cdot d\mathbf{a} + \int_V (G_\omega - \rho\beta f_\beta(\omega^2 - \omega_0^2) + D_\omega + S_\omega) dV \tag{660}$$

where S_k and S_ω are the user-specified source terms, k_0 and ω_0 are the ambient turbulence values in source terms that counteract turbulence decay, γ_{eff} is the effective intermittency provided by the Gamma ReTheta Transition model (it is unity if this model is not activated) and:

$$\gamma' = \min[\max(\gamma_{eff}, 0.1), 1] \tag{661}$$

Turbulent Production

The production G_k is evaluated as with the standard K-Omega model, Eqn. (630).

The production of ω is evaluated as:

$$G_\omega = \rho\gamma\left[\left(S^2 - \frac{2}{3}(\nabla \cdot \mathbf{v})^2\right) - \frac{2}{3}\omega\nabla \cdot \mathbf{v}\right] \tag{662}$$

where γ is a blended coefficient of the model and S is the modulus of the mean strain rate tensor given by Eqn. (632).

Cross-Derivative

The term D_ω is a cross-derivative term, defined as:

$$D_\omega = 2(1 - F_1)\rho\sigma_{\varpi 2}\frac{1}{\omega}\nabla k \cdot \nabla\omega \tag{663}$$

Relation for Turbulent Viscosity

The turbulent viscosity is computed as:

$$\mu_t = \rho k T \tag{664}$$

Turbulent Time Scale

The turbulent time scale is computed using Durbin's realizability constraint [98] as:

$$T = \begin{cases} \min\left(\dfrac{\alpha^*}{\omega}, \dfrac{a_1}{SF_2}\right) & \text{without realizable scale option} \\ \min\left(\dfrac{1}{\max(\omega/\alpha^*, (SF_2)/a_1)}, \dfrac{C_T}{\sqrt{3}S}\right) & \text{with realizable scale option} \end{cases} \tag{665}$$

and S is the modulus of the mean strain rate tensor that is defined in Eqn. (631). In Menter's original model [96], the modulus of the vorticity tensor was used in the definition of the turbulent viscosity for the SST model, rather than S. This slight modification extends the applicability of the model beyond aerodynamic applications. Furthermore, Durbin's

realizability constraint is used instead of Menter's proposal that turbulent production be limited to some multiple of the dissipation. In Eqn. (665), C_T is the realizable time scale coefficient and has a default value of 0.6.

The function F_2 is given by:

$$F_2 = \tanh(\arg_2{}^2) \tag{666}$$

where:

$$\arg_2 = \max\left(\frac{2\sqrt{k}}{\beta^* \omega y}, \frac{500\nu}{y^2\omega}\right) \tag{667}$$

The coefficient a_1 has the value:

$$a_1 = 0.31 \tag{668}$$

Model Coefficients

The coefficients in the model are calculated from the blending function F_1, such that each coefficient ϕ is given by:

$$\phi = F_1\phi_1 + (1 - F_1)\phi_2 \tag{669}$$

The coefficients of Set 1 (ϕ_1) are:

$$\begin{gathered} \beta_1 = 0.0750 \quad , \quad \sigma_{k1} = 0.85 \quad , \quad \sigma_{\omega 1} = 0.5 \\ \kappa = 0.41 \quad , \quad \gamma_1 = \frac{\beta_1}{\beta^*} - \sigma_{\omega 1}\frac{\kappa^2}{\sqrt{\beta^*}} \end{gathered} \tag{670}$$

The coefficients of Set 2 (ϕ_2) are:

$$\begin{gathered} \beta_2 = 0.0828 \quad , \quad \sigma_{k2} = 1.0 \quad , \quad \sigma_{\omega 2} = 0.856 \\ \kappa = 0.41 \quad , \quad \gamma_2 = \frac{\beta_2}{\beta^*} - \sigma_{\omega 2}\frac{\kappa^2}{\sqrt{\beta^*}} \end{gathered} \tag{671}$$

And in both Set 1 and Set 2:

$$\beta^* = 0.09 \quad , \quad \alpha^* = 1 \tag{672}$$

The blending function F_1 is defined as follows:

$$F_1 = \begin{cases} \tanh(\arg_1{}^4) & \text{without the } \gamma\text{-Re}_\theta \text{ model} \\ \max\left(\tanh(\arg_1{}^4), \exp\left[-\left(\frac{Re_T}{120}\right)^8\right]\right) & \text{with the } \gamma\text{-Re}_\theta \text{ model} \end{cases} \tag{673}$$

with:

$$\arg_1 = \min(\max\left(\frac{\sqrt{k}}{0.09\omega y}, \frac{500\nu}{y^2\omega}\right), \frac{2k}{y^2 CD_{k\omega}}) \tag{674}$$

where y is the distance to the nearest wall and $CD_{k\omega}$ is related to the cross-diffusion term. It is defined by:

$$CD_{k\omega} = \max\left(\frac{1}{\omega}\nabla k \cdot \nabla\omega,\ 10^{-20}\right) \tag{675}$$

Low-Reynolds Number Modifications

If the low-Reynolds number modifications are enabled, the coefficients β_1^*, γ_1 and α^* are replaced by the following:

$$\beta_{1,\text{ low-re}}^* = \beta_2^* \cdot \frac{4/15 + (Re_T/Re_\beta)^4}{1 + (Re_T/Re_\beta)^4} \tag{676}$$

$$\gamma_{1,\text{ low-re}} = \gamma_1 \cdot \frac{\alpha_0 + Re_T/Re_\omega}{1 + Re_T/Re_\omega} \cdot \frac{1}{\alpha^*} \tag{677}$$

$$\alpha_{1,\text{ low-re}}^* = \frac{\beta/3 + Re_T/Re_k}{1 + Re_T/Re_k} \tag{678}$$

where the constants are defined as in Eqn. (654) for the standard model.

Compressibility Correction

If the compressibility correction is enabled (to account for dilatation dissipation) the coefficient β^* is replaced by

$$\beta_{\text{comp}}^* = \beta^*[1 + \xi^* F(M_t)] \tag{679}$$

where β^* is given either by the constant coefficient of Eqn. (650) or if the low Reynolds number correction is used, the expression of Eqn. (651). Similarly, the coefficient β, computed from the blending equation of Eqn. (669) is replaced by:

$$\beta_{\text{comp}} = \beta - \beta^* \xi^* F(M_t) \tag{680}$$

The value of the coefficient is $\xi^* = 1.5$ and the compressibility function is defined as

$$F(M_t) = \max[M_t^2 - M_{t0}^2] \tag{681}$$

where $M_{t0}^2 = \frac{1}{4}$.

DES Version of the SST K-Omega Model

Based on the work by Menter and Kuntz [99], the dissipation term in the transport equation for turbulent kinetic energy (Eqn. (659)) has been modified as follows:

$$D_k = \rho\beta^* f_{\beta^*}\omega k\phi \tag{682}$$

$$\phi = \begin{cases} 1 \quad \text{if} \quad l_t < C_{\text{des}}\Delta \leftrightarrow \text{RANS mode} \\ >1 \quad \text{if} \quad l_t > C_{\text{des}}\Delta \leftrightarrow \text{LES mode} \end{cases}$$

where:

$$\phi = \max\left[\frac{l_t}{C_{\text{des}}\Delta}F, 1\right] \tag{683}$$

$$F = 1 - F_2 \tag{684}$$

$$l_t = \frac{\sqrt{k}}{\beta^*\omega} \tag{685}$$

Δ is the largest distance between the cell center under consideration and the cell centers of the neighboring cells, and F_2 is the blending function described in Eqn. (666).

The model constant C_{des}, which blends the values obtained from independent calibration of the K-Epsilon and K-Omega branches of the K-Omega SST model [101], is evaluated as follows:

$$C_{\text{des}} = C_{\text{des},k\text{-}\omega}F_1 + C_{\text{des},k\text{-}\varepsilon}(1-F_1) \tag{686}$$

where F_1 is defined in Eqn. (671).

IDDES Formulation

For the IDDES formulation, the length scale in the dissipation term in the transport equation for k is replaced with a hybrid length scale as follows:

$$D_k = \frac{\rho k^{3/2}}{l_{\mathrm{HYBRID}}} \tag{687}$$

where:

$$l_{\mathrm{HYBRID}} = \tilde{f}_d(1 + f_e)l_t + (1 - \tilde{f}_d)C_{\mathrm{des}}\Delta_{\mathrm{IDDES}} \tag{688}$$

Two more functions are introduced to the length scale calculation to add Wall-Modeled LES (WMLES) capability, a blending function f_B and the so-called "elevating" function f_e:

$$f_B = \min[2\exp(-9\alpha^2), 1.0] \tag{689}$$

$$\alpha = 0.25 - \frac{d}{\Delta} \tag{690}$$

$$f_e = \max[(f_{e1} - 1), 0]\psi f_{e2} \tag{691}$$

$$f_{e1} = \begin{cases} 2\exp(-11.09\alpha^2) & \text{if } \alpha \geq 0 \\ 2\exp(-9.0\alpha^2) & \text{if } \alpha < 0 \end{cases} \tag{692}$$

$$f_{e2} = 1.0 - \max(f_t, f_l) \tag{693}$$

$$f_t = \tanh[(C_t^2 r_{dt})^3] \tag{694}$$

$$f_l = \tanh[(C_l^2 r_{dl})^{10}] \tag{695}$$

$$r_{dt} = \frac{\nu_t}{\sqrt{\nabla\mathbf{v}:\nabla\mathbf{v}^T}\kappa^2 d^2} \tag{696}$$

$$r_{dl} = \frac{\nu}{\sqrt{\nabla\mathbf{v}:\nabla\mathbf{v}^T}\kappa^2 d^2} \tag{697}$$

The introduction of the low-Reynolds number correction function ψ in the formulation of f_e is unrelated to the low-Reynolds number correction role of this function in the LES mode of DDES, and is purely empirical.

C_t and C_l are model constants.

The WMLES and DDES branches of the model are combined using a modified version of the DDES f_d function as follows:

$$\tilde{f}_d = \max((1 - f_{dt}), f_B) \tag{698}$$

$$f_{dt} = 1 - \tanh[(C_{dt} r_{dt})^3] \tag{699}$$

C_{dt} is a model constant.

The IDDES model also uses an altered version of the mesh length scale Δ_{IDDES}, computed as:

$$\Delta_{\mathrm{IDDES}} = \min(\max(0.15*d, 0.15*\Delta, \Delta_{\min}), \Delta) \tag{700}$$

where $\Delta_{\min}$ is the smallest distance between the cell center under consideration and the cell centers of the neighboring cells.

Wall Treatment

The following equations present the wall treatment formulations. In general, the high-y^+ formulation is equivalent to the traditional wall-function approach, and the low-y^+ formulation is equivalent to the traditional low-Reynolds number approach where no modeling beyond the assumption of laminar flow is necessary in the wall cells.

For the all-y^+ formulation, a blending function g is defined in terms of the wall-distance Reynolds number:

$$g = \exp\left(-\frac{Re_y}{11}\right) \tag{701}$$

and $Re_y = \sqrt{k}y/\nu$, a wall-distance-based Reynolds number.

Reference Velocity

$$u^* = \begin{cases} \sqrt{\beta^{*1/2} k} & ; \quad \text{high-}y^+ \\ \sqrt{g\nu u/y + (1-g)\beta^{*1/2} k} & ; \quad \text{low-}y^+ \\ \sqrt{g\nu u/y + (1-g)\beta^{*1/2} k} & ; \quad \text{all-}y^+ \end{cases} \tag{702}$$

The reference velocity u^* is not needed in the calculation for low-y^+ simulations. However, it is still needed when computing the wall y^+ for post-processing. Since the assumption is that $u^* \equiv \sqrt{\tau_w/\rho} \approx \sqrt{\nu u/y}$ gives an inaccurate result for coarse meshes, the blended expression is used for the low-y^+ wall treatment as well.

Wall-Cell Production

$$G_k = \begin{cases} \frac{1}{\mu}\left(\rho u^* \frac{u}{u^+}\right)^2 \frac{\partial u^+}{\partial y^+} & ; \quad \text{high-}y^+ \\ \mu_t S^2 & ; \quad \text{low-}y^+ \\ g\mu_t S^2 + (1-g)\frac{1}{\mu}\left(\rho u^* \frac{u}{u^+}\right)^2 \frac{\partial u^+}{\partial y^+} & ; \quad \text{all-}y^+ \end{cases} \tag{703}$$

Wall-Cell Specific Dissipation

$$\omega = \begin{cases} \frac{u^*}{\sqrt{\beta^*}\kappa y} & ; \quad \text{high-}y^+ \\ \frac{6\nu}{\beta y^2} & ; \quad \text{low-}y^+ \\ g\frac{6\nu}{\beta y^2} + (1+g)\frac{u^*}{\sqrt{\beta^*}\kappa y} & ; \quad \text{all-}y^+ \end{cases} \tag{704}$$

Boundary, Region, and Initial Conditions

Wall Boundaries

At walls, a Neumann boundary condition is used for the turbulent kinetic energy k, that is, $\partial k/\partial n|_w = 0$. The specific dissipation rate ω is specified in the wall cells according to the appropriate method in the wall treatment.

Flow, Region, and Initial Conditions

When defining values for flow boundary, region and initial conditions, you have several choices for specifying the turbulent kinetic energy and specific dissipation rate:

- Enter the values directly.
- Have them derived from a specified turbulence intensity (I) and length scale (L) using the relations:

$$\begin{aligned} k &\approx \frac{3}{2}(I\mathrm{v})^2 \\ \omega &\approx \frac{\sqrt{k}}{L\beta^{*\,1/4}} \end{aligned} \tag{705}$$

where v is the turbulent velocity scale and β^* is a coefficient of the model.

- Have them derived from a specified turbulence intensity (I) and turbulent viscosity ratio μ_t/μ using the relations:

$$k \approx \frac{3}{2}(I\mathrm{v})^2$$
$$\omega \approx \frac{\rho k}{(\mu_t/\mu)\ \mu} \tag{706}$$

where v is the turbulent velocity scale.

K-Omega Turbulence Nomenclature

The following is a guide to the terms used in this section.

Scalars

d	distance to the nearest wall
k	turbulent kinetic energy
l	turbulent length scale
t	time
u^*	reference velocity that is used in wall functions
u^+	wall-parallel velocity non-dimensionalized with u^*
G_k	turbulent production
G_ω	production of specific dissipation rate
I	turbulence intensity
L	turbulent length scale
Re_T	turbulent Reynolds number
S	strain rate parameter
V	cell volume
v	velocity magnitude
ρ	density
μ	dynamic viscosity
μ_t	turbulent viscosity
ν	kinematic viscosity
σ_k, σ_ω	inverse turbulent Schmidt numbers

Tensors or Column Vectors

$\mathbf{I}$	identity matrix
S	strain rate tensor
$\mathbf{T}$	viscous stress tensor
$\mathbf{W}$	vorticity tensor

Cartesian Vectors

$\mathbf{a}$	face area vector
$\mathbf{v}$	velocity

$\mathbf{v}_g$ grid velocity

Subscripts

t turbulent

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Divergence of vector quantity

$$\nabla\cdot\mathbf{v} = \frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S}:\mathbf{T} = S_{ij}T_{ji}$$

K-Omega Turbulence Bibliography

[95] Wilcox, D.C. 1998. “Turbulence Modeling for CFD”, 2nd edition, DCW Industries, Inc.

[96] Menter, F.R. 1994. “Two-equation eddy-viscosity turbulence modeling for engineering applications”, *AIAA Journal*, **32**(8), pp. 1598-1605.

[97] Arolla, S. K., and Durbin, P. A. 2012. “Modeling rotation and curvature effects within scalar eddy viscosity model framework”, publication pending.

[98] Durbin, P.A. 1996. “On the k-e stagnation point anomaly”, *Int. J. Heat and Fluid Flow* **17,** pp. 89-90.

[99] Menter, F. R. and Kuntz, M. 2002. “Adaptation of Eddy Viscosity Turbulence Models to Unsteady Separated Flows Behind Vehicles”, in *The Aerodynamics of Heavy Vehicles: Trucks, Buses and Trains*, Springer, Asilomar, CA.

[100] Spalart, P.R. and Rumsey, C.L. 2007. “Effective Inflow Conditions for Turbulence Models in Aerodynamic Calculations”, *AIAA Journal*, **45**(10), pp. 2544-2553.

[101] Travin, A., Shur, M. L., Strelets, M. Kh., and Spalart, P. R. 2002. “Physical and numerical upgrades in the Detached-Eddy Simulation of complex turbulent flows”, in *Advances in LES of Complex Flows*, R. Friedrich and W. Rodi Eds., Kluwer Academic Press, pp. 239-254.

[102] Wilcox, D.C. 2008. “Formulation of the k-ω turbulence model revisited”, *AIAA Journal*, **46**(11), pp. 2823-2838.

Boundary and Region Types Reference

This section defines the boundary conditions that are required for each boundary type supported by the K-Omega model:

- Free stream
- Mass flow inlet
- Pressure outlet
- Stagnation inlet
- Velocity inlet

In addition, information is provided about the region conditions and values that need to be considered for each region type:

- Porous

Region and boundary types that do *not* require any conditions or values to be set are *not* listed.

Free Stream

The free stream boundary type requires the following boundary conditions when the K-Omega turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Specific dissipation rate.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulence intensity.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulent length scale.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Mass Flow Inlet

The mass flow inlet boundary type requires the following boundary conditions when the K-Omega turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Specific dissipation rate.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulence intensity.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulent length scale.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Pressure Outlet

The pressure outlet boundary type requires the following boundary conditions when the K-Omega turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Specific dissipation rate.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulence intensity.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.

- Turbulent length scale.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Stagnation Inlet

The stagnation inlet boundary type requires the following boundary conditions when the K-Omega turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Specific dissipation rate.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulence intensity.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulent length scale.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Velocity Inlet

The velocity inlet boundary type requires the following boundary conditions when the K-Omega turbulence model is active:

- Turbulence specification

The following boundary values are also required:

- Specific dissipation rate.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.

- Turbulence intensity.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulent length scale.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Porous

The porous region type requires the following conditions to be set:

- Turbulence specification

The following region values are also required:

- Specific dissipation rate.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulence intensity.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale` or `Intensity + Viscosity Ratio`.
- Turbulent kinetic energy.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `K + Omega`.
- Turbulent length scale.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Length Scale`.
- Turbulent viscosity ratio.
  This is only available if the Method property of the Conditions Turbulence Specification is set to `Intensity + Viscosity Ratio`.

Shared Conditions and Values Reference

This section details the boundary and region conditions, as well as initial conditions, that may need to be set for models in the K-Omega turbulence model family:

- Turbulence specification

Information is also provided on boundary, region, and initial condition values that are common to the K-Omega family of turbulence models:

- Specific dissipation rate

Turbulence Specification

Specifies the method for defining the turbulence profile for a boundary condition, region condition, or initial condition.

Properties

Method	The method for defining turbulence.	
	K + Omega	Direct specification of the values of k and ω. A Turbulent Kinetic Energy node and a Specific Dissipation Rate node are added to the Values node.
	Intensity + Length Scale	Calculation of k and ε from a specified turbulence intensity I and length scale L using Eqn. (705). Do not use this method if the velocity field is initialized to zero. A Turbulence Intensity node and a Turbulent Length Scale node are added to the Values node. Additionally, a Turbulent Velocity Scale node appears in the Initial Conditions node.
	Intensity + Viscosity Ratio	Calculation of k and ε from the turbulence intensity (I) and the ratio of turbulent to laminar viscosity (μ_t/μ) using Eqn. (706). Do not use this method if the velocity field is initialized to zero. A Turbulent Viscosity Ratio node and a Turbulence Intensity node are added to the Values node. Additionally, a Turbulent Velocity Scale node appears in the Initial Conditions node.

Specific Dissipation Rate

This node is added to the Values node by selecting K + Omega as the Method in the Conditions Turbulence Specification node. It is a direct specification of ω.

Properties

This value is input as a scalar profile.

Boundary Conditions and Values Reference

The following boundary conditions may need to be set for the K-Omega turbulence models, depending on the selected boundary type:

- Turbulence specification

Boundary conditions are elected by the boundary type, based on the models that have been activated in the continuum. The boundary conditions required for each boundary type when the K-Omega turbulence model is activated are listed here.

In addition, the following boundary values may need to be set for these models, again depending on the selected boundary type:

- Specific dissipation rate
- Turbulence intensity
- Turbulent kinetic energy
- Turbulent length scale
- Turbulent viscosity ratio

Boundary values are elected by the boundary type, and/or by methods selected in corresponding nodes in the boundary conditions node.

Region Conditions and Values Reference

The following region conditions may need to be set for the K-Omega turbulence model for porous regions:

- Turbulence specification

Region conditions are elected by the region type, based on the models that have been activated in the continuum. The region conditions that are required for each region type when the K-Omega turbulence model is activated are listed here.

The following region values may need to be set for this model for porous regions:

- Specific dissipation rate
- Turbulence intensity

- Turbulent kinetic energy
- Turbulent length scale
- Turbulent viscosity ratio

Information is also provided about a region value available for the K-Omega turbulence model in STAR-CCM+:

- Specific Dissipation Rate Source

Specific Dissipation Rate Source

This node is added when the SST (Menter) K-Omega Detached Eddy model is used, and the Turbulence Source Option region condition is ticked. It is entered as a scalar profile.

K-Omega Turbulence Field Functions Reference

The following primitive field functions are made available to the simulation when the K-Omega turbulence model is used. These objects all have the same base properties.

Turbulent Kinetic Energy	The transported variable k.
Specific Dissipation Rate	The transported variable ω.
Turbulent Dissipation Rate	(Only available when temporary storage is retained). The dissipation rate in the turbulent kinetic energy transport equation. This is defined as $\beta^* f_{\beta^*} \omega k$ as in Eqns (628) and (659) for the Standard and SST K-Omega models, and as $\beta^* f_{\beta^*} \omega k \phi$ as in Eqn. (682) for the DES version of the SST K-Omega model.

Using Reynolds Stress Transport Turbulence

This section explains the following:

- What are the Reynolds Stress Transport models?
- Understanding wall treatment
- Working with Reynolds Stress Transport solvers
- Formulation of the model equations

Also included is information on the following topics as they relate to the Reynolds Stress Transport model:

- The Reynolds Stress C3e coefficient
- Reynolds Stress conditions and values
- Boundary conditions and values
- Region conditions and values
- Field functions

What Are the Reynolds Stress Transport Models?

Reynolds stress transport (RST) models, also known as second-moment closure models, solve transport equations for all components of the specific Reynolds stress tensor, $\mathbf{R} = -\mathbf{T}_t/\rho \equiv \overline{\mathbf{v}'\mathbf{v}'}$. These models naturally account for effects such as anisotropy due to strong swirling motion, streamline curvature, rapid changes in strain rate and secondary flows in ducts.

The RST model carries significant computational overhead. Seven equations must be solved in three dimensions as opposed to the two equations of a K-Epsilon model. Apart from the additional memory and computational time required for these equations to be solved, there is also likely to be a penalty in the total number of iterations required to obtain a converged solution due to the numerical stiffness of the RST equations.

The reasons that the RST model requires the solution of seven equations in three dimensions are as follows. The Reynolds stress tensor is symmetric, so that only six of the nine components are unique. In addition to the six RST equations, a model equation is also needed for the isotropic turbulent dissipation ε. This is the same equation as the one used in the Standard K-Epsilon model. In two dimensions, only four of the Reynolds stress components need to be evaluated, so the total number of equations is reduced by two.

The starting point for the development of RST models is generally the exact differential transport equation for the Reynolds stresses, which is derived by multiplying the instantaneous Navier-Stokes equations by a fluctuating property and Reynolds-averaging the product (see [109] for this exact equation). In the resulting equations, only the transient, convective and molecular diffusion terms require no modeling. The terms remaining to be modeled are the diffusion term, the dissipation term and, perhaps the greatest challenge, the pressure-strain term. Appropriate models for these terms have received much attention during the past few decades; some of them are reviewed in [110].

The Quadratic Pressure Strain model is more modern and most likely the model of choice. However, the Linear Pressure Strain model is also available with a two-layer formulation, so if it is important to resolve the viscous sublayer, the linear pressure-strain model should be chosen.

STAR-CCM+ has a choice of three different Reynolds stress transport turbulence models:

- Linear Pressure Strain
- Quadratic Pressure Strain
- Linear Pressure Strain Two-Layer

What Is the Linear Pressure Strain Model?

The linear pressure-strain model of Gibson and Launder [104] uses a classic approach to modeling the pressure-strain term, splitting it up into a slow (return-to-isotropy) term, a rapid term, and a wall-reflection term. The node of this model has its own properties.

This is implemented with both a high-y^+ wall treatment (that is, using wall functions) and a two-layer approach.

Linear Pressure Strain Properties

Convection

For guidance on selecting a convection, see Diffusion Term, p. 6349.

1st-order Selects the first-order convection.

2nd-order Selects the second-order convection.

Linear Pressure Strain Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques used in STAR-CCM+, we recommend that you not make any changes within the *Expert* category. The

values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

C1	The coefficient C_1, see Eqn. (720).
C1w	The coefficient C_{1w}, see Eqn. (723).
C2	The coefficient C_2, see Eqn. (722).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (717).
C2w	The coefficient C_{2w}, see Eqn. (726).
Cl	The coefficient C_l, see Eqn. (726).
Cmu	The coefficient C_μ, see Eqn. (711)
FwMax	The coefficient $f_w^{\max}$, see Eqn. (724).
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (716) is calculated.
	None: Neglects the term $C_{\varepsilon 3}\mathrm{tr}(\mathbf{G})$.
	Boundary Layer Orientation: Computes $C_{\varepsilon 3}$ according to Eqn. (532).
	Thermal Stratification: Computes $C_{\varepsilon 3}$ according to Eqn. (533).
	Constant Coefficient: Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the node's properties.
Sarkar	The coefficient C_M, see Eqn. (717).

Secondary Gradients	Neglects or includes the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces. On — Includes both secondary gradients. Off — Excludes both secondary gradients. Interior Only — Includes the interior secondary gradients only. Boundaries Only — Includes the boundary secondary gradients only.
Sigma_e	The coefficient σ_ε, see Eqn. (717).
Sigma_k	The coefficient σ_k, see Eqn. (724).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
TkeMinimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer
Use Boussinesq Approximation	Instead of using the divergence of the computed Reynolds-stress tensor, uses the Boussinesq approximation given by Eqn. (592). Activated — Uses Boussinesq approximation. Deactivated — Uses computed Reynolds stresses.
Cs	The coefficient C_s, see Eqn. (712).
Use Daly-Harlow	Instead of using the isotropic formulation, as in Eqn. (709), uses the Generalized Gradient Diffusion model given by Eqn. (712).

Activated	Uses Daly & Harlow model (see Eqn. (712)).
Deactivated	Uses the isotropic model (see Eqn. (709)).

Wall Reflection Term	Includes or omits the wall-reflection terms $\underline{\Phi}_{1w}$ and $\underline{\Phi}_{2w}$ (see Eqn. (719)).
Activated	The terms are included.
Deactivated	The terms are omitted.

What Is the Quadratic Pressure Strain Model?

The quadratic pressure-strain model of Sarkar, Speziale, and Gatski [110] is a more modern formulation. This model uses a higher-order truncation of the expansion for the pressure-strain term. The node of this model has its own properties.

This model can only be used with a high-y^+ wall treatment (that is, using wall functions) without resolving the viscous-affected region.

Quadratic Pressure Strain Properties

Convection	For guidance on selecting a convection, see Diffusion Term, p. 6349.
1st-order	Selects the first-order convection.
2nd-order	Selects the second-order convection.

Quadratic Pressure Strain Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques used in STAR-CCM+, we recommend that you not make any changes within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

Cs1	The coefficient C_{s1}, see Eqn. (736).
Cr4	The coefficient C_{r4}, see Eqn. (736).

Cs2	The coefficient C_{s2}, see Eqn. (736).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (717).
Cr3	The coefficient C_{r3}, see Eqn. (736).
Cr3star	The coefficient C_{r3}^{*}, see Eqn. (736).
Cr4	The coefficient C_{r4}, see Eqn. (736).
Cr1	The coefficient C_{r1}, see Eqn. (736).
Cr2	The coefficient C_{r2}, see Eqn. (736).
Cmu	The coefficient C_{μ}, see Eqn. (711)
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (716) is calculated.
	None: Neglects the term $C_{\varepsilon 3}\mathrm{tr}(\mathbf{G})$.
	Boundary Layer Orientation: Computes $C_{\varepsilon 3}$ according to Eqn. (532).
	Thermal Stratification: Computes $C_{\varepsilon 3}$ according to Eqn. (533).
	Constant Coefficient: Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the node's properties.
Sarkar	The coefficient C_M, see Eqn. (717).
Secondary Gradients	Neglects or includes the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.
	On: Includes both secondary gradients.

Off	Excludes both secondary gradients.
Interior Only	Includes the interior secondary gradients only.
Boundaries Only	Includes the boundary secondary gradients only.

Sigma_e

The coefficient σ_ε, see Eqn. (717).

Sigma_k

The coefficient σ_k, see Eqn. (724).

Tdr Minimum

The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

TkeMinimum

The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer

Use Boussinesq Approximation

Instead of using the divergence of the computed Reynolds-stress tensor, uses the Boussinesq approximation given by Eqn. (592).

Activated	Uses Boussinesq approximation.
Deactivated	Uses computed Reynolds stresses.

Cs

The coefficient C_s, see Eqn. (712).

Use Daly-Harlow

Instead of using the isotropic formulation, as in Eqn. (709), uses the Generalized Gradient Diffusion model given by Eqn. (712).

Activated	Uses Daly & Harlow model (see Eqn. (712)).
Deactivated	Uses the isotropic model (see Eqn. (709)).

What Is the Linear Pressure Strain Two-Layer Model?

The advantage of the linear pressure-strain model is that it lends itself to being incorporated into a two-layer formulation which can be used to resolve the viscous sublayer for low-Reynolds number type applications. The node of this model has its own properties.

In this model's approach, suggested by Rodi [86], the computation is divided into two layers. In the layer adjacent to the wall, the turbulent dissipation rate ε and the turbulent viscosity μ_t are specified as functions of wall distance. The values of ε specified in the near-wall layer are blended smoothly with the values computed from solving the transport equation far from the wall.

Linear Pressure Strain Two-Layer Properties

Two-layer type	Selects a two-layer formulation	
	Shear Driven (Wolfstein)	The two-layer formulation of Wolfstein [90]. This is appropriate for flows that are not dominated by buoyancy forces. See Eqns. (606) to (609) in the K-Epsilon model formulation.
	Buoyancy Driven (Xu)	The two-layer formulation of Xu *et al*. [91]. This is appropriate for flows that are dominated by buoyancy forces. See Eqns. (610) to (619) in K-Epsilon model formulation.
Convection	For guidance on selecting a convection, see Diffusion Term, p. 6349.	
	1st-order	The first-order convection.
	2nd-order	The second-order convection.

Linear Pressure Strain Two-Layer Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques used in STAR-CCM+, we recommend that you not make any changes within the *Expert* category. The

values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

C1	The coefficient C_1, see Eqn. (720).
C1w	The coefficient C_{1w}, see Eqn. (723).
C2	The coefficient C_2, see Eqn. (722).
C2e	The coefficient $C_{\varepsilon 2}$, see Eqn. (717).
C2w	The coefficient C_{2w}, see Eqn. (726).
Cl	The coefficient C_l, see Eqn. (726).
Cmu	The coefficient C_μ, see Eqn. (711)
Two-Layer Delta ReY	The value of $\Delta \mathrm{Re}_y$ used in Eqn. (602) in K-Epsilon model formulation.
FwMax	The coefficient $f_w^{\max}$, see Eqn. (724).
Buoyancy Production of Dissipation	Determines how the coefficient $C_{\varepsilon 3}$ in Eqn. (716) is calculated.

None	Neglects the term $C_{\varepsilon 3}\mathrm{tr}(\mathbf{G})$.
Boundary Layer Orientation	Computes $C_{\varepsilon 3}$ according to Eqn. (532).
Thermal Stratification	Computes $C_{\varepsilon 3}$ according to Eqn. (533).
Constant Coefficient	Computes $C_{\varepsilon 3}$ as a constant coefficient. This setting activates the C3e node. The default value is `1.0`, and can be changed in the node's properties.

Two-Layer ReY*	The value of $\mathrm{Re}_y{}^*$ used in Eqn. (601) in K-Epsilon model formulation.
Sarkar	The coefficient C_M, see Eqn. (717).
Secondary Gradients	Neglects or includes the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces. On — Includes both secondary gradients. Off — Excludes both secondary gradients. Interior Only — Includes the interior secondary gradients only. Boundaries Only — Includes the boundary secondary gradients only.
Sigma_e	The coefficient σ_ε, see Eqn. (717).
Sigma_k	The coefficient σ_k, see Eqn. (724).
Tdr Minimum	The minimum value that the transported variable ε is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
TkeMinimum	The minimum value that the transported variable k is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer
Use Boussinesq Approximation	Instead of using the divergence of the computed Reynolds-stress tensor, uses the Boussinesq approximation given by Eqn. (592). Activated — Uses Boussinesq approximation. Deactivated — Uses computed Reynolds stresses.

Cs	The coefficient C_s, see Eqn. (712).
Use Daly-Harlow	Instead of using the isotropic formulation, as in Eqn. (709), uses the Generalized Gradient Diffusion model given by Eqn. (712). Activated: Uses Daly & Harlow model (see Eqn. (712)). Deactivated: Uses the isotropic model (see Eqn. (709)).
Wall Reflection Term	Includes or omits the wall-reflection terms $\underline{\Phi}_{1w}$ and $\underline{\Phi}_{2w}$ (see Eqn. (719)). Activated: The terms are included. Deactivated: The terms are omitted.

What Is the Reynolds Stress C3e Coefficient?

The C3e node represents the term $C_{\varepsilon 3}$ in Eqn. (716) as a constant coefficient. The default setting of this object is `1.0`, and can be changed with a text entry in its C3e expert property. It is activated as a subnode of any RST model node when the latter's Buoyancy Production of Dissipation expert property is set to `Constant Coefficient`.



C3e Coefficient Expert Properties

C3e	The coefficient $C_{\varepsilon 3}$, see Eqn. (716).

Understanding the Reynolds Stress Wall Treatment

The wall treatment performs the following functions for the Reynolds Stress Transport models:

- Specifies the reference velocity u^* to be used in the wall laws
- Computes a special value of turbulent production G_k in the wall cell
- Computes a special value of turbulent dissipation ε in the wall cell
- Sets the specific Reynolds stress tensor $\mathbf{R}$ in the wall cell

Two wall treatments are available as appropriate to the specific Reynolds Stress model:

- High y+ Wall Treatment
- Two-Layer All y+ Wall Treatment

The equations used for each type of Reynolds Stress wall treatment are detailed in the Reynolds Stress Transport model formulation under the wall treatment section.

What Is the High y+ Wall Treatment?

The High y+ Wall Treatment is essentially the classic wall-function approach, where wall shear stress, turbulent production and turbulent dissipation are all derived from equilibrium turbulent boundary layer theory. Apart from the specification of the Reynolds stresses in the wall cell (according to Eqn. (747)), the formulation is identical to the high-y^+ wall treatment used for the K-Epsilon model.

What Is the Two-Layer All y+ Wall Treatment?

The Two-Layer All y+ Wall Treatment is a hybrid approach that seeks to recover the behaviors of the other two wall treatments in the limit of very fine or very coarse meshes. It contains a wall boundary condition for ε that is consistent with the two-layer formulation. Apart from the specification of the Reynolds stresses in the wall cell (according to Eqn. (747)), the formulation is identical to the two-layer all-y^+ wall treatment used for the K-Epsilon model. It is a design goal that this wall treatment should give

results similar to the low-y^+ treatment as $y^+ \to 0$ and to the high-y^+ treatment for $y^+ > 30$. It will also give reasonable results for intermediate meshes where the cell centroid falls in the buffer layer.

Reynolds Stress Transport Solvers

This section describes the solvers used by the Reynolds Stress turbulence models:

- Turbulence
- Turbulent viscosity

It also lists the residual monitors that are added to the simulation by these solvers.

What Is the Reynolds Stress Turbulence Solver?

This solver controls the solution of the Reynolds Stress transport model in all the continua for which the model is activated. The node of this solver has its own properties.

For each transported variable (components of $\mathbf{R}$ and the dissipation rate ε) the basic steps involved in the solution update are as follows:

1. Update boundary conditions.
2. Compute the reconstruction gradients and cell gradients.
3. Create the linear system of Eqn. (2591) using the discretization methods described in the Reynolds stress transport model formulation.
4. Compute the residual sum:

$$R = \sum_{\text{cells}} |r|$$

   for the purpose of monitoring convergence.
5. Solve the linear system.
6. Update the transported variable field.

Reynolds Stress Turbulence Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.6. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.

Reynolds Stress Turbulence Solver Expert Properties

The properties in the *Expert* category are designed to be used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver will not update the solution.
	Deactivated	The solver will update the solution.
Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients and the cell gradients. These quantities will become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Reconstruction Frozen	Sets whether or not reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.

Deactivated	Update reconstruction gradients with each iteration.

Reconstruction Zeroed	Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) will be first-order estimates.
Activated	Set reconstruction gradients to zero at the next iteration. Gradients will remain zeroed until you clear the checkbox again.
Deactivated	Do not reset reconstruction gradients. The will be recomputed at the next iteration if they had been zeroed.

What Is the Reynolds Stress Turbulent Viscosity Solver?

This solver controls the update of the turbulent viscosity. The node of this solver has its own properties.

Reynolds Stress Turbulent Viscosity Solver Properties

Maximum Ratio	The maximum ratio of turbulent to laminar viscosity (μ_t/μ) permitted during the update process. This is used to prevent unphysically high values of turbulent viscosity from occurring on the path to convergence.
Under-Relaxation Factor	An under-relaxation factor ω_μ for the update of turbulent viscosity. The default value is 1.0. Let μ_t^n be the value of the previous iteration, and μ_t^{new} be the value computed during the current iteration. The update is controlled as follows: $\mu_t^{n+1} = \omega_\mu \mu_t^{new} + (1-\omega_\mu)\mu_t^n$

Reynolds Stress Turbulent Viscosity Solver Expert Properties

The properties in the *Expert* category are designed to be used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these

properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver will not update the solution.
	Deactivated	The solver will update the solution.

Using Residual Monitors

The Reynolds Stress turbulence solvers add the following residual monitors to the top level Monitors node of a simulation tree.

Tdr	Turbulent dissipation rate.
uu-stress	UU Reynolds stress component
uv-stress	UV Reynolds stress component
uw-stress	UW Reynolds stress component
vv-stress	VV Reynolds stress component
vw-stress	VW Reynolds stress component
ww-stress	WW Reynolds stress component

The properties of all residual monitors are the same.

Reynolds Stress Transport Turbulence Formulation

The formulation of the Reynolds Stress Transport models includes the basic transport equation and the individual terms to be modeled:

- Reynolds stress diffusion
- Turbulent production

- Buoyancy production
- Turbulent dissipation rate
- Dilitation dissipation rate
- Linear pressure-strain model
- Quadratic pressure-strain model
- Wall treatment formulations
- Formulations of boundary, region and initial conditions
- Nomenclature
- Bibliography

Reynolds Stress Transport Equation

The transport equation for the specific Reynolds stress tensor, $\mathbf{R} \equiv \overline{\mathbf{v}'\mathbf{v}'}$, is:

$$\frac{d}{dt}\int_V \rho \mathbf{R} dV + \int_A \rho \mathbf{R}(\mathbf{v}-\mathbf{v}_g) \cdot d\mathbf{a} = \int_A \mathbf{D} \cdot d\mathbf{a} + \int_V \left[\mathbf{P} + \mathbf{G} - \frac{2}{3}\rho \mathbf{I} \Upsilon_M + \underline{\phi} - \underline{\varepsilon} + \mathbf{S_R}\right] dV \tag{707}$$

where the terms on the right-hand side are diffusion $\mathbf{D}$, turbulent production $\mathbf{P}$, buoyancy production $\mathbf{G}$, turbulent dissipation $\underline{\varepsilon}$, dilatation dissipation Υ_M, pressure strain $\underline{\phi}$, and the user-specified source $\mathbf{S_R}$.

For the Linear Pressure Strain and Quadratic Pressure Strain models, the dissipation is simply:

$$\underline{\varepsilon} = \frac{2}{3}\rho \mathbf{I} \tag{708}$$

Reynolds Stress Diffusion

Two different models for the Reynolds-stress diffusion are available. By default, following [108], a simple isotropic form of the turbulent diffusion is adopted, such that:

$$\mathbf{D} = \left(\mu + \frac{\mu_t}{\sigma_k}\right)\nabla \mathbf{R} \tag{709}$$

where the turbulent Schmidt number is $\sigma_k = 0.82$ and the turbulent viscosity is computed as:

$$\mu_t = \rho C_\mu \frac{k^2}{\varepsilon} \tag{710}$$

with the turbulent kinetic energy k being defined as:

$$k = \frac{1}{2}\mathrm{tr}(\mathbf{R}) \tag{711}$$

Alternatively, according to [103], the Reynolds-stress diffusion can be written in the form:

$$\mathbf{D} = \mu\nabla\mathbf{R} + C_s\tau(\mathbf{R}\cdot\nabla)\mathbf{R} \tag{712}$$

where $C_s = 0.2$ and $\tau = k/\varepsilon$. This model is also known as the Generalized Gradient Diffusion Hypothesis, GGDH.

Turbulent Production

The turbulent production is obtained without recourse to modeling as follows:

$$\mathbf{P} = -\rho(\mathbf{R}\cdot\nabla\mathbf{v}^T + \nabla\mathbf{v}\cdot\mathbf{R}^T) = -\rho(\mathbf{R}\cdot\nabla\mathbf{v}^T + \nabla\mathbf{v}\cdot\mathbf{R}) \tag{713}$$

Buoyancy Production

For constant-density flows using the Boussinesq approximation, the buoyancy production is modeled as:

$$\mathbf{G} = \beta\frac{\mu_t}{\sigma_t}(\mathbf{g}\otimes\nabla T + \nabla T\otimes\mathbf{g}) \tag{714}$$

where β is the coefficient of thermal expansion, $\mathbf{g}$ is the gravitational vector, σ_t is the turbulent Prandtl number and ∇T is the temperature gradient vector.

For flows with varying density:

$$\mathbf{G} = \frac{\mu_t}{\rho\sigma_t}(\mathbf{g}\otimes\nabla\rho + \nabla\rho\otimes\mathbf{g}) \tag{715}$$

where $\nabla\rho$ is the density gradient vector.

Turbulent Dissipation Rate

The isotropic turbulent dissipation rate is obtained from a transport equation analogous to the K-Epsilon model (and with identical boundary conditions):

$$\frac{d}{dt}\int_V \rho\varepsilon dV + \int_A \rho\varepsilon(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_\varepsilon}\right)\nabla\varepsilon\cdot d\mathbf{a} + \int_V \left\{\frac{\varepsilon}{k}\left[C_{\varepsilon 1}\left(\frac{1}{2}\mathrm{tr}(\mathbf{P}) + \frac{1}{2}C_{\varepsilon 3}\mathrm{tr}(\mathbf{G})\right) - C_{\varepsilon 2}\rho\varepsilon\right]\right\}dv \tag{716}$$

where $C_{\varepsilon 1}$ and $C_{\varepsilon 2}$ are specified coefficients and $C_{\varepsilon 3}$ is determined as in the Standard K-Epsilon model. The coefficient $C_{\varepsilon 1}$ has the value 1.44. The coefficient $C_{\varepsilon 2}$ has the value 1.92 when the linear pressure-strain term is used, and has the value 1.83 when the quadratic pressure-strain term is used.

A two-layer formulation, which solves for k but prescribes ε algebraically with distance from the wall, is available for use with the linear pressure-strain model. This formulation is identical to the two-layer formulation using in the K-Epsilon models.

Dilatation Dissipation Rate

The dilatation dissipation Υ_M is modeled as in the K-Epsilon model, using the model of Sarkar *et al.*[109]:

$$\Upsilon_M = \frac{C_M k\varepsilon}{c^2} \tag{717}$$

where c is the speed of sound and $C_M = 2$.

Linear Pressure-Strain Model

The linear model of Gibson and Launder [104] for the pressure-strain term comprises five terms; these are the rapid part, the slow part, and their respective wall-reflection terms:

$$\underline{\phi} = \underline{\phi}_s + \underline{\phi}_r + \underline{\phi}_{r,b} + \underline{\phi}_{1w} + \underline{\phi}_{2w} \tag{718}$$

The slow pressure-strain term $\underline{\phi}_s$ is modeled as:

$$\underline{\phi}_s = -C_1\rho\frac{\varepsilon}{k}\left(\mathbf{R} - \frac{2}{3}k\mathbf{I}\right) \tag{719}$$

The rapid pressure-strain term $\underline{\phi}_r$ is modeled as:

$$\underline{\phi}_r = -C_2\left[\mathbf{P} - \frac{1}{3}\mathbf{I}\ \mathrm{tr}(\mathbf{P})\right] \tag{720}$$

The buoyancy contribution $\underline{\phi}_{r,b}$ is modeled as:

$$\underline{\phi}_{r,b} = -C_3\left[\mathbf{G} - \frac{1}{3}\mathbf{I}\ \mathrm{tr}(\mathbf{G})\right] \tag{721}$$

The slow wall-reflection term $\underline{\phi}_{1w}$ is modeled as:

$$\underline{\phi}_{1w} = \rho C_{1w}\frac{\varepsilon}{k}\left[(\mathbf{R}:\mathbf{N})\ \mathbf{I} - \frac{3}{2}(\mathbf{R}\cdot\mathbf{N} + \mathbf{N}\cdot\mathbf{R})\right]f_w \tag{722}$$

where:

$$f_w = \min\left(\frac{k^{3/2}}{C_l\varepsilon d}, f_w^{\max}\right) \qquad f_w^{\max} = 1.4 \tag{723}$$

d is the wall distance and the tensor $\mathbf{N}$ is given by:

$$\mathbf{N} = \mathbf{n}\otimes\mathbf{n} \tag{724}$$

and the "wall-normal unit vector" $\mathbf{n}$ is defined as the negative of the wall direction.

The rapid wall-reflection term $\underline{\phi}_{2w}$ is modeled as:

$$\underline{\phi}_{2w} = C_{2w}\left[(\underline{\phi}_r:\mathbf{N})\ \mathbf{I} - \frac{3}{2}(\underline{\phi}_r\cdot\mathbf{N} + \mathbf{N}\cdot\underline{\phi}_r)\right]f_w \tag{725}$$

Two possible sets of coefficients are used. In the event that wall functions are used, (that is, no two-layer model) they have the following constant values:

$$C_1 = 1.8\ ,\ C_2 = 0.6\ ,\ C_{1w} = 0.5\ ,\ C_{2w} = 0.3 \qquad C_3 = 0.5\ ,\ C_l = 2.5 \tag{726}$$

In the event that the two-layer model is used, the first four coefficients are expressed in terms of the turbulent Reynolds number and anisotropy tensor as recommended by Launder and Shima [107]:

$$C_1 = 1 + 2.58\ a\ a_2{}^{1/4}\{1 - \exp[(-0.0067Re_t)^2]\} \tag{727}$$

$$C_2 = 0.75\sqrt{a} \tag{728}$$

$$C_{1w} = -\frac{2}{3}C_1 + 1.67 \tag{729}$$

$$C_{2w} = \max\left(\frac{4C_2 - 1}{6C_2}, 0\right) \tag{730}$$

$$Re_t = \frac{k^2}{\varepsilon\nu} \tag{731}$$

The parameter a and the tensor invariants a_2 and a_3 are defined as:

$$a = 1 - \frac{9}{8}(a_2 - a_3) \tag{732}$$

$$a_2 = \mathbf{A}:\mathbf{A} \tag{733}$$

$$a_3 = A_{ik}A_{kj}A_{ji} \tag{734}$$

where the anisotropy tensor $\mathbf{A}$ is defined as:

$$\mathbf{A} = \frac{\mathbf{R}}{k} - \frac{2}{3}\mathbf{I} \tag{735}$$

Quadratic Pressure-Strain Model

The quadratic pressure-strain model of Speziale, Sarkar, and Gatski [110] is:

$$\begin{aligned}\underline{\phi} = &-[C_{s1}\rho\varepsilon + C_{r4}\mathbf{tr}(\mathbf{P} + \mathbf{G})]\,\mathbf{A} \\ &+ C_{s2}\rho\varepsilon\left(\mathbf{A}\cdot\mathbf{A} - \frac{1}{3}\mathbf{A}:\mathbf{A}\,\mathbf{I}\right) + (C_{r3} - C_{r3}^{*}\sqrt{\mathbf{A}:\mathbf{A}})\rho k\,\mathbf{S} \\ &+ C_{r1}\rho k\left(\mathbf{A}\cdot\mathbf{S} + \mathbf{S}\cdot\mathbf{A} - \frac{2}{3}\mathbf{A}:\mathbf{S}\,\mathbf{I}\right) + C_{r2}\rho k(\mathbf{A}\cdot\mathbf{W}^T + \mathbf{W}\cdot\mathbf{A})\end{aligned} \tag{736}$$

where the anisotropy tensor $\mathbf{A}$ is defined as in Eqn. (735), the strain rate tensor $\mathbf{S}$ is defined as:

$$\mathbf{S} = \frac{1}{2}(\nabla\mathbf{v} + \nabla\mathbf{v}^T) \tag{737}$$

and the rotation rate tensor $\mathbf{W}$ is defined as:

$$\mathbf{W} = \frac{1}{2}(\nabla\mathbf{v} - \nabla\mathbf{v}^T) \tag{738}$$

The coefficients have the following values:

$$C_{s1} = 1.7 \quad , \quad C_{r4} = 0.9 \quad , \quad C_{s2} = 1.05 \quad , \quad C_{r3} = 0.8$$
$$C_{r3}^{*} = 0.65 \quad , \quad C_{r4} = 0.625 \quad , \quad C_{r2} = 0.2 \tag{739}$$

Wall Treatment

The following equations present the wall-treatment formulations. In general, the high-y^+ formulation is equivalent to the traditional wall-function approach, and the all-y^+ and two-layer all-y^+ formulation.

For the all-y^+ formulation, a blending function g is defined in terms of the wall-distance Reynolds number:

$$g = \exp\left(-\frac{Re_y}{11}\right) \tag{740}$$

and $Re_y = \sqrt{k}y/\nu$, a wall-distance-based Reynolds number.

Reference Velocity

$$u^* = \begin{cases} \sqrt{C_\mu^{1/2} k} & ; \quad \text{high-}y^+ \\ \sqrt{g\nu u/y + (1-g)C_\mu^{1/2} k} & ; \quad \text{two-layer all-}y^+ \\ \sqrt{g\nu u/y + (1-g)C_\mu^{1/2} k} & ; \quad \text{all-}y^+ \end{cases} \tag{741}$$

Wall-Cell Normal Gradient

$$\frac{G_k}{\tau_w} = \begin{cases} \dfrac{u}{u^+ \kappa y} & ; \quad \text{high-}y^+ \\ g\dfrac{u}{y} + (1-g)\dfrac{u}{u^+ \kappa y} & ; \quad \text{two-layer all-}y^+ \\ 0.95\dfrac{u^2}{u^{+2}\nu}\dfrac{\nu_t^+}{(1+\nu_t^+)^2} & ; \quad \text{all-}y^+ \end{cases} \tag{742}$$

where $\nu_t^+ = \kappa y^+\left[1 - \exp\left(\frac{-y^+}{26}\right)\right]^2$.

Wall-Cell Dissipation

$$\varepsilon = \begin{cases} \dfrac{u^{*3}}{\kappa y} & ; \quad \text{high-}y^+ \\ \dfrac{k^{3/2}}{l_\varepsilon} & ; \text{ two-layer all-}y^+ \\ g\dfrac{2\nu k}{y^2} + (1-g)\dfrac{u^{*3}}{\kappa y} & ; \quad \text{all-}y^+ \end{cases} \tag{743}$$

Boundary, Region and Initial Conditions

Axis Boundaries

On an axis boundary, the procedure for setting the specific Reynolds stress tensor at the face, $\mathbf{R}_f$, is similar to the symmetry plane. In this case, however, the face values of the tensor in the coordinate system parallel to the axis are as follows:

$$\mathbf{R}'_f = \begin{bmatrix} R'_{11} & 0 & 0 \\ 0 & R'_{22} & 0 \\ 0 & 0 & R'_{33} \end{bmatrix}_c \tag{744}$$

Symmetry Boundaries

On a symmetry boundary, the procedure for setting the specific Reynolds stress tensor at the face, $\mathbf{R}_f$, is as follows. First, the tensor at the adjacent cell center, $\mathbf{R}$, is rotated such that it is parallel to the symmetry plane to obtain a new tensor termed $\mathbf{R}'$. The symmetry boundary face values of the tensor in this rotated coordinate system are explicitly set as follows:

$$\mathbf{R}'_f = \begin{bmatrix} R'_{11} & 0 & R'_{13} \\ 0 & R'_{22} & 0 \\ R'_{31} & 0 & R'_{33} \end{bmatrix}_c \tag{745}$$

The values of $\mathbf{R}_f$ are computed by rotating the $\mathbf{R}'_f$ tensor back to the original orientation of $\mathbf{R}$. (This is made easy by the fact that the inverse of the rotation tensor is simply its transpose, since it is an orthogonal tensor.)

Wall Boundaries

At walls, a Neumann boundary condition is used for the Reynolds stresses $\mathbf{R}$, that is, $\partial\mathbf{R}/\partial n|_w = 0$. In addition, a method developed by Hadzic [105] is used to impose the value of production of each stress component in accordance with a wall function approach. This method is summarized here.

In a coordinate system oriented with the wall, the production of turbulent kinetic energy is

$$G_k = \tau_w \frac{\partial \mathrm{v}_t}{\partial n} \tag{746}$$

where $\partial \mathrm{v}_t/\partial n$ is the wall-normal velocity gradient. If one obtains the production G_k and the shear stress τ_w from wall functions, and assumes that all gradients other than $\partial \mathrm{v}_t/\partial n$ are negligible, a velocity gradient tensor in the wall-oriented coordinates may be constructed as follows:

$$\nabla\mathbf{v}_{wall} = \begin{bmatrix} 0 & G_k/\tau_w & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{747}$$

The velocity gradient tensor in Cartesian coordinates is then obtained by rotating $\nabla\mathbf{v}_{wall}$ using the appropriate tensor transformation.

If this technique is used with a consistent method of evaluating ε so as to satisfy local equilibrium ($G_k = \varepsilon$), this method will ensure the correct ratios of Reynolds stress to turbulent kinetic energy in a local equilibrium situation. However, it will allow the stress values to vary whenever a non-equilibrium situation is encountered.

Flow, Region and Initial Conditions

When defining values for flow boundary, region and initial conditions, you have several choices for specifying the Reynolds stress tensor and the isotropic dissipation rate:

- Enter the values directly.
- Specify the turbulent dissipation rate but have the Reynolds stresses derived from:

$$\mathbf{R} = 2/3k\mathbf{I}$$

where k is the supplied turbulent kinetic energy and $\mathbf{I}$ is the identity matrix.

- Have them derived from turbulence intensity I and turbulent length scale L using the relations:

$$\mathbf{R} \approx \frac{2}{3}(I\mathrm{v})^2\mathbf{I}$$

$$\varepsilon \approx \frac{C_\mu^{3/4}\left[\frac{1}{2}tr(\mathbf{R})\right]^{3/2}}{L}$$

where v is the local velocity magnitude

- Have them derived from turbulence intensity and turbulent viscosity ratio μ_t/μ using the relations:

$$\mathbf{R} \approx (I\mathrm{v})^2\mathbf{I}$$

$$\varepsilon \approx \frac{C_\mu\left[\frac{1}{2}tr(\mathbf{R})\right]^2}{(\mu_t/\mu)\mu} \tag{748}$$

where v is the local velocity magnitude

Reynolds Stress Transport Turbulence Nomenclature

The following is a guide to the terms used in this section.

Scalars

d	distance to the nearest wall
k	turbulent kinetic energy
l	turbulent length scale
t	time
u^*	reference velocity used in wall functions
u^+	wall-parallel velocity non-dimensionalized with u^*
I	turbulence intensity
T	temperature
L	turbulent length scale
V	cell volume
v	velocity magnitude
β	coefficient of thermal expansion
ε	turbulent dissipation rate
ρ	density
μ	dynamic viscosity
μ_t	turbulent viscosity
σ_t	turbulent Prandtl number
ν	kinematic viscosity
σ_k, σ_ε	turbulent Schmidt numbers
Υ_M	isotropic dilitation dissipation

Tensors or Column Vectors

A	anisotropy tensor
D	diffusion of Reynolds stresses
G	turbulent production due to buoyancy
P	turbulent production
I	identity matrix
R	specific Reynolds stress tensor
S	strain rate tensor
W	vorticity tensor
$\underline{\phi}$	pressure-strain term

Cartesian Vectors

a	face area vector
g	gravitational vector
n	wall normal vector
v	velocity
$\mathbf{v}_g$	grid velocity

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Divergence of vector quantity

$$\nabla\cdot\mathbf{v} = \frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S}:\mathbf{T} = S_{ij}T_{ji}$$

- Trace of a tensor quantity:

$$tr(\mathbf{T}) = T_{ii}$$

Reynolds Stress Transport Turbulence Bibliography

[103] Daly, B.J. and Harlow, F.H. 1970. “Transport equations of turbulence”, *Physics of Fluids* **13**, pp. 2634-2649.

[104] Gibson, M.M. and Launder, B.E. 1978. “Ground effects on pressure fluctuations in the atmospheric boundary layer”, *J. Fluid Mech.*, **86**(3), pp. 491-511.

[105] Hadzic, I. 1999. "Second-moment closure modeling of transitional and unsteady turbulent flows", doctoral thesis, Technical University of Delft.

[106] Lardeau, S. and Manceau, R. 2014. "Computations of Complex Flow Configurations Using a Modified Elliptic-Blending Reynolds-Stress Model", *Proc. of the 10th Eng. Turb. Modelling and Measurement Conf.*, Marbella, Spain.

[107] Launder, B.E. and Shima, N. 1989. "Second Moment Closure for the Near-Wall Sublayer: Development and Application", *AIAA J.* **27**(10), October, pp. 1319-1325.

[108] Lien, F.S. and Leschziner, M.A. 1994. "Assessment of turbulence-transport models including non-linear RNG eddy-viscosity formulation and second-moment closure for flow over a backward-facing step", *Computers Fluids*, **23**(8), pp. 983-1004.

[109] Sarkar, S., and Balakrishnan, L. 1990. "Application of a Reynolds-stress turbulence model to the compressible shear layer", ICASE Report 90-18, NASA CR 182002.

[110] Speziale, C.G., Sarkar, S., and Gatski, T.B. 1991. "Modelling the pressure-strain correlation of turbulence: an invariant dynamical systems approach", *J. Fluid Mech.*, **227**, pp. 245-272.

Reynolds Stress Conditions and Values

This section provides reference information for conditions that are common to boundaries, regions, and initial conditions:

- Turbulence specification

Information is also provided on boundary, region and initial condition values that are common to the Reynolds Stress family of turbulence models:

- UU Reynolds stress
- UV Reynolds stress
- UW Reynolds stress
- VV Reynolds stress
- VW Reynolds stress
- WW Reynolds stress

Turbulence Specification

This condition allows you to specify the method for defining the turbulence profile for boundaries and initial conditions.

Properties

Method Choose the method for defining turbulence. The models require the Reynolds stress tensor ($\mathbf{R}$) and the turbulent dissipation rate (ε). This can done directly or derived from other input values.

K + Epsilon — The Reynolds stress tensor is calculated from the turbulent kinetic energy (k) using Eqn. (743) and the turbulent dissipation rate is specified directly. A Turbulent Kinetic Energy node and a Turbulent Dissipation Rate node will be added to the Values node.

Intensity and Length Scale — Turbulent kinetic energy (k) and turbulent dissipation rate are calculated from turbulence intensity (I) and length scale (L) using Eqn. (745). The Reynolds stress tensor is then calculated from k using Eqn. (743). Do not use this method if the velocity field is initialized to zero (that is, $\mathrm{v} = 0$). A Turbulence Intensity node and a Turbulent Length Scale node will be added to the Values node. Additionally, a Turbulent Velocity Scale node will appear in the Initial Conditions node.

Intensity + Viscosity Ratio — The Reynolds stress tensor and the dissipation rate are calculated from the turbulence intensity I and the ratio of turbulent to laminar viscosity (μ_t/μ) using Eqn. (744). Do not use this method if the velocity field is initialized to zero (that is, $\mathrm{v} = 0$). A Turbulence Intensity node and a Turbulent Viscosity Ratio node will be added to the Values node. Additionally, a Turbulent Velocity Scale node will appear in the Initial Conditions node.

RS + Epsilon — Explicit specification of the components of $\mathbf{R}$ and of the isotropic dissipation rate ε. This method is not available for porous-dominated regions.

UU Reynolds Stress

This node is added to the Values node when the Method property of the Conditions Turbulence Specification node is set to `RS + Epsilon`. It requires input of the UU component of the Reynolds stress tensor.

Properties

This value is input as a scalar profile.

UV Reynolds Stress

This node is added to the Values node when the Method property of the Conditions Turbulence Specification node is set to `RS + Epsilon`. It requires input of the UV component of the Reynolds stress tensor.

Properties

This value is input as a scalar profile.

UW Reynolds Stress

This node is added to the Values node when the Method property of the Conditions Turbulence Specification node is set to `RS + Epsilon`. It requires input of the UW component of the Reynolds stress tensor.

Properties

This value is input as a scalar profile.

VV Reynolds Stress

This node is added to the Values node when the Method property of the Conditions Turbulence Specification node is set to `RS + Epsilon`. It requires input of the VV component of the Reynolds stress tensor.

Properties

This value is input as a scalar profile.

VW Reynolds Stress

This node is added to the Values node when the Method property of the Conditions Turbulence Specification node is set to `RS + Epsilon`. It requires input of the VW component of the Reynolds stress tensor.

Properties

This value is input as a scalar profile.

WW Reynolds Stress

This node is added to the Values node when the Method property of the Conditions Turbulence Specification node is set to `RS + Epsilon`. It requires input of the WW component of the Reynolds stress tensor.

Properties

This value is input as a scalar profile.

Boundary Conditions and Values Reference

The following boundary conditions are required by the Reynolds Stress turbulence models:

- turbulence specification

In addition, the following boundary values may need to be set for these models, depending on the selected boundary type:

- Turbulent dissipation rate
- Turbulence intensity
- Turbulent kinetic energy
- Turbulent length scale
- Turbulent viscosity ratio
- UU Reynolds stress component
- UV Reynolds stress component
- UW Reynolds stress component
- VV Reynolds stress component
- VW Reynolds stress component
- WW Reynolds stress component

Boundary values are elected by the boundary type, and/or by methods selected in corresponding nodes in the boundary conditions node. The boundary types that are available are listed in the Boundary Type section of each model.

Region Conditions and Values Reference

The following region conditions may need to be set for the Reynolds Stress turbulence model, depending on the selected region type:

- Turbulence specification

Region conditions are elected by the region type, based on the models that have been activated in the continuum.

The region conditions required for each region type when the Reynolds stress turbulence model is activated are listed here.

The following region values may need to be set for the Reynolds Stress turbulence models, depending on the selected region type:

- Turbulent dissipation rate
- Turbulence intensity
- Turbulent kinetic energy
- Turbulent length scale
- Turbulent viscosity ratio
- UU Reynolds stress component
- UV Reynolds stress component
- UW Reynolds stress component
- VV Reynolds stress component
- VW Reynolds stress component
- WW Reynolds stress component

Information is also provided about region values available for the Reynolds Stress turbulence model in STAR-CCM+:

- UU Stress User Source
- UV Stress User Source
- UW Stress User Source
- VV Stress User Source
- VW Stress User Source
- WW Stress User Source

UU Stress User Source

This node is added when the when the Reynolds Stress turbulence model is used, and the Turbulence Source Option region condition is activated. It is entered as a scalar profile.

UV Stress User Source

This node is added when the when the Reynolds Stress turbulence model is used, and the Turbulence Source Option region condition is activated. It is entered as a scalar profile.

UW Stress User Source

This node is added when the when the Reynolds Stress turbulence model is used, and the Turbulence Source Option region condition is activated. It is entered as a scalar profile.

VV Stress User Source

This node is added when the when the Reynolds Stress turbulence model is used, and the Turbulence Source Option region condition is activated. It is entered as a scalar profile.

VW Stress User Source

This node is added when the when the Reynolds Stress turbulence model is used, and the Turbulence Source Option region condition is activated. It is entered as a scalar profile.

WW Stress User Source

This node is added when the when the Reynolds Stress turbulence model is used, and the Turbulence Source Option region condition is activated. It is entered as a scalar profile.

Reynolds Stress Transport Turbulence Field Functions Reference

The following primitive field functions are made available to the simulation when the Reynolds Stress turbulence model is used. These objects all have the same base properties. The same field functions can be exposed for the non-linear K-Epsilon models. See Exposing Reynolds Stresses for Standard

Non-Linear K-Epsilon Models.

Reynolds Stress uu	The specific normal stress, $\overline{u'u'}$.
Reynolds Stress vv	The specific normal stress, $\overline{v'v'}$.
Reynolds Stress ww	The specific normal stress, $\overline{w'w'}$.
Reynolds Stress uv	The specific shear stress, $\overline{u'v'}$.
Reynolds Stress uw	The specific shear stress, $\overline{u'w'}$.
Reynolds Stress vw	The specific shear stress, $\overline{v'w'}$.
Turbulent Kinetic Energy	The variable k, computed from the trace of the specific Reynolds stress tensor.
Turbulent Dissipation Rate	The transported variable ε.

Using Smagorinsky Subgrid Scale Turbulence

This section explains the following key concepts:

- What is the Smagorinsky Subgrid Scale model?
- Working with the Smagorinsky Subgrid Scale viscosity solver
- Formulation of the model equations

Field function information specific to this model is also provided.

What Is the Smagorinsky Subgrid Scale Model?

The Smagorinsky Subgrid Scale model uses a mixing length hypothesis to model the subgrid scale stresses. In its basic form, it is simple and not computationally expensive. However, a damping function is needed for the accurate simulation of wall bounded calculations. The damping function is a non-local operation, requiring the use of a KD-tree to communicate wall information to the interior. In simulations with very complex geometries, this KD-tree must be broadcast to each parallel node, resulting in a memory and communication overhead. An update policy can be used to ameliorate the communication overhead, but not the memory overhead.

It is also noted that the model coefficient C_s (see Eqn. (761)) is not universal. Typical values of C_s reported in the literature range from 0.165 for homogeneous isotropic decaying turbulence to 0.07 for channel flows. The default value of $C_s = 0.1$ is a compromise between these extremes.

Smagorinsky Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques used in STAR-CCM+, we recommend that you not make any changes within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

Limit Length Scale	Determines whether the mixing length minimum is applied to the subgrid scale viscosity, see Eqn. (761).	
	Activated	The minimum is applied.
	Deactivated	The minimum is not applied.
Cs	The coefficient C_s, Eqn. (761). Default value is 0.1.	

Ct	The coefficient C_t, Eqn. (764).	
Van Driest Damping Option	Selects the Van Driest damping function, see Eqn. (762).	
	None	No damping.
	Original	Van Driest's original damping function.
	Modified	Van Driest's damping function modified according to Piomelli and co-workers [113].
Kappa	The von Karman constant κ, Eqn. (761).	
A	Van Driest damping coefficient, Eqn. (762).	

Working with the Smagorinsky Subgrid Scale (SGS) Viscosity Solver

This section describes the solver that updates the subgrid scale viscosity μ_t in all the continua for which the Smagorinsky Subgrid Scale model is activated. Since this model does not contain any transport equations, there are no other solvers associated with the model.

Properties

Damping Factor Update Policy	Policy for updating the damping factor f_v .	
	Time Step	Updates will be triggered on a per time-step basis.
	Iteration	Updates will be triggered on a per iteration basis.
Damping Factor Update Frequency	This parameter governs the frequency at which the damping factor f_v will be updated in the free stream. The default is 1, which means that the update occurs every iteration. Increasing this parameter to a larger interval, say 10, might improve parallel performance in certain situations. This is because the free-stream values are stored in a KD tree that is broadcast to each parallel node. incurring a communication overhead.	

Expert Properties

The properties in the *Expert* category are designed to be used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the ap coefficients, the residual, the corrections, the reconstruction gradients and the cell gradients. These quantities will become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	

Activated	The solver will not update the solution.
Deactivated	The solver will update the solution.

Smagorinsky Subgrid Scale Turbulence Formulation

This section includes the following:

- Smagorinsky Subgrid Scale Viscosity
- Turbulent Scales
- Nomenclature
- Bibliography

Smagorinsky Subgrid Scale Viscosity

The Smagorinsky Subgrid Scale model [111] provides the following mixing-length type formula for the subgrid scale viscosity:

$$\mu_t = \rho \Delta^2 S \tag{759}$$

where $S = \sqrt{2\mathbf{S}:\mathbf{S}}$ is the modulus of the strain rate defined as:

$$\mathbf{S} = \frac{1}{2}(\nabla \mathbf{v} + \nabla \mathbf{v}^T) \tag{760}$$

ρ is the density and Δ is the length scale or grid filter width. The length scale is directly related to the cell volume V and the wall distance d as follows:

$$\Delta = \begin{cases} f_v\, C_s V^{1/3} & \text{if length scale limit not applied} \\ f_v \min(\kappa d, C_s V^{1/3}) & \text{if length scale limit is applied} \end{cases} \tag{761}$$

where κ is the von Karman constant (default value 0.41) and f_v is the Van Driest damping function computed as follows:

$$f_v = \begin{cases} 1 & \text{no damping} \\ 1 - \exp\left(-\frac{y^+}{A}\right) & \text{standard} \\ \sqrt{1 - \exp\left(-\frac{y^+}{A}\right)^3} & \text{modified} \end{cases} \tag{762}$$

The modified version of the Van Driest damping function [112] is attributed to Piomelli and co-workers [113]. The damping coefficient A has the default value 25. The dimensionless wall distance is defined as:

$$y^+ = \frac{u^* \mathrm{d}}{\nu} \tag{763}$$

where u^* is the friction velocity computed from the instantaneous wall shear stress at the nearest wall face, d is the wall distance, and ν is the kinematic viscosity.

The coefficient C_s has a default value of 0.1.

Turbulent Scales

Certain combustion and Lagrangian models require the computation of subgrid length and time scales.

The turbulent time scale is:

$$t = \frac{C_t}{S} \tag{764}$$

This relation is obtained from the ratio of the subgrid scale turbulent kinetic energy k_{SGS} and its dissipation rate ε_{SGS}. When computing this ratio, the subgrid scale turbulent kinetic energy is modeled as:

$$k_{SGS} = C_t \nu_t S \tag{765}$$

where S is the norm of the rate of strain tensor defined in Eqn. (760), and ν_t is the turbulent viscosity. The constant C_t is assigned a default value of 3.5. This value must be greater than $(\sqrt{3})/2$ to satisfy realizability constraints [114].

The dissipation rate ε_{SGS} is modeled as $\nu_t S^2$ from equilibrium considerations.

The turbulent length scale is:

$$l = \frac{C_t^{3/2} \nu_t^{1/2}}{S^{1/2}} \tag{766}$$

to satisfy $l = \sqrt{k_{SGS}}\, t$.

Smagorinsky Subgrid Scale Turbulence Nomenclature

The following is a guide to the terms used in this section.

Scalars

d	distance to the nearest wall
E	spectral energy
S	strain rate modulus
T	turbulent time scale
L	turbulent length scale
V	cell volume
ρ	density
μ_t	turbulent (subgrid scale) viscosity
Δ	filter width
κ	von Karman constant

Tensors or Column Vectors

S	strain rate tensor

Field Operators

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S{:}T} = S_{ij} T_{ji}$$

Smagorinsky Subgrid Scale Turbulence Bibliography

[111] Smagorinsky, J. 1963. "General Circulation Experiments with the Primitive Equations: Part I, The Basic Experiment," *Monthly Weather Review*, **91**, pp. 99-164.

[112] Van Driest, E. R. 1956. "On Turbulent Flow Near a Wall," *J. Aero. Sci.*, pp. 1007-1011.

[113] Balaras, E., Benocci, C., and Piomelli, U. 1996. "Two-Layer Approximate Boundary Conditions for Large-Eddy Simulations," *AIAA Journal* **34**(6), pp. 1111-1119.

[114] Vreman, B., Geurts, B., and Kuerten, H. 1994. "Realizability conditions for the turbulent stress tensor in Large Eddy Simulation," *J. Fluid Mech.*, **278**, pp. 351–362.

Smagorinsky Subgrid Scale Turbulence Field Functions Reference

The following primitive field function is made available to the simulation when the Smagorinsky Subgrid Scale model is used. This object has typical base properties.

SGS Wall Damping Factor	The Van Driest damping function f_v defined in Eqn. (762).

Using Dynamic Smagorinsky Subgrid Scale Turbulence

This section explains the following key concepts:

- What is the Dynamic Smagorinsky Subgrid Scale model?
- Working with the Dynamic Smagorinsky SGS viscosity solver
- Formulation of the model equations

What Is the Dynamic Smagorinsky Subgrid Scale Model?

The Dynamic Smagorinsky Subgrid Scale model has the same basic form as the Smagorinsky model, however instead of using a single user-defined constant, the model computes a local time-varying constant by test-filtering the flow field at a length scale greater than the grid length scale. This dynamic variation of the constant gives the model its name, and allows it to compute the correct result for wall-bounded flows without the use of damping functions.



Dynamic Smagorinsky Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, do not change settings within the *Expert* category. The values in that category

reflect both the design of the model and discretization approaches that have been optimized for accuracy and performance. Tampering with them can diminish the effectiveness of the model.

Ct	The time-scale coefficient C_t, Eqn. (764).
Filter Width Ratio	Ratio of test filter width to grid filter width (see Eqn. (771)).
Minimum Cs^2 Value	Clips values less than the specified value of C_s^2 (see Eqn. (772)).
Maximum Cs^2 Value	Clips values greater than the specified value of C_s^2 (see Eqn. (772)).
Average Parameter	Optional averaging in calculation of C_s^2 (see Eqn. (772)). Activated — Uses averaging (default). Deactivated — Does not use averaging.

Working with the Dynamic Smagorinsky SGS Viscosity Solver

This section describes the solver that updates the subgrid scale viscosity μ_t in all the continua for which the Dynamic Smagorinsky Subgrid Scale model is activated.



Expert Properties

The properties in the *Expert* category are designed for a temporary debug situation, at the expense of simulation accuracy and for extra memory usage. Do not change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Dynamic Smagorinsky SGS Turbulence Formulation

This section includes the following:

- Formulation of the Dynamic Smagorinsky SGS model
- Bibliography

Formulation of the Dynamic Smagorinsky SGS Model

The calculation of the turbulent viscosity in the Dynamic Smagorinsky SGS model follows the same equation as for the standard Smagorinsky SGS model:

$$\mu_t = \rho \Delta^2 S \tag{767}$$

However, the parameter Δ^2 is computed dynamically as a function of a number of test-filtered variables [115]:

$$\Delta^2 = C_s^2 V^{2/3} \tag{768}$$

If $\bar{\phi}$ is the grid-filtered LES variable in a cell, then the corresponding test-filtered value $\tilde{\bar{\phi}}$ is given as:

$$\tilde{\bar{\phi}} = \frac{1}{\sum_{n=0}^{n=N} V_n} \sum_{n=0}^{n=N} \bar{\phi}_n V_n \tag{769}$$

where the subscripts denote cell number, cell 0 is the current cell, and cells 1 to N are the face neighbors of cell 0.

Continuing with the superscript "~" notation to denote test-filtered variables, the dynamic parameter is computed as follows [116]:

$$L_{ij} = \bar{u}_i \tilde{\bar{u}}_j - \tilde{\bar{u}}_i \tilde{\bar{u}}_j \tag{770}$$

$$M_{ij} = 2\bar{L}^2 \left(|\bar{S}| \tilde{\bar{S}}_{ij} - \frac{\tilde{L}^2}{\bar{L}^2} |\tilde{\bar{S}}| \tilde{\bar{S}}_{ij} \right) \tag{771}$$

where M_{ij} is a tensor that arises during the derivation of the model by subtracting the Smagorinsky subgrid stress tensor at the grid filter scale from the same tensor at the test filter scale. $\bar{L}$ is the grid filter length $V^{2/3}$, $\tilde{L}/\bar{L}$ is the filter width ratio, and:

$$C_s^2 = \frac{\langle L_{ij} M_{ij} \rangle}{\langle M_{ij} M_{ij} \rangle} \tag{772}$$

where the brackets <> denote averaging. Since, in most simulations, there is no homogeneous direction over which to average, the same stencil is used as for the filtering operation. This averaging is optional; you can activate or deactivate it in the Average Parameter expert property of the model in the GUI.

To increase stability in the calculation, clip the value of C_s^2 using the minimum and maximum values in the expert properties of the model.

Dynamic Smagorinsky Subgrid Scale Turbulence Bibliography

[115] Germano, M., Piomelli, U., Moin, P., and Cabot, W. H. 1991. "A Dynamic Subgrid-Scale Eddy Viscosity Model", *Physics of Fluids A*, **3**(7), pp. 1760-1765.

[116] Lilly, D. K. 1992. "A proposed modification of the Germano subgrid-scale closure method", *Physics of Fluids A* **4**(3), pp. 633-635.

Using WALE Subgrid Scale Turbulence

This section explains the following key concepts:

- What is the WALE Subgrid Scale model?
- Working with the WALE subgrid scale viscosity solver
- Formulation of the model equations

What Is the WALE Subgrid Scale Model?

The WALE (Wall-Adapting Local-Eddy Viscosity) Subgrid Scale model [117] uses an algebraic formulation to model the subgrid scale stresses. This is the simplest, and hence the least computationally expensive, subgrid scale modeling approach. Similar to the Smagorinsky Subgrid Scale model, it suffers from the limitation that the model coefficient C_w (see Eqn. (774)) is not universal. Typical values of C_w commonly reported range from 0.5 for homogeneous isotropic decaying turbulence, to 0.325 for channel flows. Validations using STAR-CCM+ have shown that the WALE model is seemingly less sensitive to the value of the coefficient than the Smagorinsky model, and that the default value of $C_w = 0.544$ works very well for both homogeneous isotropic decaying turbulence and for channel flows. One advantage of the WALE model is that it automatically gives accurate scaling at walls.

WALE Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques used in STAR-CCM+, we recommend that you not make any changes within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

Cw	The coefficient C_w, Eqn. (773). Default value is 0.544.
Ct	The coefficient C_t, Eqn. (778).
Kappa	The von Karman constant κ, Eqn. (774).
Limit Length Scale	Determines whether the mixing length minimum is applied to the subgrid scale viscosity, see Eqn. (774).

Ticked	The minimum is applied.
Cleared	The minimum is not applied.

Working with the WALE Subgrid Scale (SGS) Viscosity Solver

This section describes the solver that updates the subgrid scale viscosity μ_t in all the continua for which the WALE Subgrid Scale model is activated. Since this model does not contain any transport equations, there are no other solvers associated with the model.

Expert Properties

The properties in the *Expert* category are designed to be used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Ticked	The solver will not update the solution.
	Cleared	The solver will update the solution.
Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients and the cell gradients. These quantities will become available as field functions during subsequent iterations.	
	Ticked	Retain temporary storage.
	Cleared	Do not retain temporary storage.

WALE Subgrid Scale Turbulence Formulation

This section includes the following:

- WALE Subgrid Scale Viscosity

- Turbulent Scales
- Nomenclature
- Bibliography

WALE Subgrid Scale Viscosity

The WALE Subgrid Scale model [117] provides the following mixing-length type formula for the subgrid scale viscosity:

$$\mu_t = \rho \Delta^2 S_w \tag{773}$$

where ρ is the density and Δ is the length scale or grid filter width defined in terms of the cell volume V as:

$$\Delta = \begin{cases} C_w V^{1/3} & \text{if length scale limit is not applied} \\ \min(\kappa d, C_w V^{1/3}) & \text{if length scale limit is applied} \end{cases} \tag{774}$$

where κ is the von Karman constant (default value 0.41).

The deformation parameter S_w is defined as:

$$S_w = \frac{\mathbf{S}_d : \mathbf{S}_d^{3/2}}{\mathbf{S}_d : \mathbf{S}_d^{5/4} + \mathbf{S} : \mathbf{S}^{5/2}} \tag{775}$$

where $\mathbf{S}$ is the strain rate tensor defined as:

$$\mathbf{S} = \frac{1}{2}(\nabla \mathbf{v} + \nabla \mathbf{v}^T) \tag{776}$$

and the tensor $\mathbf{S}_d$ is defined as:

$$\mathbf{S}_d = \frac{1}{2}[\nabla \mathbf{v} \cdot \nabla \mathbf{v} + (\nabla \mathbf{v} \cdot \nabla \mathbf{v})^T] - \frac{1}{3} tr(\nabla \mathbf{v} \cdot \nabla \mathbf{v}) \mathbf{I} \tag{777}$$

The coefficient C_w has a default value of 0.544.

Turbulent Scales

Certain combustion and Lagrangian models require the computation of subgrid length and time scales.

The turbulent time scale is:

$$t = \frac{C_t}{S} \tag{778}$$

This relation is obtained from the ratio of the subgrid scale turbulent kinetic energy k_{SGS} and its dissipation rate ε_{SGS}. When computing this ratio, the subgrid scale turbulent kinetic energy is modeled as:

$$k_{SGS} = C_t \nu_t S \tag{779}$$

where S is the norm of the rate of strain tensor defined in Eqn. (776), and ν_t is the turbulent viscosity. The constant C_t is assigned a default value of 3.5. This value must be greater than $(\sqrt{3})/2$ to satisfy realizability constraints [114].

The dissipation rate ε_{SGS} is modeled as $\nu_t S^2$ from equilibrium considerations.

The turbulent length scale is:

$$l = \frac{C_t^{3/2} \nu_t^{1/2}}{S^{1/2}} \tag{780}$$

to satisfy $l = \sqrt{k_{SGS}}\, t$.

WALE Subgrid Scale Turbulence Nomenclature

The following is a guide to the terms used in this section.

Scalars

S_w	deformation parameter
T	turbulent time scale
L	turbulent length scale
V	cell volume
ρ	density
μ_t	turbulent (subgrid scale) viscosity
Δ	filter width

Tensors or Column Vectors

$\mathbf{S}$	strain rate tensor
$\mathbf{S}_d$	deformation tensor

Field Operators

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S}:\mathbf{T} = S_{ij}T_{ji}$$

WALE Subgrid Scale Turbulence Bibliography

[117] Nicoud, F. and Ducros, F., 1999. "Subgrid-Scale Stress Modelling Based on the Square of the Velocity Gradient Tensor," *Flow, Turbulence and Combustion,* **62**, pp. 183-200.

Using the Turbulence Suppression Model

This section describes the use of the Turbulence Suppression model, specifically the following aspects:

- What is the Turbulence Suppression model?
- What is a transition boundary?
- How is the region defined in which turbulence is suppressed?
- How is turbulence suppressed?
- Calculating transition boundary distance
- What is the Transition Boundary Distance solver?
- Selecting the Turbulence Suppression and Transition Boundary Distance models
- Using the Transition Boundary Distance model
- Practical issues of running a simulation with the Transition model

A field function reference is also provided.

What Is the Turbulence Suppression Model?

The Turbulence Suppression model allows the user to mimic the effect of transition by suppressing the turbulence in a certain pre-defined region. Since it does not have a predictive capability, it is arguably not a true transition model. However, if the location of transition is known a priori (for instance due to a boundary layer trip) this approach is computationally efficient.

What Is a Transition Boundary?

A transition boundary is a wall boundary adjacent to which turbulence effects will be suppressed. A wall boundary in a turbulent flow may be turned into a transition boundary by first activating the Transition model in the continuum, and then selecting the turbulence suppression option for that wall. Note that the transition model will have no effect in a physics continuum unless the turbulence suppression option is selected on a wall boundary.

How Is the Region Defined in Which Turbulence Is Suppressed?

One possible approach to defining a region in space in which turbulence is suppressed would be to identify an entire region in the volume mesh. However, for certain grid topologies it might impractical to create the appropriate region in the volume mesh. Instead, the approach taken in STAR-CCM+ is to compute two distances for each field point in the interior:

1. The wall distance, defined as the distance from a field point to the nearest point on a wall that is not a transition boundary
2. The transition boundary distance, defined as the distance from the field point to the nearest point on a transition boundary

The region in space in which turbulence is suppressed is then defined as any point at which the transition boundary distance is less than the wall distance. The way in which a turbulence suppression region is created from a transition boundary is illustrated in the following figures.





How Is Turbulence Suppressed?

Turbulence suppression is achieved in a region as follows:

1. Setting the turbulent viscosity and/or the Reynolds stresses to zero
2. Setting the production terms in the turbulence transport equations to zero
3. Not using wall functions to compute wall shear stress or heat flux, but using the laminar flow methods instead

Calculating Transition Boundary Distance

Transition boundary distance is calculated in exactly the same manner as the wall distance, with the proviso that only walls that are transition boundaries are included in the calculation.

What Is the Transition Boundary Distance Solver?

This solver controls the transition boundary distance solution in all the continua for which transition boundary distance is calculated. It becomes available when the Transition Boundary Distance model is activated.

Information on the progress of this calculation will be shown in the progress and status bars beneath the STAR-CCM+ explorer window.

Transition Boundary Distance Solver Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Ticked	The solver will not update the solution.
	Cleared	The solver will update the solution.

Selecting the Turbulence Suppression and Transition Boundary Distance Models

This section describes how to select the Turbulence Suppression model. Two methods are described:

- Using auto-select
- Without using auto-select

Using Auto-Select

This section describes how to select a model using automatic selection of recommended models. This is the default mode of selection, and is indicated by the Auto-select recommended models checkbox being ticked.

To select the Turbulence Suppression model, first select the turbulence model. Then in the *Optional Models* group box, select the Turbulence Suppression radio button.



The Transition Boundary Distance model will be automatically selected.



Without Using Auto-Select

This section describes how to select the Turbulence Suppression model without using automatic selection of recommended models. This means that the Transition Boundary Distance model will not be automatically selected.

To select the Turbulence Suppression model, first select the turbulence model. Then in the *Optional Models* group box, select the Turbulence Suppression radio button.



Even if the Auto-select recommended models checkbox is cleared in the *Model Selection* dialog only one radio button, Transition Boundary Distance, will be presented in the *Transition Boundary Distance* group box. Therefore the selection procedure is identical to using auto-select.



Using the Transition Boundary Distance Model

The transition boundary distance is computed during solution initialization, and updated only after the mesh or boundary conditions have changed. There are no model properties to change, and there are no solver parameters to be set.

Running a Simulation with the Transition Model

The activation of the Transition model and the selection of the turbulence suppression option on a wall boundary are clearly essential steps. However, several other issues are important. These include:

- The choice of turbulence model and wall treatment
- The near-wall mesh resolution

- The inflow or free-stream turbulence boundary conditions

Choice of Turbulence Model and Wall Treatment

There is no restriction on the choice of turbulence model when the Transition model is activated. However, it makes the most sense to use models that are designed to resolve the laminar sublayer. These include (but are not necessarily limited to):

- The Standard Spalart-Allmaras model
- The Realizable Two-Layer K-Epsilon model
- The Lien low-Reynolds number model
- The Abe-Kondoh-Nagano low-Reynolds number model
- The Standard K-Omega model
- The SST K-Omega model
- The Linear Pressure Strain Two-Layer model

In each case the low-y^+ or all-y^+ wall treatment model should be used.

Near-Wall Mesh Resolution

Since the wall shear stress is computed using laminar flow relationships rather than wall functions adjacent to a transition boundary, the mesh should be sufficiently fine for acceptable accuracy. A sound approach would be to ensure that the mesh is adequately refined in the turbulent region and that the near-wall mesh spacing is not discontinuous.

Inflow or Free-Stream Turbulence Boundary Conditions

Suppressing turbulence in one region of the flow does not guarantee that the flow will become rapidly turbulent where it is not suppressed. The onset of transition, though not accurately predicted by turbulence models, will be affected by the Reynolds number and the free-stream turbulence values. Some low-Reynolds number turbulence models can be quite sensitive to the free-stream turbulence values, so some artful tuning of the inflow turbulence conditions might be required to get the desired transition effect.

Transition Field Functions Reference

The following primitive field functions are made available to the simulation when the transition model is used. These objects all have the same base properties.

Transition Boundary Distance	The transition boundary distance, defined as the distance from the field point to the nearest point on a transition boundary.
Transition Boundary Distance Factor	An integer number that is zero when the transition boundary distance is less than the wall distance, and unity when it is greater.

Using the Gamma ReTheta Transition Model

This section explains the following:

- What is the Gamma ReTheta transition model?
- What is the Free-Stream Edge Definition?
- How should inflow turbulence be specified?
- What are the mesh restrictions?
- What impact will the model have on convergence?
- How can field functions be used to define the correlations?
- Working with the Gamma ReTheta transition solver
- Formulation

Also included is reference information on related field functions.

What Is the Gamma ReTheta Transition Model?

Correlation-based transition models, such as that by Abu-Ghannam and Shaw [118], required the evaluation of non-local quantities such as momentum thickness Reynolds number. These quantities cannot be evaluated on unstructured meshes, which precludes their practical application in modern unstructured CFD codes.

The Gamma ReTheta transition model [119], [120] is a correlation-based transition model that has been specifically formulated for unstructured CFD codes. The evaluation of momentum thickness Reynolds number is avoided by relating this quantity to vorticity-based Reynolds number. In addition, a correlation for transition onset momentum thickness Reynolds number defined in the free stream is propagated into the boundary layer by a transport equation. An intermittency transport equation is further used in such a way that the source terms attempt to mimic the behavior of algebraic engineering correlations.

The Gamma ReTheta transition model, as originally published, is incomplete, since two critical correlations were claimed to be proprietary and hence omitted. One justification for such an omission is that the model provides a “framework” for users to implement their own correlations. Therefore, the model’s implementation is done in such a way as to allow the custom specification of correlations via field functions.

Recognizing, however, that the process of calibrating the needed correlations is complex and time consuming, the required correlations have been carefully calibrated and are provided by default within the model.

Properties

Convection	For guidance on selecting a convection scheme, see Convection Term, p. 6343.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.
Freestream Edge	Field Function that defines the free-stream edge	

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques used in STAR-CCM+, we recommend that you not make any changes within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance. Tampering with them may diminish the effectiveness of the model.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Correlation Method	Choose how the correlations for $Re_{\theta c}$ and F_{length} are to be specified	
	Suluksna-Juntasaro	Use the internally-specified correlations (see Eqns. (818), (837) and (838))
	User Field Function	Use correlations defined by user specified field function.

Sigma_f	The coefficient σ_f, see Eqn. (787).
Sigma_w	The coefficient $\sigma_{\theta t}$, see Eqn. (801).
ca1	The coefficient c_{a1}, see Eqn. (788).
ca2	The coefficient c_{a2}, see Eqn. (797).
ce1	The coefficient c_{e1}, see Eqn. (788).
ce2	The coefficient c_{e2}, see Eqn. (797).
cTheta_t	The coefficient $c_{\theta t}$, see Eqn. (802).
s1	The coefficient s_1, see Eqn. (815).
Intermittency Minimum	The minimum value that the transported variable γ is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.
ReThetaTMinimum	The minimum value that the transported variable $\overline{Re_{\theta t}}$ is permitted to have. The default value is 20.

What Is the Free-Stream Edge Definition?

Although the Gamma ReTheta transition model purports to be defined in terms of local variables, it is really semi-local since it depends on the definition of the $Re_{\theta t}$ through a correlation defined in the free stream. This free-stream value is transported through the boundary layer through the variable $\overline{Re_{\theta t}}$. In order to specify the value of $Re_{\theta t}$ using a correlation, the location of the free-stream edge must be defined.

Since there is no completely general way to define the free stream, it is up to the user to create a user field function that suitably defines the free stream in such a way that the field function takes the value of 1 in the free stream and 0 inside the boundary layer. A good starting point for this definition is in terms of Wall Distance. An example would be:

```
$WallDistance > 0.005?1:0
```

where it is estimated that the boundary layer is everywhere thinner than 5mm. The user field function, once defined and named, must then be assigned to Freestream Edge in the properties of the Gamma ReTheta transition model.

How Should Inflow Turbulence be Specified?

The Gamma ReTheta transition model attempts to reflect properly the effect of the free-stream turbulence levels on the transition location. Therefore, proper attention must be paid to the specification of inflow values of the turbulence variables so that realistic turbulence intensities are encountered in the boundary layers in question.

It is preferable that experimental data be provided for the behavior of the free-stream turbulence. Having three or more data points effectively allows the dissipation rate to be determined, and the inflow turbulence values can be determined by a regression analysis of the data.

The free-stream turbulence decay laws, derived from the SST K-Omega model transport equation and using the default coefficients, are as follows:

$$k = k_{inlet}\left[1 + 0.0828\frac{x}{U_\infty}\omega_{inlet}\right]^{-1.087} \tag{781}$$

$$\omega = \omega_{inlet}\left[1 + 0.0828\frac{x}{U_\infty}\omega_{inlet}\right]^{-1} \tag{782}$$

where x eighth streamwise distance and U_∞ is the (assumed constant) free-stream velocity.

Using the following relations for turbulence intensity I and turbulent viscosity ratio ($TVR = \mu_t/\mu$):

$$I = \frac{\sqrt{(2/3)k}}{U} \tag{783}$$

$$TVR = \frac{k}{\omega\mu} \tag{784}$$

one can transform Eqn. (781) and (782) into the following:

$$I = I_{inlet}\left[1 + \frac{0.1242 x \rho U_{\infty} I_{inlet}^{2}}{\mu\ TVR_{inlet}}\right]^{-0.5435} \tag{785}$$

$$TVR = TVR_{inlet}\left[1 + \frac{0.1242 x \rho U_{\infty} I_{inlet}^{2}}{\mu\ TVR_{inlet}}\right]^{-0.087} \tag{786}$$

Care must be taken with the specification of turbulent viscosity ratio at the inlet because large values of effective viscosity will cause excessive diffusion. The exponent in Eqn. (786) is small, so that the turbulent viscosity does not decay significantly, allowing the inlet values to propagate far downstream. Unfortunately, the smaller the value of the turbulent viscosity ratio, the faster the turbulence will decay. In the absence of measured data, the proper turbulent inflow boundary conditions become a delicate balance between an acceptable turbulent decay rate and an acceptably low turbulent viscosity ratio.

Since the turbulent viscosity ratio is, in effect, a turbulent Reynolds number, it is impossible to specify what an acceptable value is for any given flow. (Higher bulk Reynolds numbers would suggest higher turbulent Reynolds numbers, and hence higher levels of acceptable turbulent viscosity ratio.) However, a rough guideline is that the free-stream turbulent viscosity ratio should be significantly lower than the turbulent viscosity that occurs within the boundary layer. For natural transition problems and lower free-stream turbulence bypass transition, a desirable goal would be to have a TVR between 1 and 10. This will prevent the high levels of effective viscosity from contaminating the laminar boundary layer and affecting the laminar skin friction prediction.

In some situations, such as wind tunnels with very low background turbulence, it might be a challenge to obtain a suitable compromise between these two variables. This issue is often exacerbated when a far-field boundary is specified a large distance away from the body.

What Are the Mesh Restrictions?

It is clear from validations that the Gamma ReTheta transition model requires a high-quality, refined low-Reynolds number mesh. The wall-normal mesh spacing should be fine enough that $0.1 \le y^{+} \le 1$, and special attention should be paid to the streamwise mesh spacing near the transition region. Either the all-y^{+} or low-y^{+} wall treatment formulations should be used.

If transition is induced by laminar separation, then the mesh should be fine enough to resolve the curvature of the body where the separation exists, and to capture the laminar separation bubble.

What Impact Will the Model Have on Convergence?

Apart from the fact that the Gamma ReTheta transition model requires the solution of two extra transport equations, incurring an additional computational overhead as a result, an impact on solution covertness must be expected. It has been noted that the interaction between flow, turbulence and the transition model can increase the number of iterations to convergence significantly. Furthermore, the intermittency transport equation residuals are naturally noisy. This is attributed to the strong dependence of the source terms on higher powers of the mean flow strain rates, and the fact the source terms are not differentiable functions.

How Can Field Functions be Used to Define the Correlations?

A property termed Correlation Method is provided in the Gamma ReTheta transition model which allows the user to select the User Field Function method of specifying the calibrations used in the model. Once this option is selected, a node appears in under the model. The following properties appear under the node:

Properties

Transition onset Re Correlation		
	Modified	Selects the $Re_{\theta t}$ correlation of Eqn. (818).
	Modified NZPG	Selects the $Re_{\theta t}$ correlation of Eqn. (834).
	Original	Selects the $Re_{\theta t}$ correlation of Eqn. (817).
	Original NZPG	Selects the $Re_{\theta t}$ correlation of Eqn. (829).
	Abu-Ghannam & Shaw	Selects the $Re_{\theta t}$ correlation of Eqn. (816).
	Abu-Ghannam & Shaw NZPG	Selects the $Re_{\theta t}$ correlation of Eqn. (826).
	User Field Function	Allows specification of $Re_{\theta t}$ through a user field function.
F_length correlation		Field Function that defines the F_{length} correlation.
Re_crit correlation		Field Function that defines the $Re_{\theta c}$ correlation.
Re_theta_t correlation		Field Function that defines the $Re_{\theta t}$ correlation.

To emulate the internal (Suluksna-Juntasaro) calibration, the Modified option can be selected for the Transition onset Re Correlation and the field functions for F_length correlation and Re_crit correlation, respectively, would be specified as follows:

```
min(300,exp(7.168 - 0.01173*$TransitionReTheta) + 0.5)
min($TransitionReTheta, $TransitionReTheta*0.615 + 61.5)
```

If the User Field Function option is to be used to specify the $Re_{\theta t}$ correlation, the following field function could be used to emulate the internal (Modified option) implementation.

```
(100*$TurbulenceIntensity > 1.3) ?
331.5*pow(100*$TurbulenceIntensity - 0.5658, -0.671) :
(1173.51 -
589.428*100*$TurbulenceIntensity+0.2196/(10000*$Turbulenc
eIntensity*$TurbulenceIntensity))
```

Working with the GammaReTheta Transition Solver

This section describes the Gamma ReTheta transition solver, which is the solver used by the Gamma ReTheta transition model.

It also lists the residual monitors that are added to the simulation by these solvers.

What Is the Gamma ReTheta Transition Solver?

This solver controls the solution of the Gamma ReTheta transition model in all the continua for which the model is activated. For each transported variable, γ and $\overline{Re_{\theta t}}$, the basic steps involved in the solution update are as follows:

1. Update boundary conditions.
2. Compute the reconstruction gradients and cell gradients.
3. Create the linear system of Eqn. (2591) using the discretization methods described in the K-Omega model formulation.
4. Compute the residual sum for the purpose of monitoring convergence:

$$R = \sum_{\text{cells}} |r|$$

5. Solve the linear system.
6. Update the transported variable field.

Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.8. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.
Freestream Update Frequency	This parameter governs the frequency at which the correlation for Re_{θ_t} will be updated in the free stream. The default is 1, which manes that the update occurs every iteration. Increasing this parameter to a larger interval, say 10, might improve parallel performance in certain situations. This is because the free-stream values are stored in a KD tree that is broadcast to each parallel node, incurring a communication overhead.

Expert Properties

The properties in the *Expert* category are designed to be used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Ticked	The solver will not update the solution.
	Cleared	The solver will update the solution.
Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients and the cell gradients. These quantities will become available as field functions during subsequent iterations.	
	Ticked	Retain temporary storage.
	Cleared	Do not retain temporary storage.

Reconstruction Frozen

Sets whether or not reconstruction gradients are updated with each iteration.

Ticked: Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.

Cleared: Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) will be first-order estimates.

Ticked: Set reconstruction gradients to zero at the next iteration. Gradients will remain zeroed until you clear the checkbox again.

Cleared: Do not reset reconstruction gradients. The will be recomputed at the next iteration if they had been zeroed.

Using Residual Monitors

The Gamma ReTheta transition solver adds the following residual monitors to the top level Monitors node of a simulation tree.

Intermittency: Intermittency.

ReTheta_t: Transported transition momentum thickness Reynolds number.

The properties of all residual monitors are the same.

Gamma ReTheta Transition Formulation

In this section the formulation of the Gamma ReTheta model is presented:

- Intermittency Transport Equation
- Transition Momentum Thickness Reynolds Number Transport Equation

- Separation-Induced Transition Modification
- Correlations for ReTheta_t
- Correlations for ReTheta_c and F_length

Nomenclature and a bibliography are also provided.

Intermittency Transport Equation

The intermittency equation is formulated as:

$$\frac{d}{dt}\int_V \rho\gamma dV + \int_A \rho\gamma(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \left(\mu + \frac{\mu_t}{\sigma_f}\right)\nabla\gamma\cdot d\mathbf{a} + \int_V [P_\gamma - E_\lambda]\ dV \tag{787}$$

Production Term

The transition source term is defined as:

$$P_\gamma = F_{length} c_{a1} \rho S[\gamma F_{onset}]^{1/2}(1 - c_{e1}\gamma) \tag{788}$$

where:

$$F_{onset1} = \frac{Re_v}{2.193 Re_{\theta c}} \tag{789}$$

$$F_{onset2} = \min[\max((F_{onset1}^4, F_{onset1}), 2)] \tag{790}$$

$$F_{onset3} = \max[1 - (Re_T/2.5)^3, 0] \tag{791}$$

$$F_{onset} = \max(F_{onset2} - F_{onset3}, 0) \tag{792}$$

Re_v is the strain-rate Reynolds number defined as:

$$Re_v = \frac{\rho\ d^2 S}{\mu} \tag{793}$$

and S is the modulus of the mean strain rate tensor:

$$S = |\mathbf{S}| = \sqrt{2\mathbf{S}:\mathbf{S}^T} = \sqrt{2\mathbf{S}:\mathbf{S}} \tag{794}$$

where:

$$\mathbf{S} = \frac{1}{2}(\nabla\mathbf{v} + \nabla\mathbf{v}^T) \tag{795}$$

The turbulent Reynolds number R_T is defined as:

$$Re_T = \frac{k}{\nu\omega} \tag{796}$$

Note that the correlations for $Re_{\theta c}$ and F_{length} are not specified by Langtry and Menter.

Destruction Term

The destruction/relaminarization source is defined as:

$$E_\lambda = C_{a2}\rho W\gamma F_{turb}(c_{e2}\gamma - 1) \tag{797}$$

where:

$$F_{turb} = \exp[-(Re_T/4)^4] \tag{798}$$

and:

$$W = \sqrt{2\mathbf{W}:\mathbf{W}^T} \tag{799}$$

where the vorticity tensor $\mathbf{W}$ is defined:

$$\mathbf{W} = \frac{1}{2}(\nabla\mathbf{v} - \nabla\mathbf{v}^T) \tag{800}$$

Model coefficients

The model coefficients for the intermittency transport equation are defined as:

$$C_{e1} = 1.0 \ , \ C_{a1} = 2.0 \ , \ C_{e2} = 50 \ , \ C_{a2} = 0.06 \ , \ \sigma_f = 1.0$$

Transition Momentum Thickness Reynolds Number Transport Equation

The transition momentum thickness Reynolds number transport equation is formulated as:

$$\frac{d}{dt}\int_V \rho\overline{Re_{\theta t}}dV + \int_A \rho\overline{Re_{\theta t}}(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A \sigma_{\theta t}(\mu + \mu_t)\nabla\overline{Re_{\theta t}}\cdot d\mathbf{a} + \int_V P_{\theta t}\, dV \tag{801}$$

Source Term

The source term is defined as:

$$P_{\theta t} = C_{\theta t}\frac{\rho}{t}(Re_{\theta t} - \overline{Re_{\theta t}})(1 - F_{\theta t}) \tag{802}$$

where:

$$t = \frac{500\nu}{U^2} \tag{803}$$

and:

$$F_{\theta t} = \min\left\{\max\left[F_{wake}\exp\left(-\left(\frac{d}{\bar{\delta}}\right)^4\right), 1 - \left(\frac{C_{e2}\gamma - 1}{C_{e2} - 1}\right)^2\right], 1\right\} \tag{804}$$

Using:

$$\delta = 50W\frac{d}{U}\delta_{BL} \tag{805}$$

$$\delta_{BL} = \frac{15}{2}\theta_{BL} \tag{806}$$

$$\theta_{BL} = \frac{\overline{Re_{\theta t}}\nu}{U} \tag{807}$$

then:

$$\frac{d}{\bar{\delta}} = \frac{U^2}{375W\nu\overline{Re_{\theta t}}} \tag{808}$$

And F_{wake} is given by:

$$Re_{\omega} = \frac{\omega d^2}{\nu} \tag{809}$$

$$F_{wake} = \exp[-(Re_{\omega}/10^5)^2] \tag{810}$$

The model coefficients are defined as:

$$C_{\theta t} = 0.03 \ , \ \sigma_{\theta t} = 2.0 \tag{811}$$

Separation-Induced Transition Modification

The modification for separation-induced transition is:

$$\gamma_{sep} = \min\left[s_1\max\left(0, \frac{Re_{\nu}}{3.235Re_{\theta c}} - 1\right)F_{reattach}, 2\right] \tag{812}$$

where:

$$F_{reattach} = \exp\left[\left(\frac{-Re_T}{15}\right)^4\right] \tag{813}$$

$$\gamma_{eff} = \max(\gamma, \gamma_{sep}) \tag{814}$$

and the coefficient s_1 has the default value:

$$s_1 = 2 \tag{815}$$

Correlations for ReTheta_t

A well-established correlation for $Re_{\theta t}$ was developed by Abu-Ghannam and Shaw [118]. In the case of a zero pressure gradient, this correlation is:

$$Re_{\theta t} = 163 + \exp(6.91 - Tu) \tag{816}$$

where Tu is the turbulence intensity expressed as a percentage.

An alternative expression, developed for the Gamma ReTheta transition model by Menter *et al.* [120] is:

$$Re_{\theta t} = 803.73(Tu + 0.6067)^{-1.027} \tag{817}$$

and this was later revised by Langtry [119] to "improve the results at lower freestream turbulence intensities" as follows:

$$Re_{\theta t} = \begin{cases} (1173.51 - 589.428Tu + 0.2196/Tu^2) & ; \ Tu \le 1.3 \\ 331.5(Tu - 0.5658)^{-0.671} & ; \ Tu > 1.3 \end{cases} \tag{818}$$

The three zero-pressure gradient correlations are plotted in the following figure. All three are provided in STAR-CCM+ to allow the user flexibility to customize the correlations, but Eqn. (818) is the default correlation used in the internal "Suluksna-Juntasaro" calibration.



$Re_{\theta t}$ vs. Free-Stream Turbulence Intensity

The three correlations presented for zero pressure gradient in Eqns. (816) through (818) were originally presented in a form that include the effects of a pressure gradient. There is some doubt that the pressure gradient effects are indeed required in the $Re_{\theta t}$ correlation [121], since they arguably introduce a double accounting of the pressure gradient effects, but they are nevertheless optionally included in STAR-CCM+. In the presence of a pressure gradient, the correlations require the computation of the Thwaites parameter, λ_θ, and in the case of the Menter *et al.* [120] correlation, the acceleration parameter, K.

To avoid computing an integral parameter such as momentum thickness, the momentum thickness Reynolds number is estimated as:

$$Re_\theta = \frac{Re_\nu}{2.193} \tag{819}$$

Thwaite's parameter, λ_θ, is defined:

$$\lambda_\theta = \frac{\theta^2}{\nu}\frac{dU}{ds} \tag{820}$$

with:

$$\theta = \frac{\nu}{U} R_{\theta t} \tag{821}$$

The parameter K is defined as:

$$K = \frac{\nu}{U^2} \frac{dU}{ds} \tag{822}$$

The turbulence intensity is obtained from the turbulent kinetic energy as follows:

$$Tu = 100 \frac{\sqrt{2k/3}}{U} \tag{823}$$

where:

$$U = |\mathbf{v}| \tag{824}$$

The streamwise velocity gradient is calculated as:

$$\frac{dU}{ds} = \frac{(\nabla \mathbf{v}^T \cdot \mathbf{v}) \cdot \mathbf{v}}{\mathbf{v} \cdot \mathbf{v}} \tag{825}$$

Since the momentum thickness θ must be derived from the momentum thickness Reynolds number Re_θ the correlation becomes an implicit function that is solved iteratively.

The non-zero pressure gradient (NZPG) correlation of Abu-Ghannam and Shaw [118] is:

$$Re_{\theta t} = 163 + \exp\left[F(\lambda_\theta)\left(1 - \frac{Tu}{6.91}\right)\right] \tag{826}$$

where:

$$F(\lambda_\theta) = \begin{cases} 6.91 + 12.75\lambda_\theta + 63.64\lambda_\theta^2 \; ; \; \lambda_\theta \le 0 \\ 6.91 + 2.58\lambda_\theta - 12.27\lambda_\theta^2 \; ; \; \lambda_\theta > 0 \end{cases} \tag{827}$$

For numerical robustness, the following limits are applied:

$$\begin{gathered} -0.1 \le \lambda_\theta \le 0.1 \\ Re_{\theta t} \ge 20 \end{gathered} \tag{828}$$

The NZPG correlation of Menter *et al.* [120] is:

$$Re_{\theta t} = 803.73(Tu + 0.6067)^{-1.027} F(\lambda_\theta, K) \tag{829}$$

$$F(\lambda_\theta, K) = \begin{cases} 1 + F_\lambda \exp(-Tu/3) & ; \ \lambda_\theta \le 0 \\ 1 + F_K[1 - \exp(-Tu/1.5)] + 0.556[1 - e^{-23.9\lambda_\theta}]e^{-Tu/1.5} & ; \ \lambda_\theta > 0 \end{cases} \tag{830}$$

where:

$$F_\lambda = 10.32\lambda_\theta + 89.47\lambda_\theta^2 + 265.51\lambda_\theta^3 \tag{831}$$

$$F_K = [0.0962(K \cdot 10^6) + 0.148(K \cdot 10^6)^2 + 0.0141(K \cdot 10^6)^3] \tag{832}$$

For numerical robustness, the following limits are applied:

$$\begin{gathered} -0.1 \le \lambda_\theta \le 0.1 \\ -3x10^{-6} \le K \le 3x10^{-6} \\ Re_{\theta t} \ge 20 \end{gathered} \tag{833}$$

The NZPG correlation of Langtry [119] is:

$$Re_{\theta t} = \begin{cases} (1173.51 - 589.428Tu + 0.2196/Tu^2)F(\lambda_\theta) & ; \ Tu \le 1.3 \\ 331.5(Tu - 0.5658)^{-0.671} F(\lambda_\theta) & ; \ Tu > 1.3 \end{cases} \tag{834}$$

$$F(\lambda_\theta) = \begin{cases} 1 + 12.986\lambda_\theta + 123.66\lambda_\theta^2 + 405.689\lambda_\theta^3 & ; \ \lambda_\theta \le 0 \\ 1 + 0.275[1 - \exp(-35\lambda_\theta)]\exp(-2Tu) & ; \ \lambda_\theta > 0 \end{cases} \tag{835}$$

For numerical robustness, the following limits are applied:

$$\begin{gathered} -0.1 \le \lambda_\theta \le 0.1 \\ Tu \ge 0.027 \\ Re_{\theta t} \ge 20 \end{gathered} \tag{836}$$

A comparison of the three NZPG correlations is shown in the following figure for a range of free-stream turbulence intensities.



$Re_{\theta t}$ vs. λ_θ for various values of Free-Stream Turbulence Intensity

Correlations for ReTheta_c and F_length

The correlations for $Re_{\theta c}$ and F_{length} have been calibrated according to the forms suggested by Suluksna *et al.* [121]. The correlations are as follows:

$$Re_{\theta c} = \min(0.615\overline{Re_{\theta t}} + 61.5, \overline{Re_{\theta t}}) \tag{837}$$

$$F_{length} = \min[0.5 + \exp(7.168 - 0.01173\overline{Re_{\theta t}}), 300] \tag{838}$$

When combined with the correlation for $Re_{\theta t}$ in Eqn. (818), these correlations enable the Gamma ReTheta transition model to mimic the bypass transition cases T3A, T3B, T3AM [122] and the natural transition

cases TSK [123], as shown in the following figure.



Skin friction coefficient predictions for calibrated correlations

Gamma ReTheta Transition Nomenclature

The following is a guide to the terms used in this section.

Scalars

d	distance to the nearest wall
k	turbulent kinetic energy
t	time, time scale
P_γ	intermittency production
E_γ	intermittency destruction
F_{length}	function governing transition "length"
I	turbulence intensity (expressed as a fraction)
Tu	turbulence intensity (expressed as a percentage)
Re_θ	momentum thickness Reynolds number
$Re_{\theta c}$	critical momentum thickness Reynolds number
$Re_{\theta t}$	transition momentum thickness Reynolds number
$\overline{Re_{\theta t}}$	transported transition momentum thickness Reynolds number
Re_T	turbulent Reynolds number
Re_ν	strain-rate Reynolds number
Re_ω	dissipation Reynolds number

K	acceleration parameter
S	strain rate tensor modulus
W	vorticity tensor modulus
U	velocity magnitude
V	cell volume
γ	intermittency
γ_{eff}	effective intermittency
ρ	density
λ_θ	Thwaite's pressure gradient parameter
θ	boundary layer momentum thickness
μ	dynamic viscosity
μ_t	turbulent viscosity
ν	kinematic viscosity
σ_γ	turbulent Schmidt number
$\sigma_{\theta t}$	inverse turbulent Schmidt number

Tensors or Column Vectors

S	strain rate tensor
W	vorticity tensor

Cartesian Vectors

a	face area vector
v	velocity
$\mathbf{v}_g$	grid velocity

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Divergence of vector quantity:

$$\nabla\cdot\mathbf{v} = \frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

- Scalar (inner, double-dot) product of two tensors, whose result is a scalar:

$$\mathbf{S}:\mathbf{T} = S_{ij}T_{ji}$$

Gamma ReTheta Transition Bibliography

[118] Abu-Ghannam, B. J. and Shaw, R. "Natural Transition of Boundary Layers --- The Effects of Turbulence, Pressure Gradient, and Flow History," *J. Mech. Eng. Sci.*, Vol. 22, No. 5, pp213-228, 1980.

[119] Langtry, R. B. "A Correlation-Based Transition Model Using Local Variables for Unstructured Parallelized CFD Codes," Doctoral Thesis, University of Stuttgart, 2006.

[120] Menter, F. R., Langtry, R. B., Likki, S. R., Suzen, Y. B., Huang, P. G. and Volker, S. 2004. "A Correlation-Based Transition Model Using Local Variables Part 1 --- Model Formulation," *Proc. ASME Turbo Expo*, June 14-17, Vienna, Austria.

[121] Suluksna, K., Dechaumphai, P. and Juntasaro, E. "Correlations for Modeling Transitional Boundary Layers under Influences of Freestream Turbulence and Pressure Gradient," *Int. J. Heat and Fluid Flow*, in press 2008.

[122] ERCOFTAC (European Research Community on Flow, Turbulence, Combustion) Nexus. [online database], URL: http://ercoftac.mech.surrey.ac.uk/.

[123] Schubauer, G.B. and Klebanoff, P.S. "Contributions on the Mechanics of Boundary Layer Transition," *NACA Technical Note 3489*, 1955.

Gamma ReTheta Transition Field Functions Reference

The following primitive field functions are made available to the simulation when the Gamma ReTheta transition model is used. These objects all have the same base properties.

Intermittency	The transported variable γ, see Eqn. (787).
Transition Momentum Thickness Reynolds Number	The transported variable $\overline{Re_{\theta t}}$, see Eqn. (801).
Effective Intermittency	The transported variable γ_{eff}, see Eqn. (814).
Turbulence Intensity	The turbulence intensity, defined as $I = \sqrt{2/3k}/U$
Specific Dissipation Rate Reynolds Number	The specific dissipation rate Reynolds number, Re_{ω}, see Eqn. (809).
ReThetaCrit Correlation	The specified correlation for $Re_{\theta c}$, see Eqn. (837).
Flength Correlation	The specified correlation for F_{length}, see Eqn. (838).

Modeling Turbulent Heat Flux

The classical definition of heat flux in the energy equation is based on the Boussinesq approximation, where the heat flux is assumed to be proportional to the turbulent eddy-viscosity. This assumption fails either when buoyancy forces are dominant, or at locations very near the wall. A remedy proposed by Kenjeres et al. [124] is to replace the Boussinesq approximation by an algebraic formulation for the turbulent heat flux itself. This formulation is a function of the Reynolds-stress anisotropy and the temperature variance, for which an additional transport equation is solved. The performance of the algebraic heat flux (or Temperature Flux) model is strongly linked to the correct approximation of the near-wall turbulent behavior, thus requiring a low-Reynolds number turbulence model. The present model combines the heat flux model of Kenjeres et al. with the Standard Low-Reynolds Number K-Epsilon model.

The Temperature Flux model is represented in the STAR-CCM+ physics continuum by the Temperature Flux Model node. To access the properties of the model, select this node.



This section describes the following:

- Selecting the Temperature Flux model
- Setting Temperature Flux solver parameters

- Working with temperature variance

The mathematical formulation and field function references for the Temperature Flux model are also available.

Temperature Flux Model Properties

Convection	Selects the convection scheme that is used.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

Temperature Flux Model Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model, and the discretization techniques that are used in STAR-CCM+, do not change any properties within the *Expert* category. The values in that category have been optimized for accuracy and performance.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
R	The turbulent-to-thermal time-scale ratio R, see Eqns. (845) and (846).	
Sigma_Theta2	The coefficient σ_{θ^2}, see Eqns. (843) and (846).	
Ctu0	The coefficient C_{tu_0}, see Eqns. (840) and (846).	

Ctu1	The coefficient C_{tu_1}, see Eqns. (840) and (846).
Ctu2	The coefficient C_{tu_2}, see Eqns. (840) and (846).
Ctu3	The coefficient C_{tu_3}, see Eqns. (840) and (846).
Ctu4	The coefficient C_{tu_4}, see Eqns. (840) and (846).
Theta2 Minimum	The minimum value that the transported variable $\overline{\theta^2}$ is permitted to have. An appropriate value is a small number that is greater than the floating point minimum of the computer.

Selecting the Temperature Flux Model

To activate the Temperature Flux model:

- Select Segregated Flow or Coupled Flow from the *Flow* group.
- Select Constant Density from the *Equation of State* group.
- Select Segregated Fluid Temperature from the *Optional Models* group.
- Select Turbulent from the *Viscous Regime* group.

  The Reynolds-Averaged Navier-Stokes option is selected automatically.
- Deactivate the Auto-select recommended physics models option in the dialog.

  This step is necessary because the Standard Low-Reynolds Number K-Epsilon model is required instead of the realizable two-layer K-Epsilon model, which is the STAR-CCM+ default.
- Select K-Epsilon Turbulence in the *Reynolds-Averaged Turbulence* group.
- Select Standard K-Epsilon Low-Re in the *K-Epsilon Turbulence* group.
- Select All y+ Wall Treatment in the *K-Epsilon Damped Low Re Wall Treatment* group.
- Select Gravity from the *Optional Models* group.
- Select Boussinesq Model from the *Optional Models* group.

As a result of these selections, the Temperature Flux Model option becomes available in the *Optional Models* group.



Setting Temperature Flux Solver Parameters

This solver controls the solution of the Temperature Flux model in all the continua for which this model is activated. For each transported variable, the solution update includes:

- Boundary conditions
- Reconstruction gradients and cell gradients
- The residual sum for monitoring convergence

The Temperature Flux Solver node represents the Temperature Flux solver in the object tree. To access the properties of the solver, select this node.



Temperature Flux Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed velocity supplants the old one. The default value is 0.7. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.

For information about the expert properties of this solver, see Expert Properties, p. 6385.

Working with Temperature Variance

The Temperature Variance node appears as an initial condition and as a boundary condition (for flow boundaries) when the Temperature Flux model is activated. It is a direct specification of $\overline{\theta^2}$.



Properties

This value is input as a scalar profile.

Temperature Flux Formulation

The following is a mathematical overview of the Temperature Flux model. This section also includes nomenclature and a bibliography.

The classical definition of the heat flux (Eqns. (188) and (189)) in the energy equation is based on a Boussinesq approximation:

$$\dot{\mathbf{q}}'' = -\left(k + \frac{\mu_t C_p}{\sigma_t}\right)\nabla T \tag{839}$$

where σ_t is a turbulent Prandtl Number, usually taken as a constant equal to 0.9. The Temperature Flux model is based on an algebraic formulation for the turbulent heat flux [124]:

$$\overline{\mathbf{v}\theta} = -C_{tu_0}\frac{k}{\varepsilon}\left(C_{tu_1}\mathbf{R}\cdot\nabla T + C_{tu_2}\nabla\mathbf{v}\cdot\overline{\mathbf{v}\theta} + C_{tu_3}\beta\overline{\theta^2}\mathbf{g}\right) + C_{tu_4}\mathbf{A}\cdot\overline{\mathbf{v}\theta} \tag{840}$$

where $\overline{\mathbf{v}\theta}$ replaces the Boussinesq approximation in the definition of the total heat flux:

$$\dot{\mathbf{q}}'' = -\kappa\nabla T - \rho C_p\overline{\mathbf{v}\theta} \tag{841}$$

The turbulent heat flux $\overline{\mathbf{v}\theta}$ is also used for the buoyancy-related source terms G_b in the transport equation for k and ε (Eqn. (530)):

$$G_b = -\rho\beta\,\mathbf{g}\cdot\overline{\mathbf{v}\theta} \tag{842}$$

In Eqn. (840), the remaining unknown, the temperature variance $\overline{\theta^2}$, is computed by solving an additional transport equation:

$$\frac{d}{dt}\int_V \rho\overline{\theta^2}dV + \int_A \rho\overline{\theta^2}(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \int_A\left(\mu + \frac{\mu_t}{\sigma_{\theta^2}}\right)\nabla\overline{\theta^2}\cdot d\mathbf{a} + \int_V (G_\theta - \rho\varepsilon_\theta)dV \tag{843}$$

where the production of the temperature variance G_θ is:

$$G_\theta = -\rho\overline{\mathbf{v}\theta}\cdot\nabla T \tag{844}$$

The temperature variance dissipation rate ε_θ is obtained from the definition of the thermal time scale (equivalent for the temperature variance to the turbulence time scale for the K-Epsilon model):

$$\tau_\theta = \frac{\overline{\theta^2}}{2\varepsilon_\theta} \tag{845}$$

and assuming a constant turbulent-to-thermal time-scale ratio $R = \tau_\theta/\tau$:

$$\begin{gathered} C_{tu_0} = 0.12 \quad , \quad C_{tu_1} = 1 \quad , \quad C_{tu_2} = 0.6 \quad , \quad C_{tu_3} = 0.6 \\ C_{tu_4} = 1.5 \quad , \quad \sigma_{\theta^2} = 1 \quad , \quad R = 0.5 \end{gathered} \tag{846}$$

Temperature Flux Nomenclature

The following is a guide to the terms used in this section.

Scalar Quantities

Variable	Definition
C_p	Specific heat $[J/kg \cdot K]$
G_θ	Production of the temperature variance $[K^2/s]$
R	Time-scale ratio
T	Temperature $[K]$
$\overline{\theta^2}$	Temperature variance $[K^2]$
β	Thermal expansion coefficient $[1/K]$
ε_θ	Dissipation of the temperature variance $[K^2/s]$
κ	Thermal conductivity $[W/m \cdot K]$
μ	Dynamic viscosity $[Pa \cdot s]$
μ_t	Turbulent viscosity $[Pa \cdot s]$
ρ	Density $[kg/m^3]$
σ_t	Turbulent Prandtl number
τ	Turbulent time scale $[s]$
τ_θ	Thermal time scale $[s]$

Vectors

Variable	Definition
$\mathbf{g}$	Gravity $[m/s^2]$
$\mathbf{v}$	Velocity $[m/s]$
$\mathbf{v}_g$	Grid velocity $[m/s]$
$\overline{\mathbf{v}\theta}$	Turbulent heat flux $[mK/s]$
$\dot{\mathbf{q}}''$	Heat flux $[W/m^2]$

Tensors

Variable	Definition
$\mathbf{A}$	Reynolds-stress anisotropy
$\mathbf{R}$	Reynolds-stress $[m^2/s^2]$

Temperature Flux Bibliography

[124] Kenjeres, S., Gunarjo, S.B., and Hanjalic, K. 2005. "Contribution to elliptic relaxation modelling of turbulent natural and mixed convection", *Int. J. Heat and Fluid Flow*, **26**(4), pp. 569–586.

Temperature Flux Field Functions Reference

The following primitive field function is made available to the simulation when you activate the Temperature Flux model. This object has typical base properties.

Temperature Variance $\overline{\theta^2}$ in Eqns. (840), (843), and (845).

Shared Conditions and Values Reference

The following sections provide reference information for objects that are common to more than one turbulence model:

- Boundary and initial condition values
- Region conditions
- Region values

Boundary and Initial Condition Values Reference

This section details boundary and initial condition values that are available for multiple turbulence models in STAR-CCM+:

- Blended wall function
- Standard wall function
- Turbulence intensity
- Turbulent length scale
- Turbulent viscosity ratio
- Turbulent dissipation rate
- Turbulent kinetic energy
- Turbulent velocity scale

These values are elected by the boundary type, and/or by methods selected in corresponding nodes in the Conditions node.

Blended Wall Function

This boundary condition appears only in turbulent flow on walls that do not have a slip condition on them. It is added to the Values node when the all-y^+ wall treatment is selected.

Properties

E	The log law offset.
Kappa	The von Karman constant.

Standard Wall Function

This boundary condition appears only in turbulent flow on walls that do not have a slip condition on them. It is added to the Values node when the high-y^+ wall treatment is selected.

Properties

E	The log law offset.
Kappa	The von Karman constant.

Turbulence Intensity

This boundary value and initial condition value is added to the Values node by selecting `Intensity and Length Scale` or `Intensity and Viscosity Ratio` as the Method in the Conditions Turbulence Specification node. It is used in the calculation of various turbulence factors (depending on the turbulence model employed). Turbulence intensity is a measure of the root mean square of the local velocity fluctuations v' relative to the mean or reference velocity $\bar{\mathbf{v}}$:

$$I = \frac{v'}{\bar{\mathbf{v}}}$$

For a Reynolds-averaged turbulence model, the root mean square velocity fluctuations are computed as

$$v' = \sqrt{\frac{2}{3}k}$$

where k is the turbulent kinetic energy.

Properties

This value is input as a scalar profile.

Turbulent Length Scale

This boundary value and initial condition value is added to the Values node by selecting `Intensity and Length Scale` as the Method in the Conditions Turbulence Specification node. It is used in the calculation of various turbulence factors (depending on the turbulence model employed).

Properties

This value is input as a scalar profile.

Turbulent Viscosity Ratio

This boundary value and initial condition value is added to the Values node by selecting `Turbulent Viscosity Ratio` (in the case of the Spalart-Allmaras model) or `Intensity + Viscosity Ratio` (for other turbulence models) as the Method in the Conditions Turbulence Specification node. It is the ratio of turbulent to laminar viscosity (μ_t/μ).

Properties

This value is input as a scalar profile.

Turbulent Dissipation Rate

This boundary value and initial condition value is added to the Values node by selecting `K + Epsilon` or `RS + Epsilon` as the Method in the Conditions Turbulence Specification node (depending on the selected turbulence model) . It is a direct specification of ε.

Properties

This value is input as a scalar profile.

Turbulent Kinetic Energy

This boundary value and initial condition value is added to the Values node by selecting `K + Epsilon` or `K + Omega` as the Method in the Conditions Turbulence Specification node (depending on the turbulence model selected). It is a direct specification of k.

Properties

This value is input as a scalar profile.

Turbulent Velocity Scale

This initial condition value is used for estimating the initial turbulent kinetic energy. It is added to the Initial Conditions node by selecting `Intensity and Length Scale` or `Intensity and Viscosity Ratio` as the Method in the Initial Conditions Turbulence Specification node (depending on the turbulence model selected).

This value affects the initial value of k in the following equations:

- For K-Epsilon: Eqns. (623) and (624)
- For K-Omega: Eqns. (705) and (706)

For an initial value, use a representative velocity or a representative velocity scale. For example, for a pipe flow use the inlet velocity.

Properties

This value is input as a scalar profile.

Region Conditions Reference

This section details region conditions that are available for multiple turbulence models in STAR-CCM+:

- Turbulence source option

Region conditions are elected by the region type.

Turbulence Source Option

A turbulence source option can be specified, ambient or not allowed. When allowed, this option provides initial conditions for the turbulence of injected flow.

Properties

Turbulence source option	Defines whether or not you want to use a turbulence source term, as well as what type.	
	None	Does not allow a turbulence source term to be defined.

Specified — Allows a turbulence source term to be defined. A source node will be added to the Physics Values node, depending on the type of model used:

- If one of several Reynolds-Averaged Navier-Stokes Turbulence models is used, a Turbulent Dissipation Rate Source node and a Turbulent Kinetic Energy Source node will be activated.
- If the Spalart-Allmaras model is used, a Modified Diffusivity Source node will be activated.
- If the SST K-Omega model is used, a Specific Dissipation Rate Source and a Turbulent Kinetic Energy Source node will be activated.

Ambient — Adds source terms to the transport equations of K-Epsilon and K-Omega to counteract turbulence decay in external aero flows. An Ambient Turbulence Specification node will be added to the Physics Values node. This object lets you select an inflow boundary from which the turbulence source will be inferred.

Region Values Reference

This section details region values that are available for multiple turbulence models in STAR-CCM+. Note that most of these only apply to porous regions where turbulence quantities cannot be transported as usual.

- Ambient turbulence specification
- Turbulence intensity
- Turbulent length scale
- Turbulent viscosity ratio
- Turbulent dissipation rate
- Turbulent dissipation rate source
- Turbulent kinetic energy
- Turbulent kinetic energy source

These values are selected by the region type, and/or by methods selected in corresponding nodes in the Conditions node.

Ambient Turbulence Specification

This region value is added to the Values node by selecting `Ambient` in the properties of the Conditions Turbulence Source Option node. It adds source terms to the transport equations of K-Epsilon and K-Omega to counteract turbulence decay in external aero flows. The source terms are inferred from the specified inflow boundary. Provided that reasonable inflow turbulence is specified, the sources should not contaminate any boundary layers.

The benefits of this feature are:

- Easier fine tuning of turbulence intensity without having to have excessively large turbulent viscosity ratios. This will be particularly advantageous for the transition model.
- The ability to properly simulate a plane or car moving through background turbulence.

Properties

Inflow Boundary	Lets you select an inflow boundary from which the correct turbulence source will be computed. However, it is assumed that the inflow boundary is properly connected to the region, and that flow uniformly enters the domain at this boundary.

Turbulence Intensity (Porous Region)

This region value is added to the Values node by selecting Intensity and Length Scale or Intensity and Viscosity Ratio as the Method in the Physics Conditions > Turbulence Specification node of a porous region. It is used to define turbulent quantities within the porous region.

Properties

This value is input as a scalar profile.

Turbulent Length Scale (Porous Region)

This node is added to the Values node by selecting `Intensity and Length Scale` as the Method in the Physics Conditions > Turbulence Specification node of a porous region. It is used to define turbulent quantities within the porous region.

Properties

This value is input as a scalar profile.

Turbulent Viscosity Ratio (Porous Region)

This node is added to the Values node by selecting `Turbulent Viscosity Ratio` (in the case of the Spalart-Allmaras model) or `Intensity + Viscosity Ratio` (for other turbulence models) as the Method in the Physics Conditions > Turbulence Specification node of a porous region. It is the ratio of turbulent to laminar viscosity (μ_t/μ).

Properties

This value is input as a scalar profile.

Turbulent Dissipation Rate (Porous Region)

This node is added to the Values node by selecting `K + Epsilon` or `RS + Epsilon` as the Method in the Physics Conditions > Turbulence Specification node (depending on the selected turbulence model). It is a direct specification of ε within the porous region.

Properties

This value is input as a scalar profile.

Turbulent Dissipation Rate Source

This node is added when the Turbulence Source Option region condition is ticked. It is entered as a scalar profile.

Turbulent Kinetic Energy (Porous Region)

This node is added to the Values node by selecting `K + Epsilon` or `K + Omega` as the Method in the Physics Conditions > Turbulence Specification node (depending on the turbulence model selected) of a porous region. It is a direct specification of k within the porous region

Properties

This value is input as a scalar profile.

Turbulent Kinetic Energy Source

This node is added when the Turbulence Source Option region condition is ticked. It is entered as a scalar profile.

Turbulence Field Functions Reference

The following primitive field functions are made available to the simulation when turbulence models are used. These objects all have the same base properties.

Effective Viscosity	The sum of the laminar and turbulent viscosities, $\mu + \mu_t$.
Strain Rate Tensor Modulus	S in Eqn. (527), used by the K-Epsilon and K-Omega models. This is only available when the Temporary Storage Retained option is ticked for the relevant turbulence solver.
Turbulent Viscosity	The turbulent viscosity μ_t.
Turbulent Viscosity Ratio	The ratio of turbulent to laminar viscosity, μ_t/μ.
Ustar	The reference velocity, u^*, computed according to the selected wall treatment.
Virtual Local Heat Transfer Coefficient	An approximation to the true Local Heat Transfer Coefficient that does not require an energy model to be enabled. It is defined as $h = \rho C_p u^*/t^+$ where t^+ is computed according to the selected wall treatment. Since thermal material properties are not available, the specific heat, C_p, is specified as a reference value, as well as the molecular and turbulent Prandtl numbers. The fluid density, ρ, is obtained locally from the simulation.
Wall Y+	The value u^*y/ν where u^* is the reference velocity, y is the normal distance from the centroid to the wall in wall-adjacent cells and ν is the kinematic viscosity.
Wall Distance	The normal distance y from the cell centroid to the wall as calculated by the Wall Distance solver.

Modeling Radiation

The Radiation Model is the gateway, or entry point to all of the radiation modeling capabilities of STAR-CCM+. This section describes STAR-CCM+'s radiation modeling.

To solve a radiation problem:

- Select the Radiation model.
- Select one of the radiative transfer models.
- Select one of the radiation spectrum models.

Reference information is provided for field functions that the radiation models use.

Radiative Transfer Models

The radiative transfer model defines the overall solution method for the governing Radiative Transfer Equation (RTE), including how the RTE is discretized with respect to space and solid angle. Available transfer models are:

- Surface-to-Surface (S2S) radiative heat transfer for modeling non-participating media radiation.
- Participating Media Radiation (the Discrete Ordinates Method, or DOM).

Radiation Spectrum Models

The radiation spectrum model defines how the radiation wavelength spectrum is considered (and potentially discretized) within the context of the transfer model. The radiation spectrum models are:

- Gray Thermal Radiation for modeling wavelength-independent radiation properties.

- Multiband Thermal Radiation for modeling wavelength-dependent radiation properties.

For the Gray Thermal spectrum model, the full thermal wavelength domain is considered as a whole, and all radiative properties are considered invariant within this single spectrum. For the Multiband Thermal spectrum model, the full-thermal spectrum can be divided into subspectra or spectral bands. For each subspectrum, waveband-specific properties can be defined. During the radiation solution process, the transfer model is applied to each subspectrum. The total radiative transfer is the sum of the resulting solutions from the individual subspectra.

For the S2S transfer model, solar loads can additionally be considered with both the Gray and Multiband Thermal spectrum models by selecting the optional Solar Loads model.

Surface-to-surface modeling also requires view factors. These view factors are calculated using the View Factors model.

Radiation with Shell Regions

The model of radiation with solid shells is consistent with the general approach to radiation in solids, and is compatible with solar radiation, and with gray and multi-band radiation modeling. Each shell boundary has individually specified radiative properties (emissivity, reflectivity, transmissivity, and specularity).

The example given in the following figure shows a shell bounding a fluid, such as a thin-walled pipe containing liquid, applicable to engine and heating simulations. A shell can also stand alone, without containing fluid, for example as a mirror in a lighting simulation.



Selecting the Radiation Model

To select the Radiation model:

- In the *Physics Model Selection* dialog, select a space model from the *Space* group (Three Dimensional, Two Dimensional, or Axisymmetric). The Radiation radio button now becomes available.

Note: Participating Media Radiation (DOM) is available for all space models; Surface-to-Surface (S2S) Radiation is only available for three-dimensional space.

- Select the Radiation model in the *Optional Models* group.

The *Physics Model Selection* dialog with the Radiation model selected is shown below.



Radiation Characterization

Radiation is characterized in terms of the following:

- Radiation Spectrum
- Spatial Domain for Solution of Radiation Transfer
- Radiation Properties

Radiation Spectrum

Within the full electromagnetic spectrum, the thermal radiation spectrum - approximately 0.1 to 100 μm - is pertinent to heat transfer. This thermal spectrum includes a portion of the ultra-violet (UV) range and all of the visible and infra-red (IR) ranges. The thermal radiation that a given object emits generally has a wide distribution of wavelengths. Relatively hot objects, such as the sun, emit thermal radiation over the entire thermal spectrum. Relatively cooler objects (T < 800K) emit predominantly in the infra-red range.

When radiation properties are invariant with wavelength, the radiation is said to be gray. Gray radiation is modeled using the Gray Thermal Radiation model. The radiation properties are taken to be same over the entire thermal spectrum, and only a single radiative transfer solution is necessary for the full thermal spectrum.

When the radiation properties vary with wavelength, the radiation is said to be spectral. Spectral radiation is modeled using the Multiband Thermal Radiation model. Different radiation properties can be specified for each individual spectral band, and radiative transfer is considered individually in each band. The total transfer is the sum of the contributions from each individual band.

Spatial Domain for Solution of Radiation Transfer

The effects of radiation are modeled within the regions in the simulation as well as between the regions. Radiation exchange between regions takes place whenever the interface between the regions is transparent or partially transparent to radiation.

In addition, radiation exchange with the surrounding environment can be modeled with the Surface-to-Surface transfer model (see schematic below). Both terrestrial and solar emissions can be considered. The Planck

distribution governs the distributions of the solar and terrestrial emissions over the thermal spectrum. Solar emissions span the entire thermal radiation spectrum including the UV, visible, and IR ranges.



Schematic Representation of Radiation

Thermal radiative exchange with the terrestrial environment includes only diffuse effects. The terrestrial environment surrounding the regions emits diffusively as a blackbody at the environment radiation temperature. It is also possible to have thermal radiative exchange with the solar environment.

Radiation Properties

The definition of a radiation problem is completed by defining the properties of the media inside the regions and the surface properties of the boundaries. Non-participating media are transparent to radiation. The only requirement is that the surface properties of emissivity, reflectivity, and transmissivity are specified on the boundaries. In contrast, participating media can emit, absorb, and scatter radiation. These media additionally require you to specify the material absorption and scattering properties within the regions using the continuum material properties.

Modeling Surface-to-Surface Radiation

This section describes:

- The S2S (Surface-to-Surface) Radiation model
- How to select the S2S Radiation model
- How to set the relevant parameters for both the S2S Radiation model and its corresponding solver
- How to post-process radiation on non-diffuse surfaces
- The formulation of the source terms for the energy equation.

What Is the S2S Radiation Model?

The S2S (Surface-to-Surface) Radiation model makes it possible to simulate thermal radiation exchange between diffuse surfaces forming a closed set. The medium that fills the space between the surfaces is non-participating. That is, it does not absorb, emit, or scatter radiation. Under these circumstances, the radiation properties and the thermal boundary conditions that are imposed on each surface uniquely define the amount of radiation that a surface receives and emits. The surface properties are quantified in terms of emissivity, reflectivity, transmissivity, and radiation temperature. These properties are not dependent on direction; however, they can be dependent on radiation wavelength with the Multiband Radiation model.

The corresponding node (Surface-to-Surface Radiation in the following screenshot) is activated when you select an S2S simulation.



Note: The S2S Radiation model can apply to a multi-part object. Inputs must be specified per-boundary.

Note: An environment boundary condition makes it possible to model both convection and radiation heat transfer between the external sides of domain boundaries and the environment.

Summary of the STAR-CCM+Model

Within STAR-CCM+, the S2S Radiation model is operated in two steps:

1. Calculate view factors using ray tracing.
2. Apply view factors to compute radiosity and irradiation fields on all surfaces.

In the view factor calculation step, deterministic ray tracing is applied to compute the factors. Rays are traced into the hemisphere above each patch using a fixed distribution of directions over this hemisphere. By default, a distribution of 1024 rays per patch is used.

The view factors must be calculated only when the geometry and radiative properties such as transmissivity and specular reflectivity change. Therefore, they generally can be calculated once and then reused during the combined-mode heat transfer solution to update the radiative fluxes from the surface emissive powers (that is temperatures).

An advantage of the two-step approach is computational efficiency gained from doing relative infrequent ray tracing. The trade-off is that view factors must be stored, often requiring considerable amount of memory.

Impact of Patch and Angular Resolution

View factors are considered computationally efficient for combined-mode problems. However, the number of degrees of freedom (that is, the number of patches) is often limited. The limitation is because the view factors count can be $O(N^2)$ in the limit, where N is the number of patches. The STAR-CCM+ approach provides flexibility through patch and angular resolution to help address the computation time and memory issues that are associated with view factors. This approach makes it is possible to run cases having a million patches or so.

With the deterministic ray tracing approach in STAR-CCM+, the number of rays per patch limits the view factor count. Therefore, with reciprocity you can get a maximum of 2L view factors per patch, where L is the number of rays per patch. Thus, 1 million patches with 1024 rays/patch could produce a maximum of about 2 billion view factors. This number is manageable in STAR-CCM+ in terms of both computation time and memory with parallel

ray tracing and S2S algorithms. The resolution of the ray-tracing can be increased to increase accuracy, but at the expense of computation time and memory. There is thus a classic trade-off between accuracy and required computational resources (time and memory).

For some combined-mode problems, the resolution of the underlying mesh for the advection/diffusion transport can be prohibitively high to use its boundary faces directly as patches in the S2S calculations. For instance, some simulations can have tens of millions of boundary faces, making a 1-to-1 ratio between the S2S patches and the boundary faces intractable in these situations. Additionally, this fine patch resolution is not always needed with respect to accuracy, especially when the surface topology and surface temperatures do not change rapidly. STAR-CCM+ therefore provides the flexibility to use a patch resolution that is coarser than the boundary face resolution. Patches can be formed, for instance, from aggregating the boundary faces from the underlying mesh. For high-resolution meshes, you can generally use 1 patch for every 10 or so faces from the underlying mesh. There is again a trade-off between accuracy and required computational resources. A finer patch resolution generally provides better accuracy, but also requires more computational resources. You can find a balance that provides the needed accuracy within the given time and hardware constraints. In the end, most methods for solving the radiative transfer equations are constrained in a similar trade-off between performance and accuracy due to the nature of the physics (transfer at a distance).

Comparison with the Monte Carlo Method

This section briefly compares some of the merits of the STAR-CCM+ S2S model described above relative to the Monte Carlo method, which is not implemented in STAR-CCM+. The Monte Carlo method can be applied in two ways to solve the radiative transfer equations (RTE): to compute view factors in a two-step approach, or to solve the RTE equation directly.

When the Monte Carlo method is used to compute view factors, the directions for rays leaving a patch are selected randomly (non-deterministically). Because the Monte Carlo method is a statistical approach, it generally requires a high number of rays to obtain a statistically meaningful result. In contrast, STAR-CCM+ uses a deterministic approach where the directions are based on a fixed quadrature. Here, the number of traced rays determine the accuracy. The regular angular distribution of rays in the deterministic approach provides flexibility to vary angular resolution while still maintaining symmetry properties at modest resolutions. In the limit of an infinite number of rays, both methods give the same set of view factors.

When the Monte Carlo method is used to solve the RTE directly, it is invoked every time a solution of the RTE is required. For a combined-mode heat transfer analysis, the solution is normally required once every global iteration and is often prohibitively expensive for most practical problems. Therefore, the use of view factors that are calculated from either a deterministic or Monte Carlo approach is an advantage with respect to computation time: the ray tracing is done only once at the beginning of the simulation (assuming a static geometry).

Selecting the S2S Radiation Model

To select the S2S Radiation model:

- In the *Physics Model Selection* dialog, select a space model (Three Dimensional) from the *Space* group.
- Select a time model from the *Time* group.
- Select the Radiation model from the *Optional Models* group.
- Select the Surface-to-Surface Radiation model from the *Radiation* group.

  Select the appropriate material and flow models as well.
- In the *Physics Model Selection* dialog, select a single-component material model from the *Material* group.

Choose the desired solver type. For the segregated flow model:

- Select Segregated Flow in the *Flow* group.
- Select a model from the *Energy* group: Segregated Fluid Temperature, Segregated Fluid Isothermal, or Segregated Fluid Enthalpy.

- Following this selection, the *Physics Model Selection* dialog appears as shown in the following screenshot.



Alternatively, to choose the coupled flow model:

- Select Coupled Flow in the *Flow* group.
- The Coupled Energy model is selected automatically in the *Optional Models* group, provided you have the Auto-select recommended physics models checkbox activated (the default setting).
- Following this selection, the *Physics Model Selection* dialog appears as shown in the following screenshot.



Patches and View Factors

The S2S Radiation model relies on:

- A spatial discretization of the boundary surfaces into patches.
- View factors that quantify, for each patch, the proportion of surface area that the other patches illuminate.

Patches and view factors are explained in further detail in the S2S Radiation Formulation section, and are generated using the View Factors model.

Setting S2S Radiation Model Parameters

The S2S Radiation model requires specification of the radiation properties of the boundary surfaces. These radiation properties are defined in Setting Thermal Radiation Model Parameters.

However, if you use S2S, the media within all regions that have S2S activated are considered transparent to radiation transfer. All boundary surfaces are oriented inward such that radiative transfer is considered *inside* the regions. It is possible to deactivate S2S transfer in select regions. Additionally, for special situations involving solid regions, the surface orientation can be changed to outward such that only radiation external to the regions is considered.

What Is the Radiation Transfer Option?

Select the regions of your model which use S2S radiation. Deactivate radiation in a region if radiation transfer is not relevant in that region. Also deactivate radiation in a region that does not have any surfaces that "see" surfaces in other regions where radiation is active. For example, you can deactivate it in regions far away from the underhood of a car or truck.

The S2S transfer in a region is specified with the properties of the Radiation Transfer Option node under the Physics Conditions for the region.



One constraint for using this option is that either a baffle or porous baffle must be placed between any two regions for which one region has the option activated and the other does not. This baffle ensures that rays do not pass into regions where S2S is deactivated from regions where it is activated.

Radiation Transfer Option Properties

Radiation Transfer Option	Determines whether S2S transfer is solved in the region.	
	Activated	Activates S2S transfer in the region (default).
	Deactivated	Deactivates this option.

What Is the Radiation Flux Option?

This option allows you to model diffuse radiation sources at boundaries according to your own specification. For example, you can use this option to model illumination from LED bulbs or from daylight. The radiation source is non-Planck radiation unrelated to the temperature of the boundary, although this source can affect the thermal behavior of the far field boundaries and the system being modelled. The source can be either uniform over the spectrum (gray model) or spectrally variant (multi-band model). See the formulation section for more information.

To activate diffuse radiation at a boundary, select the Radiation Flux Option node the Physical Conditions for the boundary.



In the properties for the node, under Radiation Flux Option, select Diffuse radiation flux from the drop-down menu. This action creates a Diffuse Radiation Flux node under the Physics Values for the boundary. The Diffuse Radiation Flux is specified as a scalar profile.

If you are using the Multiband Thermal Radiation model, this option creates a Composite node with a node for each thermal band under it. If you are using the Gray Thermal Radiation model, there is only one set of properties. Select the methods for specifying the flux.



If you use the Environment Condition thermal specification to model a dual-sided boundary along with a diffuse radiation source term, the input value is applied to each side (the inner and the outer sides). That is, the input value is applied twice when summed over that dual-sided boundary.

Radiation Flux Option Properties

Radiation Flux Option	Determines whether or not a heat flux source is specified at the boundary.	
	No flux	No heat flux at this boundary (default).
	Diffuse radiation flux	Creates a Diffuse Radiation Flux node under Physics Values.

What Is the S2S Surface Orientation Option?

The S2S Surface Orientation Option is a deprecated option for modeling radiation between two opaque solids. It is recommended that you use the Environment Condition instead.

The S2S Surface option makes provision for the special situation where you want to model radiation exchange between two opaque *solids* that do not have a mesh in the void between. For such cases, you can set the S2S Surface Orientation property of the S2S Surface Orientation Option node to Outward on both bodies. For all other cases with mesh in between, leave the option at its default setting of Inward.

This option is only available for solid regions belonging to continua in which S2S Radiation is active.



S2S Surface Orientation Option Properties

S2S Surface Orientation	Specifies the orientation of the boundary surfaces for S2S calculations.	
	Inward	Surfaces are directed inward for S2S.
	Outward	Surfaces are directed outward for S2S.

Setting S2S Solver Parameters

This section discusses how to activate or freeze the S2S solver and how to set the algebraic multigrid parameters for that solver.

Using the S2S Solver

The S2S solver is activated by default, after Surface-to-Surface Radiation has been selected in the *Physics Model Selection* dialog and a case geometry has been imported. You can control this solver through its expert properties, which you would access by selecting the S2s Solver node under the Solvers node.



- To deactivate the solver, activate the checkbox of the Solver Frozen property.

S2s Solver Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Setting Algebraic Multigrid Parameters

To set the algebraic multigrid parameters for the S2S solver, open the S2s Solver node under the Solvers node. Then, select the node for the S2s AMG Linear Solver as shown below. The algebraic multigrid parameters can be entered using the *Properties* window.



Post-Processing Radiation on Non-Diffuse Surfaces

STAR-CCM+ allows you to post-process radiative parameters on fully transmissive and specularly reflective boundaries.

For fully transmissive boundaries, you can post-process the following:

- Irradiation - the radiative heat flux incident on a surface
- Mean radiant temperature
- Direct solar irradiation - the direct solar flux on the boundary.

For example, consider a thermal comfort case where a car driver experiences solar irradiation:



You can view the direct solar irradiation on the fully transmissive windows:



For fully specular boundaries that have non-zero specular reflectivity, you can post-process the same field functions as for fully transmissive boundaries. In addition, you can post-process radiosity: the radiative heat flux leaving a surface.

On fully specularly reflective boundaries, the total irradiation accounts for direct solar irradiation and irradiation from other boundaries. For example, the total irradiation on the boundary of a solar collector contains the contributions from the direct solar irradiation.



Fully transmissive and fully specular boundaries do not participate in the solution of S2S and energy equations; zero net radiative flux exists across them.

Note: View factors are calculated from non-diffuse surfaces to other partially diffuse surfaces, but not the other way round.

S2S Radiation Formulation

This section presents the following:

- The S2S Thermal Radiation Model
- Computation of the radiant flux Energy Equation Source Term at the boundaries
- The Mean Radiant Temperature

Also included in this section are:

- Nomenclature
- Bibliography

S2S Thermal Radiation Model

The S2S Radiation model is based on enclosure theory as described in detail in [125] and [128].

The net radiant flux on each surface is a function of the surface properties, the thermal boundary conditions that are imposed on that surface, and is calculated so that radiation is balanced. The radiation balance is enforced on the entire closed set of surfaces, by considering each surface and how it exchanges radiation with all other surfaces (see Radiation Equilibrium Formulation). Therefore, the first step in the model is to compute the surface-to-surface interaction for each surface pair, as described in the Surface Patches and View (Interaction) Factors section.

Surface Patches and View (Interaction) Factors

General Definition

The surfaces in the model can have either diffuse or specular surface properties. The surface emissivity is assumed diffuse, but the reflectivity can be either diffuse or specular, and all transmissivity is specular. Specular properties are considered in the calculation of the view factors. Hence, the view factors are not merely geometric properties, but also include the influence of the specular surface properties.

The amount of radiation that is exchanged between surfaces depends on the (diffuse) emission from the surfaces and the position and orientation of the surfaces relative to each other. The amount of radiation also depends on the presence of any specularly transmissive or reflective surfaces that can alter the transfer between the surfaces. The latter two influences -- geometry and specular properties -- are accounted for by the view (or interaction) factors which, by definition, represent the “fraction of uniform diffuse radiation leaving a surface that directly reaches another surface” [125].

For the two elemental surfaces that are depicted below,



Radiation Exchange Between Two Elemental Surfaces

The total amount of radiation power that dS_1 emits and dS_2 receives is:

$$P_{1-2} = i'_1 dS_1 \cos(\beta_1)\left(\frac{dS_2 \cos(\beta_2)}{L^2}\right) \tag{847}$$

By taking the fraction of the power that dS_2 receives over the total radiation power that is emitted from dS_1 $(P_1 = \pi i'_1 dS_1)$, the view factor for the two surfaces is obtained as:

$$dF_{1-2} = \frac{P_{1-2}}{\pi i'_1 dS_1} \tag{848}$$

Applying definition (848) to two black surfaces *i* and *j* in thermal equilibrium at the same temperature (no net energy exchange) yields the requirement that view factors must satisfy the reciprocity relation:

$$dF_{i-j}\, dS_i = dF_{j-1} dS_j \tag{849}$$

For diffuse radiation, relationship (849) applies to any two surfaces, whether black or not.

These definitions can be extended to surfaces of finite size, for which the view factors are obtained using integration over constituent elemental surfaces, as shown below.



Radiation Exchange Between Two Finite Surfaces

For the (S_1, S_2) pair:

$$F_{1-2} = \frac{\int\int_{S_1 S_2} \pi i'_1 \frac{\cos(\beta_1)\cos(\beta_2)}{\pi L^2}\, dS_1\, dS_2}{i'_1 S_1} \tag{850}$$

and:

$$F_{1-2} = \frac{1}{S_1}\int\int_{S_1 S_2} \frac{\cos(\beta_1)\cos(\beta_2)}{\pi L^2}\, dS_1\, dS_2 \tag{851}$$

Likewise:

$$F_{2-1} = \frac{1}{S_2}\int\int_{S_2 S_1} \frac{\cos(\beta_1)\cos(\beta_2)}{\pi L^2}\, dS_1\, dS_2 \quad \text{or} \quad F_{2-1} = F_{1-2}\frac{S_1}{S_2} \tag{852}$$

In STAR-CCM+, the boundary surfaces are discretized into smaller elements called patches. These patches are sets of contiguous boundary cell faces, and view factors are calculated for each patch pair. By definition, patches do not straddle boundaries and are therefore at most as large as an entire boundary or as small as a boundary cell face. The emissive power (and therefore the radiation intensity) and radiation properties are assumed to be uniform over the surface of each patch.

The view factors as in Eqns. (851) and (852) above are purely topological quantities, dependent only on the geometry of the two surfaces. However, the presence of partially transmissive and/or specularly reflective surfaces can introduce a further dependence on other surfaces and their respective surface properties. As a result, the collection of view factors is, in general, dependent on the surface geometry and the respective transmissivity and specular reflectivity of those surfaces. For a fixed geometry and fixed surface transmissivity and specular reflectivity, the view factors must be obtained only once, usually at the start of a simulation.

The double integral is approximated using a ray tracing approach. Collections of polygons represent surfaces, and the solid angle is discretized using an angular quadrature. For each surface or patch, a predefined number of rays is traced through the computational domain starting at the patch centroid. The directions and weights are based on discretization of an ideal hemisphere over the patch.

A ray shot from one patch can do any of the following:

- Hit an opaque patch.
- Pass through a semi-transparent patch (with reduced strength proportional to the transmissivity of the boundary to which the hit patch belongs).

- Reflect from a specularly reflective patch.
- Pass through a transparent boundary or an opening and be dispersed in the environment.

In all these cases a view factor is generated between the shooting patch and, respectively, the opaque patch, the semi-transparent patch, the specularly reflective patch, or the environment.

This method is fairly simple and cost-effective compared to other methods (hemicube, hemisphere, or Monte Carlo). There are some analogies to the Unit Sphere Method [126]. After the beam tracking has finished, the array of view factors obtained satisfies conservation but not reciprocity, so a post-correction procedure is necessary to enforce reciprocity. Conservation is also enforced, since as soon as reciprocity is imposed conservation is lost.

The general definitions above describe the direct, unimpeded interaction between two isolated surfaces. In reality, when more surfaces are involved, radiation energy issuing from a patch can reach another patch after having gone through one or more non-opaque surfaces. The radiation energy can also undergo reflection by one or more specularly reflective surfaces.

Note: Diffuse reflection is considered later in the surface-to-surface calculations proper.

The S2S Radiation Model makes it possible to model specular transmission, where radiation emitted from a surface is transmitted through another surface along its original direction. This situation is schematized below for boundaries B_1 to B_4 and selected patches.



Computing, for example, the irradiation from patch 2 on B_1 onto patch 8 on B_4, involves summing up the following: the energy which reaches 8 through 5 directly, the energy which reaches 8 through 6 directly, and the energy which reaches 8 through both 6 and 7. Each time a radiation ray is transmitted through a surface, the transmissivity of that surface attenuates its intensity. Therefore, the only way to account properly for all the possible paths is to account for them and the transmissivities at the time when the view factors are calculated. The calculation uses integration over small, discretized solid angles.

The influence of specularly reflective surfaces is similarly considered during the ray tracing process. Specularly reflective surfaces change the direction of rays and potentially the ray strength when the surface is less than fully reflective. Furthermore, surfaces can be both specularly reflective and transmissive. In this special case, each incident ray is split into two child rays -- one transmitted through the surface and one reflected from the surface. The strengths of the resulting two child rays are determined from the strength of the parent ray and by the respective specular transmissivity and reflectivity.

When transmissive and/or specularly reflective surfaces are present, the view factors for a patch pair are global view factors. These global view factors already account for all possible radiation interactions in the space between the two patches. The global view factors no longer are purely topological quantities, and must be recalculated every time the transmissivity or specular reflectivity of a surface in the domain is changed.

Environmental Patch

In addition to the patches on physical boundaries, a virtual patch is created when the view factors calculator detects non-opaque external surfaces. The environmental patch is internal to the code and makes it possible to collect and account for all the energy radiated to and from the environment around the regions. By definition, the environmental patch temperature is the Radiation Temperature of the Thermal Environment.

Radiation Equilibrium Formulation

After the surfaces have been decomposed into patches and view factors have been calculated for the patch pairs, radiant fluxes on each patch are obtained. The radiant fluxes are obtained by enforcing radiation equilibrium on the entire closed set of surfaces for each radiation spectrum or spectral band independently.

For a given radiation spectrum or spectral band, λ, radiation contributions on each patch can be expressed in terms of the incident, emitted, and transmitted radiation, as illustrated below for a generic patch i.



Schematic Representation of Radiation Fluxes on Patch i

For radiation to balance over the entire closed set of surfaces, all radiation received by each patch must come from the other patches. Summing over all the patches, the total radiation incident on patch i is then:

$$Q_{I,i,\lambda} = S_i I_{i,\lambda} = \sum_{j=1}^{N_p} \frac{F_{\text{j-i},\lambda}}{(1-\tau_{i,\lambda}-\rho_{s,i,\lambda})} J_{d,j,\lambda} S_j + \frac{F_{\text{e-i},\lambda}}{(1-\tau_{i,\lambda}-\rho_{s,i,\lambda})} J_{e,\lambda} S_e \quad (853)$$

where:

- $F_{\text{j-i},\lambda}$ is the view factor from patch j to i resulting from the approximation of Eqns. (847)-(852) using ray tracing.
- J_d is the diffuse component of radiosity.
- $\tau_{i,\lambda}$ is the transmissivity.
- $\rho_{s,i,\lambda}$ is specular component of reflectivity.
- The subscript e is used for the environmental patch.

The view factors include the influence of any transmissive or specularly reflective surfaces. Correspondingly, the view factors are effectively normalized using the diffuse portion of the surface properties at patch i: $(1-\tau_{i,\lambda}-\rho_{s,i,\lambda})$.

The general reciprocity relationship for view factors between surfaces that can be transmissive and/or specularly reflective is:

$$F_{\text{j-i},\lambda} = F_{\text{i-j},\lambda}\left[\frac{S_i(1-\tau_{i,\lambda}-\rho_{s,i,\lambda})}{S_j(1-\tau_{j,\lambda}-\rho_{s,j,\lambda})}\right] \quad (854)$$

By introducing the reciprocity relationship into Eqn. (853), the incoming radiation flux becomes:

$$I_{i,\lambda} = \sum_{j=1}^{N_p} \frac{F_{\text{i-j},\lambda}}{1-\tau_{j,\lambda}-\rho_{s,j,\lambda}} J_{d,j,\lambda} + \frac{F_{\text{i-e},\lambda}}{1-\tau_{e,\lambda}-\rho_{s,e,\lambda}} J_{e,\lambda} \quad (855)$$

where $J_{d,j,\lambda}$ is the diffuse radiosity:

$$J_{d,j,\lambda} = \rho_{d,j,\lambda} I_{j,\lambda} + E_{j,\lambda} + E_{d,j,\lambda} \quad (856)$$

where:

- $E_{j,\lambda}$ is the emission based on Planck's Law.
- $E_{d,j,\lambda}$ is the user-specified non-Planck emission.

After introducing the effective radiosity as:

$$J_{i,\lambda}^{eff} = \frac{J_{d,i,\lambda}}{1-\tau_{i,\lambda}-\rho_{s,i,\lambda}} \quad (857)$$

the irradiation can be written as:

$$I_{i,\lambda} = \sum_{j=1}^{N_p} F_{\text{i-j},\lambda} J_{j,\lambda}^{eff} + F_{\text{i-e},\lambda} J_{e,\lambda}^{eff} \tag{858}$$

The total diffuse radiation that patch i emits is the sum of the diffusely reflected and emitted components:

$$J_{d,i,\lambda} = \rho_{d,i,\lambda} I_{i,\lambda} + E_{i,\lambda} + E_{d,i,\lambda} \tag{859}$$

Combining Eqns. (858) and (859) for each patch yields an equation for $J_{i,\lambda}^{eff}$:

$$\begin{aligned}(1-\tau_{i,\lambda}-\rho_{i,\lambda}^{s})(1-F_{\text{i-i},\lambda})J_{i,\lambda}^{\text{eff}} - \rho_{d,i,\lambda}\sum_{(j=1),(j\neq i)}^{N_p}[F_{\text{i-j},\lambda}J_{j,\lambda}^{\text{eff}}] \\ = E_{i,\lambda} + E_{d,i,\lambda} + \rho_{d,i,\lambda}[F_{\text{i-e},\lambda}(E_{\text{env},\lambda} + J_{d,\text{solar},\lambda}) + J_{\text{dir,solar},\lambda}]\end{aligned} \tag{860}$$

where:

- $J_{d,\text{solar},e}$ is the diffuse solar flux, a specified solar parameter.
- $J_{\text{dir,solar},i,\lambda}$ is the direct solar flux impinging on the given patch, calculated using the view factors solver at the time of the view factors calculation.

To close the system of equations, the emitted energy and environment radiosity on the right-hand side of Eqn. (860) can be further specified. These terms are defined below for an arbitrary spectrum or band. As such, the definitions are applicable to either the Gray or Multiband spectrum models. Simplifications for the gray spectrum model (Gray Thermal Radiation) are noted where appropriate.

The surface emissive power for a given spectrum or band is the product of the emissivity and the blackbody emissive power for that spectrum or band:

$$E_{i,\lambda} = \varepsilon_{i,\lambda}\Delta_{(\lambda_1 \to \lambda_2, T_i)}\sigma T_i^4 \tag{861}$$

where $\Delta_{(\lambda_1 \to \lambda_2, T_i)}$ denotes the fraction of total blackbody emission for the spectrum or band, and λ_1 and λ_2 represent the lower and upper wavelength bounds of the spectrum or band. This fraction is graphically depicted in the following figure:



Fractional blackbody emission for a band

For gray radiation, the wavelength bounds are effectively zero and infinity and this factor is simply unity. For multiband radiation, this fraction is a function of the surface temperature and wavelength bounds:

$$\Delta_{(\lambda_1 \to \lambda_2, T)} = \Delta_{(0 \to \lambda_2, T)} - \Delta_{(0 \to \lambda_1, T)} = f(\lambda_2 T) - f(\lambda_1 T) \tag{862}$$

The fractional blackbody emission function can be evaluated by either integrating the Planck distribution or by interpolating from a table of based on this integration [125].

The environment is modeled as a black body with unity emissivity and zero reflectivity and transmissivity. Thus, the effective radiosity of the environment is:

$$J_{e,\lambda}^{eff} = \Delta_{(\lambda_1 \to \lambda_2,\, T_e)} \sigma T_e^4 \tag{863}$$

As before, the fraction of total blackbody emission for a given spectrum or band is dependent on the extent of the spectrum or band and the environment temperature. For Gray Thermal Radiation, the fraction reduces to unity.

If solar loads are considered, two more sources can be present -- direct solar radiation and diffuse solar radiation. For purposes of apportioning these loads by spectrum or band, the spectral distribution of the solar emissions is assumed to be represented with a blackbody at the average surface temperature of the sun. The direct and diffuse solar fluxes augment the effective radiosity of the environment:

$$F_{\text{i-e},\,\lambda} J_{e,\lambda}^{eff} = F_{\text{i-e},\,\lambda}[\Delta_{(\lambda_1 \to \lambda_2,\, T_e)} \sigma T_e^4 + \Delta_{(\lambda_1 \to \lambda_2,\, T_{sun})} J_{d,\,\text{solar},\,e}] + J_{\text{dir, solar},\,i,\,\lambda} \tag{864}$$

This patch-based quantity is dependent on the specified sun position, and the specified direct solar flux, and the bounds of the given spectral waveband (used for apportioning the total direct flux). This quantity is also dependent on the specified spectrum/band transmissivities and the model geometry.

Substituting Eqns. (861) and (864) into (860) and expanding the i to i terms out of the sum yields the general system of equations for the effective radiosity:

$$\begin{aligned}(1 - \tau_{i,\lambda} - \rho_{s,\,i,\,\lambda})(1 - F_{\text{i-i},\,\lambda})\, J_{i,\lambda}^{eff} - \rho_{d,\,i,\,\lambda} \sum_{j=1,\, j \neq i}^{N_p} F_{\text{i-j},\,\lambda}\, J_{j,\lambda}^{eff} \\ = \varepsilon_{i,\lambda} \Delta_{(\lambda_1 \to \lambda_2,\, T_i)} \sigma T_i^4 + E_{d,\,i,\,\lambda} \\ + \rho_{d,\,i,\,\lambda} \left\{ F_{\text{i-e},\,\lambda}[\Delta_{(\lambda_1 \to \lambda_2,\, T_e)} \sigma T_e^4 + \Delta_{(\lambda_1 \to \lambda_2,\, T_{sun})} J_{d,\,\text{solar},\,e}] + J_{\text{dir, solar},\,i,\,\lambda} \right\}\end{aligned} \tag{865}$$

For the gray spectrum model, only one system of equations exists, and that system represents the full thermal spectrum. For the multiband spectrum model, a separate system exists for each spectral band. In the absence of solar radiation, the solar load terms are omitted from Eqn. (865).

Numerical Solution

The patch values for emissive power, reflectivity, and transmissivity are obtained by surface-averaging the cell-face values for all boundary faces that constitute each patch:

$$E_{i,\lambda} = \varepsilon_{i,\lambda}\Delta_{(\lambda_1 \to \lambda_2,\, T_i)}\sigma T_i^4 = \frac{\sum_{f_1(i)}^{f_N(i)} \varepsilon_{f_k,\lambda}\,\Delta_{(\lambda_1 \to \lambda_2,\, T_{f_k})} T_{f_k}^4\, S_{f_k}}{\sum_{f_1(i)}^{f_N(i)} S_{f_k}}, \quad \rho_{i,\lambda} = \frac{\sum_{f_1(i)}^{f_N(i)} \rho_{f_k,\lambda}\, S_{f_k}}{\sum_{f_1(i)}^{f_N(i)} S_{f_k}},$$

$$\tau_{i,\lambda} = \frac{\sum_{f_1(i)}^{f_N(i)} \tau_{f_k,\lambda}\, S_{f_k}}{\sum_{f_1(i)}^{f_N(i)} S_{f_k}} \tag{866}$$

Surface-averaging the face emissive power rather than averaging the face emissivity and temperature independently is required to conserve energy. Given the emissive powers and properties for the entire set of patches, the set of radiation balance equations, Eqn. (865), is solved in matrix form by the S2S solver and its AMG Linear Solver for each spectrum or band.

Based on the updated values of the effective radiosities, the patch irradiation fluxes are obtained directly from Eqn. (858). The contribution to the irradiation from the environment in Eqn. (858) is based on either Eqn. (863) or (864) depending on whether solar loads are considered. Radiosities are then extracted from the effective radiosities using Eqn. (857). In the presence of solar radiation, these fluxes include the influence of the solar loads.

The radiative fluxes on the faces are then derived from the patch fluxes. The irradiation (incident flux) at each face *f* is equal to the irradiation on the associated patch *i*. The face radiosity is a combination of local face emissive power and the reflected irradiation:

$$I_{f,\lambda} = I_{i,\lambda} \tag{867}$$

$$J_{f,\lambda} = E_{f,\lambda} + \rho_{f,\lambda} I_{i,\lambda} \tag{868}$$

The reflectivity in the above equation is the full reflectivity (diffuse + specular). The use of the local face emissive power rather than the patch emissive power is critical to ensuring local energy balances when the net radiative flux is later derived from the radiosity.

The full-thermal-spectrum fluxes are calculated by summing the fluxes from the individual spectra or bands:

$$I_f = \sum_{\lambda \in specta} I_{f,\lambda} \tag{869}$$

$$J_f = \sum_{\lambda \in spectra} J_{f,\lambda} \tag{870}$$

$$(q_r)_f = \sum_{\lambda \in spectra} (q_r)_{f,\lambda} = \sum_{\lambda \in specta} [(1 - \tau_{f,\lambda}) I_{f,\lambda} - J_{f,\lambda}] \tag{871}$$

For the gray spectrum model, the summation operation is trivial; the gray fluxes are the total fluxes. For the multiband spectrum model, the summation extends over all spectral bands. Eqn. (871) defines the total or full-spectrum net radiative heat flux to the boundary (total absorption minus total emission). A positive value indicates net absorption by the boundary.

Energy Equation Source Term

The effects of radiation are accounted for in the energy equation through the net radiant flux. The total boundary heat flux is augmented with net radiative heat flux:

$$q_f = (q_{norad})_f + (q_r)_f \tag{872}$$

The direction convention for all fluxes in the above equation is into the boundary.

Mean Radiant Temperature (MRT)

The MRT is a useful variable for thermal comfort analysis as it provides "the uniform temperature of an imaginary enclosure in which radiant heat transfer from the human body equals the radiant heat transfer in the actual nonuniform enclosure" [129].

The MRT is derived by equating the actual heat flux with the idealized heat flux (that from an imaginary room at a uniform temperature and with unity emissivity). By substituting the expressions for the equated heat fluxes and rearranging, the MRT can be expressed in terms of the irradiation:

$$(T_{MRT})_f = \left(\frac{I_f}{\sigma}\right)^{\frac{1}{4}} \tag{873}$$

The irradiation or incident heat flux in the above expression is the result of the surface-to-surface solution. This formulation considers the potential influences of solar radiation, radiation from the thermal environment, transmission through non-opaque surfaces, and emission from non-black surfaces.

S2S Radiation Nomenclature

In the equations that are given above, the following nomenclature applies:

α	absorptivity - the ability of a material to take up the energy of incident radiation
β	angle between the surface normal and a line joining two surfaces [rad]
$\Delta_{(\lambda_1 \to \lambda_2, T)}$	fraction of total emission from a blackbody at temperature T within a certain wavelength interval.
E	emitted energy or emissive power $[W/m^2]$
ε	emissivity
$f_k(i)$	k th constituent boundary cell face of patch i
F_{1-2}	view or configuration factor for surface pair (S_1, S_2)
i'_k	total intensity leaving surface k, where intensity is defined as the "radiative energy passing through an area per unit solid angle, per unit of the area projected normal to the direction of passage, and per unit time" [125]. $[\frac{W}{m^2 sr}]$
I	irradiation or incident radiant flux $[\frac{W}{m^2}]$
J	radiosity or outgoing radiant flux (emitted and reflected) $[\frac{W}{m^2}]$
L	distance [m]
n	normal vector
N_p	total number of patches
P_{1-2}	total power leaving surface 1 and incident on surface 2 [W]
q_f	heat flux through boundary cell face f $[\frac{W}{m^2}]$
ρ	reflectivity
σ	Stefan-Boltzmann constant $[\frac{W}{m^2 K^4}]$
S	surface area [m^2]
T	temperature [K]
τ	transmissivity

S2S Radiation Bibliography

[125] Siegel, R., and Howell, J.R. 1992. "Thermal Radiation Heat Transfer", Third Edition, Hemisphere Publishing Co.

[126] Modest, Michael F. 2003. "Radiative Heat Transfer", Second Edition, Academic Press.

[127] Taylor, R.P. and Luck, R. 1995. "Comparison of Reciprocity and Closure Enforcement Methods for Radiation View Factors", *J Thermophysics and Heat Transfer*, **9**:4, pp. 660-666.

[128] Holman, J.P. 2001. "Heat Transfer", Eighth SI Metric Edition, McGraw Hill.

[129] ASHRAE 1997. "ASHRAE Handbook-Fundamentals", Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

Modeling Participating Media Radiation

This section describes:

- The Participating Media Radiation model and how to select it.
- How to set the relevant parameters for the model and its corresponding solver.
- The mathematical formulation of the radiation calculation method and its coupling to the fluid dynamic solution.

What Is Participating Media Radiation?

In some cases, thermal radiation can occur only as a surface phenomenon. The media separating the surfaces (such as dry air) are transparent to thermal radiation. Modeling of this type of heat transfer is available in STAR-CCM+ under the title of Surface-to-Surface (S2S) radiation. However, other applications require consideration of participating media. This term refers to media lying between the radiating surfaces that can absorb, emit, and/or scatter thermal radiation. These media can include particles belonging to a Lagrangian phase, provided the Particle Radiation model is active, or particles such as soot, using the soot emission models.

The node corresponding to the implementation of this model is activated by STAR-CCM+ as soon as the Participating Media Radiation (DOM) option is selected in the *Physics Model Selection* dialog.



The main parameter of the model is the accuracy of the angular representation of space around a given cell. The order of the ordinate set, as selected from the *Ordinate Sets* drop-down list in the corresponding *Properties* window (see Numerical Solution for an explanation of this concept), specifies this accuracy.

STAR-CCM+ accounts for participating media effects by using the Discrete Ordinate Method (DOM). DOM is a numerical model for describing radiative heat transfer in participating media (see Basic Model Equations). Some of the basic assumptions of this model are shared with the S2S model as follows:

- They both simulate thermal radiation exchange between diffuse/specular surfaces forming a closed set.
- The surface radiative properties are quantified in terms of emissivity, specular and diffuse reflectivity, transmissivity (only at interfaces for DOM), and radiation temperature. These properties do not depend on direction. However, for the Multiband spectrum model, the surface properties can depend on wavelength.

However, the medium that fills the space between the surfaces can also absorb, emit, or scatter radiation. Therefore, the amount of radiation that each surface receives and emits depends on this effect, as well as the optical properties of the surface and the thermal boundary conditions that are imposed on it. The Participating Media Radiation model lets you control these phenomena by activating material properties with related parameters.

Participating Media Radiation Properties

Sn quadratures [135] are implemented to provide angular discretization in the Discrete Ordinate Method. Sn sets are implemented for n = {2,4,6,8,12,16}. S2 through S8 are based on odd-moment formulation while S12 and S16 are based on even-moment formulation (to avoid negative weights as occur for n > 10 in odd-moment formulation).

Ordinate Sets

For a given choice Sn, the number of ordinates is n*(n+2), so increasing the size of the discrete ordinate set significantly increases both accuracy and computational time.

S2	Sets order of discrete ordinate method to 2.
S4	Sets order of discrete ordinate method to 4.
S6	Sets order of discrete ordinate method to 6.
S8	Sets order of discrete ordinate method to 8.
S12	Sets order of discrete ordinate method to 12.
S16	Sets order of discrete ordinate method to 16.

Selecting the Participating Media Radiation Model

To select the Participating Media Radiation model:

- In the *Physics Model Selection* dialog, select a space model (Three Dimensional, Two Dimensional, or Axisymmetric) from the *Space* group.
- Select a time model from the *Time* group.
- Select the Radiation model from the *Optional Models* group.
- Select the Participating Media Radiation (DOM) model from the *Radiation* group.

  A requirement for the Participating Media Radiation model is that an energy model must be activated, so you must select material and flow models as well.
- In the *Physics Model Selection* dialog, select a single-component material model from the *Material* group.

  STAR-CCM+ allows radiative material properties (absorption coefficient and scattering coefficient) of participating media to be specified as a constant or field function for any continuum. A third option is provided as the Weighted Sum of Gray Gases method when the Gray Thermal Radiation model is selected for a participating multi-component gas continuum. However, this method considers only CO_2 and H_2O as participating gases.

Choose the desired solver type. For the segregated flow model:

- Select Segregated Flow in the *Flow* group.
- Select a model from the *Energy* group: Segregated Fluid Temperature, Segregated Fluid Isothermal, or Segregated Fluid Enthalpy.

After this selection, the *Physics Model Selection* dialog appears as shown in the following screen shot:



Alternatively, to choose the coupled flow model:

- Select Coupled Flow in the *Flow* group.
- The Coupled Energy model is selected automatically in the *Optional Models* group, provided you have the Auto-select recommended physics models checkbox activated (the default setting). Otherwise select it.

Setting Participating Media Radiation Model Parameters

In a participating media simulation, a typical tree structure appearing under the Models node of a given continuum is as shown below.



STAR-CCM+ currently simulates this type of problem with the Participating Media Radiation model, which uses the Discrete Ordinate model (see Basic Model Equations). A key reason for this is that the method offers itself readily to STAR-CCM+'s parallelization techniques.

The following parameters require user input:

- The desired degree of accuracy in simulating the participating media radiation effect (see What Is Participating Media Radiation?)
- The radiative properties of the medium lying between the radiating boundaries. In the current STAR-CCM+ version, the media can absorb, emit, and scatter isotropically. The absorption coefficient controls absorption/emission, and the scattering coefficient controls scattering. These objects use constant and field function subobjects. A third option is provided as the Weighted Sum of Gray Gases method when the Gray Thermal Radiation model is selected for a participating multi-component gas continuum. However, this method considers only CO_2 and H_2O as participating gases.
- The radiative properties of the boundary surfaces. These properties are defined in the same way as for surface-to-surface models (see Setting Thermal Radiation Model Parameters).

What Is the Absorption Coefficient?

The property of a medium that describes the amount of absorption of thermal radiation per unit path length within the medium [m^{-1}]. This value node (Absorption Coefficient in the following screenshot) is activated as soon as the Participating Media Radiation model is selected. It can be defined either as a constant value or as a field function. When you have gray radiation as your thermal spectrum specification in a multi-component mixture, the Weighted Sum of Gray Gases method is also available. For the Multiband spectrum model, values of this property can be specified for each spectral band.



The emission coefficient is equal to the absorption coefficient by Kirchhoff's law.

What Is the Scattering Coefficient?

The property of a medium that describes the amount of scattering of thermal radiation per unit path length within the medium [m^{-1}]. This value node (Scattering Coefficient in the following screenshot) is activated as soon

as the Participating Media Radiation model is selected. It can be defined either as a constant value or as a field function. For the Multiband spectrum model, values of this property can be specified for each spectral band.



What Is the Constant Method?

The Constant method sub-node of the absorption coefficient and scattering coefficient objects changes dynamically in response to your thermal spectrum specifications.



When you have gray radiation, it functions as a constant. With spectral radiation (multiband), it functions as a constant array, where you would enter one comma-separated value for each band.

What Is the Field Function Method?

The Field Function method sub-node of the absorption coefficient and scattering coefficient objects changes dynamically in response to your thermal spectrum specifications. When you have gray radiation, the node is a simple field function node that defines the gray radiative property. When you have spectral radiation (multiband), the node expands to display a set of field function sub-nodes, one for each spectral band. The sub-nodes are named for the respective bands (Thermal 1 and Thermal 2 in the example

below) and define the per-band radiative properties. In each case, the field function is specified by selecting the appropriate sub-node and choosing from the drop-down list in the Scalar Function property.



What Is the Weighted Sum of Gray Gases (WSGG) Method?

The Weighted Sum of Gray Gases method sub-node of the absorption coefficient object is available when you have gray radiation as your thermal spectrum specification in a multi-component mixture. This method calculates radiative properties from local gas species concentrations and

local temperatures without you specifying them (see formulation). You specify the path length in the Optical Path Length property. This sub-node is also subject to common menu items.



Weighted Sum of Gray Gases Method Properties

Optical Path Length	Set the optical path length according to the specific case. Where material properties do not vary over distance, you can use Eqn. (895) to get an estimate.

Setting Solver Parameters

This section discusses how to activate the DO Radiation solver and how to set various parameters that affect its performance.

Using the Solver

The DO Radiation solver is activated by default, after Participating Media Radiation has been selected in the *Physics Model Selection* dialog and a case geometry has been imported. The corresponding node in the simulation tree, DO Radiation, is shown in the following screenshot.



To tune the performance of the solver, you can alter a number of the properties that are displayed in the *Properties* window after selecting the DO Radiation node.

DO Radiation Solver Properties

Under-Relaxation Factor	Applies under-relaxation to the solution of the ordinate equations. The default value (1.0) means that no under-relaxation is used.
Maximum Iteration	Controls the maximum number of internal iterations that are required to solve the ordinate equations. The default value (50) is adequate in most scenarios.
Solver Tolerance	Defines the accuracy that is required in the iterative solution of the ordinate equations. The default value (1.0E-5) is adequate in most scenarios.
Verbosity	Provides extra output for this value while the simulation is running. This property can be useful for debugging problems when they occur. none — Leaves output at its default. low — Displays the convergence history for the DOM solution. high — (Currently displays the same output as low.)

The solver can be deactivated by selecting **Solver Frozen** in the *Expert* section of the *DO radiation - Properties* window. The Solver Frozen feature is primarily a "debugging" tool and is not supported for normal simulation work.

Update Frequency Node

The Update Frequency Node has no properties, but it can have a different subnode for each time model.

Steady Node Properties

Iteration Update Frequency	The radiation terms due to participating media are recalculated after the number of iterations, specified here, of the fluid flow equation solver. The default value of 1 recalculates these terms at the same frequency.

Explicit Unsteady Properties

Time-Step Update Frequency	Runs the solver after the number of time-steps entered. The default is 1.

Implicit Unsteady Properties

Time-Step Update Frequency	Runs the solver after the number of time-steps entered. The default is 1.
Inner Iteration Update Schedule Option	Selects the inner iteration on which the solver runs.

First Iteration	Runs the solver on the first iteration.
Last Iteration	Runs the solver on the last iteration.
Every Iteration	Runs the solver in every iteration.

Under-Relaxation Factor Ramp Node

Ramp Method	Activates or deactivates the ramp for calculation.

No Ramp	Leaves off ramp calculation.
Linear Ramp	Activates linear ramp value calculation. The Linear Ramp node becomes available in the Under-Relaxation Factor Ramp node.

Setting Algebraic Multigrid Parameters

To set the algebraic multigrid parameters for the DO Radiation solver, open the DO Radiation node under the Solvers node. Then, select the node for the AMG Linear Solver as shown below. The algebraic multigrid parameters can

be entered using the *Properties* window.



Participating Media Radiation Formulation

This section describes:

- The Basic Model Equations
- The manner of their Numerical Solution in STAR-CCM+
- Discrete ordinate quadrature
- The Weighted Sum of Gray Gases (WSGG) Method

Basic Model Equations

As radiation travels through a medium, the intervening material absorbs and increments its radiant intensity I in the Ω direction. The radiative transfer equation (RTE) governs this process. When this equation is written in terms of radiant intensity for a specific wavelength λ, it is:

$$\frac{dI_\lambda}{ds} = -\beta_\lambda I_\lambda + k_{a\lambda} I_{b\lambda} + \frac{k_{s\lambda}}{4\pi}\int_{4\pi} I_\lambda(\Omega)\, d(\Omega) + k_{pa\lambda} I_{pb\lambda} + \frac{k_{ps\lambda}}{4\pi}\int_{4\pi} I_\lambda(\Omega)\, d(\Omega) \tag{874}$$

and

$$I_{b\lambda} = \frac{2C_1}{\lambda^5 (e^{C_2/\lambda T} - 1)} \tag{875}$$

and $C_1 = 0.595522 \times 10^{-16}\ Wm^2/s$, $C_2 = 0.01439\ mK$.

The in-scattering component is assumed to be isotropic.

When the medium has several components, the absorption coefficient $k_{a\lambda}$ is

$$k_{a\lambda} = \sum_i k_{a\lambda i} \tag{876}$$

where i indexes the participating components, such as CO_2, H_2O, or soot.

When the absorption and scattering coefficients of the medium are independent of wavelength, the medium is called gray. In that case, the RTE can be integrated over wavelength (or, equivalently, wavenumber) to produce a wavelength-independent equation.

The particle radiation is valid for both gray and multiband radiation for continuous media. However, for the case of multiband (spectral) radiation, the particle properties are still gray. The surface and continuum properties can be spectral, but the particle properties are gray (the absorption efficiencies are independent of wavelength). Particle scattering is isotropic.

The boundary condition that is applied to the RTE for each wavelength λ is:

$$I_\lambda(\mathbf{s}) = \varepsilon_{w\lambda} I_{b\lambda} + \frac{\rho^d_{w\lambda}}{\pi}\int_{(\mathbf{n}\cdot\mathbf{s}'<0)} I_\lambda(\mathbf{s}')\,|\mathbf{n}\cdot\mathbf{s}'|\, d\Omega + \rho^s_{w\lambda} I_\lambda(\mathbf{s}^\mathbf{s}) + \tau_{w\lambda} I_\lambda(\mathbf{s}) \tag{877}$$

where $\varepsilon_{w\lambda}$ is the diffuse emissivity, $\rho^d_{w\lambda}$ is the diffuse reflectivity, $\rho^s_{w\lambda}$ is the specular reflectivity, and $\tau_{w\lambda}$ is the transmissivity. Only interfaces can be partially transmissive; all other external boundaries are modeled as opaque.

The incoming ray direction, $\mathbf{s}^\mathbf{s}$, is determined using the laws of reflection:

$$\mathbf{s}^\mathbf{s} = \mathbf{s} - 2(\mathbf{s}\cdot\mathbf{n})\mathbf{n} \tag{878}$$

The radiant heat flux in a particular direction $\mathbf{q}_r$ is calculated using the integration of the radiant intensity over all solid angles and over the wavelength spectrum:

$$\mathbf{q}_r(r) = \int_0^\infty \int_{4\pi} I_\lambda(\mathbf{s})\ \mathbf{s}\ d\Omega\ d\lambda \tag{879}$$

The radiation solution is coupled to the fluid dynamic solution through the divergence of the radiative heat flux. This term exchanges energy between the fluid and the radiant energy field. Given the intensity field, the divergence of the heat flux is computed as:

$$\nabla \cdot \mathbf{q}_r = \int_0^\infty k_{a\lambda} \left(4\pi\, I_{b\lambda} - \int_{4\pi} I_\lambda\, d\Omega \right) d\lambda \tag{880}$$

Numerical Solution

The discrete ordinates method solves field equations for radiation intensity that is associated with a fixed direction s, represented by a discrete solid angle such as the one shown in the figure below.



Solid angle displayed as part of a hemisphere

A solid angle is the three-dimensional analog of an ordinary angle and is measured in steradians $[sr]$. One steradian is the solid angle that is subtended at the center of a sphere by an area on the surface of the sphere equal to the radius squared. As such, the solid angle of a complete sphere is 4π steradians.

The discrete ordinates calculation requires specification of the number of solid angles or "ordinates" into which the sphere surrounding any particular point is discretized. Increasing the number of ordinates leads to greater accuracy. A detailed description of the method can be found in Modest [136] and Siegel and Howell [137]. As a result of Eqn. (874), the form of these ordinate equations (for each wavelength band) is:

$$\mathbf{s}_i \cdot \nabla I_{i\Delta\lambda} = -\beta_{\Delta\lambda} I_{i\Delta\lambda} + k_{a\Delta\lambda} I_{b\Delta\lambda} + \frac{k_{s\Delta\lambda}}{4\pi} \sum_{j=1}^{n} w_j I_{j\Delta\lambda} + \bar{k}_{pa\Delta\lambda} I_{pb\Delta\lambda} + \frac{\bar{k}_{ps\Delta\lambda}}{4\pi} \sum^{n} w_j I_{j\Delta\lambda} \tag{881}$$

In the above equation, $\Delta\lambda$ represents a wavelength band from λ_m to λ_n. The black body emission in this band is:

$$I_{b\Delta\lambda} = \int_0^{\lambda_n} I_{b\lambda}\, d\lambda - \int_0^{\lambda_m} I_{b\lambda}\, d\lambda \tag{882}$$

The equivalent particle absorption and scattering coefficients are given in [136] as:

$$\bar{k}_{pa\Delta\lambda} = \sum_i Q_{a,\,i} N_i \frac{\pi d_i^2}{4} \tag{883}$$

$$\bar{k}_{ps\Delta\lambda} = \sum_i Q_{s,\,i} N_i \frac{\pi d_i^2}{4} \tag{884}$$

where i indicates the parcel index number.

The transport equation for each ordinate direction is discretized and solved independently, using the same discretization and iterative solution practices as for the other transport equations. For this reason, global iterations are necessary to include the isotropic in-scattering terms in the RTE and to compute wall boundary conditions. The angularly discretized boundary conditions take the form:

$$I_{i\Delta\lambda,w} = \varepsilon_{w\Delta\lambda} I_{b\Delta\lambda} + \frac{\rho^d_{w\Delta\lambda}}{\pi_{eff}} \sum_{(\mathbf{n}\cdot\mathbf{s}'<0)} I_{j\Delta\lambda} \left|\mathbf{n}\cdot\mathbf{s}'_j\right|_w + f\rho^s_{w\Delta\lambda} I_{\Delta\lambda}(\mathbf{s}^s) + \tau_{w\Delta\lambda} I_{i\Delta\lambda} \tag{885}$$

where π_{eff} is the half moment for the chosen ordinate set, which is defined as:

$$\pi_{eff} = \sum_{\mathbf{n} \cdot \mathbf{s}' > 0} w_j (\mathbf{n} \cdot \mathbf{s}'_j)_w \tag{886}$$

and f is the correction factor that is used when the direction of specular reflection is not aligned with any discrete ordinate direction. The diffuse reflection term represents summation over incoming (incident) ordinate directions.

Convergence is measured using the infinity norm of the normalized difference in the incident energy from successive space-angle sweeps. You set the convergence tolerance as discussed in the section on Using the Solver.

The radiation solution provides a source term to the fluid dynamic energy equation as given by Eqn. (880). The discretized form of this source term at any cell is:

$$-\nabla \cdot \mathbf{q}_r = \sum_{\lambda} k_{a\Delta\lambda} \left(\sum_{j=1}^{n} w_j I_{j\Delta\lambda} - 4\pi I_{b\Delta\lambda} \right) \tag{887}$$

The net absorption rate of radiative energy by all particles within a given cell is:

$$-\nabla \cdot \mathbf{q}_{pr} = \sum_{\lambda} \bar{k}_{pa\Delta\lambda} \left(\sum_{j=1}^{n} w_j I_{j\Delta\lambda} - 4\pi I_{pb\Delta\lambda} \right) \tag{888}$$

For Lagrangian particle transport, at convergence this quantity exactly balances the net radiative energy that all particles in the cell absorb.

Discrete Ordinate Quadrature

The ordinate set that is required to solve Eqn. (881) with boundary conditions (Eqn. (885)) is obtained using the discrete ordinate quadrature scheme. General principles for DOM quadrature schemes [134] are:

- All discrete ordinates, $\Omega_m = (\zeta_m, \eta_m, \mu_m)$, where (ζ_m, η_m, μ_m) are the direction cosines, must be placed on the unit sphere, that is:

$$(\zeta_m^2 + \eta_m^2 + \mu_m^2) = 1 \tag{889}$$

- The weights that are associated with each ordinate direction must be all positive. This requirement ensures that the error of the quadrature scheme is minimized.

- The number of photons must be preserved which is equivalent to satisfy:

$$\sum_{m=1}^{M} w_m(\mathbf{n}\overline{\Omega_m}) = 0 \tag{890}$$

  where w_m is the weight that is associated with the ordinates and $\mathbf{n}$ is an arbitrary unit vector.

- The quadrature must be invariant to any rotation of the arrangement of the discrete directions around the center of the unit sphere.

Although the first three criteria can be easily fulfilled, the fourth is impossible to be satisfied for any arbitrary rotation by a finite number of ordinates. Hence all ordinate sets must satisfy at least zeroth, first and second moments. For accuracy in radiation calculations, the ordinate set with weights must also satisfy the first moment over a half range along the three principal directions. This expression is:

$$\sum_{\mathbf{n}\overline{\Omega_m} > 0} w_m(\mathbf{n}\overline{\Omega_m}) = \pi \tag{891}$$

This additional requirement comes from the following:

- Radiative intensity can have directional discontinuity at a wall.
- Important radiative heat fluxes at the walls are evaluated through a first moment of intensity over a half range of 2π.

Level Symmetric (Sn) Quadrature

Level symmetric, also known as Sn, quadrature is a quadrature scheme where the discrete ordinates are arranged on the latitudes on the unit sphere. The corresponding latitudes have the same distance in the x-direction, y-direction, and z-direction and therefore this quadrature scheme is called Level Symmetric quadrature. A typical arrangement is

illustrated in the figure (symmetric point arrangement for n = 6) [135].



The rotational invariance to the principal octahedron is satisfied implicitly because of the arrangement on latitudes. Due to the latitude-based arrangement, Sn quadrature can be used in solution of discrete ordinate equations both in three-dimensional and axisymmetric two-dimensional problems. Currently only Sn quadratures are available in STAR-CCM+. Sn quadratures can be formulated to satisfy either odd or even moments. To find the Sn formulation that satisfies a particular moment, see the following table. Sn(O) is the odd formulation and Sn(E) is the even formulation.

	Type	
Order	**Sn(O)**	**Sn(E)**
2	0, 1	0, 2
4	0, 1, 2	0, 2, 4

	Type	
6	0, 1, 2, 3	0, 2, 4, 6
8	0, 1, 2, 3, 4	0, 2, 4, 6, 8
12	0, 1, 2, 3, 4, 5, 6	0, 2, 4, 6, 8, 10, 12
16	0, 1, 2, 3, 4, 5, 6, 7, 8	0, 2, 4, 6, 8, 10, 12, 14, 16

All Sn quadratures $(N > 2)$ with odd formulation satisfy the primary requirements of satisfying zeroth, first and second moments whereas the even formulation does not. However odd formulation of Sn incurs negative weights for $(N \geq 10)$ which are, in general, not suitable due to inaccuracy. The even formulation of Sn does not have that problem except for N = 10. Hence, in STAR-CCM+ odd formulation is available for N = 2, 4, 6, and 8, and even formulation is available for N = 12 and 16. All of the above available Sn quadratures in STAR-CCM+ satisfy the half moment criteria along principal directions.

Weighted Sum of Gray Gases (WSGG) Method

Combustion product gases do not scatter radiation significantly, but are strong selective emitters and absorbers of radiation. CO_2 and H_2O gases dominate cloud emission and absorption among the combustion gas products. These species require that the spectral variation of radiative properties be taken into account.

The Weighted Sum of Gray Gases (WSGG) method for calculating the absorption coefficient is a reasonable compromise between the oversimplified user-provided gray gas model and a complete spectral property model which takes into account detailed absorption bands. This method uses the Hottel charts [131] and sum-of-gray-gases models [132] for a mixture containing CO_2 and/or H_2O gases only. Several curve-fitted sum-of-gray-gases expressions are available [130]. In the WSGG method, the medium is assumed to consist of different fractions of gray gases with different absorption coefficients.

The relationship between the total gray gas absorption coefficient κ and the total absorptivity α for a gray medium is the Bouguer-Lambert law as:

$$\alpha = 1 - e^{-\kappa s} \tag{892}$$

where s is the optical path length (OPL).

The total absorptivity is approximated with a weighted-sum of several gray gases as:

$$\alpha \approx \sum_{k=0}^{K} a_k (1 - e^{-\kappa_k s}) \tag{893}$$

where K is the total number of gray gases, a_k is the weight factor and κ_k is the absorption coefficient of each gray gas. Based on the quality of the curve-fit of Hottel charts [131] and sum-of-gray-gases models [132], typically $K = 3$ gives satisfactory accuracy. STAR-CCM+ always uses this value of K.

After computing α from Eqn. (893), the total gray gas absorption coefficient κ is obtained using the Bouger-Lambert Law (Eqn. (892)). κ is then used to solve the RTE as k_a.

The OPL is a user-provided quantity. The OPL can be used as a tuning parameter to get the expected and reasonable distribution for the absorption coefficient of the gas medium. The OPL value can range from the cell dimensions to the dimension of the entire computational domain, depending on the optical thickness of the combustion region. For instance, if the optical thickness of the combustion medium is high then the cell thickness can be used to define the OPL. However, for an optically thin medium, use a value that corresponds to the reactor dimension.

The OPL is related to optical thickness by:

$$\theta = s\kappa \tag{894}$$

where θ is the optical thickness, s is the OPL, and κ is the absorption coefficient.

Once a value of the OPL is prescribed, the WSGG method captures local variation of the absorption coefficient (which is due to variation in the gas temperature and partial pressures of combustion gases).

For arbitrary geometries, the OPL can be approximated as [133]:

$$s = 3.6 \ \text{(volume/surface area)} \tag{895}$$

where the volume is the total volume of the gas and the surface area is its total surface area.

Participating Media Radiation Nomenclature

In the equations that are given above, the following nomenclature applies:

I_λ	radiative intensity at wavelength λ $[W/(m^2 sr\ m^{-1})]$
$k_{a\lambda}$	absorption coefficient at wavelength λ (m^{-1})

$k_{s\lambda}$	scattering coefficient at wavelength λ (m^{-1})
$I_{b\lambda}$	black body intensity at wavelength λ
Ω	solid angle
$k_{pa\lambda}$	particle absorption coefficient at wavelength λ (m^{-1})
$k_{ps\lambda}$	particle scattering coefficient at wavelength λ (m^{-1})
$I_{pb\lambda}$	particle black body intensity at wavelength λ and current particle temperature.
β_λ	extinction coefficient, which is defined as $\beta_\lambda = k_{a\lambda} + k_{s\lambda} + k_{pa\lambda} + k_{ps\lambda}$
s	distance in the Ω direction
ε_w	boundary emissivity
θ	optical thickness
τ_w	interface boundary transmissivity
ρ_w^d	boundary diffuse reflectivity
ρ_w^s	boundary specular reflectivity
$\mathbf{s}$	unit vector in the Ω direction
$\boldsymbol{n}$	unit surface vector directed outwards from the boundary
$\mathbf{s}'$	unit vector along the distance coordinate leaving the boundary
$\mathbf{s}^s$	unit vector for the incoming ray that is considered for specular reflection
$\mathbf{q}_r$	radiant heat flux in a given direction
$\beta_{\Delta\lambda}$	$k_{a\Delta\lambda} + k_{s\Delta\lambda} + \bar{k}_{pa\Delta\lambda} + \bar{k}_{ps\Delta\lambda}$
w_j	quadrature weights that depend on the chosen ordinates
n	number of ordinates
T	temperature [K]
κ	total gray gas absorption coefficient
$\bar{k}_{pa\Delta\lambda}$	equivalent particle absorption coefficient
$\bar{k}_{ps\Delta\lambda}$	equivalent particle scattering coefficient
$Q_{a\Delta\lambda i}$	particle absorption efficiency

$Q_{s\Delta\lambda i}$	particle scattering efficiency
N_i	particle number density
d_i	particle diameter

Participating Media Radiation Bibliography

[130] Coppalle, A. and Vervisch, P. 1983. "The total emissivities of high-temperature flames," *Combustion and Flame*, **49**(101), pp. 101-108.

[131] Hottel, H. C. 1954. "Radiant-Heat Transmission," *Heat Transmission*, 3rd Ed., W. H. McAdams (Ed.), McGraw-Hill, NY.

[132] Hottel, H. C. and Sarofim, A. F. 1964. *Radiative Transfer*, McGraw-Hill, NY.

[133] Incropera, F. P. and DeWitt, D. P. 1990. *Fundamentals of Heat and Mass Transfer*, 3rd Ed., Wiley, NY.

[134] Koch, R. and Becker, R. 2003. "Evaluation of quadrature schemes for the discrete ordinate method", Proceedings of Eurotherm 73 on Computational Thermal Radiation in Participating Media, pp. 15-17.

[135] Lathrop, K. D. and Carlson, B. G. 1965. "Discrete ordinates angular quadrature of the neutron transport equation", Report LA-3186, Los Alamos Scientific Laboratory.

[136] Modest, M.F. 1993. "Radiative Heat Transfer." McGraw-Hill, New York.

[137] Siegel, R. and Howell, J.R. 1992. "Thermal Radiation Heat Transfer", Third Edition, Hemisphere Publishing Co., Washington.

Modeling Thermal Radiation

This section describes:

- What a Thermal radiation model is.
- Modeling Gray Thermal radiation.
- Modeling Multiband Thermal radiation.
- How to set the relevant parameters for the thermal radiation model.

What Is a Thermal Radiation Model?

A Thermal Radiation model is used to define how thermal radiation properties are modeled -- as either wavelength independent (gray) or as wavelength dependent (spectral). The choice pertains uniformly to all radiation within the regions of the simulation and in the surrounding environment. The Thermal Radiation model is responsible for managing the radiative properties and spectrum-dependent solution variables in the continuum.

The type of Thermal Radiation model can be either Gray or Multiband (spectral). The Gray and Multiband spectrum models are available for both Surface-to-Surface and Participating Media (DOM) transfer models.

Modeling Gray Thermal Radiation

This section describes:

- The Gray Thermal radiation model
- How to select it
- The Thermal Environment and how to set thermal environment properties

What Is the Gray Thermal Radiation Model?

The Gray Thermal radiation model makes it possible to simulate diffuse radiation independently of wavelength. With this model, radiation properties of the media and surrounding surfaces are considered the same for all wavelengths.

The corresponding node (Gray Thermal Radiation in the following screenshot) is activated as soon as a Gray Thermal model is selected.



The only property for the model is an expert property to deactivate Kirchhoff's law at surfaces.

Kirchhoff's law must always be activated for modeling radiative heat transfer. However, if the radiative transfer equation is used to model molecular transport by analogy, Kirchhoff's law can be deactivated since it does not apply for molecular transport. (For radiative transfer, Kirchhoff's law can be locally violated in rare, non-equilibrium conditions; however, these conditions are not important for CAE modeling.)

Gray Thermal Radiation Properties

None.

Gray Thermal Radiation Expert Properties

This property is available only when the Surface-to-Surface model or the Viewfactors Calculator model is activated.

Kirchhoff Enabled — Controls the enforcement of Kirchhoff's law (absorptivity equals emissivity) at surfaces. This property must be activated for radiative heat transfer, and can only be deactivated in the special case of using the Surface-to-Surface model to simulate molecular transport by physical analogy. Kirchhoff's law is enforced by requiring the sum of the emissivity, reflectivity, and transmissivity to be 1.0 at surfaces.

- Activated — Activates the enforcement of Kirchhoff's law (default).
- Deactivated — Deactivates the enforcement of Kirchhoff's law.

Selecting the Gray Thermal Radiation Model

To select the Gray Thermal radiation model:

- In the *Physics Model Selection* dialog, select a space model (Three Dimensional, Two Dimensional, or Axisymmetric) in the *Space* group.
- Select the Radiation model in the *Optional Models* group.
- Select one of the models (Surface-to-Surface (S2S) Radiation, Participating Media Radiation (DOM)) in the *Radiation* group. (A requirement for the S2S and DOM radiation models is that an energy model must be activated, so you must select material and flow models as well.)
- Select the Gray Thermal Radiation model in the *Radiation Spectrum* group.

Following the selection of these models, the *Physics Model Selection* dialog appears similar to the one shown below:



What Is the Thermal Environment?

When the Surface-to-Surface (S2S) Radiation model or the Viewfactors Calculator model is used, each continuum is required to have a thermal environment. This thermal environment is a simplified representation of the environment surrounding the continuum, from the standpoint of thermal radiation. The thermal environment is modeled as a black body with unity emissivity and, hence, can be solely characterized with the Radiation Temperature. This temperature defines the energy that is

effectively radiated from the environment. However, since there has to be only one environment, the value of the Radiation Temperature in every continuum must be the same.

Setting the Thermal Environment Properties

The Thermal Environment is activated with the Surface-to-Surface (S2S) Radiation model or the Viewfactors Calculator model, and is fully defined by ascribing the desired value to the Radiation Temperature. The Radiation Temperature is a scalar quantity which is selected by opening the Thermal Environment node within the Gray Thermal Radiation model as shown below. The temperature can then be set in the *Properties* window.



Modeling Multiband (Spectral) Thermal Radiation

This section describes:

- The Multiband Thermal radiation model
- How to select it
- The spectral bands and how to set them

What Is the Multiband Thermal Radiation Model?

The Multiband Thermal radiation model makes it possible to simulate diffuse and specular radiation that is dependent on wavelength or frequency. With this model, wavelength bands specify radiation properties of the media and surrounding surfaces. The number and extents of bands are chosen to describe property variations with respect to wavelength. During the solution process, the governing radiation equations are solved for each defined band using the specified properties. The total radiation solution is then obtained by integrating (summing) the solutions from each of the individual bands. Because the governing equations are solved for each band, the radiation solution time generally increases linearly with the number of bands. As a consequence, a trade-off often exists between accuracy and computational time.

Absorption is a phenomenon which describes energy being removed from the incident radiative energy by the material or boundary. Absorptivity is a property which is the ratio of the radiative energy that is absorbed to the total radiative energy incident on a boundary. For a surface, the absorptivity dictates absorption whereas for participating media, the absorption coefficient dictates absorption. From Kirchhoff's Law, the spectral directional emissivity is equal to the spectral directional absorptivity. In STAR-CCM+, directional properties are not used so Kirchhoff's Law is implemented as spectral emissivity is equal to the spectral absorptivity. So for each band in a multiband model the emissivity equals the absorptivity. For a gray model, the properties are independent of both direction and wavelength and emissivity equals absorptivity for the entire spectrum.

The Multiband Thermal Radiation is activated as soon as a Multiband Thermal model is selected.



The bands for the Multiband model are defined using the Spectral Bands sub-node.

Multiband Thermal Radiation Properties

None.

Selecting the Multiband Thermal Radiation Model

To select the Multiband Thermal radiation model:

- In the *Physics Model Selection* dialog, select a space model (Three Dimensional, Two Dimensional, or Axisymmetric) in the *Space* group.
- Select the Radiation model in the *Optional Models* group.
- Select one of the models (Surface-to-Surface (S2S) Radiation, Participating Media Radiation (DOM)) in the *Radiation* group.
- Select the Multiband Thermal Radiation model in the *Radiation Spectrum* group.

Following the selection of these models, the *Physics Model Selection* dialog appears similar to the one shown below:



What Are Spectral Bands?

When Multiband Thermal Radiation is activated, spectral bands are used to describe the discretization of the radiation spectrum. The spectrum is divided into a number of discrete bands, each defined by a lower and upper wavelength. Surface and media properties can be different in each band. The number of bands and their extents must be chosen to represent variations in radiation properties with respect to wavelength. However, as mentioned before, radiation solution time increases linearly with the number of bands so a trade-off often exists between accuracy and runtime.

Setting the Spectral Bands

The Spectral Band node is located under the Multiband Thermal Radiation node.



Individual bands are added to this manager node by right-clicking the node and then by selecting New.

After a new band is created, its lower and upper wavelength extents can be specified under the band properties.



The rules for defining bands are:

- For an individual band, the lower wavelength must be smaller than the upper wavelength.
- Bands cannot overlap one another.

  In general, the bands must be contiguous and span the relevant part of the wavelength domain for the characteristic temperatures of the analysis (see Planck's distribution in a standard radiation textbook). The relevant domain for thermal radiation is typically located between 0.1-1000 micro meters, but it can be narrower depending on the characteristic temperatures.
- If radiation is activated in more than one continuum, the set of bands in each continuum must be the same. This rule can be met by defining the spectral bands in one continuum and then copying the bands to any other continua.

Setting Thermal Radiation Model Parameters

This section deals with the parameters required for solving fluid flows in the presence of thermal radiation. Additional input is required to characterize the radiation properties of boundary surfaces. When Kirchhoff's law is enforced, radiation properties must sum up to one on every surface where they are required. (For virtually all applications, Kirchhoff's law must be enforced, and the Kirchhoff Enabled setting of the Gray Thermal Radiation model is on by default.)

The radiation properties are defined in the Regions node, for each boundary surface. Boundaries are classified either as *Closed*, *Open*, or *Inactive*. As summarized below, *Closed* boundaries require specification of emissivity, ε, reflectivity, ρ, and transmissivity, τ. In addition, *Open* boundaries require specification of the radiation temperature, T_{rad}. The radiation does not directly see *Inactive* boundaries, so they do not require any further specifications.

	Boundary Type			Specify			
Boundary Condition	**Closed**	**Open**	**Inactive**	ε	ρ	τ	T_{rad}
Inlet: mass flow, stagnation, velocity		X		X	X	X	X
Outlet: pressure, flow-split		X		X	X	X	X
Wall	X			X	X	X	
Symmetry plane			X				

Cyclic			X		
Baffle	X			X	X	X

To specify ε, ρ, τ, and T_{rad}:

- Open the Boundaries node.
- For each of the boundary sub-nodes, specify appropriate Physics Values for the relevant optical properties as shown in the following screenshot. However, for the Discrete Ordinate Method (DOM), τ can only be specified on baffles.



The emissivity and reflectivity of the interfaces must be specified at each interface boundary, not the original boundary (for example top-solid1 [In-place 1], not top-solid1). The latter becomes defunct when the interface is created. These properties must be set in regions on each side of the interface.

In addition, the interface transmissivity is specified on the interface itself instead of the two associated interface boundaries. While emissivity and reflectivity can be different on each side of the interface, the transmissivity must be the same and is thus specified at one location, on the interface itself. In order to satisfy both the second law of thermodynamics and the reciprocity constraint, the transmissivity must be the same for both sides.

At wall boundaries with the Surface-to-Surface (S2S) model, an environment condition can also be applied. This condition effectively acts as an interface with the environment and permits:

- Absorption, reflection, and emission from both sides.
- Radiative transmission and conduction through the wall.
- Separate wall temperatures for the inner and outer sides.

Radiative surface properties must be specified for both sides of the boundary when the environment condition is activated.

What Is Emissivity?

Emissivity is the ratio of the power that a body emits to the power it would emit as a black body at the same temperature (dimensionless). This value node (Surface Emissivity in the following screenshot) is activated as soon as a thermal radiation model is selected. It is entered as a scalar profile. For the Multiband spectrum model, values of this property can be specified for each spectral band.



What Is Reflectivity?

Reflectivity is the ratio of reflected radiant energy over incident radiant energy at a given surface (dimensionless). This value node (Surface Reflectivity in the following screenshot) is activated as soon as a radiation model is selected.

It can be entered as a scalar profile or calculated automatically using the Auto-Calculate method. The Auto-Calculate method is the easiest and preferred approach. When this method is selected, the reflectivity is calculated automatically from the specified emissivity and transmissivity using Kirchhoff's law (reflectivity = 1 - emissivity - transmissivity). Anytime afterwards, when you change the emissivity and/or transmissivity, the reflectivity is automatically recalculated. If the reflectivity is entered using a scalar profile, you are responsible for ensuring Kirchhoff's law is satisfied. The only time the Auto-Calculate method must not be used is in the special situation where Kirchhoff's law is explicitly deactivated (see Gray Thermal Radiation Expert Properties).

For the Multiband spectrum model, values of this property can be specified for each spectral band.



What Is Transmissivity?

Transmissivity is the ratio of transmitted radiant energy over incident radiant energy at a given surface (dimensionless). This value node (Surface Transmissivity in the following screenshot) is entered as a scalar profile. For the Multiband spectrum model, values of this property can be specified for each spectral band.



It is activated as soon as a radiation model is selected. For the Surface-to-Surface model, transmissivity can be specified at inlet, outlet, and wall boundaries. (For the Discrete Ordinates model, these types of boundaries cannot be transparent and transmissivity specification is not permitted on them.)

Transmissivity must be same on both sides of an interface. As a result, transmissivity appears under the Physics Values node of each interface node instead of under the Physics Values nodes of the associated interface boundaries. For transmissivity to appear under the interface, continua on both sides of the interface must have the same radiation model and spectrum model (gray or multiband) activated.

If radiation or spectrum models do not match on both sides of an interface, that interface does not have any transmissivity input and is treated as opaque. The only exception is the boundary of a volume region when the Fluid Film and Surface-to-Surface Radiation models are selected. In that case, the boundary on the volume region side of the fluid film interface has transmissivity input instead of the interface. Radiation is not active on the fluid film region side.

Transmissivity becomes available in both interface boundaries when a Baffle or Porous Baffle interface is created, provided a radiative transfer model (Surface-to-Surface Radiation or Participating Media Radiation) is selected in the regions on both sides of the interface.

What Is Radiation Temperature?

Radiation temperature is the temperature of an open boundary that has a total radiant emittance identical to that of a black body radiator (K). This value node (Radiation Temperature in the following screenshot) is activated as soon as a radiation model is selected. It is entered as a scalar profile.



Setting Specular Reflectivity

Surface reflectivity can be partially or fully specular (as opposed to diffuse). Specular reflectivity is activated by selecting `Adjustable reflection specularity` in the Specularity option property of the Physics Conditions > Specularity Option node for the relevant boundary. This setting adds the Reflection Specularity node to the physics values of the boundary.



Reflection Specularity

Reflection specularity is the specular fraction of the surface reflectivity:

- A value of 0 denotes that all of the reflection is diffuse.
- A value of 1 denotes that all of the reflection is specular.
- A value in-between denotes a combination of diffuse and specular reflection.

Reflection specularity is set in the Physics Values > Reflection Specularity node. For the Gray Thermal model, only one method can be specified for the entire model in the Method property. For the Multiband spectrum model,

you can select either `Composite` or `Constant` in the Method property. Selecting the `Composite` option allows you to set a separate method for each spectral band within the Composite folder.



Specular reflectivity is then determined from the product of reflection specularity and surface reflectivity.

A special case is that of a mirrored surface, which is represented by setting the surface reflectivity and the reflection specularity to 1. This setting means that the surface is completely reflective, and all of the reflection is specular. In this case, the specular reflectivity is also equal to 1. Alternatively, specular reflective surfaces can also be partially transmissive (such as plate glass, particularly low-emissivity glass) and/or absorptive (such as high-gloss-painted surface). In such cases, the surface reflectivity is less than 1, with surface transmissivity and/or surface emissivity making up the difference to satisfy Kirchhoff's law. If part of the reflection is diffuse, the reflection specularity can be set to 1 or less. Specular reflectivity here would thus be less than 1.

Modeling Solar Radiation Loads

Thermal radiative exchange with the solar environment includes both direct and diffuse components. The direct solar flux is the irradiation from the sun incident directly on the regions without interference from the atmosphere or other objects. Diffuse solar radiation results from the atmospheric scattering, and reflection and transmission from objects such as the ground.

The coordinate system for the orientation of the direct solar flux is defined in the following figure:



Solar Coordinate System

The X-axis is aligned with the local magnetic north direction and the Y-axis is aligned with the local west direction, that is, the XY plane is the local earth-surface plane. The Z-axis then points skyward since STAR-CCM+ employs right-handed coordinate systems.

Thermal radiative exchanges with the environment include only diffuse effects, where the environment surrounding the regions emits diffusively as a blackbody at the environment radiation temperature.

Solar radiative exchanges have a direct and a diffuse component. The direct solar flux is the irradiation from the sun impacting directly on the regions, without interference from the atmosphere or other objects. Diffuse solar radiation results from the atmospheric scattering, and reflection and transmission from objects such as the ground.

This section describes:

- What the Solar Loads model is
- How to select it
- Formulation for use of date-time-location

What Is the Solar Loads Model?

The Solar Loads model is used to model radiation loads from the sun -- both direct and diffuse. This optional model is available with Surface-to-Surface Radiation and works with either the Gray or Multiband spectrum models. The solar emissions are considered together with terrestrial emissions within the context of a given spectrum model. For the Gray spectrum model, the total solar loads are applied to the single transfer system representing the full-thermal spectrum. For the Multiband spectrum model, the solar loads are divided between spectral bands that are based on Planck's distribution for a blackbody at the effective surface temperature of the sun. These per-band solar emissions are then subject to the radiation surface properties for the associated band as are all other per-band emissions.

When the Solar Loads model is selected for a continuum, a Solar Loads node becomes visible under the Models node for that continuum.



Each continuum with radiation can include a Solar Loads model. However, the coordinate and the specification method (with its properties) for the Solar Loads model in every continuum must be the same. The assumption is that only one sun exists and all continua see this single sun.

Some changes to input parameters of the solar load specification method can result in a change in solar direction and fluxes. When this happens, solar load distribution on patches is immediately recalculated.

Solar Loads Model Properties

Coordinate System	Specifies the coordinate system (laboratory by default). The drop-down list of this property includes any local coordinate systems that exist in the simulation.	
Solar Specification Option	Selects the method for specifying solar values.	
	Manual specification	Activates the Manual Specification sub-node, in which you can specify directions and fluxes.
	Solar calculator	Activates the Solar Calculator sub-node, in which you can specify values such as date, time and geographical locations, and obtain calculations of directions and fluxes.

The solar emissions are subject to the radiative surface properties (see Setting Thermal Radiation Model Parameters) specified with the use of either the Gray or Multiband Thermal Radiation models. No additional radiation property specification on boundaries is needed when the Solar Loads model is enabled.

Selecting the Solar Loads Model

To select the Solar Loads model:

- In the *Physics Model Selection* dialog, select a space model (Three Dimensional) from the *Space* group. (Currently, Surface-to-Surface (S2S) Radiation is only available for three-dimensional geometries.)
- Select the Radiation model in the *Optional Models* group.
- Select Surface-to-Surface (S2S) Radiation in the *Radiation* group. (A requirement for the S2S Radiation model is that an energy model must be activated, so you must select material and flow models as well.)

  The Solar Loads model may now be selected from the *Solar Radiation* group.

Following the selection of these models, the *Physics Model Selection* dialog appears similar to the one shown below:



What Is Manual Specification?

The Solar Loads model is fully closed by defining the following properties of the Manual Specification node: the sun direction vector in a specified coordinate system, the direct solar flux and the diffuse solar flux.



The properties can also be defined with field function syntax. For example, you can define the orientation of the sun as a function of time by defining the Azimuth and Altitude properties with expressions such as `sin(100*$Time)`.

If the definitions of the angles and/or the direct solar flux are a function of time, the solar loads are recalculated automatically during each time-step.

Manual Specification Properties

Azimuth	Sets the sun direction vector horizontal component.
Altitude	Sets the sun direction vector vertical or elevation component.
Direct Solar Flux	Sets the solar radiative energy directly from the sun along the specified direction vector.
Diffuse Solar Flux	Sets the solar radiative energy from the surrounding environment.

What Is the Solar Calculator?

For a chosen date, time and geographical location, the Solar Calculator computes sun altitude (elevation) and azimuthal angles, along with the corresponding direct and diffuse solar fluxes. Select the Solar Calculator node to access the properties required by the tool.



Additional expert inputs are a sunshine factor to accommodate the clouds and a diffuse fraction to consider a fraction of solar radiation as diffuse (the rest is direct flux). The solar calculator automatically includes a correction due to refraction of solar radiation through the atmosphere of the earth.

For unsteady simulations, the solar calculator automatically refreshes solar values at every time-step.

The following screenshots are examples of scalar displays using results from the solar calculator.





Solar Calculator Properties

Latitude	Latitude in degrees (north positive, south negative; default 51.47917, location of Greenwich, United Kingdom).
Longitude	Longitude in degrees from Greenwich, United Kingdom (east positive, west negative; default 0.0).
Time Zone	Time that is offset with respect to UTC (east positive, west negative; default 0.0 hr).
Year	Starting year in the Gregorian calendar (default 2000).
Month	Starting month of the year.
Day	Starting day of the month.
Hour	Starting local time (24-hour format; default 12:00 noon).
Minute	Starting minute.
Azimuth	Sun direction vector horizontal component (read-only).
Altitude	Sun direction vector vertical or elevation component (read-only).
Direct Solar Flux	Solar radiative energy directly from the sun along the specified direction vector (read-only).
Diffuse Solar Flux	Solar radiative energy from the surrounding environment (read-only).
Current Solar Time	Starting solar time plus elapsed solar time (read-only).

Solar Calculator Expert Properties

Diffuse Fraction	Fraction of total solar flux to be considered as diffuse (default 0.0); remainder is direct flux.
Sunshine Factor	Factor by which total solar flux is reduced to take into account a cloudy sky (default 1.0).

Solar Calculator Formulation

This formulation presents equations for solar position and flux calculations in the solar calculator [138]. A bibliography is included.

All angles that are shown in this formulation are in degrees.

User Input Range for the Solar Calculator

The user inputs for various parameters must lie within the ranges that are shown in the following table [138].

Variable	Range
Year	1950-2050
Month	1-12
Day	1-31 (depending on the month)
Hour	0 – 23
Minute	0-59
Time zone	-12.0 -+12.0 (in hours)
Latitude	- 90.0-+90.0 (in degrees)
Longitude	-180.0-+180.0 (in degrees)
Fraction for diffuse flux	0.0-1.0
Sunshine factor	0.0-1.0

Calculation of Sun Direction

Solar altitude, azimuthal angles, and related geometric parameters can be calculated as:

- Day angle [139]:

$$\alpha_d = \frac{360°}{365}(d_{year} - 1) \tag{896}$$

where d_{year} is the day of the year.

- Earth radius vector [140]:

$$r_E = 1.000110 + 0.034221\cos(\alpha_d) + 0.001280\sin(\alpha_d) + 0.000719\cos(2\alpha_d) + 0.000077\sin(2\alpha_d) \tag{897}$$

Coordinated Universal Time (UTC) in hours:

$$UTC = hour + \frac{minute}{60} - timeZone \tag{898}$$

Julian day [141]:

$$d_{jul} = 32916.5 + 365(year - 1949) + \frac{year - 1949}{4} + d_{year} + \frac{UTC}{24} \tag{899}$$

Time that is used in calculation of elliptical coordinate [141]:

$$t_{ec} = d_{jul} - 51545.0 \tag{900}$$

Mean longitude [141]:

$$L_m = 280.460 + 0.9856474 t_{ec} \qquad (0° \le L_m \le 360°) \tag{901}$$

Mean anomaly [141]:

$$A_m = 357.528 + 0.9856003 t_{ec} \qquad (0° \le A_m \le 360°) \tag{902}$$

Elliptic longitude [141]:

$$L_e = L_m + 1.915\sin(A_m) + 0.020\sin(2A_m) \qquad (0° \le L_e \le 360°) \tag{903}$$

Obliquity of the elliptic [141]:

$$O_e = 23.439 - (4 \times 10^7)\, t_{ec} \tag{904}$$

Declination [141]:

$$D_c = \sin^{-1}\{\sin(O_e)\sin(L_e)\} \tag{905}$$

Right ascension [141]:

$$R_a = \tan^{-1}\left\{\frac{\cos(O_e)\sin(L_e)}{\cos(L_e)}\right\} \qquad (0° \le R_a \le 360°) \tag{906}$$

Greenwich mean sidereal time [141]:

$$t_{gmt} = 6.697375 + 0.0657098242\, t_{ec} + UTC \qquad (0 \le t_{gmt} \le 24) \tag{907}$$

Local mean sidereal time [141]:

$$L_{mst} = 15t_{gmt} + L_o \qquad (0° \le L_{mst} \le 360°) \tag{908}$$

where L_o is the user-defined longitude of the geographical location.

Hour angle [141]:

$$\beta_h = L_{mst} - R_a \qquad (-180° \le \beta_h \le 180°) \tag{909}$$

Solar zenith Z and altitude angles without refraction [139]:

$$Z = \cos^{-1}\{\sin(D_c)\sin(L_a) + \cos(D_c)\cos(L_a)\cos(\beta_h)\} \qquad (0° \le Z \le 99°) \tag{910}$$

where L_a is the user-defined latitude of the geographical location. Z is limited to 9 degrees below the horizon to compensate for refraction.

Altitude:

$$\theta = 90° - Z \tag{911}$$

Solar azimuthal angle [139]:

$$\phi = 180° - \cos^{-1}\left\{\frac{\sin(\theta)\sin(L_a) - \sin(D_c)}{\cos(\theta)\cos(L_a)}\right\} \qquad \text{if } \beta_h > 0°, \quad \phi = 360° - \phi \tag{912}$$

Application of Refraction Correction

The refraction correction [142] is applied as follows:

cf is the correction factor due to atmospheric refraction.

$$cf = 0 \qquad \text{if } \theta > 85° \tag{913}$$

$$cf = \frac{58.1}{\tan(\theta)} - \frac{0.07}{\tan^3(\theta)} + \frac{0.000086}{\tan^5(\theta)} \qquad \text{if } 5° \le \theta \le 85° \tag{914}$$

$$cf = 1735 + \theta[-518.2 + \theta\{103.4 + \theta(-12.79 + 0.711)\}] \qquad \text{if } -0.575° \le \theta < 5° \tag{915}$$

$$cf = \frac{20.774}{\tan(\theta)} \qquad \text{if } \theta < -0.575° \tag{916}$$

$$cf = \left(\frac{cf}{3600}\right)\left(\frac{283P}{1013(273 + T)}\right) \tag{917}$$

where P equals 1013 millibars atmospheric pressure and T equals 15°C mean atmospheric temperature.

Refraction corrected altitude:

$$\theta_r = \theta + cf \qquad (-9° \le \theta_r \le 90°) \tag{918}$$

Refraction corrected zenith:

$$Z_r = 90 - \theta_r \tag{919}$$

Calculation of Solar Fluxes

Total solar flux:

$$q = C_s r_E(\cos(Z_r))f_s \qquad \text{if } \cos(Z_r) > 0 \tag{920}$$

$$q = 0 \qquad \text{if } \cos(Z_r) \le 0 \tag{921}$$

where C_s is the solar constant, equal to 1376.0 W/m^2, and f_s is the sunshine factor.

Diffuse flux:

$$q_{diff} = qf_{diff} \tag{922}$$

where f_{diff} is the diffuse factor.

Direct flux:

$$q_{dir} = q(1.0 - f_{diff}) \tag{923}$$

Solar Calculator Bibliography

[138] SOLPOS – Solar Calculator, National Renewable Energy Laboratory, http://rredc.nrel.gov/solar/codesandalgorithms/solpos/.

[139] Iqbal, M. 1983. *An Introduction to Solar Radiation*, Academic Press, NY.

[140] Spencer, J. W. 1971. “Fourier series representation of the position of the sun”, *Search* **2**(5), p. 172.

[141] Michalsky, J. 1988. “The Astronomical Almanac's algorithm for approximate solar position (1950-2050)”, *Solar Energy* **40**(3), pp. 227-235.

[142] Zimmerman, J. C. 1981. “Sun-pointing programs and their accuracy”, SAND81-0761, Experimental Systems Operation Division 4721, Sandia National Laboratories, Albuquerque, NM.

Modeling View Factors

This section describes:

- What the View Factors Calculator model is
- How to select it
- How to set the relevant parameters for both the View Factors Calculator model and its corresponding solver
- How to generate patches and view factors
- How to work with view factor data

View factors can only be calculated on three-dimensional meshes.

What Is the View Factors Calculator Model?

Selecting the View Factors Calculator model enables solvers for generating radiation patches and calculating view factors. The radiation patches and view factors are required for solving surface-to-surface radiation (see Modeling Surface-to-Surface Radiation), and is to be calculated automatically before the S2S solution when the View Factors Calculator model is enabled (see Generating Patches and View Factors).

The corresponding node (View Factors Calculator in the following screenshot) is activated as soon as the View Factors Calculator model is selected.



Selecting the View Factors Calculator Model

To select the View Factors Calculator model:

- In the *Physics Model Selection* dialog, select Three Dimensional from the *Space* group. (View factors can only be calculated on three-dimensional meshes.)
- Select the Radiation model in the *Optional Models* group.
- Select Surface-to-Surface (S2S) Radiation in the *Radiation* group.

  A requirement for using the S2S Radiation model is that an energy model must be activated, so select material and flow models as well.

  The View Factors Calculator option now becomes available under the *Radiation* group.

When you select the View Factors Calculator option, the *Physics Model Selection* dialog appears as shown below.



Setting the View Factors Calculator Model Parameters

Selecting the View Factors Calculator model enables both patch generation and view factors calculation.

Patching Parameters

The Patching feature provides you with two control parameters, the number of patches and the angularity. The control parameters (options and values) are defined by default on a region-wise basis in each region where the View Factors Calculator model is active under the Physics Conditions and Physics Values nodes.

The Patch Specification method is set in the *Properties* window.



Patch Specification Properties

Patch Specification	Sets the number of patches specification method	
	Proportion of Faces	Sets the number of patches as a proportion per face.
	Total	Sets the total number of patches.

Depending on the patch specification method chosen (Proportion of Faces or Total), the respective physics values for the region are shown below.





The number of patches and patch angularity may be further refined or customized at individual boundaries under the Physics Conditions node. Alternatively, you can let these settings be handled automatically for two boundaries by creating an internal, periodic interface between them. It carries the full transmissivity of the originating patch from one parent boundary of the interface to the other.

Number of Patches

You can specify the number of patches you would like on each relevant boundary. The number of patches may be selected as a Total Number, or, by default, as a percentage of the number of faces on the boundaries. 100% would mean that every boundary face has a distinct patch. As the specified

percentage decreases, the patch size increases until one patch is left to cover the entire boundary. It is not possible to ascribe zero patches on a boundary that requires patches.



Custom Patch Specification Properties

Custom Patch Specification	Specifies whether the number of patches is set to region or custom values.	
	Use Region Values	Uses the values set for the region.
	Proportion of Faces	Allows customization of patch angularity for the individual boundary. If this option is selected, physics values defining patches for the boundary are enabled and may be set.
	Total	Allows customization of patch angularity for the individual boundary. If this option is selected, physics values defining patches for the boundary are enabled and may be set.

Angularity

The latter parameter defines the maximum allowed angle between the normals to contiguous boundary faces for these faces to belong to the same patch. It provides a control over the patch planarity and is used to detect sharp angles in the geometry, such as would happen at a corner, and to

prevent patches from bridging such features. The angularity is set at 150 degrees by default and would require changing only in cases where specific geometric features must be captured.



Custom Patch Angularity Specification Properties

Custom Patch Angularity Specification	Specifies whether patch angularity is set to region or custom values.	
	Use Region Values	Uses the values set for the region.
	Customize Locally	Allows customization of patch angularity for the individual boundary. If this option is selected, physics values defining patches for the boundary are enabled and may be set.

View Factors Parameters

The view factors are calculated by shooting energy beams from each patch and computing the amount of energy that is deposited at intercepting patches, which is a function of the optical properties of the surface. For diffuse reflection, the view factors calculator only needs to keep track of irradiated and transmitted energy and would formally only require specification of transmissivity at the boundary. However, for completeness, and because the individual spectrum models enforce Kirchhoff's law, for Thermal spectral bands, emissivity, reflectivity and transmissivity are provided at boundaries.

In the case of transparent or partially transparent external walls, a black body, whose temperature is characterized by an environmental temperature, may absorb the beams which are transmitted out of the enclosure.

These radiative properties of the boundary surfaces, the Surface Emissivity, Surface Reflectivity and Surface Transmissivity nodes below, are defined in the same way as for thermal models (see Setting Thermal Radiation Model Parameters).



Setting View Factors Calculator Solver Parameters

This section discusses how to control the View Factors Calculator solver and how to set various parameters that affect its performance. The View Factors Calculator solver is activated by default after the View Factors Calculator model has been selected in the *Physics Model Selection* dialog and a case geometry has been imported.

You access the solver parameters by selecting the View Factors Calculator node under the Solvers node.



The solver can be deactivated by activating the Solver Frozen option in the properties. This feature should not be your first choice in most situations, as it is primarily a "debugging" tool and not supported for normal simulation work.

Patches and view factors are computed automatically at the beginning of the solution if they do not exist. However, they may also be generated manually using creation commands that are located in a pop-up menu that is accessed by right-clicking the View Factors Calculator node. The creation commands simply provide added flexibility to generate patches and view factors manually before initiating the solution. Several other useful commands are also available from the pop-up menu.

You can also control the recalculation of view factors for rigid body motion.

View Factors Calculator Solver Properties

Verbosity	Provides additional output during the calculation of the view factors.	
	None	No output.
	Low	Displays ray tracing progress and reciprocity convergence.
	High	Displays additional progress and resource usage information.

Expert Properties

Update Frequency	Controls the frequency with which the solver calculates view factors. Use this only for cases with mesh motion that is caused by rigid body motion.
Number of Beams	The number of beams affects accuracy, view factor calculation time, and view factor storage requirements. A higher number of beams increases the accuracy of the view factors. However, increasing the number also increases the view factor calculation time and the memory requirements. The default number (1024) is usually a good balance between accuracy and resource use.
Maximum Reciprocity Iterations	Sets the maximum number of iterations in the reciprocity procedure. (Reciprocity is the constraint on view factors that is imposed by the second part of Eqn. (852).)
Reciprocity Tolerance	Specifies the criterion for convergence that is used in the reciprocity procedure. (Reciprocity is the constraint on view factors that is imposed by the second part of Eqn. (852).)
Maximum Polygons per Voxel	Controls the resolution of the voxel mesh that is used in ray tracing. Lowering the value increases the resolution of the voxel mesh. This property generally should not need to be changed from the default value of 50. Increasing the number slightly decreases memory requirements, but also may slightly increase the calculation time.
Solver Frozen	Controls whether the View Factors Calculator solver is active during iteration. Currently this feature is supported for "debugging" purposes only. Activated — Freezes this solver. Deactivated — Leaves this feature active.

View Factors Calculator Solver Pop-Up Menu

The following commands are accessed by right-clicking on the View Factors Calculator node.

Generate Radiation Patches	Adds radiation patches to the simulation.
Calculate View Factors	Generates the view factors. If patches do not exist, this command generates them.
Delete View Factors	Deletes all view factors in the simulation.
Display View Factor summary	Writes a summary of the view factor data in the *Output* window.
Export ASCII View Factors	Exports the view factor data to an ASCII file in your current working directory.

See also: common menu items.

Controlling View Factor Recalculation for Rigid Body Motion

In cases where rigid body motion is present, view factors are recalculated every time-step by default. For cases where some reduction in accuracy is acceptable, you can save on computational time by recalculating the view factors less frequently. The expert property, Update Frequency, provides this control.



This update frequency is designed to control view factor recalculation only for the case of mesh motion that results from rigid body motions. For all other cases (such as addition or removal of cells, reconstruction of interfaces, and morphing motion), the view factor solver behaves similar to having the update frequency equal 1. If mesh motion occurs due to rigid body motion, view factors are recalculated based on the input value of the update frequency. Within the update frequency interval, the S2S solver uses old view factors. The maximum allowable update frequency is 100.

The update frequency is related to the time-steps, not the inner iterations for the unsteady solver. This approach represents a compromise between accuracy and the computational cost of view factor recalculation.

Note: Indiscriminate use of this feature can lead to prominent loss of accuracy.

Generating Patches and View Factors

Radiation patches and view factors can be generated automatically or manually.

Patches must be generated before view factors may be calculated. As the validity of the view factors depends on interfaces that were previously created, the Calculate View Factors pop-up menu action of the View Factors Calculator solver intersects the interfaces before the actual calculation begins.

STAR-CCM+ generates radiation patches automatically in the following ways:

- As a result of the Initialize Solution command
- After a mesh motion (unless the frequency of recalculation has been changed)
- After modifications of Patching parameters
- After a boundary type has been changed to or from symmetry or periodic
- After radiation has been activated or deactivated in a continuum

View factors are calculated automatically as a result of the following:

- In the first iteration during the solving process
- After a mesh motion (recalculating)
- After a boundary type has been changed to or from symmetry or periodic
- After radiation has been activated or deactivated in a continuum
- After modifications of transmissivities (recalculating)
- After modifications of the sun position for the Solar Radiation model (recalculating)
- After the number of wavebands has been changed (either when switching from gray to multiband or when changing the number of bands for multiband)

Manual creation is possible after the View Factors Calculator model is selected. The commands Generate Radiation Patches and Calculate View Factors are available by right-clicking on the View Factors Calculator solver.



Working with View Factor Data

There are various options for working with view factor data, including displaying, exporting, and deleting.

To display view factor data, right-click on the node of the View Factors Calculator solver and select Display View Factor Summary.



The data appears in the *Output* window.



To export the data to ASCII format, select the respective option from the pop-up menu of the solver.



This option generates a file in your current working directory.

To delete the view factor data, select the Delete View Factors option.



Radiation Field Functions Reference

The following primitive field functions are made available to the simulation when a radiation transfer model (S2S or DOM) is used.

Boundary Radiation Heat Flux	The net radiative heat flux to the boundary (absorption minus emission). A positive value indicates a net absorption by the boundary. This property is a full-spectrum quantity. For spectral radiation (multiband), it includes contributions from all wavelengths (bands).
Boundary Radiosity	The radiosity at the boundary. The radiosity is essentially the radiative heat flux leaving a surface. This property is a full-spectrum quantity. For spectral radiation (multiband), it includes contributions from all wavelengths (bands).
Boundary Radiosity of <band-name>	The radiosity for band <band-name>. Only available for spectral radiation (multiband).
Boundary Irradiation	The irradiation at the boundary. The irradiation is the radiative heat flux incident on a surface. This property is a full-spectrum quantity. For spectral radiation (multiband), it includes contributions from all wavelengths (bands).
Boundary Irradiation of <band-name>	The irradiation for band <band-name>. Only available for spectral radiation (multiband).
Boundary Emissivity	The gray emissivity at the boundary. Only available for gray radiation.
Boundary Emissivity of <band-name>	The spectral emissivity at the boundary for band <band-name>. Only available for spectral radiation (multiband).

Boundary Reflectivity	The gray reflectivity at the boundary. Only available for gray radiation.
Boundary Reflectivity of <band-name>	The spectral reflectivity at the boundary for band <band-name>. Only available for spectral radiation (multiband).
Boundary Transmissivity	The gray transmissivity at the boundary. Only available for gray radiation.
Boundary Transmissivity of <band-name>	The spectral transmissivity at the boundary for band <band-name>. Only available for spectral radiation (multiband).
Total Emissivity	The blackbody-emissive-power-weighted emissivity at the boundary. For gray radiation, this property is simply the gray surface emissivity. For spectral radiation, this property is defined as: $\varepsilon_{total} = \left(\sum_{i \in bands} \varepsilon_i E_{b,i} \right) / \left(\sum_{i \in bands} E_{b,i} \right)$ where E_b is the blackbody emissive power and the subscript i denotes the band.

The following primitive field functions are additionally made available when the Surface-to-Surface (S2S) Radiation model is enabled.

Direct Solar Irradiation	The direct solar flux on the boundary. This quantity is only available for gray radiation when the Solar Loads model is enabled.
Direct Solar Irradiation of <band-name>	The direct solar flux on the boundary for band <band-name>. This quantity is only available for spectral radiation (multiband) when the Solar Loads model is enabled.

The following primitive field functions are additionally made available when the Participating Media Radiation (DOM) model is enabled.

Incident Radiation	The zeroth angular moment of radiative intensity ($G = \int_{4\pi} I d\Omega$). This quantity is also proportionally related to the energy density ($u_e = G/c$ where c is the speed of light) and to the mean intensity ($\bar{I} = G/4\pi$). This property is a full-spectrum quantity. For spectral radiation (multiband), it includes contributions from all wavelengths (bands).
Radiative Energy Source	The volumetric source to the energy equation due to radiative absorption and emission (absorption minus emission). A positive value indicates net absorption by the media. This quantity is equal to the negative of the divergence of the radiative heat flux. This field function is available for continuous media only.
Particle Radiative Energy Source	The volumetric (integrated over dispersed media in a given cell) source due to radiative absorption and emission by dispersed media, such as particles (absorption by particles minus emission from particles). A positive value indicates net absorption by the dispersed media. This quantity is equal to the negative of the divergence of the radiative heat flux from dispersed media.
Radiative Absorption	The volumetric absorption by the media. This quantity is the absorption part of the Radiative Energy Source.

Total Absorption Coefficient

The blackbody-emissive-power-weighted absorption coefficient within the domain. For gray radiation, this property is simply the gray absorption coefficient. For spectral radiation, it is defined as:

$$\kappa_{total} = \left(\sum_{i \in bands} \kappa_i E_{b,i} \right) / \left(\sum_{i \in bands} E_{b,i} \right)$$

where E_b is the blackbody emissive power and the subscript i denotes the band.

Radiative Intensity *j*

The radiative intensity for the *j*th ordinate direction. For spectral radiation (multiband), this variable contains the intensity field for the last band that was solved (it is not a full-spectrum quantity).

Absorption Coefficient

The absorption coefficient for the media. For spectral radiation (multiband), each band has a separate variable which is differentiated by the band index, for example Absorption Coefficient 0 for the first band.

Scattering Coefficient

The scattering coefficient for the media. For spectral radiation (multiband), each band has a separate variable which is differentiated by the band index, for example Scattering Coefficient 0 for the first band.

Modeling Aeroacoustics

STAR-CCM+ offers a choice of models for simulating acoustics in liquids and gases. To date, the primary focus has been on aeroacoustics, the study of noise from flows of gases as they interact with surfaces. However, STAR-CCM+ also allows you to model hydroacoustics, the study of noise transmission in water.

The following general types of acoustics models are available in STAR-CCM+:

- Broadband Noise Source
- APE (Acoustic Perturbation Equations) based noise source
- Direct Noise Simulation
- Ffowcs Williams-Hawkings
- Lighthill Stress Tensor

What is Aeroacoustics

Aeroacoustics investigates the aerodynamic generation of sound. Modern aeroacoustics science was initiated in the 1950's by Sir James Lighthill who derived a theory for the estimation of the intensity of sound that radiates from a turbulent flow. Before this, flow-generated noise studies were focused on the relation between the frequency of the fluid fluctuations and the emitted sound.

The frequency of the aeolian tone was mathematically related to flow parameters by Strouhal who, as a result of quantitative observations, introduced the non-dimensional frequency - the Strouhal number. Lord

Rayleigh related the Strouhal number to the flow Reynolds number and further recognized that the direction of most intense sound emission coincides with the direction of the fluctuating force on the wire.

In 1952, Lighthill established the theoretical background generally referred to when investigating aerodynamic noise [172]. Lighthill first introduced the concept of aeroacoustic analogy which consists of replacing the actual flow field responsible for generating noise with an equivalent system of noise sources. The noise sources act on a uniform stagnant fluid that is governed using standard acoustic propagation equations. The aerodynamic characterization of the sources then becomes the main issue in noise prediction.

In 1955, Curle [156] extended the concepts that Lighthill developed to include the effect of flow-body interaction on sound generation. In 1969, Ffowcs Williams and Hawkings [161] extended the Lighthill analogy to account for arbitrary surface motion. This formulation was used for noise prediction in rotor blades aerodynamics, such as in the helicopter and turbine industries [151] and [157].

The main mechanisms of sound generation in the presence of solid structures can be classified as [184]:

- **Vortex shedding noise**

  Vorticity that is released from a bluff body in a flow. The time varying circulation on the body due to vortex shedding induces a fluctuating force on the body itself which is transferred to the fluid and propagates as sound.

- **Turbulence-structure interaction noise**

  Vortical structures impinging on a solid surface generate local pressure fluctuations on the body surface which feed the acoustic far field.

- **Trailing edge noise**

  This type of noise is important for all rotating blade technologies, due to the interaction of the boundary layers instabilities with the surface edges.

In current engineering practice, computations of sheared and free shear flows using the unsteady compressible Navier-Stokes equations are performed routinely. It remains challenging, however, to resolve the acoustic field of those simulations whose acoustic energy is many orders of magnitude smaller than the energy of the hydrodynamic field. There are four sources of difficulty in directly resolving the acoustic waves using Navier-Stokes computations (from [143]):

- The acoustic-field is larger than the flow-field.
- The acoustic-field has a smaller energy content than the flow field.

- The numerical discretization may act as a more significant source of sound than the simulated flow-field.
- The difficulty of imposing free space boundary conditions appropriate for acoustics in the far-field, at an artificial computational boundary positioned at a finite distance away from the source region.

Artificial boundaries must be appropriate to ensure non-reflection of acoustic waves, but also in certain regions to provide inflow and outflow of the aerodynamic field.

The aeroacoustics models in STAR-CCM+ are responsible for providing aeroacoustics equations. The Aeroacoustics model node represents the selection of aeroacoustics models.



This chapter also provides overviews of:

- Selecting an aeroacoustics model
- How to conduct an aeroacoustics analysis
- Guidelines for aeroacoustics calculations

How Do I Conduct an Aeroacoustics Analysis?

STAR-CCM+ offers a hierarchy of models to conduct an aeroacoustics analysis.



The physics in aeroacoustics are inherently transient and ultimately, must be solved using transient calculations. The computational run times that are associated with such simulations are high. As such transient simulations take a long time to run, optimize the mesh and settings so that you only capture noise sources and frequencies that are relevant to your analysis. An initial steady state analysis provides you with the data for optimizing the mesh and settings.

In steady state analysis, there are two forms of aeroacoustic modeling available:

The first category relates to models that provide aeroacoustic source information through correlations. These models are based on the assumptions implicit to RANS turbulence modeling. They are termed

Broadband Noise Source Models and can be attributed to the classical aeroacoustic categorization of dipoles and quadrupoles in shear and jet flows. They are strictly applied to turbulence-generated flow-noise--they do not apply to noise associated with large-scale flow features, such as vortex shedding. For turbulence generated noise, they will:

- Identify the location of flow generated aeroacoustic sources.
- Approximate the associated sound power or dB level.

They provide a reasonable basis on which to:

- Compare one component design with another.
- Decide where to apply mesh refinement ahead of a well-qualified transient calculation.

You can further qualify the suitability of the mesh ahead of a transient calculation by analyzing the frequency that the mesh can resolve. The suitability analysis determines where, and by how much, the mesh must be refined to capture up to the desired frequency. The resolvable frequency is deduced from the mesh frequency cutoff scalar function. If, for example, the function value shows 1 kHz and you want to resolve up to 2 kHz, then refine the mesh a factor of two.

The second category relates to synthesization of turbulent perturbations from the steady state using *Synthesized Fluctuations SNGR*. The modeled turbulence kinetic energy and length scales are used to create synthetic fluctuations in pressure and velocity through assumption of the turbulent energy spectrum and wave number decomposition. The fluctuations may then be output directly for use in acoustics propagation codes.

After you have optimized the steady state model for an aeroacoustics simulation, you can use it to initialize the transient simulation.

Recommended Analysis Procedure

It is recommended that aeroacoustics simulations are set up as three-dimensional cases, using compressible flow rather than constant density, and run:

1. Set up and run the case in steady state mode using a relatively coarse mesh to optimize the case for aeroacoustics. Use the broadband noise source models to identify where to refine the mesh. Use the mesh frequency cut-off scalar function to judge the required size of the cells in the source region.
2. Load the case setup in Step 1. Refine the mesh using information from the broadband noise models and the mesh frequency cut-off scalar function. Run the case again in steady state mode, restarting the solution from that of Step 1.

3. Load the case setup in Step 2 and modify it to be a transient simulation with an appropriate time-step. Choose LES-based or DES-based turbulence modeling and restart the solution from that of Step 2.

   To be able to post-process Lighthill Stress Tensor field functions, select the Lighthill Stress Tensor.

   For far field propagation analysis, use the Ffowcs Williams-Hawkings (FWH) model.

4. Run the simulation.

Best-Practice Guidelines

See Guidelines for Aeroacoustics Calculations for guidance on how best to apply the aeroacoustics models to your applications.

Using Broadband Noise Source Models

In near-field noise prediction, use broadband noise source models to compute the location and strength of the main noise sources. The noise sources consist of a volume distribution of quadrupole sources and a surface distribution of dipole sources. STAR-CCM+ offers a group of models capable of evaluating the broadband sources of noise. The selection of this group is represented in the object tree by the Broadband Noise Sources node.



In the frequency domain, a broadband noise has a continuous spectrum, where the acoustic energy is continuously distributed at all frequencies in a given range. These models can compute the location and strength of the main sources of noise that are generated aerodynamically. They use flow field quantities from Reynolds-Averaged Navier-Stokes (RANS) equations. Because these broadband noise models use RANS simulations, the required computational time is small.

In STAR-CCM+, the following broadband noise source models estimate the local acoustic power:

For dipole-generated sound, generated per unit surface, use:

- Curle noise source model

For quadrupole-generated sound, generated per unit volume, use:

- Goldstein Axisymmetric noise source model
- Proudman noise source model

For sound that the turbulent shear flow generates, use:

- The Lilley noise source model
- The Linearized Euler Equation (LEE) noise source model

For suppressing spurious reflections, use:

- The acoustic suppression zone model

You can use the numerically predicted loudness rankings to characterize and compare different designs.

This section also includes the following:

- Noise source model properties
- Selecting a noise source model
- Working with mesh frequency cutoff
- Formulation

Selecting an Aeroacoustics Model

To select the Aeroacoustics model:

- In the *Physics Model Selection* dialog, select a flow model from the *Flow* group (Coupled Flow or Segregated Flow). The Aeroacoustics radio button becomes available.
- Select the Aeroacoustics model in the *Optional Models* group.

Following its selection, the *Physics Model Selection* dialog shows the available aeroacoustics models.



- The Lighthill Stress Tensor is available with all turbulence models.
- The Ffowcs Williams-Hawkings model is available with both steady and implicit unsteady models.
- The Acoustic Wave model is available for three-dimensional unsteady simulations, and with any viscous regime.
- The Broadband Noise Source models are available as follows:

- Curle, Goldstein, and Proudman models can be used in steady and unsteady (transient) analyses. These models are compatible with all Reynolds-Averaged Navier-Stokes (RANS) models that can provide turbulence time and length scales; therefore they cannot be used with Spalart-Allmaras.

Linearized Euler Equation (LEE) and Lilley Noise Source Model are available only for steady-state simulations. These models can be used with Reynolds-Averaged Navier-Stokes (RANS) cases that include turbulence time and length scales.

What Is the Curle Noise Source Model?

The Curle Noise source model evaluates the noise from a turbulent boundary layer flow over a solid body at low Mach number. The model represents dipole sources of noise, that is, noise coming from the fluctuating surface pressure from solid boundaries acting on the fluid. Most system noise is generated for low-Mach number flows. Specifically, this model computes the surface acoustic power to evaluate the local contribution to the total acoustic power per unit area of the body surface. By contrast, the Goldstein and Proudman models evaluate the acoustic power on the volume, with quadrupole-generated sound sources. This model assumes isotropic turbulence.

The Curle model has been used for air handling subsystem (AHS) applications, including rotating parts (such as blowers), heat exchangers, mufflers, distribution ducts and wall-enclosed spaces.

Like the Goldstein and Proudman models, this model can be used in steady and unsteady (transient) analyses. It is compatible with all Reynolds-Averaged Navier-Stokes (RANS) models that can provide turbulence time and length scales; therefore it cannot be used with Spalart-Allmaras. This model is preferable for steady-state solutions. It is available for coupled and segregated flow. The CPU time is comparable to the times of the Goldstein and Proudman models, and calculations are done at the end of each iteration (or time-step).

This model also activates field functions.

Curle Field Functions Reference

The following primitive field functions are made available to the simulation when the Curle noise source model is used. These objects all have the same base properties.

Curle Surface Acoustic Power	The local surface acoustic power. Computed with Eqn. (930).
Curle Surface Acoustic Power DB	The local surface acoustic power in dB units. Computed with Eqn. (931).

What Is the Goldstein Axisymmetric Noise Source Model?

The Goldstein model is based on Lighthill's theory for predicting aerodynamic noise from a turbulent shear flow on an assumption of locally homogeneous and anisotropic turbulence. Based on Goldstein's theory [166], this model is a generalization of the model Ribner [182] developed for axisymmetric turbulent flow. It eliminates the Ribner assumption of the space-time correlations of Reynolds stresses.

This model is used to compute the local noise contribution per unit volume for anisotropic turbulence in axisymmetric turbulent jets.

In STAR-CCM+ this model is available for all cells from a two-dimensional axisymmetric case. It can be used in either a steady or transient simulation, with either coupled or segregated flow. Although it is preferable for steady simulations, the Goldstein model could be used for an unsteady simulation in which the correct flow field for convergent steady-state simulations is not available. It is compatible with any of the Reynolds-Averaged Navier-Stokes (RANS) turbulence models except for Spalart-Allmaras.

This model also activates field functions.

Goldstein Field Functions Reference

The following primitive field functions are made available to the simulation when the Goldstein Axisymmetric Jet noise source model is used. These objects all have the same base properties.

Goldstein Acoustic Power	The acoustic power per unit volume, which can be reported using dimensional units (W/m^3), defined in Eqn. (932).

Goldstein Acoustic Power DB	The acoustic power per unit volume, which can be reported using dB, defined in Eqn. (942).

What Is the Proudman Noise Source Model?

Like the Goldstein model, the Proudman noise source model evaluates acoustic power per unit volume, and the sound is from quadrupoles. Specifically, this model computes the acoustic power to evaluate the local contribution to the total acoustic power per unit volume from the turbulent flow. The Proudman model assumes isotropic turbulence.

The Proudman model can be used for quadrupole sources of noise, such as the areas around blower blades. This model has been used for air handling subsystem (AHS) applications, including rotating parts (such as blowers), heat exchangers and distribution ducts.

Like the Curle and Goldstein models, this model can be used in steady and unsteady (transient) analyses. It is compatible with all Reynolds-Averaged Navier-Stokes (RANS) models that can provide turbulence time and length scales; therefore it cannot be used with Spalart-Allmaras. This model is preferable for steady-state solutions. It is available for coupled and segregated flow. The CPU time taken to run the Proudman model is comparable to time required for either the Curle and Goldstein models, and calculations are done at the end of each iteration (or time-step).

This model also activates field functions.

Proudman Field Functions Reference

The following primitive field functions are made available to the simulation when the Proudman noise source model is used. These objects all have the same base properties.

Proudman Acoustic Power	The local acoustic power that is computed with Eqn. (963).
Proudman Acoustic Power DB	The local acoustic power in dB that is computed with Eqn. (965).

What Is the Lilley Noise Source Model?

The Lilley equation computes the turbulent shear flow noise from the fluctuating velocity flow field. Specifically, the Lilley Noise Source model computes three terms, as field functions:

- Lilley Self-Noise Source
- Lilley Shear-Noise Source
- Lilley Total-Noise Source

Applications include problems where local shear effects dominate, such as flow around solid bodies and shear flows in channels. This model evaluates the location and strength of aeroacoustic sound sources.

This model is available only for steady-state simulations. Use it with Reynolds-Averaged Navier-Stokes (RANS) cases that include turbulence time and length scales.

The Lilley Noise Source model is compatible with the coupled and segregated flow models. Computational costs (specifically CPU time) depend on input of two properties: Number of Realizations and Number of Fourier Modes.

The computation of the Lilley field functions involves a separate procedure with the Lilley Source solver after the simulation run is finished. Right-click the Lilley Source Solver node and select Execute solver.



The field function calculations are available for post-processing.

What Is the Lilley Source Solver?

The Lilley noise source model activates the Lilley source solver. The purpose of this solver is to control the iterative update for the calculation. The Lilley Source solver first calls the Stochastic Noise Generation and

Radiation (SNGR) solver, then it computes the Lilley Noise Source Terms. The SNGR solver computes the synthetic turbulent velocity field and gradients of that field using the SNGR method.

Although this solver has no properties, it provides a pop-up menu action to calculate the noise source values for the Lilley model.

Lilley Source Solver Pop-Up Menu

Execute solver	Controls the Lilley Source solver. This action calculates values for the Lilley field functions.

Stopping the Lilley Source Solver

To stop the solver immediately, without completing the set number of realization steps, click on the solver stop button located next to the progress bar. The solver stops and saves the partially computed field functions.



Lilley Field Functions Reference

The following primitive field functions are made available to the simulation when the Lilley noise source model is used. These objects all have the same base properties.

Lilley Self Noise Source	The Lilley source term that computes the effect of noise of the turbulent velocity components, see Eqn. (955).
Lilley Shear Noise Source	The source term which computes the noise from the interaction between the mean shear flow and turbulent velocity components, see Eqn. (956).
Lilley Total Noise Source	The summation of the Lilley Self-Noise Source and Lilley Shear Noise Source.

What Is the Linearized Euler Equation (LEE) Noise Source Model?

The Linearized Euler Equations (LEE) [147] compute the turbulent shear flow noise from the fluctuating velocity flow field. The LEE account for refraction and convection effects in any sheared mean flows. The equations can be used for a wider range of conditions than the Lilley equation, where the associated source term is a nonlinear function of the fluctuating velocity flow field.

The LEE Noise Source Model is compatible with both the coupled and segregated flow models. Computational costs (specifically CPU time) depend on the properties Number of Realizations and Number of Fourier Modes.

This model is available only for steady-state simulations. It activates the following field functions:

- LEE Self Noise Source
- LEE Shear Noise Source
- LEE Total Noise Source

Executing the LEE Source Solver

To run the LEE Source solver and obtain the field functions:

- Obtain a converged steady state solution, or a converged transient time-step.
- Right-click the LEE Source Solver node and select Execute solver.



The field function calculations are available for post-processing.

What Is the LEE Source Solver?

The LEE noise source model activates the LEE Source solver. Its primary function is to control the iterative update for the calculation.

The Lee Source solver first calls the Stochastic Noise Generation and Radiation (SNGR) solver, then computes the LEE Noise Source Terms. The SNGR solver computes the synthetic turbulent velocity field and gradients of that field using the SNGR method. Although the LEE Source solver has no properties, it provides a pop-up menu from which to run the solver.

LEE Source Solver Pop-Up Menu

Execute solver	Controls the LEE Source solver. This action calculates values for the LEE noise source model's field functions.

Stopping the LEE Source Solver

To stop the solver immediately, without completing the set number of realization steps, click on the solver stop button located next to the progress bar. The solver stops and saves the partially computed field functions.



LEE Field Functions Reference

The following primitive field functions are made available to the simulation when the LEE noise source model is used. These objects all have the same base properties.

LEE Self Noise Source	The LEE source term that computes the sound generated by turbulent velocity components. See Eqn. (947).
LEE Shear Noise Source	The LEE source term that computes the acoustic effects from interaction between the mean shear flow and turbulent velocity components. See Eqn. (943).

LEE Total Noise Source	Summarizes the LEE Self Noise Source and LEE Shear Noise Source.

What is the Acoustic Suppression Zone Model?

Use the Acoustic Suppression Zone model to suppress spurious reflective disturbances in unsteady aeroacoustic simulations.

Running unsteady flow simulations produces time-dependent variations in the flow. These flow disturbances convect with the flow (for vortical disturbances) or propagate at the speed of sound (for acoustic disturbances) outward to exterior boundaries.

If the applied exterior boundary types have reflective properties, these disturbances can be spuriously reflected back into the domain. This reflection affects the flow solution, causing problems for acoustic-type calculations.

To satisfy the requirements for good quality aeroacoustic calculations, the numerical algorithm must allow these pressure waves and vortices to travel through the boundary smoothly. In other words, the algorithm must give the boundary quasi-non-reflective properties.

For example, a pressure outlet with non-reflecting properties allows the defined constant pressure to represent an average pressure at the outlet boundary. At the same time, it allows the flow pressure at the boundary to vary and disturbances to pass out of the domain without reflecting off the boundary.

A practical approach to constructing a non-reflective boundary treatment is to complement the respective characteristics-based boundary conditions with an Acoustics Suppression Zone treatment.

Use the Acoustic Suppression Zone when:

- Boundary disturbances are creating spurious acoustic effects.
- The spurious acoustic effects are propagating upstream.
- There is enough distance (at least several wavelengths) adjacent to the boundary in which to suppress the spurious acoustic effects.

The following screenshot shows an instantaneous pressure field of sound radiated by a monopole source at 6800 Hz in a 1.6 m square domain. The simulation ran using the coupled solver. An Acoustic Suppression Zone 0.3 m deep was applied to the boundaries of the domain to suppress echoes.



The Acoustic Suppression Zone is compatible with both the coupled flow solver and the segregated flow solver.

Expert users can use the Acoustic Suppression Zone model to model complex far-field conditions (for example, the simulation of wind gust effects on an airfoil), beyond the boundary conditions documented here. See Kim *et al.* [170].

Acoustic Suppression Zone Model Properties

Verbose	Displays the operations performed to compute Continuity, Momentum, or Energy source.	
	Activated	Displays operations
	Deactivated	Does not display operations. This is the default.
Continuity source	Specifies $-\sigma(\rho - \rho_{ref})$ as a damping term. See Eqn. (967).	
	Activated	Applies acoustic damping to the Continuity source. This is the default.
	Deactivated	Does not use Continuity source damping term.
Momentum source	Specifies $-\sigma(\rho v - \rho_{ref} v_{ref})$ as a damping term. See Eqn. (967).	
	Activated	Uses the Momentum source damping term. This is the default.
	Deactivated	Does not use Momentum source damping term.
Energy source	Specifies $-\sigma(\rho E - \rho_{ref} E_{ref})$ as a damping term. See Eqn. (967). You are advised to avoid using the Energy source in conjunction with the segregated solver.	
	Activated	Uses the Energy source damping term. This is the default.
	Deactivated	Does not use Energy source damping term.
Flow Conditions Option	Specifies the flow conditions within the Acoustic Suppression Zone.	

	Specified Per-Region	Adds three nodes under Physics Values for the region: Acoustic Suppression Zone Pressure, Acoustic Suppression Zone Velocity, and Acoustic Suppression Zone Temperature. In each of these nodes, you can specify the values by Constant, Field Function, Table, or User Code. This method is the default.
	Computed Average	The relaxation coefficient is defined as the average per boundary of pressure, temperature, or velocity, for all boundaries using Acoustic Suppression.
Relaxation Coefficient Option	Specifies the relaxation coefficient σ used in Eqn. (967).	
	Use Boundary Distance	Computes the relaxation coefficient σ, based on the boundary distance from the selected boundaries and the maximum specified value for the relaxation coefficient. Adds a node under the Acoustic Suppression Zone model node, called Computed Relaxation Coefficient, which lets you specify Maximum Relaxation Coefficient (default value, 0) and Acoustic Suppression Zone Thickness (default value, 0). This is the default.
	Specify Per-Region	Specifies the relaxation coefficient σ for any selection of regions, positioning the Acoustic Suppression Zone anywhere in the domain. Adds a node under Physics Values for each region, the Acoustic Suppression Zone Relaxation Coefficient node. In each of these nodes, you can specify the values by Constant, Field Function, Table, or User Code.

When Flow Conditions Option is set to Specified Per-Region, set the values for p, v, and T to the local flow mean values of the unsteady flow field. The model assumes that a statistically steady-state flow is achieved before starting to accumulate the flow-mean values.

The Acoustic Suppression Zone model is unlikely to work properly with ill-posed Flow Conditions values, or with aggressive or even discontinuous distributions (in space and time) of the relaxation coefficient σ. It is best that the Suppression Zone be at least 5-7 (representative) wave-lengths wide and that the value of Maximum Relaxation Coefficient be increased gradually and limited so that it does not introduce numerical instabilities in the simulation.

Selecting the Acoustic Suppression Zone Model

To select the Acoustic Suppression Zone model, first select a flow model and an unsteady model. This model is optional, and is compatible with compressible ideal gas and real gas models.

Noise Source Model Properties

The Curle, Goldstein and Proudman noise source models share the following properties, except where otherwise noted.

Alpha	The alpha coefficient, used by the Proudman model only.
Far-Field Density	The density of the far-field fluid. The default value is 1.225 kg/m^3.
Far-Field Sound Speed	The speed of sound in the far-field. See Eqn. (924) The default value is 340 m/s.

Reference Acoustic Power	An input to a model-specific acoustic power equation: If the Curle model is selected, the acoustic power per unit surface area is calculated in decibels using Eqn. (930). The default reference acoustic power is 1.0E-12 W/m^2. If the Goldstein model is selected, the acoustic power per unit volume is calculated in decibels using Eqn. (942). The default reference acoustic power is 1.0E-12 W/m^3. If the Proudman model is selected, the acoustic power per unit volume is calculated in decibels using Eqn. (965). The default reference acoustic power is 1.0E-12 W/m^3.

The LEE and Lilley noise source models share the following properties.

Number of Realizations	The number of samples used to compute the average of the synthetic turbulent velocity field based on the Stochastic Noise Generation and Radiation (SNGR) approach. The default value is 200. You can use up to 1000 realizations (see Karweit *et al.* [169]).
Number of Fourier Modes	The number of the Fourier modes used for the spectral resolution for isotropic turbulence given by a von Karman-Pao spectrum, when computing the synthetic turbulent velocity field (see n from Eqn. (958)). The default value is 50. Use values between 30 and 50 Fourier modes.

Selecting a Noise Source Model

To select a noise source model, first select the Reynolds-Averaged Navier-Stokes (RANS) turbulence modeling approach along with the Aeroacoustics model. The Broadband Noise Sources option is available in the *Aeroacoustics Models* group box.



Once you select this option, the *Broadband Noise Source Models* group box appears, with the noise source models available.



These models are:

- Curle
- LEE
- Lilley
- Proudman

If you have selected Axisymmetric as your option for space modeling, the Goldstein model is also in this group.



Working with Mesh Frequency Cutoff

The mesh frequency cutoff scalar function can be used to qualify the suitability of the mesh ahead of a transient calculation by analyzing the frequency that the mesh can resolve. This analysis determines where and by

how much to refine the mesh to capture up to the desired frequency. If, for example, the function value shows 1 kHz and you want to resolve up to 2 kHz, then refine the mesh by a factor of two.

The mesh frequency cutoff field function becomes available when the Curle, Proudman and Goldstein broadband noise source model are activated.

Create a scalar scene to display the mesh frequency cutoff contour at a desired location and select Mesh Frequency Cutoff Hz for the scalar field function.



The following screenshots provide an example of working with the mesh frequency cutoff scalar function. See the Flow Generated Noise Around a Cylinder tutorial for further details.



The above plot shows the mesh frequency cutoff contour for a coarse mesh. This plot indicates that the mesh is able to resolve up to 2 kHz (approximately) in the main volume source regions indicated by the Proudman correlation (in this example). Typically, however, the peak of the

human hearing range is between 1-3 kHz. Therefore, refine the mesh further to increase the resolvable frequency. The plot that is shown in the following screenshot shows the same scalar function with the refined mesh.



From the above plot, it is clear that the fine mesh adequately resolves higher frequencies close to the cylinder.

Broadband Noise Source and Mesh Frequency Cutoff Formulation

This section contains formulation for the broadband noise source models as well as the mesh frequency cutoff.

- Speed of Sound
- Curle
- Goldstein
- Linearized Euler Equations (LEE)
- Lilley
- Proudman

- Mesh Frequency Cutoff
- Acoustic Suppression Zone

Speed of Sound

The far-field speed of sound a_0 is the distance a sound wave travels in a unit of time.

$$a_0 = \sqrt{(\gamma R T_0)} \tag{924}$$

where:

γ is the adiabatic index or the ratio of heat capacity at constant pressure to the heat capacity at constant volume for the fluid.

R is the universal gas constant

T_0 is the absolute temperature, measured in Kelvin.

In dry air at 293K, the speed of sound is 343.2 m/s.

Curle

The Curle boundary layer model is the Curle integral [156] based on Lighthill's theory, which determined the equations of sound propagation in a medium at rest from the equation of continuity and momentum.

The Curle surface integral computes the sound generated by dipoles representing the fluctuating surface pressure with which the solid boundaries act on the fluid. The model evaluates the far field noise emitted by turbulent boundary layer flow over a solid body at low Mach number.

The Curle surface integral is:

$$\rho'(\grave{x}, t) = \frac{1}{(4\pi a_0^3)} \int_S \left[\frac{(\grave{x} - \grave{y})}{r^2} \frac{\partial p}{\partial t}\left(\grave{y}, t - \frac{r}{a_0}\right) \right] \grave{n} dS(\grave{y}) \tag{925}$$

where:

$t - r/a_0$ is the emission time

p the surface pressure

ρ' the acoustic pressure

a_0 the far-field sound speed.

On the assumption of small perturbations and an adiabatic problem, then:

$$\frac{p}{\rho^\gamma} = ct \tag{926}$$

which can be used to relate variations in acoustic pressure with density perturbations:

$$p' = a_0^2 \rho' \tag{927}$$

Then Eqn. (925) becomes:

$$p'(\grave{x}, t) = \frac{1}{(4\pi a_0)} \int_S \left[\frac{(\grave{x} - \grave{y})}{r^2} \frac{\partial p}{\partial t}\left(\grave{y}, t - \frac{r}{a_0}\right) \right] \grave{n} dS(\grave{y}) \tag{928}$$

The acoustical directional intensity per unit surface of the solid body on the far field prediction is approximated with:

$$\overline{p'^2} \approx \frac{1}{16\pi^2 a_0^2} \int_S \frac{(\cos\theta)^2}{r^2} \overline{\left[\frac{\partial p}{\partial t}\left(\grave{y}, t - \frac{r}{a_0}\right) \right]^2} A_c(\grave{y}) dS(\grave{y}) \tag{929}$$

where:

A_c is the correlation area

$r = (\grave{x} - \grave{y})$

θ is the angle between r and the $\grave{n}$ wall-normal direction.

The measure of the local contribution to acoustic power per unit surface area can be computed from:

$$SAP = \frac{1}{\rho_0 a_0} \left[\int_0^{(2\pi)} \int_0^{\pi} \overline{p'^2} r^2 \sin\theta \, d\theta \, d\gamma \right] = \int_S I(\grave{y}) dS(\grave{y}) \tag{930}$$

$$= \int_S \frac{A_c(\grave{y})}{(12\rho_0 \pi a_0^3)} \overline{\left(\frac{\partial p}{\partial t}\right)^2} dS(\grave{y})$$

where $I(\grave{y})$ is the directional acoustic intensity per unit surface.

The model can be enabled for steady and unsteady cases with Reynolds-Averaged Navier-Stokes (RANS) models which can provide turbulence time scale, turbulence length scale and wall shear stress necessary to compute $\overline{(\partial p/(\partial t))^2}$, the mean-square time derivative of the source surface pressure [167].

The acoustic power per unit surface can be reported in dimensional units ($\mathrm{W/m^2}$) and in dB:

$$SAP(\mathrm{dB}) = 10\log\frac{SAP}{P_{ref}} \tag{931}$$

where P_{ref} is the reference acoustic power.

Goldstein

In Goldstein's model, the local contribution to the total acoustic power emitted per unit volume of an axisymmetric turbulent jet, for cases where the axis of symmetry is along the direction of flow, is computed from:

$$G_{AP}(\mathbf{y}) = \int_0^{(2\pi)} \int_0^{(\pi)} I(r, \theta/\mathbf{y}) r^2 (\sin\theta)(d\theta) d\gamma = (2\pi) r^2 \int_0^{(\pi)} I(r, \theta/\mathbf{y})(\sin\theta)(d\theta) \tag{932}$$

where r and θ are the radial and angular coordinates of the receiver location, with the receiver assumed to be sufficiently far away from the flow. $I(r, \theta/\mathbf{y})$ is the directional acoustic intensity per unit volume of a jet:

$$I(r, \theta/\mathbf{y}) = \frac{12\rho_0 \omega_f^4 L_1 L_2^2 \left(\overline{u_{tl}^2}\right)^2}{5\pi a_0^5 r^2} \frac{D_{\text{self}}}{C^5} + \frac{24\rho_0 \omega_f^4 L_1 L_2^4 \left(\overline{u_{tl}^2}\right)}{\pi a_0^5 r^2} \left(\frac{\partial U}{\partial r}\right)^2 \frac{D_{\text{shear}}}{C^5} \tag{933}$$

C is the modified convection factor:

$$C = 1 - M_c(\cos\theta) \tag{934}$$

The decomposition of the sound intensity into shear and self noise (with no coupling terms) occurs under the assumption that turbulence is locally homogeneous:

$$I(r, \theta/\mathbf{y}) = I_{\text{self}}(r, \theta/\mathbf{y}) + I_{\text{shear}}(r, \theta/\mathbf{y}) \tag{935}$$

$$I_{\text{self}}(r, \theta/\mathbf{y}) = \frac{12\rho_0 \omega_f^4 L_1 L_2^2 (\overline{u_{tl}^2})^2 D_{\text{self}}}{5\pi a_0^5 r^2 C^5} \tag{936}$$

$$I_{\text{shear}}(r, \theta/\mathbf{y}) = \frac{24\rho_0 \omega_f^4 L_1 L_2^4 (\overline{u_{tl}^2}) v_{\text{grad}}^2 D_{\text{shear}}}{\pi a_0^5 r^2 C^5} \tag{937}$$

where ρ_0 is the far-field density and a_0 is the far-field speed of sound.

$$D_{\text{self}} = 1 + 2\left(\frac{M}{9} - N\right)(\cos^2\theta)(\sin^2\theta) + \frac{1}{3}\left(\left[\frac{M^2}{7} + M - 1.5N\left(3 - 3N + \frac{1.5}{\Delta^2} - \frac{\Delta^2}{2}\right)\right](\sin^4\theta)\right) \tag{938}$$

$$D_{\text{shear}} = (\cos^2\theta)\left[(\cos^2\theta) + \frac{1}{2}\left(\frac{1}{\Delta^2} - 2N\right)(\sin^2\theta)\right] \tag{939}$$

The effect of the anisotropic structure of the turbulence appears through the following parameters:

$$M = \left[\frac{3}{2}\left(\Delta - \frac{1}{\Delta}\right)\right]^2 \quad , \quad N = 1 - \frac{\left(\overline{u_{t2}^2}\right)}{\left(\overline{u_{t1}^2}\right)} \quad , \quad \Delta = \frac{L_2}{L_1} \tag{940}$$

$$L_1 = \frac{\left(\overline{u_{t1}^2}\right)^{\frac{3}{2}} T^3}{L^2} \quad , \quad L_2 = \frac{\left(\overline{u_{t2}^2}\right)^{\frac{3}{2}} T^3}{L^2} \quad , \quad \omega_f = 2\pi\left(\frac{1}{T}\right)$$

where L_1 and L_2 are the longitudinal and transversal correlation lengths, which are not equal. The longitudinal and transversal turbulence velocities $\overline{u_{t1}^2}$ and $\overline{u_{t2}^2}$ are not equal in all directions, so $N \neq 0$ and $\Delta \neq 1$. In the limiting case of isotropic turbulence the directivity expressions become identical to those of the Ribner model, so $N = M = 0$ and $\Delta = 1$.

The longitudinal and transversal turbulence velocities $\overline{u_{t1}^2}$ and $\overline{u_{t2}^2}$ are computed using the following relations:

$$\overline{u_{t1}^2} = \frac{8}{9}\left(\frac{L^2}{T^2}\right) \quad , \quad \overline{u_{t2}^2} = \frac{4}{9}\left(\frac{L^2}{T^2}\right) \tag{941}$$

where T is the turbulence time scale and L is the turbulence length scale.

In STAR-CCM+ the total acoustic power per unit volume can be reported in dimensional units ($\mathrm{W/m^3}$) and in dB:

$$G_{AP}(\mathrm{dB}) = 10\log\left(\frac{G_{AP}}{P_{ref}}\right) \tag{942}$$

where P_{ref} is the reference acoustic power.

See also Bechara *et al.* [147].

Linearized Euler Equations (LEE)

The linearized form of the Euler equations [147], derived from the non-linear Euler equations, can be achieved by superimposing small perturbations of turbulent and acoustic components on the mean flow as follows:

$$\begin{aligned} p_i(\vec{x}, t) &= P_i(\vec{x}) + p_i'(\vec{x}, t) + p_{ai}(\vec{x}, t) \\ u_i(\vec{x}, t) &= U_i(\vec{x}) + u_i'(\vec{x}, t) + u_{ai}(\vec{x}, t) \\ \rho_i(\vec{x}, t) &= \overline{\rho}_i(\vec{x}) + \rho_i'(\vec{x}, t) + \rho_{ai}(\vec{x}, t) \end{aligned} \tag{943}$$

where the acoustic components p_{ai}, u_{ai}, ρ_{ai} are small compared to the mean components P_i, U_i, $\bar{\rho}_i$ and the turbulent components p_i', u_i', ρ_i'.

Including the decomposition of variables in Eqn. (943), and considering the assumption that the acoustic perturbations are isentropic:

$$p_a = a_0^2 \rho_a \tag{944}$$

and that turbulent density fluctuations (ρ') can be neglected, then the Linearized Euler equations for velocity components become:

$$\begin{aligned} S &= \frac{\partial u_{ai}}{\partial t} + U_j \frac{\partial u_{ai}}{\partial x_j} + \frac{1}{\bar{\rho}} \frac{\partial p_a}{\partial x_i} + u_{aj} \frac{\partial U_i}{\partial x_j} - \frac{\rho_a}{\bar{\rho}^2 a_0^2} \frac{\partial P}{\partial x_i} \\ &= -U_j \frac{\partial u_i'}{\partial x_j} - u_j' \frac{\partial U_i}{\partial x_j} - u_j' \frac{\partial u_i'}{\partial x_j} - \frac{1}{\bar{\rho}} \frac{\partial p'}{\partial x_i} - \frac{\partial u_i'}{\partial t} + \frac{\partial}{\partial x_j} (\overline{u_j' u_i'}) \end{aligned} \tag{945}$$

The second set of terms includes the source terms responsible for sound generation.

The decomposition of the main sources of sound generated aerodynamically can be done onto the shear noise term $S_{\text{ShearNoiseTerm}}$ (representing the interaction between the mean shear velocities and turbulent velocity components of the flow) and the self noise term $S_{\text{SelfNoiseTerm}}$ (representing turbulence-turbulence interactions):

$$S_{\text{ShearNoiseTerm}} = -U_j \frac{\partial u_i'}{\partial x_j} - u_j' \frac{\partial U_i}{\partial x_j} \tag{946}$$

$$S_{\text{SelfNoiseTerm}} = -u_j' \frac{\partial u_i'}{\partial x_j} \tag{947}$$

The turbulent velocity components are synthesized by the Stochastic Noise Generation and Radiation model.

The source terms in the Linearized Euler Equations, similar to source terms in the Lilley Equation, are available for steady-state solutions computed using a Reynolds-Averaged Navier-Stokes (RANS) model that can provide turbulence length scale and turbulence time scale.

Lilley

Lilley [174] derived a third-order wave equation from conservation of mass and momentum of compressible flows. The Lilley source terms are available for steady-state solutions computed using a Reynolds-Averaged Navier-Stokes (RANS) model. A RANS model is the only type that can provide turbulence length scale and turbulent time scale dominated by local shear effects, as is the case of flow around solid bodies and shear flows in

channels. With this model STAR-CCM+ provides the capability to evaluate the location and strength of aeroacoustic sound sources. The Lilley equation can be written as:

$$\frac{D}{Dt}\left[\frac{D^2\Pi}{Dt^2} - \frac{\partial}{\partial x_j}\left(a^2\frac{\partial\Pi}{\partial x_j}\right)\right] + 2\left[\frac{\partial u_k}{\partial x_j}\frac{\partial}{\partial x_k}\left(a^2\frac{\partial\Pi}{\partial x_j}\right)\right] = -2\left[\frac{\partial u_k}{\partial x_i}\frac{\partial u_j}{\partial x_k}\frac{\partial u_i}{\partial x_j}\right] + \psi \tag{948}$$

$$\psi = 2\left[\frac{\partial u_j}{\partial x_i}\frac{\partial}{\partial x_j}\frac{1}{\rho}\frac{\partial\tau_{ik}}{\partial x_k}\right] - \left[\frac{D}{Dt}\frac{\partial}{\partial x_i}\frac{1}{\rho}\frac{\partial\tau_{ij}}{\partial x_j}\right] + \left[\frac{D^2}{Dt^2}\frac{1}{c_p}\frac{DS}{Dt}\right] \tag{949}$$

where a is the speed of sound, $\partial\tau_{ik}$ is the stress tensor and ψ represents the effects of entropy fluctuations and fluid viscosity.

The entropy fluctuation is negligible for problems without heat transfer or combustion. Viscous effects are also negligible for a high Reynolds number flow. It is therefore usual to neglect these terms from the right-hand side. Then Eqn. (948) becomes:

$$\frac{D}{Dt}\left[\frac{D^2\Pi}{Dt^2} - \frac{\partial}{\partial x_j}\left(a^2\frac{\partial\Pi}{\partial x_j}\right)\right] + 2\left[\frac{\partial u_k}{\partial x_j}\frac{\partial}{\partial x_k}\left(a^2\frac{\partial\Pi}{\partial x_j}\right)\right] = -2\left[\frac{\partial u_k}{\partial x_i}\frac{\partial u_j}{\partial x_k}\frac{\partial u_i}{\partial x_j}\right] \tag{950}$$

where:

$$\Pi = \left(\frac{c_v}{c_p}\right)\ln\frac{p}{p_0} \tag{951}$$

and p_0 is reference pressure.

By assuming that the flow satisfies the small perturbation assumption, the flow is locally linearized as follows:

$$u_i(\vec{x}, t) = U_i(\vec{x}) + u_i'(\vec{x}, t) \tag{952}$$

where $U_i(\vec{x})$ is the mean velocity field and $u_i'(\vec{x}, t)$ is the turbulent velocity field. The right-hand side of Eqn. (950) can be expanded as follows:

$$\begin{aligned}-2\left[\frac{\partial u_k}{\partial x_i}\frac{\partial u_j}{\partial x_k}\frac{\partial u_i}{\partial x_j}\right] = &-2\left[\frac{\partial U_k}{\partial x_i}\frac{\partial U_j}{\partial x_k}\frac{\partial U_i}{\partial x_j}\right] - 2\left[\frac{\partial u_k'}{\partial x_i}\frac{\partial u_j'}{\partial x_k}\frac{\partial u_i'}{\partial x_j}\right] - \\ &6\left[\frac{\partial U_k}{\partial x_i}\frac{\partial U_j}{\partial x_k}\frac{\partial u_i'}{\partial x_j}\right] - 6\left[\frac{\partial u_k'}{\partial x_i}\frac{\partial u_j'}{\partial x_k}\frac{\partial U_i}{\partial x_j}\right]\end{aligned} \tag{953}$$

$$A = -2\left[\frac{\partial U_k}{\partial x_i}\frac{\partial U_j}{\partial x_k}\frac{\partial U_i}{\partial x_j}\right] \tag{954}$$

$$B = S_{\text{SelfNoiseTerm}} = -2\left[\frac{\partial u_k'}{\partial x_i}\frac{\partial u_j'}{\partial x_k}\frac{\partial u_i'}{\partial x_j}\right] \tag{955}$$

$$C = S_{\text{ShearNoiseTerm}} = -6\left[\frac{\partial U_k}{\partial x_i}\frac{\partial U_j}{\partial x_k}\frac{\partial u_i'}{\partial x_j}\right] - 6\left[\frac{\partial u_k'}{\partial x_i}\frac{\partial u_j'}{\partial x_k}\frac{\partial U_i}{\partial x_j}\right] \tag{956}$$

where A represents the mean flow shape change, the Self Noise Term B contains effects of the turbulent (fluctuating) velocity components, and the Shear Noise Term C contains the interaction between the mean shear flow and turbulent velocity components.

The noise source terms of Lilley's equation are evaluated based on mean velocity fields provided by a steady-state RANS solution and turbulent velocity components, which are generated by the Synthetic Turbulence method developed by Bailly and Juve [144], Bechara *et al.* [147], and Kraichnan [171].

Synthetic Turbulence

A time-space turbulent velocity field is generated using random Fourier modes. This method, called the Stochastic Noise Generation and Radiation (SNGR) Model, was proposed by Kraichnan [171] and Karweit *et al.* [169], and afterward developed by Bailly and Juve [144] and Bechara *et al.* [147]. In STAR-CCM+ the SNGR method, outlined below, requires user input of two properties of the Lilley noise source model: Number of Realizations and Number of Fourier Modes.

Consider a three-dimensional Fourier decomposition of a turbulent homogeneous isotropic field at a given point at location x:

$$\vec{u}_t(\vec{y}) = \int u_t(k) e^{(jk\vec{y})} dk \tag{957}$$

where $\vec{u}_t$ is the turbulent velocity field, k is the wave vector and j is the complex number.

This Fourier integral can be transformed into a limited sum of n Fourier modes:

$$\vec{u}_t(\vec{x}) = 2\sum_{n=1}^{N} u_n \cos(\vec{k}_n \vec{x} + \psi_n)\vec{\sigma}_n \tag{958}$$

where u_n, ψ_n and $\vec{\sigma}_n$ are the amplitude, phase and directional unit vector respectively of the n-th Fourier mode associated with the wave-number vector $\vec{k}_n$.

The amplitude u_n of each mode is computed so the turbulence kinetic energy $E(k_n)$ corresponds to the von Karman-Pao energy spectrum for isotropic, homogeneous turbulence:



$$u_n = \sqrt{E(k_n)(\Delta k_n)} \tag{959}$$

where u_n is the amplitude of mode and Δk_n is the small interval in the spectrum at k_n.

The expression for the von Karman-Pao energy spectrum can be written as:

$$E(k_n) = \frac{SNGR_\alpha \overline{u^2}}{k_e} \left(\frac{k}{k_e}\right)^4 \frac{1}{\left[1 + \left(\frac{k}{k_e}\right)^2\right]^{17/6}} e^{\left(-2\left(\frac{k}{k_\eta}\right)^2\right)} \tag{960}$$

Where k is the wave number, k_η is the Kolmogorov wave number, $\overline{u^2}$ is the root-mean-square value of the velocity fluctuations corresponding to the turbulent kinetic energy, $SNGR_\alpha$ is a numerical constant, k_e is the wave number corresponding to the most energy containing eddies at the peak in the spectrum.

Proudman

The analytical result of Proudman [181] estimated the local acoustic power generated by unit volume of isotropic turbulence having no mean flow. Proudman considered the generation of noise by isotropic turbulence and using statistical models of various two-point moments, using the Lighthill analogy. In Proudman's high-Reynolds model for isotropic turbulence in near incompressible flow, Lilley added the effects of retarded time in the evaluation of the two-point covariance of Lighthill's tensor (an effect previously neglected by Proudman), and obtained the following expression for acoustic power, AP per unit volume:

$$AP = \alpha\rho_0 \frac{u^3}{l}\frac{u^5}{a_0^5} \tag{961}$$

where α is a constant related to the shape of the longitudinal velocity correlation, u is the root mean square of one of the velocity components, l is the longitudinal integral length scale of the velocity, ρ_0 is the far-field density and a_0 is the far-field sound speed.

In Proudman's original derivation [181], α is approximately 13. Lilley [174] found it to be about 10.96. In a DNS simulation done by Sarkar and Hussaini [185] α = 2.6, and in the LES study done by Witkowska and Juve [186] α = 2.5. In Proudman's paper, the terms of u and ε can be written as:

$$u = \sqrt{\left(\frac{2}{3}k\right)} \quad , \quad \varepsilon = \frac{1.5u^3}{l} \tag{962}$$

In terms of STAR-CCM+ of the turbulence velocity scale and of the turbulence length scale, the local acoustic power due to the unit volume of isotropic turbulence (in $\mathrm{W/m^3}$) becomes:

$$AP = \alpha_c\, \rho_0 \frac{U^3}{L}\frac{U^5}{a_0^5} \tag{963}$$

with:

$$U = \frac{L}{T} \quad , \quad \alpha_c = 0.629 \tag{964}$$

where ρ_0 is the far-field density, U is the turbulence velocity, L is the turbulence length scale, T is the turbulence time scale and a_0 is the far-field sound speed.

The rescaled constant is based on Direct Numerical Simulation for isotropic turbulence done by Sarkar and Hussaini [185].

The total acoustic power per unit volume can be reported in dimensional units ($\mathrm{W/m^3}$) and in dB:

$$AP(\mathrm{dB}) = 10\log\left(\frac{AP}{P_{ref}}\right) \tag{965}$$

where P_{ref} is the reference acoustic power.

Mesh Frequency Cutoff

The following section is adapted from Wagner, et al [187] on the Mesh Frequency Cutoff (MFC) measure of Mendonça.

The main value in performing a mesh-frequency cut-off analysis as post-processed from a steady-state solution is to decipher in advance of a transient calculation whether the chosen grid has sufficient resolution to capture the turbulence flow structures in the frequency range of interest.

Mendonça [179] shows that, given a cell dimension D and local turbulent kinetic energy k, the smallest length scale of a turbulent eddy structure captured by the mesh is $2D$; its associated isotropic fluctuation velocity is $\sqrt{((2/3)k)}$. Therefore the maximum frequency f_{MC} reasonable resolved by the local grid spacing D is:

$$f_{MC} = \frac{\sqrt{\left(\frac{2}{3}k\right)}}{2\Delta} \tag{966}$$

The choice of k, the mean turbulent kinetic energy contained by the mean turbulent structures, to represent the turbulence velocity fluctuation is easily justified. Most widely used RANS turbulence models solve transport equations for it directly; for those that do not, it is easily derivable. The transport-equation for k is derived directly from the unsteady Navier-Stokes equations, therefore contains contributions from all the important mechanisms: convection, diffusion, production and dissipation. So long as the mean flow features are captured in the modeling, the k solution tends also to be reasonably grid independent.

This measure is used in preference to one derived from the turbulence time scale, that is, the ratio of kinetic energy to its dissipation rate, k/e. This is better representative of the dissipative scales which, in frequency and energy content, are hopefully orders removed from the human hearing range or peak sensitivity.

Since this measure is derived from a steady-state solution, some limitations are inherent. The frequencies associated with time varying large-scale motions such as vortex-shedding, which convect through the mesh, is not accounted for. Instead, its usefulness is to approximate the frequencies of the turbulence scales which are modeled in RANS, and which become resolved in Large Eddy Simulation. In other words, this measure is more valid for the broadband and less so for narrowband excitations.

f_{MC} represents an important step in the analysis process. To be correctly applied and interpreted, use the following conditions:

1. Use an isotropic mesh (unit aspect ratio). In this case, the cube-root of the volume becomes an appropriate measure of the hexahedral cell dimension D (appropriate modifications are required for perfect tetrahedral, and so on). Obviously it is not the case when cells are stretched excessively; stretching is also detrimental to the preservation of eddy structures convecting through the grid.

2. Use realistic boundary conditions that are applied to k, especially at the inflow.
3. The mesh must be fine enough to resolve the important mean-flow features.

The mesh frequency cutoff measure has been applied to typical industrial applications at low Mach numbers flow.

Acoustic Suppression Zone

The Acoustic Suppression Zone model (also known as Sponge Zone or Buffer Zone) adds appropriate damping term(s) to the right-hand side of the governing equations, of the form:

$$-\sigma(Q - Q_{ref}) \tag{967}$$

where:

- σ is the relaxation coefficient.

  When the boundary distance is less than or equal to the zone thickness, the specific variation is a smooth transition from 0 to the Maximum Relaxation Coefficient at the boundary:

  $$\sigma = \frac{\sigma_{max}}{2}\left(1 + \cos\pi\frac{d}{w}\right)$$

  where:

  - σ_{max} is the maximum relaxation coefficient. It is the inverse of the relaxation time and is measured in sec^{-1}. The value must be greater than zero, and values of several hundred are typical.
  - d is the distance from the cell in question to the boundary using Acoustic Suppression.
  - w is the width of the Acoustic Suppression Zone.
- $Q = \{\rho, \rho v, \rho E\}$.
  - v is the flow velocity vector.
  - E is the total energy, $E = C_v T + 0.5|v|^2$, where C_v is the heat capacity at constant volume.

These additional terms damp flow disturbances (computed with regard to the mean or target flow conditions), before they reach the boundary in question. This Acoustic Suppression Zone treatment helps characteristics-based boundary conditions to produce a quasi non-reflective response, that is, to allow disturbances to travel through the boundary without producing spurious back-reflections.

Using the Acoustic Wave Model

This section contains the following:

- What is the Acoustic Wave Model?
- Acoustic Wave Simulation Workflow
- Selecting the Acoustic Wave Model
- Obtaining the Mean Flow Pressure Field
- Setting Up the Noise Source for the Acoustic Wave Simulation
- Reducing Spurious Acoustic Reflections
- Acoustic Wave Conditions and Values
- Acoustic Wave Solver Properties
- Best Practices for Using the Acoustic Wave Model
- Acoustic Wave Solver Field Functions Reference
- Acoustic Wave Model Formulation

What is the Acoustic Wave Model?

The Acoustic Wave model uses the theory of Acoustic Perturbation Equations (APE) to model noise that is generated by the flow of an incompressible fluid. The Acoustic Wave solver is one-way coupled with the CFD solver through this APE-based noise source term. In this approach, STAR-CCM+ models the acoustic pressure field inside a specified zone of interest only. A key part of the setup for this model is that you define the zone within which noise is generated.

By adopting a hybrid approach, the Acoustic Wave model produces better results (with fewer spurious effects) than a pure compressible flow simulation.

For the formulation of the Acoustic Perturbation Equation, see Acoustic Wave Model Formulation.

Acoustic Wave Simulation Workflow

The Acoustic Wave model is intended for use in unsteady flow simulations, typically incompressible (constant density). Perform acoustic wave simulations in double precision.

Preliminary Analysis

Before you begin the acoustic wave simulation:

1. Obtain a fully developed flow solution:
   - Run the unsteady flow simulation and complete several cycles of the oscillating flow pattern.
   - (Recommended) In the flow solution, use a Field Mean of Pressure monitor to compute the mean of flow pressure.

     The Acoustic Wave solver requires the mean of flow pressure in order to compute the instantaneous flow pressure fluctuations that are then used to compute the APE-based noise source.

     See Using the Mean Flow Pressure Field from a Fully Developed Flow Solution.
2. Create the appropriate field functions to define the noise source filtering and acoustic damping zones.
   - The Acoustic Wave solver uses a field function ($0 \le f_{source} \le 1$) to specify the noise source or act as a filter to specify the region of interest in the flow.
   - Physical damping is needed in the Acoustic Wave solver, since the CFD mesh typically becomes coarse away from region of interest.

     The Acoustic Wave solver requires a field function ($0 \le f_{damp} \le 1$) to define the damping region that is used to avoid spurious acoustic reflections from mesh transition.

   See Setting Up an Acoustic Perturbation-Based Noise Source and Reducing Spurious Acoustic Reflections.

Acoustic Wave Analysis

The acoustic wave simulation workflow is as follows:

1. Activate the Acoustic Wave model.

   See Selecting the Acoustic Wave Model.
2. Specify how the Mean Flow Pressure is obtained.

   See Obtaining the Mean Flow Pressure Field.
3. Set up the noise source to use in the model.

   See Setting Up the Noise Source for the Acoustic Wave Simulation.
4. If necessary, set the boundary conditions in the flow region.

   The Acoustic Wave model supports the following boundary types: Wall, Symmetry Plane, Free Stream, Inlet, Outlet, and Periodic.

   See Acoustic Wave Conditions and Values.

5. Set up the Acoustic Wave solver.
   - Use implicit time marching with sufficient inner iterations to ensure convergence.
   - If you are computing the mean flow pressure value, you must specify the appropriate start time delay.
   - If necessary to suppress simulation instability due to mesh imperfections, adjust the Newmark Delta value. Setting this value to slightly higher than the default, for example 0.502 or higher, can help to stabilize the solution.

   See Acoustic Wave Solver Properties.
6. Set up the necessary stopping criteria, derived parts, and reports: the usual options to create probes for acoustic pressure monitoring in the region of interest.
7. Run the acoustic wave simulation (with the relevant flow and Acoustic Wave models activated) to obtain the acoustic pressure values throughout the flow region.

Post Processing

To post-process the acoustic wave simulation results:

1. Output the acoustic data in a .trn file.

   Use the File > Auto Export option to export the simulation results. For more information, see the section *Exporting Surface Spectral Data* in the *User Guide*.
2. Perform post-processing of the acoustic data to compare the simulation results against other data.

   STAR-CCM+ provides built-in data set functions, such as Fourier Transforms, that you can use to apply signal processing to the acoustic data.

   For more information, see the section *Using Data Set Functions* in the *User Guide*.

Selecting the Acoustic Wave Model

The Acoustic Wave model is available for three-dimensional unsteady simulations, and with any viscous regime. The Acoustic Wave model is not compatible with Explicit, Two-Dimensional, or Axisymmetric models. The Acoustic Wave model is compatible with stationary cases only - the simulation cannot include any moving regions (such as fan blades).

Note: When you use the Acoustic Wave model, it is recommended that you also activate the Cell Quality Remediation model.

When you select the Aeroacoustics model from the *Optional Models* group box, the Acoustic Wave option becomes available in the *Aeroacoustics Models* group box.



Acoustic Wave Model Properties

The Acoustic Wave node represents the model in the physics continuum. This object has properties.



Acoustic Wave Model Properties

Far-Field Sound Speed	The sound speed in the far-field. See Eqn. (924). The default value is 340 m/s.
Noise Source	Defines the source of acoustic pressure. The specified source is used as the right-hand side of the Acoustic Perturbation Equation (APE) based noise source equation (Eqn. (970)).

None	No noise source is defined in the flow region. You use this option when you model the acoustic waves that are generated by pressure fluctuations at a boundary wall. For more information, see Setting Up a Wall Pressure Fluctuation Noise Source.
Acoustic Perturbation Based Source	Models the acoustic waves that are generated by the flow pressure fluctuations in the fluid flow. This method is the recommended option for hybrid flow and aeroacoustics simulations. For more information, see Setting Up an Acoustic Perturbation-Based Noise Source.
User-Specified Source	Models a user-defined pressure field that is applied to a single cell and then propagates into the domain. This option is used for testing purposes. For more information, see Setting Up a User-Defined Noise Source.

Mean Flow Pressure Reference	Specifies the method by which the mean pressure field is obtained. For more information, see Obtaining the Mean Flow Pressure Field.

Specified	Use the pressure field that is specified in the Mean Flow Pressure Reference node. This pressure is typically the Mean of Pressure field function that is obtained from a fully developed flow solution. For more information, see Using the Mean Flow Pressure Field from a Fully Developed Flow Solution.
Compute	Compute the mean pressure field with the Acoustic Wave solver already set up to run. For more information, see Computing the Mean Flow Pressure Field.

Obtaining the Mean Flow Pressure Field

There are two ways that you can obtain the mean flow pressure field:

- Use the Mean of Pressure field function that is obtained from a fully developed flow solution.
- Compute the mean pressure field with the Acoustic Wave solver already set up to run.

Using the Mean Flow Pressure Field from a Fully Developed Flow Solution

For hybrid flow and aeroacoustics simulations, it is recommended that you obtain a fully developed flow solution and a mean pressure field before you start the Acoustic Wave solver.

To obtain the mean flow pressure field from a fully developed flow solution:

1. In the flow simulation, set up a Field Mean of Pressure monitor:
   a. Right-click the Monitors node and then click New Monitor > Field Mean.
   b. Select the new Field Mean node and set the following properties:

Property	Value
Parts	[Flow Region]
Field Function	Pressure

2. Run the flow simulation to obtain the fully developed solution.
3. Activate the Acoustic Wave model and, under the Acoustic Wave node, set the Mean Flow Pressure Reference property to Specified.

4. Set the Mean Flow Pressure Reference to use the Mean of Pressure field function.



Computing the Mean Flow Pressure Field

In the absence of a Field Mean Pressure monitor, you can compute the mean pressure field with the Acoustic Wave solver already set up to run.

To ensure a stable mean pressure field, you must specify a Start Time delay for the Acoustic Wave solver. The time that is required to obtain a stable mean pressure field depends on the simulation parameters, particularly the geometry and the flow velocity. A reasonable guideline is to allow a delay time of ten flow-through periods.

To compute the mean pressure field:

1. In the Acoustic Wave model, set the Mean Flow Pressure Reference property to Compute.
2. In the Acoustic Wave solver, set the appropriate Start Time.

Clearing Acoustic Data from the Simulation

When you clear acoustic data, the acoustic pressure and all of the related variables that are calculated in the Acoustic Wave solver are removed from the simulation. The mean flow pressure is reset to the field that was originally specified. The flow and pressure data that was calculated by the flow solver is not affected.

To clear the acoustic data from a simulation:

- Right-click the Acoustic Wave model and then click Clear Solution.

Setting Up the Noise Source for the Acoustic Wave Simulation

You can specify the noise source in the Acoustic Wave simulation in the following ways:

- Acoustic perturbation-based noise source

  Models the acoustic waves that are generated by the fluid flow in the physics continuum.

  See Setting Up an Acoustic Perturbation-Based Noise Source.

- User-defined noise source

  Models the acoustic waves that are generated by a user-defined pressure field in a specified location in the flow region.

  See Setting Up a User-Defined Noise Source.

- Wall pressure fluctuation noise source

  Models the acoustic waves that are generated by user-defined pressure fluctuations at a boundary wall.

  See Setting Up a Wall Pressure Fluctuation Noise Source.

Setting Up an Acoustic Perturbation-Based Noise Source

In the Acoustic Wave method, you define a particular volume within which you expect most of the noise generation to occur. Outside the noise source volume, you do not expect the flow to generate any appreciable noise. You typically use a field function to specify the volume of interest and to filter out the quiescent region.

This field function for the noise source zone must return a value $0 \leq f_{source} \leq 1$. You can apply a transition modifier, such as a Hanning window function, to provide a smooth transition between the value of 1 in the interior of the noise source zone, and 0 on the outside.

Note: Do not use the name *APE Source* as the name for the noise source field function. This name is used by an existing field function.

An example noise source field function is illustrated below.



The noise source is the region around the cylinder. In this example, there is a strong left-to-right airflow that generates acoustic waves by vortex shedding from the cylinder. Note the fine mesh throughout the noise source zone.

To set up an acoustic perturbation-based noise source:

1. Create the appropriate *Noise Source Filter* field function (f_{source}) to define the noise source enclosure space.

   You can optionally create the appropriate *Acoustic Damping* field function (f_{damp}) to suppress acoustic reflections from outside the area of interest.

For more information, see Reducing Spurious Acoustic Reflections.

2. In the physics continuum, select the Acoustic Wave model and set the Noise Source property to Acoustic Perturbation Based Source.



3. Under the Regions node, expand the appropriate region. In the Physics Conditions for the region, select the Noise Source Weighting Option node and set the Noise Source Region Option property to Specified.



4. In the Physics Values for the region, select the Noise Source Weighting Coefficient node and set the Method property to Field Function.

5. Select the Field Function node and set the Scalar Function property to the *Noise Source Filter* field function that you created earlier.



If you want to apply damping to the region outside the area of interest, perform the following steps:

6. In the Physics Conditions for the region, select the Acoustic Damping Option node and set the Acoustic Damping Option property to Specified.
7. In the Physics Values for the region, select the Acoustic Damping Coefficient node and set the Method property to Field Function.

8. Select the Field Function node and set the Scalar Function property to the *Acoustic Damping* field function that you created earlier.



Setting Up a User-Defined Noise Source

You can create a user-defined noise source for testing or calibrating the model. You specify the pressure field for a particular cell (or cells) in the flow mesh. The pressure field then propagates into the flow region.

Monopole (the pressure field is defined for a single cell), dipole (two cells), and quadrupole (four cells) noise sources are typical examples of user-defined noise sources.

For example, the analytical expression of the monopole source is:

$$P = A\frac{R}{r}\sin(\omega t - kr) \tag{968}$$

The corresponding field function that you need to apply for the monopole source is:

$$\text{Input Pressure} = 4\pi AR\sin(\omega t) \tag{969}$$

Note that a scaling factor may be required in the acoustic source.

The acoustic pressure field that is generated by the monopole source is illustrated below:



To set up a user-defined noise source:

1. Create the appropriate *Acoustic User Source* field function.

2. In the physics continuum, select the Acoustic Wave node and set Noise Source to User-Specified Source.



3. Select Regions > [flow region] > Physics Conditions > Noise Source Weighting Option and set Noise Source Region Option to Specified.

4. Within the [flow region] > Physics Values node, select Noise Source Weighting Coefficient and set Method to Field Function.

5. In the Field Function node, set the Scalar Function to the *Acoustic User Source* field function.



Setting Up a Wall Pressure Fluctuation Noise Source

You can specify a pressure fluctuation at a boundary wall as the source of acoustic waves.

This noise source is used for testing purposes.

To set up a wall pressure fluctuation noise source:

1. Create the appropriate *Wall Pressure* field function.
2. If the wall pressure fluctuation is the only noise source in the physics continuum, select the Acoustic Wave model and set the Noise Source property to None.

   You can have both a wall pressure fluctuation noise source and an acoustic perturbation-based or user-specified noise source in the same simulation. In that case, leave the Noise Source property set to Acoustic Perturbation Based Source or User-Specified Source, whichever is appropriate.
3. Under the Regions node, expand the appropriate wall boundary. In the Physics Conditions for the boundary, select Wall Acoustic Pressure Specification and set Method to Specified.

4. In the Physics Values, select Wall Acoustic Pressure and set the Method to Field Function.
5. Select the Field Function node and set the Scalar Function to your *Wall Pressure* field function.

An example of the acoustic pressure field that is produced by a wall pressure fluctuation noise source is shown below:



Reducing Spurious Acoustic Reflections

For practical applications, the flow mesh is typically coarsened away from the region of interest. A strong vortical flow crossing through this mesh transition in the downstream region can generate spurious acoustic waves that propagate upstream. These spurious acoustic waves interfere with those inside the noise source zone, resulting in erroneous results.

You typically use a field function (f_{damp}) to suppress acoustic reflections from outside the noise source zone. This field function must return a value of 0 within the noise source zone, and a value of 1 at the boundaries of the flow domain. Within the transition from 0 to 1, you can apply a Hanning windowing modifier to give a smooth transition.

An example of a damping field is illustrated below.



In this example, there is a strong left-to-right air flow which generates acoustic waves by vortex shedding from the cylinder. The region of interest (in blue) is undamped and everywhere else is fully damped (in red). Note the fine mesh throughout the undamped zone and the damping that is applied at the mesh transitions.

The recommended points per wavelength (PPW) must be maintained inside the undamped zone.

For more information on the recommended PPW, see Best Practices for Using the Acoustic Wave Model.

Acoustic Wave Conditions and Values

If necessary, depending on your geometry, you can set acoustic boundary conditions at walls and symmetry planes.

Initial Conditions for the Acoustic Wave Model

When you activate the Acoustic Wave model, the following nodes are added under the Initial Conditions node in the physics continuum:

Acoustic Pressure	The value p^a that is used in the Acoustic Perturbation Equation (Eqn. (970)).
Acoustic Pressure Second Time Derivative	The second time derivative of acoustic pressure, $\frac{\partial^2 p^a}{\partial t^2}$. This value is used for the Newmark solver.
Acoustic Pressure Time Derivative	The first time derivative of acoustic pressure, $\frac{\partial p^a}{\partial t}$. This value is used for the Newmark solver.

You do not need to set these initial conditions in a typical acoustic wave simulation. However, you can use particular settings for testing purposes.

For more information on the corresponding field functions, see Acoustic Wave Solver Field Functions Reference.

Boundary Conditions for the Acoustic Wave Model

When you activate the Acoustic Wave model, the following acoustic boundary conditions become available:

Wall Acoustic Pressure Specification	Applies to Wall boundaries only.	
	None	The wall boundary has no specific acoustic pressure but is a reflecting boundary.

	Specified	The acoustic pressure is specified in the Physics Values > Wall Acoustic Pressure node. For more information on setting this boundary condition, see Setting Up a Wall Pressure Fluctuation Noise Source.
	Non-reflecting	The wall is a non-reflecting boundary. Acoustic waves pass through the boundary.
Symmetry Acoustic Pressure Specification	Applies to Symmetry Plane boundaries only. By default the boundary is reflecting, but you can make it non-reflecting if necessary.	
	Non-reflective Boundary Condition Treatment	• Activated - the symmetry plane boundary is non-reflecting. Acoustic waves pass through the boundary. • Deactivated - the symmetry plane boundary is a reflecting boundary.

Acoustic Wave Solver Properties

The Acoustic Wave solver is activated when the Acoustic Wave model is selected. Similar to the flow solvers, the Acoustic Wave solver uses the AMG Linear Solver infrastructure that is provided in STAR-CCM+.

The Acoustic Wave solver is represented by the Acoustic Wave node in the Solvers tree.

Acoustic Wave Solver Properties

Under-Relaxation Factor	The under-relaxation factor is set to 1.0 and cannot be modified. It is not appropriate to use under-relaxation, as that would be equivalent to applying damping to the simulation.
Start Time	The elapsed simulation time before the solver begins calculating. By default, the solver starts with a five time step delay (if the start time is set to 0.0s) in order to pre-compute the noise source time derivative. If the Mean Flow Pressure Reference property in the Acoustic Wave model is set to Compute, you must specify a start time delay. The time that is required depends on the simulation parameters, particularly the geometry and the flow velocity. A reasonable guideline is to allow a delay time of ten flow-through periods.

Acoustic Wave Solver Expert Properties

Newmark Delta

This property specifies a parameter value for the Newmark-beta method of numerical integration that is used in the Acoustic Wave solver. The default value of 0.5 is suitable for most cases.
However, to suppress simulation instability due to mesh imperfections, you can set this value to slightly higher than the default. For example, 0.502 or higher, as required to stabilize the solution. Some trial and error can be required to determine the appropriate value.

Solver Frozen

Prevents the solver from updating the solution while iterating.

Activated — The solver does not update the solution.

Deactivated — The solver updates the solution.

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

Activated — Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.

Deactivated — Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

Activated — Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.

Deactivated — Do not reset reconstruction gradients. They are recomputed at the next iteration if they had been zeroed.

Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. These quantities become available as field functions during subsequent iterations. Activated — Retain temporary storage. Deactivated — Do not retain temporary storage.

Best Practices for Using the Acoustic Wave Model

Some best practice recommendations for using the Acoustic Wave model are as follows:

- When you use the Acoustic Wave model, it is recommended that you also activate the Cell Quality Remediation model.
- Use a high-resolution turbulence model to ensure good noise source predictions.
- Use the appropriate mesh resolution in the zone of interest.

  The mesh resolution must be based on the recommended points per wavelength (PPW) to maintain the accuracy of acoustic predictions, where PPW = (speed of sound/frequency) / grid spacing.

  It is recommended that you use PPW ≥ 30–40, and this value must be maintained throughout the zone of interest. The recommended Acoustic CFL depends on the PPW: for PPW of 40, it is recommended that Acoustic CFL ≤ 4 inside the zone of interest. However, for higher PPW values, a higher Acoustic CFL is acceptable.
- Apply damping to improve acoustic prediction. There are two damping methods that you can use to improve simulation stability:
  - Numerical damping that is applied to the Acoustic Wave solver

    The implementation of the Acoustic Wave model takes into account the solver robustness, but some instability can still arise due to mesh imperfections. Over time in a long simulation, localized hot spots of acoustic pressure can appear. These hot spots can lead to the accumulation and spreading of instability.

    To suppress simulation instability due to mesh imperfections, you can adjust the Newmark delta value in the Acoustic Wave solver.
  - Physical damping to eliminate spurious acoustics from mesh transition or boundary effects

    In order to eliminate acoustic wave reflections into the area of interest, you can apply a damping function (f_{damp}) to the region outside the area of interest. The reflections are attenuated and their effects are minimized.

For more information, see Reducing Spurious Acoustic Reflections.

- When you specify field functions to define the noise source filtering (f_{source}) and acoustic damping (f_{damp}) zones, a Hanning windowing is recommended for the transition from a value of 0–1.

Acoustic Wave Solver Field Functions Reference

When the Acoustic Wave Solver has the Temporary Storage Retained property activated, the following field functions are made available to the simulation.

Acoustic CFL in Wave Equation	The Acoustic CFL ($c\frac{\Delta t}{\Delta x}$, where c is the speed of sound) is computed in the whole domain, based on the mesh size and time step.
Acoustic Pressure	The Acoustic Pressure is the variable being solved.
Acoustic Pressure Gradient Acoustic Pressure Recon	The computed gradient fields using STAR-CCM+ gradient options.
Acoustic Pressure Second Time Derivative Acoustic Pressure Time Derivative	The fields that are used for the Newmark solver.
APE Pressure	The fluctuating flow pressure field.
APE Source	The acoustic perturbation-based source term that is generated in the Acoustic Wave equation.
Mean Flow Pressure	The mean flow pressure that is obtained from the flow reference.

Acoustic Wave Model Formulation

The Acoustic Perturbation Equation (APE) based acoustic wave equation with physical damping is:

$$\frac{1}{c^2}\frac{\partial^2}{\partial t^2}p^a - \nabla \cdot \nabla\left(p^a + \tau\frac{\partial p^a}{\partial t}\right) = -\frac{1}{c^2}\frac{\partial^2}{\partial t^2}P' \tag{970}$$

where:

- p^a is the acoustic pressure
- $P' = P - P_{mean}$ is the flow pressure fluctuation
- c is the speed of sound
- $\tau = f_{damp}\left(\frac{\Delta t}{\pi CFL^a}\right)$
  where:
  - f_{damp} = damping field
  - Δt = time step
  - CFL^a = local acoustic CFL

Hanning Window

The Hanning function has the general form:

$$\mathrm{H}(r) = 0.5\left(1 - \cos\left(\pi\left(\frac{\mathbf{r}}{\mathbf{t}}\right)\right)\right)$$

where:

- $\mathbf{r}$ is the direction vector
- $\mathbf{t}$ is the thickness

Using the Direct Noise Simulation

Direct noise simulation, or the Computational Aeroacoustics approach, requires solving the entire flow field for the full unsteady simulation. To calculate the aerodynamically generated acoustics, record the time-varying static pressure on points (probes) or surfaces. It requires high numerical resolution to capture the spatial and transient effects of the acoustic pressure perturbations. For post-processing, the Fourier Transform (FT) tools can be used for spectral analysis. (See the section on data set functions for details.).

Predicting noise levels near to the source of sound generation can be done directly within a CFD model. Direct simulation can be too expensive to predict noise levels some distance from the source, however. For such far-field predictions, STAR-CCM+ provides the Ffowcs Williams-Hawkings (FW-H) acoustics integral formulation.

Using the Ffowcs Williams-Hawkings Aeroacoustics Models

Modern CFD adopts two strategies for prediction of far-field acoustics: one is based on the Direct Noise Simulation of sound propagation and the other is based on integral formulations.

Both approaches are known Computational Aero-Acoustics (CAA) methods, and have been extensively used to study the sound generation and sound propagation in the fluid flow.

The Ffowcs Williams-Hawkings (FW-H) acoustics integral formulation is the preferred strategy for far-field noise prediction. This model is able to calculate the far-field sound signal that is radiated from near-field flow data from a CFD solution. The goal is to predict small amplitude acoustic pressure fluctuations at the locations of each receiver.



Index	Description
1	Flow direction
2	Cylinder surface / FW-H Impermeable source surface
3	FW-H Permeable source surface
4	Computational domain boundary / Limit of Direct Noise Simulation near-field solution
5	FW-H receiver in the far field
6	FW-H quadrupole noise visualization

The Ffowcs Williams-Hawkings Models

The Ffowcs Williams-Hawkings (FW-H) acoustic model is the preferred strategy for far-field noise prediction such as overhead aircraft noise. This model is able to calculate the far-field sound signal that is radiated from near-field flow data from a CFD solution. The goal is to predict small amplitude acoustic pressure fluctuations at the location of each receiver.

The FW-H acoustic model for stationary motion can simulate wind tunnel aeroacoustic measurements with stationary sources such as aircraft landing gear or an airfoil trailing edge. The FW-H acoustic model is used only to predict the propagation of sound in free space. It does not include effects such as sound reflections, refraction, or material property change.

The formulation is based on Farassat's Formulation 1A, that is the non-convective form of FW-H [151]. In the FW-H approach, propagation of sound into the far field is not explicitly calculated, but computed using an analytical integral solution to the generalized wave equation (see [143] for validation).

Therefore, numerical accuracy is an issue of much less concern than in the full Computational Aeroacoustics approach, which explicitly resolves propagation of sound waves all the way to the far field. With the FW-H approach, a second order discretization scheme, both spatial and temporal, together with judicious choices for mesh resolution and time-step size, can give accurate far-field noise prediction.

As an input of FW-H model, time-accurate data is needed on the near field flow, corresponding to a fine mesh on the near field. The fine meshes used for accurate prediction of the boundary layer or other requirements for turbulence models (such as y+ for LES) are sufficient. The computational domain for this model does not need to include the location of the receivers. However, to use non-reflecting boundary conditions requires a larger domain so a buffer layer can be added to reduce and dissipate the acoustic waves.

The depth of the computational domain must be greater than, or equal to, the acoustic correlation length. This depth may be different from the depth that is required for a good LES simulation. Currently, the FW-H model uses the same solution domain for acoustic sources estimation. When three-dimensional CFD solutions with the FW-H model are periodic in the spanwise direction (for example an infinite cylinder), replicate the solution domain in the spanwise direction. For two-dimensional simulations, a correlation length can be used to overcome this issue.

The following sections cover the three main Ffowcs Williams-Hawkings sub-models:

- The Ffowcs Williams-Hawkings Steady model, for 3-D rotating cases with impermeable FW-H surfaces, using the FW-H Steady Solver.
- The On-The-Fly FW-H model, an unsteady time model for far-field noise prediction and stationary cases, using the FW-H Unsteady Solver.
- The Post FW-H model, an unsteady time model, for collecting flow field data in `.trn` files, in stationary cases with impermeable FW-H surfaces, using the Post FW-H Solver.

The following sections also cover the following:

- Working with FW-H receivers
- Working with FW-H surfaces
- Visualizing the FW-H Noise Sources
- FW-H Field Functions Reference
- Formulation

The Ffowcs Williams-Hawkings Steady Aeroacoustics Model

The FW-H Steady model is based on the convergent steady-state RANS approach for aerodynamic computation of propeller blade loads for a single or multiple rotating reference frames with the FW-H Acoustic Analogy Model. It is for three-dimensional simulations with impermeable FW-H surfaces. It cannot be used for non-rotating cases. (To see applications of this model refer to [151], [158], [162], [163] and [177].) The Ffowcs Williams-Hawkings steady node represents the model, which has properties

and is subject to common menu items.



This method deals with the case where an unsteady motion problem is converted to a steady-state RANS problem by imposing a rotating frame of reference on a static mesh. The FW-H Steady acoustics model (for convergent steady-state RANS simulations with MRF) uses a single rotating reference frame when all parts selected for an FW-H impermeable surface have an identical rotation rate value. In this case, even though all parts execute an identical rotating movement, different FW-H impermeable surfaces can have different rotation rates. When different parts rotate at different rotation rates, the FW-H Steady model implements multiple rotating reference frames.

Thickness surface noise is dependent only on the shape and motion of the blade. Conceptually, the displacement of the air by the rotor blade causes this noise. Loading noise is an aerodynamic adverse effect due to the acceleration of the force distribution on the air around the rotor blade due to the blade passing through it. Brentner and Farassat [151] listed various sound generation mechanisms. However, steady-loading noise and blade-displacement noise are perceived as the major contributions in the absence of installation effects.

The thickness and loading noise from relations of the FW-H acoustics model are computed from steady-state computation of the pressure field, and the time derivative of the surface pressure is set to zero.

The steady formulation allows prediction of an unsteady pressure time signal from a steady simulation by defining an artificial unsteady time-step. Specify this artificial time-step in the Number of Time Steps Per Revolution and the Number of Revolutions properties of the FW-H Steady solver.

Note: When the flow direction is not aligned with the axis of rotation, it is not possible to use FWH for steady simulations because the loading on the blade is unsteady. Therefore, perform a transient rigid body motion simulation (with a transient sliding mesh) with the FW-H Unsteady model.

Previous aeroacoustic results for a few propeller cases show good agreement with experimental measurements for high rotational speed (over 2500 rpm) and larger under-prediction of the noise at low rotational speed. Thus it is not possible to predict vortex noise at low tip Mach numbers (low rotational speed). For low tip Mach numbers, it is suggested to choose accurate simulations such as a transient LES simulation with rotating motion, to simulate all transient features and broadband noise.

If the boundary is part of the region where motion is stationary and the reference frame is rotating, then the model uses the rotating reference frame properties to compute the rotating movement of the FW-H Impermeable Surface. The reference frame properties are: origin, direction, and rotation rate. To determine the position of the blade on rotating movement, the face centroid of the blade and the surface area of the blade are evaluated at each artificial unsteady time-step.

Periodic domain is considered in the same way as in the On-The-Fly FW-H model for rotating motion.

Ffowcs Williams-Hawkings Steady Acoustics Model Properties

Far-Field Density	The density of the far-field fluid. The default value is 1.225 kg/m^3.
Far-Field Sound Speed	The sound speed in the far-field. See Eqn. (924) The default value is 340 m/s.

Selecting the Ffowcs Williams-Hawkings Steady Model

The steady model becomes available in the *Physics Model Selection* dialog after you select the Steady time model.



Using the FW-H Steady Solver

The solver provides an approximation of the results that you can expect from a steady state simulation. On computation of the thickness and loading noise, the term corresponding to the time derivative of the surface pressure is set to zero.

The Ffowcs Williams-Hawkings steady model activates the FW-H Steady solver. This solver implements the advanced time approach, to account for the time delay between the emission time and reception time. The FW-H Steady model is able to compute noise from propellers in the subsonic as well as transonic and low supersonic domain at low numerical cost. It is

suited for propeller optimization because it is inexpensive and reliable, though it is an approximate method. The FW-H Steady Solver node represents the solver. The node has properties.



To run the solver:

1. Enter the Number of Time Steps per Revolution and Number of Revolutions, as described in the following properties table.
2. Right-click the FW-H Steady Solver node and select Execute solver.

FW-H Steady Solver Properties

Number of Time Steps per Revolution	Sets the number of equivalent transient time-steps for the rotating region to complete one revolution.
Number of Revolutions	Specifies the number of revolutions that simulated by the FW-H Steady Model.

FW-H Pop-up Menu

Execute Solver	Runs the FW-H Steady Solver.

Stopping the FW-H Steady Solver

To stop the solver immediately, without completing the set number of time-steps, click on the solver stop button located next to the progress bar. The solver stops and saves the partial solution on the point receivers.



The On-The-Fly FW-H Model

The On-The-Fly FW-H acoustic model provides the aerodynamic noise prediction in parallel with the aerodynamic computation. The On-The-Fly FW-H node represents the model, which has properties and is subject to common menu items. It becomes available when the Ffowcs Williams-Hawkings Unsteady model is active.



A transient simulation using the On-The-Fly FW-H model does the following:

1. It computes the time-accurate data on the emission surfaces (sources of noise), which are the integration surfaces for the FW-H solver. Time-accurate solutions can be obtained from an unsteady turbulent or laminar simulation (DES/LES simulations are suggested).
2. Time variation of the sound pressure at receiver locations is computed based on the FW-H solver.

On-The-Fly FW-H Model Expert Properties

The expert properties below are common for both stationary simulations or simulations with rigid body motion:

Quadrupole Noise	Determines if the volumetric quadrupole radiated noise is calculated.	
	Activated	The quadrupole noise is calculated using the collapsing sphere formulation. The Quadrupole Volume Term field function is activated. An option to gather quadrupole noise data at each FW-H receiver is activated.
	Deactivated	The quadrupole noise is not calculated.

Using the On-The-Fly FW-H Model with Rigid Body Motion

The On-The-Fly FW-H model can also be used with a rigid body motion (see [149], [151], [158]). This method involves actual displacement of mesh vertices in real time. The FW-H Surface can have a rigid body motion, where it can be available for rotating motion as well as rotating and translating motion. The On-The-Fly FW-H model for rotating sliding mesh (as subset of Rigid Body Motion) could be used in the prediction of noise from rotating blades as helicopter rotors and aircraft propellers.

For rotating motion, a periodic domain is considered when the coupled/segregated solver computes flow in one sector of a rotating body. The On-The-Fly FW-H model reproduces noise sources for all sectors of the body. The model searches for a Periodic Interface configuration and determines the number of sectors from the Rotation angle that appears in

the Periodic Transformation properties. The Rotation angle is automatically computed when the Periodicity property is set to Rotational or Rotational+Translational, with Translation equal to zero.

The On-The-Fly FW-H model computes the sound pressure signal for all boundary zones of the FW-H surface, where each boundary zone can be part of a region with different number of sectors. You can use the model with periodic domains by selecting FW-H Impermeable, but only for three-dimensional simulations.



What Is the On-The-Fly FW-H Workflow?

Before using the On-The-Fly FW-H model, first solve your unsteady simulation for noise sources in the near-field using the aeroacoustics recommended analysis procedure. Then, do the following:

1. Add the FW-H Unsteady and On-The-Fly FW-H models to your analysis.
2. Create FW-H surfaces.
3. Create FW-H receivers and associate each with an FW-H surface.
4. Set the FW-H Unsteady Solver properties, especially the time when receivers start collecting data.

5. When the solution has finished, export the receiver noise data and reimport as tables to visualize and post-process the data.

Note: For quasi-steady rotating machines, use the FW-H Steady Solver.

What Is the FW-H Unsteady Solver?

The On-The-Fly FW-H model activates the FW-H Unsteady Solver. This solver implements the advanced time approach for subsonic and transonic flows, to account for the time delay between the emission time and reception time. The FW-H Unsteady Solver node represents the model and has properties and is subject to common menu items.



The advanced time approach is also known in the literature as a "source-time-dominant" approach. It was developed by Brentner [150], based on Formulation 1A of the FW-H analogy by Farassat [160]. It was extended to permeable (or penetrable) integration surfaces by di Francescantonio [157] and by Brentner and Farassat [151].

This approach looks forward in time to see when the observer receives the currently generated sound waves. In the source-time-dominant algorithm, it fixes the source time and determines when the sound from each point on the FW-H Surface reached the observer. The arrival time of the acoustic signal is different for each point on the FW-H Surface. Then the sound pressure at the receiver is obtained by accumulating the arriving signals in time slots. The overall observer acoustic signal is obtained from the summation of the acoustic signal that is radiated from each source element of the FW-H Surface during the same source time.

Note: It is important to collect a time history of the sound pressure at the receiver location long enough to do spectral analysis using Fourier Transforms.

FW-H Unsteady Solver Properties

Start Time — The elapsed simulation time before the solver begins calculating.

FW-H Unsteady Solver Expert Properties

Solver Frozen — Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work.

- Activated — The solver does not update the solution.
- Deactivated — The solver updates the solution.

Selecting the On-The-Fly FW-H Unsteady Model

The On-The-Fly FW-H model becomes available in the *Physics Model Selection* dialog after you select the Aeroacoustics model and the Ffowcs Williams-Hawkings Unsteady model, provided you are running a transient implicit simulation.



Working with FW-H Receivers

The FW-H Receivers node is available for the Ffowcs Williams-Hawkings steady model and the On-The-Fly FW-H model. The FW-H Receivers node is the manager node for all (point) receivers in the STAR-CCM+ simulation tree. It has no properties, but has a pop-up menu for controlling its contents.



FW-H Receivers Node Properties

Far-Field Density	The density of the far-field fluid. The default value is 1.225 kg/m^3.
Far-Field Sound Speed	The sound speed in the far-field. See Eqn. (924) The default value is 340 m/s.
Acoustic Correlation Length	Sets the correlation length to include perfect spanwise correlation of the flow over a finite span, for two-dimensional simulations only.

What is a Point Receiver?

A point receiver is a point inside or outside the domain (located in the mid to far field), where the sound pressure is predicted. The total acoustic signal that is measured at a receiver is the sum of the acoustic signals that are

radiated from each source element of the FW-H surface. A point receiver is represented by a node in the object tree, which has properties and a pop-up menu.



If more than one of these objects are added to the simulation, then the second and subsequent ones have numbers at the end of the labels. You can define many receivers, each with a specific FW-H surface to watch.

Receivers can be renamed, copied, and deleted in the same way as other simulation objects.

Creating Point Receivers

To add a point receiver to the simulation, right-click the FW-H Receivers manager node and select New Point Receiver.



Point Receiver Properties

Index	Index reference for the point receiver (read-only).
Coordinate System	Combined with the setting for the Position property, your selection of a local coordinate system permits definition of the point location of the receiver.
Position	The position of the receiver.
FW-H Surface	Selects which FW-H surface to watch from a drop-down list of surfaces in the simulation. Only one of these surfaces (impermeable or permeable) can be selected at a time for each point receiver, though each surface can contain multiple boundaries.

Point Receiver Expert Properties

Regions	Defines the regions or user-defined volumes to contribute to the computation of quadrupole noise. By default, no regions are selected.

Point Receiver Pop-Up Menu

Export Surface Total...	Exports the summation of the overall surface acoustic signal, $p'_S(\grave{x}, t)$, that is radiated from each element of the FW-H Surface during the source time. See Eqn. (980) for formulation.
Export Surface Thickness...	Exports the summation of the acoustic thickness signal, $p'_T(\grave{x}, t)$, that is radiated from each element of the FW-H Surface during the source time. See Eqn. (981) for formulation.
Export Surface Loading...	Exports the summation of the acoustic loading signal, $p'_L(\grave{x}, t)$, that is radiated from each element of the FW-H Surface during the source time. See Eqn. (982) for formulation.
Export Volume Quadrupole...	Exports the summation of the acoustic quadrupole signal, $p'_Q(\grave{x}, t)$, that is radiated from each cell in the regions defined in the point receiver property during the source time. See Eqn. (990) for formulation.
Clear Receiver	Clears the acoustic recorded data of the point receiver.

See also: Common menu items.

Clearing Point Receiver Data

To clear the solution data of the point receivers, right-click the FW-H Receivers manager node and select Clear Receivers.



Exporting Point Receiver Data

To export the solution data of a point receiver, right-click the node of that receiver and select one of the export options.



This action opens a standard dialog for you to save the data of the receiver as a `.csv` file in a location of your choice. The options available are listed in its pop-up menu.

For each type of export, the sound pressure data is pruned before export. The pruning accounts for the fact that signals from source surface faces closest to the receiver arrive before signals from faces furthest from the receiver. The pruning action clips the receiver data at the beginning and end of the collection time to ensure that all faces on all the source surfaces were contributing to the received data. To account for this pruning of exported data, run the case for a longer period.

Pop-Up Menu of the FW-H Receivers Manager

New Point Receiver	Adds a point receiver to the object tree.
Clear Receivers	Clears the acoustic recorded data of all the point receivers.

See also: Common menu items.

Working with FW-H Surfaces

FW-H Surfaces define the source locations of the noise. Think of FW-H as being similar to a speaker in a radio -- they are the source of noise. Greschner *et al.* [164] suggest that there are two types of FW-H Surfaces: impermeable and permeable. Surfaces can be renamed, copied, and deleted in the same way as other simulation objects.

Impermeable FW-H Surfaces

An impermeable FW-H surface acts as a filter of wall boundary conditions from which noise originates. Locate impermeable surfaces at one or more wall boundaries.



Note: For the FW-H steady model, use only impermeable surfaces.

Permeable FW-H Surfaces

A permeable FW-H surface acts as a filter of internal interface boundaries. Locate permeable surfaces on a fine region of the mesh that surrounds all sources of noise: monopoles, dipoles, and quadrupoles.



Creating Surfaces

To add a surface to the simulation, right-click the FW-H Surfaces manager node and select Impermeable Surface or Permeable Surface from the New submenu.



FW-H Surfaces Manger

The FW-H Surfaces manager node is available for the Ffowcs Williams-Hawkings steady model and the On-The-Fly FW-H model. The FW-H Surfaces node is the manager node for all FW-H surfaces

(impermeable and permeable) in the STAR-CCM+ simulation tree. It has no properties, but has a pop-up menu for adding surfaces.



This manager can hold any number of surfaces up to the memory limits of the computer.

Surface Properties

Boundaries	List of boundaries used for FW-H surfaces. Edit this property to add or remove boundaries with the in-place object selector .

Pop-Up Menu of the FW-H Surfaces Manager

New	Adds a surface (impermeable or permeable) to the object tree, depending on your choice from the sub-menu.

See also: Common menu items.

Visualizing the FW-H Noise Sources

To view the Loading and Thickness surface noise components, or the Quadrupole volume noise, select the scalar function relating to the correct part. For a surface term, display the scalar function on the same parts you used to define the associated FW-H Surface. The image below shows the Total Surface Noise for a cylinder in a crossflow:



For a volume term, display the scalar function on the region you specified in the Point Receiver properties. The image below shows the quadrupole noise for the same cylinder as shown above:



Note: The FW-H Noise Sources field functions can be visualized after the FW-H solver begins calculating, provided that the physical time is larger than the start time of the FW-H solver. (The start time of the FW-H solver is set on the Solvers > FW-H Solver node.)

FW-H Field Functions Reference

The following field functions are made available when either the Ffowcs Williams-Hawkings steady model or the On-The-Fly FW-H model is used:

Total Surface Term: [Receiver Name]	$p'_S(\grave{x}, t)$ in Eqn. (980). The summation of the Loading Surface Term and the Thickness Surface Term for a particular source surface generates the Total Surface Term.
Thickness Surface Term: [Receiver Name]	$p'_T(\grave{x}, t)$ in Eqn. (981) for general flows and in Eqn. (986) for flows with rigid body motion or moving reference frames. In each equation, the field function is normalized with face-area. The displacement of fluid as the body passes generates the monopole noise. For stationary bodies (such as side view mirrors and circular cylinders) the thickness surface term is zero.
Loading Surface Term: [Receiver Name]	$p'_L(\grave{x}, t)$ in Eqn. (982) for general flows and in Eqn. (987) for flows with rigid body motion or moving reference frames. In each equation, the field function is normalized with face-area. The unsteady motion of the force distribution on the body surface generates the dipole noise.
Quadrupole Volume Term: [Receiver Name]	$p'_Q(\grave{x}, t)$ in Eqn. (990), normalized with cell volume. Non-linearities in the flow generate quadrupole noise. The field function is only enabled when you activate the Quadrupole Noise expert property.

The Post FW-H Model

The Post FW-H model collects data in `.trn` files for post-processing. This model is presently limited to stationary motion with permeable or impermeable surfaces. This is an unsteady model, used by the Post FW-H Solver.

This model uses the advanced time approach, as described for the FW-H Unsteady Solver. In the advanced time approach, the surface integrals are computed at source time, but signals are accumulated in discrete bins, at receiver time.

The model calculates the far-field noise contribution at each point receiver on request, when you use the Execute Solver action in the Post FW-H Solver. For each selected imported surface and each receiver, at each collected flow field timestep, the model:

- Computes the thickness and loading terms. It computes all surface integrals based on reading pressure surface data at the required time levels for accurate time derivatives.
- Accumulates the contribution of each individual integral in discrete "time bin(s)," each bin correspond with an actual reception time.

The noise level at each individual receiver, for the time bin of that receiver, is a superposition of signals emitted by sound sources on the selected acoustic surfaces, lagged in time.



What Is the Post FW-H Workflow?

Before using the Post FW-H model, solve the unsteady simulation for noise sources in the near-field using the aeroacoustics recommended analysis procedure:

1. Add the Post FW-H model to your analysis.

2. Create Post FW-H Surfaces.
3. Set the Post FW-H Solver properties, especially the Start Time Step and Collection Frequency.
4. Run the simulation and collect source data in `.trn` file format.
5. Import the `.trn` file format and create Post FW-H Receivers, then set the properties Point Part and Imported FW-H Surfaces.
6. Right-click on the Post FW-H Solver node and select Execute Solver.
7. When the Post FW-H Solver has finished, export the receiver noise data and re-import it as tables to visualize and post-process the data.



What Is the Post FW-H Solver?

The Post FW-H model activates the Post FW-H Solver. The solver has two actions:

1. Define the method for collecting source data in `.trn` files, before the simulation starts.

   The solver collects source data for Impermeable Post FW-H Surfaces, in `.trn` file format, on a schedule based on the properties Start Time Step and Collection Frequency.

2. Compute the FW-H sound pressure signals after the simulation is finished.

   The solver computes the sound pressure at the locations of the selected point receivers, for the selected imported surfaces.

To run the solver:

1. Right-click on the Post FW-H Solver node.
2. Select Execute Solver.

The sound pressure data is saved on each point receiver.

The export of noise source data in `.trn` file format is completely separate from the computation of the sound pressure signal at different receivers, at different noise sources.



Post FW-H Solver Properties

Start Time Step	The timestep at which the Collection Frequency comes into effect. The value can range from 0 up. The default is 0.
Collection Frequency	The number of timesteps between updates. The value can range from 1 up. The default is 1.
Base Name	The base `.trn` file name. The default is the name of the `.sim` file.
Export Directory	The directory for exported `.trn` files. The default is the directory of the `.sim` file.

Stopping the Post FW-H Solver

To stop the solver immediately, without completing far-field computation for all time-steps in the .trn file, click on the solver stop button located next to the progress bar. The solver stops and saves the partial solution on the post point receivers.



Working with Post FW-H Surfaces

Post FW-H surfaces define the boundaries from which field data is exported to the .trn file, so that it becomes available for a subsequent analysis using the Post FW-H solver. The surfaces define the source locations of noise for far-field noise prediction using the Post FW-H solver. You can create permeable and impermeable forms of Post FW-H surfaces. For permeable surfaces, the Post FW-H model exports pressure, density, and velocity data. For impermeable surfaces, the model exports pressure data. See FW-H Surface Terms.

To create a post FW-H surface:

1. Right-click on the Post FW-H Surfaces node.
2. Select New > Post Permeable Surface or New > Post Impermeable Surface.

Post FW-H surfaces can be renamed, copied, and deleted in the same way as other simulation objects.

Note: The Post FW-H acoustic model is currently available for stationary simulations only.



Post FW-H Surface Properties

Boundaries	Sets the boundaries defined as permeable or impermeable.

Working with Post FW-H Point Receivers

The post FW-H point receivers collect acoustic data. Use them to compute the far-field noise prediction at each receiver location. You must select a point derived part for each receiver and a list of imported FW-H surfaces from a single `.trn` file.

To set up a point receiver:

1. Import a `.trn` file containing all transient surface data for the FW-H surfaces generated during a preliminary transient run:

   a. Import the `.trn` file: select File > Import > Import CAE Model...
   b. Select the `.trn` file and click Open.
   c. Select a list of imported surfaces from the same `.trn` file.

2. Set up a point probe under Derived Parts.

3. For the Input Parts property of the probe, select all Imported FW-H surfaces from the `.trn` file.

4. Create the new post point receiver: right-click the Post FW-H Receivers node and select New...

5. Set the Point Part property to the name of the point probe.

6. Set the Imported FW-H Surfaces property using only the surfaces generating the sound for that specific post point receiver.

After you run the Post FW-H Solver, the data from the post-processing point receiver is available for export.

FW-H post-processing point receivers can be renamed, copied, and deleted in the same way as other simulation objects.

Post FW-H Receivers Node Properties



Far-Field Density	The density of the far-field fluid. The default value is 1.225 kg/m^3.
Far-Field Sound Speed	The sound speed in the far-field. See Eqn. (924) The default value is 340 m/s.
Acoustic Correlation Length	Sets the correlation length to include perfect spanwise correlation of the flow over a finite span, for two-dimensional simulations only.

Post Point Receiver Properties



Imported FW-H Surfaces	Selects the FW-H surfaces used as inputs.
Point Part	Selects the point probe used as receiver, representing a single microphone.

Post Point Receiver Pop-Up Menu

Export...	Exports the sum of surface acoustic signals radiated from each element of the imported FW-H surface during the collected time in the `.trn` file. See Eqn. (980) for the formulation.

Ffowcs Williams-Hawkings Aeroacoustics Formulation

The Ffowcs Williams-Hawkings (FW-H) formulation is based on Farassat's Formulation 1A, that is the non-convective form of FW-H [151].

The FW-H equation is an exact rearrangement of the continuity and the momentum equations into the form of an inhomogeneous wave equation. The FW-H equation gives accurate results even if the surface of integration lies in the nonlinear flow region. It is based on the free-space Green's function to compute the sound pressure at the observer location, x. The FW-H equation for pressure that is radiated into a medium at rest by a flow in a region or a set of surfaces is:

$$p'(\grave{x}, t) = p'_T(\grave{x}, t) + p'_L(\grave{x}, t) + p'_Q(\grave{x}, t) \tag{971}$$

- $p'_T(\vec{x}, t)$ -- the monopole term is:

$$p'_T(\vec{x}, t) = \frac{1}{4\pi}\left(\left(-\frac{\partial}{\partial t}\right)\int_S \left[\frac{Q}{(r(1-M_r))}\right]_{ret} dS\right) \tag{972}$$

- $p'_L(\vec{x}, t)$ -- the dipole term is:

$$p'_L(\vec{x}, t) = \frac{1}{4\pi}\left(\left(-\frac{\partial}{\partial x_i}\right)\int_S \left[\frac{L_i}{(r(1-M_r))}\right]_{ret} dS\right) \tag{973}$$

- $p'_Q(\vec{x}, t)$ -- the quadrupole term is:

$$p'_Q(\vec{x}, t) = \frac{1}{4\pi}\left(\left(\frac{\partial^2}{(\partial x_i)(\partial x_j)}\right)\int_V \left[\frac{T_{ij}}{(r(1-M_r))}\right]_{ret} dV\right) \tag{974}$$

with:

$$Q = \rho_0 U_i n_i \tag{975}$$

$$U_i = \left(1 - \frac{\rho}{\rho_0}\right) v_i + \frac{\rho u_i}{\rho_0} \tag{976}$$

$$L_i = P_{ij} n_i + \rho u_i (u_n - v_n) \tag{977}$$

$$P_{ij} = (p - p_0)\delta_{ij} - \sigma_{ij} \tag{978}$$

$$T_{ij} = \rho u_i u_j + \delta_{ij}[(p - p_0) - c_0^2(\rho - \rho_0)] - \sigma_{ij} \tag{979}$$

where:

- u_i represents fluid velocity components in the x_i direction.
- u_n is the fluid velocity component normal to the surface.
- v_i represents surface velocity components in the x_i direction.
- v_n is the surface velocity component normal to the surface.
- n_i is the surface normal vector.
- σ_{ij} is the viscous stress tensor.
- ρ_0 is the far field density.
- P_{ij} is the compressive stress tensor.
- T_{ij} is the Lighthill stress tensor.

The space derivatives from Eqn. (972) and Eqn. (973) are transformed into time derivatives and afterwards, the time derivatives at the observer locations are moved into the integrals.

Eqn. (981) describes the monopole source term using the advanced-time formulation. Eqn. (982) describes the dipole source term using the advanced-time formulation.

FW-H Surface Terms

When the integration surface coincides with the body, these terms are called:

- $p'_T(\grave{x}, t)$ -- known as the Thickness Surface Term, resulting from the displacement of fluid as the body passes. The term is defined in Eqn. (981) for general flows and in Eqn. (986) for flows with rigid body motion or moving reference frames.
- $p'_L(\grave{x}, t)$ -- known as the Loading Surface Term, resulting from the unsteady motion of the force distribution on the body surface. The term is defined in Eqn. (982) for general flows and in Eqn. (987) for flows with rigid body motion or moving reference frames.
- $p'_S(\grave{x}, t)$ -- known as the Total Surface Term, resulting from the sum of the Thickness Surface Term and the Loading Surface Term. The term is defined in Eqn. (980).

$$p'_S(\grave{x}, t) = p'_T(\grave{x}, t) + p'_L(\grave{x}, t) \tag{980}$$

$$p'_T(\grave{x}, t) = \frac{1}{4\pi}\left(\int_{(f=0)} \left[\frac{\rho_0(\dot{U}_n + U_n)}{r(1-M_r)^2} \right]_{ret} dS + \int_{(f=0)} \left[\frac{\rho_0 U_n[r\dot{M}_r + a_0(M_r - M^2)]}{r^2(1-M_r)^3} \right]_{ret} dS \right) \tag{981}$$

$$p'_L(\grave{x}, t) = \frac{1}{4\pi}\left(\frac{1}{a_0} \cdot \int_{(f=0)} \left[\frac{\dot{L}_r}{r(1-M_r)^2} \right]_{ret} dS + \int_{(f=0)} \left[\frac{(L_r - L_M)}{r^2(1-M_r)^2} \right]_{ret} dS + \frac{1}{a_0} \cdot \int_{(f=0)} \left[\frac{L_r[r\dot{M}_r + a_0(M_r - M^2)]}{r^2(1-M_r)^3} \right]_{ret} dS \right) \tag{982}$$

and:

$$M_i = U_i / a_0 \tag{983}$$

$$r = x_{\text{observer}} - y_{\text{face}} \tag{984}$$

where:

- $f = 0$ denotes a mathematical surface to embed the exterior flow problem $f > 0$ in an unbounded space.
- $f = 0$ represents the emission surface and is made coincident with a body, impermeable surface, or permeable surface.

If the data surface coincides with a solid surface, then the normal velocity of the fluid is the same as the normal velocity of the surface:

$$u_n = v_n \tag{985}$$

In this case, Eqn. (981) and Eqn. (982) correspond to the Impermeable FW-H Surface type and some of the terms are eliminated.

Surface Terms for Rigid Body Motion and Moving Reference Frames

The most prolific proponent of time-domain methods for propellers and rotors has been Farassat. Farrassat's Formulation 1A was used in both ANOPP [188] and DFP-ATP [159] for subsonic source regions to compute the acoustic pressure $p'(t)$. The same Formulation 1A is coded in the STAR-CCM+, where an aerodynamic noise prediction can be processed parallel to the aerodynamic computation. The acoustic pressure $p'(t)$ is computed for Impermeable formulation, 3D, rotating sliding mesh (and similar for FW-H steady with rotating reference) using Eqn. (986) and Eqn. (987) based on [168].

For the Permeable FW-H Surface type, the FW-H Surface integrals are computed on the internal interface boundary, from an in-place interface.

$$p'_T(\vec{x}, t) = \frac{1}{4\pi}\left(\int_{(f=0)} \left[\frac{\rho_0 v_n [r(r\dot{M}_r) + a_0(M_r - M^2)]}{r^2(1 - M_r)^3} \right] dS \right) \tag{986}$$

$$\begin{aligned} p'_L(\vec{x}, t) = \frac{1}{4\pi}\Bigg(& \frac{1}{a_0} \cdot \int_{(f=0)} \left[\frac{\dot{L}_i r_i}{r(1 - M_r)^2} \right] dS \\ & + \int_{(f=0)} \left[\frac{(L_r - L_i M_i)}{r^2(1 - M_r)^2} \right] dS \\ & + \frac{1}{a_0} \cdot \int_{(f=0)} \left[\frac{L_r [r(r\dot{M}_r) + a_0(M_r - M^2)]}{r^2(1 - M_r)^3} \right] dS \Bigg) \end{aligned} \tag{987}$$

and:

$$\dot{M}_r = \dot{M}_i r_i \tag{988}$$

where:

- L_i is the blade load vector.
- $\dot{L}_i$ is the time derivative of the blade load vector with respect to source time.
- v_i is the local velocity of the blade surface with respect to the quiescent fluid.
- v_n is the surface normal velocity.
- r_i is the distance from a source point to the observer.
- $\dot{L}_i$ is the time derivative of the blade load vector with respect to source time.
- M_r is the Mach number of the source toward the observer.

Accounting for Time-Lag

The major task in evaluating the FW-H integrals is how to account for the time-lag between emission and reception times.

The advanced time algorithm looks forward in time to see when the observer perceives the currently generated sound waves:

- The procedure starts with a sequence of emission times (conveniently taken as the flow times).
- The source strengths are calculated (thickness surface noise and loading surface noise) at all source elements (faces of the integration surfaces) for a given emission time.
- The contribution of the sources is interpolated in the far-field time domain to build the sound signal.

The total sound pressure the observer perceives consists of the contribution from all source elements. The sound pressure at the receiver is obtained by accumulating the arriving signals in time slots. The overall observer acoustic signal is found from the summation of the acoustic signal from each source element of the FW-H surface during the same source time.

FW-H Volume Terms

The quadrupole noise is a volume distribution of sources, which accounts for nonlinearities in the flow. These nonlinearities are of two types, as Lighthill described [173]. First, the local speed of sound is not constant, but varies due to particle acceleration. Second, the finite particle velocity near the body (for example blade) influences the velocity of sound propagation. When strong shear layers exist in the flow or when the Mach number increases, the quadrupole term is not negligible.

Farassat and Brentner [151] have shown that the noise contribution from the quadrupole, $p'_Q(\grave{x}, t)$, may be expressed as a "collapsing-sphere" formulation. Using this formulation, the space derivatives from Eqn. (974) are transformed into time derivatives:

$$p'_Q(\grave{x}, t) = \frac{1}{4\pi}\left(\left(\frac{1}{c}\right)\left(\frac{\partial^2}{\partial t^2}\right)\int_{-\infty}^{t}\left[\int_{(f>0)}\frac{T_{rr}}{r}d\Omega\right]d\tau + \left(\frac{\partial}{\partial t}\right)\int_{-\infty}^{t}\left[\int_{(f>0)}\frac{3T_{rr}-T_{ii}}{r^2}d\Omega\right]d\tau\right) + \left(c\int_{-\infty}^{t}\left[\int_{(f>0)}\frac{3T_{rr}-T_{ii}}{r^3}d\Omega\right]d\tau\right) \tag{989}$$

where:

- T_{rr} is the double contraction of T_{ij}.
- r_i and r_i are the components of the unit vector in the direction of radiation.
- T_{ij} is the Lighthill stress tensor.

Eqn. (989) is transformed from a collapsing-sphere formulation to an advanced time formulation using equations (990), (991), (992) and (993). In these equations, the time derivatives at the observer are moved into the integrals to prevent numerical time differentiation of the integrals. The "source-time-dominant" algorithm from [154] is used to allow the estimation of the $p'_Q(\grave{x}, t)$ volume term of the FW-H equation as follows:

$$p'_Q(\grave{x}, t) = \frac{1}{4\pi}\left(\int_{(f>0)}\left[\frac{K_1}{c^2 r} + \frac{K_2}{cr^2} + \frac{K_3}{r^3}\right]_{ret} dV\right) \tag{990}$$

$$K_1 = K_{11} + K_{12} + K_{13} \tag{991}$$

$$= \left[\frac{\ddot{T}_{rr}}{(1-M_r)^3}\right] + \left[\frac{\ddot{M}_r T_{rr} + 3\dot{M}_r \dot{T}_{rr}}{(1-M_r)^4}\right] + \left[\frac{3\dot{M}_r^2 T_{rr}}{(1-M_r)^5}\right]$$

$$K_1 = K_{21} + K_{22} + K_{23} + K_{24} = \tag{992}$$

$$\left[\frac{-\dot{T}_{ii}}{(1-M_r)^2}\right] + \left[\frac{4\dot{T}_{M_r} + 2T_{\dot{M}_r} + \dot{M}_r T_{ii}}{(1-M_r)^3}\right]$$

$$+ \left[\frac{3[(1-M^2)\dot{T}_{rr} - 2\dot{M}_r T_{Mr} - M_i \dot{M}_i T_{rr}]}{(1-M_r)^4}\right]$$

$$+ \left[\frac{6\dot{M}_r (1-M^2) T_{rr}}{(1-M_r)^5}\right]$$

$$K_3 = K_{31} + K_{32} + K_{33} \tag{993}$$

$$= \left[\frac{2T_{MM} - (1-M^2)T_{ii}}{(1-M_r)^3}\right] - \left[\frac{6(1-M^2)T_{Mr}}{(1-M_r)^4}\right] + \left[\frac{3(1-M^2)^2 T_{rr}}{(1-M_r)^5}\right]$$

and:

$$T_{rr} = T_{ij} r_i r_j \tag{994}$$

$$T_{MM} = T_{ij} M_i M_j \tag{995}$$

$$T_{Mr} = T_{ij} M_i r_j \tag{996}$$

$$T_{\dot{M}r} = T_{ij} \dot{M}_i r_j \tag{997}$$

$$\dot{T}_{Mr} = \dot{T}_{ij} M_i r_j \tag{998}$$

$$\dot{T}_{rr} = \dot{T}_{ij} r_i r_j \tag{999}$$

$$\ddot{T}_{rr} = \ddot{T}_{ij} r_i r_j \tag{1000}$$

$$M_i = \frac{v_i}{c} \tag{1001}$$

$$M_r = M_i r_i \tag{1002}$$

$$\dot{M}_r = \dot{M} r_i \tag{1003}$$

$$\ddot{M}_r = \ddot{M}_i r_i \tag{1004}$$

$$\ddot{M}_{rr} = \ddot{M} r_i r_j \tag{1005}$$

$$r_i = \frac{x_i - y_i}{r} \tag{1006}$$

where:

- r_i denotes the unit vector in the direction of radiation.
- A dot above a variable denotes the time derivative with respect to source time of that variable.

Using the Lighthill Stress Tensor

This section explains the following:

- What is the Lighthill Stress Tensor?
- Selecting the Lighthill Stress Tensor

Also included is formulation and a reference section on related field functions.

What Is the Lighthill Stress Tensor?

In STAR-CCM+ the Lighthill Stress Tensor is purely for post-processing. Although it is represented by the Lighthill Stress Tensor node as a continuum model, it does not actually model acoustics in its own right. This object has properties and is subject to common menu items.



This optional feature interrogates the flow field and compiles velocity derivatives which make up the Lighthill Stress Tensor. It is a scalar with vector components. By observing Lighthill results, it is possible to identify the location of flow-generated aeroacoustics sources.

Lighthill Stress Tensor Model Properties

Far-Field Sound Speed	The sound in the far-field: $a_0 = \sqrt{(\gamma R T_0)}$ The default value is 340 m/s.

Selecting the Lighthill Stress Tensor

The Lighthill Stress Tensor is available for steady and unsteady simulations, and with any viscous regime. To make the Lighthill Stress Tensor option available, first select the Aeroacoustics model from among the *Optional Models* group.



Lighthill Stress Tensor Formulation

Lighthill's inhomogeneous wave equation governing the sound propagation is derived from conservation of mass and momentum:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i}(\rho v_i) = 0 \tag{1007}$$

$$\frac{\partial}{\partial t}(\rho v_i) + \frac{\partial}{\partial x_j}(\rho v_i v_j + p_{ij}) = 0 \tag{1008}$$

where:

- ρ is density
- v_i and v_j are the velocity components
- p_{ij} is the stress tensor.

$$p_{ij} = -\sigma_{ij} + (p - p_0)\delta_{ij} \tag{1009}$$

where

- p is the thermodynamic pressure of the flow field,
- δ_{ij} the Kronecker delta
- σ_{ij} is the viscous stress tensor.

$$\sigma_{ij} = \mu\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} - \frac{2}{3}\frac{\partial v_i}{\partial x_i}\delta_{ij}\right) \tag{1010}$$

Using Einstein notation, Lighthill's equation can be written as:

$$\frac{\partial^2 \rho}{\partial t^2} - c_0^2(\nabla^2 \rho) = \frac{\partial^2 T_{ij}}{(\partial x_i)(\partial x_j)} \tag{1011}$$

Where T_{ij} is the "Lighthill turbulence stress tensor" for the acoustic field, and it is denoted by T.

$$T_{ij} = \rho v_i v_j + p_{ij} - c_0^2(\rho - \rho_0)\,\delta_{ij} \tag{1012}$$

substitute Eqn. (1009) into Eqn. (1012) to get:

$$T_{ij} = \rho v_i v_j - \sigma_{ij} + \delta_{ij}[(p - p_0) - c_0{}^2(\rho - \rho_0)] \tag{1013}$$

where:

- δ_{ij} is the Kronecker delta.
- σ_{ij} is the viscous stress tensor.

- ρ_0 is the far field density.

In laminar flow, the stress tensor is given by:

$$\sigma_{ij} = \sigma_{ij,\,\mathrm{lam}} = \mu\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} - \frac{2}{3}\frac{\partial v_i}{\partial x_i}\delta_{ij}\right) \tag{1014}$$

In turbulent flow, the stress tensor is given by:

$$\sigma_{ij} = \sigma_{ij,\,\mathrm{lam}} + \sigma_{ij,\,\mathrm{turb}} = \mu_{eff}\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} - \frac{2}{3}\frac{\partial v_i}{\partial x_i}\delta_{ij}\right) \tag{1015}$$

where the effective viscosity is $\mu_{eff} = \mu + \mu_t$, the sum of the laminar and turbulent viscosities.

In inviscid flow, the stress tensor is $\sigma_{ij} = 0$.

Each of the acoustic source terms from the Lighthill stress tensor T_{ij} in Eqn. (1013) may describe an significant role in the generation of sound as:

- $\rho v_i v_j$ -- describes unsteady convection of flow (or Reynold's Stress)
- σ_{ij} -- describes sound shear
- $\delta_{ij}(p - c_0^2\rho)$ -- describes nonlinear acoustic generation processes

Lighthill Stress Tensor Field Functions Reference

The following field functions are made available to the simulation when the Lighthill Stress Tensor is used.

Lighthill Stress Tensor XX	The sound the specific normal stress, shear, nonlinear acoustic generation processes, T_{xx}
Lighthill Stress Tensor XY	The sound the specific shear stress, shear, nonlinear acoustic generation processes, T_{xy}
Lighthill Stress Tensor XZ	The sound the specific shear stress, shear, nonlinear acoustic generation processes (activated only in three-dimensional simulations), T_{xz}
Lighthill Stress Tensor YY	The sound the specific normal stress, shear, nonlinear acoustic generation processes, T_{yy}
Lighthill Stress Tensor YZ	The sound the specific shear stress, shear, nonlinear acoustic generation processes (activated only in three-dimensional simulations), T_{yz}
Lighthill Stress Tensor ZZ	The sound the specific normal stress, shear, nonlinear acoustic generation processes (activated only in three-dimensional simulations), T_{zz}

Aeroacoustics Bibliography

[143] Aybay, O., Caraeni, M. and Holst, S. 2011."Tandem Cylinder and Idealized Side Mirror Far-Field Noise Predictions Using DES and An Efficient Implementation of FW-H Equation", 17th AIAA/CEAS Aeroacoustics Conference, AIAA, Portland, Oregon (USA).

[144] Bailly, C. and Juve, D. 1999. "A Stochastic Approach to Compute Subsonic Noise using Linearized Euler's Equations", *AIAA Journal*, AIAA-99-1872.

[145] Baris, O., and Mendonca, F. 2011. "Automotive Turbocharger Compressor CFD and Extension Towards Incorporating Installation Effects", *Proceedings of ASME Turbo Expo 2011: Power for Land, Sea, and Air*. **GT2011-46796**.

[146] Bechara, W., Bailly, C., Lafon, P., and Candel, S. 1994. "Stochastic approach to Noise Modeling for Free Turbulent Flows", *AIAA Journal*, **32**:3, pp. 455-463.

[147] Bechara, W., Lafon, P., Bailly, C., and Candel, S. 1994. "Application of a k-epsilon turbulence model to the prediction for simple and coaxial free jets", *Journal Acoustics Soc. Am*, **97**:6, pp. 3518-3531.

[148] Bodony, D. J. 2006. "Analysis of Sponge Zones for Computational Fluid Mechanics", *Journal of Comp. Physics*, **212-2**, pp. 681-702.

[149] Bres, G.A., Brentner, K.S., Perez, G. and Jones, H.E. 2004. "Maneuvering rotorcraft noise prediction", *J Sound and Vibration*, **275**:3-5, pp. 719-738.

[150] Brentner, K.S. 1994. "A new algorithm for computing acoustic integrals", *Proceedings of the 14th IMACS World Congress*.

[151] Brentner, K.S. and Farassat, F. 1998. "Analytical comparison of the acoustic analogy and Kirchhoff formulation for moving surfaces", *AIAA Journal*, **36**:8, pp. 1379-1386.

[152] Brentner, K.S. and Farassat, F. 2003. "Modeling aerodynamically generated sound of helicopter rotors", *Progress in Aerospace Sciences*, **39**:2-3, pp. 83-120.

[153] Caridi, D. 2008. "Industrial CFD Simulation of Aerodynamic Noise", Doctoral Thesis, University of Naples Federico II.

[154] Casalino, D. 2003. "An advanced time approach for acoustic analogy predictions", *J Sound and Vibration*, **261**:4, pp. 583-612.

[155] Colonius, T. 2004. "Modeling Artificial Boundary Conditions for Compressible Flow", *Annual Review of Fluid Mechanics*, **36**, pp. 315-345.

[156] Curle, N. 1955. "The Influence of Solid Boundaries upon Aerodynamic Sound", *Proceedings of the Royal Society of London, series A, Mathematical and Physical Sciences*, **231**, pp. 505-514.

[157] di Francescantonio, P. 1997. "A new boundary integral formulation for the prediction of sound radiation" *J. Sound and Vibration*, **202**, 491-509.

[158] Farassat, F. 2007. "Derivation of formulations 1 and 1A of Farassat", Technical Memorandum TM-2007-214853, NASA, NASA Langley Research Center (USA).

[159] Farassat F., Padula S.L. and Dunn M.H. 1987. "Advanced Turboprop Noise Prediction Based on Recent Theoretical Results" *J. Sound and Vibration*, **119**:1 pp. 53-79.

[160] Farassat F., Succi G. P. 1983. "The prediction of helicopter discrete frequency noise" *Vertica*, **7**:4 pp. 309-320.

[161] Ffowcs Williams, J. and Hawkings, D. 1969. "Sound Generation by Turbulence and Surfaces in Arbitrary Motion", *Philosophical Transactions of Royal Society of London, Series A, Mathematical and Physical Sciences (1934-1990)*, **264**:1151, pp. 321-342.

[162] Frota, J., Lempereur, P. and Roger, M. 1998. "Computation of the noise of a subsonic propeller at an angle of attack", 4th AIAA/CEAS Aeroacoustic Conference, AIAA, Toulouse (France).

[163] Gennaro, M. and Caridi, D. 2010. "Ffowcs Williams-Hawkings Acoustic Analogy for Simulation of NASA SR2 Propeller Noise in Transonic Cruise Condition", ECCOMAS CFD 2010.

[164] Greschner, B., Thiele, F., Jacob, M.C., and Casalino, D., 2008. "Prediction of sound generated by a rod-airfoil configuration using EASM DES and the genralise Lighthill/FW-H analogy", *Computers and Fluids*, **37**, pp. 402-413.

[165] Goldstein, M.E. 1976. *Aeroacoustics*, McGraw-Hill, New York.

[166] Goldstein, M.E. and Rosenbaum, B. 1973. "Effect of Anisotropic Turbulence on Aerodynamic Noise", *J. Acoustical Society of America*, **54**, pp. 630-645.

[167] Hinze, J. O. 1975. *Turbulence*, McGraw-Hill, pp. 672-673.

[168] Hubbard, H.H. 1995. "Aeroacoustics of Flight Vehicles, Theory and Practice, Volume 1: Noise Sources" NASA Langley Research Center, Hampton, Virginia, pp. 13.

[169] Karweit, M., Blanc-Benon, P., Juve, D., and Comte-Bellot, G. 1991. "Simulation of the Propagation of an Acoustic Wave Through a Turbulent Velocity Field: A Study of Phase Variance", *J. Acoustical Society of America*, **89**:1, pp. 52-62.

[170] Kim, J. W., Lau, A. S.H., Sandham, N.D. 2010. "Proposed Boundary Conditions for Gust-Airfoil Interaction Noise", *AIAA Journal*, **48-11**.

[171] Kraichnan, R.H. 1970. "Diffusion by a Random Velocity Field", *J. Computational Physics*, **13**:1, pp. 22-31.

[172] Lighthill, M. 1952. "On Sound Generated Aerodynamically. I. General Theory", *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1190)*, **211**:1107, pp. 564-587.

[173] Lighthill, M. 1954. "On Sound Generated Aerodynamically. I. Turbulence as a Source of Sound", *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1190)*, **222**:1148, pp. 1-32.

[174] Lilley, G.M. 1993. "The radiated noise from isotropic turbulence revisited", NASA Contract Report No. 93-75, NASA Langley Research Center.

[175] Lin, P.H. 2011. "Computational Aero- and Hydro-acoustics by Space-Time Conservation Element and Solution Element Method", Ph. D. Thesis, Ohio State University.

[176] Mankbadi, R.R. 1993. "Computational aero- and hydro-acoustics" ASME United Engineering Center.

[177] Marinus, B. and Roger, M. 2009. "Truncated Method for Propeller Noise Prediction Up to Low Supersonic Helical Tip Mach Numbers", AIAA-2009-3330.

[178] Mendonca, F., Baris, O., and Capon, G. 2012 "Simulation of Radial Compressor Aeroacoustics Using CFD.", Effects", *Proceedings of ASME Turbo Expo 2011: Power for Land, Sea, and Air*. **GT2012-70028**.

[179] Mendonça, F., Read, A., Caro, S., Debatin, K. and Caruelle, B. 2005. "Aeroacoustic Simulation of Double Diaphragm Orifices in an Aircraft Climate Cooling System", AIAA-2005-2976.

[180] Pilon A. and Lyrintzis A. 1998. "Development of an Improved Kirchhoff Method for Jet Aeroacoustics", *AIAA Journal*, **36**:5, pp.783-790.

[181] Proudman, I. 1952. "The Generation of Noise by Isotropic Turbulence", *Proc. R. Soc. Lond.* **214**, pp. 119-132.

[182] Ribner, H.S. 1964. "The Generation of Sound by Turbulent Jets", in *Advances in Applied Mathematics*, Academic Press, New York.

[183] Robinson, A.R. and Lee, D. 1994. "Oceanography and acoustics: prediction and propagation models", American Institute of Physics.

[184] Roger, M. 2006 "Fundamentals of aeroacoustics. In Experimental Aeroacoustics", M. L. Riethmuller and M. R. Lema editors, LS 2007-01. von Karman Institute, Rhode-St-Genese.

[185] Sarkar, S., Hussaini, M.Y. 1993. "Computation of the Sound Generated by Isotropic Turbulence", NASA Report 191543, ICASE Report No. 93-74.

[186] Witkowska, A., and Juve, D. 1993. "Numerical Simulation of Noise Generated by Homogeneous and Isotropic Turbulence,' Proc. 13th Colloquium on Aero- and Hydro-Acoustics, Lyon.

[187] Wagner, Hüttl, Sagaut (2007). "LES for Acoustics", Cambridge University Press, ISBN-13: 9780521871440.

[188] Zorumski W.E. and Weir D.S. 1986. "Aircraft Noise Prediction Program Theoretical Manual- Propeller Aerodynamics and Noise" *NASA TM-83199*, **3**.

Modeling Hydroacoustics

Hydroacoustics is the study of sound generation and propagation in water [191]. Sound comes from marine wildlife, by SONAR, and by ships [190]. A sound pressure wave moves through water at 1,484 m/s, over 4 times faster than the speed of sound in air. Also, due to slower dissipation sound can travel much longer distances underwater than is possible in air.

Setting up a Hydroacoustics Analysis

Modeling hydroacoustics in STAR-CCM+ is similar to the process for modeling aeroacoustics. Both types of simulation require that you model the full unsteady flow surrounding the geometry from which sound propagates [189]. However, for your hydroacoustics analysis, use the double precision version of STAR-CCM+ to simulate the generation of acoustic waves in turbulent water flows. Also, in hydroacoustics, use a compressible liquid such as IAPWS-IF97 (Water) to model the transmission of sound.

To model sound wave propagation through water:

- Launch the double precision version of STAR-CCM+.
- Choose an equation of state that supports compressibility.

  An IAPWS-IF97 (Water) model or a User Defined EOS model with compressibility enabled are both appropriate for hydroacoustics analyses.
- Use the recommended analysis procedure for aeroacoustics to define a suitable mesh and time-step and solve the case in both the near-field and far-field.

Caution: When setting up a hydroacoustics simulation pay special attention to the default acoustics properties, many of which use values for air. For instance, the Far-Field Sound Speed FW-H property is set to 340 m/s by default.

Hydroacoustics Bibliography

[189] Lin, P.H. 2011. "Computational Aero- and Hydro-acoustics by Space-Time Conservation Element and Solution Element Method", Ph. D. Thesis, Ohio State University.

[190] Mankbadi, R.R. 1993. "Computational aero- and hydro-acoustics" ASME United Engineering Center.

[191] Robinson, A.R. and Lee, D. 1994. "Oceanography and acoustics: prediction and propagation models", American Institute of Physics.

Modeling Combustion and Other Reactions

This chapter describes the current STAR-CCM+ facilities for modeling combustion and other chemical reactions. The fundamental principles are described in an introductory section. Details are also given on the following specific model types:

- Eddy Break-Up (EBU) models
- The Premixed Eddy Break-Up model
- The Eddy Contact Micromixing model
- The Surface Chemistry model
- Homogeneous Reactor models
- Turbulence-chemistry interactions (TCI)
- The Coherent Flame model (CFM)
- The Partially-Premixed Coherent Flame Model (PCFM)
- The Turbulent Flame Speed Closure (TFC) model
- The Partially-Premixed Turbulent Flame Speed Closure (PTFC) Model
- Presumed Probability Density Function (PPDF) models
- The Progress Variable model (PVM)
- The Thickened Flame model
- Nox model
- Fuel Nox model
- Coal Combustion modeling
- Soot modeling
- Particle Reaction modeling

- Polymerization modeling

Information is also provided about the following:

- Working with the fuel mass fraction
- Defining chemical reactions
- Post-processing reaction rates
- Importing Chemkin-formatted data
- Using ignitors
- Using fluid streams
- Using laminar flame speed methods
- Field functions

Using Combustion Models

This section contains:

- A description of basic reacting flow, combustion flame type, and combustion model concepts.
- A brief guide on how to choose a model that is appropriate to a specific application.
- An outline of the combustion model selection process with the STAR-CCM+ GUI.

What Is a Reacting Flow System?

A system containing a multi-component fluid mixture whose constituents react chemically with each other. A node representing such a system (Reacting in the following screenshot) appears in the simulation tree when this type of flow problem is specified in the *Physics Model Selection* dialog.



What Are Combustion Flame Types?

A combustion flame type is the object that represents how you choose to model your combustion: premixed, non-premixed or partially premixed.



This choice becomes available after you specify a reacting flow problem in the *Physics Model Selection* dialog. In turn, the combustion flame type that you select affects the availability of specific combustion models.

What Is a Combustion Model?

Combustion models calculate the reaction state space: the concentrations of the various species present in a chemical reaction, and the quantities they influence: density, viscosity, and temperature. To this end, STAR-CCM+ provides several models that are based on various physical and chemical approximations. The general principles underlying these models are described below.

A large chemical reaction set, such as found in hydrocarbon combustion, can span a wide range of time scales. In addition to these time scales, the turbulent flow field imposes its own limits on length and time scales. These scales range from the Kolmogorov and Batchelor time scales at the low end to the large, energy-containing eddy time scales at the high end. Resolving all the length and time scales affecting the grid-mean properties in a reacting flow system demands computational resources beyond the currently available resources. Therefore, combustion models are required to account for the processes that occur at length and time scales below what can be resolved on a numerical simulation grid.

The four basic types of combustion model currently available in STAR-CCM+ are:

- Eddy break-up (EBU) models

- The Homogeneous Reactor model
- The Coherent Flame model (CFM)
- Presumed Probability Density Function (PPDF) models

The choice of combustion model is decided by knowing the Damkohler number, that is defined as:

$$Da = \frac{t_{mix}}{t_{rxn}} \tag{1016}$$

where t_{mix} is the mixing time scale and t_{rxn} is the reaction time scale.

When the Damkohler number is large, the turbulent mixing that brings reactants together at the molecular scale controls the reaction rate. In this limit, the standard EBU and the equilibrium PPDF models are fairly accurate because they assume that the reaction occurs instantaneously upon micromixing.

The PPDF model parameterization depends on whether the flow is assumed to be locally adiabatic or non-adiabatic:

- Under the adiabatic assumption, only the scalar mixing process and the initial temperature of the control volume affect chemical equilibrium within a control volume.
- Under the non-adiabatic assumption, chemical equilibrium within a control volume can also change due to heat losses or gains. This assumption requires the solution of a transport equation for enthalpy. Typically, the most important contributor to this loss or gain is radiative heat transfer in participating media.

When the Damkohler number is of order 1, finite rate kinetics must be considered. The EBU model can be extended to account for these effects using reaction rates from finite-rate kinetics that are based on mean species concentrations and mean temperatures. The adiabatic PPDF model can be extended to account for finite-rate kinetic effects by using the laminar flamelet approach instead of the equilibrium approach. Currently, only the adiabatic flamelet method has been implemented.

How Do I Decide on a Combustion Model?

The EBU approach is more general than PPDF as it can be used for premixed, non-premixed, and partially premixed flames.

Within the EBU model family, if kinetic data exist, you can choose between the combined time scale model and hybrid model because they account for finite-rate kinetic effects. However, if the problem is purely turbulence-controlled, the standard EBU model is adequate and efficient, since no finite-rate calculations are required.

The Homogeneous Reactor model offers the most general approach in modeling finite-rate kinetics. However, it is computationally expensive as the chemical step is solved through the ordinary differential equation (ODE) which is stiff. Therefore, use this model if other models do not capture the finite-rate effects that drive the process. If the problem is transient, then either the Operator Splitting model or the Chemistry ADI model can be used.

For time-steps smaller than the local CFL number, the operator splitting algorithm is preferred for computational efficiency reasons, as the algorithm separates flow-field and chemistry time scales. If the problem is steady-state, use the Chemistry ADI model, as it is less sensitive to the pseudo-time-step used to integrate the ODE for the chemistry step.

For diffusion flames, PPDF methods are preferred because they account for subgrid heterogeneity due to turbulence. The flamelets method is better than the equilibrium method for non-premixed flames because it gives a better representation of the underlying physics. However, this method requires creation of a suitable flamelets table. The non-adiabatic PPDF option is necessary if significant heat losses, such as radiation, are present.

The current CFM combustion model is designed for premixed flames where the fuel and oxidizer are perfectly mixed before entering the computational domain. Other combustion models in STAR-CCM+ are designed primarily for non-premixed flames.

Selecting a Combustion Model

To use the STAR-CCM+ facilities for solving combustion problems:

- In the *Physics Model Selection* dialog, select either a Multi-Component Gas or a Multi-Component Liquid material model from the *Material* group.
- Specify that your problem requires chemical reaction simulation by selecting Reacting in the *Reaction Regime* group.

- Specify whether your combustion is premixed, non-premixed or partially premixed by selecting the corresponding radio button in the *Combustion* group.



The availability of specific combustion modeling options depends on this choice and on the material model that was chosen earlier. Choose a turbulence model and any optional models that reflect their compatibility with your choice of combustion model.

Accessing Combustion Models

After selecting a combustion flame type, you will be presented with a choice of combustion models valid for use with the flame type. If the combustion model you wish to use is not available for the flame type, deactivate the flame type from the *Enabled Physics Models* section and choose the correct alternative. You can switch between combustion models that are presented for a particular flame type without this step. For example, when you choose Homogeneous Reactor, you can switch among the three flame types even if you do not deactivate the combustion model. The relationships among material models, combustion flame types, and specific combustion modeling options are outlined in the following table.

		Material	
		Multi-Component Gas	Multi-Component Liquid
Combustion	Non-Premixed	• EBU, including Ebu coefficients, Arrhenius coefficients, and the user reaction coefficient • Homogeneous Reactor • PPDF (both variants) • PVM • TFM	• EBU, including Ebu coefficients and the user reaction coefficient • Homogeneous Reactor • Polymerization
	Partially-Premixed	• EBU, including Ebu coefficients, Arrhenius coefficients, and the user reaction coefficient • Homogeneous Reactor • PCFM (both variants) • PTFC (both variants) • PVM • TFM	• EBU, including Ebu coefficients and the user reaction coefficient • Homogeneous Reactor • Polymerization
	Premixed	• PEBU • Homogeneous Reactor • CFM • TFC • TFM • PVM	

Most of the other features available with gas phase reactions, such as the user-defined species source and specific heat with NASA polynomials, are available with liquid phase reactions as well.

Combustion Models and Turbulence

The following table outlines which turbulence models are appropriate for combustion:

Turbulence Model	Combustion Model					
	Std Ebu	Kin Ebu	PPDF	DARS	CFM	PCFM
Spalart-Allmaras		X		X		
K-Epsilon	X	X	X	X	X	X
K-Omega	X	X	X	X	X	X
Reynolds Stress	X	X	X	X	X	X
Spalart-Allmaras DES		X		X		
K-Omega DES	X	X		X		
LES	X	X	X	X		

Specifically, the following restrictions apply to the use of turbulence models with combustion models:

- Standard EBU requires velocity and length scale.
- The following combustion models require velocity and length scale, and can only be used with Reynolds-Averaged Navier-Stokes turbulence:
  - CFM
  - PCFM

Combustion Models and Optional Models

- Lagrangian evaporation (Quasi-Steady and User-Defined) models work only with Non-Premixed and Partially-Premixed Combustion models. This restriction is because, in principle, a system is not premixed when Lagrangian evaporation occurs.
- Lagrangian evaporation models work only with the Non-Adiabatic option for PPDF and PCFM models.

Using Eddy Break-Up (EBU) Models

This section describes:

- The general concept of an EBU combustion model for non-premixed and partially premixed combustion
- How to select an EBU model
- The mathematical formulation of the available EBU model variants

Other sections describe:

- The Premixed Eddy Break-Up model
- Setting the EBU ignitor

Using the EBU models requires the presence of multiple species, which are defined as described in the section, Working with Species.

What Is an EBU Model?

EBU combustion models track individual mean species concentrations on the grid through transport equations. The reaction rates that are used in these equations are calculated as functions of the mean species concentrations, turbulence characteristics and, depending on the specific model that is used, temperature. A mean enthalpy equation is solved in addition to the species transport equations. The mean temperature, density, and viscosity are then calculated knowing the mean enthalpy and species concentrations.

The corresponding node (Eddy Break-up in the following screenshot) is activated as soon as a combustion simulation using the EBU model is selected.



Four variants are available for this model, selectable from the Reaction Control property in the *Properties* window.

EBU Model Properties

Reaction Control	Selects the variant of EBU combustion model.	
	Standard EBU	Selects Standard EBU — assumes that the reaction rate is dictated solely by the turbulent mixing time scale.
	Hybrid	Selects Hybrid — assumes that the reaction rate is also affected by finite-rate chemical kinetics.
	Combined Time Scale	Selects Combined Time Scale — makes the same assumption as the hybrid model but implements it in a different way.
	Kinetics Only	Selects Kinetics Only — assumes that the reaction rate is dictated solely by finite-rate chemical kinetics. This option is the ***only*** valid choice if the flow is laminar.
Store Reaction Rates	Activates the capability to post-process reaction rates, forward and backward (where applicable).	
	Activated	Stores reaction rates during the simulation run and makes them available among scalar functions.
	Deactivated	Does not store reaction rates (default).
Begin	Describes the number of iterations or time-steps when the sources from combustion become enabled.	
Source Enabled Trigger	Controls the combustion sources.	
	Never	Leaves combustion sources off. The Begin property is inactive if this option is selected.

Iteration	Turns on combustion sources. This option is available in steady-state cases.
Timestep	Turns on combustion sources. This option is available in unsteady-state cases.

Expert Properties

Source Term Limiting	For the EBU combustion model (particularly with the Kinetics Only option), the specified reaction rate can be large enough to drive the reactant concentration below zero. To keep all reactant concentrations in a physically correct range, activate this limit on the final reaction rate that is applied to the species transport equation. For most applications, leave this option deactivated. Use it only when a reasonable convergence cannot be obtained and the species concentrations clearly show unphysical values.

Activated	Limits the final reaction rate that is applied to the species transport equation.
Deactivated	Does not limit the final reaction rate (default).

Selecting the EBU Model

This procedure is only for non-premixed and partially premixed combustion. When you select premixed combustion, you work with the Premixed Eddy Break-Up model.

To solve combustion problems using an EBU model:

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas, but Multi-Component Liquid is also acceptable.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Specify whether your combustion is non-premixed or partially premixed by selecting the corresponding radio button in the *Combustion* group.

- Select the Eddy Break-up option from the particular combustion group, as shown below.



Choose the desired solver type. For the coupled flow model:

- Select Coupled Flow in the *Flow* group.

  The Coupled Energy and Coupled Species models, along with the Ideal Gas (with combustion) model in the *Equation of State* group box, are selected automatically.

The *Physics Model Selection* dialog looks as shown below.



Alternatively, to choose the segregated flow model:

- Select Segregated Flow in the *Flow* group.

  The Segregated Species and Segregated Fluid Enthalpy models, along with the Ideal Gas (with Combustion) model in the *Equation of State* group box, are selected automatically.

The *Physics Model Selection* dialog looks as shown below.



As an alternative to Ideal Gas, you can use Real Gas models with the EBU model:

- Before choosing an option from the *Flow* group, deactivate the Auto-select recommended Physics models checkbox.
- To make the *Equation of State* group box visible, select a species option.
- Select a Real Gas model.

  Use the Thermodynamic Polynomial Data material properties method with this combination of models.

Real Gas is helpful in combustion applications that include high pressure (such as liquid rockets) and low temperature (transcritical and supercritical environments). These applications involve considerable departure from ideal-gas behavior.

The remaining selections in this dialog depend on problem characteristics that are not specifically related to simulating combustion:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- Make an appropriate selection from the *Viscous Regime* group. In most cases, this choice is Turbulent, as most of the available combustion models are only valid for turbulent flows.
- If radiative heat transfer is important, select one of the options in the *Radiation* group.
- If gravity forces significantly influence the solution, for example in fire simulations, select Gravity from the *Optional Physics models* group.

The choice of the desired EBU model variant is made within the Eddy Break-up node appearing in the simulation tree once this type of model is selected.

EBU Formulation

EBU models characterize the reacting flow system by a specified number of species and chemical reactions.(see Finite-Rate Kinetics Convention). Three variants of this type of model are currently available in STAR-CCM+:

- The standard EBU model which assumes that the rate of reaction is dictated solely by the turbulent mixing time scale.
- The hybrid kinetics/EBU model which assumes that the rate of reaction is also affected by the reaction rate from finite rate kinetics.
- The combined time scale model which also makes the same assumption as the hybrid model.

Finite-Rate Kinetics Convention

A reacting flow system can be characterized by a set of chemical reactions that can be written as:

$$\sum_i \nu_{ij} C_i = 0 \quad , \qquad \forall j = 1, \ldots\ldots, n_R \tag{1017}$$

in which n_R is the number of reactions and ν_{ij} the stoichiometric coefficient of species i in reaction j. For example, for a system of two reactions,

$$CH_4 + 2O_2 \rightarrow CO_2 + 2H_2O \tag{1018}$$

$$C + O_2 \rightarrow CO_2 \tag{1019}$$

the stoichiometric coefficients can be written as

Species	Reaction 1	Reaction 2
CH_4	1	0
O_2	2	1
C	0	1
CO_2	1	1
H_2O	2	0

Each reaction proceeds at its own rate hence, at a given location, it has a unique time scale that depends on the concentrations of different species at that location.

In the STAR-CCM+ GUI, all stoichiometric coefficients are entered as positive numbers.

The Standard EBU Model

In this model, individual species in the global reaction are assumed to be transported at different rates according to their own governing equations. These equations are derived from the instantaneous governing equations for species i:

$$\frac{\partial}{\partial t}(\rho\chi Y_i) + \nabla\bullet(\rho U Y_i - F_i) = S_i \tag{1020}$$

in which F_i is the diffusion flux, and is defined as:

$$F_i = D_i \nabla Y_i \tag{1021}$$

where D_i is the molecular (multi-component) diffusivity of species i, and S_i is the reaction rate of species i.

The standard EBU model uses the time-average of the instantaneous governing equations:

$$\frac{\partial}{\partial t}(\rho\chi \bar{Y}_i) + \nabla\bullet(\rho U \bar{Y}_i - \bar{F}_i) = \bar{S}_i \tag{1022}$$

where $\bar{F}_i$ now includes the effect of turbulent diffusion and is defined as:

$$\bar{F}_i = \left(D_i + \frac{\mu_t}{Sc_i}\right)\nabla \bar{Y}_i \tag{1023}$$

where μ_t is the turbulent diffusivity and Sc_i is the turbulent Schmidt number.

The reaction rate is modeled through an expression that takes the turbulent micromixing process into account. This is done through dimensional arguments. Thus, for a reaction of the form:

$$\nu_F F + \nu_O O \rightarrow \nu_{P1} P_1 + \nu_{P2} P_2 + \ldots\ldots + \nu_{Pj} P_j \tag{1024}$$

the rate of fuel depletion is assumed to be:

$$R_F = -\frac{\rho}{M_F}\left(\frac{1}{\tau_R}\right) A_{ebu} \min\left[\bar{Y}_F, \frac{\bar{Y}_O}{s_O}\right] \quad \text{moles}/(\text{m}^3\text{s}) \tag{1025}$$

or, optionally:

$$R_F = R_{F,\,mix} = -\frac{\rho}{M_F}\left(\frac{1}{\tau_R}\right) A_{ebu} \min\left[\bar{Y}_F, \frac{\bar{Y}_O}{s_O}, B_{ebu}\left(\frac{\bar{Y}_{P1}}{s_{P1}} + \frac{\bar{Y}_{P2}}{s_{P2}} + \ldots + \frac{\bar{Y}_{Pj}}{s_{Pj}}\right)\right] \tag{1026}$$

In the above equation,

$$s_O \equiv \nu_O M_O / \nu_F M_F \tag{1027}$$

$$s_{Pi} \equiv |\nu_{Pi}| M_{Pi} / \nu_F M_F \tag{1028}$$

and

$$\tau_R = \tau_{mix} = \frac{k}{\varepsilon} \tag{1029}$$

where k is the turbulent kinetic energy and ε its dissipation rate. The min operator on the right-hand side of Eqn. (1025) indicates that the concentration of the limiting reactant should be used to determine a mass fraction scale when calculating the reactant consumption rate. Eqn. (1025) essentially states that the integrated micromixing rate is proportional to the mean (macroscopic) concentration of the limiting reactant divided by the time scale of the large eddies. Eqn. (1026) is an optional modification to Eqn. (1025), in which the inhibition of the forward reaction given by Eqn. (1024) at low temperatures is accounted for by a low product concentration. In the EBU model, $\bar{Y}_F$, $\bar{Y}_O$, and $\bar{Y}_{Pi}$ evolve by means of individual transport equations of the form of Eqn. (1022), each with their own diffusivities (binary or multicomponent). The reaction rate that is obtained from Eqn. (1025) or Eqn. (1026) is also used, with the proper stoichiometric coefficients, for each species. For example, for product P_1, the reaction rate is:

$$S_{P1} = M_{P1} \nu_{P1} R_F \tag{1030}$$

The EBU model can be used for multiple reactions. For each reaction j, a turbulence-limited reaction rate R_{Fj} is calculated based on its own reactants and products, as in Eqn. (1026). The reaction source term for each species can then be obtained from:

$$S_i = M_i \sum_{j=1}^{n_R} \nu_{ij} R_j \quad \mathrm{kg/(m^3 s)} \tag{1031}$$

EBU Model Extensions to Account for Finite-Rate Kinetics

As mentioned earlier, the EBU model assumes that the reaction rate is determined solely by the turbulent mixing rate. There have been attempts in the literature to add the effects of chemical kinetics to the eddy break-up model. Before describing these, we need to explain how chemical reaction rates are obtained from finite-rate kinetics.

Chemical Reaction Rates from Finite-Rate Kinetics

The reaction rate of each reaction can be written down (from the point of view of a reactant) as a function of composition and temperature in the modified Arrhenius form:

$$R_j = R_{i,kin} = -A_j T^{\beta_j} \prod_{all\ \mathrm{reactants}} \left(\frac{\rho Y_i}{M_i}\right)^{p_{ij}} e^{-E_{aj}/R_u T} \quad \mathrm{kmol/(sm^3)} \tag{1032}$$

In the above equation, A_j, β_j and E_{aj} are the pre-exponential factor, temperature exponent, and activation energy for the jth reaction, respectively, and R_u is the universal gas constant. The p_{ij} exponents of the reacting species i are the same as ν_{ij} when the reactions are elementary, as in the case of full complex chemistry. When a reduced or a global chemical mechanism is used instead of the full known mechanism, the p_{ij} exponents are different from the stoichiometric coefficients.

The expression

$$\prod_{all\ \mathrm{reactants}} \left(\frac{\rho Y_i}{M_i}\right)^{p_{ij}}$$

has units of $[(\mathrm{kmol/(m^3)})^{P_j}]$ where P_j is the sum of all p_{ij}.

User Reaction Coefficient

The user-defined reaction coefficient K_f is a profile--it can be given as a constant, field function, table, or user-defined subroutine. For a particular reaction $\mathbf{R}$, given K_f as the user-defined reaction coefficient, the overall reaction rate is defined as:

$$R_j = R_{i,kin} = K_f \bullet \prod_{all\ \text{reactants}} \left(\frac{\rho Y_i}{M_i}\right)^{p_{ij}} \quad \text{kmol}/(\text{sm}^3) \tag{1033}$$

where ρ is density, Y_i is the mass fraction of species i, M_i is its molecular weight, and p_{ij} is its rate exponent.

The expression

$$\prod_{all\ \text{reactants}} \left(\frac{\rho Y_i}{M_i}\right)^{p_{ij}}$$

has units of $[(\text{kmol}/(\text{m}^3))^{N_j}]$ where N_j is the sum of all p_{ij}.

From here, the source term of species i in reaction $\mathbf{R}$ is defined as:

$$\Omega_i^R = M_i \bullet \nu_{ij} \bullet R_j \tag{1034}$$

The overall source term of species i is defined as a summation of source terms of species i for all reactions in the system:

$$\Omega_i = \sum_R \Omega_i^R \tag{1035}$$

For linearizing the species source term, two derivative terms (listed below) are used. For both derivatives, the assumption is that the user-defined reaction coefficient K_f is not in the functional relationship with either temperature (for the coupled solver only) or mass fractions of species.

- For the coupled solver only, the species source term with respect to temperature is:

  $$\frac{\partial \Omega_i}{\partial T}$$

  For the case where the user-defined reaction coefficient is used, this term is zero. If the user-defined reaction coefficient contains temperature, this term is not calculated properly.

- The source term of species i with respect to the mass fraction of species j is:

$$\frac{\partial \Omega_i}{\partial Y_j}$$

If the user-defined reaction coefficient contains mass fractions, this term is not calculated properly.

Hybrid Kinetics/EBU Model

This model accounts for finite rate effects by assuming that the actual reaction rate R_i, used in Eqn. (1031), is the minimum of the reaction rates from Eqn. (1032), with $\bar{Y}$, $\bar{T}$, and $\bar{\rho}$, in place of Y_i, T, and ρ, and Eqn. (1025) or Eqn. (1026), as the case may be. This assumption can be expressed as:

$$R_i = -min(\left|R_{i,\,kin}(\bar{\rho}, \bar{Y}_1, \bar{Y}_2, \bar{Y}_3, \ldots, \bar{Y}_N, \bar{T})\right|, \left|R_{i,\,mix}\right|) \tag{1036}$$

Combined Time Scale Model

The turbulent mixing time scale, τ_{mix}, used in the standard EBU model, decreases with decreasing distance from solid surfaces. This leads to over-prediction of reaction rates in near-wall regions. To alleviate this problem, the mixing time scale is augmented by a time scale that is derived from the chemical reaction rate from finite-rate kinetics, Eqn. (1032), defined as:

$$\tau_{kin} = \frac{\rho Y_F}{M_F}\left(\frac{1}{\left|R_{i,\,kin}(\bar{\rho}, \bar{Y}_1, \bar{Y}_2, \bar{Y}_3, \ldots, \bar{Y}_N, \bar{T})\right|}\right) \tag{1037}$$

As a result, the time scale τ_R used in the standard EBU model (Eqn. (1025) or Eqn. (1026)) can now be modified to:

$$\tau_R = \tau_{mix} + \tau_{kin} \tag{1038}$$

Clearly, the reaction rate predicted by using this model is smaller than that from the standard EBU model. This may lead to difficulty in sustaining the reaction during the numerical simulation.

Finite-Rate Kinetics Nomenclature

The following is a guide to the terms and units that are used in this section.

R_j	final reaction rate [kmol/(s m^3)]
A_j	pre-exponential factor $[([\text{kmol}/(\text{m}^3)]^{(1-P_j)})/\text{s}]$ The units for the pre-exponent A_j change according to the order of the reaction (sum of p_{ij}).
T^{β_j}	temperature to the exponent β_j [dimensionless]
E_{aj}	activation energy [J/kmol]
$\exp(-E_{aj}/RT)$	[dimensionless]

$(\rho Y_i)/M_i$	$[\mathrm{kmol/m^3}]$
R_u	universal gas constant [J/kmol K]
p_{ij}	rate exponents of reacting species [dimensionless]
P_j	sum of all p_{ij} for reaction j [dimensionless]
K_f	user-defined reaction coefficient $[([\mathrm{kmol/m^3}]^{(1-N_j)})/\mathrm{s}]$
$\mathbf{R}$	reaction [dimensionless]
ρ	density $[\mathrm{kg/m^3}]$
Y_i	mass fraction of species i [dimensionless]
M_i	molecular weight of species i [kg/kmol]
ν_{ij}	stoichiometric coefficients of reacting species
N_j	sum of all ν_i for reaction j [dimensionless]
τ_R	time scale [s]
τ_{mix}	turbulent mixing time scale [s]
τ_{kin}	time scale that is derived from chemical reaction rate from finite-rate kinetics [s]
Ω_i	source term of species i $[\mathrm{kg/(s\ m^3)}]$
T	temperature [K]

Using the Premixed Eddy Break-Up Model

This section describes the following about the Premixed Eddy Break-Up Model (PEBU):

- What Is the PEBU?
- Selecting the PEBU
- Preparing the PEBU for a combustion simulation
- The mathematical formulation of the PEBU

Also included is a reference section on solvers that are related to the PEBU.

The following related topics are described elsewhere:

- Using the Fuel Mass Fraction Ignitor
- Fuel Mass Fraction

What Is the Premixed Eddy Break-Up Model (PEBU)?

The current PEBU combustion model is designed for premixed flames where the fuel and oxidizer are perfectly mixed before entering the computational domain. The corresponding node (Eddy Break-up in the following screenshot), which has its own properties and is subject to common menu items, is activated as soon as the PEBU is selected.



In this model, a fuel mass fraction on the grid is tracked through transport equations. The mean species concentrations are obtained as functions of the mean fuel mass fraction and a one-step global reaction scheme, which is

internally calculated based on the unburnt gas composition. A mean enthalpy transport equation is also solved. The mean temperature, density, and viscosity are then calculated knowing the mean enthalpy and species concentrations.

PEBU Properties

The PEBU model node uses the Premixed EBU Rate Coefficient, A subnode.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that are used in STAR-CCM+, we recommend that you do not change any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance.

Convection	Selects the convection scheme to be used. For guidance as to convection scheme selection, see Diffusion Term, p. 6349.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Using the Premixed EBU Rate Coefficient

With the Premixed EBU Rate Coefficient, A node, you can specify the coefficient a (see Eqn. (1042)) for the premixed EBU reaction rate just like a typical scalar profile.



This node, which is subject to common menu items, can use any of the following methods:

- Constant
- Field function
- Tables
- User code

Selecting the Premixed Eddy Break-Up Model

To solve combustion problems using the Premixed Eddy Break-Up model (PEBU):

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas but Multi-Component Liquid is also acceptable.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select Premixed Combustion from the *Combustion* group.
- Select the option Eddy Break-up from the *Premixed Combustion* group.
- Select the desired solver type: Coupled Flow, or Segregated Flow, from the *Flow* group.

Since most of the available combustion models are only valid for turbulent flows, the Reynolds-Averaged Navier-Stokes option is selected automatically as the viscous regime. The *Physics Model Selection* dialog then looks as shown below. In this example, the option Segregated Flow was selected, which also activates the Segregated Fluid Enthalpy model.



The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select an appropriate model from the *Equation of State* group, depending on whether the fluid in your problem is a gas or a liquid.
- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- If radiative heat transfer is important, select one of the options in the *Radiation* group.

- If gravity forces significantly influence the solution, for example in fire simulations, select Gravity from the *Optional Physics models* group.

Preparing the PEBU for a Combustion Simulation

After model selection, the next steps for preparing the PEBU are:

- Specifying mixtures
- Setting properties of the PEBU
- Working with the unburnt gas fluid stream
- Setting common reference values:
  - Minimum Allowable Temperature
  - Maximum Allowable Temperature
  - Reference Pressure
- Setting boundary and initial condition values
- Setting the Fuel Mass Fraction Ignitor

Working with the Unburnt Gas Fluid Stream

A fluid stream named Unburnt Gas enables you to define the temperature and composition of the gas. The Unburnt Gas node appears whenever the PEBU is selected.



The unburnt temperature may be determined from an isentropic relationship that is based on pre-ignition values of T and p, as follows:

$$T_u = T_0 \left(\frac{p}{p_0} \right)^{\frac{\gamma - 1}{\gamma}} \tag{1039}$$

The reference unburnt temperature set here corresponds to T_0 in this equation.

PEBU Formulation

The PEBU assumes that the turbulent combustion rate is mixing controlled and is inversely proportional to the turbulent time scale. Implementation of the PEBU includes the following steps:

- Solving a transport equation for the fuel mass fraction
- Species calculation

Solving a Transport Equation for the Fuel Mass Fraction

The Y_f transport equation for the PEBU can be written:

$$\frac{\partial \rho Y_f}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} Y_f) - \nabla \cdot (\Gamma_{Y_f} \nabla Y_f) = \omega_{pebu} \tag{1040}$$

Here, b is defined as:

$$b = \frac{Y_f - Y_{res}}{Y_{ft} - Y_{res}} \tag{1041}$$

Y_f is the fuel mass fraction, Y_{ft} is the mass fraction in the unburnt gas. For a pure premixed case, Y_{ft} can be obtained by a transport equation for the mixture fraction. Y_{res} is the amount of fuel that is left at the end of combustion. It equals 0 in lean and stoichiometric mixtures and is positive in rich mixtures.

The source term is calculated using:

$$\omega_{pebu} = -\rho \bullet a \bullet b(1-b)\frac{(Y_{ft} - Y_{res})}{\tau_t} \tag{1042}$$

where a is the model coefficient, b is the regress variable and τ_t is the turbulent time scale.

Species Calculation

The species mass fractions are obtained based on internal algebra relations. The fuel is calculated with Eqn. (1041) based on the regress variable b. The other species are then calculated based on the global reaction.

Solver Reference

This section details the properties for the various solvers that are associated with the PEBU:

- PEBU Combustion Solver

PEBU Combustion Solver

Controls the solution update for the PEBU.

Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.8.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

Using Eddy Contact Micromixing

Micromixing is defined as molecular-scale mixing, dominated by viscous-convective and viscous-diffusive processes. The existing Eddy Break-Up (EBU) model does not account for the micromixing phenomenon that is present in liquid-liquid reaction systems. Specifically, liquid mixtures have diffusivity much smaller than kinematic viscosity, so the molecular Schmidt number S_c is much greater than 1, that is, $S_c \gg 1$. Therefore liquid reactions are more strongly influenced by the concentration gradient of the reactants. The density variation in liquid is insignificant as compared to gas-phase combustion. In chemical reactors, micromixing can change product quality due to selectivity.

For very slow reactions (the Damkohler number in Eqn. (1016) $D_a \ll 1$, that is, $t_{mix} \ll t_{rxn}$), the chemical kinetics determine the net reaction rate. So micromixing is not important and can be neglected. For fast ($D_a \sim 1$) and instantaneous ($D_a \gg 1$) reactions, micromixing is important for liquid reactions. In liquids, the rate of mixing, that is, the scalar dissipation rate, does not necessarily scale like ε/k. So the assumption that is used in the eddy break-up concept for the gas-phase, that scalar mixing time scales can be directly related to the dissipation of turbulent eddies (ε/k), is not valid for liquids due to the high Schmidt number.

The Eddy Contact Micromixing model [192] has the advantages of being numerically efficient, reasonably accurate, and easy to use for multiple reactions. It can be used for both steady and unsteady cases. This model is only available with the multicomponent liquid model, and requires

turbulence modeling. It works with all turbulent models that are available for the EBU hybrid model.



Eddy Contact Micromixing Model Properties

Source Enabled Trigger	Controls the combustion sources.	
	Never	Leaves combustion sources off. The Begin property is inactive if this option is selected.
	Iteration	Turns on combustion sources. This option is available in steady-state cases.
	Timestep	Turns on combustion sources. This option is available in unsteady-state cases.
Begin	Describes the number of iterations or time-steps when the sources from combustion become enabled.	
Store Reaction Rates	Activates the capability to post-process reaction rates, forward and backward (where applicable).	
	Activated	Stores reaction rates during the simulation run and makes them available among scalar functions.
	Deactivated	Does not store reaction rates (default).

See also:

Selecting the Eddy Contact Micromixing Model

Working with the Micromixing Timescale

Eddy Contact Micromixing Formulation

Selecting the Eddy Contact Micromixing Model

Activate the Eddy Contact Micromixing model as a partially premixed or non-premixed combustion model from the physics model selection dialog of the relevant physics continuum.

- In the physics model selection dialog, select Multi-Component Liquid from the *Material* group box. This model is for multi-component liquid only.

- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group box.
- Select Non-Premixed Combustion or Partially-Premixed Combustion from the *Combustion* group box.
- Select Eddy Contact Micromixing from the *Non-Premixed Combustion* or *Partially-Premixed Combustion* group box.
- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group box.

  Since this combustion model is designed for turbulent flows, the Turbulent and Reynolds-Averaged Navier-Stokes options are selected automatically, provided the Auto-select recommended Physics models option is activated.

  Other models that are selected automatically with Segregated Flow are Segregated Species and Segregated Fluid Enthalpy. With Coupled Flow, the models that are selected automatically are Coupled Species and Coupled Energy.

The physics model selection dialog contains your selections similar to the example shown below. In this example, the options Non-Premixed Combustion and Segregated Flow were selected.



Working with the Micromixing Timescale

When you use the Eddy Contact Micromixing model, the Micromixing Timescale node is available among the objects for defining chemical reactions.



To work with the micromixing timescale:

- Create a reaction within the Reacting > Reactions node.
- Open the [Reaction] > Properties manager node.
- Choose a method from those that are available for the Select method property.

  You can set any of the methods that you choose as a scalar profile.

Micromixing Timescale Properties

Select method	Allows you to select the method for calculating the micromixing timescale.	
	Classical Scalar Dissipation	Selects the classical scalar method (considers the high-Schmidt number effect) [193].
	Kolmogorov	Selects the Kolmogorov method (corresponds to Engulfment type of mixing).
	No Micromixing (Kinetics Only)	Selects the assumption that the reaction rate is dictated solely by finite-rate chemical kinetics.
	User Micromixing	Selects the User Defined method, allowing you to work with a user-defined field function.

Eddy Contact Micromixing Formulation

The reaction rate is modeled through an expression that takes the molecular micromixing process into account. This is done through dimensional arguments. So for a reaction of the form:

$$\nu_F F + \nu_O O \rightarrow \nu_{P_1} P_1 + \nu_{P_2} P_2 + \ldots + \nu_{P_j} P_j \tag{1043}$$

The rate of fuel depletion is assumed to be:

$$R_F = -\frac{\rho}{M_F}\left(\frac{1}{\tau_R}\right)\min\left[\bar{Y}_F, \frac{\bar{Y}_O}{s_O}\right] \quad \text{moles}/(\text{m}^3\text{s}) \tag{1044}$$

In the above equation:

$$s_O \equiv (\nu_O M_O)/(\nu_F M_F) \tag{1045}$$

See also:

Kolmogorov Method

Classical Scalar Mixing

Eddy Contact Micromixing Bibliography

Finite-Rate Kinetics Nomenclature

Kolmogorov Method

For this method:

$$\frac{1}{\tau_R} = A\sqrt{\frac{\varepsilon}{\nu}} \tag{1046}$$

Where ε is the turbulent dissipation rate, τ_R is the time scale, ν is the kinematic viscosity, and A is a constant (Mixing Coefficient) with a default value of 0.04.

Classical Scalar Mixing

$$\frac{1}{\tau_R} = A\frac{2 + Sc^{-1}}{2}\left[\frac{1}{2}\frac{k}{\varepsilon} + \frac{1}{2}\ln(Sc)\sqrt{\frac{\nu}{\varepsilon}}\right]^{-1} \tag{1047}$$

Where k is the turbulent kinetic energy, τ_R is the time scale, and A is a constant (Mixing Coefficient) with a default value of 1.0.

Eddy Contact Micromixing Bibliography

[192] Forney, L. J., and Nafia, N. 2000. "Eddy Contact Model: CFD Simulations of Liquid Reactions in Nearly Homogeneous Turbulence", *Chem. Eng. Sci.*, **55**(24), pp. 6049-6058.

[193] Hjertager, L.K., Hjertager, B.H., Solberg, T., 2002. "CFD modelling of fast chemical reactions in turbulent liquid flows", *Computers and Chem. Eng.* **26** pp. 507–515.

Using Surface Chemistry

The Surface Chemistry model is used to represent surface reaction mechanisms and surface species on the reacting walls or surfaces in porous regions.

This model becomes available for use in any one of the following situations:

- The Homogeneous Reactor with Surface Chemistry model is activated.
- The Eddy Break-Up (EBU) model is activated.
- The Non-reacting model and Segregated Species model are activated.

In the case of the Homogeneous Reactor with Surface Chemistry, the Dars-CFD Reaction model is used to solve the ODEs (Ordinary Differential Equations) that arise from the reaction mechanism. For the other two model selections, a Surface-Gas Interaction model is used for resolving the reaction system. Also, with the EBU model and with the Non-reacting option, surface chemistry is not applicable on porous regions but only on fluid regions with reacting surfaces.

For more information, refer to the following:

- Selecting the Surface Chemistry Model
- Managing Surface Species and Mechanisms -- how to manage surface reaction mechanisms and species
- Setting Wall Boundary Conditions
- Setting Porous Region Conditions
- General Surface Chemistry Formulation

Selecting the Surface Chemistry Model

Activate the Surface Chemistry model as an optional physical model from the *Physics Model Selection* dialog of the relevant physics continuum.

Before you begin, select any of the following prerequisite physical models:

- The Homogeneous Reactor with Surface Chemistry model
- The Eddy Break-Up (EBU) model
- The Non-reacting and Segregated Species models

The Surface Chemistry appears in the *Optional Models* group box.

To select the Surface Chemistry model:

- In the *Physics 1 Model Selection* dialog, select Surface Chemistry from the *Optional Models* group box.



The Surface Chemistry node appears under the Models node in the simulation tree.



Managing Surface Species and Mechanisms

The Surface Mechanism Manager is a placeholder for the nodes of the surface reaction mechanisms that are used in the simulation. Each mechanism has a corresponding set of bulk species, site species, and reactions.



The surface mechanism (CATALYST in the above screenshot) is a placeholder for relevant models that define it. The models are selected automatically when the surface mechanism is created and cannot be deselected. Models that are selected are: the Reacting Surface model and the Solid Composition model.

The properties of the surface mechanism are:

Id	The unique ID number of this phase (read-only).	
Non-Conservation On	The option for checking conservation.	
	Activated	Checks every reaction for the conservation of the number of sites, as well as stoichiometry.
	Deactivated	Checks every reaction for stoichiometry only.

When the Surface-Gas Interaction model is activated, it is possible to create a surface mechanism for a surface chemistry definition.

- To create a new surface mechanism, right-click the Surface Mechanism Manager node and select New Surface Mechanism.



To have alternate surface chemistry definitions in the same simulation, create additional surface mechanisms.

The Reacting Surface Model

The Reacting Surface model defines chemical reactions for the surface mechanism in the same manner that the Reacting model defines reactions for the continuum gas phase. It uses reaction definition objects.



Use this node to define chemical reactions occurring between mixture components that were defined earlier, either by importing Chemkin-formatted data or manually (see Defining Chemical Reactions).

The Solid Composition Model

This is a multi-component material model that allows you to specify two distinct types of surface species:

- Site Solid Mixture (surface species that actively participate in the reactions)
- Bulk Solid Mixture (surface species that can participate in reactions and of which there is an infinite supply)



In the case when the Dars-CFD model is used to resolve surface chemistry reaction mechanism, these managers cannot be manipulated explicitly -- no adding, removing, or renaming species is possible. Only when a complex chemistry definition is imported do these manager nodes become populated with available mechanisms and the surface species that are involved in surface-gas reactions.

However, when the Surface-Gas Interaction model is used, you can define these mechanisms, together with their species and reactions, either manually, or by importing the Chemkin-formatted surface chemistry definition.

Setting Wall Boundary Conditions

Only boundary regions that are designated as Wall can be reactive surfaces. To allow the wall boundary to become a reactive surface, select the Physics Conditions > Surface Mechanism Option sub-node of that boundary. Change the Mechanism property to one of the available mechanisms for your simulation (CATALYST in the following screenshot).



Mechanisms are available only after you import the complex chemistry definition or after you manually define the surface chemistry mechanism and surface species/reactions.

For reactive walls, the following values are entered in the *Properties* window as scalar array profiles:



- Surface Bulk Fraction -- the surface molar fraction for species of the given mechanism on the reactive wall
- Surface Site Fraction -- the fraction of the total number of sites that a given species occupies. This fraction only applies to species of the given mechanism on the interface between the reactive wall and gas phase.

Setting Porous Region Conditions

To activate porous chemistry on a porous region, select the Physics Conditions > Surface Mechanism Option sub-node of that region. Change the Mechanism property to one of the available mechanisms for your simulation (CATALYST in the following screenshot).



Mechanisms are available only after you import the complex chemistry definition (on porous regions, surface chemistry is available only through Dars-CFD).

Once you activate porous chemistry, you can control the area-to-volume ratio through a scalar profile among the objects in the Physics Values node of the region.



For reactive walls, the following values are entered in the *Properties* window as scalar array profiles:



- Surface Bulk Fraction -- the surface molar fraction for species of the given mechanism on the reactive wall

- Surface Site Fraction -- the fraction of the total number of sites that a given species occupies. This fraction only applies to species of the given mechanism on the interface between the reactive wall and gas phase.

The Surface-Gas Interaction Model

This model is activated when Surface Chemistry Model is chosen together with either the EBU model or the Non-reacting and Segregated Species models.



This model solves the ODEs that arise from the surface reaction system using the CVODE solver from the Sundials kit.

Surface-Gas Interaction Model Properties

Enabled Trigger	Controls the activation of the model.	
	Never	The model is not activated. The Begin property is inactive if this option is selected.
	Iteration	Activates the model in a steady-state case.
	Timestep	Activates the model in an unsteady-state case.
Begin	Sets the number of iterations or time-steps when the model is activated.	
verbose on error	Provides details about solver errors.	
	Activated	During iterations, this option generates maximum details in the *Output* window about solver errors.
	Deactivated	Uses standard iteration output.

Surface-Gas Interaction Model Expert Properties

Absolute Tolerance	Absolute tolerance for the CVODE ODE solver.
Relative Tolerance	Relative tolerance for the CVODE ODE solver.

Surface Chemistry Definition for the Surface-Gas Interaction Model

You can define the surface site and bulk species together with reactions manually; but if the surface mechanism is involved, it is better to import it from Chemkin-formatted files.

If you choose to import surface site and bulk species manually, the database materials that are selected must satisfy the following requirements:

- For surface site materials, define the elemental composition and thermodynamic polynomial for specific heat.
- For surface bulk materials, define the elemental composition.

Note: Manual definition of chemistry can lead to reactions with unrealistic reaction rates, causing the ODE solver to fail. To identify why the solver failed, activate the checkbox of the verbose on error property. Details about the solver failure appear in the *Output* window during subsequent iterations.

Importing Surface Chemistry Definition with Chemkin Files for the Surface-Gas Interaction Model

To import the surface chemistry definition with Chemkin files for the Surface-Gas Interaction model:

1. Check if the standard or extended material database contains all the species involved in a mechanism:
   - If it contains all the relevant species, skip to Step (3).
   - If not, create your own material database in Step (2).
2. Create a new material database.

   The Chemkin mechanism must contain at a minimum the description of the NASA polynomials for all the species involved in the mechanism. Put this information into the mechanism file itself or in separate files.

   There can be more than one file describing the NASA polynomials (usually one for gas and one for surface species).

   - Import these files into the new material database.
3. Define the gas species:
   - Follow the procedure for importing mixture components from a file.
   - In the *Select file and material database* dialog, select the material database that contains the data for the mechanism.
4. Define the site species, bulk species, and (if applicable) reactions:

- To define all of these phenomena at once, right-click the node of the surface mechanism and select Import Species and Reactions. Use this option if you need only one material database for the surface mechanism species and the reactions.



  - To select multiple files for the Select File text box in the *Select file and material database* dialog, use the `<Ctrl><Click>` method.
  - Select the material database that contains the data for the mechanism.
- Alternatively, if you need to specify another material database for the reactions:
  - Select Import Species from the same context menu.
  - Proceed to Step 5 to define the reactions.

5. Define reactions:
   - Right-click the node of the surface mechanism and select Import Species and Reactions.
   - Choose again the file that contains the surface species information and the material database that is relevant for this case.

General Surface Chemistry Formulation

With the Surface-Gas Interaction model, three types of species have to be defined (as outlined in the diagram of surface chemistry setup below):

- Species in the gas phase G_i
- Species on the surface at the gas-surface interface S_i. Some sites may be "open" (denoted as a generic species $OS^{(s)}$), meaning that they do not contain any species and that they can be used to model adsorption from the gas phase.

- Species in the solid layer -- bulk species B_i (bulk is often optional).



Limitations to surface chemistry modeling:

- Surface chemistry can only be applied to wall and baffle boundary conditions.
- No energy equations are solved for the reactive surface. Fixed temperature boundary conditions apply.
- Surface growth and shrinking are not taken into account.

You specify the surface reaction mechanism, which you can import from Chemkin files. The surface reaction rate is defined per unit surface area; this is different for the fluid-phase reactions, which are based on unit volume.

Basic Setup

The general surface reaction form for reaction r is:

$$\sum_{i=1}^{N_g} g'_{i,r} G_i + \sum_{i=1}^{N_s} s'_{i,r} S_i + \sum_{i=1}^{N_b} b'_{i,r} B_i = \sum_{i=1}^{N_g} g_{i,r''} G_i + \sum_{i=1}^{N_s} s_{i,r''} S_i + \sum_{i=1}^{N_b} b_{i,r''} B_i \qquad (1048)$$

where:

- G_i, S_i, and B_i are gas, site and bulk species.
- N_g, N_s, and N_b are the numbers of gas, site, and bulk species.
- $g'_{i,r}$, $s'_{i,r}$, and $b'_{i,r}$ are stoichiometry coefficients for reactants.
- $g_{i,r''}$, $s_{i,r''}$, and $b_{i,r''}$ are stoichiometry coefficients for products.

The rate of the rth reaction is then:

$$R_r = k_{f,r} \prod_{i=1}^{N_g} [G_i]_{\text{wall}}^{p_{g(i,r)}} \prod_{i=1}^{N_s} [S_i]_{\text{wall}}^{p_{s(i,r)}} \tag{1049}$$

where $p_{g(i,r)}$ and $p_{s(i,r)}$ are rate exponents for reactants, and $[*]_{\text{wall}}$ represents molar concentration on the wall. The forward rate constant for reaction r is given by the Arrhenius expression. For the simplest type of the reaction, it would be:

$$k_{f,r} = A_r T^{\beta r} e^{-E_r/RT} \tag{1050}$$

It is assumed that the reaction rate does not depend on the concentrations of the bulk species.

The net molar rate of change (production or consumption) of each species i is given by:

$$\dot{R}_{i,\text{gas}} = \sum_{r=1}^{N_{\text{reac}}} (g_{i,r''} - g'_{i,r}) R_r \tag{1051}$$

$$\dot{R}_{i,\text{site}} = \sum_{r=1}^{N_{\text{reac}}} (s_{i,r''} - s'_{i,r}) R_r \tag{1052}$$

$$\dot{R}_{i,\text{bulk}} = \sum_{r=1}^{N_{\text{reac}}} (b_{i,r''} - b'_{i,r}) R_r \tag{1053}$$

Equations

In this generalized surface chemistry implementation, it is assumed that the gas species mass fractions at the reacting surfaces are the same as the gas species mass fractions at the center of the cell volume. As the goal of the surface reaction is to compute concentrations of gas and site species at the wall, there is the following system of equations to solve:

$$\frac{\partial \rho Y_{i,\text{wall}}}{\partial t} = \dot{R}_{i,\text{gas}}, \; i = 1, 2, 3, \ldots, N_g \tag{1054}$$

$$\frac{\partial \rho S_{i,\,\mathrm{wall}}}{\partial t} = R_{i,\,\mathrm{site}},\ i = 1, 2, 3, \ldots, N_s \tag{1055}$$

$Y_{i,\,\mathrm{wall}}$ is the mass fraction of the gas species i at the wall and $S_{i,\,\mathrm{wall}}$ is the site species concentration, given as:

$$S_{i,\,\mathrm{wall}} = \rho_{\mathrm{site}} z_i \tag{1056}$$

where ρ_{site} is site density and z_i is site coverage of site species i.

Together, Eqns. (1054) and (1055) form a system of ODEs that are solved with the CVODE package.

Integrating the Gas Species Solution back to STAR-CCM+

The solution of this surface reaction model is the composition of the gas and surface species on the wall, $Y^*_{i,\,\mathrm{wall}}$ and $S^*_{i,\,\mathrm{wall}}$. The surface site species fraction is updated directly. For the gas species, the source term of the species transport equation is updated with the addition of averaged source term, Ω:

$$\Omega_i = \frac{(Y^*_{i,\,\mathrm{wall}} - Y_{i,\,\mathrm{wall}})}{dt} i = 1, 2, 3, \ldots, N_g \tag{1057}$$

This term contributes to the region source term S_{Y_i} from Eqn. (425).

Using Homogeneous Reactor Models

This section describes:

- The two homogeneous reactors available in STAR-CCM+:
  - The Homogenous Reactor model
  - The Homogeneous Reactor model with surface chemistry
- What is the Operator Splitting Complex Chemistry Model?
- What is the Chemistry ADI Complex Chemistry Model?
- What is the Dars-CFD Reaction model?
- What is the CFL Scaling Factor?
- Transitioning from the User Reaction model to the Passive Scalar model

What Is a Homogeneous Reactor Model?

The Homogenous Reactor model node, which is subject to common menu items, serves as a placeholder for chemistry models. This model is appropriate when you want to include a detailed chemistry description in the chemistry mechanism of your STAR-CCM+ simulation.



This version of STAR-CCM+ provides the Dars-CFD Reaction model as the choice for complex chemistry. To run this model, appropriate input files that describe the chemistry mechanism are needed (either ASCII text files or a case-specific FORTRAN library, the latter also requiring a FORTRAN compiler). You can create these files in a Dars software application or using the Chemkin import mechanism in STAR-CCM+.

What Is the Homogeneous Reactor Model with Surface Chemistry?

The Homogeneous Reactor with Surface Chemistry includes modeling of the reactive surfaces. This model requires a Surface Chemistry model for handling of the surface species and surface chemistry mechanisms.



Similar to the homogeneous reactor model, this version of STAR-CCM+ provides the Dars-CFD Reaction model as the choice for complex chemistry. From version 5.06 onwards, both ASCII and binary (case-specific FORTRAN library) inputs are allowed.

See also Surface Chemistry Formulation for Dars-CFD.

What Is the Operator Splitting Complex Chemistry Model?

This complex chemistry model employs the Operator Splitting Algorithm.



Properties

Source Enabled Trigger	Controls the combustion sources.	
	Never	Leaves combustion sources deactivated. The Begin property is inactive if this option is selected.
	Iteration	Activates combustion sources. This option is available in steady-state cases.
	Timestep	Activates combustion sources. This option is available in unsteady-state cases.
Begin	Describes the number of iterations or time-steps when the sources from combustion get activated.	

What Is the Chemistry ADI Complex Chemistry Model?

This complex chemistry model employs the Chemistry ADI Algorithm. This model is only compatible with the Segregated Flow model.



Properties

Source Enabled Trigger	Controls the combustion sources.	
	Never	Leaves combustion sources deactivated. The Begin property is inactive if this option is selected.
	Iteration	Activates combustion sources. This option is available in steady-state cases.
	Timestep	Activates combustion sources. This option is available in unsteady-state cases.
Begin	Describes the number of iterations or time-steps when the sources from combustion become enabled.	

What Is the Dars-CFD Reaction Model?

The Dars-CFD Reaction model is a complex chemistry model which is suitable for simulations where a detailed prediction of non-fuel and oxidizer species is required. It uses the Dars-CFD library to solve ODEs (Ordinary Differential Equations) for each cell in the regions where it is activated. These ODEs arise from the complex chemistry definition of the particular reaction system.

The Dars-CFD library includes the stiff ODE solver that is used to integrate the chemical composition equations, thus advancing the chemistry for each time-step or iteration.

As these stiff systems are often difficult to converge, the Dars-CFD Reaction model is computationally expensive.

The Dars-CFD Reaction node, which has properties and a pop-up menu, contains the objects for working with this type of chemical reaction.



This section explains the following:

- Selecting the Dars-CFD Reaction model
- What is the Dars-CFD Ideal Gas model?
- Setting the Dars-CFD properties
- Setting the approximation option
- What is In-Situ Adaptive Tabulation (ISAT)?
- Setting up a complex chemistry definition
- Packaging Fortran 90 files into a case-specific library
- Importing the complex chemistry definition in Dars-CFD format
- Importing precompiled chemical libraries
- Moving a simulation file with a complex chemistry model to another file system
- Formulation of the complex chemistry model
- Formulation of the surface chemistry model

See also:

DarsCFD Reports

Pop-Up Menu of the Reports Manager

Setting the Dars-CFD Properties

The Dars-CFD Properties node, which has properties, allows specification of a few parameters that the Dars-CFD Reaction Model requires.



The Dars-CFD Properties node also allows you to activate user-coding to modify the species source terms or reaction rates using the Dars-CFD FORTRAN files.



Once you have selected your particular user coding option, modify the correct Dars-CFD FORTRAN file to implement your desired changes in the species source terms or reaction rates.

The user-coding template files are in the following directories:

- `star\darsFiles` folder for Windows
- `star/bin/darsCfd` folder for Linux

To implement your changes:

- Copy the correct Dars-CFD FORTRAN template file from the STAR-CCM+ installation folder to the local folder where your mechanism is located.
- To implement your modifications, edit the file. The files are commented and self-explanatory.
- In STAR-CCM+, do the following:
  - Remove the Dars-CFD library if it has already been imported (without deleting species).
  - Import the new Dars-CFD library.

Dars-CFD Properties

User Coding

Controls the user modification of the species source terms or reaction rates using the Dars-CFD FORTRAN files.

The following options are available if the Binary Problem Description is activated:

None	Does not use user coding.
Calculate Source Term	Calculates the species source terms with the Dars-CFD FORTRAN file `SetDARSGasSrcTerms.f`
Modify Reaction Rates	Modifies the species reaction rates with the Dars-CFD FORTRAN file `ModDARSGasReacRates.f`
Modify Source Terms	Modifies the species source terms with the Dars-CFD FORTRAN file `ModDARSGasSrcTerms.f`

The following option is available if the Binary Problem Description is deactivated:

None	Does not use user coding.

Absolute Tolerance

Absolute tolerance for the Dars-CFD ODE solver.

Relative Tolerance

Relative tolerance for the Dars-CFD ODE solver.

Use Reference Values for Temperature Range.

Activated by default. If deactivated, a separate Dars-CFD Temperature Range node appears which allows you to set the high and low temperature limit.

Activated	The low and high temperature limits for the Dars-CFD ODE solver are taken from Minimum and Maximum Allowable Temperature of the Continuum.
Deactivated	A separate, Dars-CFD Temperature Range node appears that allows you to manually set the low and high temperature limit to the Dars-CFD ODE solver

Dars-CFD Temperature Range properties.

Minimum Temperature	Low temperature limit for the Dars-CFD ODE solver. The default value, 250K, is the Minimum Allowable Temperature.
Maximum Temperature	High temperature limit for the Dars-CFD ODE solver. The default value 4000.0K is the Maximum Allowable Temperature.

Setting the Approximation Option

The Approximation Option node controls the option of using the ISAT approximation to speed up chemistry calculations where possible. To access the properties of the approximation option, select this node.



This object also facilitates load balancing (for parallel cases) through one of its sub-nodes: Equilibrium Time Scale, ISAT, or No Approximation.

Approximation Option Properties

Method	Controls the option of using ISAT approximation.	
	Equilibrium Time Scale	Turns on the Equilibrium Time Scale approximation and adds the child node to the object tree. This option is available for the Operator Splitting complex chemistry model without surface chemistry, as well as for porous media. The Equilibrium Time Scale method is not available when both the Implicit Unsteady and Eddy Dissipation Concept models are activated.
	ISAT	Turns on ISAT approximation and adds the ISAT child node to the object tree. This option is available for the Operator Splitting complex chemistry model without surface chemistry, as well as for porous media. The ISAT method is not available when both the Implicit Unsteady and Eddy Dissipation Concept models are activated.
	No Approximation	Turns off ISAT approximation and adds a No Approximation child node to the object tree.

Balancing the Computational Load for Complex Chemistry

The dynamic load balancing feature is on by default.

In most combustion problems that use Dars-CFD, combustion is confined to one part of the computational domain. Since the mesh partitioning does not take this fact into account, this part of the computational domain can get distributed to few processors. Hence these processors have a large computational overhead, while others become idle, since they would finish their complex chemistry computation in a much shorter time period. This distribution problem causes the loss of scalability for simulations that involve complex chemistry.

When dynamic load balancing is turned on, an algorithm is employed which attempts to redistribute computation from processors that own the combustive parts of the domain to the ones that do not. To control this feature, use the Dynamic Load Balancing expert property of either the ISAT or No Approximation node.

If it is known that most of the computational domain is combustive, turn off the load balancing. In this case, redistributing the computation does not help the load balancing, but it incurs some overhead.

What Is the No Approximation Option?

The No Approximation node appears as the sub-node of the Approximation Option node when ISAT approximation is turned off. It serves as a placeholder for the dynamic load balancing option in its expert properties.



No Approximation Expert Properties

Dynamic Load Balancing	For parallel cases, redistributes computational load of complex chemistry by randomization.	
	Activated	Uses the load balancing option.
	Deactivated	Leaves the load distribution in its default state.
Ode Solver	Currently, Standard is the only option for the Ode Solver.	

What Is In-Situ Adaptive Tabulation (ISAT)?

To approximate functions that are expensive to calculate, Dr. Stephen B. Pope developed the In-Situ Adaptive Tabulation (ISAT) method. Within the Dars-CFD reaction model framework, the ISAT method approximates the solution of the ODE (Ordinary Differential Equation) for a given initial condition through:

- The thermodynamic variable (sensible enthalpy)
- The composition variable (specific molar number)
- The possible absolute pressure and time-step (depending on the type of simulation)

Note: The ISAT method is not available when both the Implicit Unsteady and Eddy Dissipation Concept models are activated.

The ISAT node, which has properties and a pop-up menu, allows the specification of the parameters that the ISAT approximation method requires.



Unless restart is activated, ISAT starts from a blank table and gradually adds points to its table based on the criteria that come from the absolute and relative error of ISAT approximation. When Dars-CFD is used to calculate a chemistry step in a particular cell, ISAT first searches the table for a point with which to approximate the value of the new point.

If ISAT finds such a point, then the new value is not calculated directly with the ODE solver, but rather approximated based on the function value of the point in the table. Otherwise, if ISAT cannot find such a suitable point in the table, the ODE solver inside Dars-CFD is called, and a function value for this point is calculated. Then this point is stored in the table. The size of the table thus grows as the simulation progresses, although the Table Size property of the ISAT node limits its overall size. When the table reaches its limit, none of the points are added to it.

For more details on the ISAT method, see the ISAT formulation.

ISAT Properties

Restart

Controls whether to load the ISAT table from a file.

- **Activated** — The simulation resumes from the ISAT table that is given in the work folder. The ISAT file name is `isat_1.tab`. Otherwise, it is started from a blank table.
- **Deactivated** — The simulation does not resume from the ISAT table.

Absolute Tolerance

Absolute tolerance for the ISAT approximation.

Table Size (MB)

Maximum size for the ISAT table.

Pressure Specification

Controls whether pressure remains constant or not.

- **Constant** — Pressure remains constant.
- **Variable** — Pressure is variable.

Time Step Specification

In the current version of STAR-CCM+, this property is preset to one of the following:

- In unsteady simulations, this property is `Constant`.
- In most steady simulations, variation of the time-step is not suitable for ISAT tabulation. Therefore STAR-CCM+ uses a special `Iterative tables` technique to accommodate this variation. An ISAT Tables node is within the ISAT node.

Expert Properties

Dynamic Load Balancing

For parallel cases, redistributes computational load of complex chemistry by randomization.

- **Activated** — Uses the load balancing option.
- **Deactivated** — Leaves the load distribution in its default state.

Save on Master Only	When activated, this option saves ISAT tables only on the master in a parallel server.
Ode Solver	Selects a type of Dars-CFD ODE solver. Only the standard type is available for the current version of STAR-CCM+.

ISAT Pop-Up Menu

Save ISAT table	Saves the table in the work folder under the `isat_1.tab` file name.
Delete ISAT table	Deletes the internal ISAT table and makes ISAT start from a blank table.

What Are the ISAT Tables?

The ISAT Tables node appears within the ISAT node when the simulation uses the Iterative table technique for the time-step specification. This node acts as a placeholder for table nodes. To control the upper and lower time limits of the ISAT tables, use the properties of the table nodes.



ISAT Table Node Properties

Time Lower Limit	Sets the lower time limit for the ISAT table. This is the time-step at which the chemistry is resolved for this table.
Time Upper Limit	Sets the upper time limit for the ISAT table.
Index	Index reference for the ISAT table object (read-only).

The time-step in the range `[Lower limit, Upper limit]` is achieved by calling this ISAT table iteratively, in each iteration starting from the previous solution, until the desired time-step is reached.

If the pseudo-time step for chemistry in a given cell is within the range of a particular table, then this table will be used for the approximation.

You can add more ISAT tables, but be careful that the time-step domains of the multiple tables are consecutive and do not overlap.

There are two special cases when the time-step is outside of the range of all the tables:

- The time-step is less than the lowest value of Time Lower Limit among the ISAT table nodes. In this case, the table with the smallest value for this property is called once.
- The time-step is greater than the highest value of Time Upper Limit. In this case, ISAT is not called and the chemistry step is calculated directly through Dars-CFD.

What Is the Equilibrium Time Scale?

Use the Equilibrium Time Scale option for detailed chemistry computations within the framework of the Dars-CFD model. This method is designed to provide a quick approximate solution for detailed chemistry computations. The converged solution obtained using Equilibrium Time Scale can then be further refined using the more expensive ISAT or No Approximation options.

The chemistry time scale is computed based on the specified fuel species. See Equilibrium Time Scale Formulation.

The Equilibrium Time Scale node, which has properties and a pop-up menu, allows the specification of the parameters that the Equilibrium Time Scale approximation method requires.

The Fuel Components node appears under the Equilibrium Time Scale node. Right-click the Fuel Components node and select a primary fuel species. (For premixed flames, include some radical species also to get a better estimate of the chemical time scale.)



Equilibrium Time Scale Properties

ISAT Absolute Tolerance	Absolute tolerance for the ISAT approximation.
ISAT Table Size (MB)	Maximum size for the ISAT table.
Flow Time Constant	Time constant for the flow c_{flow}. See Eqn. (1077) and the following explanation.
Chemistry Time Constant	Time constant for the chemical reaction c_{chem}. See Eqn. (1077) and the following explanation.
Include Chemistry Time?	Includes the chemistry time scale in the computation. To reduce computational expense, do not activate when the flow time scale is expected to dominate (for instance, in non-premixed mixing dominated flows). Activated: Compute the chemistry time scale. Deactivated: Do not compute the chemistry time scale. This is the default.

Expert Properties

Restart ISAT Table	Controls whether to load the ISAT table from a file. Activated: The simulation resumes from the ISAT table that is given in the work folder. The ISAT file name is `isat_10.tab`. Otherwise, it is started from a blank table. Deactivated: The simulation does not resume from the ISAT table.
Save on Master Only	When activated, this option saves ISAT tables only on the master in a parallel server.

Equilibrium Time Scale Pop-Up Menu

Save Equilibrium ISAT table	Saves the table in the work folder under the `isat_10.tab` file name.
Delete Equilibrium ISAT table	Deletes the internal ISAT table and makes ISAT start from a blank table.

Fuel Components Pop-Up Menu

Add Fuel Component	Displays a menu of available fuel components. Selecting a component adds a node under the Fuel Components node.
Refresh	Refreshes the Add Fuel Component menu.

Fuel Node Properties

Mass Fraction	Mass fraction of the fuel selected for the run, between 0.0 and 1.0. This information is used elsewhere, but not in the calculation of the equilibrium time scale.

Dars-CFD Reaction Model Properties

Binary Problem Description

Controls whether to use a binary library file.

Activated	Uses binary format.
Deactivated	Does not use binary format.

Case Path

The name and path of the folder containing the complex chemistry description (either ASCII text files or a FORTRAN library). The path is set when you navigate to the folder while importing complex chemistry or importing case from STAR-CCM+ repository.

Is Dars-CFD library loaded

Indicates with a read-only checkbox whether the Dars-CFD library is loaded into the simulation. The library maps to the type of complex chemistry model you select. It is loaded either when you import it or when the solution is initialized.

Activated	The library is loaded.
Deactivated	The library is not loaded.

Expert Properties

Cold Flow Temperature Limit

If temperature in the cell is lower than this number, then cold flow is assumed - no chemistry is calculated.

Allow Switch to ASCII

Allows you to use the ASCII format. This option helps when a binary case description is not available.

Activated	Activates this option.
Deactivated	Leaves this option deactivated.

Run Time Compile

Controls whether to use a temporary set of input files for the chemistry import when restoring the simulation file.
If this option is deactivated and the Binary Problem Description option is activated, then only the first item in the pop-up menu is available to import files (Import Complex Chemistry Definition (Dars-CFD format)...). In this situation, select the directory that contains the actual pre-compiled library.

Activated

Instructs STAR-CCM+ to import the input files as follows:

- Create a temporary folder.
- Place copies of the locally stored (in the `.sim` file) input files to that folder.
- Carry out the steps for the importing mechanism in Dars-CFD format.
- Delete the temporary folder and copies of input files once the data is imported into the simulation file.
The temporary folder is never visible to you, since it exists only during the import process.

Deactivated

Leaves this option deactivated.

Dars-CFD Reaction Model Pop-Up Menu

Import Complex Chemistry Definition (Dars-CFD format)...

Activates a standard *Open* dialog that imports a library file for the complex chemistry definition in Dars-CFD format.

Import Complex Chemistry Definition (Chemkin format)...

Activates a standard *Open* dialog that imports a library file for the complex chemistry definition in Chemkin format.

Import Case From STAR-CCM+ Repository

Selects from among two types of Jacobian libraries.

Libraries with Numerical Jacobian

Selects a numerical Jacobian library. Currently the repository contains the following numerical library files:

- Methane combustion
- CO/H2 Flame
- Hydrogen combustion

Libraries with Analytical Jacobian

Selects an analytical Jacobian library. Currently the repository contains the following analytical library files:

- Hydrogen 13-species
- Methane 17-species skeletal
- Methane 30-species skeletal
- Methane 53-species GRI3
- Methane 32-species skeletal
- n-Heptane 88-species skeletal
- n-Heptane 188-species skeletal
- n-Dodecane 106-species skeletal
- PRF 171-species skeletal

Reload Dars-CFD Case

Reimports the Dars-CFD format library file. Use this option if you modify the file after importing it.

Save Complex Chemistry Definition Input...	Activates a standard *Save* dialog to save the complex chemistry definition in the Dars-CFD format.
Delete Complex Chemistry Case...	Removes the library that is used for the complex chemistry definition.

Selecting the Dars-CFD Reaction Model

To solve combustion problems using the Dars-CFD Reaction model:

- In the *Physics Model Selection* dialog, select the Multi-Component Gas model from the *Material* group. (The Dars-CFD Reaction model is not available with liquid.)
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Specify whether your combustion is premixed, non-premixed or partially premixed by selecting the corresponding radio button in the *Combustion* group.
- Select the Homogeneous Reactor or Homogeneous Reactor with Surface Chemistry option from the particular combustion group. If you select the latter option with the Auto-select recommended physics models checkbox activated (the default setting), the Surface Chemistry model is also selected.
- Select the Operator Splitting Complex Chemistry model or the Chemistry ADI Complex Chemistry model from the *Complex Chemistry Model* group. The Dars-CFD Reaction model is automatically selected.

  A requirement for either Complex Chemistry model is that species and enthalpy models be activated, so select a flow model (Coupled or Segregated) from the *Flow* group as well. If you select Chemistry ADI with the Auto-select recommended physics models checkbox activated, the following models are also selected, since they are the recommended flow options for this type of complex chemistry model:

  - Segregated Flow
  - Segregated Species
  - Segregated Fluid Enthalpy
- Select an appropriate model from the *Equation of State* group. Currently, only the Dars-CFD Ideal Gas model is available. This model has an Ideal gas equation property with two options: `Dars-CFD` and `STAR-CCM+`.

The difference between these options is that the Dars-CFD equation of state uses functions from the Dars-CFD library to convert temperature to enthalpy, as well as enthalpy to temperature.

The *Physics Model Selection* dialog contains your selections similar to the example shown below. In this example, the options Coupled Flow and Dars-CFD Ideal Gas were selected.



The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- Make an appropriate selection from the *Viscous Regime* group.
- If radiative heat transfer is important, select Radiation from the *Optional Physics models* group.

- Select Gravity from the *Optional Physics models* group if gravity forces significantly influence the solution.

After you import the complex chemistry definition, make sure that Read-Only Thermodynamic Polynomial Data is the active method for calculating specific heat. It is the required method for the Dars-CFD Reaction model.

In any physics continuum where the Dars-CFD Reaction model is selected, a [Continuum] > Partially Stirred Reactor node appears. This node allows you to set a variety of properties. See What is the Partially-Stirred Reactor (PaSR)?

See also:

DarsCFD Reports

What Is the Dars-CFD Ideal Gas Model?

This special version of the Ideal Gas model is available when the Dars-CFD Reaction model is used in the simulation. The corresponding node (Dars-CFD Ideal Gas in the following screenshot), which has properties and is subject to common menu items, is activated automatically when the Dars-CFD Reaction model is selected.



The difference between Ideal Gas and Dars-CFD Ideal Gas is that the Dars-CFD Ideal Gas model uses functions from the Dars-CFD library to convert between temperature and enthalpy.

Dars-CFD Ideal Gas Model Properties

Ideal gas equation	Selects the ideal gas equation for the simulation.	
	Dars-CFD	Uses functions from the Dars-CFD library.
	STAR-CCM+	Uses the standard ideal gas equation in STAR-CCM+.

What is the Partially-Stirred Reactor (PaSR)?

The Partially-Stirred Reactor (PaSR) is a computationally inexpensive tool that you can use to pre-test Dars-CFD model simulations with detailed chemistry before you perform complex full simulations that are computationally expensive.

Within STAR-CCM+, the PaSR focuses on the implementation of detailed chemistry using two reactors simultaneously.

Within each reactor, there are three fluid streams of discrete particles which flow together, mix, and react. Notional particles which carry information about the species compositions, enthalpy, and pressure, represent the thermo-chemical composition of the fluid mixture. This information is combined from all particles and characterizes the overall thermo-chemical state of the fluid. The flow inside the reactor is assumed to be homogeneous, so there is no need to track the position of particles or the fluid flow inside the reactor. The fluid flow is not solved, it is the change in chemical composition of each particle due to mixing and chemical reaction that is solved (See equations (1058), (1059), (1060)).

You can specify the composition of species in each stream to represent:

- Fuel

  This stream is often composed of fuel that is mixed with oxygen or air.
- An oxidizer

  Oxygen, or a mixture of oxygen and nitrogen (air).
- A pilot stream

  The pilot stream consists of hot unburnt fuel and products of burning.

Fuel and oxidizer streams are typically at low temperatures of 300-500 K. Hot pilot streams are often over 1500 K and are used to ignite the flame in the PaSR and keep it stable.



The PaSR runs two reactors simultaneously using the following combustion models:

- Dars-CFD Reaction model with No Approximation (Direct Integration)
- Dars-CFD Reaction model with ISAT

The PaSR pre-analysis results allow you to monitor the temperature and chemical composition of various species in the two reactors. You can then compare these results to assess the error that is incurred using ISAT approximation, compared to using no approximation. By evaluating the error incurred using the PaSR, you can decide what models and ISAT error tolerance to use in the full simulation.

Typically, the PaSR is useful to test and monitor the following possibilities:

- Different chemical mechanisms
- Relative performance and accuracy of Dars-CFD model (with and without ISAT)
- ISAT approximation error in individual species composition

Setting up the PaSR

In any physics continuum where a model that can use ISAT is selected, a [Continuum] > Partially-Stirred Reactor node appears:



- To activate the PaSR, select a physics model that allows you to use ISAT. The Dars-CFD Reaction model is used in this example.

  The Physics 1 > Partially-Stirred Reactor node appears.

- Import the Complex Chemistry Definition.

- Select the Dars-CFD Reaction > Dars-CFD Properties > Approximation Option node and set Method to ISAT.



- Select the Approximation Option > ISAT node and set the ISAT properties as required.

Specifying the PaSR Stream Compositions

To specify the composition of a fluid stream for the PaSR:

- Select the Physics 1 > Fluid Stream Manager > Stream 1 node.
- Set the properties that you require for Stream 1.
- Right-click the Stream 1 > Components node and select a component for that stream. Repeat this process until all of the required components are selected for that stream.
- Select each of the Stream 1 > Components > [Component] nodes in turn and set the Mass Fraction property as required.
- Repeat the steps above to specify the compositions of Stream 2 and Stream 3.

Specifying the PaSR Error and Statistics Variables

- Select the Physics 1 > Partially-Stirred Reactor > PaSR Error node and set the properties that you require.
- Select the Physics 1 > Partially-Stirred Reactor > PaSR Statistic node and set the properties that you require.

If you also require species variables:

- Right-click the PaSR Error node and select the species that you require.

The species that you add are listed in the *PaSR Error - Properties* window and are available to select when you plot an XY-plot.

- Right-click the PaSR Statistics node and select the species that you require.

  The species that you add are listed in the *PaSR Statistics - Properties* window and are available to select when you plot an XY-plot.

Running the PaSR

To set the length of time for the PaSR run:

- Select the Physics 1 > Partially-Stirred Reactor node.
- Set the Number of Residence Times to the required value.



To run the PaSR:

- Right-click the Partially-Stirred Reactor node and select Run PaSR.



- You can stop and resume the PaSR if necessary.

To view the effects of the error and statistics variables:

- Right-click the Partially-Stirred Reactor > PaSR XY Plots node.
- Select the type of plot that you require. See Plotting a PaSR XY Plot.

In the output window, a variety ISAT statistics are printed. Towards the end of the output, a summary of overall ISAT tabulation error and the estimated speedup that is achieved using ISAT is also printed.

After reviewing the plots:

Use the information from the plots and output window to assess the performance and accuracy of ISAT, then:

- Adjust the ISAT reactor parameters.
- Rerun the PaSR.
- Repeat the previous steps until you achieve the level of accuracy that you require.

You can use the results to adjust the error tolerance appropriately with ISAT in the real simulation.

PaSR Properties

Partially-Stirred Reactor Properties

Residence Time	τ_{res} The average time that a PaSR particle spends in the reactor.
Mixing Time Scale	τ_{mix} Controls the extent of mixing between particles. A large mixing time scale implies very slow mixing. The smaller the mixing time scale, the faster the mixing. See Eqn.(1058) and (1059).
Pairing Time Scale	τ_{pair} Determines the number of particles that are randomly selected for pairing with inflowing particles. Used to model turbulence effects and molecular mixing in the reactor.
Minimum Pressure	The minimum pressure within the PaSR. If the minimum and maximum pressure values are equal, the pressure within the PaSR remains constant.
Maximum Pressure	The maximum pressure within the PaSR. If the minimum and maximum pressure values are equal, the pressure within the PaSR remains constant.
Pressure Variation Period	The time period in which the pressure inside the PaSR varies sinusoidally from the specified maximum pressure to the specified minimum pressure.
Number of Residence Times	The number of residence times for which you want to run the PaSR. The total duration of the PaSR run is: Residence Time x Number of residence times.
Number of Particles	N Number of particles that are used to represent the chemistry in the PaSR.
Relative Flow Rates of Streams	The relative mass flow rates (kg/s) of the three inflowing streams.

Partially-Stirred Reactor Expert Properties

The Expert Properties are for information and are set to read only.

Number of Reactors	Specifies the number of reactors for the PaSR. Set to 2 by default.
Reactor 1	Specifies the first reactor that is used in the PaSR. Implements the chemistry using Dars-CFD (No Approximation).
Reactor 2	Specifies the second reactor that is used in the PaSR. Implements the chemistry using Dars-CFD with ISAT.

Partially-Stirred Reactor Pop-up menu

Allows you to Start, Stop, and Resume the PaSR.

PaSR Error

The PaSR monitors the error between:

- A baseline reactor which uses the Dars-CFD Reaction combustion model with No Approximation, and
- A second reactor which uses the Dars-CFD Reaction combustion model with ISAT

You can use the PaSR Error node to select specific variables for monitoring error. If the errors are high, you can use a smaller ISAT error tolerance when running the full simulation.

PaSR Error Properties

Variables Selected	Lists the variables that are selected for computing error.
Baseline Reactor	Specifies the baseline reactor relative to which errors are computed. Therefore, the errors in the baseline reactor are zero.

Note: If you want to monitor error in individual species, then select those species before running the PaSR.

PaSR Statistics

Use the PaSR to monitor statistics such as temperature (mean, maximum, and minimum), enthalpy, and species composition.

You can use the PaSR Statistics node to view the variables that are selected for monitoring statistics. The statistics are collected independently for each reactor.

PaSR Statistics Properties

Variables Selected	Lists the variables that are selected for computing statistics.

Note: By default, statistics are only collected for temperature and enthalpy. If statistics are needed for other species, select these species before running the PaSR.

PaSR XY Plots

You can use the PaSR XY Plots node to create error and statistics plots which allow you to compare the selected variables in the two reactors:

- A baseline reactor which uses the Dars-CFD Reaction combustion model with No Approximation, and
- A second reactor which uses the Dars-CFD Reaction combustion model with ISAT

Pop-up Menu of PaSR XY Plots Node

If you select Add New PaSR Plot, the following options are available:

Timestep Error

Absolute Error	Plots the minimum, maximum, or mean difference (specified in the select points for plotting dialog) of the variable that is monitored between the two reactors at each time-step. Options appear that show the variables that are selected to be monitored.

Relative Error	Calculates the minimum, maximum, or mean (specified in the select points for plotting dialog) relative error of the variable that is monitored between the two reactors at each time-step. The baseline reactor quantity normalizes the relative error. The resulting values are plotted for each time-step. Options appear that show the variables that are selected to be monitored.
L2 Species Error	The vector 2-norm error used in all species compositions.

Time-step Stats

You can plot the minimum, maximum, or mean statistics of a variable in each reactor. Options appear that show the variables that are selected to be monitored. The statistics are collected independently for each reactor.

See Plotting a PaSR XY Plot

See also: Specifying the PaSR Error and Statistics Variables

Plotting a PaSR XY Plot

There are a variety of plots that you can produce to monitor the performance of the PaSR.

To plot temperature on a PaSR Time-step Statistics XY plot:

- Expand the Continua > Physics 1 > Partially-Stirred Reactor node.
- Right-click the PaSR XY Plots Node.

- Select Add New PaSR Plot > Timestep Stats > Temperature.



- In the *Graphics* window, right-click and select Add New Table View for Plotting.



- In the Select points for plotting dialog, select the reactor and type of variable that you want to plot.



- Click OK.

The following PaSR XY plot shows the mean temperature of reactor 1, which represents the Dars-CFD (No Approximation) chemistry:



Select Points For Plotting Dialog

The following options are available to select from the *Select points for plotting* dialog:

Select reactor	1. Dars-CFD (No Approximation) 2. Dars-CFD with ISAT	
Select type of variable to plot	MIN	Plots the minimum value over all of the particles for the variable that is selected.
	MAX	Plots the maximum value over all of the particles for the variable that is selected.
	MEAN	Plots the mean value over all of the particles for the variable that is selected.

PaSR Formulation

The PaSR is used to test implementations of combustion chemistry.

There are three inflowing streams of particles. Usually, one stream consists of gaseous fuel, another stream contains an oxidizer (which could be air). The third stream is known as a pilot stream. The pilot stream consists of a mixture of hot burnt products, unburnt fuel, and air. This stream is used to ignite the flame.

At any time, there are an even number of particles and at the specified time step, Δt, the ith particle has the composition $\phi^{(i)}(t)$. However, this composition of species, enthalpy, and pressure, can change discontinuously at discrete times $k\Delta t$. Where k is an integer depending on the inflow, outflow, and pairing of particles. Between these discrete times, there are three sub-steps that alter the composition. Each of these sub-steps includes a mixing fractional time-step and a reaction fractional time-step.

In the mixing fractional step, particles are arranged in pairs (for example, particles 1 and 2, 3 and 4, or N-1 and N). Each pair of particles, say p and q, evolves by:

$$\frac{d\phi^{(p)}}{dt} = -(\phi^{(p)} - \phi^{(q)})/\tau_{mix} \tag{1058}$$

$$\frac{d\phi^{(q)}}{dt} = -(\phi^{(q)} - \phi^{(p)})/\tau_{mix} \tag{1059}$$

where τ_{mix} is the specified mixing timescale.

In the reaction fraction step, each particle evolves by the reaction equation:

$$\frac{\phi^{(i)}}{\Delta t} = S(\phi^{(i)}) \ . \tag{1060}$$

During the specified residence time τ_{res} that a PaSR particle spends in the reactor, outflow and inflow consist of selecting $\frac{1}{2}N\Delta t/\tau_{res}$ pairs of particles at random. The compositions of the selected pairs are replaced with inflow compositions which are drawn from a specified distribution.

The reaction fractional time-step (Δt) is taken to be 1/3rd (for each sub step) x 1/10th of the minimum of residence and pairing times.

If residence time = 0.01 s and pairing time scale = 0.001 s, the reaction fractional time step (Δt) = 1/3 x 1/10 x 0.001 s = 3.33E-5 s. This value is displayed in the output window when the PaSR is run.

τ_{pair} is the specified pairing time scale, when $\frac{1}{2}N\Delta t/\tau_{pair}$ pairs of particles (other than the inflowing particles) are randomly selected for pairing. The selected pairs of particles are randomly shuffled with the inflowing particles and the pairs most likely change partners.

Bibliography

[194] Pope, S.B. 1997. "Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation", *Combust. Theory and Modelling*, **1**, pp. 41-63.

Setting Up a Complex Chemistry Definition in Dars-CFD Format

The Dars-CFD library requires a specific form of input files for the complex chemistry definition. To create these input files, you must have the DARS tool available.

There are two types of input files for Dars-CFD:

- ASCII input files -- Two input files in a specific format, `InputRedKinMec.txt` and `InputRedKinTherm.txt`, are needed.
- Binary input files -- The DARS tool creates five FORTRAN 90 files with specific subroutines that contain the hard-coded chemistry information. These FORTRAN files are then recompiled with the Dars-CFD library to provide a case-specific library.

The obvious drawback of the binary input is that you must have a FORTRAN compiler available. However, due to the way the Dars-CFD library is designed, cases with ASCII input are on average two to three times slower than their counterparts with the binary input. For gas phase chemistry, both modes of using this library are supported in STAR-CCM+.

The Binary Problem Description property of the Dars-CFD Reaction model determines the form of the complex chemistry description.



When its checkbox is activated, the binary-input description is used. Otherwise the ASCII-input approach is required.

Packaging the Fortran 90 Files into a Case-Specific Library

You must have a FORTRAN compiler available for this step.

Note: On Linux systems only, if the FORTRAN compiler is in the PATH environment variable, and the Run Time Compile expert property of the Dars-CFD Reaction model is turned on, STAR-CCM+ compiles the input files into a library. This compiling occurs either during the importing procedure or during the solution initialization.

The default FORTRAN compiler on windows is Intel Fortran 12.1 and on Linux is Intel Fortran 11.1.

To compile the Dars-CFD files manually:

The script `darscfdPack` is provided in the folder `star/bin/darsCfd`. From the command line, navigate to the folder where the Fortran 90 files are located, and call the script:

```
darscfdPack [-f fortranCompilerOption]
```

Option `-f` is not always necessary, since most of the supported architectures have only one FORTRAN compiler available. This option is only useful on Linux architectures (the options are the same as described above). The scripts `darscfdPackWin32` or `darscfdPackWin64` should be used on Windows.

The script creates the folder with the name derived from the computer architecture and the FORTRAN compiler, and put the case-specific library in that folder.

Manual compilation can be useful when you want to run the simulation on a cluster in which some nodes do not have a FORTRAN compiler installed. (The compilation is performed on the master node, but all nodes on a cluster have an equal chance of being a master node.) In this case, pre-compile the library on a different computer and use that library (with the Run Time Compile expert property deactivated) on the cluster.

Importing the Complex Chemistry Definition

Right-click the Dars-CFD Reaction Model node and select Import Complex Chemistry Definition (Dars-CFD format)...



The *Open* dialog appears. If the checkbox of the Binary Problem Description property is activated:

- If the checkbox of the Run Time Compile expert property is activated, then select the folder where the `*f90` files are located.
- If the checkbox of the Run Time Compile expert property is deactivated, then select the folder where the case-specific pre-compiled library (`libdarscfdstarccm+.so` or `libdarscfdstarccm+.dll`) is located.

To use ASCII input files for case definition, select the folder that contains appropriate input files (`InputRedKinMec.txt` and `InputRedKinTherm.txt`).



To replace the definition, select Delete Complex Chemistry Case... from the pop-up menu of the Dars-CFD Reaction Model node, and then repeat the import step. If you try importing without deleting, the following error message appears:



The Dars-CFD Reaction model only works when the Read-Only Thermodynamic Polynomial Data method is selected for specific heat calculation of the mixture components. If you are converting some other chemistry model to this one, it is possible that this method is not selected

automatically. In this case, examine all the mixture components and change the active method of calculation of the specific heat to Read-Only Thermodynamic Polynomial Data.

It is also possible to import complex chemistry definitions directly from a Chemkin-formatted file.

Importing Precompiled Chemical Libraries

STAR-CCM+ provides a selection of chemical mechanisms for different fuels as precompiled chemical libraries.

- To import a precompiled chemical library, right-click the Dars-CFD Reaction Model node and select Import Case From STAR-CCM+ Repository.

  The Dars-CFD solver makes use of Jacobian matrices to perform ODE integration of chemical kinetic equations. These Jacobian matrices can be computed numerically or analytically.

- In the submenus, select from the two types of libraries that are available: numerical or analytical.



Numerical Jacobian Libraries

The numerical approach uses the divided difference approach. Jacobian matrices tend to be sparse (that is, many entries in the Jacobian matrix are zero). However, evaluating the Jacobian matrices using the divided

difference approach does not exploit this sparsity, which makes the computations expensive. Also, the numerical Jacobian is less accurate due to divided difference approximation.

Analytical Jacobian Libraries

For a given chemical mechanism, the non-zero elements of the Jacobian can be computed analytically by considering all of the species and reactions that participate in the mechanism (see Reference below). This approach produces more accurate Jacobians, and involves less computation time since only the non-zero entries of the Jacobian are evaluated.

It is recommended to use the analytical Jacobian libraries rather than the numerical libraries wherever applicable, since the analytical libraries significantly reduce the computation time.

Reference

Lu, T. 2012. "Computational Tools for Diagnostics and Reduction of Detailed Chemical Kinetics", 2012 Princeton-CEFRC Summer School on Combustion, Princeton University.

Moving a Simulation File with a Complex Chemistry Model to Another File System

To move a complex-chemistry simulation file to another file system (the computer of a colleague, or a cluster of machines), update the path to the complex chemistry description. To update this path, edit the properties of the Dars-CFD Reaction model. The library is imported when the solution is initialized.

This dependency exists because the path to the complex chemistry description, ASCII, or binary, is recorded as an absolute path in the simulation. If the simulation file is moved to a different file system, that path is no longer valid.

Complex Chemistry Formulation

A detailed complex chemistry model is implemented in STAR-CCM+ with the Dars-CFD library. This library includes the stiff ODE (Ordinary Differential Equation) solver to integrate the chemical composition equations using gas mechanisms. There are two ways of decoupling the general species transport equation in STAR-CCM+:

- The Operator Splitting Algorithm
- The Chemistry ADI Algorithm

For both models when run with Steady State Model, an artificial chemistry time step is introduced that is based on convection and diffusion fluxes in that cell.

The Operator Splitting Algorithm

The general species transport equation can be formulated as follows:

$$\frac{\partial}{\partial t}\rho\chi Y_k + \frac{\partial}{\partial x_j}(\rho u_j Y_k + F_{k,j}) = \omega_k \tag{1061}$$

where $F_{k,j}$ is the diffusion flux component and source term ω_k is the rate of production of species k.

This algorithm takes advantage of the different time scales involved for the chemical reactions and the flow field. The time integration of the chemical state (species mass fractions Y_k and enthalpy h) is performed in two steps:

- At the beginning of each time-step, the chemical state is integrated from state $(Y_k, h)^n$ to $(Y_k, h)^*$, accounting only for the chemical source terms:

  $$\frac{\partial}{\partial t}\rho\chi Y_k = \omega_k \tag{1062}$$

  The system of Eqn. (1062) is solved using the stiff ODE solver from the Dars-CFD library.
- The flow field is then integrated from $(Y_k, h)^*$ to $(Y_k, h)^{n+1}$ without the chemical source term. The following systems of equations are solved:

  $$\frac{\partial}{\partial t}\rho\chi Y_k + \frac{\partial}{\partial x_j}(\rho u_j Y_k + F_{k,j}) = 0 \tag{1063}$$

  $$\frac{\partial}{\partial t}\rho\chi h + \frac{\partial}{\partial x_j}(\rho u_j h + F_{h,j}) = \frac{\partial p}{\partial t} + u_j\frac{\partial p}{\partial x_j} + \tau_{ij}\frac{\partial u_i}{\partial x_j} \tag{1064}$$

This algorithm is more suitable for transient simulations, although it can be used in steady-state simulations as well. It is essential to maintain a low Courant number to ensure that the flow field can be decoupled from the chemical reactions.

The Chemistry ADI Algorithm

The general species transport equation can be formulated as follows:

$$\frac{\partial}{\partial t}\rho Y_k + \frac{\partial}{\partial x_j}(\rho u_j Y_k + D_{k,j}) = \omega_k \tag{1065}$$

where $D_{k,j}$ is the diffusion flux component and source term ω_k is the rate of production of species k.

Within each time-step n, the fluid solver goes through a number of iterations. The current iteration is denoted with k. For steady-state simulations, n and k denote the same index.

Chemistry Step

At the beginning of each iteration k the chemical state is integrated from state $(Y, h)^{n,0}$ to $(Y^*, h^*)^{n,k}$, accounting for the convection and diffusion fluxes given at the current iteration k through a point-implicit formulation.

In the subsequent equations, the following notation is used for clarity:

- Y_n represents the value at the beginning of the time-step, $Y^{n,0}$.
- Y_k represents the value at the k iteration of the n time-step, $Y^{n,k}$.
- Y_k^* represents the intermediate solution from the Dars-CFD solver after the chemistry step at iteration k.

The species subscript is not explicitly stated, but these equations apply to all the species. Where summation is needed, it is performed over all the species.

$$\rho\frac{Y_k^* - Y_n}{dt} + \left[F_k + \frac{\partial F_k}{\partial Y}(Y_k^* - Y_k)\right] = \omega_k^* + Y_k\left[\sum_{all_species} F_k + \frac{\partial F_k}{\partial Y}(Y_k^* - Y_k)\right]$$

$$\frac{h_k^* - h_n}{dt} + Fh_k + \frac{\partial Fh_k}{\partial h}(h_k^* - h_k) = 0 \tag{1066}$$

where F_k and Fh_k are the explicit part of species and enthalpy fluxes (both convective and diffusive) at iteration k and $\partial F_k/\partial Y$ is evaluated at Y_k, the current value for Y at iteration k. The system of Eqn. (1066) is solved using the stiff ODE solver from the Dars-CFD library.

Output -- Species Production Rate

Output from Dars-CFD is the source term at iteration k. Using the Taylor expansion around point Y_k^* we get:

$$\omega_{k+1} = \omega_k^* + \frac{\partial \omega_k}{\partial Y}(Y_{k+1} - Y_k^*) \tag{1067}$$

Since STAR-CCM+ is solving a delta form of the equation, the term must be expanded on the right-hand side of the expression for ω_k:

$$\omega_{k+1} = \omega_k^* + \frac{\partial \omega_k}{\partial Y}(Y_{k+1} - Y_k) + \frac{\partial \omega_k}{\partial Y}(Y_k - Y_k^*) \tag{1068}$$

In this case, the contribution to the source term is:

$$\omega_k^* + \frac{\partial \omega_k}{\partial Y}(Y_k - Y_k^*)$$

and the contribution to the matrix diagonal is $-\partial \omega_k / \partial Y$.

Fluid Step

The flow field is then solved for the iteration k to get $(Y, h)_{n,\, k+1}$ with the source term from the Dars-CFD solver.

For more details about equations solved in STAR-CCM+ for species and energy, see the following sections:

- Species Transport Formulation
- Segregated Fluid Energy Formulation

Surface Chemistry Formulation for Dars-CFD

With the surface chemistry model, three types of species are defined:

- Species in the gas phase with mass fractions denoted with $Y_i^{(g)}$
- Species on the surface at the gas-surface interface with site fractions denoted as $Z_i^{(s)}$. Some sites are "open" (denoted as a generic species $OS^{(s)}$) when they do not contain any species, and they can be used to model absorption from the gas phase.
- Species in the solid layer--bulk species with their mole fractions denoted as $W_i^{(b)}$

Limitations to surface chemistry modeling:

- Surface chemistry can only be applied to wall and baffle boundary conditions.
- Surface reactions converse the number of surface sites.
- No energy equations are solved for reactive surface.
- Surface growth and shrinking are not taken into account.

It is assumed that the gas species mass fractions at the reacting surfaces are the same to the gas species mass fractions at the center of the cell volume. Then species and energy transport equations for the gas phase have one more source term ($\omega_k^{(s)}$ in the species transport equation and $src_h(s)$ in the enthalpy equation) corresponding to surface reactions:

$$\frac{\partial}{\partial t}\rho\chi Y_k^{(g)} + \frac{\partial}{\partial x_j}(\rho u_j Y_k^{(g)} + D_{k,j}) = \omega_k^{(g)} + \omega_k^{(s)} \tag{1069}$$

$$\frac{\partial}{\partial t}(\rho\chi E) + \frac{\partial}{\partial x_j}(\rho u_j H + Dh_j) = s + src_h(s) \tag{1070}$$

It is assumed that no convection or diffusion of surface or bulk species is present. Therefore, the governing transport equations for surface and bulk species are the following:

$$\frac{\partial}{\partial t}Z_k^{(s)} = S_{Z_k^{(s)}} \tag{1071}$$

$$\frac{\partial}{\partial t}W_k^{(b)} = S_{W_k^{(b)}} \tag{1072}$$

These equations are solved simultaneously in Dars-CFD. Output for the gas phase depends on the algorithm that is used to transform the gas species transport equation from the PDE to ODE form (Operator Splitting or Chemistry ADI). Output for the surface phase consists of the new molar fraction of bulk species and site fraction of surface site species. However, there are two more source terms that are related to the surface/gas interaction — $\omega_k^{(s)}$ and $src_h(s)$ from Eqns. (1069) and (1070). Their form depends on the underlying complex chemistry algorithm. Unlike the source terms from the gas phase, these source terms are not linearized.

Operator Splitting Algorithm

This algorithm requires mean values for the source terms over the time-interval dt.

$$\omega_k^{(s)} = \rho\frac{Y_k^{*(g)} - Y_k^{(g)}}{dt} \tag{1073}$$

$$src_h(s) = \rho\frac{\sum_{k=1,g} Y_k^{*(g)} h_k^* - Y_k^{(g)} h_k}{dt} \tag{1074}$$

Index g denotes that summation is performed over the gas species. The h_k is the partial enthalpy of the gas species k. Unlike in the case of gas phase only, the source terms of species do not sum to 0 so this summation term has to be added to the mass conservation equation.

Chemistry ADI Algorithm

This algorithm obtains reaction source terms for the time-interval dt directly from the Dars-CFD interface and applies it to the species transport equation during the CFD step. Again, the source terms of species do not sum to 0 so this summation term has to be added to the mass conservation equation. The enthalpy source term is calculated as follows:

$$src_h(s) = \sum_{k=1,sb} \Omega_k h_k \tag{1075}$$

Index sb denotes that summation is performed over the surface site and bulk species. Ω_k is the surface site/bulk explicit source and h_k is the partial enthalpy of the surface site/bulk species k. The temperature of the reactive surface remains constant during the integration.

See also:

DarsCFD Reports

ISAT Approximation Formulation

To approximate functions that are expensive to calculate, Dr. Stephen B. Pope developed the In-Situ Adaptive Tabulation (ISAT) method [195]. Within the Dars-CFD reaction model framework, the ISAT method approximates the solution of the ODE (Eqn. (1062)) for a given initial condition through:

- Thermodynamic variable (temperature or total enthalpy).
- Composition variable (mass fraction or molar number).
- Possibly pressure and time step (depending on the type of the simulation).

These values constitute an independent variable or a query point (ϕ). The solution f of the ODE equation (Eqn. (1062)) is called reaction mapping.

The ISAT method starts with an empty table and gradually adds points to it. To solve the first ODE equation for ϕ_0 (a first query), ISAT does the following:

- It calls the ODE solver (Dars-CFD in this case) to calculate f_0.
- It creates an entry and adds it to the table.

Every table entry consists of:

- An independent variable ϕ

- A reaction mapping variable f
- A mapping gradient matrix, $A = \frac{\partial f}{\partial \phi}$
- An Ellipsoid Of Accuracy (EOA), derived by a singular value decomposition of the mapping gradient.

For all subsequent queries, ISAT does the following for query point ϕ_n:

- It traverses the table and finds a suitable point to check for the possible approximation.
- For the suitable point ϕ_t, it checks the value: $f_q = f_t + A_t(\phi_n - \phi_t)$

The following two scenarios can happen:

- Either f_q lies in the EOA, which means that the value f_q approximates, within the user-specified tolerance, the true value of $f(\phi_n))$. In other words, query ϕ_n is retrieved.
- Or f_q lies outside of the EOA. In this case, the true value of $f(\phi_n))$ is calculated and the difference between the two values is deduced. If this difference is smaller than the user-defined tolerance, this means that the EOA at point ϕ_n is too strict. In this case, the ellipsoid is grown to accept point f_q. If the difference is higher than the specified tolerance, the point is added to the table.

Points in the table are stored in the leaves of the binary space partition tree. This tree divides the region in the n-dimensional space (n being the dimension of the query point) currently occupied by the points into smaller subregions for faster search. When a new query point is added to the tree, an original leaf becomes now two leafs, split by a hyperplane. One new leaf contains the old point, which lies on one side of the hyperplane. The second new leaf, which contains the new query point, lies on the other side of the hyperplane.

Therefore the following properties characterize the ISAT method:

- Number of retrieves: the number of queries for which the reaction mapping was approximated
- Number of grows: the number of queries for which the reaction mapping was calculated but the point was not added to the table because the EOA of the old point in the table was grown
- Number of adds: the number of queries for which the reaction mapping was calculated and the point was added to the table

Growing of the EOAs can be deactivated through a property within the ISAT approximation method. Since EOA is only an approximation of the actual region of accuracy, growing of the EOA can introduce errors in the approximation table. Furthermore, it is possible that adding a point to the table will cover more future queries than growing the existing EOA.

ISAT Bibliography

[195] Pope, S.B. 1997. "Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation", *Combust. Theory and Modelling*, **1**, pp. 41-63.

Equilibrium Time Scale Formulation

The Equilibrium Time Scale method assumes the chemical composition relaxes to the local equilibrium composition at a time scale determined by flow and chemistry time scales. The equilibrium composition is computed using ISAT.

The chemistry time scale is computed based on the specified fuel species.

$$\frac{dY_i}{dt} = -\frac{Y_i - Y_i^{eq}}{\tau_{char}} \tag{1076}$$

where:

- Y_i is the mass fraction for species i.
- Y_i^{eq} is the local instantaneous thermodynamic equilibrium mass fraction for species i.
- τ_{char} is a characteristic time scale.

$$\tau_{char} = \tau_{flow} + \tau_{chem}\text{, where:} \tag{1077}$$

  - $\tau_{flow} = k/(\varepsilon c_{flow})$, where:
    - ε is turbulent dissipation rate.
    - c_{flow} is a turbulent rate constant with a default value of 4.

    For laminar flow, $\tau_{flow} = \Delta/(Uc_{flow})$, where:
    - Δ is the grid size.
    - U is the local velocity.
  - $\tau_{chem} = c_{chem} min_f\left(\frac{Y_f}{\omega_f}\right)$, where:
    - c_{chem} is a model constant with a default value of 2.
    - f is an index for a user-specified list of fuel species.
    - min_f is taking the minimum of all the user-specified fuel species.
    - Y_f is the fuel species mass fraction in a CFD cell.

- ω_f is the net reaction rate of species f.

What Is the CFL Scaling Factor?

For steady-state simulations, an artificial time-step is used to advance chemical reactions (see Complex Chemistry Formulation). This time-step can be observed through the Chemistry Time Step field function. Use CFL scaling to reduce the size of chemistry time-steps for steady-state simulations.

This feature appears as the Cfl Scaling Factor node, which becomes available with either complex chemistry model: Operator Splitting or Chemistry ADI. To access the properties of this node, select it.



The chemistry time-step is equal to the local time-step multiplied by the CFL Factor. The local time-step is based on the inverse of accumulated diffusion and convection fluxes of the species and the Under-Relaxation Factor of the species solver.

Since this object is a profile, you can use a constant value. However, you can also use a field function to manipulate the time-step. For example, to limit the time-step to 0.01:

```
if ($LocalTimeStep > 1e-2) $CflScalingFactor =
1e-2/$LocalTimeStep
```

CFL Scaling Factor Properties

Value	Sets the CFL scaling factor.

Transitioning from the User Reaction Model to the Passive Scalar Model

Earlier versions of STAR-CCM+ provided the User Reaction model. The initial purpose of user reacting scalars was handling user source terms. From STAR-CCM+ 7.04 onwards, the Passive Scalar model can be used to achieve the same objective.

The following sections demonstrate how you can convert legacy cases from the User Reaction model to the Passive Scalar model:

- Update the physics models
- Create the passive scalars
- Define the individual sources
- Define the passive scalar boundary conditions
- Rename the incompatible field functions
- Retaining temporary storage of passive scalars variable
- Additional notes

Update the Physics Models

Load your simulation file and open the physics model selection dialog. In the example case that is shown here, the original model selection is as shown in the following image:



In this example, the following models are removed:

- Laminar Flame Concept
- User Reacting Scalar
- User Reaction Model
- Operator Splitting

Add the following models:

- Eddy Break-up model from the *Partially-Premixed Combustion models* group box
- Passive Scalar model from the *Optional Models* group box

Create the Passive Scalars

The next step in the transition is to create the passive scalars and define them using the same names and properties as the original reacting scalars. By doing this action, you can reuse the new passive scalars in your original user-defined field functions.



User Reacting Model



Passive Scalar Model

Passive scalars do not have a separate trigger option or a positivity factor. You can still clip the minimum and maximum value of the scalars.



User Reacting Model



Passive Scalar Model

Define the Individual Sources

In order to preserve this behavior, activate the Passive Scalar Source Option in the Physics Conditions node of the region.



Define the passive scalar sources so that they are identical to the original user reacting scalar sources. The Passive Scalar Source node is equivalent to the previously available User Reacting Scalar Sources, while the Passive Scalar Source Derivative node is equivalent to the User Reacting Scalar Sources Jacobian node.



User Reacting Model



Passive Scalar Model

Define the Passive Scalar Boundary Conditions

Define all boundary conditions so that they match the original case. Most of the original values are retained, so only define the Passive Scalar > Constant value.



User Reacting Model



Passive Scalar Model

Rename Incompatible Field Functions

The original diffusion coefficient field function was called `DiffusionCoefficient`. This field function is now called `DiffCoef`. The syntax remains the same.



User Reacting Model



Passive Scalar Model

Retaining Temporary Storage of Passive Scalar Variables

For this example case, the temporary storage of the passive scalar variables must be retained by activating the Temporary Storage Retained option in the Passive Scalar node.



The following field functions are exposed by activating the temporary storage of the passive scalar variables:



Additionally, the temporary storage of the segregated species must also be activated by selecting the Temporary Storage Retained option in the Segregated Species node.



Additional Notes

The `LocalTimeStep` field function is not available for the first iteration, when the sources are evaluated. Any field functions that reference the `LocalTimeStep` generate an error message when you run the simulation. To resolve this issue, you can deactivate the Species Source Term for the first iteration, and then re-enable it for the remaining iterations.



Using Turbulence-Chemistry Interactions

CFD modeling is playing an increasingly important role in improving combustor design to reduce fuel consumption and pollutant emissions. Usually it is essential to include detailed chemistry mechanisms, particularly for emission predictions.

In RANS or LES simulations, the computational grid size is not fine enough to resolve the thin reaction zone where the chemical reaction takes place. The *turbulence-chemistry interactions* (TCI) require additional modeling. The Eddy Dissipation Concept (EDC) model has been developed to treat the TCI when considering a detailed chemical mechanism. The alternative to the EDC model is the Laminar Flame Concept model where TCI is not considered. When you select the EDC model, some of its calculations can be post-processed as field functions.

What Is the Eddy Dissipation Concept (EDC) Model?

In the Eddy Dissipation Concept (EDC) model, in each cell the space is divided into two zones: the fine structure and the bulk gas zone. All the chemical reactions are assumed to occur in the fine structures which represent the smallest turbulence scales, where all turbulence energy is dissipated into heat. The reactions are neglected in the bulk gas zone, and the mass transfer exchanges the reactant and product gases between the fine structures and the bulk zone. Therefore, in the EDC model, the subgrid chemistry scales do not need to be resolved at a cell level.

The EDC model is represented in the object tree by the Eddy Dissipation Concept node, which has properties and is subject to common menu items.



Eddy Dissipation Concept Model Properties

Fine Structure Time Constant	C_t in Eqn. (1081).
Fine Structure Length Constant	C_l in Eqn. (1082).

What Is the Laminar Flame Concept Model?

The Laminar Flame Concept model, represented with a node by that name, has no properties. Its activation simply tells STAR-CCM+ that you are setting up a Homogeneous Reactor simulation that does not include TCI.



Selecting a Turbulence-Chemistry Interaction Model

To specify the type of turbulence-chemistry interaction for a Homogeneous Reactor combustion simulation:

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas but Multi-Component Liquid is also acceptable.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Specify whether your combustion is partially premixed, non-premixed, or premixed (gases only) by selecting the corresponding radio button in the *Combustion* group.

- Select Homogeneous Reactor or Homogeneous Reactor Model with Surface Chemistry from the particular combustion group. If you select the latter model, the Laminar Flame Concept model is automatically activated provided the automatic selection of recommended models is enabled. The Laminar Flame Concept model is for homogeneous reactor simulations without considering TCI.
- Select the Operator Splitting Complex Chemistry model or the Chemistry ADI Complex Chemistry model from the *Complex Chemistry Model* group. Selecting the latter option activates the Dars-CFD Reaction model automatically.

  A requirement for either Complex Chemistry model is that species and enthalpy models be activated, so you need to select a flow model (Coupled or Segregated) from the *Flow* group as well. If you select Chemistry ADI with the Auto-select recommended physics models checkbox ticked, the models Segregated Flow, Segregated Species, and Segregated Fluid Enthalpy are also selected, since they are the recommended flow options for this type of complex chemistry model.

As a result of the Complex Chemistry selection, the *Turbulence Chemistry Interactions* group box becomes available. In the following example, the Eddy Dissipation Concept (EDC) model is being selected.



The EDC model is designed for turbulent flow, mainly Reynolds-Averaged Navier-Stokes (RANS). Therefore the Turbulent and RANS models are automatically selected.

Eddy Dissipation Concept Formulation

In the fine structures, the governing equations for the species and enthalpy are:

$$\frac{dY_\alpha^*}{dt} = -\frac{1}{\tau_r}(Y_\alpha^* - <Y_\alpha>) + \frac{\omega_\alpha^* W_\alpha}{\rho^*} \tag{1078}$$

$$\frac{dh^*}{dt} = -\frac{1}{\tau_r}\sum_{\alpha=1}^{n}(Y_\alpha^* h_\alpha^* - <Y_\alpha><h_\alpha>) \tag{1079}$$

Here $*$ represents the quantities in the fine structures, and $<>$ represents the cell mean values. Y_α, h_α, ω_α and W_α are the mass fraction, enthalpy, reaction rate, and molecular weight of species α. ρ is the density.

The residence time which indicates how long the species remain in the fine structures can be computed as:

$$\tau_r = \frac{1-\chi r^*}{\dot{m}^*} \tag{1080}$$

$\dot{m}^*$ is the mass exchange rate between fine structures and the bulk gas zone:

$$\dot{m}^* = \frac{1}{\tau^*} = \frac{1}{C_t}\left(\frac{\varepsilon}{\nu}\right)^{0.5} \tag{1081}$$

where ε and ν are the turbulence dissipation rate and the kinematic viscosity respectively, and C_t is the fine structure time constant with a default value of 0.4082.

r^* is the mass fraction of fine structures in the cell:

$$r^* = \left[C_l\left(\frac{\nu\varepsilon}{k^2}\right)^{0.25}\right]^3 \tag{1082}$$

C_l is the fine structure length constant with a default value of 2.1377.

χ is the fraction of the fine structures that is heated sufficiently and may react. χ equals 1 when a detailed chemistry is employed.

EDC Field Functions Reference

The following field functions are made available to the simulation when the Eddy Dissipation Concept (EDC) model is used. These objects all have the same base properties.

EDC Fine Structure Length Fraction	The fine structure length constant C_l in Eqn. (1082).
EDC Fine Structure Time	The fine structure time constant C_t in Eqn. (1081).
EDC Mean Reaction Rate Ratio	The ratio that is used to modify the chemical reaction source term.

Using the Coherent Flame Model (CFM)

This section describes the following about the Coherent Flame Model (CFM):

- What Is the CFM?
- Selecting the CFM
- Preparing the CFM for a combustion simulation
- The mathematical formulation of the CFM

Also included are reference sections on solvers and shared boundary and initial condition values related to the CFM.

What Is the Coherent Flame Model (CFM)?

The current CFM combustion model is designed for premixed flames where the fuel and oxidizer are perfectly mixed before entering the computational domain. The corresponding node (Coherent Flame Model in the following screenshot), which has its own properties and is subject to common menu items, is activated as soon as the CFM is selected.



In this model, a flame area density and a fuel mass fraction on the grid are tracked through transport equations. The mean species concentrations are obtained as functions of the mean fuel mass fraction and a one-step global reaction scheme, which is internally calculated based on the unburnt gas composition. A mean enthalpy transport equation is also solved. The mean temperature, density, and viscosity are then calculated knowing the mean enthalpy and species concentrations.

CFM Properties

The following properties allow you to set up the CFM parameters:

CFM Constant, A	The coefficient a, see Eqn. (1088).
CFM Constant, Alpha	The coefficient α, see Eqn. (1088).
CFM Constant, B	The coefficient B, see Eqn. (1094).
CFM Constant, Beta	The coefficient β, see Eqn. (1088).
Laminar Flame Speed Control	Selects the method for laminar flame speed calculation.

Gulder	Selects the Gulder method.
Metghalchi	Selects the Metghalchi method.
User Defined Profile	Selects the User-Defined method.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that STAR-CCM+ uses, do not change any properties within the *Expert* category. The values in that category reflect both the design of the model and the discretization approaches that have been optimized for accuracy and performance.

Convection	Selects the convection scheme. For guidance as to convection scheme selection, see Diffusion Term, p. 6349.

1st-order	Selects the first-order upwind convection scheme.
2nd-order	Selects the second-order upwind convection scheme.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.

Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Selecting the Coherent Flame Model

To solve combustion problems using the Coherent Flame model (CFM):

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas but Multi-Component Liquid is also acceptable.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select Premixed Combustion from the *Combustion* group.
- Select the option Coherent Flame Model from the *Premixed Combustion* group.
- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group.

Since most of the available combustion models are only valid for turbulent flows, the Reynolds-Averaged Navier-Stokes option is selected automatically as the viscous regime. The *Physics Model Selection* dialog then looks as shown below. In this example, the option Segregated Flow is selected, which also activates the Segregated Fluid Enthalpy model.



The remaining selections in this dialog depend on problem characteristics that are not directly related to simulating combustion:

- Select an appropriate model from the *Equation of State* group, depending on whether the fluid in your problem is a gas or a liquid.
- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- If radiative heat transfer is important, select one of the options in the *Radiation* group.
- Select Gravity from the *Optional Physics models* group if gravity forces significantly influence the solution, for example, in fire simulations.

Preparing the CFM for a Combustion Simulation

After model selection, the next steps for preparing the CFM are:

- Specifying mixtures
- Setting properties of the CFM
- Working with the unburnt gas fluid stream
- Setting common reference values:
  - Minimum Allowable Temperature
  - Maximum Allowable Temperature
  - Reference Pressure
- Setting boundary and initial condition values
- Setting the Flame Area Density Ignitor

Working with the Unburnt Gas Fluid Stream

A fluid stream named Unburnt Gas allows you to define the temperature and composition of the gas. The Unburnt Gas node appears whenever the CFM is selected.



The unburnt temperature can be determined from an isentropic relationship that is based on pre-ignition values of T and p, as follows:

$$T_u = T_0 \left(\frac{p}{p_0} \right)^{\frac{\gamma - 1}{\gamma}} \tag{1083}$$

The reference unburnt temperature set here corresponds to T_0 in this equation.

Coherent Flame Formulation

CFM assumes that:

- Combustion occurs in the flamelet region.
- The mean turbulent reaction rate can be expressed as the product of the flame surface density by the consumption rate of fuel per unit of flame area.

The implementation of CFM includes the following steps:

- Solving a transport equation for the fuel mass fraction
- Solving a transport equation for the flame area density
- Net flame stretch calculation
- Species calculation
- Bibliography

See also Laminar Flame Speed Calculation.

Solving a Transport Equation for the Fuel Mass Fraction

The Y_f transport equation for the CFM can be written:

$$\frac{\partial \rho Y_f}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} Y_f) - \nabla \cdot (\Gamma_{Y_f} \nabla Y_f) = \omega_{cfm} \tag{1084}$$

Here b is defined as:

$$b = \frac{Y_f - Y_{res}}{Y_{ft} - Y_{res}} \tag{1085}$$

Y_f is the fuel mass fraction and Y_{ft} is the mass fraction in the unburnt gas. For a pure premixed case, it is possible to obtain Y_{ft} by a transport equation for mixture fraction. Y_{res} is the amount of fuel that is left at the end of combustion. It equals 0 in lean and stoichiometric mixtures and is positive in rich mixtures.

The source term is calculated as follows:

$$\omega_{cfm} = -(\rho_u S_u \Sigma)(Y_{ft} - Y_{res}) \tag{1086}$$

where ρ_u and S_u are the unburnt density and laminar flame speed, respectively.

Solving a Transport Equation for the Flame Area Density

The equation for σ is:

$$\frac{\partial(\rho\sigma)}{\partial t} + \nabla \cdot (\rho\, \boldsymbol{u}\, \sigma) - \nabla \cdot (\rho\, \Gamma_\sigma\, \nabla\sigma) = S_\Sigma + \nabla \cdot (\Gamma_\sigma\, \sigma\, \nabla\rho) \tag{1087}$$

The source term S_σ includes the flame area production by stretch and destruction by fuel consumption and is calculated using the following equation:

$$S_\Sigma = \alpha\, K_t\, \Sigma - \beta\, \frac{\rho_u\, Y_{ft}\, S_u\, (1 + a\sqrt{k}/S_u)}{\rho\, Y_f}\, \Sigma^2 \tag{1088}$$

α, β and a are model parameters. The method for K_t calculation is described below.

Net Flame Stretch K_t Calculation

The methods that are developed in [196] are used. K_t is tabulated with two parameters:

$$\Gamma_K = \frac{K_t}{\varepsilon/k} = f\left(\frac{u'}{S_u}, \frac{I_l}{\delta_l}\right) \tag{1089}$$

u' is the turbulence intensity that is obtained from the following equation:

$$u' = \sqrt{2/3k} \tag{1090}$$

I_l is the integral length scale that is defined as:

$$I_l = \frac{C_\mu^{3/4}}{\kappa} \cdot \frac{k^{3/2}}{\varepsilon} \tag{1091}$$

where C_μ and κ are the turbulent viscosity coefficient and von Karman constant of the K-Epsilon model, respectively.

δ_l is the thermal laminar flame thickness that is calculated as follows:

$$\delta_l = 2\frac{\mu_b}{Pr\rho_u S_u} \tag{1092}$$

In this expression, Pr is the laminar Prandtl number of the burnt gas and assumed to be 0.9. μ_b is the molecular viscosity of the burnt gas that is calculated from Sutherland's law:

$$\mu_b = \frac{\mu_1 T_b^{1.5}}{\mu_2 + T_b} \tag{1093}$$

Here, μ_1 and μ_2 are constants with the values of 1.457e-6 and 110, respectively. T_b is the burnt gas temperature. Γ_k has the following form:

$$\Gamma_k = \Gamma_p - B \cdot \Gamma_q \tag{1094}$$

Γ_p and Γ_q are the flame production and quench due to the stretch; they follow the empirical correlation in [196]. B is a model parameter.

Species Calculation

The species mass fractions are obtained based on internal algebra relations. The fuel is calculated with Eqn. (1085) based on the progress variable b. The other species then are calculated based on the global reaction.

Coherent Flame Bibliography

[196] Meneveau, C., and Poinsot, T. 1991 'Stretching and quenching of flamelets in premixed turbulent combustion', Combust. Flame, 86, pp. 311-332.

Solver Reference

This section details the properties for the various solvers that are associated with the CFM:

- CFM Combustion Solver

CFM Combustion Solver

Controls the solution update for the CFM.

Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.9.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

Shared Boundary and Initial Condition Values Reference

This section details boundary and initial condition values that are common to the CFM:

- Flame Area Density

Fuel Mass Fraction is covered in a separate section.

Flame Area Density

The flame area density describes the flame distribution in the CFM and is defined as flame area per unit mass.



This boundary value and initial condition value is added to the Values node by selecting `Specified Value` as the Method in the Physics Conditions > Wall Combustion Scalar Option node.

Properties

This value is input as a scalar profile.

Using the Partially-Premixed Coherent Flame Model

This section describes:

- What is the Partially-Premixed Coherent Flame Model (PCFM)?
- Selecting the adiabatic and non-adiabatic variants of the PCFM
- Preparing the PCFM for a combustion simulation
- The mathematical formulation of the PCFM

What Is the Partially-Premixed Coherent Flame Model?

The Partially-Premixed Coherent Flame Model (PCFM) model is designed for simulating flames where the fuel and the oxidizer are partially premixed, that is, they are neither perfectly premixed (the premixed flames) nor completely segregated (the non-premixed flames).



A partially-premixed flame has the characteristics of both premixed and non-premixed flames. In the PCFM model, the premixed combustion process in a partially-premixed flame is handled with the Coherent Flame Model (CFM) and the subsequent non-premixed combustion process is simulated with the Presumed Probability Density Function (PPDF) model (with either the equilibrium chemistry model or the laminar flamelet chemistry model). The Adiabatic PCFM is for adiabatic reacting flow conditions in which adiabatic PPDF model is used, while the Non-Adiabatic PCFM is for non-adiabatic reacting flow conditions in which non-adiabatic PPDF model is used. This model also invokes the PCFM Ideal Gas model.

What Is the PCFM Ideal Gas Model?

This special version of the Ideal Gas model is used when the PCFM model is used in the simulation. The corresponding node (PCFM Ideal Gas in the following screenshot) is activated when the PCFM model is selected.



Selecting the Adiabatic PCFM

To solve combustion problems using an adiabatic PPDF model, follow the steps below. If the Auto-select recommended physics models checkbox is ticked, some of the selections are made automatically.

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select Partially-Premixed Combustion from the *Combustion* group.
- Select Adiabatic Partially-Premixed CFM from the *Partially-Premixed Combustion* group.
- Select Coherent Flame Model from the *Premixed Models for PCFM* group.
- Select Adiabatic PPDF from the *Non-Premixed Adiabatic Models for PCFM* group.
- Select the desired adiabatic model type, PPDF Equilibrium, or PPDF Flamelet, from the *Partially-Premixed PPDF Reaction* group.
- Select PCFM Ideal Gas from the *PCFM Equation of State* group.

- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group.

The *Physics Model Selection* dialog contains your selections similar to the example shown below. In this example, the options PPDF Equilibrium and Coupled Flow were selected.



The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- Select Turbulent from the *Viscous Regime* group as this type of combustion model is only valid for turbulent flows.
- Select Gravity from the *Optional Physics models* group if gravity forces significantly influence the solution, for example, in fire simulations.

Selecting the Non-Adiabatic PCFM

To solve combustion problems using a non-adiabatic PPDF model, follow the steps below. If the Auto-select recommended physics models checkbox is ticked, some of the selections are made automatically.

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas.

- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select Partially-Premixed Combustion from the *Combustion* group.
- Select Non-Adiabatic Partially-Premixed CFM from the *Partially-Premixed Combustion* group.
- Select Coherent Flame Model from the *Premixed Models for PCFM* group.
- Select Non-Adiabatic PPDF from the *Non-Premixed Non-Adiabatic Models for PCFM* group.
- Select the desired non-adiabatic model type, PPDF Equilibrium, or PPDF Flamelet, from the *Partially-Premixed PPDF Reaction* group.
- Select PCFM Ideal Gas from the *PCFM Equation of State* group.
- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group.

The *Physics Model Selection* dialog contains your selections similar to the example shown below. In this example, the options PPDF Equilibrium and Coupled Flow were selected.



The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- Select Turbulent from the *Viscous Regime* group as this type of combustion model is only valid for turbulent flows.
- Select Gravity from the *Optional Physics models* group if gravity forces significantly influence the solution, for example, in fire simulations.

Preparing the PCFM for a Combustion Simulation

Once the PCFM is selected, a CFM and PPDF model are also registered under Models. So preparations of a PCFM model are really to set up a CFM model and a PPDF model. The following steps can be followed:

- Set up the PPDF model. This mainly includes the importing/constructing a PPDF table. When constructing a PPDF table for equilibrium chemistry, it also requires to define the composition of a multi-component reacting mixture and working with PPDF fluid streams.
- Set up the CFM model. This mainly includes the CFM model parameters and setting the Flame Area Density Ignitor.
- Set up initial conditions, boundary conditions, and solver parameters for the transport equations that are involved in the simulation.

PCFM Formulation

In a partially premixed flame, the averaged value for any quantity Q at location $\underline{x}$ and time t, $\tilde{Q}(\underline{x}, t)$, can be calculated with the following form:

$$\tilde{Q}(\underline{x}, t) = \int_0^1 \int_0^1 Q(z, b) P(z, b; \underline{x}, t) dz db \tag{1095}$$

Here $P(z, b; \underline{x}, t)$ is the joint pdf of the mixture fraction z and regress variable b.

The joint pdf can be decomposed to a conditional pdf $P_z(z|b; \underline{x}, t)$ and a marginal pdf $P_b(b; \underline{x}, t)$:

$$P(z, b; \underline{x}, t) = P_z(z|b; \underline{x}, t) P_b(b; \underline{x}, t) \tag{1096}$$

Using the classical flamelet approximation of a bimodal form:

$$P_b(b; \underline{x}, t) = \tilde{b}\delta(1 - b) + (1 - \tilde{b})\delta(b) \tag{1097}$$

where δ is a δ function, we can get:

$$\tilde{Q}(\underline{x}, t) = \tilde{b}\tilde{Q}^u(\underline{x}, t) + (1 - \tilde{b})\tilde{Q}^b(\underline{x}, t) \tag{1098}$$

where $\tilde{Q}^u(\underline{x}, t)$ is the unburnt state and $\tilde{Q}^b(\underline{x}, t)$ is the burned state. In this model, the burned state is from the equilibrium or laminar flamelet chemistry model.

PCFM Nomenclature

The following is a guide to the terms used in this section.

b	regress variable
$P(z, b;\underline{x}, t)$	joint pdf of the mixture fraction z and regress variable b
$P_z(z\|b;\underline{x}, t)$	a conditional pdf on b
$P_b(b;\underline{x}, t)$	a marginal pdf
$\tilde{Q}$	Favre average of quantity Q
$\tilde{Q}^u(\underline{x}, t)$	Favre averaged quantity at unburnt state
$\tilde{Q}^b(\underline{x}, t)$	Favre averaged quantity at burned state
t	physical time
$\underline{x}$	physical position
z	mixture fraction

Using the Turbulent Flame Speed Closure (TFC) Model

The Turbulent Flame Speed Closure model is used as a method to model premixed combustion. To set up this model, use the properties of the Turbulent Flame Speed Closure Model node.

The TFC Reaction Model node is an inactive model object which indicates that the TFC model is used in a combustion simulation. However, the TFC Reaction Model has no properties to set.



Turbulent Flame Speed Closure Model Properties

The following properties allow you to set up the TFC parameters:

Turbulent Flame Speed Control	Selects the control mechanism option for TFC combustion.	
	Zimont	Selects the Zimont method.

Peters	Selects the Peters method.
User Defined Profile	Selects the User-Defined method.

Turbulent Flame Speed Closure Model Expert Properties

CD-adapco highly recommends that the expert properties remain at the default settings, unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that STAR-CCM+ uses. The values in that category reflect both the design of the model and the discretization approaches that have been optimized for accuracy and performance.

Secondary Gradients	Neglects or includes the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Includes both secondary gradients.
	Off	Excludes both secondary gradients.
	Interior Only	Includes the interior secondary gradients only.
	Boundaries Only	Includes the boundary secondary gradients only.
Convection	Selects the convection scheme. For guidance with convection scheme selection, see Diffusion Term.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

See also:

Selecting the TFC Model

Preparing the TFC Model for the Simulation

The TFC Combustion Solver

TFC Formulation

Selecting the TFC Model

This procedure explains how to activate the Turbulent Flame Speed Closure (TFC) model in a relevant physics continuum.

To select the TFC model:

- In the physics model selection dialog, select Multi-Component Gas from the *Material* group box.
- To specify that the problem involves the simulation of chemical reactions, select Reacting in the *Reaction Regime* group box.
- Select Premixed Combustion from the *Combustion* group box.
- Select Turbulent Flame Speed Closure Model from the *Premixed Combustion* group box.

  Since this combustion model is designed for turbulent flows, the Turbulent and Reynolds-Averaged Navier-Stokes options are selected automatically, provided the Auto-select recommended Physics models option is activated.

  Other models selected automatically are the TFC Reaction Model and the Ideal Gas (with combustion) model.

The physics model selection dialog looks as shown below.



Preparing the TFC Model for the Simulation

To prepare the TFC model, set the following properties:

- To choose the turbulent flame speed control method, select the Turbulent Flame Speed Closure Model node and set Turbulent Flame Speed Control to one of the following:
  - Zimont
  - Peters
  - User-Defined Profile
- Specify the scalar profile method, TFC Rate Coefficient, A.
- Work with laminar flame properties. See Using Laminar Flame Speed Methods.
- Set up the Fuel Mass Fraction Ignitor. See Using Ignitors.
- Activate and set the Fuel Mass Fraction value as an initial condition and as a boundary condition. See Working with the Fuel Mass Fraction.

Working with Zimont Turbulent Flame Speed

The Zimont turbulent flame speed correlation is the most widely used for this type of combustion simulation. Use the Zimont Turbulent Flame Speed node to set the parameters of this option.



- The Zimont option is selected by default when you activate the TFC model. However, to activate the Zimont Turbulent Flame Speed node:
  - Select the Turbulent Flame Speed Closure Model node.
  - Set Turbulent Flame Speed Control to `Zimont`.
- To define unburnt thermal diffusivity as a scalar profile:
  - Select the Zimont Turbulent Flame Speed node and set Unburnt Thermal Diffusivity Control to `User Defined Profile`.
  - Set the method of the User Defined Unburnt Thermal Diffusivity > Unburnt Thermal Diffusivity Profile node.
- To work with the flame stretch effect (Eqn. (1102)):
  - Select the Zimont Turbulent Flame Speed node and activate Flame Stretch Effect.
  - If desired, modify the μ_{str} constant (Eqn. (1103)) in the properties of the Flame Stretch Effect node.
  - By default, the critical strain rate effect (Eqn. (1104)) uses the chemical time scale method. To set it as a scalar profile, select the Flame Stretch Effect node and set Critical Strain Rate Control to `User Defined Profile`.

Zimont Turbulent Flame Speed Properties

Wall Effect Constant	Uses wall effects to model the quenching of the flame at walls. Select a setting from `0` (fully extinguished) to `1` (no effect).	
Unburnt Thermal Diffusivity Control	Selects the calculation technique for unburnt thermal diffusivity.	
	Laminar Flame Thickness Option	Uses the laminar flame calculation technique (see Laminar Flame Thickness).
	User Defined Profile	Allows you to set this diffusivity as a scalar profile.
Flame Stretch Effect	Controls the use of the flame stretch factor G (Eqn. (1102)).	
	Activated	Takes the flame stretch effect into account by representing the probability of unquenched flamelets.
	Deactivated	Leaves the flame stretch effect out of the calculations.

Working with Peters Turbulent Flame Speed

The Peters turbulent flame speed correlation is commonly used for internal combustion engines. Use the Peters Turbulent Flame Speed node to set the parameters of this option.



- To activate the Peters Turbulent Flame Speed node, select the Turbulent Flame Speed Closure Model node and set Turbulent Flame Speed Control to `Peters`.
- To set terms in the Peters correlation, as well as the wall effects, work with the properties of the Peters Turbulent Flame Speed node.

Peters Turbulent Flame Speed Properties

Wall Effect Constant	Uses wall effects to model the quenching of the flame at walls. Select a setting from `0` (fully extinguished) to `1` (no effect).
Constant, A1	The coefficient A_1, see Eqn. (1107).
Constant, A4	The coefficient A_4, see Eqn. (1107).
Constant, B1	The coefficient B_1, see Eqn. (1107).
Constant, B3	The coefficient B_3, see Eqn. (1107).

Expert Properties

Ewald's Corrector Constant	c_{ew}, see Eqn. (1107).

Working with User-Defined Turbulent Flame Speed

Use the User Defined Turbulent Flame Speed node to set the turbulent flame speed as a scalar profile.



- To activate the User Defined Turbulent Flame Speed node, select the Turbulent Flame Speed Closure Model node and set Turbulent Flame Speed Control to `User Defined`.
- To set wall effects, work with the properties of the User Defined Turbulent Flame Speed node.
- To work with this object as a scalar profile, select the User Defined Turbulent Flame Speed > Turbulent Flame Speed Profile node and choose a method.

User Defined Turbulent Flame Speed Properties

Wall Effect Constant	Uses wall effects to model the quenching of the flame at walls. Select a setting from `0` (fully extinguished) to `1` (no effect).

Working with the TFC Rate Coefficient, A

The TFC Rate Coefficient, A node represents a coefficient that is used by the TFC model (Eqn. (1101)). You can set it as a scalar profile.

The TFC Combustion Solver

This solver controls the solution update for the TFC combustion model.

Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.8.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

TFC Formulation

The TFC combustion model assumes that in a premixed combustion system, the reaction takes place in a thin layer which separates the whole domain to unburnt and fully burnt mixtures. The reaction rate is closed using the turbulent flame speed. A transport equation for the fuel mass fraction is solved to obtain the reaction progress:

$$\frac{\partial \rho Y_f}{\partial t} + \nabla \cdot (\rho u Y_f) = \nabla \cdot (\Gamma_{Y_f} \nabla Y_f) + \omega_{tfc} \tag{1099}$$

The source term is calculated as follows:

$$\omega_{tfc} = \rho_u S_t |\nabla Y_f| \tag{1100}$$

where ρ_u and S_t are the unburnt density and turbulent flame speed, respectively.

There are three options for the turbulent flame speed:

- Zimont
- Peters
- User-defined (a scalar profile)

An explanation of wall effects and a bibliography are included.

Zimont Turbulent Flame Speed

Zimont used the following correlation for the turbulent flame speed [197]:

$$S_T = AG(u')^{3/4} S_L^{1/2} \alpha_u^{-1/4} I_l^{1/4} \tag{1101}$$

Here, A is the model constant, which has a default value of 0.5. u' is the turbulent velocity. S_L is the laminar flame speed (see Laminar Flame Thickness). α_u is the unburnt thermal diffusivity of the unburnt mixture. I_l is the integral turbulent length scale.

The stretch factor G takes the stretch effect into account by representing the probability of unquenched flamelets which is obtained by integrating the log-normal distribution of the turbulent dissipation rate:

$$G = \frac{1}{2} erfc\left[-\frac{1}{\sqrt{2\sigma}}\left(\ln\frac{\varepsilon_{cr}}{\varepsilon} + \frac{\sigma}{2}\right)\right] \tag{1102}$$

where $erfc$ is a complementary error function and σ is the standard deviation of the distribution of ε computed with the following equation:

$$\sigma = \mu_{str} \ln(I_l/\eta) \tag{1103}$$

where I_l is the integral turbulent length scale, η is the Kolmogorov micro-scale, and μ_{str} is an empirical model coefficient with a default value of 0.28.

ε_{cr} is the turbulent dissipation rate at the critical strain rate g_{cr}:

$$\varepsilon_{cr} = 15\nu g_{cr}^2 \tag{1104}$$

where ν is the kinematic viscosity of the fluid.

A high value for g_{cr} suggests no occurrence of the flame stretch. One method to compute g_{cr} is to assume that it is proportional to the chemical time scale:

$$g_{cr} = \frac{BS_L^2}{\alpha_u} \tag{1105}$$

Peters Turbulent Flame Speed

The Peters correlation [198] for turbulent flame speed has the following form:

$$s_T = s_L(1 + \sigma_t) \tag{1106}$$

where:

$$\begin{aligned} \sigma_t &= -AB + \sqrt{(AB)^2 + C\frac{u'I_l}{s_L\delta_L^0}} \\ A &= \frac{A_4B_3^2}{2B_1} \\ B &= \frac{I_l}{\delta_L^0}c_{ew} \\ C &= A_4B_3^2 \end{aligned} \tag{1107}$$

where δ_L^0 is the laminar flame thickness, and c_{ew} is Ewald's corrector which has a default value of 1.0.

A_1, A_4, B_1, and B_3 are model constants with default settings of 0.37, 0.78, 2.0, and 1.0, respectively.

Wall Effects

A wall effect constant is used to model the quenching of the flame at walls by multiplying the constant when computing the turbulent flame speed. By modifying the constant from 0 to 1, the wall quenching effect on the flame in the vicinity of wall boundaries can be adjusted from fully extinguished to no effect.

TFC Bibliography

[197] Zimont, V.L., Polifke, W., Bettelini, M., and Weisenstein, W. 1998. "An Efficient Computational Model for Premixed Turbulent Combustion at High Reynolds Numbers Based on a Turbulent Flame Speed Closure", *ASME J. Eng. Gas Turbines and Power*, **120**(3), pp. 526-532.

[198] Peters, N. 2000. *Turbulent Combustion*, Cambridge University Press, Cambridge.

Using the Partially-Premixed Turbulent Flame Speed Closure (PTFC) Model

The Partially-Premixed Turbulent Flame Speed Closure (PTFC) model is designed for simulating flames where the fuel and the oxidizer are partially premixed, that is, they are neither perfectly premixed (the premixed flames) nor completely segregated (the non-premixed flames).



A partially-premixed flame has the characteristics of both premixed and non-premixed flames. In the PTFC model, the premixed combustion process in a partially-premixed flame is handled with the Turbulent Flame Speed Closure Model (TFC). The subsequent non-premixed combustion process is simulated with the Presumed Probability Density Function (PPDF) model (with either the equilibrium chemistry model or the laminar flamelet chemistry model). The Adiabatic PTFC model is for adiabatic reacting flow conditions in which the Adiabatic PPDF model is used, while the Non-Adiabatic PTFC model is for non-adiabatic reacting flow conditions in which the Non-Adiabatic PPDF model is used. This model also invokes the PTFC Ideal Gas model.

See also:

Selecting the Adiabatic PTFC Model

Selecting the Non-Adiabatic PTFC Model

Preparing the PTFC Model for a Combustion Simulation

PTFC Formulation

The PTFC Ideal Gas Model

When a Partially-Premixed TFC model is activated, the PTFC version of the Ideal Gas model is automatically selected.



Selecting the Adiabatic PTFC Model

This procedure explains how to activate the Adiabatic PTFC model in a relevant physics continuum.

To select the Adiabatic PTFC model:

- In the physics model selection dialog, select Multi-Component Gas from the *Material* group box.
- To specify that the problem involves the simulation of chemical reactions, select Reacting in the *Reaction Regime* group box.
- Select Partially-Premixed Combustion from the *Combustion* group box.
- Select Adiabatic Partially-Premixed TFC from the *Partially-Premixed Combustion* group box.

The following models are selected automatically, provided the Auto-select recommended Physics models option is activated:

- Adiabatic PPDF
- Turbulent Flame Speed Closure Model
- TFC Partially-premixed Reaction Model

- Select the adiabatic model type, PPDF Equilibrium, or PPDF Flamelet, from the *Partially-premixed PPDF Reaction* group box.

  Since this combustion model is designed for turbulent flows, the Turbulent and Reynolds-Averaged Navier-Stokes options are selected automatically.

  Other models that are selected automatically are Segregated Flow and PTFC Ideal Gas.

The physics model selection dialog displays the models that are selected. In the example shown below, the option PPDF Equilibrium is selected.



Selecting the Non-Adiabatic PTFC Model

This procedure explains how to activate the Non-Adiabatic PTFC model in a relevant physics continuum.

To select the Non-Adiabatic PTFC model:

- In the physics model selection dialog, select Multi-Component Gas from the *Material* group box.
- To specify that the problem involves the simulation of chemical reactions, select Reacting in the *Reaction Regime* group box.

- Select Partially-Premixed Combustion from the *Combustion* group box.
- Select Non-Adiabatic Partially-Premixed TFC from the *Partially-Premixed Combustion* group box.

  The following models are selected automatically, provided the Auto-select recommended Physics models option is activated:

  - Non-Adiabatic PPDF
  - Turbulent Flame Speed Closure Model
  - TFC Partially-premixed Reaction Model
- Select the adiabatic model type, PPDF Equilibrium, or PPDF Flamelet, from the *Partially-premixed PPDF Reaction* group box.

  Since this combustion model is designed for turbulent flows, the Turbulent and Reynolds-Averaged Navier-Stokes options are selected automatically.

  Other models that are selected automatically are Segregated Flow, Segregated Fluid Enthalpy, and PTFC Ideal Gas.

The physics model selection dialog contains your selections similar to the example shown below. In this example, the option PPDF Equilibrium was selected.



Preparing the PTFC Model for a Combustion Simulation

When the PTFC model is selected, TFC and PPDF models are selected automatically. To prepare the PTFC model, set up the TFC and PPDF models:

- Set up the PPDF model by importing or constructing a PPDF table.

  When constructing a PPDF table for equilibrium chemistry, also do the following:

  - Define the composition of a multi-component reacting mixture.
  - Work with PPDF fluid streams.
- Set up the TFC model. See Preparing the TFC Model for the Simulation.
- Set up initial conditions, boundary conditions, and solver parameters.

PTFC Formulation

The formulation for the PTFC uses the same principles as that of the Partially-Premixed Coherent Flame Model (PCFM) model.

Using PPDF Models

This section describes:

- The general concept of a PPDF combustion model
- How to select the adiabatic and non-adiabatic variants
- The role of mixture fraction variance
- If an Equilibrium PPDF model has been selected, how to define the composition of a multi-component reacting mixture
- The PPDF Weighted Sum of Gray Gases method
- Working with PPDF fluid streams
- How to set up a look-up table for mapping chemical species and other thermodynamic parameters to the quantities computed using this model. Depending on the selected model, the table may be either internally generated or constructed using data from an external program. The table size can be controlled, and the table contents may be inspected graphically by entering the PPDF Plots node.
- How to set up initial conditions, boundary conditions and solver parameters for the transport equations that are involved in the simulation
- The mathematical formulation of the adiabatic and non-adiabatic PPDF model variants

What Is a PPDF Model?

In a PPDF model, a few parametric variables such as the mean mixture fraction (and in some cases, the mean enthalpy) are tracked on the grid. The mean species concentrations, temperature, and density are obtained as functions of the mean parametric variables, after an averaging process around a presumed probability distribution (PPDF) that is considered to represent turbulent fluctuations.

The corresponding nodes appearing in the simulation tree depend on whether an adiabatic or non-adiabatic PPDF model is selected. For either an adiabatic or non-adiabatic selection, you make a further choice between an equilibrium or flamelet variant.

Density calculations can be performed by a PPDF-specific implementation of the ideal gas model, or for non-adiabatic cases, an ideal gas law model.

What Is an Adiabatic PPDF Model?

A class of PPDF combustion models that assume locally adiabatic reacting flow conditions (see Adiabatic PPDF Formulation). The corresponding node (Adiabatic PPDF in the following screenshot) is activated as soon as this type of model is selected. Select this node to access the properties of the Adiabatic PPDF model.



In cases with large eddy simulation, this model includes a sub-node that lets you choose methods for mixture fraction variance.

Adiabatic PPDF Properties

Dissipation constant	Represents the ratio of velocity and chemical species fluctuations (see The PPDF Model Applied to the Adiabatic Equilibrium, Eqn. (1115)).
Number of streams	Allows the number to be chosen from a drop-down list. You can have either two or three streams.

Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Convection	The convection scheme that is used. For guidance on selecting a convection scheme, see Convection Term, p. 6343.	
	1st-order	First-order convection scheme.
	2nd-order	Second-order convection scheme.

What Is an Equilibrium PPDF Model?

A class of PPDF combustion models that assume local instantaneous chemical equilibrium conditions. The corresponding node (PPDF Equilibrium in the following screenshot) may be activated when an adiabatic PPDF model is selected (see The Adiabatic Equilibrium Model, Non-Adiabatic

Equilibrium).



This node also serves as the starting point for generating a look-up table that contains all information about chemical species and reactions that are needed by the PPDF equilibrium model.

PPDF Equilibrium Properties

Variable Species Update Allowed	When this read-only property is activated, you can control the frequency at which STAR-CCM+ updates the species components. See What Is a PPDF Species Update?

What Is the PPDF Flamelet Model?

A PPDF combustion model that accounts for non-equilibrium and finite-rate chemistry effects in turbulent non-premixed flames (see The Laminar Flamelet Model). The corresponding node (PPDF Flamelet in the following screenshot) may be activated when an adiabatic PPDF model is selected.



This node also serves as the starting point for importing or constructing a look-up table that contains all information about chemical species and reactions that are needed by the PPDF flamelet model.

PPDF Flamelet Properties

Variable Species Update Allowed	When this read-only property is activated, you can control the frequency at which STAR-CCM+ updates the species components. See What Is a PPDF Species Update?

What Is a Non-Adiabatic PPDF Model?

A class of PPDF combustion models that assume non-adiabatic reacting flow conditions, as for example when there are heat losses due to radiation (see Non-Adiabatic PPDF Formulation). The corresponding node

(Non-Adiabatic PPDF in the following screenshot) is activated as soon as this type of model is selected. Select this node to access the properties of the Non-Adiabatic PPDF model.



In cases with large eddy simulation, this model includes a sub-node that lets you choose methods for mixture fraction variance.

Non-Adiabatic PPDF Properties

Dissipation constant	Represents the ratio of velocity and chemical species fluctuations (see The PPDF Model Applied to the Adiabatic Equilibrium, Eqn. (1115)).
Number of streams	Allows the number to be chosen from a drop-down list. You can have either two or three streams.

Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Convection	The convection scheme that is used. For guidance on selecting a convection scheme, see Convection Term, p. 6343.	
	1st-order	First-order convection scheme.
	2nd-order	Second-order convection scheme.

What Is the PPDF Ideal Gas Model?

This model is a special version of the Ideal Gas model, used when PPDF combustion models are activated in the simulation. In such cases, local density values are obtained from the PPDF look-up table rather than by directly invoking the Ideal Gas model. This model can only be applied to cases with negligible pressure changes. For example, it cannot be used to

simulate combustion in a closed system where pressure is expected to have an obvious rise. The corresponding node (PPDF Ideal Gas in the following screenshot) is activated when this model is selected.



What Is the Ideal Gas (Law) Model?

This Ideal Gas model is an alternative to the PPDF Ideal Gas model. It is available for use only with the Non-Adiabatic PPDF model.



The difference between these models is that with PPDF Ideal Gas, the species, density, and temperature are interpolated from the PPDF table. With this Ideal Gas model, only species are obtained from the table. The density is from the Ideal Gas law, and the temperature is computed from the enthalpy and species.

Since currently the PPDF table is generated under the assumption of constant pressure, the PPDF Ideal Gas is only applicable to constant pressure flows. With this Ideal Gas model, the pressure effects can be accounted for temperature and density; this is important for compressible flows.

This Ideal Gas option is only recommended for high-speed compressible flows.

Selecting an Adiabatic PPDF Model

To solve combustion problems using an adiabatic PPDF model:

- In the *Physics Model Selection* dialog, select Multi-Component Gas from the *Material* group.

- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select the Non-Premixed Combustion option in the *Combustion* group.
- Select the Adiabatic PPDF option from the *Non-Premixed Combustion* group, as shown below.



- Select the desired adiabatic model type, PPDF Equilibrium, or PPDF Flamelet, from the *PPDF Reaction* group. Your selection determines the number of streams that are available in the properties of the model object.
- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group.

If you choose the first option in each of the above groups, the *Physics Model Selection* dialog looks as shown below.



The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- Select Turbulent from the *Viscous Regime* group as this type of combustion model is only valid for turbulent flows.
- If gravity forces significantly influence the solution, such as in fire simulations, select Gravity from the *Optional Models* group.

Selecting a Non-Adiabatic PPDF Model

To solve combustion problems using a non-adiabatic PPDF model:

- In the model selection dialog (*Physics 1 Model Selection* in the following screenshot), select Multi-Component Gas from the *Material* group.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select the Non-Premixed Combustion option in the *Combustion* group.

- Select the Non-Adiabatic PPDF option from the *Non-Premixed Combustion models* group, as shown below. You can select either the equilibrium PPDF or flamelet PPDF model.



- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group. If you choose Coupled Flow, the model selection dialog appears as shown below.



This type of combustion model is only valid for turbulent flows, so the Turbulent and Reynolds-Averaged Navier-Stokes options are selected automatically along with the related energy model.

With the non-adiabatic PPDF model, you can use a Real Gas model as an alternative to the Ideal Gas model:

- Before choosing an option from the *Flow* group, deactivate the Auto-select recommended Physics models checkbox.
- To make the *Equation of State for PPDF* group box visible, select a *PPDF Reaction* option.
- Select a *Real Gas Equation of State* model.

Use the Thermodynamic Polynomial Data material properties method with this combination of models.

The Real Gas models are helpful in combustion applications that include high pressure (such as liquid rockets) and low temperature (transcritical and supercritical environments). These applications involve considerable departure from ideal-gas behavior.

The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- If radiative heat transfer is important, select Radiation in the *Optional Models* group, and then select one of the options in the *Radiation* group.
- If gravity forces significantly influence the solution, such as in fire simulations, select Gravity from the *Optional Models* group.

What Is Mixture Fraction Variance?

The mixture fraction variance is a measure of turbulent fluctuations in the local mixture fraction values (see The PPDF Model Applied to the Adiabatic Equilibrium, Eqn. (1116)). It is a scalar quantity whose initial and boundary condition values are entered in the *Properties* window as a scalar array profile.



Mixture fraction variance has additional controls when used in large-eddy simulations (LES).

Using Mixture Fraction Variance with LES

In PPDF cases (adiabatic or non-adiabatic) that use large eddy simulation (LES), it is possible to select one of two methods for mixture fraction variance: Transport Equation or Algebraic Relationship. The mixture

fraction variance is represented in the STAR-CCM+ physics continuum by the Mixture Fraction Variance node. Select this node to access the properties of the mixture fraction variance.



Mixture Fraction Variance Properties

Method	Selects one of two methods for mixture fraction variance.	
	Transport Equation	Uses standard STAR-CCM+ calculations.
	Algebraic Relationship	Allows user controls (default). Activates the Algebraic Relationship node.

What Is the Algebraic Relationship Method?

The Algebraic Relationship node represents this method. Select this node to access the properties of the algebraic relationship method.



Algebraic Relationship Method Properties

Cv constant	Sets a value for the C_v constant (Eqns. (1121) and (1122)).	
Internal Dissipation Constant Method	Selects the method for computing the dissipation constant C_d.	
	Activated	Uses a separate method (Eqn. (1122)) for computing the dissipation constant.
	Deactivated	Uses the dissipation constant value in the properties of the Adiabatic PPDF or Non-adiabatic PPDF (Eqn. (1115)).

Specifying Mixture Components

The Models node of a given continuum in a typical simulation contains a multi-component gas model and a gas mixture material that are represented by the Multi-Component Gas and Gas Mixture nodes as shown below.



STAR-CCM+ automatically creates the gas mixture, which has components and bulk properties (Gas Components and Mixture Properties in the above example) that need to be defined. For the PPDF Flamelet model, STAR-CCM+ generates these components automatically once the PPDF Flamelet table is constructed.

For PPDF Equilibrium models, the Gas Components node is initially empty. You define the chemical species making up your combusting mixture by adding new components to this node, as described in the section on adding components to a gas mixture.

The name and material properties of each new components are retrieved from the database and (Constant) property values are assigned by default. You should review the default material property settings of the mixture as well as the settings of each new component. If necessary, change the settings of the mixture properties and then the settings of each new component accordingly (see Managing Multi-component Materials).

Note that:

- If the database does not include a component that is needed for your combusting mixture, make a local copy of the STAR-CCM+ material database (file `props.mdb`) and add the required component and its physical properties to it.
- For a combustion simulation, you most likely want to select Thermodynamic Polynomial Data (if it is available) as the method for the Specific Heat property of each mixture component. If you choose an alternative method for Specific Heat (for example, Constant or Polynomial in T), then you are required to provide Heat of Formation and Standard State Temperature values for that component as well.

The PPDF Weighted Sum of Gray Gases Method

The PPDF Weighted Sum of Gray Gases (WSGG) method is a variant of the non-reacting WSGG method. It calculates the absorption coefficient using the weighted sum of gray gases method and a PPDF table.

In this method, concentrations of CO2 and/or H2O are derived directly from the table. This technique allows the frequency of species update in PPDF models to be variable without loss of accuracy.

The PPDF WSGG method is only available with non-adiabatic PPDF, and also requires radiation modeling. Work with the properties of this object as you would with the WSGG method.



Working with PPDF Fluid Streams

Fluid streams enable you to define the temperature and composition of each of the streams entering the combustion chamber. The Fuel Stream node(s) and Oxidizer Stream node appear whenever a PPDF model is selected.

If you choose two fluid streams for the number of streams in either the Adiabatic or Non-Adiabatic PPDF model, they are named Fuel Stream and Oxidizer Stream.



If you choose three fluid streams for the number of streams in either the Adiabatic or Non-Adiabatic PPDF model, they are named Fuel Stream 0, Fuel Stream 1, and Oxidizer Stream.



The properties of fluid streams cannot be changed as long as there is a PPDF flamelet or equilibrium table that is not empty. To change the properties, first delete the tables.

What Is a PPDF Equilibrium Table?

This table is a look-up table containing all information about chemical species and reactions that are needed by an Equilibrium PPDF model. The corresponding node (PPDF Equilibrium Table in the following screenshot) is activated when either an adiabatic or non-adiabatic equilibrium PPDF

model is selected. Select this node to access the properties of the PPDF Equilibrium Table. Right-click this node to access the pop-up menu of the PPDF Equilibrium Table.



If you want to construct a PPDF equilibrium table, make sure that either the thermodynamic polynomial or polynomial in T is chosen for calculation of specific heat for each component of the mixture.

The following sub-nodes are created under this node:

- Non PPDF Variables and PPDF Variables -- These variables allow you to control table size along various dimensions.
- PPDF XY Plots -- This feature enables you to inspect the results of the look-up table calculation graphically.

For defining properties of PPDF fluid streams, you would work with separate objects. The properties of fluid streams cannot be changed as long as there is a PPDF equilibrium table that is not empty. To change the properties, first delete the tables.

STAR-CCM+ does not duplicate large tables when copying continua; see section, Copying Continua with Tables for the recommended procedure.

PPDF Equilibrium Table Properties

The following parameters determine the characteristics of the table to be computed.

Type of the equilibrium algorithm	Two algorithms are available, Standard Equilibrium and Dars Equilibrium.
Relative Pressure of the mixture	Sets the relative pressure of the mixture, which is relative to the Reference Pressure of the continuum.
Default material database	The database from which the mixture components are imported into the mixture when importing the table. You can choose between available (open) material databases. This feature is flexible, in that if the mixture component is not present in the selected database, the standard database (if it is not the one selected) is checked.

PPDF Equilibrium Table Expert Properties

Identical Mixture Fraction Space	Selects the approach for the mixture fraction spacing of table entries. Activated: All entries in the PPDF table have the same mixture fraction spacing. This allows the updating of species to be vectorized and thus speed up the computation. Deactivated: Each entry has its own mixture fraction spacing, optimized to fit the values best for this entry.
High Mixture Fraction Limit	For mixture fraction values above a certain limit, the mixture can fail to burn because it is too fuel-rich. STAR-CCM+ does not perform the equilibrium calculation above this limit, but can obtain a composition for a non-reacting state based on pure mixing. To impose this non-reacting state, set this property at less than 1. The default setting is 1.0, which is the maximum.
High Limit Window Size	Allows for a gradual change in flow properties from the reacting to the non-reacting state. Specifically, when you set the High Mixture Fraction Limit expert property to a value lower than 1, the transition from the reacting equilibrium state to the non-reacting pure mixing state can be abrupt as well as discontinuous. A value that is set greater than 0.0 (default) for High Limit Window Size (between 0 and 1) allows a smoother transition between the two states.
Delete Extra Entries	When activated, deletes the species entries that do not exist in the mixture.

Table Not Complete — A read-only property that indicates whether the table contains only part of the species with which it was created.

PPDF Equilibrium Table Pop-Up Menu

Import — Reads a pre-computed table into the memory of STAR-CCM+ from a file with extension `.tbl`. When you import, make sure that all components defined in the file are present in the material database file `props.mdb` or in the default material database (if different from the standard material database). Otherwise you see an error message like this one.



Construct — Creates a table from scratch using the information that is required for import.

Export — Saves a table that is stored in the memory of STAR-CCM+ to a file with extension `.tbl`.

Delete — Clears a stored table from memory.

See also: common menu items.

What Is a PPDF Flamelet Table?

A look-up table containing all information about chemical species and reactions that are needed by the PPDF Flamelet model. The corresponding node (PPDF Flamelet Table in the following screenshot) is activated when an

adiabatic PPDF Flamelets model is selected. Select this node to access the properties of the PPDF Flamelet Table. Right-click this node to access the pop-up menu of the PPDF Flamelet Table.



How Is a PPDF Flamelet Table Constructed?

The current method of constructing a PPDF look-up table employs library data that are generated by the DARS package based on user-specified information on fuel and oxidizer types and combustion conditions. The data consists of a series of n files of form flameletxx.lib01 (where xx starts at 01 and goes up to n), stored in a separate subdirectory within the user's working directory.

The DARS data is then integrated over a PPDF beta function that is computed for each of the entries of mixture fraction and mixture fraction variance in order to generate the table. Its characteristics are determined by parameters in the properties and the following sub-nodes:

- Non PPDF Variables and PPDF Variables -- These sub-nodes allow you to control table size along various dimensions.
- PPDF XY Plots -- This feature enables you to inspect the results of this calculation graphically.

The properties of fluid streams cannot be changed as long as there is a PPDF flamelet table that is not empty. To change the properties, first delete the tables.

STAR-CCM+ does not duplicate large tables when copying continua; see section, Copying Continua with Tables for the recommended procedure.

PPDF Flamelet Table Properties

Default material database	The database from which the mixture components are imported into the mixture when importing or constructing the table. You can choose between available (open) material databases. This feature is flexible, in that if the mixture component is not present in the selected database, the standard database (if it is not the one selected) is checked.

PPDF Flamelet Table Expert Properties

Identical Mixture Fraction Space	Selects the approach for the mixture fraction spacing of table entries.	
	Activated	All entries in the PPDF table have the same mixture fraction spacing. This allows the updating of species to be vectorized and thus speed up the computation.
	Deactivated	Each entry has its own mixture fraction spacing, optimized to fit the values best for this entry.
Delete Extra Entries	When activated, deletes the species entries that do not exist in the mixture.	
Table Not Complete	A read-only property that indicates whether the table contains only part of the species with which it was created.	

PPDF Flamelet Table Pop-Up Menu

Import

Reads a pre-computed table into the memory of STAR-CCM+ from a file with extension `.tbl`. When you import, make sure that all components defined in the file are present in the material database file `props.mdb` or in the default material database (if different from the standard material database). Otherwise you see an error message like this one.



Construct

Creates a table from scratch using externally generated data files.

Export

Saves a table that is stored in the memory of STAR-CCM+ to a file with extension `.tbl`.

Delete

Clears a stored table from memory.

What Is a PPDF Species Update?

You can save considerable computing time by controlling the frequency at which STAR-CCM+ updates the solution for the species. By default, STAR-CCM+ updates the solution for every species component at every iteration.

When the update of variable species is allowed, the PPDF model contains a PPDF Species Update sub-node.



For the variable species update to be allowed, certain requirements must be met as follows, depending on which PPDF model you are using:

- The material properties of species components must be independent of the mixture properties. That is, the mixture properties must not have any multi-component parameters.
- With the PPDF Equilibrium model, the PPDF Ideal Gas model must be used.
- With the PPDF Flamelet (Adiabatic) model, the PPDF Ideal Gas model must be used.
- With the PPDF Flamelet (Non-Adiabatic) model, the Ideal Gas (with non-adiabatic flamelet) model must be used.

To control the species update:

- Open the PPDF Species Update sub-node of the PPDF model.
- For an unsteady simulation, choose between iterations and time-steps in the PPDF Species Update Frequency Trigger property.
- Select the PPDF Species Update > PPDF Species Update Frequency node.
- Enter a value for Frequency.

PPDF Species Update Properties

PPDF Species Update Frequency Trigger	Selects the frequency trigger type.	
	Iteration	Triggers an update every n iterations. Specify the update frequency in the PPDF Species Update Frequency node.
	Timestep	Triggers an update every time-step. This trigger type is only available for transient simulations.

Working with Table Size

The look-up tables that are used by certain combustion models in STAR-CCM+ have sub-nodes that let you control the table size in the properties. This interface lets you adjust table size along either the Mixture Fraction or Mixture Fraction Variance dimension. For two-stream PPDF equilibrium tables, there is a lower limit on the number of mixture fraction points (61) and mixture fraction variance points (31).

Additionally it includes Heat Loss Ratio control for Equilibrium models, and Scalar Dissipation Rate specification for Flamelet models. The nodes that represent these variables are grouped between two placeholder manager nodes in the object tree: Non PPDF Variables and PPDF Variables.



What Is the Heat Loss Ratio?

The Heat Loss Ratio node is used with the PPDF Equilibrium table and the Nox Equilibrium table. It determines the table size along the Heat Loss Ratio dimension (see Eqn. (1140)). This value should always be an odd number.

If you are using the Adiabatic PPDF model, only one heat loss ratio point is needed. In this case, you can create a PPDF equilibrium table with multiple heat loss ratio points, but only the zero-heat loss ratio point is used for calculations. If you are using the Non-Adiabatic PPDF model, a minimum of three heat loss ratio points is needed.

This object is editable as long as the table is empty (that is, before importing or constructing). If the table is populated, the property cannot be changed. The only way to change it is to delete the table first.

What Is the Scalar Dissipation Rate?

The Scalar Dissipation Rate node is used with the PPDF Flamelet table and the Nox Source Flamelet table. It determines the table size along the Scalar Dissipation Rate dimension (Eqn. (1134)). This object is editable as long as the table is empty (that is, before importing or constructing). If the table is populated, the property cannot be changed. The only way to change it is to delete the table first.

What Is the Mixture Fraction Variable?

The Mixture Fraction node determines the table size along the Mixture Fraction dimension. It is common to all of the look-up tables.

What Is the Mixture Fraction Variance Variable?

The Mixture Fraction Variance node determines the table size along the Mixture Fraction Variance dimension. It is common to all of the look-up tables.

Properties of Table Variables

Table Dimension	Specifies the table size for all the variable nodes. This property is editable as long as the table is empty (that is, before importing or constructing). If the table is populated, the property cannot be changed. The only way to change it is to delete the table first.
Axis Index	Index reference for the Mixture Fraction and Mixture Fraction Variance nodes (read-only).

Working with PPDF XY Plots

The PPDF Plots manager node, which has its own pop-up menu, contains all the PPDF plot nodes in the simulation.



It is a sub-node of a PPDF Equilibrium table or PPDF Flamelet table. This manager object can contain any number of plots up to the memory limits of the computer and the plots are saved in the database when the simulation is saved.

This section includes information on creating a PPDF plot.

What Is a PPDF Plot?

PPDF plots allow you to visualize the data in the PPDF Equilibrium table or PPDF Flamelet table. A TablePlot node, which has its own properties and pop-up menu, represents the plotting of a particular variable. The node name includes the variable plotted.



You can modify elements of this plot, such as the axes, legend and symbols, just like other plots by customizing the common properties.

PPDF Plot Properties

Variable Name	The name of the variable plotted (read-only).
X-Axis Variable	You can choose the x-axis variable from the drop-down list: • Mixture Fraction • Mixture Fraction Variance • Scalar Dissipation (flamelet table) or Heat Loss Ratio (equilibrium table). You can choose the variable only when the plot is empty. It is your responsibility to make sure that the variable makes sense —for example, that you have more than one heat loss ratio point in the table or that you are not plotting mixture fraction variance for zero mixture.
Title	The overall title of the plot. This label appears centered at the top of the plot. Its position and format attributes cannot currently be modified. To remove the title, set this value to blank.

PPDF Plot Expert Properties

Footer	A text box in which you can type a footer for the plot.
Title Font	The font family, font style, and font size. Use the customizer to set this property (changing the text has no effect). The property value is displayed using the selected font family, style, and size.

PPDF Plot Pop-Up Menu

These pop-up menu items are available when you right-click a PPDF plot node.

Add New Table View for Plotting...	Completes the process of creating a PPDF plot by adding data to the plot display.

See also: common menu items.

Creating a PPDF Plot

Prerequisites

- Before you begin creating a PPDF plot, finish the process of creating a PPDF Equilibrium table or PPDF Flamelet table.

Creating the Plot

- To create the plot, right-click the PPDF XY Plots manager node and select a variable from the submenu of the New Plot menu item.



A new node for the plot appears in the object tree, and an empty plot display is opened.



Adding Data to the PPDF Plot Display

To add data to the plot display:

- Select the node of the PPDF XY plot, and choose a setting for the X-Axis Variable property.



- Right-click the PPDF XY plot node and select Add New Table View for Plotting...



- To control the data that appears in the plot, select one of the tabs in the *Select points for plotting* dialog. Depending on your choice of tab, make entries in the text boxes as follows:



  - Select values to plot -- Specify any values within a given range of available table data, and STAR-CCM+ interpolates values for the plot.
  - Select table points to plot -- Specify particular table entries to plot.

After you click OK in the dialog, the plot display contains the data.



- To overlay plots for multiple sets of points in the same plot display, repeat the procedure, entering different values in the *Select points for plotting* dialog.



Pop-Up Menu of the PPDF Plots Manager

This pop-up menu is available when you right-click the PPDF Plots manager node.

New Plot	Adds a PPDF plot with a name that includes the variable plotted.

Setting Initial Conditions

Following a PPDF model selection in the *Physics Model Selection* dialog, the Initial Conditions node contains two additional sub-nodes, one for the Mixture Fraction and the other for the Mixture Fraction Variance, as shown below.



Use the corresponding Value nodes to define initial conditions for these variables. Note that:

- The mixture fraction value may be left at its default value of 0.0 unless an initial distribution of fuel within the combustor needs to be specified.
- The mixture fraction variance value is normally left at its default value of 0.0.
- An initial temperature distribution can be specified only if a Non-Adiabatic PPDF model is selected.

What Is Mixture Fraction?

The mixture fraction is a measure of the local fuel/oxidizer ratio in a combustion system consisting of separate fuel and oxidizer streams (see The Adiabatic Equilibrium Model, Eqn. (1108)). It is a scalar quantity whose initial and boundary condition values are entered in the *Properties* window as a scalar array profile.



Setting Boundary Conditions

Following a PPDF model selection in the *Physics Model Selection* dialog, the node of every **open** boundary contains two additional sub-nodes, one for the Mixture Fraction and the other for the Mixture Fraction Variance, as shown below.



Use the corresponding Value nodes to define boundary conditions for these variables. Note that:

- Mixture fraction values at fuel inlets should be set equal to 1.0. Values at oxidizer inlets should be set equal to 0.0
- Mixture fraction variance values at both types of inlet are normally left at their default value of 0.0

Boundary condition input at **closed** (wall) boundaries for mixture fraction and its variance depends on the setting that is chosen for the Wall Combustion Scalar Option:

- Zero Flux -- no input is required.
- Specified Value -- separate nodes for each variable appear under the Physics Values node. Boundary flux conditions are then entered in the corresponding *Properties* window as a scalar array profile.

In addition, selecting an Adiabatic PPDF model automatically imposes adiabatic conditions on all wall boundaries.

What Is the Wall Combustion Scalar Option?

This option uses zero or a choice of values as a wall boundary condition for combustion-related scalars. You can select these scalars in the properties of this object.



Wall Combustion Scalar Option Properties

Wall Combustion Scalar	Selects the scalars for wall combustion calculation.	
	Zero Flux	No calculation. No input is required and no additional sub-nodes appear under the Physics Values node.
	Specified Value	When the Partially-Premixed Combustion or Non-Premixed Combustion models are selected, the following sub-nodes appear: Mixture Fraction Profile, Mixture Fraction Variance Profile, and Progress Variable boundary value nodes. When the Premixed Combustion model (available for the PVM) is selected, the following sub-nodes appear: Progress Variable and Progress Variable Variance.
	Specified Flux	Activates the Mass Flux and Mixture Fraction Profile boundary value nodes. This option is not available when the Premixed Combustion model is selected with the PVM.

Using the PPDF Combustion Solver

The PPDF combustion solver is activated by default, after any PPDF-type combustion model has been selected in the *Physics Model Selection* dialog and a case geometry has been imported. It is used to solve the transport equations for the mixture fraction and mixture fraction variance (see The PPDF

Model Applied to the Adiabatic Equilibrium).



Adiabatic PPDF Formulation

The adiabatic PPDF model encompasses two underlying concepts:

- The adiabatic model: few scalar variables in a locally adiabatic reacting flow system can represent the entire species set and temperature at any point in space.
- The PPDF model: a probability density function can represent the turbulent micromixing process.

The PPDF approach forms the basis of two models:

- The adiabatic equilibrium model.
- The laminar flamelet model.

Note that these models are only applicable to non-premixed combustion processes.

The Adiabatic Equilibrium Model

In this model, the instantaneous value of any scalar at any spatial location and time is uniquely determined by the atomic concentrations and initial temperature at that location and time. Since, from mass conservation, the number of atoms in the reactor as whole does not change, a single conserved scalar can characterize the reaction state space. In other words, what is

mixed has reacted and reached chemical equilibrium. This approach accommodates both forward and reverse reactions. A basic assumption of the model is that all species diffuse at the same rate, which is reasonable for turbulent flows where turbulent diffusivity is much greater than molecular diffusivity.

Under this assumption, if one defines a mixture fraction

$$f = \frac{m_p}{m_p + m_s} \tag{1108}$$

where m_p is the total mass of all atoms originating from the primary inlet (typically, fuel) at any spatial location, and m_s is the total mass of all atoms originating from the secondary inlet (typically, oxidizer), it can be seen that f is a conserved scalar. The equation for such a scalar is simple; f is transported by convection and diffusion, and can accumulate locally, and there is no f production when there is no phase change. Therefore:

$$\frac{\partial}{\partial t}(\rho\chi f) + \nabla\bullet[\rho U f - \rho D_f \nabla f] = 0 \tag{1109}$$

Note that for steady-state flows there is no accumulation, so the first term on the left-hand side of Eqn. (1109) is zero. Boundary condition values are also needed for f. After obtaining the latter, we can calculate the value of any conserved scalar ϕ at any spatial location according to the following formula:

$$\phi = f\phi_p + (1 - f)\phi_s \tag{1110}$$

where ϕ_p and ϕ_s are the values of the conserved scalar ϕ in the primary and secondary streams. In particular, if ϕ stands for the concentration of atoms of a particular element, then knowing f automatically gives the concentration of that element's atoms at any point. Once the atomic concentrations are known, STAR-CCM+ feeds them and the initial temperature to an equilibrium routine that yields:

- The mass fractions of all species;
- the density;
- the temperature at that point using a Gibbs free-energy minimization technique.

The only inputs that are needed for the equilibrium model are the boundary condition values for f and the specification of the $f = 1$ and $f = 0$ streams in every respect, that is, temperature, molecular weight, atomic composition...

This model has been extended for multiple streams.

Extension of the Adiabatic Equilibrium Model to Multiple Streams

If we have n streams, with the total mass of each stream i represented by m_i, then the mixture fraction j is defined as:

$$f_j = \frac{m_j}{\sum_{i=1}^{n} m_i} \tag{1111}$$

In this case, we have to solve $n-1$ transport equations that are defined in Eqn. (1109) and the n-th mixture fraction can be calculated as:

$$f_n = 1 - \sum_{i=1}^{n-1} f_i \tag{1112}$$

As in the case of the two fluid streams, the instantaneous values of the conserved scalar ϕ can be calculated at any spatial location according to the following formula:

$$\phi = 1 - \sum_{i=1}^{n} f_i \phi_i \tag{1113}$$

where ϕ_i is the value of the conserved scalar ϕ at fluid stream i.

The PPDF Model Applied to the Adiabatic Equilibrium

The adiabatic equilibrium model provides the instantaneous value of any variable of interest in the reaction state-space, such as density and temperature, if we provide it with the instantaneous mixture fraction. However, we also need averaged values of reaction state-space variables that account for turbulent fluctuations at any spatial location. Four steps are necessary for obtaining these averaged values:

- Perform a time average of Eqn. (1109), as is done with the momentum equations. This yields the time-averaged value (the first moment) of f, $\bar{f}$, from the following equation:

$$\frac{\partial}{\partial t}(\rho\chi\bar{f}) + \nabla\bullet\left[\rho U\bar{f} - \left(\rho D_f + \frac{\mu_t}{Sc_f}\right)\nabla\bar{f}\right] = 0 \tag{1114}$$

  where μ_t, is the turbulent viscosity and Sc_f the turbulent Schmidt number for the mixture fraction.

- Derive an equation for the second moment g_f, also known as the variance, of the mixture fraction:

$$\frac{\partial}{\partial t}(\rho\chi g_f) + \nabla\bullet\left(\rho U g_f - \frac{\mu_t}{Sc_g}\nabla g_f\right) = \frac{2\mu_t}{Sc_g} tr\ (\nabla f{:}\nabla f) - C_d\rho\frac{\varepsilon}{k}g_f \tag{1115}$$

where

$$g_f = \overline{(f-\bar{f})^2} \tag{1116}$$

Sc_g is the turbulent Schmidt number for the variance and C_d is a dissipation constant.

- Assume a shape for the PDF that depends only on these two moments. This presumed PDF of the mixture fraction is a beta function of the form:

$$P(f) = \frac{f^{a-1}(1-f)^{b-1}}{\int_0^1 f^{a-1}(1-f)^{b-1}\,df} \tag{1117}$$

where a and b are related to the mean and variance by:

$$a = \frac{\bar{f}}{g_f}[\bar{f}(1-\bar{f}) - g_f] \tag{1118}$$

and

$$b = \frac{(1-\bar{f})}{\bar{f}}a \tag{1119}$$

- Obtain the averaged value of any quantity from the PDF as:

$$\bar{\phi}(\bar{f}) = \int \phi(f)P(f)\,df \tag{1120}$$

The mapping of the mean value of any scalar to the mean value of the mixture fraction and its variance is accomplished through a pre-computed table. A simple look-up and interpolation is all that is then necessary when these mean values are required.

For PPDF combustion models, although a transport equation for variance of mixture fraction is usually solved in RANS (Eqn. (1115), it is common practice to use an algebraic relationship in LES combustion. When assuming that production equals dissipation in the mixture fraction variance equation, the following relationship can be used:

$$g_f = C_v \Delta^2 (\nabla \tilde{f})^2 \tag{1121}$$

C_v is a user–adjustable constant with a default value of 1/12 [201]. Δ is the subgrid scale.

In order to fulfill the assumption that the production equals dissipation, the dissipation constant C_d (Eqn. (1115)) can be computed internally based on the following equation:

$$C_d = \frac{2 C_s^2 C_l}{C_v} \tag{1122}$$

Here C_l is the constant when computing the turbulent time scale (Eqn. (764) for Smagorinsky Subgrid Scale Turbulence or Eqn. (778) for Wale Subgrid Scale Turbulence). C_s is the constant for computing the length scale (C_s in Eqn. (761) for Smagorinsky Subgrid Scale Turbulence or C_w in Eqn. (774) for Wale Subgrid Scale Turbulence).

This model has been extended for multiple streams.

Extension of the PPDF Model to Multiple Streams

To obtain the averaged value of any quantity from the PPDF for multiple streams, we can use the following equation:

$$\phi = \int_0^1 P(f_1) \int_0^{(1-f_i)} P(f_2) \ldots \int_0^{(1-\Sigma_{i=1}^{n-2} f_i)} P(f_{n-1}) \phi(f_1, f_2, \ldots, f_{n-1}) df_1 df_2 \ldots df_{n-1} \tag{1123}$$

The complexity of computing the multiple integral comes from the limits that are not constant. To simplify this computation, we introduce the *Scaled Mixture Fraction* space:

$$\begin{aligned} s_1 &= f_1 \\ s_2 &= \frac{f_2}{1 - f_1} \\ &\vdots \\ s_{n-1} &= \frac{f_{n-1}}{1 - \Sigma_{i=1}^{n-2} f_i} \end{aligned} \tag{1124}$$

Now Eqn. (1123) can be reformulated as:

$$\phi = \int_0^1 P(s_1) \int_0^1 P(s_2) \ldots \int_0^1 P(s_{n-1}) \phi(s_1, s_2, \ldots, s_{n-1}) ds_1 ds_2 \ldots ds_{n-1} \tag{1125}$$

Four steps that are described in Eqns. (1114)-(1119) can be reformulated as follows to encapsulate the multiple stream PPDF approach:

- Since the scaled mixture fraction is not a conserved value, it is still necessary to perform a time average of the transport equations for the original mixture fraction space for each mixture fraction i:

$$\frac{\partial}{\partial t}(\rho \chi \bar{f}_i) + \nabla \bullet \left[\rho U \bar{f}_i - \left(\rho D_f + \frac{\mu_t}{Sc_f} \right) \nabla \bar{f}_i \right] = 0 \tag{1126}$$

  where μ_t is turbulent viscosity and Sc_f is the turbulent Schmidt number for the mixture fraction.

- Compute the scaled mixture fraction space based on Eqn. (1113).
- Derive an equation for the second moment of every scaled mixture fraction s_i, $g_s i$, also known as scaled mixture fraction variance:

$$\frac{\partial}{\partial t}(\rho \chi g_{si}) + \nabla \bullet \left(\rho U g_{si} - \frac{\mu_t}{Sc_f} \nabla g_{si} \right) = \frac{2\mu_t}{Sc_f} tr\ (\nabla \bar{s}_i : \nabla \bar{s}_i) - C_d \rho \frac{\varepsilon}{k} g_{si} \tag{1127}$$

  where:

$$g_{si} = \overline{(s_i - \bar{s}_i)^2} \tag{1128}$$

  Sc_g is the turbulent Schmidt number for the variance and C_d is a dissipation constant.

- Assume a shape for the PDF that depends only on these two moments. This presumed PDF of the scaled mixture fraction is a beta function of the form:

$$P(s_i) = \frac{s_1^{a_1 - 1} (1 - s_i)^{b_1 - 1}}{\int_0^1 s_1^{a_1 - 1} (1 - s_i)^{b_1 - 1} ds_i} \tag{1129}$$

  where a and b are related to the mean and variance by:

$$a_1 = \frac{\bar{s}_i}{g_{si}} [\bar{s}_i (1 - \bar{s}_i) - g_{si}] \tag{1130}$$

  and:

$$b_i = \frac{(1 - \bar{s}_i)\, a_i}{\bar{s}_i} \tag{1131}$$

The Laminar Flamelet Model

The main shortcoming of the equilibrium chemistry model is that non-equilibrium effects, such as flame stretch, cannot be considered. An overprediction of the CO level in hydrocarbon-air flames can also occur, as reported in [199]. In order to account for non-equilibrium and finite-rate chemistry effects in turbulent non-premixed flames, the laminar flamelet model is implemented in STAR-CCM+.

This model assumes that the thickness of the reaction zone is smaller than the smallest turbulence length scale, the Kolmogorov length scale. Therefore, the local instantaneous reaction zone structure is considered to be the same as that in a quasi-steady one-dimensional laminar flame that is called a laminar flamelet, as shown in the Figure below, albeit convected and distorted by the turbulent velocity field.



Idealized axisymmetric counterflow combustion system

The turbulent flame may thus be regarded as an ensemble of laminar flamelets [200]. The main effect of the turbulent velocity field on the structure of the laminar flamelet is through the stretch of the reaction zone. The scalar dissipation rate (or strain rate) describes this. At low values of strain rate, laminar flamelet structures approach the equilibrium state; at high strain rates, flame extinction can occur. Through solving for the quasi-steady, one-dimensional laminar flames under different strain rates and with detailed reaction mechanisms, non- equilibrium and finite-rate chemistry effects can be introduced into the turbulent combustion modeling.

With this assumption, the mixture fraction can replace the space coordinate in the one-dimensional system that is shown above, to yield the following equations [202]:

$$\frac{\partial Y_i}{\partial t} = R_{i,\,kin} + D_i \chi \frac{\partial^2 Y_i}{\partial f^2} \tag{1132}$$

$$\frac{\partial T}{\partial t} = -\sum_{i=1}^{N} \frac{h_i R_{i,\,kin}}{C_p} + \alpha\chi\left(\frac{\partial^2 T}{\partial f^2}\right). \tag{1133}$$

Here, χ is the scalar dissipation rate that is given by:

$$\chi = |\nabla f|^2 = \left(\frac{\partial f}{\partial x_K}\frac{\partial f}{\partial x_K}\right) \tag{1134}$$

and α is the thermal diffusivity, C_p the specific heat and h_i the enthalpy of species i. In a counterflow diffusion flame, what is known is the velocity strain rate.

The steady laminar flamelet method (SLFM) omits the transient terms and solves for the balance between diffusion and reaction along the one-dimensional flamelet as a function of the scalar dissipation rate and the mixture fraction. A flamelet library is generated by solving Eqn. (1132) and Eqn. (1133) for different values of χ and tabulating the results as functions of the mixture fraction along the line, f, and the scalar dissipation rate, χ. Therefore, any variable of interest, ϕ, can be expressed as:

$$\phi = \phi(f, \chi) \tag{1135}$$

The PPDF Model Applied to Laminar Flamelets

The application of the PPDF approach to the laminar flamelet model is similar to that for the equilibrium model. The essential difference here is that, given Eqn. (1135) for any state-space variable, we require a joint PDF of f and χ:

$$\bar{\phi}(\dot{f}, \bar{\chi}) = \iint \phi(f, \chi) P(f, \chi)\, df\, d\chi \tag{1136}$$

In resolving the issue of the joint probability over f and χ, note that in a one-dimensional laminar counterflow flame, χ is a function only of the strain rate and the mixture fraction, that is:

$$\chi = \chi_0 F(f) \tag{1137}$$

where χ_0 is proportional to the strain rate. We can integrate this equation with respect to f and obtain the integrated scalar dissipation over the PDF of f. This is then parametrized by the mean stoichiometric scalar

dissipation, $\bar{\chi}_{st}$, which is related to the mean scalar dissipation rate, $\bar{\chi}$, obtained from the CFD analysis. The latter is calculated by applying a simple relaxation model for the evolution equation of the scalar dissipation:

$$\bar{\chi} = \frac{g_f}{D_f(k/\varepsilon)}. \tag{1138}$$

To complete the description, an independence hypothesis for the joint probability of f and $\bar{\chi}_{st}$ is used, expressed as:

$$\bar{\phi}(\bar{f}, \bar{\chi}) = \int \phi(f, \bar{\chi}_{st}) P(f) df \tag{1139}$$

Again, as in the equilibrium PPDF model, the mapping of mean scalar values to mean values of scalar dissipation rates and mixture fractions is accomplished through a pre-computed table.

Extension of the PPDF Model for Laminar Flamelets to Multiple Streams

The laminar flamelet model is available for multiple streams. The formulation is similar to the formulation of multiple streams with adiabatic equilibrium. The only difference is that, as with the two-stream laminar flamelets case, the joint PDF of f and χ is calculated.

Adiabatic PPDF Bibliography

[199] Jones, W.P. 1980. 'Prediction methods for turbulent flames', in "Prediction Methods for Turbulent Flow" (Ed. W. Kollman). Hemisphere, Washington, D.C., pp. 1-45.

[200] Peters, N. 1986. 'Laminar flamelet concepts in turbulent combustion', 21st Symp. (Int.) on Combustion, The Combustion Institute, pp. 1231-1250.

[201] Pitsch, H., Desjardins, O., Balarac, G. and Ihme, M. 2008. "Large-Eddy Simulation of Turbulent Reacting Flows", *Prog. Aerospace Sci.*, pp. 466–478.

[202] Veynante D. and Vervisch L. 2002. 'Turbulent Combustion Modeling' in 'Progress in Energy and Combustion Science', Vol. ***28***, pp. 248-252.

Non-Adiabatic PPDF Formulation

This model is currently available only with STAR-CCM+'s non-adiabatic equilibrium model or with a laminar flamelet approach.

Non-Adiabatic Equilibrium

A heat loss ratio γ is used to indicate the amount of heat loss with the following equation:

$$\gamma = \frac{h_{ad}(f) - h}{h_{sens}(f)} \tag{1140}$$

where h_{ad} is the adiabatic enthalpy, h is the cell enthalpy and h_{sens} is the sensible enthalpy h_{ad}. and h_{sens} are a function of mixture fraction f.

PPDF Model Applied to Non-Adiabatic Equilibrium

In the non-adiabatic case, all properties are functions of the mixture fraction and enthalpy. Therefore, the averaged value of any scalar in a turbulent flow field is the instantaneous value that is integrated over the joint PPDF of f and h:

$$\bar{\phi}(\dot{f}, \bar{h}) = \int \phi(f, h) P(f, h) df dh \tag{1141}$$

STAR-CCM+ makes a statistical independence assumption for $P(f, h)$ so that:

$$P(f, h) = P(f) P(h) \tag{1142}$$

and further assumes a Dirac delta function for $P(h)$. In other words, the mean value of enthalpy is used in the calculation of:

$$\bar{\phi}(\dot{f}, \bar{h}) = \int \phi(f, \bar{h}) P(f) df \tag{1143}$$

The integrals in Eqn. (1120), Eqn. (1139) and Eqn. (1143) are pre-computed. When the value of an integral is needed during the calculation, a simple interpolation is all that is needed to provide the correct value.

The non-adiabatic PPDF model can work with Lagrangian evaporation models. In that case, a source term resulting from the evaporation of the Lagrangian phase is added to the right-hand side of the transport equation for mixture fraction f (Eqns. (1114) and (1126)).

Using the Progress Variable Model (PVM)

This section describes the following about the Progress Variable Model (PVM):

- What Is the Progress Variable Model (PVM)?
- What Is the PVM Ideal Gas Model?
- What Is the PVM Table?
- What is the Progress Variable Variance?
- Selecting the Progress Variable Model
- Preparing the PVM for a Combustion Simulation
- PVM Formulation
- PVM Bibliography

What Is the Progress Variable Model (PVM)?

The PVM combustion model is designed for premixed, non-premixed, and partially-premixed flames. The corresponding Progress Variable Model node, which has its own properties and is subject to common menu items, is activated as soon as the PVM is selected. When the Progress Variable Model is selected along with the Premixed Combustion model, a sub-node which allows you to specify a method for computing the progress variable variance, appears under the Progress Variable Model node.



For simulations involving combustion, it is often important to include a detailed reaction mechanism, which has many reactions and many species. Inclusion of such a mechanism in a CFD simulation with realistic geometries, however, inevitably leads to a large computational cost.

The PVM is designed to incorporate complex chemical mechanisms with a reduced computational cost. The effects of detailed chemistry on combustion are stored in a precomputed table, and the needed values are

retrieved during the CFD simulation. A progress variable is used to facilitate interactions between CFD computation and the precomputed table.

At the same time, in the table, STAR-CCM+ can identify species that have source terms and solve them with a transport equation. You can either define the source terms when generating the Standard PVM table in DARS-Basic, or use the TUe FGM Table with predefined source terms. The TUe FGM (Flamelet Generated Method) table consists of pre-defined information about fuel and oxidizer types and combustion conditions. It is only available for premixed combustion models. These species are accessible through Transported Mass Fraction Profile nodes.

A PVM-specific implementation of the ideal gas model performs density calculations.

PVM Properties

The following properties allow you to set up the PVM parameters.

Number of streams (Only available for Partially Pre-mixed or Non-Premixed Combustion models)	Allows you to choose the number of streams. Make the selection in this property before you begin constructing a PVM table.	
	2	two streams
	3	three streams
PVM constant, Alpha	A scaling factor to the table source term for the chemical enthalpy.	
Dissipation constant	The ratio of velocity and chemical species fluctuations.	

Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Convection	Selects the convection scheme to be used.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

What Is the PVM Ideal Gas Model?

This special version of the Ideal Gas model is used when PVM combustion is activated in the simulation. In such cases, the molecular weight that is stored in the PVM look-up table is used to obtain local density values, rather than by directly invoking the Ideal Gas Model. The corresponding node

(PVM Ideal Gas in the following screenshot) is activated when this model is selected.



PVM Ideal Gas Properties

Table Density — This option, when activated, uses the density value that is stored in the PVM table rather than computing density using the ideal gas law. For high-speed flows, deactivate this option to include pressure effects.

Activated	Gets the value for density directly from the PVM table.
Deactivated	Calculates density using molecular weight from the PVM table in conjunction with cell values of pressure and temperature.

What Is the PVM Table?

A PVM table is a look-up table that contains information about chemical species and reactions that the PVM needs. The corresponding node (PVM Table in the following screenshot), which has its own properties and pop-up menu, is activated when the PVM model is selected.



- When the Partially-Premixed or Non-Premixed Combustion models are selected, you can specify the information about chemical species and reactions in the Standard PVM table using library data that the DARS Basic package generates.
- When the Premixed Combustion model is selected, you can specify the information about chemical species and reactions with:
  - The Standard PVM table using library data that the DARS Basic package generates.
  - The data that the TUe FGM Table generates.

The data is then integrated over a PPDF beta function that is computed for each of the variable entries, in order to generate the table. Parameters in the sub-nodes and the *Properties* window determine the PVM table characteristics.

The following sub-nodes are created under the PVM Table node:

Non PPDF Variables	Displays independent variables that determine the number of points that are used from the data for that parameter. The values are specified upon import.	
	Progress Variable (Only for Partially-Premixed or Non-Premixed Combustion models)	The number of progress variable data points in the read-in data.
	Total Enthalpy (Only for the Premixed Combustion model)	The number of total enthalpy data points in the read-in data.
PPDF Variables	A placeholder for the following nodes that let you specify the table size. The Axis Index property is read-only. Table Dimension represents the number of points that are used from the imported table data. You can edit Table Dimension as long as the table is empty (that is, before importing or constructing). If the table is populated, the property cannot be changed unless you delete the table.	
	Mixture Fraction (Only for Partially-Premixed or Non-Premixed Combustion models)	Determines the table size along the Mixture Fraction dimension. See What is the Mixture Fraction Profile?
	Mixture Fraction Variance (Only for Partially-Premixed or Non-Premixed Combustion models)	Determines the table size along the Mixture Fraction Variance dimension. See What is the Mixture Fraction Variance Profile?
	Progress Variable (Only for the Premixed Combustion model)	An independent variable that the data from the TUe FGM Table determines. See What is the Progress Variable?
	Progress Variable Variance (Only for the Premixed Combustion model)	An independent variable that the data from the TUe FGM Table determines. See What is the Progress Variable Variance?

PVM XY Plots This feature allows you to inspect the results of the look-up table calculation graphically in the same way as the PPDF XY Plots.

STAR-CCM+ does not duplicate large tables when copying continua; see section, Copying Continua with Tables for the recommended procedure.

What is the Standard PVM Table?

The Standard PVM Table consists of data from a single file named `PVM.txt`. The DARS Basic package creates this file based on user-specified information about fuel and oxidizer types and combustion conditions.

What Is the TUe Flamelet Generated Manifold (FGM) PVM Table?

The TUe FGM PVM Table consists of pre-defined information about fuel and oxidizer types and combustion conditions. Selecting this table import option automatically defines initial conditions and boundary conditions for each type of boundary. The TUe FGM PVM table is only available for premixed combustion models and provides an alternative option for importing data.

The TUe FGM PVM Table is produced by the Eindhoven University of Technology.

PVM Table Properties

Absolute pressure of the table (Only available when Partially-Premixed or Non-Premixed Combustion models are selected)	The pressure that is used when generating the table in the DARS Basic package.
Default Material Database	The database from which the mixture components are imported into the mixture when importing or constructing the table. You can choose between available (open) material databases. This feature is flexible, -- if the mixture component is not present in the selected database, the standard database (if it is not the one selected) is activated.
Import Format (Only available when the Premixed Combustion model is selected)	Allows you to specify the source of the data that is used to generate the PVM table.

Standard	Allows you to import data that the DARS Basic package generates.
TUe FGM Table	Allows you to import data that the TUe FGM Table generates.

Expert Properties

Identical Mixture Fraction Space	Determines if all table entries are assigned identical mixture fraction space.

Activated	All table entries that are imported are assigned equal mixture fraction space.
Deactivated	Default setting. All table entries maintain the mixture fraction space that is defined in the imported data.

PVM Table Pop-Up Menu

Import...	This option allows you to read a pre-integrated table into STAR-CCM+'s memory from an external file with extension `.tbl` that the Export option saves. When importing, ensure that all components defined in the file are present in the selected material database file.
Construct...	This option allows you to create a table from scratch using a chemistry table that DARS-Basic generates. Before constructing the table, set the Number of streams property of the PVM to the correct option, 2 or 3, to reflect the number of streams in the case.
Export....	This option allows you to save a table that is stored in STAR-CCM+'s memory to a file with extension `.tbl`.
Delete....	This option allows you to remove a stored table from memory.

What is the Progress Variable Variance?

When the Progress Variable Model is selected along with the Premixed Combustion model, a Progress Variable Variance sub-node appears. The Progress Variable Variance node allows you to specify a method for computing the progress variable variance. Select this node to access the Progress Variable Variance Properties.



Progress Variable Variance Properties

Method	Allows you to select one of two methods for progress variable variance.	
	Transport Equation	Uses standard STAR-CCM+ calculations.
	Algebraic Relationship	Activates the Algebraic Relationship node which allows you to access the algebraic relationship properties.

Algebraic Relationship Properties

Cv constant	Allows you to set a value for the C_v constant in Eqn. (1146).

Selecting the Progress Variable Model

To solve combustion problems using the Progress Variable model (PVM), in the *Physics 1 Model Selection* dialog do the following:

- Select the material model Multi-Component Gas from the *Material* group.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select Non-Premixed Combustion, Partially-Premixed Combustion, or Premixed Combustion from the *Reactive Flow* group.
- Select Progress Variable Model from the appropriate *Combustion models* group.
- Select the desired solver type, Coupled Flow, or Segregated Flow, from the *Flow* group.

As this combustion model is only valid for turbulent flows, the Reynolds-Averaged Navier-Stokes option is selected automatically for the viscous regime. The *Physics 1 Model Selection* dialog then looks as shown below. In this example, Non-Premixed Combustion and Segregated Flow are selected; the latter option also activates the Segregated Fluid Enthalpy model.



The remaining options in this dialog depend on problem characteristics that are not related to simulating combustion. For example:

- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.

- Select Gravity from the *Optional Models* group if gravity forces significantly influence the solution, such as in fire simulations.

Preparing the PVM for a Combustion Simulation

After model selection, the next steps for preparing the PVM are:

- Specifying mixture components
- Setting properties of the PVM
- Setting properties of the PVM table
- Setting common reference values:
  - Minimum Allowable Temperature
  - Maximum Allowable Temperature
  - Reference Pressure
- Setting initial conditions
- Setting the Progress Variable Ignitor
- Working with the fluid stream
- Setting boundary conditions
- Solving for species with transport equations
- Setting solver parameters

Note: The only prerequisite to constructing the PVM table is setting the Number of Streams from the PVM Properties. Component information is imported with the PVM table that the DARS package generates.

Specifying Mixture Components

The Models node of a given continuum in a typical simulation contains a multi-component gas model that the Multi-Component Gas node represents.



The gas components, that STAR-CCM+ creates automatically, have components and bulk properties (Gas Components and Material Properties in the above example) that you define. For the PVM model, STAR-CCM+ automatically generates these components when the PVM table is constructed or imported.

Setting Initial Conditions

After selecting the Progress Variable Model in the *Physics 1 Model Selection* dialog, the Initial Conditions node is populated with extra sub-nodes:

For Partially-Premixed or Non-Premixed Combustion models, these sub-nodes are:

- Mixture Fraction Profile
- Mixture Fraction Variance Profile
- Progress Variable.



For Premixed Combustion models, these sub-nodes are:

- Progress Variable
- Progress Variable Variance



To define initial conditions for these variables, use the corresponding sub-nodes.

What is the Mixture Fraction Profile?

By default, the Method property is set to Constant. Maintain the default value for the Constant sub-node, `0.0`, unless you are specifying an initial fuel distribution. To set a scalar array profile value, set Method to Composite.

See What Is Mixture Fraction for the PPDF Model.

What is the Mixture Fraction Variance Profile?

For most cases, maintain the default value, 0.0.

See What Is Mixture Fraction Variance for the PPDF Model.

What is the Progress Variable?

The progress variable is a dimensionless parameter, which Eqn. (1145) defines, that shows the progress of the combustion. The progress variable is a scalar quantity, with Method set to Constant by default. The Constant value is 0 for the unburnt state and 1 for the fully burnt state.



You can set a value for the Progress Variable > Constant node or select a scalar array profile for the Method property and define the scalar profile values appropriately.

Maintain the default value, 0.0 for the Progress Variable > Constant Method unless you specify a non-zero combustion progress.

What is the Progress Variable Variance?

By default, Method is set to Constant. You can edit the progress variable variance profile in a similar way to the progress variable profile.

Setting Boundary Conditions

Inlet and Outlet Boundaries

After selecting the Progress Variable Model in the *Physics 1 Model Selection* dialog, each Inlet and Outlet boundary node is populated with extra sub-nodes.

When Partially-Premixed Combustion or Non-Premixed Combustion models are selected, the following sub-nodes appear:

- Mixture Fraction Profile
- Mixture Fraction Variance Profile
- Progress Variable



When the Premixed Combustion model is selected, the following sub-nodes appear:

- Progress Variable
- Progress Variable Variance

To define boundary conditions for these variables, follow the same guidelines for setting the initial conditions -- set a value for the corresponding Constant nodes, or edit the scalar array profiles.

Wall Boundaries

The Wall Combustion Scalar Option physics condition selection that you choose for a wall boundary determines if separate sub-nodes for each variable appear under the Physics Values node for that wall boundary.

See What Is the Wall Combustion Scalar Option?

Solving for Species with Transport Equations

The Transported Mass Fraction Profile node appears as an initial condition when the following requirements are met:

- The Progress Variable Model (PVM) is activated.
- The PVM table includes the source term for at least one species (which indicates that the species is solved through a transported equation).

This profile is also available for an individual flow boundary.



Transported Mass Fraction Profile Properties

This value is entered as a scalar array profile.

What Is the Wall Combustion Scalar Option?

See What Is the Wall Combustion Scalar Option for the PPDF Model.

Setting Solver Parameters

The PVM Combustion node, which has its own properties, controls the solution update for the PVM.



PVM Combustion Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.9.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

PVM Formulation

Transport equations for the mixture fraction and its variance are solved for premixed, partially-premixed, and non-premixed combustion models, similar to PPDF Models.

For Partially-Premixed and Non-Premixed Combustion Models:

In addition to solving transport equations for the mixture fraction, the PVM also solves a transport equation for chemical enthalpy h_{298}.

The h_{298} transport equation for the chemical enthalpy can be written:

$$\frac{\partial \rho h_{298}}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} h_{298}) - \nabla \cdot (\Gamma_{h_{298}} \nabla_{h_{298}}) = \omega_{h_{298}} \tag{1144}$$

The source term is taken from the precomputed table.

A progress variable is then defined as:

$$c = \frac{h_{298} - h_{298_{c=0}}}{h_{298_{c=1}} - h_{298_{c=0}}} \tag{1145}$$

Where:

- h_{298} is the chemical enthalpy of the mixture.
- $h_{298_{c=0}}$ is the chemical enthalpies at the initial state.
- $h_{298_{c=1}}$ is the chemical enthalpies at the equilibrium state.

For Premixed Combustion Models:

When using the PVM with the Premixed Combustion model, there are two options available to compute the progress variable variance:

- Algebraic Relationship

  This option computes the progress variable variance using:

$$c_{var} = c_v * \Delta^2 * (\nabla c)^2 \tag{1146}$$

where,

c_v = model constant

Δ = mesh size

- Transport Equation

  This option computes the progress variable variance using a transport equation:

$$\frac{\partial}{\partial t}(\rho\chi c_{var}) + \nabla\bullet\left(\rho U c_{var} - \frac{\mu_t}{Sc_g}\nabla c_{var}\right) = \frac{2\mu_t}{Sc_g}\,tr\,(\nabla c{:}\nabla c) - C_d\rho\frac{\varepsilon}{k}c_{var} \tag{1147}$$

where

$$c_{var} = \overline{(c-\bar{c})^2} \tag{1148}$$

Sc_g = turbulent Schmidt number for the variance

C_d = dissipation constant

PVM Bibliography

[203] J.A. van Oijen, L.P.H. de Geoy. Modelling of premixed laminar flames using flamelet-generated manifolds. *Combust. Sci. Technol.*, 161:113, 2000.

Using the Thickened Flame Model

The Thickened Flame model, represented with a node by that name in the continuum object tree, is for a three-dimensional, transient large-eddy simulation. Applications are mostly for gas-turbine combustors.



This model does not have properties. Instead, it has three child manager nodes that all influence the behavior of the model:

- Efficiency Factor
- Flame Thickening Factor
- Laminar Flame Properties

If you are using imported reaction data, the selection of this model automatically activates the Global Chemistry Kinetics model, which helps control that import.

This model uses a solver when it is working with non-premixed or partially-premixed combustion. The solver is not needed for premixed combustion since its purpose is to calculate mixture fraction. Some of the calculations can be post-processed as field functions.

What Is the Global Chemistry Kinetics Model?

The Global Chemistry Kinetics model is for thickened flame modeling only. It is represented with a node by that name in the continuum object tree. This model's properties control the timing of imported reaction data in the simulation.



This model is activated automatically along with the Thickened Flame model, provided you have the auto-select feature enabled.

Global Chemistry Kinetics Model Properties

Source enabled trigger — Controls whether to enable this model.

- Never — Disables this model.
- Timestep — Calculates starting at the time-step that is specified in the Begin property.

Begin — The solution step at which to begin collecting data. Samples will be collected at time-steps after this start time.

Expert Properties

Source Term Limiting — For the Thickened Flame model, the specified reaction rate can be large enough to drive the reactant concentration below zero. To keep all reactant concentrations in a physically correct range, activate this limit on the final reaction rate that is applied to the species transport equation.
For most applications, leave this option deactivated. Use it only when a reasonable convergence cannot be obtained and the species concentrations clearly show unphysical values.
An example of unphysical values would be large variations of the N2 mass fraction in a premixed flame. In fact, the mass fraction of N2 is expected to be uniform in the domain, since N2 does not participate in the chemical reaction of a premixed flame.

- Activated — Limits the final reaction rate that is applied to the species transport equation.
- Deactivated — Does not limit the final reaction rate (default).

Selecting the Thickened Flame Model

To solve combustion problems using the Thickened Flame model:

- In the *Physics Model Selection* dialog, select a material model from the *Material* group. Typically, this model is Multi-Component Gas.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.

- Specify whether your combustion is premixed, non-premixed or partially premixed by selecting the corresponding radio button in the *Combustion* group. The choice for most simulations with thickened flame modeling is premixed combustion.
- Select Three Dimensional from the *Space* group.
- Select the Implicit Unsteady or Explicit Unsteady option from the *Time* group. The typical choice is Implicit Unsteady.
- Select the Thickened Flame Model option from the particular combustion group, as shown in the following screenshot.



When you make this selection, the Global Chemistry Kinetics model is automatically activated provided the automatic selection of recommended models is enabled.

The next step is to choose the flow solver type. For the Segregated Flow model (the typical choice with the Thickened Flame model):

- Select Segregated Flow in the *Flow* group. The Segregated Species and Segregated Fluid Enthalpy models are automatically selected.

  Since this model is designed for large-eddy simulation, a set of turbulence models is automatically selected as well:

  - Large Eddy Simulation
  - WALE Subgrid Scale
  - All y+ wall treatment

The *Physics Model Selection* dialog then looks as shown in the following screenshot.



Alternatively, to activate the Coupled Flow model:

- Select Coupled Flow in the *Flow* group. The Coupled Energy and Coupled Species models are selected automatically.

The *Physics Model Selection* dialog then looks as shown in the following screenshot.



A further selection should be made from the Equation of State group.

The remaining selections in the dialog depend on problem characteristics that are not related to simulating combustion.

Understanding the Thickened Flame Model Methods

When you select the Thickened Flame model, a set of method nodes become available within the model node. They are designed to let you control how the model is calculated.



The following objects are within the Thickened Flame Model node:

- Efficiency Factor
- Flame Thickening Factor
- Laminar Flame Properties

Setting the Efficiency Factor

The Efficiency Factor manager node lets you choose one of two known methods in its properties: Power Law and Wrinkling Factor Ratio.



Efficiency Factor Properties

Method	Selects the method for the efficiency factor calculation.	
	Power Law	Selects the Power Law method. Activates the Power Law node.
	Wrinkling Factor Ratio	Selects the Wrinkling Factor Ratio method. Activates the Wrinkling Factor Ratio node.

Working with the Power Law

The exponents α (Eqn. (1163)) and b (Eqn. (1164)), as well as the constant C_k (Eqns. (1165) through (1167)), can be set using the Power Law node.



Power Law Properties

Constant alpha	The coefficient α, see Eqn. (1163).
Constant b	The coefficient b, see Eqn. (1164).
Constant Ck	The coefficient C_k, see Eqns. (1165) through (1167).

Working with the Wrinkling Factor Ratio

The coefficient c_{ms} (Eqn. (1160)) can be set using the Wrinkling Factor Ratio node.



Wrinkling Factor Ratio Properties

Constant cms	The coefficient c_{ms}, see Eqn. (1160).

Setting the Flame Thickening Factor

The Flame Thickening Factor manager node lets you choose one of three methods in its properties: Arrhenius, Progress Variable and User-Defined.



Flame Thickening Factor Properties

Method	Selects the method for the flame thickening factor calculation.	
	Arrhenius	Selects the Arrhenius method. Activates the Arrhenius node.
	Progress Variable	Selects the Progress Variable method. Activates the Progress Variable node.
	User Defined	Selects the User-Defined method. Activates the User Defined node.

Working with the Arrhenius Method

For one-step chemistry, set F_{max}, β and Γ using the properties of the Arrhenius node.



Arrhenius Properties

Thickening Factor Maximum	Sets an upper limit (F_{max} in Eqn. (1153)) for the flame thickening factor.
Constant Beta	The coefficient β, see Eqns. (1153), (1159) and (1160).
Constant Gamma	The coefficient Γ, see Eqns. (1152), (1159), (1161), (1163) and (1164).

Working with the Progress Variable

For multi-step chemistry, set F_{max} (Eqn. (1154)) using the properties of the Progress Variable node.



Progress Variable Properties

Thickening Factor Maximum	Sets an upper limit (F_{max} in Eqn. (1154)) for the flame thickening factor.

Working with the User-Defined Method

Set the dynamic thickening factor F (Eqns. (1150), (1151), (1153), (1154), (1157) and (1158)) using the Flame Thickening Factor Profile node, which is a sub-node of the User-Defined node. The Flame Thickening Factor Profile node functions as a scalar profile.



Setting Laminar Flame Properties

The Laminar Flame Properties node has no properties. Rather, it serves as a placeholder for the manager nodes Laminar Flame Speed and Laminar Flame Thickness, which let you control laminar flame parameters.



Working with the Laminar Flame Speed

The Laminar Flame Speed node provides options for controlling flame speed in its properties.



Laminar Flame Speed Properties

Method	Selects the method for laminar flame speed calculation.	
	Gulder Laminar Flame Speed	Selects the Gulder method. Activates the Gulder Laminar Flame Speed node.
	Metghalchi Laminar Flame Speed	Selects the Metghalchi method. Activates the Metghalchi Laminar Flame Speed node.
	User Defined Laminar Flame Speed	Selects the User-Defined method. Activates the User Defined Laminar Flame Speed node.

Working with the Laminar Flame Thickness

The Laminar Flame Thickness node provides options for controlling flame thickness in its properties.



Laminar Flame Thickness Properties

Method	Selects the method for laminar flame thickness calculation.	
	Laminar Flame Speed Relation	Selects the Laminar Flame Speed Relation method. Activates the Laminar Flame Speed Relation node.
	User Defined Laminar Flame Thickness	Selects the User-Defined method. Activates the User Defined Laminar Flame Thickness node.

Using the Laminar Flame Speed Relation

Set the Prandtl number P_r (Eqn. (1169)) along with the coefficients μ_1 and μ_2 (Eqn. (1170)) using the Laminar Flame Speed Relation node.



Laminar Flame Speed Relation Properties

Laminar Prandtl Number	Specifies the value for the Prandtl number P_r in Eqn. (1169).
Constant mu1	The coefficient μ_1 in Eqn. (1170).
Constant mu2	The coefficient μ_2 in Eqn. (1170).

Using the User-Defined Laminar Flame Thickness

Set the laminar flame thickness using the Laminar Flame Thickness Profile node, which is a sub-node of the User Defined Laminar Flame Thickness node. The Laminar Flame Thickness Profile node functions in a manner similar to a scalar profile, but it also includes a table with laminar flame thickness as a function of equivalence ratio.



What Is the Thickened Flame Solver?

The Thickened Flame solver is represented by the TFM Combustion node which has properties. This object appears in the simulation only for non-premixed and partially premixed combustion since its purpose is to calculate mixture fraction.



Thickened Flame Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the old solution is supplanted by the newly computed solution. The default value is `0.9`.

Expert Properties

For information about properties in the *Expert* category, see Expert Properties.

Thickened Flame Formulation

This section provides an overview of how the Thickened Flame model is calculated, along with formulation of the dynamic thickening factor, the efficiency function, and the laminar flame thickness.

A bibliography is also provided.

Overview

In a large-eddy (LES) simulation of turbulent reactive flows, the subgrid-scale filtered balance equation for chemical species has the form:

$$\frac{\partial(\rho)\tilde{Y}_\alpha}{\partial t} + \frac{\partial(\rho)\tilde{Y}_\alpha\,\tilde{u}_i}{\partial x_i} = \frac{\partial}{\partial x_i}\left[\left(D + \frac{\mu_t}{\sigma_k}\right)\frac{\partial \tilde{Y}_\alpha}{\partial x_i}\right] + (\dot{\omega}_\alpha) \tag{1149}$$

In practice, the computational mesh is too coarse to resolve the actual flame, which usually has a size around 0.1 to 1.0 mm. Thus the filtered reaction rate cannot be directly evaluated in terms of resolved quantities.

One approach to overcome this difficulty is known as the Thickened Flame model. The thickness of the computed flame front structure is artificially increased to the size which can be resolved on a LES filter, while the laminar flame speed remains unchanged. This can be achieved by multiplying species and heat diffusion coefficient by a factor F, and decreasing the reaction rate by the same factor F. This is based on the well-established theories of laminar premixed flame that laminar flame speed s_L^0 and the laminar flame thickness δ_L^0 scale as $s_L^0\ \alpha\sqrt{D\dot{\omega}}$ and $\delta_L^0\ \alpha(D/s_L^0)$, where D is the diffusion coefficient and ω the reaction rate. For gas turbines to resolve the flame on a LES grid, F needs to be approximately 10 to 100.

The thickening procedure above reduced the Damkohler number D_α to D_α/F, which makes the flame less sensitive to the small turbulent motions, that is, eddies smaller than $F\delta_L$. This subgrid scale effect can be incorporated in the modeling by using an efficiency factor E_F evaluated from the subgrid scale wrinkling factor. Then the underestimation of the flame front wrinkling by the thickened flame approach can be corrected by increasing the flame speed by the factor E_F.

While the thickening factor F can be constant, the dynamic forms of the thickened flame model have been proposed. There the thickening factor F is dynamically determined based on time and spatial location. Usually it takes large values inside the reaction zone where flame structure is small, while it is set to be unity away from the flame front. A sensor Ω is usually used to detect the presence of the reaction zone.

The resulting species and enthalpy equation is:

$$\frac{\partial(\rho)\tilde{Y}_\alpha}{\partial t} + \frac{\partial(\rho)\tilde{Y}_\alpha\,\tilde{u}_i}{\partial x_i} = \frac{\partial}{\partial x_i}\left[E_F F\left(D + \frac{\mu_t}{\sigma_k}\right)\frac{\partial \tilde{Y}_\alpha}{\partial x_i}\right] + \frac{E_F(\dot{\omega}_\alpha)}{F} \tag{1150}$$

Dynamic Thickening Factor

There are two ways to compute the dynamic thickening factor F:

- Arrhenius expression
- Reaction progress variable expression

Arrhenius Expression

For a simple one–step chemistry scheme:

$$\gamma_F F + \gamma_O O = \gamma_P P \tag{1151}$$

The sensor Ω is formulated as:

$$\Omega = Y_F^{\gamma F} Y_O^{\gamma O} \exp\left(-\Gamma \frac{T_a}{T}\right) \tag{1152}$$

Γ is a modeling parameter $(\Gamma < 1)$.

The thickening factor F takes this form [208]:

$$F = 1 + (F_{max} - 1)\tanh\left(\beta \frac{\Omega}{\Omega_{max}}\right) \tag{1153}$$

Reaction Progress Variable Expression

The Arrhenius expression is only valid for one-step chemistry and is hard to extend to multi-step mechanisms. Durand and Polifke [207] proposed a different expression in which the sensor Ω is related to the reaction progress variable. When c is close to 0 or 1, it is relatively far from the flame front, so the flame thickening process is not needed and the thickening factor F is close to 1, while the thickening factor takes the maximum value F_{max} at the location where $c = 1$. It is computed using the following form:

$$F = 1 + (F_{max} - 1)\Omega(c) \tag{1154}$$

$$\Omega(c) = 16[c(1-c)]^2 \tag{1155}$$

$$c = 1 - Y_f/(Y_f(z)) \tag{1156}$$

Here the progress variable is evaluated based on the fuel mass fraction Y_f in the domain and the unburnt fuel mass fraction $Y_f(z)$ that is solely dependent on the mixture fraction z.

In addition, the effective diffusivity D_{loc} is also represented in terms of the sensor function $\Omega(c)$ as follows:

$$D_{loc} = DEF + (1 - \Omega(c))\frac{\mu_t}{\sigma_k} \tag{1157}$$

and the species transport equation is modified to:

$$\frac{\partial(\rho)\tilde{Y}_\alpha}{\partial t} + \frac{\partial(\rho)\tilde{Y}_\alpha \tilde{u}_i}{\partial x_i} = \frac{\partial}{\partial x_i}\left[E_F F D_{loc}\frac{\partial \tilde{Y}_\alpha}{\partial x_i}\right] + \frac{E_F(\dot{\omega}_\alpha)}{F} \tag{1158}$$

Efficiency Function

There are two ways to compute the efficiency function E_F:

- Wrinkling factor ratio
- Power-law flame wrinkling

Wrinkling Factor Ratio Model

A dimensionless wrinkling factor Ξ is defined as the ratio of the flame surface to its projection in the direction of propagation. Based on DNS data and spectral analysis, Colin *et al.* [206] proposed the following expression for modeling the wrinkling factor:

$$\Xi = 1 + \beta \frac{u'_{\Delta_e}}{s_L^0} \Gamma\left(\frac{\Delta_e}{\delta_L^0}, \frac{u'_{\Delta_e}}{s_L^0}\right) \tag{1159}$$

Here u'_{Δ_e} is the local subgrid scale turbulent velocity, s_L^0 is the laminar flame speed, Δ_e is the local filter size, and δ_L^0 is the laminar flame thickness.

β is computed as:

$$\beta = \frac{2\ln 2}{3c_{ms}(Re_t^{1/2} - 1)}, c_{ms} = 0.28, Re_t = \frac{u' l_t}{\mu} \tag{1160}$$

where Re_t is the turbulent Reynolds number.

The function Γ represents the integration of the effective strain rate that is induced by all scales affected due to artificial thickening, and is estimated as:

$$\Gamma\left(\frac{\Delta_e}{\delta_L^0}, \frac{u'_{\Delta_e}}{s_L^0}\right) = 0.75\exp\left[-1.2\left(\frac{u'_{\Delta_e}}{s_L^0}\right)^{-0.3}\right]\left(\frac{\Delta_e}{\delta_L^0}\right)^{2/3} \tag{1161}$$

Finally, the efficiency function takes the following form as defined by the ratio between the wrinkling factor Ξ, of laminar flame $(\delta_L = \delta_L^0)$ to thickened flame $(\delta_L = \delta_L^l)$:

$$E = \frac{\Xi|_{\delta_L = \delta_L^0}}{\Xi|_{\delta_L = \delta_L^l}} \geq 1 \tag{1162}$$

Power-Law Flame-Wrinkling Model

Charlette *et al.* [205] proposed a formulation for estimating the efficiency function which relates the flame surface area to a cutoff length and limits the wrinkling at the smallest length-scales of the flame. Based on the asymptotic analysis, the efficiency function is evaluated using a power-law expression:

$$E = \left(1 + \min\left[\frac{\Delta_e}{\delta_L^0}, \Gamma\frac{u'_{\Delta_e}}{s_L^0}\right]\right)^{\alpha} \tag{1163}$$

where Γ is defined as:

$$\Gamma\left(\frac{\Delta_e}{\delta_L^0}, \frac{u'_{\Delta_e}}{s_L^0}, Re_\Delta\right) = \left[((f_u^{-a} + f_\Delta^{-a})^{-1/a})^{-b} + f_{Re}^{-b}\right]^{-1/b} \tag{1164}$$

$$f_u = 4\left(\frac{27C_k}{110}\right)\left(\frac{18C_k}{55}\right)\left(\frac{u'_{\Delta_e}}{s_L^0}\right)^2 \tag{1165}$$

$$f_\Delta = \left[\frac{27C_k\pi^{4/3}}{110} \times \left(\left(\frac{\Delta_e}{\delta_L^0}\right)^{4/3} - 1\right)\right]^{1/2} \tag{1166}$$

$$f_{Re} = \left[\frac{9}{55}\exp\left(-\frac{3}{2}C_k\pi^{4/3}Re_\Delta^{-1}\right)\right]^{1/2} \times \; Re_\Delta^{1/2} \tag{1167}$$

The constants a, b, and C_k control the sharpness of the transitions between the asymptotic behaviors. The suggested values are $b = 1.4$, $C_k = 1.5$:

$$a = 0.6 + 0.2\exp[-0.1(u'_{\Delta_e}/s_L^0)] - 0.20\exp[-0.01(\Delta_e/\delta_L^0)], Re_\Delta = 4\frac{\Delta_e}{\delta_L^0}\frac{u'_{\Delta_e}}{s_L^0} \tag{1168}$$

Currently, $\alpha = 0.5$ is used, resulting in the non-dynamic formulation.

Laminar Flame Thickness

The laminar flame thickness is computed using a method that is known as the "Laminar Flame Speed Relation", as follows: [204]:

$$\delta_L^0 = 2\frac{\mu_b}{P_r\,\rho_u s_L^0} \tag{1169}$$

Here P_r is the laminar Prandtl number of the burnt gas, assumed to be constant and the default value is 0.9. μ_b is the molecular viscosity of the burnt gas calculated from Sutherland's law:

$$\mu_b = \frac{\mu_1 T_b^{1.5}}{\mu_2 + T_b} \tag{1170}$$

where μ_1 and μ_2 are coefficients with the default values of 1.457e-6 $\mathrm{kgm}^{-2}K^{-1.5}$ and 110K. T_b is the burnt gas temperature.

Thickened Flame Bibliography

[204] Blint, R. 1986. "The relationship of the laminar flame width to flame speed", *Combust. Sci. and Tech.*, **49**:79.

[205] Charlette, F., Meneveau, C. and Veynante, D. 2002. "Wrinkling Model for LES of Premixed Turbulent Combustion Part I: Non–Dynamic Formulation and Initial Tests", *Combust. Flame*, **131**:159.

[206] Colin, O., Ducros, F., Veynante, D. and Poinsot, T. 2000. "A Thickened Flame Model for Large Eddy Simulation of Turbulent Premixed Combustion", *Phys. Fluids*, **12**:1843.

[207] Durand, L. and Polifke, W. 2007. "Implementation of the Thickened Flame Model for Large Eddy Simulation of Turbulent Premixed Combustion in a Commercial Solver", ASME Paper No. GT2007-28188.

[208] Legier, J.P., Poinsot, T. and Veynante, D. 2000. "Dynamically thickened flame LES model for premixed and non-premixed turbulent combustion", Proceedings of the Summer Program 2000, Center for Turbulence Research, pp. 157-168.

Thickened Flame Field Functions Reference

The following primitive field functions are made available to the simulation when the Thickened Flame model is used. These objects all have the same base properties.

Efficiency Factor	This field function lets you capture eddies that are smaller than the thickened flame. It is expressed as E in Eqns. (1150), (1157), (1158), (1162) and (1163).
Flame Thickening Factor	F in Eqns. (1150), (1151), (1153), (1154), (1157) and (1158).

Using the Nox Emission Model

This section describes:

- What is the Nox Emission model?
- What is the Nox Thermal model?
- What is the Nox Prompt model?
- What is the Nox Flamelet Library model?
- What is the NOX PPDF model?
- What is the Nox Zeldovich model?
- What is the Nox PPDF Equilibrium model?
- What is the Nox PPDF Flamelet model?
- The Nox solver and Nox PPDF Combustion solver
- Selecting the Nox Emission model
- Selecting the Nox Thermal model
- Selecting the Nox Prompt model
- Setting initial and boundary conditions
- Best practices with Nox emissions
- Formulation of the Nox Emission model

Also included is reference information on the following objects that are related to the Nox Emission model:

- Shared boundary and initial condition values
- Field functions

What Is the Nox Emission Model?

The Nox Emission model is one of the optional physics models that are activated by all Partially-Premixed and Non-Premixed combustion models, except for the Homogeneous Reactor. It is represented in the object tree by the Nox Emission node, which has its own properties and is subject to common menu items.



The Nox Emission model provides a framework for modeling the NOx transport equation. This model provides a tool to predict the Nox production from various sources and help in the design of Nox control measures. To calculate a reactive source term that is needed for assembling the NOx transport equation, STAR-CCM+ offers a choice of the Fuel Nox model and/or the Nox Thermal model. The Fuel Nox Model contributes the Nox emission arising from the fuel portion. The Nox Thermal model contributes to this source term by using some variant of the three-step extended Zeldovich reaction mechanism (the exact variant depends on the submodel chosen next).

NOx is considered to be a passive scalar, not influencing density calculations. Therefore, in steady-state simulations, use the Nox Emission model as a post-processing tool, that is, activate it after the flow field has converged. In transient simulations, to model the evolution of NOx accurately, activate this model from the beginning.

Nox Emission Model Properties

Source Enabled Trigger — Controls the combustion sources.

Never	Leaves combustion sources off. This can be useful for steady-state simulations when you want to set up the Nox Emission model but run it only when convergence is achieved.
Iteration	Turns on combustion sources. This option is available in steady-state cases.
Timestep	Turns on combustion sources. This option is available in unsteady-state cases.

Begin — Describes the number of iterations or time-steps when the Nox Emission model becomes enabled.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that STAR-CCM+ uses, we recommend that you do not change any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance.

Secondary Gradients — Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

On	Include both secondary gradients.
Off	Exclude both secondary gradients.
Interior Only	Include the interior secondary gradients only.
Boundaries Only	Include the boundary secondary gradients only.

Convection — Selects the convection scheme to be used. For guidance as to convection scheme selection, see Convection Term, p. 6343.

1st-order	Selects the first-order upwind convection scheme.
2nd-order	Selects the second-order upwind convection scheme.

What Is the Nox Thermal Model?

The Nox Thermal model presents the three-step extended Zeldovich reaction mechanism for computation of the NOx source term. There are four models that the Nox Thermal model may employ to calculate this source term (depending on the flow type):

- Nox Flamelet Library model
- Nox PPDF Equilibrium model
- Nox PPDF Flamelet model
- Nox Zeldovich model

The Nox Thermal Model becomes available for selection when the Nox Emission model is activated. It is represented in the object tree by the Nox Thermal model node, which is subject to common menu items.



What Is the Nox Prompt Model?

Prompt NOx is formed by the reaction of hydrocarbon fragments and molecular nitrogen. The term “prompt” comes from the fact that it is characterized by short-residence time. Prompt NOx can be produced under

certain combustion environments with relatively low-temperature and fuel-rich conditions. Such conditions may exist in staged combustion systems and gas turbines. The significance of predicting and reducing prompt NOx has increased as NOx emission levels are reduced continuously by new combustion technologies and stricter regulations.

The Nox Prompt model becomes available for selection when the Nox Emission model is activated. It is represented in the object tree by the Nox Prompt model node, which has properties and is subject to common menu items.



Nox Prompt Model Properties

Fuel component	Selects a component from a drop-down list of materials added to the continuum.
Pre-exponent correction factor	If you have the kinetics data of a particular hydrocarbon, you can adjust this parameter for it, which is f in Eqn. (1192).
Activation energy, Ea	E_A in Eqn. (1192).

What Is the Nox Flamelet Library Model?

The Nox Flamelet Library model is used for computation of the reactive source term that is needed for assembling the NOx transport equation. It computes this source term from the flamelet-based library. The Nox Flamelet Library model becomes available for selection when the Nox Emission model is activated. It is represented in the object tree by the Nox Flamelet Library node, which has its own properties and is subject to common menu items.



This model contains the Nox Flamelet manager object.

For the Nox Flamelet Library model, the current framework allows you to use a built-in flamelet library or import a user-defined flamelet library for computing the source term for NOx production. A NOx flamelet library is based on specific fuel/oxidizer assumptions. It would be best for the accuracy of the calculations to use a flamelet library that corresponds to the fuel used in the simulation. However, if such a flamelet library is not available, a different flamelet library based on some other fuel composition can be used with reasonable accuracy.

Nox Flamelet Library Model Properties

Nox Flamelet	Selects the Nox flamelet tables available in the simulation from a drop-down list.

What Is the Nox Flamelet Manager?

The Nox Flamelet manager is a collection of available NOx flamelet libraries. The corresponding node is a placeholder for Nox flamelet table objects, and imports those tables with its pop-up menu.



Currently the only available library is based on N-heptane. However, the capability to import a user-defined flamelet library exists. Right-clicking the Nox Flamelet Manager node makes the action Import User-Defined Nox Flamelet Table... available.

NOx flamelet libraries are created by the company LOGE AB (Lund Combustion Engineering), which controls the library format.

Nox Flamelet Manager Pop-Up Menu

Import User-Defined Nox Flamelet Table...	Imports a Nox flamelet table from a separate file that you create.

What Is a Nox Flamelet Table?

Nox Flamelet tables are imported through the Nox Flamelet manager. Each table is represented by a node (HEPTANE in the following screenshot), which has its own properties and is subject to common menu items.



Nox Flamelet Table Properties

The Nox Flamelet Table consists of the integrated source term of NO (Eqn. (1171)) for given temperature T, pressure P, exhaust gas recirculation percentage EGR and scalar dissipation rate χ. Properties of the Nox flamelet table determine the range of values for temperature, pressure, and EGR ratio for which the source term of NO is presented. These properties are read-only, since they are values from the imported data file.

Egr Lower Limit	Lower limit for the exhaust gas ratio factor.
Egr Upper Limit	Upper limit for the exhaust gas ratio factor.
Oxidizer Temperature Lower Limit	Lower limit for the oxidizer temperature.

Oxidizer Temperature Upper Limit	Upper limit for the oxidizer temperature.
Pressure Lower Limit	Lower limit for the pressure.
Pressure Upper Limit	Upper limit for the pressure.

What Is the NOX PPDF Model?

The NOX PPDF model is activated automatically with the Nox Flamelet Library model or the Nox PPDF Equilibrium model in simulations with the EBU combustion model.

This model exposes fluid streams that have to be defined before the simulation starts running. This model is represented in the object tree by the NOX PPDF node, which has its own properties and is subject to common menu items.



NOX PPDF Properties

Dissipation constant — Sets the ratio of velocity and chemical species fluctuations.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the discretization techniques that STAR-CCM+ uses, we recommend that you do not change any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance.

Secondary Gradients — Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.

- On — Include both secondary gradients.
- Off — Exclude both secondary gradients.
- Interior Only — Include the interior secondary gradients only.
- Boundaries Only — Include the boundary secondary gradients only.

Convection — Selects the convection scheme to be used. For guidance as to convection scheme selection, see Convection Term, p. 6343.

- 1st-order — Selects the first-order upwind convection scheme.
- 2nd-order — Selects the second-order upwind convection scheme.

What Is the Nox Zeldovich Model?

The Nox Zeldovich model is used for computation of the instantaneous reactive source term that is needed for assembling the NOx transport equation. It is represented in the object tree by the Nox Zeldovich node, which has its own properties and is subject to common menu items.



Nox Zeldovich Model Properties

When setting the concentrations of species to PPDF Table, the concentrations of species are taken directly from the PPDF Table for each iteration. This option is computationally less expensive than depending on the PPDF model to update the species list at every iteration.

O concentration	Selects the approximation model for the concentration of oxygen.	
	Equilibrium	Selects the approach that is used in Eqns. (1193) through (1195).
	Partial Equilibrium	Selects the approach that is used in Eqns. (1196) and (1197).
	Species List	This option is recommended if the O species is present in the gas-phase reaction mechanism.
	PPDF Table	Obtains the concentration of the O species directly from the PPDF Table for each iteration.

OH concentration	Selects the approximation model for the concentration of the OH group.	
	Exclusion	Selects the approach that is used in Eqn. (1198).
	Partial Equilibrium	Selects the approach that is used in Eqn. (1199).
	Species List	This option is recommended if the OH species is present in the gas-phase reaction mechanism.
	PPDF Table	Obtains the concentration of the OH species directly from the PPDF Table for each iteration.

What Is the Nox PPDF Equilibrium Model?

The Nox PPDF Equilibrium model is used for computation of the averaged reactive source term that is needed for assembling the NOx transport equation. It is represented in the object tree by the Nox PPDF Equilibrium node, which has its own sub-nodes and is subject to common menu items.



What Is a Nox Equilibrium Table?

This table contains the equilibrium coefficients A, B, C, D (as described in the formulation): temperature, enthalpy, and sensible enthalpy. The corresponding node (Nox Source Equilibrium Table in the following screenshot) is activated when the Nox PPDF Equilibrium model is selected. Select this node to access the properties of the Nox Source Equilibrium Table. Right-click this node to access the pop-up menu of the Nox Source Equilibrium Table.



When constructing the Nox source PPDF table, STAR-CCM+ searches to see if there is a suitable PPDF table in the combustion model that has non-integrated data available. If so, then this table is used to compute the coefficients that are needed for the Nox source term calculation. If not, then the standard algorithm (used so far) is employed. This has no impact on the equilibrium calculations: if the suitable table with necessary data is not available, the procedure of constructing the Nox source table takes a bit longer, but the process is the same.

The Nox Source Equilibrium Table node has the following sub-nodes:

- Non PPDF Variables and PPDF Variables -- These sub-nodes allow you to control table size along various dimensions.
- PPDF XY Plots -- This feature enables you to inspect the results of the look-up table calculation graphically.

For defining properties of PPDF fluid streams, you would work with separate objects. The properties of fluid streams cannot be changed as long as there is a PPDF equilibrium table that is not empty. To change the properties, delete the tables first.

STAR-CCM+ does not duplicate large tables when copying continua; see section, Copying Continua with Tables for the recommended procedure.

Nox Source Equilibrium Table Properties

The characteristics of the table to be computed are determined by the following parameter.

Default material database	The database from which the mixture components are imported into the mixture when importing or constructing the table. You can choose between available (open) material databases. This feature is flexible, in that if the mixture component is not present in the selected database, the standard database (if it is not the one selected) is checked.

Nox Source Equilibrium Table Pop-Up Menu

Construct	Creates a table from scratch using the information required. Make sure that all components defined in the file are present in the material database file `props.mdb` or in the default material database (if different from the standard material database).
Export	Saves a table that is stored in the memory of STAR-CCM+ to a file with extension `.tbl`.
Delete	Clears a stored table from memory.

See also: common menu items.

What Is the Nox PPDF Flamelet Model?

The Nox PPDF Flamelet Model is used for computation of the averaged reactive source term that is needed for assembling the NOx transport equation. This model becomes available for selection when the Nox

Thermal Model is activated. It is represented in the object tree by the Nox PPDF Flamelet node, which has its own sub-nodes and is subject to common menu items.



This model is similar to the Nox Flamelet Library model in the sense that it is using the same PPDF flamelet approach for calculating the coefficients that are used in evaluating the thermal Nox source term (A, B, C, and D coefficients, see formulation).

The difference is that the Nox PPDF Flamelet model uses the flamelet table (generated, for example, by the Dars Basic tool) and calculates and tabulates coefficients that are used in calculating the Nox source term.

The Nox Flamelet Library model uses the flamelet library approach in which the coefficients used in evaluating the source term are already given through the flamelet library as a polynomial fit based on mixture fraction, variance, and some other parameters.

What Is a Nox Source Flamelet Table?

This table contains the flamelet coefficients for A, B, C, and D as described in the formulation, temperature, enthalpy, and sensible enthalpy. The corresponding node (Nox Source Flamelet Table in the following screenshot)

is activated when the Nox PPDF Flamelet model is selected. Select this node to access the properties of the Nox Source Flamelet Table. Right-click this node to access the pop-up menu of the Nox Source Flamelet Table.



How Is a Nox Source Flamelet Table Constructed?

When constructing the Nox source PPDF table, STAR-CCM+ searches to see if there is a suitable PPDF table in the combustion model that has non-integrated data available. If so, then this table is used to compute the coefficients that are needed for the Nox source term calculation. If not, then the standard algorithm (used so far) is employed. This has no impact on the equilibrium calculations: if the suitable table with necessary data is not available, the procedure of constructing the Nox source table takes a bit longer, but the process is the same. However, in flamelet, if the suitable table with the necessary data is available, you are not asked to select the folder that contains the flamelet library data. In cases when such a table is not available (or the suitable table exists but it does not contain non-integrated data), you are asked to choose the folder that contains flamelet data.

The characteristics of the table to be computed are determined by parameters in the properties and the following sub-nodes:

- Non PPDF Variables and PPDF Variables -- These sub-nodes allow you to control table size along various dimensions.
- PPDF XY Plots -- This feature enables you to inspect the results of the look-up table calculation graphically.

The properties of fluid streams cannot be changed as long as there is a PPDF flamelet table that is not empty. To change the properties, delete the tables first.

STAR-CCM+ does not duplicate large tables when copying continua; see section, Copying Continua with Tables for the recommended procedure.

Nox Source Flamelet Table Properties

Default material database	The database from which the mixture components are imported into the mixture when importing or constructing the table. You can choose between available (open) material databases. This feature is flexible, in that if the mixture component is not present in the selected database, the standard database (if it is not the one selected) is checked.

Nox Source Flamelet Table Pop-Up Menu

Construct	Creates a table from scratch.
Export	Saves a table that is stored in the memory of STAR-CCM+ to a file with extension `.tbl`.
Delete	Clears a stored table from memory.

What Is the Nox Solver?

The Nox solver is activated by default, after the Nox Emission model has been selected. It is represented in the object tree by the Nox node, which has its own properties and is subject to common menu items.



Nox Solver Properties

Under-Relaxation Factor

At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.8. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

Ticked — Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.

Cleared — Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

Ticked — Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.

Cleared — Do not reset reconstruction gradients. Gradients are recomputed at the next iteration if they had been zeroed.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

Ticked — Retain temporary storage.

Cleared — Do not retain temporary storage.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating.

Ticked — The solver does not update the solution.

Cleared — The solver updates the solution.

What Is the Nox PPDF Combustion Solver?

The Nox PPDF Combustion solver is activated by default whenever the NOX PPDF model is used. It is represented in the object tree by the NOX PPDF Combustion node, which has its own properties and is subject to common menu items.



Nox PPDF Combustion Solver Properties

Under-Relaxation Factor

At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.8. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

Ticked — Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be ticked.

Cleared — Update reconstruction gradients with each iteration.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

Ticked — Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you clear the checkbox again.

Cleared — Do not reset reconstruction gradients. Gradients are recomputed at the next iteration if they had been zeroed.

Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations. Ticked — Retain temporary storage. Cleared — Do not retain temporary storage.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. Ticked — The solver does not update the solution. Cleared — The solver updates the solution.

Selecting the Nox Emission Model

The Nox Emission model is available as an optional model for non-premixed and partially premixed combustion that does not use a homogeneous reactor model. When using a homogeneous reactor model, the detailed NOx mechanism can be included in the reaction definition and thus the Nox Emission model is not needed.

The Nox Emission model solves the transport equation for the concentration of NOx, with the Fuel Nox model and/or the Nox Thermal model contributing to the NOx source.

To solve combustion problems using the Nox Emission model:

- In the *Physics Model Selection* dialog, select option Multi-Component Gas from the *Material* group.
- Specify that your problem involves simulation of chemical reactions by selecting Reacting in the *Reaction Regime* group.
- Select Non-Premixed Combustion or Partially-Premixed Combustion in the *Combustion* group.

- Select the combustion models other than a homogeneous reactor from the particular combustion group, as shown in the following screenshot.



- Select the Nox Emission option from the *Optional Physics Models* group.
- Select one of three options from the *Specific Nox Models* group: Nox Fuel, Nox Prompt and/or Nox Thermal.

  A requirement for the Nox Emission model is that species and enthalpy models be activated, so you need to select a flow model (Coupled or Segregated) from the *Flow* group as well, unless one has already been selected automatically with your particular combustion model. Segregated is the recommended choice when using the Nox Emission model.

The *Physics Model Selection* dialog contains your selections similar to the example shown in the following screenshot. In this example, the option Segregated Flow was selected.



The remaining selections in this dialog depend on problem characteristics that are not related to simulating combustion:

- Select an appropriate model from the Equation of State group, unless your choice of combustion model causes a selection to be made automatically.
- Select the Implicit Unsteady or Steady option from the *Time* group, depending on the nature of your simulation.
- Select a model from the *Viscous Regime* group.
- Select Gravity from the *Optional Physics models* group if gravity forces significantly influence the solution, such as in fire simulations.

Selecting the Nox Thermal Model

The Nox Thermal Model is one of the three specific Nox models offered by STAR-CCM+. It contributes to the thermal part of the Nox source term by employing a specific Nox reaction model. All of these models use a variant of the extended three-step Zeldovich mechanism.

After selecting Nox Thermal Model, the following options appear in the Thermal Nox Models group:

- Nox Flamelet Library
- Nox PPDF Equilibrium
- Nox PPDF Flamelet
- Nox Zeldovich

The first three models use the ppdf approach to account for chemistry-turbulence interactions through modeling mixture fraction and mixture fraction variance space. They can work for laminar/inviscid cases as well - in this case, the mixture fraction variance case is zero all throughout the fluid domain.

These ppdf-based models are using the mixture fraction space distribution to calculate the thermal Nox source term. If the Nox Emission model is not used with PPDF or PCFM models with mixture fraction space available, then the mixture fraction distribution has to be defined and solved based on the species mass fraction distribution. The NOX PPDF model provides a solution for the mixture fraction space. If the automatic model selection is turned on, this model is automatically selected when any of these models are selected with, for example, the Eddy Break-up model.

The Nox Zeldovich model has no chemistry-turbulence interaction model and uses the three-step mechanism as described in the formulation to contribute to the Nox source term.

Selecting the Nox Prompt Model

The Nox Prompt Model is one of the three specific Nox models offered by STAR-CCM+. It is possible to activate this along with the Nox Thermal model. If the conditions are fuel-rich, and you suspect that you have hydrocarbon fragments, then prompt NOx may become important. In cases such as gas turbines and staged combustion, you might want to include the Nox Prompt model to get more accurate NOx predictions.

Setting Initial and Boundary Conditions

A value of special interest to be defined for the Nox Emission model is a NO mass fraction for the fluid domains and boundaries of interest.

If the NOX PPDF model is used to calculate the mixture fraction distribution, then all the initial values pertaining to the PPDF combustion models have to be defined here as well, together with the fuel and oxidizer fluid streams.

Also see the shared boundary and initial condition values reference.

Best Practices with Nox Emissions

This section provides guidance on the following topics:

- Defining streams when using Nox PPDF models with the Ebu combustion model
- Creating flamelet library files in Dars Basic
- Creating a Nox Source equilibrium/flamelet table

Defining Streams When Using Nox PPDF Models with the Ebu Combustion Model

Boundary conditions for mixture fraction are derived from the definition of the fuel and oxidizer stream and boundary values for species (on the inlet and outlet) -- either through perfect mixing assumption or through the ppdf equilibrium table that is internally created. Therefore, defining fuel and oxidizer streams properly is essential to a successful thermal Nox prediction.

Oxidizer composition should correspond to the species composition at the air inlet (if there is one). In most cases, oxidizer composition is air composition (either mixture of O2 and N2, with possibly H2O and CO included). The temperature of the oxidizer stream should also correspond to the temperature of the air inlet.

Fuel stream composition should correspond to the species composition at the fuel inlet (if there is one). Fuel stream can consist of pure fuel (for example, CH4), or of a mixture of air and fuel. The temperature of the fuel stream should also correspond to the temperature of the fuel inlet.

The species composition on all other boundaries should either be a perfect mixture of fuel and oxidizer streams or correspond to a value in the equilibrium table.

Creating Flamelet Library Files in Dars Basic

When creating the flamelet library files in Dars-basic, apart from having a working reaction mechanism, setup of the following solver parameters is essential:

- Fuel temperature (part of defining streams)
- Oxidizer temperature (part of defining streams)
- Pressure (Reference pressure)
- Fuel and oxidizer stream composition
- Number of output points (150 or larger)

Creating a Nox Source Equilibrium/Flamelet Table

In STAR-CCM+, when constructing a Nox source PPDF equilibrium/flamelet table, make sure that there are enough mixture fraction points to capture the possible sharp gradient of values near the stoichiometry mixture fraction.

If possible, create two tables side-by-side and see if the plots for the A, B, C, and D coefficients are resolution-independent.

Nox Emission Formulation

This section describes the formulation of the following options for calculating the reactive source term, which is needed for assembling the NOx transport equation:

- The Nox Flamelet Library model
- The Nox Thermal model
- The Nox Prompt model
- The Nox Zeldovich model
- The Nox PPDF Equilibrium model
- The Nox PPDF Flamelet model

A bibliography is also included.

The Nox Flamelet Library Model

NOx emissions mainly consist of nitrogen oxide (NO) although there can be trace amounts of nitrogen dioxide (NO2) and nitrous oxide (N2O). In this case, we consider NOx emissions to consist only of nitrogen oxide (NO). The NO source term expresses the production of NO and is modeled as a

function of pressure P, temperature T, scalar dissipation rate χ, exhaust gas recirculation percentage EGR and mixture fraction Z. To account for turbulence chemistry interactions, this term is additionally integrated over mixture fraction distribution for the fluid cell:

$$\Omega_{NO} = \int_0^1 S(Z)P(Z)dZ \tag{1171}$$

where Ω_{NO} is the NO source term, $S(Z)$ is the source term for a particular mixture fraction Z and the $P(Z)$ is the probability distribution of the mixture fractions in that cell.

Since there exists a simple analytical expression for the integral of the product of two beta distributions, $S(Z)$ is fitted to a sum of beta distributions.

The reactions that are used for NO are:

$$N_2 + O \leftrightarrow NO + N \tag{1172}$$

$$N + O_2 \leftrightarrow NO + O \tag{1173}$$

$$N + OH \leftrightarrow NO + H \tag{1174}$$

Assuming nitrogen concentrations are in steady state, we get the following expression for $S(Z)$:

$$S(Z) = A - B[NO] + \frac{A\frac{C}{D}}{[NO] + \frac{C}{D}} - [NO]^2 \frac{B}{[NO]\frac{C}{D}} \tag{1175}$$

where:

$$A = [N2][O]K_{1,b}$$

$$B = [O]K_{2,b} + [H]K_{3,b}$$

$$C = [O_2]K_{2,f} + [OH]K_{3,f}$$

$$D = K_{1,b}$$

and K denotes the reaction rate coefficient. The subscripts of K denote the reaction number and the direction (backward or forward).

In the library, these four parameters (A, B, C and D) are fitted to a scaled beta distribution in Z for every T, P, EGR and χ.

The Nox Thermal Model

Formation of nitric oxide during combustion of hydrocarbons is temperature-dependent for the Nox Thermal model. It is produced by the reaction of atmospheric nitrogen with oxygen at elevated temperatures.

Thermal NOx is formed by high temperature oxidation of atmospheric nitrogen. For thermal nitric oxide, the principal reactions are generally recognized to be those proposed by the following three extended Zeldovich mechanisms:

$$N_2 + O \underset{K_{-1}}{\overset{K_1}{\Leftrightarrow}} NO + N \tag{1176}$$

$$N + O_2 \underset{K_{-2}}{\overset{K_2}{\Leftrightarrow}} NO + O \tag{1177}$$

$$N + OH \underset{K_{-3}}{\overset{K_3}{\Leftrightarrow}} NO + H \tag{1178}$$

The rate constants for these reactions have been measured in numerous experimental studies, and the data obtained from these studies has been critically evaluated by Baulch *et al.* [209]. The expressions for rate coefficients of reactions are given below (in units of $\mathrm{m}^3(\mathrm{kmol})^{-1}s^{-1}$):

$$K_1 = 1.8 \times 10^{11} \exp\left(\frac{-38370}{T}\right) \tag{1179}$$

$$K_{-1} = 3.8 \times 10^{10} \exp\left(\frac{-425}{T}\right) \tag{1180}$$

$$K_2 = 1.8 \times 10^{7} T \exp\left(\frac{-4680}{T}\right) \tag{1181}$$

$$K_{-2} = 3.8 \times 10^{6} T \exp\left(\frac{-20820}{T}\right) \tag{1182}$$

$$K_3 = 7.1 \times 10^{10} \exp\left(\frac{-450}{T}\right) \tag{1183}$$

$$K_{-3} = 1.7 \times 10^{11} \exp\left(\frac{-24560}{T}\right) \tag{1184}$$

where K_1, K_2 and K_3 are the forward rate constants and K_{-1}, K_{-2} and K_{-3} are the backward rate constants for Nox reactions, respectively. The rate of formation of NOx is significant only at high temperatures since the thermal fixation of nitrogen requires the breaking of a strong N2 bond. This effect is represented by the high activation energy of the first reaction of the mechanism, which makes this reaction the rate-limiting step of the Zeldovich mechanism.

The net rate of formation of NO via these three reactions is given by:

$$\frac{d[NO]}{dt} = K_1[O][N2] + K_2[N][O2] + K_3[N][OH] - K_{-1}[NO][N] - K_{-2}[NO][O] - K_{-3}[NO][H] \tag{1185}$$

where $[X]$ is molar concentration of X (in units of $\text{kmol m}^{-3}s^{-1}$).

The activation energy for oxidation of N is small, hence a quasi-steady state can be established:

$$\frac{d[N]}{dt} = 0 = K_1[O][N2] - K_2[N][O2] - K_3[N][OH] - K_{-1}[NO][N] + K_{-2}[NO][O] + K_{-3}[NO][H] \tag{1186}$$

$$\Rightarrow [N] = \frac{K_1[O][N2] + K_{-2}[NO][O] + K_{-3}[NO][H]}{K_{-1}[NO] + K_2[O2] + K_3[OH]}$$

Based on this assumption, the instantaneous rate of formation of NOx is (in units of $\text{kmol m}^{-3}s^{-1}$):

$$\frac{d[NO]}{dt} = K_1[O][N2] - K_{-2}[NO][O] - K_{-3}[NO][H] + \frac{K_1[O][N2] + K_{-2}[NO][O] + K_{-3}[NO][H]}{K_{-1}[NO] + K_2[O2] + K_3[OH]}(K_2[O2] + K_3[OH] - K_{-1}[NO]) \tag{1187}$$

It can be simplified further to get the following expression:

$$\frac{d[NO]}{dt} = K_1[O][N2] - K_{-2}[NO][O] - K_{-3}[NO][H] + \frac{K_1[O][N2](K_2[O2] + K_3[OH])}{K_{-1}[NO] + K_2[O2] + K_3[OH]} + \frac{K_{-1}[NO]^2(K_{-2}[O] + K_{-3}[H])}{K_{-1}[NO] + K_2[O2] + K_3[OH]} \tag{1188}$$

Therefore, the instantaneous rate of Nox production can be expressed as:

$$R_{NO} = A - B[NO] + \frac{A\frac{C}{D}}{[NO] + \frac{C}{D}} - [NO]^2 \frac{B}{[NO] + \frac{C}{D}} \tag{1189}$$

where:

$$\begin{aligned} A &= K_1[O][N2] \\ B &= K_{-2}[O] - K_{-3}[H] \\ C &= K_2[O2] + K_3[OH] \\ D &= K_{-1} \end{aligned} \tag{1190}$$

These are the basic equations that are further adjusted depending on the viscous regime that you select.

If both the Nox Thermal and Nox Prompt models are activated, then the net production rate of NOx is the sum of the thermal and prompt NOx.

The Nox Prompt Model

The formation of prompt NOx involves complex reactions with many intermediate species such as CH, CH2, HCN, CN, and others. Due to many uncertainties and a lack of data on the detailed reaction pathways and their kinetics rates, it is necessary to simplify the mechanism to a global rate. For most hydrocarbon fuels, the production rate of prompt NOx is taken from the global kinetics parameters derived by De Soete [210] and represented by the following equations:

$$r_{NO}(\mathbf{kg\ m^{-3}s^{-1}}) = k_r Mw_{NO} \tag{1191}$$

Where Mw_{NO} is the molecular weight of NO and k_r is the rate constant given by:

$$k_r(\mathbf{kgmol\ m^{-3}s^{-1}}) = fA[\mathbf{O}_2]^a[\mathbf{N}_2][\mathbf{Fuel}]\exp\left(\frac{-E_A}{R_u T}\right) \tag{1192}$$

E_A is the activation energy. A is the pre-exponential factor and is given by $6.4e6(R_u T/p)^{a+1}$. f is the correction factor to change the value of the pre-exponential factor for a particular hydrocarbon whose kinetics data is known to you. The value of b varies with the oxygen mole fraction as follows:

$$a = 1 \qquad X_{O2} \leq 4.1 \times 10^{-3}$$
$$a = -3.95 - 0.9 \ln X_{O2} \qquad 4.1 \times 10^{-3} < X_{O2} \leq 1.11 \times 10^{-2}$$
$$a = 0.35 - 0.1 \ln X_{O2} \qquad 1.11 \times 10^{-2} < X_{O2} < 0.03$$
$$a = 0 \qquad X_{O2} \geq 0.03$$

If both the Nox Thermal and Nox Prompt models are activated, then the net production rate of NOx is the sum of the thermal and prompt NOx.

The Nox Zeldovich Model

For laminar/inviscid flow, the instantaneous Nox rate expressed in Eqn. (1189) is used for calculating the Nox source term. Various methods to estimate concentrations of O, OH, and H radicals are described below.

If the concentrations of radicals O, OH and H are calculated by the combustion model, they are used in the NOx model. Otherwise, certain assumptions have to be made to obtain their values. The concentration of radical H is assumed to be zero. There are several ways of modeling O- and OH-concentrations.

Modeling the O-Concentration

Currently, two approaches are being used.

Equilibrium Approach

$$[O] = K_p [O_2]^{1/2} \tag{1193}$$

where $[O]$ is given in units of kmol m^{-3}, and

$$Y_O = [O] \frac{M_O}{\rho} \tag{1194}$$

where:

$$K_p = 1.255 \times 10^4 T^{-1/2} \exp\left(\frac{-31090}{T}\right) \tag{1195}$$

Partial Equilibrium Approach

This approach considers a third-body reaction:

$$O_2 + M \Leftrightarrow O + O + M \tag{1196}$$

which leads to a different expression for K_p:

$$K_p = 1.16 \times T^{1/2} \exp\left(\frac{-27123}{T}\right) \tag{1197}$$

This model generally predicts higher O-concentrations than the equilibrium approach.

Modeling the OH-Concentration

Currently, two approaches are being used.

Exclusion Approach

In this approach, the third reaction in the extended Zeldovich mechanism is assumed to be negligible through the following observation:

$$K_2[O_2]_{eq} >> K_3[OH]_{eq} \tag{1198}$$

Partial Equilibrium Approach

With this approach, the OH-concentration is given as:

$$[OH] = 6.7331 * T^{-0.57} \exp\left(\frac{-4595}{T}\right)[O]^{1/2}[H_2O]^{1/2} \tag{1199}$$

The Nox PPDF Equilibrium Model

To account for turbulence/chemistry interactions, it is necessary to obtain an average value of the Nox production rate. This is done by utilizing the two-stream PPDF equilibrium approach. The only difference is that the values stored in the table are not mass fractions of species but A, B, C, and D coefficients that are needed for calculating the Nox production rate:

$$\begin{aligned} \bar{A} &= A(f) \\ \bar{B} &= B(f) \\ \bar{C} &= C(f) \\ \bar{D} &= D(f) \\ \bar{R}_{NO} &= \bar{A} - \bar{B}[NO] + \frac{\bar{A}\frac{\bar{C}}{\bar{D}}}{[NO] + \frac{\bar{C}}{\bar{D}}} - [NO]^2 \frac{\bar{B}}{[NO] + \frac{\bar{C}}{\bar{D}}} \end{aligned} \tag{1200}$$

where $[NO]$ is the molar concentration of NO at a given time-step/iteration.

If the Nox Emission model is used with the Non-Adiabatic Combustion Model or the EBU Model, then the Nox source table contains values for various heat loss ratios. If it is used with the Adiabatic Combustion Model, only one heat loss ratio point is needed.

The Nox PPDF Flamelet Model

To account for turbulence/chemistry interactions, it is necessary to obtain an average value of the Nox production rate. This is done by utilizing the flamelet approach. The only difference is that the values stored in the table are not mass fractions of species but A,B,C, and D coefficients that are needed for calculating the Nox production rate:

$$\begin{aligned} \bar{A} &= A(f) \\ \bar{B} &= B(f) \\ \bar{C} &= C(f) \\ \bar{D} &= D(f) \\ \bar{R}_{NO} &= \bar{A} - \bar{B}[NO] + \frac{\bar{A}\frac{\bar{C}}{\bar{D}}}{[NO] + \frac{\bar{C}}{\bar{D}}} - [NO]^2 \frac{\bar{B}}{[NO] + \frac{\bar{C}}{\bar{D}}} \end{aligned} \tag{1201}$$

where $[NO]$ is the molar concentration of NO at a given time-step/iteration.

Nox Emission Bibliography

[209] Baulch, D. L., Drysdall, D. D., Horne, D. G., and Lloyd, A. C. 1973. *Evaluated Kinetic Data for High Temperature Reactions*, vols. 1-3. Butterworth.

[210] De Soete, G. G. 1975. "Overall Reaction Rates of NO and N2 Formation from Fuel Nitrogen", in 15th Symp. (Intl.) on Combustion, pp. 1093-1102. The Combustion Institute.

Shared Boundary and Initial Condition Values Reference

This section details boundary and initial condition values that are common to the Nox Emission model:

- Nitrogen Oxide

Nitrogen Oxide

When the Nox Emission model is activated, this object is added among initial conditions, as well as among the boundary values of the free stream and flow boundaries. Also, it can be activated in wall boundaries when the Wall Combustion Scalar Option boundary condition is set to `Specified Value`.

Properties

This value is input as a scalar profile.

Nox Emission Field Functions Reference

The following primitive field function is made available to the simulation when the Nox Emission model is used. This field function has a typical set of base properties.

Nitrogen Oxide	Represents the mass fraction of NO. Since we consider that NOx emissions consist only of NO, this is the transported scalar of the Nox Emission model.

Nox Flamelet Field Function

The following primitive field function is made available to the simulation when the Nox Flamelet Library model is used:

Oxidation Temperature	Current temperature of the oxidizer stream. This temperature is used as an entry for the Nox Flamelet Table when extracting the NO source term.

Nox Thermal Field Functions

The following primitive field functions are made available to the simulation when the Nox Thermal model is used:

Molar Concentration (O)	The molar concentration of oxygen.
Molar Concentration (OH)	The molar concentration of the OH group.

Using the Fuel Nox Model

This section explains the following about the Fuel Nox model:

- What is the Fuel Nox model?
- Selecting a Fuel Nox model
- Formulation

What Is the Fuel Nox Model?

Modeling of fuel NOx is an important issue in the combustion of fuels with some nitrogen content, such as coal, biomass and heavy oil. In such cases, fuel NOx is produced in addition to the thermal NOx. During coal combustion, fuel NOx is formed from the gas-phase oxidation of volatile nitrogen from coal and from the heterogeneous combustion of char-bound nitrogen.

There are also two types of fuels for which fuel Nox modeling is available in Lagrangian phases (after you select the Fuel Nox model):

- Liquid fuel
- Coal

In addition to the above fuels, STAR-CCM+ permits you to model fuel NOx for a purely gaseous fuel. In this scenario, you can introduce the intermediate species, HCN and/or NH3, arising from the fuel directly at the inlet boundary. The Fuel Nox model takes care of the reactions to produce NOx. The applicable calculations are represented by Eqns. (1207) and (1208).

The Fuel Nox model can be chosen individually or with the Nox Prompt and/or the Nox Thermal model.

The Nox Fuel model node appears when you activate the Nox Emission continuum model. It is subject to common menu items.



What Is the Liquid Fuel Nox Model?

For liquid fuel Nox, the source of fuel-bound nitrogen comes from the evaporating liquid fuel. The applications for the Liquid Fuel Nox model are for all liquid fuels with nitrogen content, such as heavy oil.

This model is represented in the Lagrangian phase by the Liquid Fuel Nox node, which has properties and is subject to common menu items. Single- or Multi-Component Liquid, Evaporation, and Two-Way Coupling models are the prerequisites for this model.



This node has a sub-node, Nox Intermediates.

Liquid Fuel Nox Properties

These properties let you control the proportion of nitrogen.

Fuel Vapor Component	Selects the material component for the fuel vapor from among the components of the gas mixture in the simulation.
Mass Fraction of N	Specifies the mass fraction of nitrogen (Y_N in Eqn. (1203)).

What Are the Nox Intermediates?

The Nox Intermediates node lets you specify intermediate fractions of HCN and NH3 separately for Nox.



Nox Intermediates Properties

Conversion Factor, CF	Specifies what fraction of the fuel-bound nitrogen is converted to Fuel Nox (CF in Eqn. (1203)).
Partition Fraction for HCN	Specifies the fraction of the intermediate species HCN (f_{HCN} in Eqn. (1203)).
Partition Fraction for NH3	Specifies the fraction of the intermediate species NH3 (f_{NH3} in Eqn. (1203)).
Partition Fraction for NO	Indicates the fraction of intermediate species NO. This read-only value is determined by the entries for HCN and NH3, as follows: $NO = 1 - HCN - NH3$ (f_{NO} in Eqn. (1205)).

What Is the Coal Fuel Nox Model?

The Coal Fuel Nox model accommodates the fact that, for coal, the source of nitrogen can be either the devolatilized volatile or the char, or a combination of the two. The applications for this model are all coal-fired or biomass combustion processes.

This model is represented in the Lagrangian phase by the Coal Fuel Nox node, which has properties and is subject to common menu items. The Multi-Component Coal and Coal Combustion models are the prerequisites for this model.



This node has two sub-nodes: Volatile Nox Intermediates, and Char Nox Intermediates.

Coal Fuel Nox Properties

These properties let you control the proportions of volatile and char sources of nitrogen.

Mass Fraction of N in Volatile	Specifies the mass fraction of nitrogen in coal volatile ($Y_{N,\,\mathrm{vol}}$ in Eqn. (1216)).
Mass Fraction of N in Char	Specifies the mass fraction of nitrogen in char ($Y_{N,\,\mathrm{char}}$ in Eqn. (1216)).

What Are the Volatile Nox Intermediates?

The Volatile Nox Intermediates node lets you specify intermediate fractions of HCN and NH3 separately for volatile Nox.



Volatile Nox Intermediates Properties

Conversion Factor, CF	Specifies what fraction of the fuel-bound nitrogen in Volatile is converted to Fuel Nox (CF in Eqn. (1216)).
Partition Fraction for HCN	Specifies the fraction of the intermediate species HCN (f_{HCN} in Eqn. (1216)).
Partition Fraction for NH3	Specifies the fraction of the intermediate species NH3 (f_{NH3} in Eqn. (1218))
Partition Fraction for NO	Specifies the fraction of the intermediate species NO. This read-only value is determined by the entries for HCN and NH3, as follows: $NO = 1 - HCN - NH3$ (f_{NO} in Eqn. (1220)).

What Are the Char Nox Intermediates?

The Char Nox Intermediates node lets you specify intermediate fractions of HCN and NH3 separately for char Nox.



Char Nox Intermediates Properties

Conversion Factor, CF	Specifies what fraction of the fuel-bound nitrogen in Volatile is converted to Fuel Nox (CF in Eqn. (1216)).
Partition Fraction for HCN	Specifies the fraction of the intermediate species HCN (f_{HCN} in Eqn. (1216)).
Partition Fraction for NH3	Specifies the fraction of the intermediate species NH3 (f_{NH3} in Eqn. (1218))
Partition Fraction for NO	Specifies the fraction of the intermediate species NO. This read-only value is determined by the entries for HCN and NH3, as follows: $NO = 1 - HCN - NH3$ (f_{NO} in Eqn. (1220)).

Selecting a Fuel Nox Model

When you activate the Nox Emission model, the Nox Fuel, the Nox Prompt,

and Nox Thermal models appear in the *Specific Nox Models* group box.



You can use these three models with each other. However in the dialog, when one is activated the other options move to the *Optional Models* group box.



After you select these continuum models, you will be ready to activate the liquid fuel or coal fuel Nox model in a Lagrangian phase.

Activating the Liquid Fuel Nox Model

The Liquid Fuel Nox model becomes available in the *Optional Phase Models* group box of the Lagrangian *Phase Model Selection* dialog, once you have selected both the Two-Way Coupling and Droplet Evaporation phase models.



Activating the Coal Fuel Nox Model

The Coal Fuel Nox model becomes available in the *Optional Phase Models* group box of the Lagrangian *Phase Model Selection* dialog, once you have selected the Multi-Component Coal and Coal Combustion phase models.



Fuel Nox Formulation

This section provides formulation of the following Fuel Nox models:

- Liquid Fuel
- Coal

Bibliographic information is also included.

Liquid Fuel Nox

In Lagrangian phases, you can use the Fuel Nox model for liquid fuel or coal.

For liquid fuel, the only source of nitrogen is the evaporating fuel. The intermediates formed can be HCN or NH3. The fuel can also form NO through other pathways. There is also a possibility that N2 is formed as one of the intermediate species. The following figure shows the possible routes.



It is necessary to solve the transport equations of HCN, NH3, and NO, and get their respective source terms S_{HCN}, S_{NH3} and S_{NO}:

$$\begin{aligned} S_{HCN} &= S_{HCN,\,\text{produced}} + S_{HCN,\,\text{consumed}} \\ S_{NH3} &= S_{NH3,\,\text{produced}} + S_{NH3,\,\text{consumed}} \\ S_{NO} &= S_{NO,\,\text{produced}} + S_{NO,\,\text{consumed}} \end{aligned} \tag{1202}$$

You need to identify the fuel species (since there can be a multi-component droplet such as $C_xH_yN_z + H_2O$). Only the $C_xH_yN_z$ source is wanted.

$$\begin{aligned} S_{HCN,\,\text{produced}} &= \frac{(S_{\text{fuel}})(Y_N)(CF)(f_{HCN})(M_{W,\,HCN})}{(M_{W,\,N})(V)} \\ S_{NH3,\,\text{produced}} &= \frac{(S_{\text{fuel}})(Y_N)(CF)(f_{NH3})(M_{W,\,NH3})}{(M_{W,\,N})(V)} \end{aligned} \tag{1203}$$

where S_{fuel} is the rate of evaporation of liquid fuel (kg/s) and V is the cell volume.

$S_{NO,\,\text{produced}}$ consists of NO produced from the fuel directly and/or NO produced by the oxidation of HCN/NH3:

$$S_{NO,\,\text{produced}} = S_{NO,\,\text{direct}} + S_{NO,\,\text{oxidation}} \tag{1204}$$

$$S_{NO,\,\text{direct}} = \frac{(S_{\text{fuel}})(Y_N)(CF)(f_{NO})(M_{W,\,NO})}{(M_{W,\,N})(V)} \tag{1205}$$

There are four types of reactions:

$$\begin{aligned} HCN + O2\ (R_{1-HCN})\ \text{leads to}\ NO_{\text{production}} \\ HCN + NO\ (R_{2-HCN})\ \text{leads to}\ NO_{\text{consumption}} \\ NH3 + O2\ (R_{1-NH3})\ \text{leads to}\ NO_{\text{production}} \\ NH3 + NO\ (R_{2-NH3})\ \text{leads to}\ NO_{\text{consumption}} \end{aligned} \tag{1206}$$

The kinetics rates are obtained from De Soete [211].

For intermediate species HCN:

$$\begin{aligned} R_{1-HCN} &= A_{1-HCN}\, X_{HCN}\, X_{O2}^{a}\, e^{-E_{1,\,HCN/(RT)}} \\ R_{2-HCN} &= A_{2-HCN}\, X_{HCN}\, X_{NO}^{a}\, e^{-E_{2,\,HCN/(RT)}} \end{aligned} \tag{1207}$$

where:

$$R_1, R_2 = \text{conversion rates of } HCN(S^{-1})$$

$$X = \text{mole fractions}$$

$$A_{1,\,HCN} = 1.0 \times 10^{10}(S^{-1})$$

$$A_{2,\,HCN} = 3.0 \times 10^{12}(S^{-1})$$

$$E_{1,\,HCN} = 280451.95\ \text{J/gmol}$$

$$E_{2,\,HCN} = 251151\ \text{J/gmol}$$

$$\begin{aligned} a &= 1 & X_{O2} &\le 4.1 \times 10^{-3} \\ a &= -3.95 - 0.9 \ln X_{O2} & 4.1 \times 10^{-3} &< X_{O2} \le 1.11 \times 10^{-2} \\ a &= 0.35 - 0.1 \ln X_{O2} & 1.11 \times 10^{-2} &< X_{O2} < 0.03 \\ a &= 0 & X_{O2} &\ge 0.03 \end{aligned}$$

For intermediate species NH3:

$$R_{1-NH3} = A_{1,\,NH3}\, X_{NH3}\, X_{O2}^{a}\, e^{-E_{1,\,NH3/(RT)}}$$
$$R_{2-NH3} = A_{2,\,NH3}\, X_{NH3}\, X_{NO}^{a}\, e^{-E_{2,\,NH3/(RT)}} \tag{1208}$$

$$R_1, R_2 = \text{conversion rates of } NH3(S^{-1})$$
$$A_{1,\,NH3} = 4.0 \times 10^{6}(S^{-1})$$
$$A_{2,\,NH3} = 1.8 \times 10^{8}(S^{-1})$$
$$E_{1,\,NH3} = 133947.2 \text{ J/gmol}$$
$$E_{2,\,NH3} = 113017.95 \text{ J/gmol}$$

From Eqn. (1202):

$$S_{HCN,\,\text{consumed}} = S_{HCN-1} + S_{HCN-2} = -(R_{1,\,HCN} + R_{2,\,HCN})\rho\frac{M_{W,\,HCN}}{M_{W,\,\text{mix}}} \tag{1209}$$

$$S_{NH3,\,\text{consumed}} = S_{NH3-1} + S_{NH3-2} = -(R_{1,\,NH3} + R_{2,\,NH3})\rho\frac{M_{W,\,NH3}}{M_{W,\,\text{mix}}} \tag{1210}$$

From Eqns. (1202) and (1204):

$$S_{NO,\,\text{oxidation}} = (R_{1,\,HCN} + R_{1,\,NH3})\rho\frac{M_{W,\,NO}}{M_{W,\,\text{mix}}} \tag{1211}$$

$$S_{NO,\,\text{consumed}} = (R_{2,\,HCN} + R_{2,\,NH3})\rho\frac{M_{W,\,NO}}{M_{W,\,\text{mix}}} \tag{1212}$$

So putting together all equations:

$$S_{HCN} = \frac{(S_{\text{fuel}})(Y_N)(CF)(f_{HCN})(M_{W,\,HCN})}{(M_{W,\,N})(V)} - \frac{(R_{1,\,HCN} + R_{2,\,HCN})\rho M_{W,\,HCN}}{M_{W,\,\text{mix}}} \tag{1213}$$

$$S_{NH3} = \frac{(S_{\text{fuel}})(Y_N)(CF)(f_{NH3})(M_{W,\,NH3})}{(M_{W,\,N})(V)} - \frac{(R_{1,\,NH3} + R_{2,\,NH3})\rho M_{W,\,HCN}}{M_{W,\,\text{mix}}} \quad (1214)$$

$$S_{NO} = \frac{(S_{\text{fuel}})(Y_N)(CF)(f_{NO})(M_{W,\,NO})}{(M_{W,\,N})(V)} + \frac{(R_{1,\,HCN} + R_{1,\,NH3} - R_{2,\,HCN} - R_{2,\,NH3})\rho M_{W,\,NO}}{M_{W,\,\text{mix}}} \quad (1215)$$

Coal Nox

Unlike the liquid fuel Nox, coal nitrogen is divided into volatile bound “N” and char bound “N”.



$$S_{HCN,\,\text{produced}} = \frac{(S_{\text{devol}})(Y_{N,\,\text{vol}})(CF)(f_{HCN})(M_{W,\,HCN})}{(M_{W,\,N})(V)} + \frac{(S_{\text{char-oxi}})(Y_{N,\,\text{char}})(CF)(f_{HCN})(M_{W,\,HCN})}{(M_{W,\,N})(V)} \quad (1216)$$

$$S_{HCN,\,\text{consumed}} = -(R_{1,\,HCN} + R_{2,\,HCN})\rho\frac{M_{W,\,HCN}}{M_{W,\,\text{mix}}} \quad \text{(as in Eqn. (1209))} \quad (1217)$$

$$S_{NH3,\,\text{produced}} = \frac{(S_{\text{devol}})(Y_{N,\,\text{vol}})(CF)(f_{NH3})(M_{W,\,HCN})}{(M_{W,\,N})(V)} + \frac{(S_{\text{char-oxi}})(Y_{N,\,\text{char}})(CF)(f_{NH3})(M_{W,\,HCN})}{(M_{W,\,N})(V)} \tag{1218}$$

$$S_{NH3,\,\text{consumed}} = -(R_{1,\,NH3} + R_{2,\,NH3})\rho\frac{M_{W,\,NH3}}{M_{W,\,\text{mix}}} \quad \text{(as in Eqn. (1210))} \tag{1219}$$

Note that the values for CF, f_{HCN}, f_{NH3} and f_{NO} can be specified differently for char and volatile.

$$S_{NO,\,\text{direct}} = \frac{(S_{\text{devol}})(Y_{N,\,\text{vol}})(CF)(f_{NO})(M_{W,\,NO})}{(M_{W,\,N})(V)} + \frac{(S_{\text{char-oxi}})(Y_{N,\,\text{char}})(CF)(f_{NO})(M_{W,\,NO})}{(M_{W,\,N})(V)} \tag{1220}$$

The values for $S_{NO,\,\text{oxidation}}$ and $S_{NO,\,\text{consumed}}$ are the same as in Eqns. (1211) and (1212).

There is some restriction in choosing values for $Y_{N,\,\text{char}}$ and $Y_{N,\,\text{vol}}$: they are not independent of each other. You need ultimate and proximate analysis. Ultimate analysis indicate the percentage of N in coal, that is, the mass fraction of N in daf coal.

So in daf coal:

$$(\text{mass fr. of vol})(Y_{N,\,\text{vol}}) + (\text{mass fr. of char})(Y_{N,\,\text{char}}) = \text{mass fr. of N in daf coal} \tag{1221}$$

The mass fractions of vol and char come from proximate analysis. Always make sure that Eqn. (1221) is satisfied.

Fuel Nox Bibliography

[211] De Soete, G. G. 1975. “Overall Reaction Rates of NO and N2 Formation from Fuel Nitrogen”, in 15th Symp. (Intl.) on Combustion, pp. 1093-1102. The Combustion Institute.

Using Coal Combustion

Coal is an important fuel that is used in electricity generation worldwide. It is pulverized to a fine powder before being blown into the combustion chambers of high-temperature furnaces. STAR-CCM+ provides models to simulate the transport, gasification, and combustion of coal particles as they pass through such equipment. These models are provided within the Lagrangian Multiphase framework of STAR-CCM+, since the particle loading within pulverized coal combustion systems is low.

The Coal Combustion model node, that is located within the Models manager node of the relevant Lagrangian phase, indicates that the coal combustion phase models are available for use.



Coal is a composite material consisting of organic and inorganic substances that may also be moist. In addition to an appropriate combustion model in the continuous phase, the modeling of coal combustion requires the following three submodels (which appear as phase models for the coal Lagrangian phase):

- Coal Moisture Evaporation

- Raw Coal Devolatilization (so that the combustible substances are available for combustion)--one of the following:
  - Two-Step
  - User-Defined
- Heterogeneous Char Oxidation--one of the following:
  - First-Order
  - Half-Order
  - User-Defined

When these Lagrangian phase models are used with a combustion model in the continuous phase, and if the Two-Way Coupling phase model is also active, the complete coal combustion process can be realized. It is also possible to model a non-reacting Lagrangian phase containing the multi-component coal material model only, without activating coal combustion.

Coal Combustion Properties

Depletion Criterion	A user-set limit, expressed as a field function, outside which a parcel is depleted from the simulation.

Understanding the Coal Combustion Phase Models

The following Lagrangian phase models are provided for coal combustion and are selected using the Lagrangian Phase Model dialog:

- Coal Moisture Evaporation model
- Raw Coal Devolatilization models:
  - Two-Step
  - User-Defined
- Char Oxidation models:
  - First-Order
  - Half-Order
  - User-Defined

What Is the Coal Moisture Evaporation Model?

The Coal Moisture Evaporation node represents the Quasi-Steady Evaporation model for the moisture component (H_2O) of the coal particle. It allows the H_2O component of the coal particle to lose its mass through a quasi-steady evaporation process.



The coal particle is assumed to be initially coated with moisture so that a single-component evaporation treatment is appropriate. The driving force for evaporation of this layer is the departure from equilibrium of the liquid-vapor system. In coal combustion, the moisture evaporation is usually heat-transfer limited and takes place rapidly. The devolatilization and char oxidation reactions occur after the moisture has evaporated.

There are no properties that are associated with this model. The formulation is that of the Quasi-Steady Single-Component Droplet Evaporation model as applied to a water droplet, and the associated Sherwood and Nusselt numbers are calculated using the Ranz-Marshall correlation.

What Is the Two-Step Devolatilization Model?

This model, along with User-Defined Devolatilization, becomes available when you select the Coal Combustion model among the Lagrangian phase models.

The Two-Step Devolatilization node represents the Two-Step Devolatilization model that controls the process by which raw coal is converted to gaseous combustible material (volatile matter) and solid char residue. The Volatile Vapor Component property allows you to select the species to which the raw coal devolatilizes. If no species is selected, an error message (Error: no CoalVolatile vapor selected for Devolatilization model in coal) is issued when the solution field is initialized. The benefit of this property is that it allows the modeling of two different coal types with two different types of CoalVolatiles.

This node contains two sub-nodes that provide the properties for each step in the devolatilization process.



Experimental studies have shown that, under certain heating conditions, the volatile matter yield is higher than the value obtained from standard proximate analysis. This effect is accounted for by allowing the coal material to devolatilize via two temperature-dependent competing reactions--the so-called Two-Step Devolatilization Model [212].

Two-Step Devolatilization Properties

Volatile Vapor Component	Select the species to which the raw coal devolatilizes.

First Step Devolatilization Properties

Activation energy, Ea	Term E_{pn} in Eqn. (1227).
Pre-exponential factor, A	Factor A_{pn} in Eqn. (1227).
Mass stoichiometric coefficient, YY	Coefficient YY_{pn} in Eqn. (1225).

Second Step Devolatilization Properties

Activation energy, Ea	Term E_{pn} in Eqn. (1225).
Pre-exponential factor, A	Factor A_{pn} in Eqn. (1227).
Mass stoichiometric coefficient, YY	Coefficient YY_{pn} in Eqn. (1225).

The mass stoichiometric coefficient, YY, for the first step is based on the volatile matter content that is obtained from proximate analysis. If the Second Step Devolatilization properties are set to zero, a one-step process is obtained. If both steps are used, the YY for the second step is based on the known higher yield value of the volatile matter. Typically, the activation energy of the first step is lower than the activation energy of the second step.

What Is the User-Defined Devolatilization Model?

This model, along with Two-Step Devolatilization, becomes available when you select the Coal Combustion model among the Lagrangian phase models.

The User-Defined Devolatilization node represents the User-Defined Devolatilization model. This model supplies the rate of devolatilization for the raw coal component of the coal.



The User-Defined Devolatilization model allows you to define the devolatilization rate with the User Devolatilization Rate object. The Volatile Vapor Component property allows you to select the species to which the raw coal devolatilizes. If no species is selected, an error message (Error: no CoalVolatile vapor selected for Devolatilization model in coal) is issued when the solution field is initialized. The benefit of this property is that it allows the modeling of two different coal types with two different types of CoalVolatiles. In addition, you can set the volatile yield in the properties of this model.

User-Defined Devolatilization Properties

Volatile Vapor Component	Select the species to which the raw coal devolatilizes.
Volatile Yield	Coefficient YY_{pn} in Eqns. (1225) and (1226).

What Is the User Devolatilization Rate?

The User Devolatilization Rate node functions much like a typical profile. It lets you define the devolatilization rate by selecting a constant, a table or a field function from the list in its properties.



The devolatilization rate should be positive, and the units are in kg/s.

User Devolatilization Rate Properties

Method	Selects the method to use for entering the value.	
	Constant	Specifies the value as a single, constant value. A Constant node is added as a child to this node.
	Field Function	Defines the value using a field function (typically user defined). A Function node is added as a child to this node.
	Table (iteration)	Defines the value as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (r)	Defines the value as a function of radius. A Table (r) node is added as a child to this node.
	Table (time)	Defines the value as a function of physical time. A Table (time) node is added as a child to this node.
	Table (x,y,z)	Defines the value as a function of position. A Table (x,y,z) node is added as a child to this node.
	Table (xyz, time)	Defines the value as a function of position and time. A Table (xyz, time) node is added as a child to this node.
Dimensions	The dimensionality of the value (read-only). Clicking the customizer executes the define dimensionality action.	

What Is the First-Order Char Oxidation Model?

This model, along with Half-Order and User-Defined Char Oxidation, becomes available when you select the Coal Combustion model among the Lagrangian phase models.

The First-Order Char Oxidation node represents the First-Order Char Oxidation model that is used to predict the rate at which the char content of the coal burns. The time that is required for this char burnout is a significant part of the coal combustion process.



The First-Order Char Oxidation model consists of one or more First-Order Combined Rate reactions that account for gas phase diffusion of the oxidizer as well as the first-order surface reaction kinetics.

The node for this model contains three reaction nodes, one for each of the possible char oxidation processes. Only the first reaction, Char Reaction With O2, is active by default. The other two reactions can be optionally activated depending on the following:

- How accurate the simulation needs to be
- The availability of reaction rates
- The particular application being modeled

What Is the Half-Order Char Oxidation Model?

This model, along with First-Order and User-Defined Char Oxidation, becomes available when you select the Coal Combustion model among the Lagrangian phase models.

The Half-Order Char Oxidation model consists of one or more Half-Order Combined Rate reactions that account for gas phase diffusion of the oxidizer as well as the half-order surface reaction kinetics. The Half-Order Char Oxidation node contains three reaction nodes, one for each of the possible char oxidation processes.



Only the first reaction, Char Reaction With O2, is active by default. The other two reactions can be optionally activated.

What Is the User-Defined Char Oxidation Model?

This model, along with First-Order and Half-Order Char Oxidation, becomes available when you select the Coal Combustion model among the Lagrangian phase models.

This model is an alternative to the two other char oxidation rates. You can input the rate as a profile or as a constant. The User-Defined Char Oxidation node contains three reaction nodes, one for each of the possible char oxidation processes.



Only the first reaction, Char Reaction With O2, is active by default. The other two reactions can be optionally activated.

What Is the User Reaction Rate?

The User Reaction Rate node functions as a scalar profile value. It lets you define the char oxidation rate by selecting a constant, a table, or a field function from the list in its properties.



The reaction rate should be positive, and the units are in kg/s.

What Are Char Oxidation Reactions?

Three reactions are defined for char oxidation; the following nodes represent these reactions:

- Char Reaction With O2
- Char Reaction With H2O
- Char Reaction With CO2

The reaction nodes, which have properties, are named according to their oxidizing agent. It is not necessary to include all reactions within an analysis; each reaction node has an Active property that is used to indicate whether the reaction is included or not.



The stoichiometric coefficients for the first reaction of Char with O2 can be edited to form either CO and/or CO2. The default oxidation product for Char oxidation by O2 is CO. The stoichiometric coefficients for the other two reactions cannot be edited as their oxidation products are fixed.

Char Oxidation Reaction Properties

Temperature exponent, beta	Temperature exponent n in Eqn. (1234).
Activation energy, Ea	Energy term E_{pl} in Eqn. (1234).
Pre-exponential factor, A	Factor A_{pl} in Eqn. (1234).
Active	Includes or excludes the reaction from the analysis. The reactions corresponding to each of the oxidizers are given by Equations (1230) to (1232). Ticked — The reaction is included. Cleared — The reaction is not included.

Setting Up a Coal Combustion Simulation

Several steps are involved in setting up a coal combustion simulation:

- Create the required gas components in the continuous phase.
- Select an appropriate combustion model for the continuous phase.
- Activate Lagrangian Multiphase modeling and choose appropriate phase models.
- Set properties for the multi-component coal material.
- Set properties on the coal combustion phase models.
- Define injectors for the coal particles.
- Ensure that the coal particles ignite.

Creating Gas Components in the Continuous Phase

A number of components must be created in the multi-component gas mixture in order for the coal combustion model to operate. The mandatory components are:

- O2
- CO2
- CO
- H2

- H2O
- CoalVolatile

It may be necessary to add more components than the above, depending on the reaction mechanisms that are employed in the gas phase.

The CoalVolatile component contains the volatile matter that is transferred from the coal to the gas phase during coal devolatilization. It is of vital importance to set the properties of CoalVolatile in accordance with the coal material type being modeled. The elemental composition of CoalVolatile must be based on the proximate and ultimate analyses. In order to have correct heat release and flame temperature, the calorific value of the coal must be used to determine the heat of formation of CoalVolatile. The heat of formation can be calculated from the difference between the calorific value (heat of reaction) of coal and the heat of formation of the products that are obtained from coal combustion.

Components are created as follows:

- Right-click on the Multi-Component Gas > Gas Mixture > Gas Components node and choose Select Mixture Components.



The *Select Mixture Components* dialog appears.

- Tick the boxes of the desired components and click OK.

Selecting a Combustion Model in the Continuous Phase

Coal combustion simulations require an appropriate combustion model to be present in the continuous phase (that is, the physics continuum). This model is selected in the *Physics Model Selection* dialog for the appropriate continuum, as described in the section, Selecting a Combustion Model.

The following combustion models are appropriate for use with the Lagrangian Coal Combustion phase model:

- EBU Non-Premixed Combustion Models
- EBU Partially-Premixed Combustion Model

Activating Lagrangian Multiphase Modeling and Associated Phase Models

Several steps are required to set up the coal particle phase for use in combustion:

- The Lagrangian Multiphase model must be activated by selecting it from the *Optional Physics Models* group box of the *Physics Model Selection* dialog for the physics continuum. This option is only available when the flow model has been selected from the *Flow* group box.



A Lagrangian Multiphase node is added to the Models manager for the physics continuum.

- At least one Lagrangian phase should be added to the Lagrangian Multiphase > Lagrangian Phases node by right-clicking the Lagrangian Phases node and choosing New. This command adds a phase node, Phase 1, which can be renamed if desired.



Lagrangian phase models should be selected for the new phase as described in the section Selecting Lagrangian Phase Models. For coal combustion, the following models should be included:

- Material Particles from the *Particle Type* group box.
- Multi-Component Coal from the *Material* group box.
- Coal Combustion from the *Mass Transfer* group box.
- Two-Way Coupling from the *Optional Phase Models* group box.

If the Auto-select recommended Phase models option is activated, the following additional phase models are included:

- Energy
- Species
- Coal Moisture Evaporation

Then, make the following selections pertaining to coal combustion:

- From the *Char Oxidation* group box:

  - First-Order
  - Half-Order
  - User-Defined
- From the *Raw Coal Devolatilization* group box:
  - Two-Step
  - User-Defined

An example of a completed dialog is shown below.



Setting Properties of the Multi-Component Coal Material

The default properties of coal components RawCoal, Char, Ash, and H2O should be checked to ensure that they are appropriate for the coal type being modeled. The Multi-Component Coal model node contains these components.

Setting Properties on the Coal Combustion Phase Models

Two types of coal combustion phase models require your further attention: raw coal devolatilization model and char oxidation.

The Two-Step Devolatilization model provides two reaction steps, each with its own properties. The default property values should be checked to ensure that they are appropriate for your scenario.



Alternatively, you can use the User-Defined Devolatilization model, which allows you to define the devolatilization rate with the User Devolatilization Rate object.



The User Devolatilization Rate node functions much like a typical profile.

The First-Order Char Oxidation model provides three possible reactions by which char is oxidized. Only the Char Reaction With O2 reaction is active by default. Each active Char Oxidation reaction needs the specification of temperature exponent, activation energy, and pre-exponential factor.



The Half-Order and User-Defined options for char oxidation use similar objects.

Defining Injectors and the Composition of the Coal Particles

Particle injectors specify the manner in which coal particles enter the flow domain, and the points at which they enter. The top-level Injectors manager node in the simulation tree manages these injectors. A full description of their creation and use is given in the section Working with Injectors.

As the Species model is active for the coal Lagrangian phase, the associated injector contains a Species Mass Fraction node. Use this node to set the initial mass fraction for each of the four components in the coal particles entering through the injector. The initial RawCoal mass fraction is the sum of the volatile matter and char from the proximate analysis.



Obtaining Coal Ignition

To achieve gas phase combustion, the coal particles should first devolatilize to form combustible gaseous fuel. It is necessary to ensure that the local temperature field surrounding the particles provides a sufficiently high temperature for devolatilization to commence. The high temperature need only exist in the initial stages of the simulation as when gas phase combustion begins the process becomes self-sustaining.

An easy way to achieve ignition is to set a high initial temperature (around 800K). If this method does not ignite the mixture successfully, two other methods can be used. One method is to raise the inlet gas temperature (typically to around 800K) so that the devolatilization and gas phase combustion can initiate. When combustion begins, the inlet gas temperature

can be changed back to its original value. The other method is to raise the injection temperature of the coal particles temporarily to initiate the devolatilization process.

Coal Combustion Field Functions

The following field functions are made available to the simulation when the Temporary Storage Retained property in the Lagrangian Multiphase Solver Expert Properties is selected.

Char Production Rate	A simple mass balance results in an expression for the rate of formation of char from particle *p* of the *nth* devolatilization reaction. Eqn. (1233)
CoalVolatile Production Rate	The kinetic rate of production of volatiles from the *nth* devolatilization reaction. Eqn. (1226)
Char Oxidation Rate	For first-order char oxidation model, the overall first-order combined char reaction rate. Eqn. (1235) For half-order char oxidation model, the overall half-order combined char reaction rate for reaction. Eqn. (1237)

Coal Combustion Formulation

This section contains the formulation underlying the models that are used in coal combustion. A coal particle consists of four components: raw coal, char, ash, and moisture. Ash represents the chemically inert part of the particle. Char is the residue that is left when the volatile matter is released from the raw coal, plus that portion of the particle that cannot undergo devolatilization.

Coal Moisture Evaporation

The moisture content of the coal particle is assumed to coat the particle. Moisture evaporation occurs rapidly before other mass transfer processes take place. The continuity equation for the moisture component in particle p is as follows:

$$\frac{d\alpha_{wp}}{dt} = -r_{wp} \tag{1222}$$

The moisture evaporation rate r_{wp} required in Eqn. (1222) is obtained by solving for both heat and mass transfer. The formulation for this evaporation rate is identical to that of the Quasi-Steady Single-Component Droplet Evaporation model. In coal combustion, the heat transfer rate commonly controls the evaporation rate.

Raw Coal Devolatilization

The continuity equation for the raw coal component in particle p is described as:

$$\frac{d\alpha_{cp}}{dt} = -r_{cp} \tag{1223}$$

The net rate of raw coal consumption is given by:

$$r_{cp} = k_{pn}\alpha_{cp} \tag{1224}$$

For the User-Defined model, you specify r_{cp} as the User Devolatilization Rate.

The raw coal, that is, the dry, ash-free portion of the coal, undergoes devolatilization to form volatiles and char by one or more reactions. The nth reaction is represented by:

$$(raw\ coal)_p \xrightarrow{k_{pn}} YY_{pn}(volatile\ matter)_g + (1 - YY_{pn})(char)_p \tag{1225}$$

The volatile matter consists of combustible hydrocarbons and oxygen, and may also contain inorganic species (for example, N, Cl, and S). This is represented by the CoalVolatile component in the gas phase.

The kinetic rate of production of volatiles from the nth reaction (Eqn. (1225)) is given by:

$$r_{vpn} = k_{pn}YY_{pn}\alpha_{cp} \tag{1226}$$

where the reaction rate constant k_{pn} is given by:

$$k_{pn} = A_{pn} \exp\left(-\frac{E_{pn}}{R_u T_p}\right) \tag{1227}$$

The two-step mechanism [212] is recommended for computing the devolatilization rate. Using this mechanism is a reasonable compromise between a simple one-step mechanism, which does not account for the effect of particle history on volatile yield, and a more complex, multi-step mechanism which increases computational time.

Char Oxidation

Raw coal devolatilization produces char according to Eqn. (1225). The continuity equation for the char component in particle p is as follows:

$$\frac{d\alpha_{hp}}{dt} = -r_{hp} \tag{1228}$$

Char reacts heterogeneously, after diffusion of the reactants (that is, O2, CO2 and H2O) to the particle surface, by one or more reactions of the form:

$$(char)_l + (oxidizer)_l \xrightarrow{k_{pl}} (products)_l \tag{1229}$$

Depending on the char oxidation model that is selected, the oxidation reaction is assumed to be first-order or half-order with respect to the oxidizer concentration. Three oxidizers are currently considered for char reactions: O_2, H_2O and CO_2. These correspond to the following reactions:

$$\varphi C + O_2 \rightarrow (2\varphi - 2)CO + (2 - \varphi)CO_2 \tag{1230}$$

$$C + H_2O \rightarrow CO + H_2 \tag{1231}$$

$$C + CO_2 \rightarrow 2CO \tag{1232}$$

First-Order Char Oxidation

A simple mass balance results in an expression for the rate of formation of char from particle p of the nth devolatilization reaction as shown below:

$$r_{hpn} = r_{vpn}(1 - YY_{pn})/YY_{pn} \tag{1233}$$

The net rate of raw coal consumption is the rate of char oxidation plus volatile formation.

Char is assumed to react heterogeneously with gaseous oxidizers that diffuse to the particle surface from the bulk gas phase. Two rate-limiting steps are considered for this process:

- Gas phase diffusion
- Heterogeneous reaction

The surface reaction is assumed to be first-order in the surface concentration of the oxidizer. The Arrhenius expression for the rate constant is then given by:

$$k_{pl} = A_{pl} T_p^n \exp\left(-\frac{E_{pl}}{R_u T_p}\right) \tag{1234}$$

The overall combined char reaction rate is calculated from the following implicit equation [213]:

$$r_{hpl} = \frac{A_p^2 M_{hp} M_g \varphi_l k_{cpl} k_{pl} \zeta_{pl} C_{olg} C_g}{M_g A_p C_g (\zeta_{pl} k_{pl} + k_{cpl}) + r_p} \tag{1235}$$

The mass transfer coefficient without the correction from blowing, k_{cpl}, is computed from the Sherwood number, which in turn is based on the Ranz-Marshall correlation. The value of particle pore area coefficient, ζ_{pl}, has been chosen as 1.0. Since the char oxidation reaction is zero-order with respect to char concentration, heterogeneous ignition is permitted.

Half-Order Char Oxidation

Char is assumed to react heterogeneously with gaseous oxidizers that diffuse to the particle surface from the bulk gas phase. Two rate-limiting steps are considered for this process:

- Gas phase diffusion
- Half-Order Heterogeneous reaction

The surface reaction is assumed to be half-order in the surface concentration of the oxidizer. The Arrhenius expression for the rate constant is then given by:

$$k_{pl} = A_{pl} T_p^n \exp\left(-\frac{E_{pl}}{R_u T_p}\right) \tag{1236}$$

The overall half-order combined char reaction rate for reaction l is calculated from the following implicit equation:

$$r_{hpl} = \phi_l k_{cpl} A_p M_{hp} \left[C_{olg} - \left(\frac{r_{hpl}}{\phi_l k_{pl} A_p M_{hp} \zeta_{pl}} \right)^2 \right] \tag{1237}$$

The mass transfer coefficient without the correction from blowing, k_{cpl}, is computed from the Sherwood number, which in turn is based on the Ranz-Marshall correlation. The value of particle pore area coefficient, ζ_{pl}, has been chosen as 1.0.

Coal Combustion Nomenclature

α	mass fraction of coal component
T_p	particle temperature
YY_{pn}	Mass stoichiometric coefficient for particle devolatilization reaction, n
A_p	particle surface area (m^2)
A_{pl}	Arrhenius pre-exponential factor for char oxidation reaction l
A_{pn}	Arrhenius pre-exponential factor for devolatilization reaction n
r_{vpn}	production rate of volatiles for particle p and nth devolatilization step (kg/s)
r_{hpn}	production rate of char for particle p and nth devolatilization step (kg/s)
r_{wp}	moisture evaporation rate (kg/s)
M_{hp}	char molecular weight (kg/kmol)
M_g	molecular weight of the gas mixture
R_u	universal gas constant [8314.4621 J/kmol K]
φ_l	stoichiometric coefficient for the char reaction
k_{cpl}	mass transfer coefficient
ζ_{pl}	particle pore area coefficient
C_{olg}	concentration of oxidizer ($kmol/m^3$)
C_g	overall gas concentration ($kmol/m^3$)
r_{hpl}	overall combined char reaction rate
r_p	total reaction rate of coal particle p (kg/s)

Subscripts

w	moisture component
c	raw coal component
h	char component
v	volatile matter
n	devolatilization step
l	char oxidation reaction
g	gas phase
p	particle phase

Coal Combustion Bibliography

[212] Kobayashi, H., Howard, J.B., and Sarofim, A.F. 1977. "Coal Devolatilization at High Temperatures", 18th Symposium (International) on Combustion, The Combustion Institute, Pittsburgh, PA, p. 411.

[213] Smoot, D.J., and Smith, P.J. 1985. "Coal Combustion and Gasification", The Plenum Chemical Engineering series, New York.

Using the Soot Emission Models

The formation and emission of carbonaceous particles is often observed during the combustion of hydrocarbons. These particulates, called soot, are identified in flames and fires as yellow luminescence. In gas turbines, internal combustion engines and other practical combustion devices, the formation of soot is mostly a product of incomplete combustion. Apart from the resulting loss in combustion efficiency, a serious effect of soot formation is the health hazard. On the other hand, there are situations where the presence or generation of soot is required. For example, generation of carbon black is needed in the production of automobile tires. In furnaces for industrial application or in heat generators, the intermediate formation of soot is required to augment the heat transfer by radiation. The soot, however, must undergo oxidation before the exhaust is released into the environment.

It is widely accepted that the formation of soot is a complex process which consists of the following:

- Fuel pyrolysis and oxidation reactions
- Formation of polycyclic aromatic hydrocarbons (PAH)
- Inception of first particles
- Growth of soot particles due to reaction with gas-phase species
- Coagulation of particles
- Oxidation of soot particles and intermediates

See also:

What Is the Soot Moments Model?

What Is the Soot Two-Equation Model?

Selecting the Soot Emission Models

To solve combustion problems using the Soot Moments model or the Soot Two-Equation model, in the model selection dialog activate the following models:

- In the *Material* group box, select the Multi-Component Gas model.
- In the *Reaction Regime* group box, select Reacting.

  This selection specifies that you require simulation of chemical reactions.
- In the *Reactive Flow* group box, select Non-Premixed Combustion.

- In the *Non-Premixed Combustion* group box, select the specific combustion models that you require.
- In the *Optional Models* group box, select the Soot Emissions option. Depending on your particular combustion model selection, this may automatically activate the Emissions PPDF model.
- In the *Soot Emissions Model* group box, select either:
  - Soot Moments or,
  - Soot Two-Equation
- If a solver type has not already been selected automatically with your particular combustion model, in the *Flow* group box, select the solver type (Coupled or Segregated) that you require.

The model selection dialog, *Physics 1 Model Selection*, contains your selections similar to the example shown. In this example, the Segregated Flow model is selected.



The remaining options in this dialog depend on problem characteristics that are not related to simulating combustion. For example:

- In the *Time* group box, select the Implicit Unsteady or Steady option depending on the nature of your simulation.
- In the *Viscous Regime* group box, select a suitable model.
- If gravity forces significantly influence the solution, such as in fire simulations, select Gravity from the *Optional Models* group box.

Soot Emissions Model Properties

The following Soot Emissions model properties are available in the *properties* window when you select either the Soot Moments or Soot Two-Equation model nodes.

Source Enabled Trigger	Controls the combustion sources.	
	Never	Leaves combustion sources deactivated. This option can be useful for steady-state simulations when you want to set up the Soot Moments model but run it only when convergence is achieved.
	Timestep	Turns on combustion sources. This option is available in unsteady-state cases.
	Iteration	Turns on combustion sources. This option is available in steady-state cases.
Begin	Specifies the number of iterations or time-steps when the Soot Moments model becomes enabled.	

Soot Emission Models Expert Properties

Avoid changing the properties within the *Expert* category unless you are thoroughly familiar with the theoretical aspects of these models and the discretization techniques that STAR-CCM+ uses. The values in the *Expert* category reflect both the model's design and discretization approaches that are optimized for accuracy and performance. These properties are available in the *Properties* window when you select either the Soot Moments or Soot Two-Equation model nodes.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Includes both secondary gradients.
	Off	Excludes both secondary gradients.
	Interior Only	Includes the interior secondary gradients only.
	Boundaries Only	Includes the boundary secondary gradients only.

Convection	Selects the convection scheme to use. See Convection Term.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.

What Is the Soot Moments Model?

STAR-CCM+ provides a soot model that is based on a technique that is called Method of Moments (see Soot Model Bibliography). This method is used to calculate the soot particle size distribution function (PSDF). The method of moments is based on the fact that solving an infinite set of equations for the statistical moments of the PSDF is equivalent to the direct simulation of the PSDF. Additionally, it can be shown that often the accuracy using only a few moments is sufficiently high. Usually, a set of equations for the first few moments is applied, where the accuracy of the approach increases with the number of moments that are used. The main advantages are:

- Computational efficiency
- The ability to extract major features of the PSDF from the moments, such as:
  - Soot volume fraction
  - Mean number density
  - Total surface
  - Mean diameter

The Soot Moments model in STAR-CCM+ can solve up to four moments: Moment 0, 1, 2, and 3. Soot is treated as a passive scalar. The model is compatible with all non-premixed combustion models.

The Soot Moments node in the continuum object tree represents this type of modeling. To access the properties of the Soot Moments model, select the Soot Moments node.

See Setting Up the Soot Moments Model.



Setting Up the Soot Moments Model

The overall procedure for a case with the Soot Moments model is as follows:

- First construct a flamelet library in DARS-Basic with the appropriate Soot Chemistry Mechanism.
- Select the continuum models that you require.
- Define the soot material properties as needed.
- Set the Soot Moments Model Properties.
- Construct the PPDF Soot Flamelet Table using the flamelet library. Important combustion variables such as mixture fraction are accessed by this flamelet table.
- Once the table is constructed, the Soot Moments model is ready for STAR-CCM+ to use in your simulation.

Important variables in the formulation can be post-processed as field functions.

Setting Parameters of the Soot Moments Model

The Soot Moment Model Properties node allows you to control moment-related properties of the Soot Moments model. Select the Soot Moment Model Properties node to access the properties.



Properties of the Soot Moment Model Properties

Number of Moments	Selects the number of moments to solve. Possible values range from 2 to 4. Solving for 4 moments is more computationally expensive than 2. However, if the problem and mesh size are not too large, and you are interested in soot size dispersion, then run 4 moments.

Expert Properties

Sg Scale	Scales surface growth ($S_{r,\,\mathrm{sr}}$ in Eqn. (1240)).
Pi Scale	Scales condensation ($S_{r,\,\mathrm{cond}}$ in Eqn. (1240)).
Fr Scale	Scales fragmentation ($S_{r,\,\mathrm{fr}}$ in Eqn. (1240)).
Ox Scale	Scales oxidation ($S_{r,\,\mathrm{sr}}$ in Eqn. (1240)).
Condensation	Specifies whether condensation ($S_{r,\,\mathrm{cond}}$ in Eqn. (1240)) is permitted. Activated — Condensation is permitted. Deactivated — Condensation is not permitted.
Cold Temperature Limit	If temperature in the cell is lower than this number, then cold flow is assumed - no soot source calculations are performed in that cell.

What Is the Soot Two-Equation Model?

STAR-CCM+ provides a Soot Two-Equation model that is based on a technique that is called the Moss-Brookes-Hall (MBH) soot model. (See Soot Model Bibliography).

The Soot Two-Equation model is a semi-empirical soot model that is based on four physical processes of soot formation:

- Nucleation / Inception
- Coagulation
- Surface growth
- Oxidation

The Soot Two-Equation model allows you to implement any two-equation soot model by setting the built-in source terms to zero, and then specifying new source terms for each of the four physical processes of soot formation.

The Soot Two-Equation model solves transport equations for scaled number density $\varnothing_N$ and scaled mass density $\varnothing_M$.

The Soot Two-Equation model uses the soot particle number density (N) and the mass density (M) computed from $\varnothing_N$ and $\varnothing_M$ (see equations (1246) and (1247)) to describe the soot formation in confined turbulent jet diffusion flames.

When you activate the Soot Two-Equation model with the Eddy Break-up (EBU) model, the fluid stream manager becomes available. The fluid stream manager lets you define the fuel and oxidizer streams.

In STAR-CCM+ the Soot Two-Equation model offers a choice of nucleation options in its properties: Acetylene-based and PAH-based.

To access the properties of the Soot Two-Equation model, select the Soot Two-Equation node.



To set up a Two-Equation Model, see Setting Up the Soot Two-Equation Model.

Setting Up the Soot Two-Equation Model

The overall procedure for setting up a case with the Soot Two-Equation Emission model is as follows:

- Select the continuum models that you require.
- Select the Soot Two-Equation continuum model.
- Define the following properties as you require:
  - Soot Two-Equation - Properties and Expert Properties
  - Soot Material Properties
  - Soot Species Option
- Once you have set the property values you require, the Soot Two-Equation Emission model is ready for STAR-CCM+ to use in your simulation.

Important variables in the formulation can be post-processed as field functions.

What Are the Two-Equation Soot Model Properties?

Selecting the Soot Two-Equation model > Two-Equation Soot Model Properties node gives you access to more properties in addition to the properties and expert properties of the Soot Two-Equation model.



Two-Equation Soot Model Properties

Nucleation Submodel	Chooses one of two nucleation options.	
	Acetylene-based	This nucleation model pathway (Eqns. (1253) and (1254)) is based on the acetylene concentration. It provides a popular pathway for soot modeling that is in line with the PPDF flamelet combustion model.
	PAH-based	This nucleation model pathway (Eqn. (1252)) is based on polycyclic aromatic hydrocarbons (PAH). This option appears only when the required PAH species are available. PAH nucleation requires more species than acetylene nucleation, and is therefore more computationally expensive due to the large chemical mechanisms involved.

Two-Equation Soot Model Expert Properties

Surface-growth Scale	Default value is 1.0
Coagulation Scale	Default value is 1.0
Nucleation Scale	Default value is 1.0
Oxidation Scale	Default value is 1.0
Cold Temperature Limit	Default value is 250.
Soot User Rates	Deactivated as default

Defining Soot Materials

Define the material properties of the soot separately from the material properties of the gases that are used in combustion. To access the Soot material properties, expand the Soot Moments > Soot > Material Properties

node.



What Is the PPDF Soot Flamelet Table?

This look-up table contains all information about chemical species and reactions that are needed by the Soot Moments model. The Soot Flamelet table provides the source terms in Eqn. (1240) as a function of mixture fraction, its variance, and scalar dissipation rate. The corresponding node (PPDF Soot Flamelet Table in the following screenshot) is activated when the Soot Moments model is selected.

To access the properties of the PPDF Soot Flamelet Table, expand the PPDF Soot Flamelet Table node.

To access the pop-up menu of the PPDF Soot Flamelet Table, right-click the PPDF Soot Flamelet Table node.



The PPDF Soot Flamelet Table node has the following sub-nodes:

- Non PPDF Variables and PPDF Variables -- These sub-nodes allow you to control table size along various dimensions.
- PPDF XY Plots -- This feature enables you to inspect the results of the look-up table calculation graphically.

The properties of fluid streams cannot be changed as long as there is a PPDF flamelet table that is not empty. To change the properties, delete the tables first.

STAR-CCM+ does not duplicate large tables when copying continua; see Copying Continua with Tables for the recommended procedure.

PPDF Soot Flamelet Table Properties

Default material database	The database from which the mixture components are imported into the mixture when importing or constructing the table. You can choose between available (open) material databases. This feature is flexible, in that if the mixture component is not present in the selected database, the standard database (if it is not the one selected) is checked.

PPDF Soot Flamelet Table Pop-Up Menu

Import	Imports a Soot Flamelet Table
Construct	Creates a table from scratch.
Export	Saves a table that is stored in the memory of STAR-CCM+ to a file with extension `.tbl`.
Delete	Clears a stored table from memory.

What is the Soot Species Option?

The Soot Species Option node appears in the simulation tree when the Soot Two-Equation model is selected.

Selecting the Soot Species Option node allows you to control the selection of different methods that are used to compute the O, OH, and C2H2 species in the properties window. These species are required to compute Eqn. (1250), Eqn. (1256), and Eqn. (1257) for the Soot Two-Equation model.



Properties of the Soot Species Option

When setting the concentrations of species to PPDF Table, the concentrations of species are taken directly from the PPDF Table for each iteration. This option is computationally less expensive than depending on the PPDF model to update the species list at every iteration.

C2H2 Concentration

Empirical	This option allows you to compute the concentration of C2H2 empirically from the mixture fraction. This option is only available for the non-premixed flamelet-based method or for cases where C2H2 is not a gas-phase species.
Species List	This option is recommended if the C2H2 species is present in the gas-phase reaction mechanism.
PPDF Table	Obtains the concentration of the C2H2 species directly from the PPDF Table for each iteration.

OH Concentration — Selects the approximation model for the concentration of the OH group.

Exclusion	Selects the approach that is used in Eqn. (1198).
Partial Equilibrium	Selects the approach that is used in Eqn. (1199).
Species List	This option is recommended if the OH species is present in the gas-phase reaction mechanism.
PPDF Table	Obtains the concentration of the OH species directly from the PPDF Table for each iteration.

O Concentration — Selects the approximation model for the concentration of oxygen.

Equilibrium	Selects the approach that is used in Eqns. (1193) through (1195).
Partial Equilibrium	Selects the approach that is used in Eqns. (1196) and (1197).
Species List	This option is recommended if the O species is present in the gas-phase reaction mechanism.
PPDF Table	Obtains the concentration of the O species directly from the PPDF Table for each iteration.

What Is the Soot Solver?

The Soot solver is activated by default when one of the two Soot Emission Models are selected. When either the Soot Moments Model or the Soot Two-Equation Model is selected, the Solvers > Soot node appears in the object tree.

To access the properties of the solver, select the Solvers > Soot node and adjust the properties in the *Properties* window.



Soot Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value is 0.5. For the theoretical background, see Eqn. (2592) in Implicit Iteration, p. 6362.

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. It is not recommend that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in

STAR-CCM+.

Reconstruction Frozen

Sets whether reconstruction gradients are updated with each iteration.

- Activated: Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
- Deactivated: Update reconstruction gradients with each iteration.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

- Activated: Retain temporary storage.
- Deactivated: Do not retain temporary storage.

Reconstruction Zeroed

Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.

- Activated: Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you deactivate the checkbox again.
- Deactivated: Do not reset reconstruction gradients. Gradients are recomputed at the next iteration if they had been zeroed.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating.

- Activated: The solver does not update the solution.
- Deactivated: The solver updates the solution.

The Soot Emission Models and Radiation Models

The soot emission models influence the Participating Media Radiation (DOM) and Gray Thermal Radiation models by contributing to the absorption coefficient of the continuous phase (the absorption coefficient describing both absorption and emission). STAR-CCM+ models the soot as a gray medium, treating the soot as one component of the continuous phase. The total absorption coefficient of the gas includes the soot absorption coefficient as part of an algebraic sum. You can specify the value of the absorption coefficient as a constant, as a field function, or by the Planck Mean Absorption Coefficient method.

Soot Model Formulation

This section provides formulation for the two soot models, and the absorption coefficient that is available in conjunction with participating media radiation.

- Soot moments model formulation
- Soot two-equation model formulation
- Absorption coefficient

Soot Moments Model Formulation

In the Method of Moments formulation, the r-th moment of the soot particle size distribution function (PSDF) is defined as:

$$M_r = \sum_{i=1}^{\infty} i^r N_i \tag{1238}$$

where N_i is the number density of soot particles of size class i. The zero-th moment is related to the mean number density, whereas the first moment is related to the mean volume fraction of the soot particles. The second and third moments are related to the variance and skewness, respectively.

The transport equation for the r-th moment is:

$$\frac{\partial}{\partial t}(\rho M_{mr}) + \frac{\partial}{\partial x_j}\left(\rho u_j M_{mr} - \left(\frac{\mu}{Sc} + \frac{\mu_t}{Sc_t}\right)\frac{\partial M_{mr}}{\partial x_j}\right) = \overline{S}_r \tag{1239}$$

where $M_{mr}(= M_r/\rho)$ is the mass-based moment, Sc_t is the turbulent Schmidt number for Soot (specified by you as the soot property), and $\overline{S}_r$ is the average source term.

Soot Moment Source

The average soot moment source $\overline{S_r}$ is obtained from a Soot flamelet table that is created by a PPDF integration of the instantaneous source term S_r over the mixture fraction and scalar dissipation rate. The instantaneous source terms are available from a Soot flamelet library that is generated through DARS-Basic using an appropriate Soot Mechanism. S_r is given by:

$$S_r = S_{r,\,\mathrm{nu}} + S_{r,\,\mathrm{cond}} + S_{r,\,\mathrm{coag}} + S_{r,\,\mathrm{sr}} + S_{r,\,\mathrm{fr}} \tag{1240}$$

where $S_{r,\,\mathrm{nu}}$, $S_{r,\,\mathrm{cond}}$, $S_{r,\,\mathrm{coag}}$, $S_{r,\,\mathrm{sr}}$, and $S_{r,\,\mathrm{fr}}$ are the instantaneous rates of particle nucleation, condensation, coagulation, heterogeneous surface reactions (that is, surface growth and oxidation) and fragmentation, respectively. Currently, the Soot flamelet table can be constructed only for two streams.

Based on the moments the Soot characteristics are computed as follows:

- Soot Number Density $[1/\mathrm{m}^3]$:

$$N = M_0 N_A \tag{1241}$$

  where M_0 is the zero-th moment.

- Soot Volume Fraction:

$$f_v = M_1 \frac{W_{soot}}{\rho_{soot}} \tag{1242}$$

- Soot Mass Density $[\mathrm{kg}/\mathrm{m}^3]$:

$$M = M_1 W_{soot} \tag{1243}$$

- Soot Mean Diameter:

$$d = \frac{M_{1/3}}{M_0}\left(\frac{6W_{soot}}{\pi\rho_{soot}N_A}\right)^{1/3} \tag{1244}$$

- Soot Surface Density:

$$S = \pi N_A M_{2/3}\left(\frac{6W_{soot}}{\pi\rho_{soot}N_A}\right)^{1/3} \tag{1245}$$

Soot Two-Equation Model Formulation

This section explains the formulations of the Soot Two-Equation model.

The soot two-equation model developed here is based on the modeling framework provided by [222] and [223].

This model is a semi-empirical soot model where the overall soot emission has been modeled with four key physical processes:

- Nucleation
- Coagulation
- Surface-growth
- Oxidation

This model is particularly useful for non-premixed turbulent jet flames.

In this model, transport equations are solved for two variables:

- Scaled soot number density (N) [/m3]

  The soot number density is expressed as:

$$N = \rho N_A \varnothing_N \tag{1246}$$

- Scaled soot mass density (M) [kg/m3]

  The soot mass density is expressed as:

$$M = \rho \varnothing_M \tag{1247}$$

The generic transport equation for scalar $\varnothing$ is of the form:

$$\frac{\partial \rho \varnothing}{\partial t} + \frac{\partial \rho u_j \varnothing}{\partial x_j} = \frac{\partial}{\partial x_j}\left(D \frac{\partial \varnothing}{\partial x_j}\right) + S_\varnothing \tag{1248}$$

For scaled soot number density, $\varnothing = \varnothing_N$

For scaled soot mass density, $\varnothing = \varnothing_M$

Although soot number density and mass density are the primary variables of interest, for numerical convenience scaled variables are preferred for the transport equations.

The soot number density and mass density are Soot Two-Equation Source terms which are expressed as:

$$S_{\varnothing N} = \frac{d\rho \varnothing_N}{dt} = \frac{1}{N_A}\frac{dN}{dt} \tag{1249}$$

$$S_{\varnothing M} = \frac{d\rho \varnothing_M}{dt} = \frac{dM}{dt} \tag{1250}$$

To improve the convergence, the source terms are linearized as:

$$S_{\varnothing} = S_{\varnothing}^{0} + \left(\frac{dS_{\varnothing}}{d\varnothing}\right)^{0} \Delta\varnothing \tag{1251}$$

$\left(\frac{dS_{\varnothing}}{d\varnothing}\right)^{0}$ must be negative, otherwise the convergence rate is compromised. If this value is positive, $\left(\frac{dS_{\varnothing}}{d\varnothing}\right)^{0} = 0$.

Physical processes

The source terms $(S_{\varnothing N}$ and $S_{\varnothing M})$ are split into four physical processes:

- Nucleation

  One nucleation submodel is based on polycyclic aromatic hydrocarbons (PAH), and may be expressed as:

$$\begin{aligned}\left(\frac{dN}{dt}\right)_{Inc.} &= 8 \cdot c_2 \frac{N_A}{M_P}\left[\rho^2\left(\frac{Y_{C_2H_2}}{W_{C_2H_2}}\right)^2 \frac{Y_{C_6H_5}W_{H_2}}{W_{C_6H_5}Y_{H_2}}\right]e^{-\frac{4378}{T}} \\ &+ 8 \cdot c_3 \frac{N_A}{M_P}\left[\rho^2 \frac{Y_{C_2H_2}Y_{C_6H_6}Y_{C_6H_5}W_{H_2}}{W_{C_2H_2}W_{C_6H_6}W_{C_6H_5}Y_{H_2}}\right]e^{-\frac{6390}{T}}\end{aligned} \tag{1252}$$

  See [220].

  Alternatively, an acetylene-based nucleation submodel can be used which requires the soot precursor – C2H2.

  - The inception rate for soot number density is expressed as:

$$\left(\frac{dN}{dt}\right)_{nu} = C_1 N_A \rho \frac{Y_{C2H2}}{W_{C2H2}} e^{\frac{-21000}{T}} \tag{1253}$$

  - The inception rate for mass density is expressed as:

$$\left(\frac{dM}{dt}\right)_{nu} = \frac{M_p}{N_A}\left(\frac{dN}{dt}\right)_{nu} \tag{1254}$$

$$M_p = 1200\frac{kg}{kmol}$$

$$N_A = 6.023\times10^{26}/\text{kmol}$$

$$C_1 = 54/s$$

- Coagulation

$$\left(\frac{dN}{dt}\right)_{co} = -\left(\frac{24R}{\rho_{soot}N_A}\right)^{1/2}\left(\frac{6}{\pi\rho_{soot}}\right)^{1/6} T^{1/2} M^{1/6} N^{11/6} \tag{1255}$$

$$\rho_{soot} = 2000\frac{kg}{m^3}$$

R is the universal gas constant.

$$\left(\frac{dM}{dt}\right)_{co} = 0$$

- Soot surface growth (requires C2H2)

$$\left(\frac{dN}{dt}\right)_{sg} = 0$$

$$\left(\frac{dM}{dt}\right)_{sg} = C_4\rho\frac{Y_{C2H2}}{W_{C2H2}}e^{\frac{-12100}{T}}\left[(\pi N)^{1/3}\left(\frac{6M}{\rho_{soot}}\right)^{2/3}\right] \tag{1256}$$

$$C_4 = 9000.6\frac{kg.\mathrm{m}}{kmol.\mathrm{s}}$$

- Soot oxidation (requires O2 and OH)

$$\left(\frac{dN}{dt}\right)_{ox} = 0$$

$$\left(\frac{dM}{dt}\right)_{ox} = -C_5\rho\eta\frac{Y_{OH}}{W_{OH}}T^{\frac{1}{2}}\left[(\pi N)^{\frac{1}{3}}\left(\frac{6M}{\rho_{soot}}\right)^{\frac{2}{3}}\right] - C_6\rho\frac{Y_{O2}}{W_{O2}}T^{\frac{1}{2}}e^{\frac{-19778}{T}}\left[(\pi N)^{\frac{1}{3}}\left(\frac{6M}{\rho_{soot}}\right)^{\frac{2}{3}}\right] \tag{1257}$$

$$\eta = 0.13$$

$$C_5 = 105.81\frac{kg.\mathrm{m}}{kmol.\mathrm{K}^{1/2}s}$$

$$C_6 = 8903.51\frac{kg.\mathrm{m}}{kmol.\mathrm{K}^{1/2}s}$$

The total source terms for number density and mass are:

$$\frac{dN}{dt} = \left(\frac{dN}{dt}\right)_{nu} + \left(\frac{dN}{dt}\right)_{co} + \left(\frac{dN}{dt}\right)_{sg} + \left(\frac{dN}{dt}\right)_{ox} \tag{1258}$$

$$\frac{dM}{dt} = \left(\frac{dM}{dt}\right)_{nu} + \left(\frac{dM}{dt}\right)_{co} + \left(\frac{dM}{dt}\right)_{sg} + \left(\frac{dM}{dt}\right)_{ox} \tag{1259}$$

Equations (1258) and (1259) are substituted in (1249) and (1250) to obtain the source terms for the transport equations.

The two-equation model developed here allows you to specify your own rates for the four physical processes (Nucleation, Coagulation, Surface growth and Oxidation).

The generalized framework for the computation of source term for $\varnothing = \varnothing_N, \varnothing_M$ is:

$$(S_{\varnothing})_{total} = A_1 + A_2 S_{\varnothing} \tag{1260}$$

A_2 is the scale factor that needs to be specified for all four processes, and A_1 is the user specified rate. If A_1 =0 and A_2 =1, the total source term is computed from equations (1253) – (1257). By setting A_2 =0 you can disable the STAR-CCM+ inbuilt source computation. You can then implement your own source term (through A_1) by specifying the rates for each of the physical processes. When you opt to specify your own rates, Jacobians that you can specify values for are also included as additional options for better numerical stability.

By default, the user-specified rate option is disabled and all the scaling factors are set 1.0.

Species Options

The complete description of the physical processes requires specification of the mass fraction of a few species. These species may or may not be the part of the reaction set up. So choices are provided for you to compute the species concentrations. These options are available for O, OH and C2H2.

See What is the Soot Species Option?

Fluid Streams

Based on the fluid stream definition, an algebraic expression is used to compute the local mixture fraction with Bilger's mixture fraction definition [218]:

$$F = \frac{\dfrac{2(Z_C - Z_{C,O})}{W_C} + \dfrac{(Z_H - Z_{H,O})}{2W_H} - \dfrac{(Z_O - Z_{O,O})}{W_O}}{\dfrac{2(Z_{C,f} - Z_{C,O})}{W_C} + \dfrac{(Z_{H,f} - Z_{H,O})}{2W_H} - \dfrac{(Z_{O,f} - Z_{O,O})}{W_O}} \tag{1261}$$

Z refers to the elemental mass fraction. The first subscript refers to elemental C, H, and O. The second subscript refers to the fuel and oxidizer stream.

Soot Absorption Coefficient

You can set the absorption coefficient as a constant (default value zero), as a field function, or as the Planck Mean Absorption Coefficient, K_a:

$$K_a = \frac{C_0(f_v)T}{0.0037535\text{m-K}} \tag{1262}$$

where:

- C_0 is a constant you can set, with a default value of 4.9.
- f_v is the soot volume fraction from Eqn. (1242).
- T is the temperature.

See [219], [220], and [221].

Soot Model Nomenclature

The following is a guide to the terms and units that are used in this section.

M_r	r-th Moment
N_A	Avogadro Number (6.02214179^{26} 1/kmol)
W_{soot}	Soot Molecular Weight (default value = 24 kg/kmol)
ρ_{soot}	Soot Density (default value = 1800 kg/m^3)
d	$\left(\frac{6M}{\pi N \rho_{soot}}\right)^{\frac{1}{3}}$ Soot Mean Diameter
S	Soot Surface Density (1/m)
K_a	Planck Mean Absorption Coefficient
f_v	$\left(\frac{M}{\rho_{soot}}\right)$ Soot volume fraction
$\varnothing_N$	Soot scaled number density
$\varnothing_M$	Soot scaled mass density
N	Soot number density
M	Soot mass density

Soot Model Bibliography

[214] Kazakov, A., Wang, H. and Frenklach, M. 1995. "Detailed Modeling of Soot Formation in Laminar Premixed Ethylene Flames at a Pressure of 10 Bar", *Comb. & Flame*, 100, pp. 111-120.

[215] Kazakov, A. and Frenklach, M. 1998. "Dynamic Modeling of Soot Particle Coagulation and Aggregation: Implementation With the Method of Moments and Application to High-Pressure Laminar Premixed Flames", *Comb. & Flame*, 114, pp. 484-501.

[216] Mauss, F., Trilken, B., Breitbach, H., and Peters, N. 1994. "Soot Formation in Partially Premixed Diffusion Flames at Atmospheric Pressure", *Soot Formation in Combustion: Mechanisms and Models*, H. Bockhorn ed., Springer Verlag, pp. 325-349.

[217] Balthasar, M., Mauss, F., and Wang, H. 2002. "A computational study of particle thermal ionization and its effect on soot mass growth in laminar premixed flames", *Comb. & Flame*, 129, pp. 204-216.

[218] Bilger, R.W. 1988. “The structure of turbulent nonpremixed flames”, *Proc. Combust. Inst.* **23**, pp. 475–488.
[219] Brookes, S. J. and Moss, J. B. 1999. "Predictions of soot and thermal radiation in confined turbulent jet diffusion flames," *Comb. & Flame*, 116, pp. 486-503.
[220] Wen, Z., Yun, S., Thomson, M. J., and Lightstone, M. F. 2003. "Modeling soot formation in turbulent kerosene/air jet diffusion flames," *Comb. & Flame*, 135, pp. 323-340.
[221] Modest, Michael F., 2003. *Radiative Heat Transfer*, Academic Press.
[222] Wen, Z., Yun, S., Thomson, M. J., and Lightstone, M. F. 2003. "Modeling soot formation in turbulent kerosene/air jet diffusion flames," Comb. & Flame, 135, pp. 323-340.
[223] Brookes, S. J. and Moss, J. B. 1999. "Predictions of soot and thermal radiation in confined turbulent jet diffusion flames," Comb. & Flame, 116, pp. 486-503.

Soot Model Field Functions Reference

The following primitive field functions are made available to the simulation when the Soot Moments model or Soot Two-Equation model is used. These objects all have the same base properties.

Absorption Coefficient	Available when the Participating Radiation model and the Gray Thermal Radiation model are activated.
Density Weighted Soot Moment	M_{mr} in Eqn. (1239) with $r = 0$ and $r = 1$ for moment 0 and 1, respectively.
Soot Coagulation Rate	$\left(\frac{dN}{dt}\right)_{co}$ in Eqn. (1255).
Soot Growth Rate	$\left(\frac{dM}{dt}\right)_{sg}$ in Eqn. (1256).
Soot Moment 0	M_0 in Eqn. (1241).
Soot Moment 1	M_1 in Eqns. (1242) and (1243).
Soot Moment 2	Activated when the Number of Moments property is set to 3 or 4.

Soot Moment 3	Activated when the Number of Moments property is set to 4.
Soot Moment Source 0	$r = 0$ in Eqn. (1240).
Soot Moment Source 1	$r = 1$ in Eqn. (1240).
Soot Moment Source 2	Activated when the Number of Moments property is set to 3 or 4.
Soot Moment Source 3	Activated when the Number of Moments property is set to 4.
Soot Mass	m_s in Eqn. (1243) and (1247).
Soot Mean Diameter	d in Eqn. (1244).
Soot Nucleation Rate For Scaled Mass Density	$\left(\frac{dM}{dt}\right)_{nu}$ in Eqn. (1254).
Soot Nucleation Rate For Scaled Number Density	$\left(\frac{dN}{dt}\right)_{nu}$ in Eqn. (1253).
Soot Number Density	N in Eqn. (1241) and (1246). The units are [/m^3]
Soot Oxidation Rate	$\left(\frac{dM}{dt}\right)_{ox}$ in Eqn. (1257).
Soot Size Dispersion	Activated when the Number of Moments property is set to 4. This value is calculated as follows: $D_{soot} = \frac{(M_2 M_0)}{(M_1)^2}$
Soot Surface Density	S in Eqn. (1245).

Soot Volume Fraction	f_v in Eqn. (1242).
Soot Mass Density	M in Eqn. (1247). The units are $[\mathrm{kg/m^3}]$
Soot Scaled Mass Density	$\varnothing_M$ Obtained from equation (1248). There are no units as this entity is dimensionless.
Soot Scaled Mass Density Rate	Computed from equation(1250). There are no units as this entity is dimensionless.
Soot Scaled Number Density	$\varnothing_N$ Obtained by equation (1248). The units are $\left(\frac{kmol}{kg}\right)$
Soot Scaled Number Density Rate	Computed from equation (1249). The units are $\left(\frac{kmol}{m^3 - s}\right)$
Soot Mean Diameter	$d = \left(\frac{6M}{\pi N \rho_{soot}}\right)^{\frac{1}{3}}$ The units are meters.

Using Lagrangian Particle Reaction

There are many applications that involve the transport of reacting particles, such as calcination and char oxidation. The particle reaction models let you simulate the process of such reacting solid particles through Lagrangian Multiphase models. The particle components can devolatilize or react with a material in a gas phase to form desired products in gas/solid phases.

The calculation of particle reaction requires the Multi-Component Gas model for the continuous phase, along with either of the following submodels:

- Particle Combustion
- Particle Devolatilization

These appear as phase models for the multicomponent solid Lagrangian phase. It is possible to activate both these models.

The Particle Reaction model node, that is located within the Models manager node of the relevant Lagrangian phase, indicates that the particle reaction phase models are available for use. This object has properties and is subject to common menu items.



The use of these models involves a detailed procedure for setting up the simulation. This model is only for use with a multi-component solid.

Particle Reaction Model Properties

Depletion Criterion	A user-set limit, expressed as a field function, outside which a particle in the Lagrangian phase is depleted from the simulation. The default criterion is zero, which deactivates depletion.

Example forms of depletion criteria:

- `$ParticleMass` < *[value]*
- `$ParticleDiameter` < *[value]*
- `$ParcelDEMTimestep` < *[value]*

Understanding the Particle Reaction Phase Models

The following Lagrangian phase models are provided for particle reaction and are selected using the Lagrangian Phase Model dialog:

- Particle Combustion
- Particle Devolatilization

What Is the Particle Combustion Model?

The Particle Combustion model can simulate various particle reactions such as char oxidation, biomass particle oxidation, and ferric oxide reduction. In particle combustion modeling, a solid particle reacts with a gas-phase species to form solid and/or gas-phase products. The Particle Combustion node represents this model. When the Particle Combustion model is chosen,

you can set up multiple reactions (see Eqn. (1268)).



This node contains a Particle Reactions sub-node for setting up particle reactions.

There are no properties that are associated with this model. This model, along with Particle Devolatilization, becomes available when you select the Particle Reaction model among the Lagrangian phase models.

What Is the Particle Devolatilization Model?

Particle devolatilization is a gasification process where a solid particle is converted to volatile matter and/or a residue. The Particle Devolatilization model can simulate various particle mass transfer processes, such as gasification, sublimation, phase transformation, and biomass/coal devolatilization. The Particle Devolatilization node represents this model. When the Particle Devolatilization model is chosen, you can set up

devolatilization reactions (see Eqn. (1263)).



This node contains a Particle Reactions sub-node for setting up particle reactions.

There are no properties that are associated with this model. This model, along with Particle Combustion, becomes available when you select the Particle Reaction model among the Lagrangian phase models.

Setting Up a Particle Reaction Simulation

Several steps are involved in setting up a particle reaction simulation:

- Create the required gas components in the continuous phase.
- Activate Lagrangian Multiphase modeling and choose appropriate phase models.
- Set properties for the multi-component solid material.
- Set up particle reactions -- reactants and products, along with the reaction rate.
- Define injectors for the reacting particles.

Creating Gas Components in the Continuous Phase

Components from the gas phase can be used as both reactants and as products in the particle reactions. Therefore, these components need to be defined in the multi-component gas phase mixture.

Activating Lagrangian Multiphase Modeling and Associated Phase Models

Several steps are required to set up the particle reaction phase for use in combustion:

- Activate the Lagrangian Multiphase model.
- Add at least one Lagrangian phase.

Lagrangian phase models should be selected for the new phase as described in the section Selecting Lagrangian Phase Models. For particle reaction, the following models should be included:

- Material Particles from the *Particle Type* group box.
- Multi-Component Solid from the *Material* group box.
- Particle Reaction from the *Mass Transfer* group box.

If the Auto-select recommended Phase models option is activated, the following additional phase models are included:

- Spherical Particles
- Energy
- Species
- Two-Way Coupling from the *Optional Phase Models* group box.

  Two-Way Coupling is needed because the reaction includes interaction with a gas phase.

Finally, select one or both of the following models from the *Particle Reaction* group box:

- Particle Combustion
- Particle Devolatilization

An example of a completed dialog is shown below.



Setting Properties of the Multi-Component Solid Material

Use the Multi-Component Solid model node to add the necessary solid components for the reacting particle.

For each component that you add, the Molecular Weight material property's method is set to `Elemental Composition` by default. Therefore, the molecular weight depends on the settings within the Elemental

Composition material property node. This default choice of method reflects the fact that elemental composition is vital in checking the stoichiometry for a particle reaction.



Setting Up Particle Reactions

The Particle Combustion node and the Particle Devolatilization node each contain a Particle Reactions sub-node. Rather than using properties, this sub-node has a pop-up menu with which you can create objects for setting up particle reactions. (These objects function in a manner similar to those

used in defining chemical reactions for the continuous phase.) The Particle Reactions sub-node for the Particle Combustion model is shown in the following screenshot.



Particle Reactions Pop-Up Menu

New Particle Reaction	Creates a particle reaction. This new object contains sub-nodes for reactants, products, and the reaction rate.

See also: common menu items.

Use the Particle Reactions node to define particle reactions occurring between the mixture components, that were defined earlier, as follows:

- Right-click the Particle Reactions node and select New Particle Reaction to set up a new (blank) reaction definition.



- A node named Reaction is added. Reactions that are added after the first one have a number at the end of the label indicating the total number of reactions in the Lagrangian phase.



- Define the particle reaction in two parts:
  - The relationship between reactants and products
  - The reaction rate, represented by a manager node by that name
- Repeat the above steps until all reactions required to describe the problem's chemistry fully have been defined.

Setting Up Particle Reactants and Products

Each reaction must include an expression specifying the relationship between reactants and products -- that is, which reaction components occur on the left- or right-hand side of the reaction equation and in what proportions. The participants on each side are defined separately within two sub-nodes that are called Reactants and Products.

Defining a Reactant

A reactant is any reaction component appearing on the left-hand side of a chemical reaction equation, created within the Reactants manager node as shown below. It can be defined with its pop-up menu.

To define a reactant:

- Right-click the Reactants node and select Add Gas Reactant or Add Solid Reactant.
- Choose a component from the submenu. The menu items correspond to components that have been defined as part of the Lagrangian particle in the case of solid reactants, or defined as part of the gas mixture in the case of gas reactants, with the exception that the Particle Devolatilization model does not have any gas-phase reactant.
- Define the required parameters for each component.
- Repeat the above steps until all components on the left-hand side of the equation have been specified.

A rule for particle devolatilization reactions: There can be only one reactant per reaction because there is only one compound per reaction.



Reactants Manager Pop-Up Menu

Add Solid Reactant	Opens a submenu of all materials of the same phase in the material database.
Add Gas Reactant	Opens a submenu of all materials of the same phase in the material database.

See also: common menu items.

Defining a Product

A product is any reaction component appearing on the right-hand side of a particle reaction equation, created within the Products manager node. It can be defined with its pop-up menu.

To define a product:

- Right-click the Products node and select Add Gas Product or Add Solid Product.

- Choose a component from the submenu.

  The menu items correspond to components that have been defined as part of the Lagrangian particle in the case of solid products, or defined as part of the gas mixture in the case of gas products.
- Define the required parameters for each component.
- Repeat the above process until all components on the right-hand side of the equation have been specified.

For particle devolatilization reactions, in contrast to the restriction of only one reactant per reaction, you can add as many products as there are components in the phase, provided the reaction equation is balanced.



Products Manager Pop-Up Menu

Add Solid Product	Opens a submenu of all materials of the same phase in the material database.
Add Gas Product	Opens a submenu of all materials of the same phase in the material database.

See also: common menu items.

Working with Reaction Components

A reaction component is any constituent of a multi-component mixture that has already been defined as a reactant or product in a particle reaction equation, as shown below. The nodes of these components have properties and are subject to common menu items.



To specify the required parameters for a given reaction component:

- Select the component.
- In the *Properties* window, specify a stoichiometric coefficient appropriate to the current reaction.

Reaction Component Properties

Stoich. Coeff.	Sets the stoichiometric coefficient.

Defining Reaction Rates

The rate of the particle reaction is controlled by the Reaction Rate manager node, which provides a choice of methods for calculating the rate in its properties. The example in the following screenshot shows the Particle Combustion model using the First-Order Combined Rate node.



Reaction Rate Properties

Method	Determines the method for calculating the reaction rate.	
	First-Order Combined Rate	Activates the First-Order Combined Rate node (Particle Combustion only).
	Half-Order Combined Rate	Activates the Half-Order Combined Rate node (Particle Combustion only).
	First-Order Rate	Activates the First-Order Rate node (Particle Devolatilization only).
	User Reaction Rate	Activates the User Reaction Rate node.

What Is the First-Order Combined Rate?

The first-order combined rate, available for the Particle Combustion model, calculates the effective particle reaction rate by accounting for both chemical kinetics at the particle surface and diffusion of gas-phase reactant species. The term first-order means that the kinetics reaction rate is first-order with respect to gas-phase reactant species concentration. The First-Order

Combined Rate node uses its properties to control part of this calculation method, while its sub-nodes Pre-exponent and Diffusion Coefficient provide inputs as scalar profiles.



What Is the Half-Order Combined Rate?

The half-order combined rate, available for the Particle Combustion model, calculates the effective particle reaction rate by accounting for both chemical kinetics at the particle surface and diffusion of gas-phase reactant species. The term half-order means that the kinetics reaction rate is half-order with respect to gas-phase reactant species concentration. The Half-Order Combined

Rate node uses its properties to control part of this calculation method, while its sub-nodes Pre-exponent and Diffusion Coefficient provide inputs as scalar profiles.



What Is the First-Order Rate?

The first-order rate, available for the Particle Devolatilization model, calculates the particle reaction rate where the chemical kinetics is first-order with respect to solid-particle reactant concentration. The First-Order Rate node uses its properties to control part of this calculation method, while its sub-node Pre-exponent provides inputs as a scalar profile.



What Is the User Reaction Rate?

The user-reaction rate, available for both Particle Combustion and Particle Devolatilization models, allows you to specify your own reaction rates. The User Reaction Rate node functions as a scalar profile (see Eqn. (1275)).



Reaction Rate Method Properties

These properties are common to the following method nodes:

- First-Order Combined Rate
- Half-Order Combined Rate
- First-Order Rate

Temperature exponent, beta	Sets the temperature exponent β for this reaction (Eqns. (1265), (1270) and (1273)).
Activation energy, Ea	Sets the activation energy E_a for this reaction (Eqns. (1265), (1270) and (1273)).

What Is the Pre-Exponent?

The Pre-exponent node, which is a sub-node of reaction rate nodes other than the User Reaction Rate node, functions as a scalar profile. It lets you use the Arrhenius form (Eqns. (1265), (1270) and (1273)).



What Is the Diffusion Coefficient?

The Diffusion Coefficient node is included in the First-Order Combined Rate and Half-Order Combined Rate nodes (Particle Combustion only). It is needed because this reaction includes a gas oxidizer, and you have to know the rate at which it diffuses. It represents D_m in Eqn. (1272), and functions as a scalar

profile.



Defining Injectors and the Composition of the Particles

The manner in which particles enter the flow domain, and the points at which they enter, are specified by particle injectors. These are managed by the top-level Injectors manager node in the simulation tree. A full description of their creation and use is given in the section Working with Injectors.

As the Species model is active for the particle reaction Lagrangian phase, the associated injector contains a Species Mass Fraction node. Use this node to set the initial mass fraction for each of the components in the particles entering through the injector.



Particle Reaction Formulation

The Particle Reaction model depends on multi-component gas and reacting flow being activated for the continuum, and then calculates particle reaction using one or both particle reaction types:

- Particle Devolatilization
- Particle Combustion

Terms that are used in this formulation are defined in the nomenclature. A bibliography is also included.

Particle Devolatilization

Devolatilization is a gasification process where a solid is converted to volatile matter and a residue. Typically, the process can be represented as:

$$A(s) \rightarrow B(g) + C(s) \tag{1263}$$

Here s and g stand for solid and gas, respectively. When this model is chosen, you can set up devolatilization reactions similar to the one above. There is a restriction on the choice of reactants: only one solid component can be chosen as a reactant in one reaction. The products can be both solid or gas. Either $B(g)$ or $C(s)$ can be optionally chosen. Any number of solid or gas-phase products can be chosen. So, the most general representation of the devolatilization would be:

$$A(s) \rightarrow B1(g) + \ldots + Bn(g) + C1(s) + \ldots + Cn(s) \tag{1264}$$

The particle devolatilization can also be used to model sublimation of a solid component to a gas phase. You should stoichiometrically balance the reaction, otherwise STAR-CCM+ generates an exception.

Once the reaction is set up, you can choose one of two methods for specifying the reaction rate: First Order Rate or User Reaction Rate.

First Order Rate Formulation

In the First Order Rate method, the reaction rate coefficient is of the Arrhenius form given by:

$$k = AT^{\beta}\exp\left(\frac{-E_a}{R_u T}\right) \tag{1265}$$

The rate of consumption of the devolatilizing solid component is given by:

$$\frac{dm_i}{dt} = -k\alpha_i m \qquad \text{kg/s} \tag{1266}$$

User Reaction Rate Formulation

In the User Reaction Rate method, you specify the reaction rate directly:

$$\frac{dm_i}{dt} = -(\text{User Reaction Rate}) \qquad \text{kg/s} \tag{1267}$$

Particle Combustion

Particle Combustion is a process where a solid component reacts with a gas-phase component to form solid and/or gas-phase products.

Typically this process can be represented as follows:

$$A(s) + B(g) \rightarrow C(g) + D(s) \tag{1268}$$

When this model is chosen, you can set up multiple reactions similar to the one above. There is a restriction on the choice of reactants: only one solid component and only one gas component can be chosen as reactants in one reaction. The presence of both is required. The products can be both solid and/or gas. Either $C(g)$ or $D(s)$ can be optionally chosen. Any number of solid or gas-phase products can be chosen. So, the most general representation of the combustion would be:

$$A(s) + B(g) \rightarrow C1(g) + \ldots + Cn(g) + D1(s) + \ldots + Dn(s) \tag{1269}$$

The particle combustion can also be used to model sublimation of a solid component to a gas phase. You should stoichiometrically balance the reaction, otherwise STAR-CCM+ generates an exception.

Once the reaction is set up, you can choose one of three methods for specifying the reaction rate:

- First-Order Combined Rate
- Half-Order Combined Rate
- User Reaction Rate

When the First-Order Combined Rate or Half-Order Combined Rate method is chosen, you also have to specify the Diffusion Coefficient D for the gas-phase reactant species. The Diffusion Coefficient can be specified as a scalar profile.

First-Order Combined Rate Formulation

In the First-Order Rate method, the reaction rate coefficient is of the Arrhenius form given by:

$$k = AT^{\beta}\exp\left(\frac{-E_a}{R_u T}\right) \tag{1270}$$

The rate of consumption of the solid component is dictated by the combined effect of the Arrhenius rate and gas-reactant diffusion rate to the particle surface given by [224]:

$$\frac{dm_i}{dt} = -\frac{k_m k}{k + k_m}\phi C_g M_w A_p \qquad \text{kg/s} \tag{1271}$$

where:

$$k_m = \frac{(Sh)(D_m)}{d} \tag{1272}$$

Half-Order Combined Rate Formulation

In the Half-Order Rate method, the reaction rate coefficient is of the Arrhenius form given by:

$$k = AT^{\beta}\exp\left(\frac{-E_a}{R_u T}\right) \tag{1273}$$

The rate of consumption of the solid component is dictated by the combined effect of the Arrhenius rate and gas-reactant diffusion to the particle surface given by:

$$r = k_m C_g M_w \phi A_p - \frac{k_m r^2}{(M_w \phi A_p) k^2} \qquad \mathrm{kg/s} \tag{1274}$$

User Reaction Rate Formulation

In the User Reaction Rate method, you specify the reaction rate directly:

$$\frac{dm_i}{dt} = -(\text{User Reaction Rate}) \qquad \mathrm{kg/s} \tag{1275}$$

Particle Reaction Nomenclature

The following is a guide to the terms used in this section.

k	reaction rate coefficient (units vary)
α_i	mass fraction of ith component in particle
m_i	mass of ith component in particle
m	mass of particle
ϕ	ratio of stoichiometric coefficients of solid and gas reactants
k_m	mass transfer coefficient (m/s)
D_m	diffusion coefficient $(\mathrm{m^2/s})$
C_g	reactant gas concentration (kmol/kg)
M_w	molecular weight of solid reactant (kg/kmol)
A_p	surface area of the particle $(\mathrm{m^2})$
β	temperature exponent
r	rate of consumption of solid reactant (kg/s)
d	particle diameter (m)
Sh	Sherwood number

Particle Reaction Bibliography

[224] Smoot, D.J. and Smith, P.J. 1985. “Coal Combustion and Gasification”, The Plenum Chemical Engineering series, New York.

Using the Polymerization Model

Polymerization is the process of creating long-chain molecules by chemically combining large numbers of small monomer molecules. This process starts by mixing monomer and initiator molecules in a solvent.

The process of polymerization is comprised of several chemical reactions such as, initiation, propagation, chain transfer, chain branching, scission, and termination, that involve radicals of different chain lengths.

The Polymerization model in STAR-CCM+ simulates free-radical polymerization using Method of Moments.

The following sections provide further information about the Polymerization model:

- Setting Up the Polymerization Model
- Polymerization Model Properties
- What are Free Radical Polymer Reactions?
- User Moment Rates Sub-Nodes
- Polymerization Formulation
- Polymerization Nomenclature
- Polymerization Bibliography

Setting Up the Polymerization Model

The polymerization processes involve various chemical reactions which require some of the combustion physics models in STAR-CCM+.

To access the Polymerization model:

1. Select appropriate physics models for the simulation. These models must include:
   - Multi-component liquid
   - Reacting
   - Non-Premixed Combustion or Partially-Premixed Combustion

- Polymerization



2. Set the Polymerization Model Properties.
3. Specify the Liquid Components and Material Properties.

   See Managing Mixture Components.



4. Set the Boundary Conditions.

   Expand the Physics Values node and set concentrations for the Polymer Initiator, Polymer Monomer, Polymer Radical, Polymer Solvent, and Polymer Modifier using scalar array profiles.



5. Set the Initial Conditions for the Polymer Initiator, Polymer Monomer, Polymer Radical, Polymer Solvent, and Polymer Modifier scalars.

Polymerization Model Properties

The Polymerization model properties are available when you select the [Continuum] > Models > Polymerization node.



The Polymerization model contains a sub-node which allows you to specify parameters for the free-radical polymer reactions that occur during polymerization. See What are Free Radical Polymer Reactions?

Expert Properties

Secondary Gradients	Allows you to activate cross diffusion secondary gradients.	
	On	Activates secondary gradients throughout the solution.
	Off	Secondary gradients are deactivated throughout the solution.
	Interior Only	Activates secondary gradients within interior regions.
	Boundaries Only	Activates secondary gradients at boundaries only.
Convection	Allows you to specify the convection scheme for the reactions.	
	1st-order	Specifies the convection scheme for the free-radical polymer reactions as 1st-order.
	2nd-order	Specifies the convection scheme for the free-radical polymer reactions as 2nd-order.
User Moment Rates	Allows you to add user sources to Moment equations, in addition to the built in sources. Activating this property creates extra Polymerization sub-nodes. See User Moment Rates Sub-Nodes	

Initial Conditions

When the Polymerization model is selected, the Initial Conditions node is automatically populated with a set of sub-nodes. Define the initial conditions according to the specific polymerization reaction that is being simulated.



Define the following polymerization related properties for each initial condition:

Polymer Dead Moment	Define the properties using a scalar array profile. This initial condition is solved using Eqn. (1296), Eqn. (1298), and Eqn. (1300).
Polymer Initiator	Define the properties using a scalar profile. This initial condition is solved using Eqn. (1277).
Polymer Live Moment	Define the properties using a scalar array profile. This initial condition is solved using Eqn. (1290), Eqn. (1292), and Eqn. (1294).
Polymer Modifier	Define the properties using a scalar profile. This initial condition is solved using

Polymer Monomer	Define the properties using a scalar profile. This initial condition is solved using Eqn. (1280).
Polymer Radical	Define the properties using a scalar profile. This initial condition is solved using Eqn. (1278).
Polymer Solvent	Define the properties using a scalar profile. This initial condition is solved using Eqn. (1283).

What are Free Radical Polymer Reactions?

The process of polymerization is comprised of several sub-reactions which include free radical polymer reactions. Free radicals are small, charged molecules that react with monomers and create charged monomers. The charged monomers then combine with other monomers and form a polymer chain. Throughout the polymerization process, free radical polymer reactions consume and create free radicals at different rates. See Eqn. (1277) to Eqn. (1283). In STAR-CCM+ you can specify which free radical polymer reactions occur and define the rates.

The free-radical polymer reaction properties are available when you select the [Continuum] > Models > Polymerization > Free-Radical Polymer Reactions node.



Free-Radical Polymer Reactions Properties

Initiator Efficiency	Allows you to specify the efficiency of the initiator that initiates the chain in Eqn. (1291), Eqn. (1293), and Eqn. (1295).
Chain Transfer to Modifier	Activating this property creates an extra Free-Radical Polymer Reactions sub-node that represents a chain transfer to modifier reaction. You can then edit this node to represent the reaction rate of such a reaction.
Chain Transfer to Solvent	Activating this property creates an extra Free-Radical Polymer Reactions sub-node that represents a chain transfer to solvent reaction (Eqn. (1282)). You can then edit this node to represent the reaction rate of such a reaction.
Chain Transfer to Polymer	Activating this property creates an extra Free-Radical Polymer Reactions sub-node that represents a chain transfer to polymer reaction (Eqn. (1284)). You can then edit this node to represent the reaction rate of such a reaction.
Beta Scission	Activating this property creates an extra Free-Radical Polymer Reactions sub-node that represents a Beta Scission reaction (Eqn. (1286) and Eqn. (1287)). You can then edit this node to represent the reaction rate of such a reaction.
Heat of Propagation	Allows you to specify a value for the heat of propagation which defines the heat that is released due to propagation reactions.
Heat of Initiation	Allows you to specify a value for the heat of initiation which defines the heat that is released due to initiation reactions.

Free-Radical Polymer Reactions Sub-Nodes

The Free-Radical Polymer Reactions node contains a set of sub-nodes that you can use to specify the reaction rates of the individual free-radical polymer

reactions that are involved in polymerization.



To define the reaction rate for a free-radical polymer reaction:

- Select the required free-radical polymer reaction sub-node and set the Method that you require as a scalar profile.
- Define the Unit System property appropriate to the reaction.

If you do not wish to specify the reaction rate for any of the free-radical polymer reactions, select the relevant sub-node and:

- Set the Method property for the sub-node to Constant.
- Select the [sub-node] > Constant node.
- Set Value to `0.0`.

Free-Radical Polymer Reaction Sub-Node Properties

Method	Allows you to specify the Method as a scalar profile to define for the free-radical polymer reaction.
Unit System	Allows you to define units that are appropriate to the reaction. kmol, m, s, K mol, cm, s, K

See also:

What are Free Radical Polymer Reactions?

User Moment Rates Sub-Nodes

The User Moment Rates sub-nodes allow you to specify source terms directly to the Moment equations.

To access the User Moment Rates sub-nodes:

- Activate the Polymerization property, User Moment Rates.

  The following sub-nodes appear under the Polymerization model node:

  - Live Polymer Moment User Rate
  - Live Polymer Moment User Rate Jacobian
  - Dead Polymer Moment User Rate
  - Dead Polymer Moment User Rate Jacobian
- Set the properties of these sub-nodes as scalar array profiles.

Polymerization Formulation

This feature is based on the Method of Moments Model for polymerization.

Primary properties like chain length distribution of polymer molecules are obtained by classical method of moments. STAR-CCM+ solves for the scalar transport of three Live polymers, λ_0, λ_1, and λ_2, and three Dead polymers, μ_0, μ_1, and μ_2. In addition to the source terms described below, STAR-CCM+ allows you to add moment sources that you define. STAR-CCM+ also solves for the scalar transport of the Initiator, Monomer, Radical, Solvent, and Modifier.

The source terms of the moment transport equations depend on the sub processes of polymerization:

- Chain Initiation

$$R^{\cdot} + M \xrightarrow{k_I} R_1 \tag{1276}$$

  where,

  the initiator consumption rate is found using:

$$r_{Ii} = -k_{di}[I_i],\ i = 1 \tag{1277}$$

  and the primary radical formation rate is found using:

$$r_{R^{\cdot}} = \sum_{i=1}^{N_d} 2 f_i k_{di}[I_i] - k_I[R^{\cdot}][M] \tag{1278}$$

where,

f_i = initiator efficiency

- Propagation

$$R_x + M \overset{k_p}{\rightarrow} R_{x+1} \tag{1279}$$

where the monomer consumption rate is found using:

$$r_M = -k_p[M]\lambda_0 \tag{1280}$$

- Chain transfer to monomer

$$R_x + M \overset{k_{tm}}{\rightarrow} P_x + R_1 \tag{1281}$$

- Chain transfer to solvent

$$R_x + S \overset{k_{ts}}{\rightarrow} P_x + R_1 \tag{1282}$$

where the solvent consumption rate is found using:

$$r_S = -k_{ts}[S]\lambda_0 \tag{1283}$$

- Chain transfer to polymer

$$R_x + P_y \overset{k_{ts}}{\rightarrow} R_y + P_x \tag{1284}$$

- Branching

$$R_x \overset{k_{bi}}{\rightarrow} R_x \tag{1285}$$

- β-Scission

$$R_x + P_y \overset{k_\beta}{\rightarrow} P_x + R_z + P^{=}_{y-z} \tag{1286}$$

- Reaction with terminal double bond

$$R_x + P^{=}_y \overset{k_{db}}{\rightarrow} R_{x+y} \tag{1287}$$

- Termination by disproportionation

$$R_x + R_y \overset{k_{td}}{\rightarrow} P_x + P_y \tag{1288}$$

- Termination by combination

$$R_x + R_y \overset{k_{tc}}{\rightarrow} P_{x+y} \tag{1289}$$

STAR-CCM+ solves scalar transport equations for moments as follows:

- For Live Polymers

  λ_0 is solved using r_{λ_0} as the source term

$$\lambda_0 = \sum_{n=1}^{\infty} R_n \tag{1290}$$

  where r_{λ_0} is defined by:

$$r_{\lambda_0} = \sum_{i=1}^{N_d} 2fk_{d_i}[I_i] - (k_{tc} + k_{td})\lambda_0^2 \tag{1291}$$

  λ_1 is solved using r_{λ_1} as the source term

$$\lambda_1 = \sum_{n=1}^{\infty} nR_n \tag{1292}$$

  where r_{λ_1} is defined by:

$$r_{\lambda_1} = \sum_{i=1}^{N_d} 2fk_{d_i}[I_i] - (k_{tm}[M] + k_{ts}[S])\lambda_1 + k_p[M]\lambda_0 - (k_{tc} + k_{td})\lambda_0\lambda_1 \tag{1293}$$

$$+ k_{tp}(\lambda_0\mu_2 - \lambda_1\mu_1) + k_{db}\lambda_0\mu_1 + k_\beta\left(\frac{1}{2}\lambda_0\mu_2 - \lambda_1\mu_1\right)$$

  λ_2 is solved using r_{λ_2} as the source term

$$\lambda_2 = \sum_{n=1}^{\infty} n^2 R_n \tag{1294}$$

  where r_{λ_2} is defined by:

$$r_{\lambda_2} = \sum_{i=1}^{N_d} 2fk_{d_i}[I_i] - (k_{tm}[M] + k_{ts}[S])\lambda_2 + 2k_p[M]\lambda_1 - (k_{tc} + k_{td})\lambda_0\lambda_2 \tag{1295}$$

$$+ k_{tp}(\lambda_0\mu_3 - \lambda_2\mu_1) + k_{db}(2\lambda_1\mu_1 + \lambda_0\mu_2)$$

$$+ k_\beta\left(\frac{1}{3}\lambda_0\mu_3 - \frac{1}{2}\lambda_0\mu_2 + \frac{1}{6}\lambda_0\mu_1 - \lambda_2\mu_1\right)$$

- For Dead Polymers

μ_0 is solved using r_{μ_0} as the source term

$$\mu_0 = \sum_{n=1}^{\infty} P_n \tag{1296}$$

where r_{μ_0} is defined by:

$$r_{\mu_0} = (k_{tm}[M] + k_{ts}[S])\lambda_0 + \left(\frac{1}{2}k_{tc} + k_{td}\right)\lambda_0^2 - k_{db}\lambda_0\mu_0 + k_\beta\lambda_0\mu_1 \tag{1297}$$

μ_1 is solved using r_{μ_1} as the source term

$$\mu_1 = \sum_{n=1}^{\infty} nP_n \tag{1298}$$

where r_{μ_1} is defined by:

$$r_{\mu_1} = (k_{tm}[M] + k_{ts}[S])\lambda_1 + (k_{tc} + k_{td})\lambda_0\lambda_1 \tag{1299}$$
$$- k_{tp}(\lambda_0\mu_2 - \lambda_1\mu_1) - k_{db}\lambda_0\mu_1 + k_\beta\left(\lambda_1\mu_1 - \frac{1}{2}\lambda_0\mu_2\right)$$

μ_2 is solved using r_{μ_2} as the source term

$$\mu_2 = \sum_{n=1}^{\infty} n^2 P_n \tag{1300}$$

where r_{μ_2} is defined by:

$$r_{\mu_2} = (k_{tm}[M] + k_{ts}[S])\lambda_2 + (k_{tc} + k_{td})\lambda_0\lambda_2 + k_{tc}\lambda_1^2 \tag{1301}$$
$$-k_{tp}(\lambda_0\mu_3 - \lambda_2\mu_1) - k_{db}\lambda_0\mu_2 + k_\beta\left[\lambda_2\mu_1 - \lambda_0\left(\frac{2}{3}\mu_3 + \frac{1}{2}\mu_2 - \frac{1}{6}\mu_1\right)\right]$$

The closure of the higher moment, μ_3, is adopted from Pladis and Kiparissides [225].

Based on the specified Heat of Propagation and Heat of Initiation, appropriate source terms are added to the energy equation.

By solving the transport of the above six Moment equations, you can obtain information about the polymer molecular structure, such as Polydispersity and the Molecular Weight Distribution of the polymer. STAR-CCM+ provides post-processing field functions for Polydispersity Index, Polymer Number-Averaged, and Mass-Averaged Chain Lengths. All Live and Dead Moments are also available as post-processing field functions.

Polydispersity Index is given by:

$$\frac{\mu_2 \mu_0}{\mu_1^2} \tag{1302}$$

Polymer Number-Averaged Chain Length (NACL) is given by:

$$\frac{\mu_1}{\mu_0} \tag{1303}$$

Polymer Mass-Averaged Chain Length (MACL) is given by:

$$\frac{\mu_2}{\mu_1} \tag{1304}$$

Polymerization Nomenclature

M	Monomer
I	Initiator
R	Live polymer
λ_i	Live polymer with i th moment
R_n	Live polymer with n monomer units
P	Dead polymer

μ_i	Dead polymer with i th moment
P_n	Dead polymer with n monomer units
$R^{\cdot}$	Radical
r	Reaction rate
S	Solvent
f	Initiator frequency
f_i	Initiator efficiency
k	Rate coefficient for different polymer reactions

Polymerization Bibliography

[225] Pladis and Kiparissides, Chem Eng Sc, 53, No.18, pp 3315, 1998

Working with the Fuel Mass Fraction

This initial condition value is used by the Premixed Eddy Break-Up model (PEBU), the Coherent Flame model (CFM), and the Turbulent Flame Speed Closure (TFC) model. It can also be activated for an individual boundary by selecting `Specified Value` as the Method in the Physics Conditions > Wall Combustion Scalar Option node.



The fuel mass fraction is the actual mass fraction of the combustible material (fuel) in a cell or on a boundary. For a fully burnt condition, the fuel is consumed, so this value is 0.

For an unburnt condition, the fuel mass fraction is the mass fraction of

- The fuel in the unburnt gas stream (PEBU or CFM), or
- The mixture fraction times the mass fraction of the fuel in the fuel stream (PCFM).

For example, consider a simulation that uses CH4, O2 and N2:

- With the PEBU or CFM activated, the mass fraction of CH4 is the fuel mass fraction in a CH4-O2-N2 unburnt gas stream.
- With the PCFM activated, using a CH4-N2 fuel stream and an O2-N2 oxidizer stream, the mass fraction of CH4 in the whole boundary is the fuel mass fraction. So if the boundary has a mixture fraction of 0.5, and CH4 is 50% (mass based) of the fuel stream, then the fuel mass fraction is 0.25.

This value is Y_f in Eqns. (1040) through (1042) for the PEBU and (1084) through (1086) for the CFM.

Properties

This value is input as a scalar profile.

Defining Chemical Reactions

Once an EBU model is selected, the Models node of the continuum contains a Reacting node that has a Reactions sub-node, as shown below.



Use this node to define chemical reactions occurring between mixture components that were defined earlier, either by importing Chemkin-formatted data or manually, as follows:

- Right-click the Reactions node and select New Reaction to set up a new (blank) reaction definition.

  A node that is named Reaction n is added at the bottom of current set, where **n** is the sequence number.
- Define the reaction components and other physical parameters.
- Repeat the above steps until all reactions required to describe the problem's chemistry fully have been defined.

What Is a Reaction Definition?

The Reaction n node represents a chemical reaction in the simulation tree, as shown below.



The reaction definition consists of two parts:

- An expression specifying the relationship between reactants and products —that is, which reaction components occur on the left— or right-hand side of the reaction equation and in what proportions. The participants on each side are defined separately within two sub-nodes that are called Reactants and Products.
- For the EBU model only, properties that are needed for calculating the reaction rate: these properties are defined separately within a sub-node that is called Properties.

What Is a Reactant?

A reactant is any reaction component appearing on the left-hand side of a chemical reaction equation, created within the Reactants node as shown below.



To define a reactant:

- Right-click the Reactants node and select Add Reactant.
- Choose a name from the submenu.

  Note that the menu items correspond to components that have been defined as part of the mixture and not already selected as a reactant in this reaction.
- Define the required parameters for each component.
- Repeat the above steps until all components on the left-hand side of the equation have been specified.

What Is a Product?

A product is any reaction component appearing on the right-hand side of a chemical reaction equation, created within the Products node as shown below.



To define a product:

- Right-click the Products node and select Add Product.
- Choose a name from the submenu.

  Note that the menu items correspond to components that have been defined as part of the mixture and not already selected as a product in this reaction.
- Define the required parameters for each component.
- Repeat the above process until all components on the right-hand side of the equation have been specified.

What Is a Reaction Component?

A reaction component is any constituent of a multi-component mixture that has already been defined as a reactant or product in a chemical reaction equation, as shown below.



To specify the required parameters for a given reaction component:

- Select the component.
- In the *Properties* window, specify:
  - A stoichiometric coefficient ν_{ij} appropriate to the current reaction (see Finite-Rate Kinetics Convention)
  - For the EBU model and surface chemistry reactions, specify a rate exponent p_{ij} (see Chemical Reaction Rates from Finite-Rate Kinetics, Eqn. (1032)).

    This parameter is used to calculate chemical reaction rates from finite-rate kinetics and its value is not required if the Standard EBU model has been selected in the Eddy Break-up node.
- Repeat the above process until all components on both sides of the reaction equation have been fully specified.

Setting Reaction Properties

In a reaction object, the Properties manager node is subject to common menu items, although it has no properties of its own. Rather, it is a placeholder for objects representing the properties of a reaction.



Reactions in STAR-CCM+ can be reversible.

The properties objects are of two types:

- The Ebu coefficients
- The reaction coefficient

What Are Reversible Reactions?

Reversible reactions are used to specify a reaction that progresses in both directions. For example, the reversible reaction that is given by:

$$A + B \leftrightarrow C + D \tag{1305}$$

represents two one-way reactions; the forward reaction:

$$A + B \rightarrow C + D \tag{1306}$$

and the backward, or reverse, reaction:

$$C + D \rightarrow A + B \tag{1307}$$

The forward and backward reactions progress at different rates. In a reversible reaction, you specify the forward rate coefficients and the backward rate is calculated from the forward rate and the equilibrium constants that are determined from the thermodynamics.

If you define a reversible reaction, make sure that either the thermodynamic polynomial or polynomial in T is chosen for calculation of specific heat for each component of the mixture.

Third Body

Some reactions require a "third body" in order to proceed, as in:

$$H + O2 + M \leftrightarrow HO2 + M \tag{1308}$$

where M represents the third body. Third bodies are somewhat of a catalyst. All species in the mixture can act as third bodies. Third bodies affect the rate at which the reaction proceeds.

By default, each species in the mixture is assumed to contribute with equal efficiency as a third body. Hence all species are given a third body efficiency of unity. If one or more species contribute more or less as a third body, then the efficiency coefficient of that component can be specified explicitly.

What Are the Ebu Coefficients?

The Ebu Coefficients node is a placeholder for Ebu rate coefficient nodes. This node is activated when the Reaction Control property of the EBU model is set to `Standard EBU`, `Hybrid`, or `Combined Time Scale`. It is the only properties node available if the EBU model is set to `Standard EBU`. When `Kinetics Only` is selected, the EBU Coefficients node is unavailable.



The Ebu rate coefficient nodes specify coefficients for the EBU reaction rate. They are scalar profiles.



Settings in Ebu rate coefficient, A affect Eqn. (1025) and (1026). Settings in Ebu rate coefficient, B affect Eqn. (1026).

Ebu Coefficients Manager Properties

Use products for the rate	Determines whether Eqn. (1025) or Eqn. (1026) is to be used in the calculations. Ticked — Uses Eqn. (1026). Cleared — Uses Eqn. (1025).

What Is the Reaction Coefficient?

The Reaction Coefficient node, which is activated when the EBU model is set to `Hybrid`, `Combined Time Scale`, or `Kinetics Only`, is a placeholder for the Arrhenius Coefficients and User Reaction Coefficient nodes. It is the only properties node available if the Reaction Control property of the EBU model is set to `Kinetics Only`.



The Arrhenius and user reaction parameters are used to calculate chemical reaction rates from finite-rate kinetics, and their values are not required if `Standard EBU` has been selected in the Reaction Control property of the Eddy Break-up node. Therefore when `Standard EBU` is selected, the Reaction Coefficient node is unavailable.

Reaction Coefficient Properties

Select method	Determines the method for calculating the reaction.	
	Arrhenius Coefficients	Activates the Arrhenius Coefficients node.
	User Reaction Coefficient	Activates the User Reaction Coefficient node.

What Are the Arrhenius Coefficients?

The Arrhenius Coefficients node, which has properties and is subject to common menu items, is activated when the Select method property of the Reaction Coefficient manager node is set to `Arrhenius Coefficients`.



The Pre-exponent subnode lets you work with the pre-exponential factor A_j in Eqn. (1032).

Arrhenius Coefficients Properties

Temperature exponent, beta	Sets the temperature exponent β_j for this reaction in Eqn. (1032).
Activation energy, Ea	Sets the activation energy E_{aj} for this reaction in Eqn. (1032).
Reversible	Specifies whether this reaction is reversible. Ticked — Makes the reaction reversible. Cleared — Leaves the reaction irreversible.
Third bodies	Specifies whether this reaction involves third bodies. Ticked — Involves third bodies, and activates the Third Body Efficiencies node. Cleared — Excludes third bodies from the reaction.

What Is the User Reaction Coefficient?

The User Reaction Coefficient node, which is a scalar profile and is subject to common menu items, is activated when the Select method property of the Reaction Coefficient manager node is set to `User Reaction Coefficient`.



This object uses the scalar profile functionality to specify K_f in Eqn. (1033). This setting does not include the molar concentrations raised to the rate exponent, since the STAR-CCM+ solver works with that internally.

What Are Third Body Efficiencies?

Some reactions require a "third body". The Third Body Efficiencies node, which is controlled through a pop-up menu, is activated when the checkbox of the Third bodies property of the Arrhenius Coefficients node is ticked.



What Are Third Body Coefficients?

Third body coefficient nodes, which have properties and are subject to common menu items, are objects that correspond to a reaction component that you select for this reaction. Each node's name is the same as the material symbol.



Third Body Coefficient Properties

Third Body Efficiency Coefficient	Sets the third body efficiency coefficient for the corresponding reaction component.

Adding Third Body Coefficients

To add a third body coefficient, right-click the Third Body Efficiencies node. From the pop-up menu, select Add Third Body Coefficient and then an available reaction component from the submenu.



A third body coefficient node (H2S in the following screenshot) appears in the object tree.



Post-Processing Reaction Rates

You can visualize forward and backward (where applicable) reaction rates for each reaction in a simulation. In the current version of STAR-CCM+, this feature is available with the Eddy Break-Up (EBU) and Eddy Contact Micromixing models.

To activate this feature:

- Before you run the simulation, select the Eddy Break-up or Eddy Contact Micromixing node.
- Activate Store Reaction Rates.

  After the simulation run, each reaction rate is available in the scalar function list under the item `Reaction Rate Forward` or `Reaction Rate Backward`.

Note: If you set the Reaction Control property of the EBU model to `Standard EBU`, you can only visualize forward reaction rates.

In the following example, the forward reaction rate $H2 + 0.5O2 \rightarrow H2O$ appears in a scalar scene.



Importing Chemkin-Formatted Data

You can specify mixture components plus define chemical reactions by importing Chemkin-formatted data into a STAR-CCM+ simulation. This technique allows you to use legacy data in this format efficiently.

The file format that is referenced here is known as Chemkin II. This section provides the following information about working with it:

- Importing mixture components from a file
- Importing chemical reactions from a file
- Importing the complex chemistry definition directly from a file
- Restrictions in the use of the Chemkin II format

Importing Mixture Components from a File

To import gas components from a Chemkin-formatted (`.chm`) file:

- Open the Continua > [Continuum Name] > Models > Multi-Component Gas node.
- Right-click the Gas Components node and select Import Species...



- In the *Select file and material database* dialog that appears, provide the path to the `.chm` file. Enter the path in the Select File text box or click Browse.



When you click Browse, a standard *Open* dialog is activated. Use it to navigate to your file.



- Click Open.
- Back in the *Select file and material database* dialog, the file path is in the Select File text box. Click OK to import the Chemkin data.

The Gas Components node is populated with species data.



Importing Chemical Reactions from a File

To import chemical reactions from a Chemkin-formatted (`.chm`) file:

- Open the Continua > Physics n > Models > Reacting node.

- Right-click the Reactions node and select Import Reactions from a File...



A standard *Open* dialog is activated. Use it to navigate to your file.



- Click Open to import the Chemkin data.

The Reactions and Gas Components nodes are populated with reaction and species data, respectively.



The Gas Components node is populated automatically by this operation, since the combustion modeling requires gas components to be specified before reactions can be defined.

Importing the Complex Chemistry Definition Directly from a File

You can import complex chemistry definitions directly from a file. Two file formats (DarsCfd and Chemkin) are currently available. These options are activated by right-clicking on the DarsCfd Reaction Model node.

Importing from a DarsCfd-Formatted File

If the Run Time Compile expert property is activated (checkbox ticked), ensure that case-specific F90 files (`*.f90` and `*.txt`) exist. If these files were created by STAR-CCM+ (see Importing from a Chemkin-formatted file), they are in a directory that is called `DarsInputFiles`.

If that expert property is not activated, make sure that you are selecting a folder that contains the actual case-specific pre-compiled library.

Right-click the DarsCfd Reaction Model node and select Import Complex Chemistry Definition (DarsCfd format)...



In the *Open* dialog, navigate to and select the `DarsInputFiles` directory and click Open.

If the Run Time Compile expert property is activated, STAR-CCM+ will then:

- Call the Fortran compiler to compile the case-specific F90 files and package them into a case-specific Dars-CFD library.
- Read in the case-specific Dars-CFD library.

Importing from a Chemkin-Formatted File

Right-click on the DarsCfd Reaction Model node and select Import Complex Chemistry Definition (Chemkin format)...



When this option is selected, a dialog appears.



The fields in the dialog are:

- Project name (default is `StarccmProject`)
- Chemkin files -- clicking Browse... opens a file chooser. Multiple selections are enabled (useful when dealing with gas and surface mechanisms).
- Thermal properties files -- the default is the file that comes with the installation. If you do not have any additional species for which the thermodynamical polynomial has not been defined in the default file, or you do not want to override the polynomials for the existing species, you do not have to change this file. Otherwise, clicking Browse... opens a file chooser with multiple selections enabled.
- Molecular properties files -- the default is the file that comes with the installation. If you do not have any additional species for which molecular properties have not been defined in the default file, or you do not want to override the molecular properties for the existing species, you do not have to change this file. Otherwise, clicking the button Browse... opens a file chooser with multiple selections enabled.

Multiple selections are possible (for example, it is possible to choose two Chemkin files). However, the files have to be in the same folder. With multiple selections and certain properties defined in both files, the properties that are recorded in the last file to be read remain in the database. For example, if `generalTherm.dat` and `surfTherm.dat` are selected (in that order) as thermodynamic files that define NASA polynomials, and solid phase carbon C is defined in both of them, the NASA polynomial recorded in the material database comes from the file `surfTherm.dat`.

For a pure gas chemistry simulation, if the Binary Problem Description checkbox is NOT ticked, STAR-CCM+ converts the Chemkin files to Dars-CFD ASCII input files and import them. If the Binary Problem Description option is on, together with the Run Time Compile option, then the following actions take place:

- STAR-CCM+ converts Chemkin files to Dars-CFD ASCII input files.
- The Dars-redkin library is called to convert the ASCII input files into case-specific F90 files.
- The Fortran compiler is called to compile these F90 files into a case-specific Dars-CFD library.
- The case-specific Fortran library is then read in.

If the Binary Problem Description option is on but the Run Time Compile option is not, the import procedure is incomplete. Dars-CFD files (both ASCII and `*.f90`) are created, but since in this case STAR-CCM+ does not perform the run-time compile, there is no library to import. This option can be useful if this procedure is performed on a Windows machine or on a machine that does not have a FORTRAN compiler. In such instances, you can take the `*.f90` files and compile them either manually or on a different machine. Then you can import the library with the first import option (Import Complex Chemistry Definition (Dars-CFD format)...).

If you wish to use the Binary Problem Description option with Chemkin files, make certain that the Fortran compiler is provided and installed in the standard location (that is, in the `PATH` variable).

The file format that is read is compatible with the Chemkin-II format subject to restrictions listed in the following section.

Restrictions in the Use of the Chemkin II Format

Restrictions exist on importing the following in Chemkin II format:

- Surface mechanisms
- Species
- Reactions

Restrictions on Importing Surface Mechanisms

- Only one site and one bulk phase per surface mechanism is allowed.
- Species in the site phase must have suffix (S) and species in the bulk phase must have the suffix (B).

Restrictions on Importing Species

- Thermodynamical polynomials and molecular data also must be present for the imported species, either in the supplied files or in the user-defined ones.
- The species phase in the user-defined thermdat.dbs file must be G(gas), L(liquid), or S(solid).

Restrictions on Importing Reactions

- There should be no space between the stoichiometric coefficient and reactant/product.

  For example, this format is invalid:

  ```
  CH4 + 2 O2 => CO2 + 2 H2O
  ```

  The proper format is:

  ```
  CH4 + 2O2 => CO2 + 2H2O
  ```

Using Ignitors

The Ignitors node, which has its own pop-up menu, is the manager object for ignitors that are used in the continuum.



This section explains the following:

- Using the Ebu ignitor
- Using the Fixed Temperature ignitor
- Using the Flame Area Density ignitor
- Using the Fuel Mass Fraction ignitor
- Using the Progress Variable ignitor
- What is the pulse activator?
- What is the always activator?

What Are Ignitors?

Ignitors allow you to control the simulation of combustion when using any of the following model types:

- Eddy break-up (EBU)
- Coherent Flame Model (CFM)
- Turbulent Flame Speed Closure (TFC)
- Progress Variable Model (PVM)

When any of these models are added to the simulation, an Ignitors manager node appears under the Continua > Physics node.

When you create an ignitor, the node that appears has its own properties and is subject to common menu items.



The supported ignitors are:

- The Ebu ignitor and Fixed Temperature ignitor for use with the EBU models
- The Flame Area Density ignitor for use with the CFM
- The Fuel Mass Fraction ignitor for use with the CFM or TFC model
- The Progress Variable ignitor for use with the PVM

Ignitor Properties

All ignitors have the following properties.

Activator — The timing controller which specifies when the ignitor is "on", where what "on" means is Ignitor-dependent.

- **Always** — Indicates that the Ignitor is always on. The Always activator node that corresponds to this option has no additional properties.
- **Pulse** — Allows the ignitor to be turned on for a user-specified duration during the simulation. The Pulse activator node corresponds to this option.

Shapes	The volume shapes (selected from the Volume Shape manager under Tools) which represent the area or volume for the ignitor to “ignite”, where “ignite” is dependent on the type of ignitor. Mesh cells which are contained within these volume shapes have their properties changed as specified by the ignitor when it is on.

Using the Ebu Ignitor

This section explains the following:

- What is the Ebu ignitor?
- Creating the Ebu ignitor

What Is the Ebu Ignitor?

This ignitor has no additional properties.



The Ebu Ignitor forces reaction rates to be mixing limited without using products. This is limited to the cells which are in the EBU ignitor volume shapes while the ignitor is on.

This ignitor requires that the Reaction Control property for the Eddy break-up (EBU) Combustion model is set to either EBU or Hybrid with turbulent flow.

Creating an Ebu Ignitor

To create an Ebu ignitor, you must have selected one of the Eddy break-up (EBU) models. Right-click the Continua > [physics continuum node] > Ignitors node and select New > Ebu Ignitor.



This menu item may simply appear as New, since the fixed temperature ignitor is not available when the EBU model's Reaction Control property is set to `Standard EBU`.

An EbuIgnitor node appears in the object tree. Its label has a number at the end indicating the total of Ebu ignitors in the continuum.



Using the Fixed Temperature Ignitor

This section explains the following:

- What is the Fixed Temperature ignitor?
- Creating the Fixed Temperature ignitor

What Is the Fixed Temperature Ignitor?

The Fixed Temperature ignitor has an additional property of Temperature value. This property is a constant which represents the temperature (K) at which to set the cells contained within the ignitor's volume shapes while the ignitor's activator is on.



This ignitor is supported by the Eddy break-up (EBU) Combustion model. It is not available when the EBU model's Reaction Control property is set to `Standard EBU`.

Creating a Fixed Temperature Ignitor

To create a Fixed Temperature ignitor, you must have selected one of the Eddy break-up (EBU) models. Right-click the Continua > [physics continuum node] > Ignitors node and select New > Fixed Temperature Ignitor.



This menu item simply appears as New if the EBU model's Reaction Control property is set to `Kinetics Only`.

A TemperatureIgnitor node appears in the object tree. Its label has a number at the end indicating the total of Fixed Temperature ignitors in the continuum.



Using the Flame Area Density Ignitor

This section explains the following:

- What is the Flame Area Density ignitor?
- Creating the Flame Area Density ignitor

What Is the Flame Area Density Ignitor?

This Ignitor has additional properties. While it is on, the Flame Area Density ignitor sets the value of the cells contained within the volume shapes to that of the Flame Density Value property. When it is off, this ignitor does not affect the values of flame area density.



Flame Area Density Ignitor Properties

In addition to the standard ignitor properties, the Flame Area Density ignitor has the following:

Ignitor Variable Value	A floating point value for flame density.
Ignitor Set Method	Sets the method to determine the value of flame area density (FAD) in each cell.
Maximum of Cell FAD or Ignitor Constant	Uses the greater of the cell FAD obtained from the transport equation, and the ignitor constant set with the Ignitor Variable Value property.
Ignitor Constant	The latter uses the ignitor constant regardless of the value from the transport equation. This option is helpful when you know the exact FAD value for one cell.

Creating a Flame Area Density Ignitor

To create a Flame Area Density ignitor, you must have selected the Coherent Flame model (CFM). Right-click the Continua > [physics continuum node] > Ignitors node and select New.



A FlameAreaDensityIgnitor node appears in the object tree. Its label has a number at the end indicating the total of Flame Area Density ignitors in the continuum.



Using the Fuel Mass Fraction Ignitor

This section explains the following:

- What is the Fuel Mass Fraction ignitor?
- Creating the Fuel Mass Fraction ignitor

What Is the Fuel Mass Fraction Ignitor?

This Ignitor has additional properties. While it is on, the Fuel Mass Fraction ignitor sets the value of the cells contained within the volume shapes to that of the Ignitor Variable Value property. When it is off, this ignitor does not affect the values of the fuel mass fraction.



Fuel Mass Fraction Ignitor Properties

In addition to the standard ignitor properties, the Fuel Mass Fraction ignitor has the following:

Ignitor Variable Value	A floating point value for the fuel mass fraction.	
Ignitor Set Method	Sets the method to determine the value of the fuel mass fraction in each cell.	
	Maximum of Cell Value or Ignitor Constant	Uses the greater of the cell value obtained from the transport equation, and the ignitor constant set with the Ignitor Variable Value property.
	Ignitor Constant	The latter uses the ignitor constant regardless of the value from the transport equation. This option is helpful when you know the exact value for one cell.

Creating a Fuel Mass Fraction Ignitor

To create a Fuel Mass Fraction ignitor, you must have selected the Coherent Flame model (CFM) or the Turbulent Flame Speed Closure (TFC) model. Right-click the Continua > [physics continuum node] > Ignitors node and select New.



A FuelMassFractionIgnitor node appears in the object tree. Its label has a number at the end indicating the total of Fuel Mass Fraction ignitors in the continuum.



Using the Progress Variable Ignitor

This section explains the following:

- What is the Progress Variable ignitor?
- Creating the Progress Variable ignitor

What Is the Progress Variable Ignitor?

This Ignitor has additional properties. While it is on, the Progress Variable ignitor sets the value of the cells contained within the volume shapes to that of the Ignitor Variable Value property. When it is off, this ignitor does not affect the values of the progress variable.



Progress Variable Ignitor Properties

In addition to the standard ignitor properties, the Progress Variable ignitor has the following:

Activator	The timing controller which specifies when the ignitor is "on", where what "on" means is Ignitor-dependent.	
	Always Activator	Indicates that the Ignitor is always on. The Always activator node that corresponds to this option has no additional properties.
	Pulse Activator	Allows the ignitor to be turned on for a user-specified duration during the simulation. The Pulse activator node corresponds to this option.
Ignitor Variable Value	A floating point value for the progress variable.	
Ignitor Set Method	Sets the method to determine the value of the progress variable in each cell.	
	Maximum of Cell value or Ignitor Constant	Uses the greater of the cell Progress Variable obtained from the transport equation and the ignitor constant set with the Ignitor Variable Value property.
	Ignitor Constant	The latter uses the ignitor constant regardless of the value from the transport equation. This option is helpful when you know the exact Progress Variable value for one cell.
Shapes	The volume shapes (selected from the Volume Shape manager under Tools) which represent the area or volume for the ignitor to "ignite", where "ignite" is dependent on the type of ignitor. Mesh cells which are contained within these volume shapes have their properties changed as specified by the ignitor when it is on.	

Creating a Progress Variable Ignitor

To create a Progress Variable ignitor, you must have selected the Progress Variable Model. Right-click the Continua > [physics continuum node] > Ignitors node and select New.



A ProgressVariableIgnitor node appears in the object tree. Its label has a number at the end indicating the total of Progress Variable ignitors in the continuum.



What Is the Pulse Activator?

The Pulse Activator node corresponds to the setting of the same name of an ignitor's Activator property.



It has two properties: "begin" and "end". These indicate the time (either iteration or physical time depending on how stepping is specified in the simulation) which the Ignitor is to be turned on and off, respectively. The activator is on for iterations when begin <= iteration <= end, and for physical time when begin < time <= end.

What Is the Always Activator?

The Always Activator node corresponds to the setting of the same name of an ignitor's Activator property. It indicates that the Ignitor is always on.



Pop-Up Menu of the Ignitors Manager

One of these specific pop-up menu items are available when you right-click the Ignitors manager node, depending on which combustion model you are using.

New	Creates an ignitor of the selected type. This menu appears only when one of the Eddy break-up (EBU) models is selected.	
	Ebu Ignitor	Creates an Ebu ignitor.
	Fixed Temperature Ignitor	Creates a Fixed Temperature ignitor.
New	Creates an ignitor, depending on the specific type of combustion model selected.	

See also: common menu items.

Using Fluid Streams

The Fluid Stream Manager node, which has properties and a pop-up menu, contains all the fluid streams of a continuum.



This node, along with the appropriate sub-nodes, appears when one of the following models is selected:

- PPDF
- Coherent Flame
- Progress Variable
- Thickened Flame

Although the default configuration of the fluid streams is considered optimal, it is possible to delete and add them.

What Is a Fluid Stream?

Fluid stream objects allow you to specify the temperatures and composition of the streams entering the combustion chamber. They are represented by nodes in the Fluid Stream Manager node; these nodes have properties and are subject to common menu items.



If you use a PPDF model, the Fuel Stream and Oxidizer Stream objects (shown above) are used. If you select the Coherent Flame model or the Premixed Eddy Break-Up model, the Unburnt Gas object is the single fluid stream used.

The required user input in the properties for this node is the stream temperature and the stream type (fuel, oxidizer, or other). Where a choice of multiple stream types is available, it does not affect the computation, but it aimed to help you distinguish between various fluid streams.

The composition of each fluid stream is represented by component objects, which can be added and removed.

The properties of fluid streams cannot be changed as long as there is a PPDF flamelet or equilibrium table that is not empty. To change the properties, delete the tables first.

Managing Fluid Stream Components

The Components manager node, which has a pop-up menu, contains all the components of the fluid stream.



Initially, this node is empty. You define the chemical species present in the stream by adding new components.

Components Manager Pop-Up Menu

Add Fluid Stream Component	Adds components from a submenu of materials that have been defined as part of the mixture, but not yet added to that particular fluid stream.

See also: common menu items.

What Is a Fluid Stream Component?

Component nodes, which have properties and are subject to common menu items, represent constituents of a multi-component mixture that have been included in the fluid stream.



The only user input that is required in the *Properties* window for each component that is selected is specification of its Mass Fraction in the stream.

Component Properties

Mass Fraction	The value and units of the mass fraction. Clicking the customizer executes the define value action.

Adding Fluid Stream Components

To add a component to a fluid stream, right-click the Components manager node and select Add Fluid Stream Component. Choose a name from the submenu.



The submenu items correspond to components that have been defined as part of the mixture, but not those already selected as a component of this fluid stream.

Removing Fluid Stream Components

Components can be removed by right-clicking that component's node and selecting Delete. This action removes the component from the fluid stream, but not from the mixture.

Fluid Stream Properties

Temperature of the stream	Sets the temperature of a particular stream.	
Fluid Stream Type	Selects one of the available fluid stream types for a particular stream.	
	Fuel	Selects fuel -- available for every stream but the last (which by default has to be an oxidizer).
	Oxidizer	Selects oxidizer -- available for every stream but the first (which by default has to be a fuel).
	Other	(Available for every stream but the first and last.)
Index	A read-only property.	

Adding Fluid Streams

To add a fluid stream, right-click the Fluid Stream Manager node and select New Fluid Stream.



This menu item is disabled if the maximum number of fluid streams for the continuum has been reached.

Deleting Fluid Streams

Similar to other objects, fluid streams can be deleted using the standard methods.

Properties of the Fluid Stream Manager

Maximum number of fluid streams	This indicates the limit of fluid streams that can be in the continuum. It is a read-only property that is determined by the selection of the combustion model.

Pop-Up Menu of the Fluid Stream Manager

New Fluid Stream	Creates a fluid stream. This menu item is disabled (grayed out) when the maximum number of fluid streams for the continuum is reached.

Using Laminar Flame Speed Methods

STAR-CCM+ provides two established methods for calculating laminar flame speed, Gulder and Metghalchi, along with a user-defined method. These methods are available with the Coherent Flame Model (CFM), the Thickened Flame model, and the Turbulent Flame Speed Closure (TFC) model.

See also:

Laminar Flame Speed Calculation

Using the Gulder Laminar Flame Speed

The laminar flame speed correlation that Guelder proposed is given in Eqn. (1310). As it uses fuel-dependent constants, you must specify the fuel type.

Use this method for general applications.

Gulder Laminar Flame Speed Properties

Fuel Name	Specifies the fuel name control mechanism option for CFM combustion. Select an existing fuel that has the closest chemical structure to the real fuel. Any differences have effects on the solution, but the approximation is still good.	
	Methane	Specifies methane (CH4).
	Propane	Specifies propane (C3H8).
	Methanol	Specifies methanol (CH4O).
	Ethanol	Specifies ethanol (C2H6O).
	Iso-Octane	Specifies iso-octane (C8H18).

Using the Metghalchi Laminar Flame Speed

Unlike the Gulder method, the Metghalchi method (given in Eqn. (1309)) has no properties.

If you choose this method, set the equivalent ratio of the fuel, ϕ, to be greater than or equal to 0.65 for the following reasons:

- None of the combustion models can accurately predict burning in lean mixtures.
- The laminar flame speed correlation function progressively diverges from the experimentally observed data in lean mixtures.

Using the User-Defined Laminar Flame Speed

The User-Defined Laminar Flame Speed node lets you control laminar flame speed just like a typical scalar profile.

Working with the Laminar Flame Speed Profile node, you can specify laminar flame speed using any of the following methods:

- Constant
- Field function
- Tables, including a table with laminar flame speed as a function of equivalence ratio
- User code

Laminar Flame Speed Calculation

STAR-CCM+ incorporates two correlations for calculating the unstrained laminar flame speed. The one proposed in Metghalchi and Keck [226] is calculated as follows:

$$S_u = S_{u0}\left(\frac{T_u}{T_0}\right)^{\alpha}\left(\frac{p}{p_0}\right)^{\beta} \qquad \text{m/s} \tag{1309}$$

where p is the pressure, T is the temperature, the subscripts 0 and u denote reference and unburnt gas properties, respectively, and S is the laminar flame speed. The reference laminar flame speed S_{u0} and the exponents α and β depend on the equivalence ratio ϕ of the fuel. If $0.4 \le p \le 50$ atm, $300 \le T \le 700$ K and $0.8 \le \phi \le 1.5$, the authors claim that the laminar flame speed is within 10% of the measured data. $\phi \ge 0.65$ is recommended because:

- None of the combustion models can accurately predict burning in lean mixtures.
- The laminar flame speed correlation function progressively diverges from the experimentally observed data in lean mixtures.

The second laminar flame speed correlation, which Guelder proposed [227], is calculated as follows:

$$S_u = Z\,W\phi^{\eta}\exp[-\xi(\phi-1.075)^2]\left(\frac{T_u}{T_0}\right)^{\alpha}\left(\frac{P}{P_0}\right)^{\beta} \tag{1310}$$

where Z, W, η, ξ, α and β are fuel-dependent constants that are defined in the table below:

Table 5-1:

Fuel	***Z***	***W***	η	ξ	α	β	
						ϕ **< 1**	ϕ **> 1**
Methane	1	0.422	0.15	5.18	2.00	–0.5	
Propane	1	0.446	0.12	4.95	1.77	–0.2	
Metha-nol	1	0.492	0.25	5.11	1.75	$-0.2/\sqrt{\phi}$	$-0.2\,\phi$
Ethanol	1	0.465	0.25	6.34	1.75	$-0.17/\sqrt{\phi}$	$-0.17\,\phi$
Iso-octan e	1	0.465 8	–0.32 6	4.48	1.56	–0.22	

As the correlation by Metghalchi and Keck is only valid for ϕ close to 1, you are advised to use the correlation by Guelder for general applications.

Laminar Flame Speed Bibliography

[226] Metghalchi, M. and Keck, J.C. 1982 'Burning velocities of mixtures of air with methanol, isooctane, and indelene at high pressure and temperature', Combust. Flame, 48, pp. 191-210.

[227] Guelder, O.L. 1990 'Turbulence premixed flame propagation models for different combustion regimes', 23rd Symp. (Int.) on Combustion, The Combustion Institute, pp. 743-750.

Combustion Field Functions Reference

The following primitive field function is common to more than one combustion model. It has typical base properties.

Chemistry Heat Release Rate

The sum of the reaction rate multiplied by the formation enthalpy for each species. The formula is:

$$Q_r = \sum_{i=1}^{N} (\omega_i H_{f,i}^0)$$

where Q_r is the chemistry heat release rate (W), ω_i is the reaction rate for species i, $H_{f,i}^0$ is the formation enthalpy for species i, and N is the number of species. It is available for two types of combustion models: Eddy Break-Up (EBU) and Homogeneous Reactor.

To calculate total thermal power, select this field function for a sum report. For parts, the sum report should use the region where the reaction continuum is active.

Modeling Multiphase Flow

Multiphase flow is a term which refers to the flow and interaction of several phases within the same system where distinct interfaces exist between the phases. The term 'phase' usually refers to the thermodynamic state of the matter: solid, liquid, or gas. In modeling terms, a phase is defined in broader terms, and can be defined as a quantity of matter within a system that has its own physical properties to distinguish it from other phases within the system. For example:

- Liquids of different density
- Bubbles of different size
- Particles of different shape

Multiphase flows are different from multi-component flows. In multi-component flows, the different species are mixed at the molecular level. These species have the same convection velocity. In multiphase flows, the different phases are mixed at the macroscopic scale. These phases have different convection velocity. Many flows are multiphase multi-component flows.

Multiphase flows can be classified into two categories:

- Dispersed flows, such as bubbly, droplet, and particle flows
- Stratified flows, such as free surface flows, or annular film flow in pipes.

A phase is considered dispersed if it occupies disconnected regions of space—otherwise it is continuous.

STAR-CCM+ provides six distinct models to meet the requirements of these two categories of flow:

- The Multiphase Segregated Flow model: this model is commonly known as the Eulerian Multiphase model in the literature, but that term has been given a wider significance in STAR-CCM+. The Multiphase Segregated Fluid model solves conservation equations for mass,

momentum, and energy for each phase. Phase interaction models are provided to define the influence that one phase exerts upon another across the interfacial area between them.

- The Lagrangian Multiphase model: this model solves the equation of motion for representative parcels of the dispersed phase as they pass through the system. It is intended for systems that consist mainly of a single continuous phase carrying a relatively small volume of discrete particles, droplets, or bubbles. It is suited where the interaction of the discrete phase with physical boundaries is important.
- The Dispersed Multiphase Model (DMP): this model simulates dispersed phases in a Eulerian manner. The Dispersed Multiphase model combines aspects of both the Lagrangian Multiphase (LMP) model and the Segregated (Eulerian) Multiphase (EMP) models. The DMP model and the Volume of Fluid (VOF) model can be activated in the same simulation.
- The Discrete Element Model (DEM): this model is an extension of the Lagrangian Multiphase model, but where individual particles are modeled rather than representative parcels, and where inter-particle contact forces are explicitly accounted for.
- The Fluid Film model: this model predicts the dynamic characteristics of wall films using boundary layer approximations and assumed velocity and temperature profiles across the depth of the film. Film transport is predicted using thin shells that lie across the surface of solid walls on which the film is formed.
- The Volume of Fluid (VOF) model: this model is provided for systems containing two or more immiscible fluid phases, where each phase constitutes a large structure within the system (such as typical free surface flows). This approach captures the movement of the interface between the fluid phases, and is often used for marine applications.

In mathematical terms, both the Volume of Fluid model and the Segregated Flow model utilize a Eulerian (as opposed to a Lagrangian) framework in their formulation. Therefore, setting up simulations for both these approaches involve use of distinct Eulerian phases.

Both the Multiphase Segregated model and the Volume of Fluid model use distinct Eulerian phases, as opposed to a Lagrangian phase and a dispersed phase. Therefore, a common framework has been developed to specify the phases for both models. This framework is called the Eulerian Multiphase model in STAR-CCM+.

A common framework has also been developed to specify the models that define how pairs of phases interact with each other. These models are known as phase interactions.

What Are Multiphase Flows?

Multiphase flow is a term which refers to the flow and interaction of several phases within the same system where distinct interfaces exist between the phases. The term 'phase' usually refers to the thermodynamic state of the matter: solid, liquid, or gas. In modeling terms, a phase is defined in broader terms, and can be defined as a quantity of matter within a system that has its own physical properties to distinguish it from other phases within the system. For example:

- Liquids of different density
- Bubbles of different size
- Particles of different shape

Multiphase flows are different from multi-component flows. In multi-component flows, the different species are mixed at the molecular level. These species have the same convection velocity. In multiphase flows, the different phases are mixed at the macroscopic scale. These phases have different convection velocity. Many flows are multiphase multi-component flows.

Multiphase flows can be classified into two categories:

- Dispersed flows, such as bubbly, droplet, and particle flows
- Stratified flows, such as free surface flows, or annular film flow in pipes.

A phase is considered dispersed if it occupies disconnected regions of space—otherwise it is continuous.

STAR-CCM+ provides five distinct models to meet the requirements of these two categories of flow:

- The Lagrangian Multiphase model: this model solves the equation of motion for representative parcels of the dispersed phase as they pass through the system. It is intended for systems that consist mainly of a single continuous phase carrying a relatively small volume of discrete particles, droplets, or bubbles. It is suited where the interaction of the discrete phase with physical boundaries is important.
- The Multiphase Segregated Flow model: this model is commonly known as the Eulerian Multiphase model in the literature, but that term has been given a wider significance in STAR-CCM+. The Multiphase Segregated Fluid model solves conservation equations for mass, momentum, and energy for each phase. Phase interaction models are provided to define the influence that one phase exerts upon another across the interfacial area between them.

- The Volume of Fluid (VOF) model: this model is provided for systems containing two or more immiscible fluid phases, where each phase constitutes a large structure within the system (such as typical free surface flows). This approach captures the movement of the interface between the fluid phases, and is often used for marine applications.
- The Discrete Element Model (DEM): this model is an extension of the Lagrangian Multiphase model, but where individual particles are modeled rather than representative parcels, and where inter-particle contact forces are explicitly accounted for.
- The Fluid Film model: this model predicts the dynamic characteristics of wall films using boundary layer approximations and assumed velocity and temperature profiles across the depth of the film. Film transport is predicted using thin shells that lie across the surface of solid walls on which the film is formed.

In mathematical terms, both the Volume of Fluid model and the Segregated Flow model utilize a Eulerian (as opposed to a Lagrangian) framework in their formulation. Therefore, setting up simulations for both these approaches involve use of distinct Eulerian phases.

What Is a Multiphase Mixture?

A multiphase mixture is a fluid that is composed of multiple phases. The corresponding node (Eulerian Multiphase > Mixture in the following screenshot) is activated as soon as the option Eulerian Multiphase is selected in the *Physics Model Selection* window.



The Multiphase Interaction model is added automatically in the *Physics Model Selection* dialog when the Eulerian Multiphase model is selected. The Eulerian Multiphase model is used to define phases and mixture properties, and the Multiphase Interaction model is used to define phase interactions.

Once these models are activated, you can choose one of three approaches for modeling the multiphase flow:

- The Volume of Fluid (VOF) homogenous multiphase model, which assumes a common velocity, pressure and temperature field for all phases.

- The Multiphase Segregated Flow model which solves separate conservation equations for each phase (apart from the pressure field, which is common).
- The Eulerian Multiphase Mixture model, which can be used with the Coupled Flow model or the Segregated Flow model.

If you select the Volume of Fluid or the Eulerian Multiphase Mixture model, the Eulerian Multiphase > Mixture node allows you to specify the bulk properties of the mixture. When using the Multiphase Segregated Flow approach, there are no settings to define under the Eulerian Multiphase > Mixture node.

Defining Multiphase Mixture Properties

To define the mixture for the Volume of Fluid (VOF) model:

1. Set up Eulerian phases to specify the materials and the equation of state.
2. Expand the Eulerian Multiphase node and its sub-node, Mixture, to access the Mixture Properties node.
3. Expand the Mixture Properties node.

   The calculation method for key properties such as Dynamic Viscosity is set by default to Volume Weighted Mixture and you do not normally change it.



What Is a Free Surface?

The interface between two immiscible fluids is called a free surface. The VOF multiphase model allows you to resolve the position and shape of the free surface, which is also important in evaluating surface tension effects.

What Is a Phase?

A phase is matter whose properties continuously change in space. Typically the word "properties" brings to mind the physical properties, and this phase definition is relevant for the VOF multiphase model. However a generalization of the concept of phases also includes "properties" like velocity, temperature, or characteristic length scale of the matter. For example, if a problem involves spherical water droplets which are very close to each other and have different diameters, and/or move with different velocities, and/or have different temperatures, they could also be considered different phases.

STAR-CCM+ uses the Eulerian Multiphase model to manage materials and equation of state specifications in multiphase mixtures.



This phase is not the same as a Lagrangian phase, which is restricted to dispersed multiphase flows that are modeled using the Lagrangian Multiphase model.

What Are Phase Conditions?

Some physics models expose their own region or boundary conditions for the phase to which they apply. In this case, a Phase Conditions manager node is added to the regions and boundaries for which the phase is active. For example, when you select a multi-component liquid in a Eulerian phase, STAR-CCM+ adds the Species Source Option within a phase node for regions that are associated with the phase.



In some circumstances, you may need to define different boundary conditions for a specific phase. To set these conditions, *phase conditions* are applied at the boundary level and are used to define phase-specific attributes. For example, in a Multiphase Segregated Flow case you can specify that a wall boundary is permeable for a particular phase, while all other phases still see a wall. Similarly, in a Lagrangian Multiphase flow case, the settings that are used to define a Lagrangian phase in the physics continuum can be overridden at a particular boundary by specifying different settings at the boundary level.

The Phase Conditions node has no properties and is subject to common menu items.



Defining Phase Interactions

The Multiphase Interaction model node represents The process of how Eulerian phases interact with each other in STAR-CCM+. This node contains a Phase Interactions sub-node for managing phase interaction objects. The phase interaction objects define the submodels that predict the interchange of conserved quantities between phases in a multiphase simulation. The Multiphase Interaction model is selected automatically after the Eulerian Multiphase model.



This section explains how to create phase interactions and select phase interaction models. The phase interaction models available for selection depend on the physics models that are selected in the physics continuum. For further information on the models available, refer to the following sections:

- Modeling phase interactions in Multiphase Segregated Flow
- Modeling phase interactions in VOF
- Modeling phase interactions in the Discrete Element Method
- Modeling phase interactions in Multiphase Mixtures
- Modeling phase interactions in Dispersed Multiphase

What Is a Phase Interaction?

A phase interaction is an object that defines the choice of models that are used to predict the influence of one phase upon another phase within a multiphase flow simulation. Define a phase interaction for every possible significant coupling.

Once created, phase interactions are stored under the Phase Interactions manager node. Each Phase Interaction node has properties and is subject to common menu items. It contains a Models manager node that is used to select phase interaction models for each phase pair.



For Eulerian multiphase flows, these interactions effectively provide the closures that are required for the conservation equations in each phase. It is only necessary to find closures (and hence to define phase interactions) where significant phase interactions are expected. For example, it can be possible to neglect interactions between a number of phases that are in dilute concentration, while retaining interactions between each dilute phase and a dominant phase. Thus, for a simulation where all N phases are strongly interacting, it would be necessary to define all possible N(N-1)/2

pairs. For multiple dilute phases that are carried in one dominant continuous phase, however, it can be necessary to define only N-1 phase interactions.

Standard phase interaction models are provided for the flow pattern where one dispersed phase (particles or bubbles) is carried in a continuous phase (liquid or gas). For modeling other kinds of multiphase flow patterns, you can define appropriate phase interaction models using field functions.

Phase Interaction Properties

Index	A unique identifier for this Phase Interaction (read-only).

Phase Interaction Manager Pop-Up Menu

This pop-up menu item is available when you right-click the Phase Interactions manager node.

New	Creates a phase interaction with a default name and a Models sub-node.

Creating a Phase Interaction

To create a phase interaction, right-click the Phase Interactions manager node and select New.



The new Phase Interaction node appears as a child of the Phase Interactions manager node, and contains a Models sub-node. It is given a default name that is based on the number of existing Phase Interaction nodes, if any.



Selecting Phase Interaction Models

The Models manager node, which has its own pop-up menu, contains the nodes of all the models in a phase interaction.



This section deals with the selection of models that is required to define a phase interaction. The primary tool for defining the phase interaction is the *Phase Interaction Model Selection* dialog, which presents the available models.

Using the Phase Interaction Model Selection Dialog

The *Phase Interaction Model Selection* dialog is used to select and manage phase interaction models for a simulation using the Multiphase Segregated Flow or the Volume of Fluid model.

To activate the dialog for a phase interaction, for example Phase Interaction 1:

1. Open the Phase Interaction 1 node and right-click on the Models sub-node.

2. In the pop-up menu, choose item Select Models...



The appearance of the *Phase Interaction Model Selection* dialog when accessed for the first time depends on the type of multiphase flow you are modeling. For a multiphase segregated flow, it appears as follows:



Models that have been selected are shown in the *Enabled Models* list on the right. Models that are still available to be selected are in the various group boxes that appear on the left.

The dialog is "live" in that the server is updated each time that you change the state of a checkbox. If you have the relevant nodes open, the change is reflected in the tree. In the screenshot below, for example, you have selected several models from the dialog, and their corresponding nodes are shown under the Phase Interaction parent node.



Managing Phase Interaction Models

The *Phase Model Selection* dialog is used to manage the selection of models subsequent to the initial selection. It can be used to view the models that are already chosen, and to modify that selection if desired.

Phase Interaction Models Manager Pop-Up Menu

The options that appear when right-clicking on the Models sub-node are:

Select Models... Opens the *Phase Interaction Model Selection* dialog.

See also: common menu items.

Using the Lagrangian Multiphase Model

A wide variety of flow processes involve the transport of solid particles, liquid droplets, or gas bubbles—known as *dispersed phases*—by a gaseous or liquid *continuous phase*. Examples include:

- Cyclone dust separators
- Vehicle soiling
- Spray coating
- Aerosol dispersion
- Spray cooling
- Liquid fuel combustion

The generic term *particle* is used in STAR-CCM+ to describe solid particles, liquid droplets, gas bubbles and massless (virtual) particles. If interaction of particles with the continuous phase, rather that with each other, dominate their motion, the dispersed phases are *dilute*. The Lagrangian Multiphase model is designed for such flows.

In STAR-CCM+, the Lagrangian Multiphase node is activated as soon as the model is selected in the *Physics Model Selection* dialog for a physics continuum. This node is shown below.



The documentation on the Lagrangian Multiphase model is divided into the following overall sections:

- Introducing Lagrangian Multiphase Modeling
- Working with Lagrangian Phase Models
- Working with Primary Atomization

- Working with Secondary Breakup
- Working with Lagrangian Phase Boundary Conditions
- Working with Injectors
- Working with Particle Tracks
- Load Balancing for Lagrangian/Eulerian Flows
- Understanding the Lagrangian Multiphase Solver
- Formulation
- Field Functions Reference

Introducing Lagrangian Multiphase Modeling

The physics continuum is a continuous phase whose governing equations are expressed in Eulerian form. The Lagrangian Multiphase model permits solving an arbitrary number of dispersed phases, each modeled in a Lagrangian framework [237]. In this framework, particle-like elements that are known as parcels are followed through the continuum. The state of each parcel is updated according to a selected set of models and can be optionally recorded as a track. Dispersed phases modeled using the Lagrangian Multiphase model are known as Lagrangian phases.

In the same way as models are selected for physics continua, the models that are provided to predict the state of each parcel can be selected per phase. Either (real) material particles or (virtual) massless particles can be modeled. Material models are available to represent single-component gases, single-, or multi-component liquids and single-component solids, together with constant density, or polynomial density equations of state. The momentum conservation equation governs the flow of material particles and can optionally include drag, virtual mass, and user-defined forces. Similarly, an energy conservation equation can be solved, including convective heat transfer and an optional user-defined energy source. Specific models for droplets in a gaseous medium ("sprays") include a mass conservation equation to account for the change in droplet mass from evaporation or condensation. They also include models for primary atomization and secondary breakup.

Boundary conditions define the interaction between particles and boundaries. Particle behavior at impermeable boundaries, in particular, requires modeling. You can control this behavior, depending on the active models. Particles can optionally be considered to cause erosion at these boundaries.

Injectors define the initial conditions of each particle, namely how and where they are introduced into the simulation.

In general, the state of the continuous phase influences the state of the dispersed phases. Two-way coupling can therefore be applicable, in which the state of the continuous phase depends on the dispersed phases, through inter-phase mass, momentum, and energy transfer effects.

If the flow of the continuous phase is laminar, a parcel that is released from a point at a given instant follows a smooth unique trajectory, that is, the motion is deterministic. On the other hand, parcels that are introduced into a turbulent carrier flow each have their own random path due to interaction with the fluctuating turbulent velocity field. A turbulent dispersion model is provided to account for this phenomenon.

What Is a Lagrangian Phase?

A Lagrangian phase in STAR-CCM+ is a dispersed phase (droplets, bubbles, or particles) modeled in a Lagrangian framework. The Lagrangian phase node, which has properties, represents this model. The Lagrangian phase node contains two manager nodes: a Models node that facilitates selection of phase models, and a Boundary Conditions node that facilitates setting of phase boundary conditions.



The definition of each Lagrangian phase includes the set of models applicable to its parcels. By definition, the same models govern particles in the same Lagrangian phase. Therefore, for example, particles of two different materials constitute two different Lagrangian phases.

Lagrangian Phase Properties

Id	A unique identifier for this phase (read-only).

What Is a Parcel?

For flows involving a comparatively small number of dispersed phases, it is possible to formulate and solve governing equations for every droplet, bubble, or particle. However, if the number of dispersed phases is large, a

statistical approach is more practical. In this approach, a smaller number of computational *parcels* represents the total population of dispersed phases. Each parcel represents a localized group (cluster) of dispersed phases having the same properties. In effect, parcels are a discretization of the population of dispersed phases in the same way that cells are a discretization of continuous space. As with cells, the number of parcels is not arbitrary; it must be large enough so that the properties of the full population of dispersed phases are represented. This number can be assessed, in the absence of any other measures, by performing calculations with different numbers of parcels and comparing the results.

Visualizing Lagrangian Phase Parcels

It is possible to visualize the parcels of a Lagrangian phase in a scene.



Parcels are available to be visualized in this way provided they have not left the computational domain, which usually means only in unsteady simulations.

As with any other parts in a simulation, you can manage the parcels through the scene, or through an individual displayer as shown in the following screenshot.



Using Lagrangian Phase Parcels

A parcel can also serve as an input part for certain derived parts: a cell surface, section, threshold, or iso surface. When a derived part has parcels that are specified as input parts, parcel fields can be visualized and analyzed on it, and parcel-specific reports can operate on it.

Parcel Interaction with Boundaries

A parcel centroid determines the intersection of a Lagrangian parcel with a boundary face. The parcel centroid is the center of mass of all the dispersed phases within the parcel.



STAR-CCM+ does not provide the distribution of the dispersed phases in a parcel.

Note: If you require a higher level of resolution for dispersed phase, use DEM.

Setting Up a Lagrangian Multiphase Simulation

Several actions are required to set up a simulation that uses the Lagrangian Multiphase model. These actions are outlined briefly in the following steps:

- Select the Lagrangian Multiphase Model
- Create one or more Lagrangian phases
- Select appropriate models for each Lagrangian phase
- Define material properties for each Lagrangian phase
- Define boundary conditions for each Lagrangian phase
- Create and define injectors
- Adjust parameters for the Lagrangian Multiphase solver

Selecting the Lagrangian Multiphase Model

The Lagrangian Multiphase model is activated from the *Physics Model Selection* dialog of the relevant physics continuum. It is necessary to first select a Liquid or Gas Material model in the dialog (which is accessed by right-clicking on the Models node of the physics continuum

and choosing item Select Models...). Continue by selecting either the Coupled Flow model or the Segregated Flow model. The Lagrangian Multiphase option then appears in the *Optional Models* group box.



The remaining selections in this dialog depend on problem characteristics that are not related to dispersed phases. When the Lagrangian Multiphase option has been selected in the dialog, the corresponding model node becomes visible in the object tree.



Creating One or More Lagrangian Phases

Having activated the Lagrangian Multiphase model itself, it is necessary to create one or more Lagrangian phases to represent each distinct dispersed phase within the simulation. To create a Lagrangian phase:

- Open the Models > Lagrangian Multiphase model node. This action exposes the Lagrangian Phases manager node.

- Right-click the Lagrangian Phases manager node and select New from the pop-up menu.



A new Lagrangian phase node appears with the default name, Phase 1.



The Lagrangian phase node can be renamed if desired. The Lagrangian phases manager can hold any number of phases up to the memory limits of the computer.

Lagrangian Phase Manager Pop-Up Menu

This pop-up menu item is available when you right-click the Lagrangian Phases manager node.

New	Creates a Lagrangian phase with a default name with a Models sub-node and a Boundary Conditions sub-node.

Copying a Lagrangian Phase

To copy the contents of a Lagrangian phase from one phase to another, use the clipboard or the drag-and-drop method. The copy action activates all the phase models from the source Lagrangian phase in the destination one, with the properties of each phase model also being copied.

This copy feature can be used between continua of the same simulation, provided the destination continuum already has the Lagrangian Multiphase model activated, with at least one Lagrangian phase created. Such copying is also possible between two different simulations, if both simulations have been opened in the same STAR-CCM+ session.

Additionally, you can right-click on the phase node and select Copy, and then right-click on the phase manager node and select Paste.

Selecting Appropriate Models for Each Lagrangian Phase

The Models manager node, which has its own pop-up menu, contains the nodes of all of the models in a Lagrangian phase.



This section deals with selecting the models that are required to define a Lagrangian phase. The primary tool for this selection is the Phase Model Selection dialog, which guides you through the model selection process and assists with model recommendations.

Lagrangian phase model selection is not identical to model selection for a continuum, since Lagrangian phases implicitly inherit certain model choices such as those for space, time and motion.

To select phase models for a Lagrangian phase:

- Right-click the Models node within the phase and choose Select Models... from the pop-up menu.



The *Phase Model Selection* dialog appears.



Lagrangian Phase Models Manager Pop-Up Menu

Select Models...	Activates selection of the Lagrangian phase models.

Using the Phase Model Selection Dialog

The *Phase Model Selection* dialog is used to select Lagrangian phase models, and change the selection in a simulation that uses Lagrangian Multiphase modeling. Lagrangian phase model selection and deletion can only be done through the dialog:

- This dialog is similar to a wizard in that it can assist you through a fairly complex process selection network.
- It is unlike a wizard in that it does not have the traditional Next and Back buttons.
- It does not restrict you to answering questions in a particular order.
- It has an auto select feature that adds recommended models automatically in response to a particular model selection by you. While this feature is on by default, you can deactivate it.

The *Phase Model Selection* dialog first appears as shown below for a new Lagrangian phase in a steady simulation.



For a steady simulation, Residence Time is selected automatically in the *Enabled Phase Models* group box. For an unsteady simulation, no initial selections are made.

Two types of model are available for selection: *Particle Type* and *Optional Models*. The <Optional> labels indicate that the selection of these items is not mandatory. The Auto-select recommended Phase models checkbox provides an option for the automatic selection of recommended models.

The *Phase Model Selection* dialog changes dynamically in response to previous selections. For instance if Residence Time and Material Particles are selected, it appears as shown below, where a selection from the *Material* group is required. The <Additional selections are required> label tells you that the model selection is incomplete. An incomplete model selection means that the Lagrangian phase cannot be initialized in its current state.



Trying to iterate when you have not completed the model selection results in an error similar to the exception shown below.



Selections that are made in this dialog can have an impact on the models that are required in the continuous phase. For example, selecting Quasi-Steady Evaporation for a Liquid particle in the *Phase Model Selection* dialog requires the presence of a multi-component gas in the continuous phase (that is, the physics continuum). If this selection has not been made in the continuous phase, the simulation fails to initialize, and an information dialog is displayed describing the continuous phase models that are missing.

You can check the compatibility of models that are chosen in the *Phase Model Selection* dialog with the models of the underlying continuum. Reopen the *Physics Model Selection* dialog for the physics continuum and look for any `<Select One>` indicators on the left-hand side. If these indicators are present, it is necessary to make further selections that satisfy the Lagrangian phase requirements. This process can involve deselecting existing models to expose appropriate group boxes on the left.

An example of a *Physics Model Selection* dialog with `<Select One>` indicators is shown below. In this example, Quasi-Steady Evaporation is selected for the Lagrangian phase without having first selected a multi-component gas in the physics continuum. To rectify this issue, it

would be necessary to clear models to the point where the Gas model is on the left-hand side. You can then select the Multi-Component Gas model from the appropriate group box.



More details on the *Phase Model Selection* dialog, as well as descriptions of the individual models available, are given in section, Working with Lagrangian Phase Models.

Using Auto-Select Recommended Phase Models

The Auto-select recommended Phase models checkbox is activated by default, which indicates that when there is a choice of models, the recommended ones are selected automatically.



Defining Material Properties for Each Lagrangian Phase

If a Material Particle model is chosen in the *Phase Model Selection* dialog, set the material properties, using the corresponding material node in the Lagrangian phase object tree. For example, the tree below has been expanded to show the Density and Specific Heat node for Lagrangian phase, Phase 1.



Further information on the material models available for a Lagrangian phase is given in section, Modeling Lagrangian Phase Materials.

Defining Boundary Conditions for Each Lagrangian Phase

How particles respond on impact with a solid boundary greatly depends on the properties of the particle. For this reason, each Lagrangian phase contains a Boundary Conditions manager node that defines the default response at solid boundaries for particles belonging to that phase.



It is possible to override the default behavior on a particular solid boundary using the Method property of the Physics Conditions > Lagrangian Specification node belonging to the boundary. In the example that is shown below, the Method property is being set to Specify for Boundary instead of Use Phase Defaults.



This change results in a Phase Conditions manager appearing for the boundary. The Phase Conditions node contains sub-nodes for each Lagrangian phase, and these sub-nodes can be used to customize the particle impact behavior for the boundary. An example is shown below.



Further information on specifying impact behavior is given in section, Working with Lagrangian Phase Boundary Conditions.

Creating and Defining Injectors

The locations at which particles enter the fluid continuum, and the manner in which they enter, are defined using injectors. Each injector generates new parcels for a given Lagrangian phase. Any number of injectors can be associated with the same Lagrangian phase.

Full information on creating and defining injectors is given in section, Working with Injectors.

Adjusting Parameters for the Lagrangian Multiphase Solver

The default settings of the Lagrangian Multiphase solver are a good compromise between accuracy, CPU time, and stability. The section, Understanding the Lagrangian Multiphase Solver, provides more detail on how accuracy and stability can be increased when necessary.

Working with Lagrangian Phase Models

A combination of Lagrangian phase models defines a Lagrangian phase. Each individual model represents one of the characteristics of that phase.



A Lagrangian phase model plays the same role in a Lagrangian phase as a conventional model does in a continuum. It defines how a physical phenomenon in a phase is represented, what the primary variables of the simulation are and what mathematical formulation is used to generate the solution.

The phase models available in STAR-CCM+ are:

Group Box	Phase Model
Particle Types	Massless Particles
	Material Particles
Material	Gas
	Liquid
	Multi-Component Coal
	Multi-Component Liquid
	Multi-Component Solid
	Solid
Equations of State	Constant Density
	Polynomial Density
Optional particle forces	Drag Force

Group Box	Phase Model
	Pressure Gradient Force
	User-Defined Body Force
	Virtual Mass
	Electrical Charge
Mass transfer	Quasi-Steady Single-Component Droplet Evaporation
	User-Defined Single-Component Droplet Evaporation
	Multi-Component Droplet Evaporation
	Coal Combustion
	Particle Reaction
Energy	Lagrangian Energy Model
Secondary Breakup	KHRT Breakup
	Reitz-Diwakar Breakup
	TAB Breakup
	SSD Breakup
Primary Atomization	LISA
Track Sampling	Track File
	Boundary Sampling
Wall Impingement	Bai-Gosman Wall Impingement
	Satoh Wall Impingement
	Ice Accretion
Species	Lagrangian Species
Optional Models	Lagrangian Energy Model
	Particle Radiation
	Turbulent Dispersion
	Residence Time
	Coal Fuel Nox
	Erosion
	Two-Way Coupling
	TAB Distortion

Group Box	Phase Model
	Passive Scalar
Optional Energy Sources	User-Defined Energy Source
Collision Detection	NTC Collision Model
Moisture Evaporation	Coal Moisture Evaporation
Char Oxidation	First-Order Char Oxidation
	Half-Order Char Oxidation
Raw Coal Devolatilization	Two-Step Devolatilization
	User-Defined Devolatilization

Modeling Massless Particles

This section describes what the Massless Particle model is and how to select it.

What is the Massless Particles Model?

The Massless Particles node represents the Massless Particles model. This phase does not represent a material; rather it represents massless virtual particles, sometimes known as fluid particles, which simply move with the continuous phase, without influencing each other or the continuous phase, see Massless Particle Equation of Motion, p. 4820.



It can be useful to include massless particles in a simulation to visualize or quantify:

- Turbulent diffusion of the continuous phase, with the Turbulent Dispersion model;
- The residence time of fluid in the continuous phase, with the Residence Time model;
- Mixing of fluid streams in the continuous phase, by associating a different massless Lagrangian phase with each fluid stream.

Properties

None

Selecting the Massless Particles Model

Select the Massless Particles option from the *Particle Type* group, as shown below.



The model selection dialog shows the Massless Particles model activated.



Modeling Material Particles

This section describes what the Material Particles model is and how to select it.

What Is the Material Particles Model?

The Material Particles node represents the Material Particles model. This phase represents real particles of a physical substance, such as sand, water droplets, or air bubbles.



Material particles have mass and volume, and physical conservation laws govern their behavior. In the most basic setup, these laws are conservation of mass and momentum. The particles are assumed to be internally homogenous and without internal motion.

Properties

None

Selecting the Material Particles Phase Model

Select the Material Particles option from the *Particle Type* group, as shown below.



The model selection dialog shows the activated Material Particles model, as well as various other modeling options made available by this selection.



Modeling Lagrangian Phase Materials

The section describes what Lagrangian phase Material Models are and how to select them.

What Are Lagrangian Phase Material Models?

Any Lagrangian phase that includes material particles also has its own material model, since the phase represents a distinct physical substance. (Lagrangian phases that use massless particles do not have their own material models.) There are currently six selectable material models for a Lagrangian phase:

- Gas, representing a single-component gas bubble
- Liquid, representing a single-component liquid droplet
- Multi-Component Coal, representing a coal particle
- Multi-Component Liquid, representing a multi-component liquid droplet
- Multi-Component Solid, representing a solid particle
- Solid, representing a single-component solid particle

Each of these material models behaves in an identical manner to their continuous phase counterpart. Appropriate links are provided in the individual sections below.

What Is the Single-Component Gas Bubble Model?

This material model node represents the Lagrangian phase material model for a single-component gas.



The object behaves in an identical manner to the Gas Model for the continuous phase. Its child-node is a single-component Gas material to which the conventions for managing a single-component material apply.

Properties

None

What Is the Single-Component Liquid Droplet Model?

This material model node represents the Lagrangian phase material model for a single-component liquid.



The object behaves in an identical manner to the Liquid Model for the continuous phase. Its child-node is a single-component Liquid material to which the conventions for managing a single-component material apply.

Properties

None

What Is the Multi-Component Coal Mixture Model?

This material model node represents the Lagrangian phase material model for a coal particle. The particle is a mixture that is composed of four constituent components: raw coal, char, ash, and H_2O. Raw coal, that is, the dry ash-free (daf) portion of the coal, consists of volatile matter (VM) and char. In coal combustion literature, char is also sometimes referred to as fixed carbon (C_s). You can either split the VM and char content present in raw coal and inject char separately, or VM and char can be injected together as raw coal. In the latter case, which is recommended, the char mass fraction specified at the injector would be zero and the char content in the domain would evolve due to the devolatilization process. The Coal Components

node manages these four constituents. Default properties for each of these components are supplied from the property database. Modify the properties according to the type of coal particle being simulated.



The proportions in which components are present for a particular phase of coal particles are set using the Species Mass Fraction node on the corresponding injector. Using this node is equivalent to setting the proximate analysis of a given coal mixture. If a particular component is not present, set its mass fraction to zero.

Material properties of each component can be set using the property sub-nodes within the Component Properties nodes for each of the four components. Relevant properties of the overall mixture can be set using the property nodes within the Mixture Properties manager node.

See the section Using Coal Combustion for details on using coal particles in a coal combustion simulation.

What Is the Single-Component Solid Particle Model?

This material model node represents the Lagrangian phase material model for a single-component solid.



The object behaves in an identical manner to the Solid Model for the continuous phase. Its child-node is a single-component Solid material to which the conventions for managing a single-component material apply.

Properties

None

What Is the Multi-Component Liquid Droplet Model?

This material model node represents the Lagrangian phase material model for a multi-component liquid.



The object behaves in an identical manner to the Multi-Component Liquid Model for the continuous phase. Its child-node is a Liquid Mixture material to which the conventions for managing multi-component materials apply.

If the Auto-selected recommended Phase models option has been ticked in the *Phase Model Selection* dialog, selecting the Multi-Component Liquid model results in the Species model also being selected.

Properties

None

What Is the Multi-Component Solid Model?

This material model node represents the Lagrangian phase material model for a multi-component solid.



Material properties of each component can be set using the property sub-nodes within the Component Properties nodes for each of the components. Relevant properties of the overall mixture can be set using the property nodes within the Mixture Properties manager node.

If the Auto-selected recommended Phase models option has been ticked in the *Phase Model Selection* dialog, selecting the Multi-Component Solid model results in the Species model also being selected.

See the section Using Particle Reaction for details on simulating the combustion of a compound of two solids.

Properties

None

Selecting a Lagrangian Phase Material Model

To expose the Lagrangian phase material model selections, activate the Material Particles model, then select the appropriate material model from the *Material* group, as shown below.



The model selection dialog shows the Lagrangian phase material model activated.



Modeling the Lagrangian Phase Equation of State

Any Lagrangian phase that includes material particles also requires an Equation of State model, to define how the density of the material is computed. This section explains the two equation of state models available for Lagrangian phases:

- What is the Constant Density Lagrangian Phase Model?
- Selecting the Constant Density Lagrangian Phase Model
- What is the Polynomial Density Lagrangian Phase Model?
- Selecting the Polynomial Density Lagrangian Phase Model

What Is the Constant Density Lagrangian Phase Model?

The Constant Density node represents the Constant Density Lagrangian phase model. This model is the same as the Constant Density model for the STAR-CCM+ simulation in general, except that it applies to the Lagrangian phase rather than throughout the continuum.



Properties

None

Selecting the Constant Density Lagrangian Phase Model

To expose the *Equation of State* model selections, activate a material model, then select the Constant Density option from the *Equation of State* group, as shown below.



The model selection dialog shows the Constant Density model activated.



What Is the Polynomial Density Lagrangian Phase Model?

The Polynomial Density node represents the Polynomial Density Lagrangian phase model. This model is the same as the Polynomial Density model for the STAR-CCM+ simulation in general, except that it applies to the Lagrangian phase rather than throughout the continuum.



Properties

None

Selecting the Polynomial Density Lagrangian Phase Model

To expose the Equation of State model selections, activate a material model, then select the Polynomial Density option from the *Equation of State* group, as shown below.



The model selection dialog shows the Polynomial Density model activated. In addition, the Energy option is selected automatically, provided you have the Auto-select recommended Phase models checkbox ticked.



Modeling Drag Force

This section describes the Drag Force model and how to select it.

What Is Drag Force?

The Drag Force node represents the Drag Force model. This models the force on a material particle due to its velocity relative to the continuous phase, Eqn. (1332).



Although the Drag Force node has no properties of its own, its sub-node Drag Coefficient does.

What Is the Drag Coefficient?

The Drag Coefficient node, which has properties, is a sub-node of the Drag Force node.



The drag coefficient is used to formulate the drag force, Eqn. (1332).

Drag Coefficient Properties

Dimensions — The dimensionality of the drag coefficient (read-only).

Method — Selects the method that is used for specifying the drag coefficient.

- Schiller-Naumann — Defines the drag coefficient using the Schiller-Naumann correlation. This correlation is available only when the continuous phase is viscous. It is suitable for spherical solid particles, liquid droplets, and small-diameter bubbles.
- Field Function — Defines the drag coefficient using a field function (typically user defined). A field function method node is added as a child to this node.
- Liu — Defines the coefficient using the Liu correlation. This correlation is only available when the TAB Distortion phase model has also been selected.

Selecting Drag Force

To expose the Optional Particle Forces model selections, activate the Material Particles model, then select the Drag Force option from the *Optional Particle Forces* group, as shown below.



The model selection dialog shows the Drag Force model activated.



Modeling Pressure Gradient Force

This section describes the Pressure Gradient Force model and how to select and deactivate it.

What Is Pressure Gradient Force?

The Pressure Gradient Force node represents the Pressure Gradient Force model. This model activates the pressure gradient forces acting on the surface of the particle, as given by Eqn. (1335).



Selecting Pressure Gradient Force

Whether the Auto-select recommended models option is activated or deactivated, the Pressure Gradient Force model is selected automatically when you select the Material Particles phase model.

You can deactivate the Pressure Gradient Force model from the *Enabled Models* group box.



The model selection dialog shows the Pressure Gradient Force model deactivated.



Modeling Virtual Mass

This section describes the Virtual Mass model and how to select it.

What Is Virtual Mass?

The Virtual Mass node, which has properties, represents the Virtual Mass model. This models the force that is exerted on a material particle as it accelerates the surrounding continuous phase, Eqn. (1336).



The name virtual mass comes from the fact that it is equivalent to increasing the particle mass; the terms *apparent* or *added* mass are sometimes also encountered. The virtual mass is proportional to the mass of the continuous phase the particle displaces. Therefore, virtual mass is more important when the continuous phase density is significant compared to the particle density. For example, it can be important for gas bubbles in a liquid.

Virtual Mass Properties

Cvm	The coefficient C_{vm} in Eqn. (1336)

Selecting Virtual Mass

To expose the Optional Particle Forces model selections, activate the Material Particles model, then select the Virtual Mass option from the *Optional Particle Forces* group, as shown below.



The model selection dialog shows the Virtual Mass model activated.



Modeling Charged Particles

Particles can enter the domain bearing electrical charges or, under some circumstances, can acquire electric charges after entering, resulting in additional forces acting on the particles. Charged particles appear in many applications, such as electrostatic spray painting and electrostatic precipitation in air pollution control.

To model particles that enter the domain already charged, STAR-CCM+ uses the Charged Particles model in combination with the Coulomb Force model.

To model particles that acquire a charge after entering the domain, STAR-CCM+ uses the Charged Particles model in combination with the Field Charging model.

These models are described in the following sections.

What is the Charged Particles Model?

The Charged Particle model is the base model that must be selected in order to simulate charge on particles. See Selecting Charge Models for the relationships to the Coulomb Force and Field Charging models.



This model initializes charge on the parcels and adds Particle Charge to the injector values. As the base model, this model is required by the other two models and does not need any electrical properties.

Once the particles have a charge, a force acts on them if they are in an electric field, using the Coulomb Force model. If the particle charge changes over time by field charging, use the Field Charging model.

The implementation of this model in the current version of STAR-CCM+ is subject to the following assumptions and restrictions:

- If the Two-Way Coupling model is not activated, the charge on the particles does not influence the electric field. This condition can be desirable to save CPU time, provided the charge on the particles makes a negligible contribution to the electric field distribution.

  The Two-Way Coupling model does not get selected automatically with Coulomb Force—it must be selected manually.
- If the Two-Way Coupling model is not activated, particles do not attract or repel each other based on their charge.
- Once initialized, the charge per unit mass of each particle is fixed unless the Field Charging model is selected. During droplet breakup, this value is conserved.

What Is the Coulomb Force Model?

Coulomb Force models the electrostatic force acting on charged particles. This force is modeled as $q\mathbf{E}$ where q is the particle charge and $\mathbf{E}$ is the electric field.

This model has no properties.

What Is the Field Charging Model?

Field Charging models the charging of particles by an electric field that is produced by a unipolar charge distribution in the carrier phase. It is the dominant mechanism of particle charging for particles larger than 1 micron.

When the Field Charging model is active, it requires the electrical permittivity property ε. This property applies to the material used in the Lagrangian phase. The model also requires inputs from the user in the form of the ion mobility and space charge density.

What Is the Ion Mobility?

Ion mobility K is the terminal velocity of ions through a given medium in response to a unit electric field, and is a factor in field charging. See Eqn. (1346).

Ion Mobility Properties

Dimensions — Dimensions of the units associated with the modeled quantity (read-only).

Method — Selects the method that is used for specifying the ion mobility, K in Eqn. (1346), from a subset of scalar profile options.

What Is the Ion Space Charge Density?

Ion space charge density ρ is the electrical charge from ions, per unit volume. It is a factor in field charging. See Eqn. (1346).

Ion Space Charge Density Properties

Dimensions — Dimensions of the units associated with the modeled quantity (read-only).

Method — Selects the method that is used for specifying the ion space charge density, ρ in Eqn. (1346), from a subset of scalar profile options.

Selecting Charge Models

To select one of the charged particle models:

- Observe the following prerequisites:
  - Make sure no mass transfer models are activated.
  - Make sure the Electrostatic Potential model is already selected in the physics continuum models.
  - Make sure the Material Particles model is already selected in the phase models.
- To select the Coulomb Force model, select Coulomb Force from the *Optional Particle Forces* group. The Charged Particles model is selected automatically.
- To select the Field Charging model, first select Charged Particles from the *Optional Models* group, then select Field Charging from the *Optional Models* group.

Modeling User-Defined Body Force

This section describes the User-Defined Body Force model and how to select it.

What Is User-Defined Body Force?

The User-Defined Body Force node represents the User-Defined Body Force model. This model allows you to customize the particle momentum equation by adding a term to it, Eqn. (1338).



Although the User-Defined Body Force node has no properties of its own, its sub-node User Body Force does.

What Is User Body Force?

The User Body Force node, which has properties, is a sub-node of the User-Defined Body Force node.



The User Body Force dictates how the user-defined body force is defined; see Eqn. (1338). It is a vector profile, with a reduced set of method options.

User Body Force Properties

Coordinate System	The coordinate system to use.	
	Cartesian	Specifies the Cartesian coordinate system.
	Cylindrical	Specifies the cylindrical coordinate system.
	Laboratory	Specifies the laboratory coordinate system.
	Spherical	Specifies the spherical coordinate system.
Method	Selects the method that is used for entering the force.	
	Composite	Defines the force as three (for example, X, Y, Z) components (or two (for example, X, Y) for two-dimensional problems).
	Constant	Specifies the force as a single three-part (or two-part for two-dimensional), comma-separated value. A Constant node is added as a child to this node.
	Field Function	Defines the force using a field function (typically user defined). A field function method node is added as a child to this node.
	Table (iteration)	Defines the force as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (r)	Defines the force as a function of radius. A Table (r) node is added as a child to this node.
	Table (time)	Defines the force as a function of physical time. A Table (time) node is added as a child to this node.
	Table (x,y,z)	Defines the force as a function of position. A Table (x,y,z) node is added as a child to this node.

Table (xyz, time)	Defines the force as a function of position and time. A Table (xyz, time) node is added as a child to this node.

Dimensions	The dimensionality of the force (read-only).

Selecting User-Defined Body Force

To expose the Optional Particle Forces model selections, activate the Material Particles model, then select the User-Defined Body Force option from the *Optional Particle Forces* group, as shown below.



The model selection dialog shows the User-Defined Body Force model activated.



Modeling Lagrangian Energy

This section describes the Lagrangian Energy model and how to select it.

What Is Lagrangian Energy?

The Energy node represents the Lagrangian Energy model. This model activates the energy conservation equation for the particles.



Although the Energy node has no properties of its own, its sub-node Heat Transfer Coefficient does.

What Is the Heat Transfer Coefficient?

The Heat Transfer Coefficient node, which has properties, is a sub-node of the Energy node.



The heat transfer coefficient is used to formulate the surface heat transfer, Eqn. (1355).

Heat Transfer Coefficient Properties

Dimensions — The dimensionality of the heat transfer coefficient (read-only).

Method — Selects the method that is used for specifying the heat transfer coefficient.

- Ranz-Marshall — Defines the heat transfer coefficient using the Ranz-Marshall correlation. This correlation is available only when the continuous phase is viscous. It is suitable for spherical particles up to a Reynolds number of ~5000.
- Field Function — Defines the heat transfer coefficient using a field function (typically user defined). A field function method node is added as a child to this node.

Selecting Lagrangian Energy

To expose the Lagrangian Energy model among the optional models, activate the Material Particles model in the Lagrangian phase. Next, select the Energy option from the *Optional Models* group, as shown below. This action creates an additional requirement that an Energy model is selected in the physics continuum.



The model selection dialog shows the Lagrangian Energy model activated.



If you select the Polynomial Density model, the Energy option is selected automatically provided you have Auto-select recommended Phase model activated.



Trying to iterate when you have not selected an Energy model in the physics continuum results in an error similar to what is shown below.

Modeling the User-Defined Energy Source

This section describes the User-Defined Energy Source model and how to select it.

What Is the User-Defined Energy Source?

The User-Defined Energy Source node represents the User-Defined Energy Source model. This model activates an optional user-defined source in the particle energy equation, Eqn. (1361).



Although the User-Defined Energy Source node has no properties of its own, its sub-node User Energy Source does.

What Is the User Energy Source?

The User Energy Source node is a sub-node of the User-Defined Energy Source node.



The User Energy Source dictates how the user-defined energy source is defined; see Eqn. (1361). It is a scalar profile, with a reduced set of method options.

Selecting the User-Defined Energy Source

To expose the Optional Energy Sources model selections, activate the Lagrangian Energy model, then select the User-Defined Energy Source option from the *Optional Energy Sources* group, as shown below.



The model selection dialog shows the User-Defined Energy Source model activated.



Modeling Droplet Collision with the NTC Model

This section describes the NTC (No Time Counter) collision detection model and how to select it.

What Is the NTC Collision Model?

The NTC Collision Model node represents the NTC collision detection model in the simulation tree.



NTC Collision Model Properties

None.

Selecting the NTC Collision Model

The NTC Collision model is available only in transient simulations.

This model is made available after a Liquid or Multi-Component Liquid material model is selected.

Select NTC Collision Model from the *Collision Detection* group box as shown below.



The model selection dialog shows the NTC Collision Model activated.



Modeling Quasi-Steady Single-Component Droplet Evaporation

This section describes the Quasi-Steady Single-Component Droplet Evaporation model and how to select it.

What Is Quasi-Steady Single-Component Droplet Evaporation?

The Quasi-Steady Evaporation node, which has properties, represents the Quasi-Steady Single-Component Droplet Evaporation model. This model adds the quasi-steady evaporation rate to the mass conservation equation for the droplets. This model is only available when the droplet material is a single-component liquid.



The Quasi-Steady Single-Component Droplet Evaporation model allows liquid particles (that is, droplets) to lose mass through evaporation. It also optionally allows droplets to gain mass through condensation. The driving force for evaporation and condensation is the departure from equilibrium of the liquid-vapor system. This system consists of the droplet and the gas component of the continuous phase that is nominated as the droplet vapor. This equilibrium is defined on an idealized phase diagram that is shown

below. Two material properties determine the vaporization line: the saturation pressure and the critical temperature, which are added to the material properties of the Lagrangian phases.



Besides the driving force, two other important parameters in determining the droplet evaporation rate are the molecular diffusivity of the gas component and the Sherwood Number of the droplet. The Quasi-Steady Evaporation node has a sub-node for the latter.

Additionally, you can delete particles from the system based on their mass, using the Depletion Criterion property of the model. For example, you can removed particles of negligible mass. This prevents waste of CPU time and improves stability as well. You can specify the minimum mass as a constant or by a field function.

When this model is active, the enthalpy difference method for specifying latent heat of vaporization becomes available.

Quasi-Steady Evaporation Properties

Vapor Component	The gas component in the continuous phase which represents the droplet vapor.
Condensation	Specifies whether condensation is permitted in addition to evaporation. Activated — Condensation is permitted. Deactivated — Condensation is not permitted.
Depletion Criterion	A user-set limit, expressed as a constant or field function, outside which a particle in the Lagrangian phase is depleted from the simulation. The mass of the depleted particle is removed from the simulation. The default criterion is zero, which deactivates depletion. Example: `$ParticleDiameter <` *[value]*

Selecting Quasi-Steady Droplet Evaporation

To expose the Quasi-Steady Evaporation model among the optional models, activate the Liquid model in the Lagrangian phase. Select the Quasi-Steady Evaporation option from the *Mass Transfer* group, as shown

below. This setting creates more requirements for models in the physics continuum, namely Multi-Component Gas, the Ideal Gas equation of state and an Energy model.



The model selection dialog shows the Quasi-Steady Evaporation model activated. In addition, the Energy option is selected automatically, provided you have Auto-select recommended Phase models checkbox activated.



Trying to iterate when you have not selected a Multi-Component Gas, Ideal Gas, or Energy model in the physics continuum results in an error similar to what is shown below.



Modeling User-Defined Single-Component Droplet Evaporation

This section describes the User-Defined Single-Component Droplet Evaporation model and how to select it.

What Is User-Defined Single-Component Droplet Evaporation?

The User-Defined Evaporation node, which has properties, represents the User-Defined Single-Component Droplet Evaporation model. This model supplies the rate of mass transfer in the mass conservation equation for the particles. This model is only available when the particle material is a single-component liquid.



The User-Defined Evaporation model allows you to define the evaporation rate using a constant, a table, or a field function. Make the evaporation rate positive and the units in kg/s. If the Two-Way Coupling model is active, the evaporated mass enters the continuous phase, provided the destination gas component has been set in the Vapor Component property.

Additionally, you can delete particles from the system based on their mass, using the Depletion Criterion property of the model. For example, you can removed particles of negligible mass. This prevents waste of CPU time and improves stability as well. You can specify the minimum mass as a constant or by a field function.

When this model is active, the enthalpy difference method for specifying latent heat of vaporization becomes available.

User-Defined Evaporation Properties

Vapor Component	The gas component in the continuous phase which represents the droplet vapor.
Depletion Criterion	A user-set limit, expressed as a constant or field function, outside which a particle in the Lagrangian phase is depleted from the simulation. The mass of the depleted particle is removed from the simulation. The default criterion is zero, which deactivates depletion. Example: `$ParticleDiameter` < *[value]*

What Is the User Evaporation Rate?

The User Evaporation Rate node, is a sub-node of the User-Defined Evaporation node.



The evaporation rate that this node provides is used in Eqn. (1319) as the rate of mass transfer. It is entered as a scalar profile.

User Evaporation Rate Properties

Method	Selects the method for entering the scalar value.
Dimensions	The dimensionality of the value (read-only).

Selecting User-Defined Droplet Evaporation

To expose the User-Defined Evaporation model among the optional models, activate the Liquid model in the Lagrangian phase. Select the User-Defined Evaporation option from the *Mass Transfer* group, as shown below. This setting creates more requirements for models in the physics continuum, namely Multi-Component Gas, and an Energy model.



The model selection dialog shows the User-Defined Evaporation model activated. In addition, the Energy option is selected automatically, provided you have ticked the Auto-select recommended Phase models checkbox.



Trying to iterate when you have not selected a Multi-Component Gas or Energy model in the physics continuum results in an error similar to what is shown below.



Modeling Multi-Component Droplet Evaporation

This section describes the Multi-Component Droplet Evaporation model and how to select it.

What Is Multi-Component Droplet Evaporation?

The Droplet Evaporation node, which has properties, represents the Multi-Component Droplet Evaporation model. This model adds the multi-component evaporation rate to the component mass conservation equations for the droplets. This model is only available when the droplet material is a multi-component liquid.



The Multi-Component Droplet Evaporation model allows some or all of the components within the droplet to lose mass through evaporation. It also optionally allows droplets to gain mass through condensation. The driving force for evaporation and condensation is the departure from equilibrium of the liquid-vapor system. This system consists of the droplet and the gas components that are nominated as exchanging mass, that is gas components can be nominated as representing the vapor of selected droplet components. The equilibrium state of each component is defined on an idealized phase diagram as for the equivalent single-component droplet evaporation model.

Besides the driving force, two other important parameters in determining the droplet evaporation rate are the Sherwood Number of the droplet and the molecular diffusivity of the gas. (The molecular diffusivity is assumed equal for all components.) The Droplet Evaporation node has a sub-node for the latter.

Additionally, you can delete particles from the system based on their mass, using the Depletion Criterion property of the model. For example, you can removed particles of negligible mass. This prevents waste of CPU time and improves stability as well. You can specify the minimum mass as a constant or by a field function.

When this model is active, the enthalpy difference method for specifying latent heat of vaporization becomes available.

Droplet Evaporation Properties

Condensation	Specifies whether condensation is permitted in addition to evaporation.	
	Activated	Condensation is permitted.
	Deactivated	Condensation is not permitted.
Vapor Components	The gas components in the continuous phase which represent the droplet vapors.	
Depletion Criterion	A user-set limit, expressed as a constant or field function, outside which a particle in the Lagrangian phase is depleted from the simulation. The mass of the depleted particle is removed from the simulation. The default criterion is zero, which deactivates depletion. Example: `$ParticleDiameter` < *[value]*	

Mapping Droplet Components to Gas Components

The Vapor Components property in the *Properties Window* for the Droplet Evaporation node defines how droplet components are mapped to gas components in the continuous phase. The mapping is defined as follows:

- Click the ... (custom editor) button to the right of the Vapor Components property.

The *Droplet Evaporation - Vapor Components* dialog appears.



The liquid components in the droplet are listed on the right side of the dialog. These components can be mapped to available components in the continuous gas phase by using the drop-down menus next to each droplet component as shown below.



- Click OK to close the dialog when all components have been mapped.

Selecting Multi-Component Droplet Evaporation

To expose the Droplet Evaporation model among the optional models, activate the Multi-Component Liquid model in the Lagrangian phase. Select the Droplet Evaporation option from the *Optional Models* group, as shown below. This setting creates more requirements for models in the physics continuum, namely Multi-Component Gas, and the Ideal Gas equation of state.



The model selection dialog shows the Droplet Evaporation model activated. In addition, the Energy option is selected automatically, provided you have ticked the Auto-select recommended Phase models checkbox.



Trying to iterate when you have not selected a Multi-Component Gas, Ideal Gas, or Energy model in the physics continuum results in an error similar to what is shown below.



Modeling Lagrangian Species

This section describes the Lagrangian Species model and how to select it.

What Is the Lagrangian Species Model?

The Species node represents the Lagrangian Species model. This model is required when a multi-component material model is selected for the Lagrangian phase. It indicates that the multi-component particles consist of a discrete set of components.



When the Species model is active, the mass fraction of each component can describe the composition of the particles. To calculate the mass fraction, a mass conservation equation is solved for each component.

Selecting the Lagrangian Species Model

The Species option is automatically selected in the *Phase Model Dialog* when a multi-component material model is selected (if the Auto-select recommended Phase models option is activated). It is not exposed unless a

multi-component material model, such as the Multi-Component Liquid model, is selected. The *Phase Model Selection* dialog appears as shown below when the Lagrangian Species model is activated.



Modeling Two-Way Coupling

This section describes the Two-Way Coupling model and how to select it.

What Is Two-Way Coupling?

The Two-Way Coupling node represents the Two-Way Coupling model. This model indicates that the Lagrangian phase can exchange mass, momentum, and energy with the continuous phase, depending on the active Lagrangian phase models.



Without the Two-Way Coupling model, the default "one-way coupling" means that the continuous phase influences the particles through terms such as drag in the momentum equation and heat transfer in the energy equation. However, there is no reverse effect. With the Two-Way Coupling model, the reverse effect is accounted for, and Lagrangian source terms appear in continuous phase equations. Typically, this effect becomes more important at higher particle loadings (see [237], for example).

Another use for the Two-Way Coupling model is with Coulomb Force. When both these models are activated, the electric charge on the particles influences the electric field.

Consistent with the initial assumption of a dilute phase, particles appear to the continuous phase as point sources. That is, they contribute source terms (provided this model is active) but are not seen to occupy any volume. This assumption is likely to be reasonable for volume fractions less than ~0.1. Take care with volume fractions above this level: results are likely to suffer from shortcomings in the physical modeling, not only because the point-source approximation can become invalid but, for example, because inter-particle collisions are no longer negligible.

Properties

None

Selecting Two-Way Coupling

To expose the Two-Way Coupling model among the optional models, activate the Material Particles model; massless particles are one-way coupled by definition. Select the Two-Way Coupling option from the *Optional Models* group, as shown below.



The model selection dialog shows the Two-Way Coupling model activated.



Modeling Turbulent Dispersion

This section describes the Turbulent Dispersion model and how to select it.

What Is Turbulent Dispersion?

The Turbulent Dispersion node represents the Turbulent Dispersion model, which collaborates with the Reynolds-Averaged Navier Stokes turbulence model in the physics continuum.



This model employs a random-walk technique to synthesize the fluctuating nature of the turbulent velocity field in the continuous phase. The introduction of a random element means that turbulent flows are likely to need a higher number of parcels than similar laminar flows, to provide statistically meaningful results.

Properties

Exact Eddy Interaction Time

With this property active, the model restricts sub-step durations so that they exactly sum to the eddy interaction time. Although this restriction improves accuracy, computational speed is reduced.
Without this property, the model does not allow the sub-step durations to become less than the limit implied by the minimum Courant number. (See Lagrangian Multiphase Solver Properties.) However, the model still aims for the closest possible approximation to the exact sum.

Activated	Sub-steps sum exactly to interaction time.
Deactivated	Sub-steps sum to a close approximation of interaction time. This is the default.

Selecting Turbulent Dispersion

Select the Turbulent Dispersion option from the *Optional Models* group, as shown below.



The model selection dialog shows the Turbulent Dispersion model activated.



Trying to iterate when you have not selected a Reynolds-Averaged Navier Stokes turbulence model in the physics continuum results in an error similar to what shown below.



Modeling Erosion

This section describes the Erosion model and how to select it.

What Is the Erosion Model?

The Erosion node represents the Erosion model in the simulation tree.



The erosion model predicts the rate of erosion from particle impact on solid walls. Erosion rate is defined as the mass of wall material removed per unit area per unit time, and is provided as a field function when the Erosion model is activated.

To calculate the erosion rate, it is necessary to accumulate the damage that each particle impact causes. This calculation is done by selecting a correlation for the erosion ratio, that is, the mass of wall material eroded per unit mass of impinging particles. Erosion damage depends on both the material of the walls and the material of the particles. Hence, the erosion ratio correlation is a boundary condition which can optionally be specified differently for each boundary. Further details are given in the formulation.

Any correlation for the erosion ratio is specific to a particular wall material, particle material, and carrier fluid. It can also be specific to a particular range of flow conditions. Correlations that are used outside their intended application can be inaccurate. For this reason, take care to ensure that a suitable correlation is chosen and its coefficients are specified accordingly. It is prudent to compare the predictions of a number of suitably chosen correlations, unless there is a strong reason to trust a particular one.

It is also important to appreciate that erosion modeling does not necessarily predict material loss with a consistent degree of accuracy and it is not recommended that you estimate component lifetimes using CFD modeling alone. If estimating component lifetime is important, then additional experimental validation should be obtained to assess the effectiveness of the erosion model under appropriate conditions.

Properties

None

Selecting the Erosion Model

To expose the Erosion model among the optional models, activate the Material Particles model. Although erosion is more often associated with solid particles, this association is not enforced. Select the Erosion option from the *Optional Models* group, as shown below.



The model selection dialog shows the Erosion model activated.



Modeling Residence Time

This section describes the Residence Time model and how to select it.

What Is the Residence Time Model?

The Residence Time node represents the Residence Time model. This model makes particle residence time available as a field function.



Properties

None

Selecting the Residence Time Model

To select the Residence Time model, open the node of the Lagrangian phase that you wish to configure. Select the Select models... item from the pop-up menu.



The *Phase Model Selection* dialog appears. This option can be selected automatically when the phase was created. Otherwise you can click its radio button.



Modeling Particle Radiation

This section describes the Particle Radiation model and how to select it.

What Is the Particle Radiation Model?

The Particle Radiation model node, which has properties, represents the Particle Radiation model for a Lagrangian phase.



The Particle Radiation model accounts for the scattering and absorption of thermal radiation due to particles when the Participating Media Radiation model is active in the continuous phase. It also calculates the radiation that is emitted from the particles themselves. Two properties, the absorption efficiency and the scattering efficiency, govern absorption and scattering for each Lagrangian phase. These properties are specified without respect to wavelength, and so even if surface and continuum radiative properties are spectral, the particle properties are gray. Scattering is modeled as isotropic. You can use the particle radiative energy source field function to post-process radiative contribution from particle phases.

Particle Radiation Properties

Qap	Particle absorption efficiency, Q_a in Eqn. (883). This property does not change with wavelength.
Qsp	Particle scattering efficiency, Q_s in Eqn. (884). This property does not change with wavelength.

Selecting the Particle Radiation Model

To activate the Particle Radiation model, first select the Energy model from the *Optional Models* group, as the Particle Radiation model depends on it. The Particle Radiation model becomes available for selection in the *Optional*

Models group as shown below. Selecting this model requires that an Energy model is selected in the continuous phase, along with the Participating Media Radiation model.



The Particle Radiation model appears under the *Enabled Phase Models* group when selected.



Trying to iterate without having selected the appropriate Energy or Participating Radiation models in the continuous phase results in an error message similar to the following:



Modeling TAB Distortion

This section describes the TAB Distortion model and how to select it.

What Is the TAB Distortion Model?

The TAB Distortion model is used to calculate the distortion of liquid droplets under the action of aerodynamic forces. The TAB Distortion node represents this model in the Lagrangian phase model tree.



As discussed in the formulation, this model uses the so-called Taylor analogy [252]. The analogy represents a distorting droplet as a damped spring-mass system; it considers only the fundamental mode of oscillation of the droplet. The displacement and velocity of the mass in the spring-mass system correspond to representative distortion and rate of distortion quantities for the droplet. These quantities are made available as the TAB Distortion and TAB Distortion Rate field functions. While these quantities are sometimes of interest themselves, the main purpose of the TAB Distortion model is to make these quantities available to other models, such as:

- The TAB Breakup model, which uses the TAB Distortion to determine when and how droplets breakup under the action of aerodynamic forces.
- The Liu drag coefficient method, which uses the TAB Distortion to modify the drag coefficient of distorted droplets consistent with their non-spherical form. If the drag model is active, activating the TAB distortion model also activates the Liu drag coefficient method.

TAB Distortion Properties

WeCrit	Critical Weber number, see Eqn. (1420)
Cd	Damping coefficient, see Eqn. (1420)
Ck	Stiffness coefficient, see Eqn. (1420)

What Is the Sherwood Number?

The Sherwood Number node, which has properties, is a sub-node of both the Quasi-Steady Evaporation node and the multi-component Droplet Evaporation Node. The Quasi-Steady Evaporation node is shown below



The Sherwood number is used to formulate the quasi-steady evaporation rate for single-component droplets and multi-component droplets.

Sherwood Number Properties

Method	Selects the method for specifying the Sherwood number.	
	Ranz-Marshall	Defines the Sherwood number using the Ranz-Marshall correlation. This correlation is available only when the continuous phase is viscous. It is suitable for spherical particles up to a Reynolds number of ~5000.
	Field Function	Defines the Sherwood number using a field function (typically user defined). A field function method node is added as a child to this node.

Modeling Lagrangian Passive Scalars

In the Lagrangian framework, passive scalars are user-defined variables of arbitrary value, assigned to fluid phases or individual particles. They are passive because they do not affect the physical properties of the simulation. An intuitive way to think of passive scalars is as marker paint on a particle or parcel, but with numerical values instead of colors, and with no appreciable mass or volume.

More elaborate physical effects can be modeled with passive scalars on particles, such as change of composition through radioactive decay.

You can use more than one passive scalar in a simulation, for example when analyzing the mixing of multiple streams of particles with the same properties. Each particle can be tagged with a scalar indicating its source. Each particle can also be tagged with a scalar that changes steadily over time, allowing you to track residence time in different regions of the simulation.

Lagrangian passive scalars can also interact with Eulerian passive scalars, to allow transfer of scalar values from one to the other.

The value of a passive scalar on a particle is set on the injector from which the particle enters the simulation. As a result, particles of the same phase can be given different passive scalars at different injectors. See the Passive Scalar entry in the Injector Values Reference.

For uses outside the scope of Lagrangian modeling, see Modeling Passive Scalar.

Selecting the Lagrangian Passive Scalar Model

To select the Passive Scalar model in a Lagrangian physics continuum:

- Select the Lagrangian multiphase model.
- Under the Lagrangian Phases node, create phase nodes for each Lagrangian phase.
- Select models for each Lagrangian Phase.
- Select the Passive Scalar model from the *Optional Models* group box.

  The Passive Scalar radio button becomes available for selection under *Optional Models* for each phase.



- Within the model tree for the Lagrangian phase, right-click the Passive Scalar > Passive Scalars node and choose New.
- Set the properties of the newly created Passive Scalar node. See The Lagrangian Passive Scalar Node.

The Lagrangian Passive Scalar Node

The Passive Scalar model node has a child node, the Passive Scalars manager, which allows you to organize multiple passive scalars and their properties. This node has properties and is subject to common menu items.

The node appears as follows:



Passive Scalars Properties

Source Active — Creates a new node, Passive Scalar Source, under the Passive Scalar model node, used to specify a source term.

- Activated — Use a source term.
- Deactivated — Use no source term. The default.

Managing Lagrangian Passive Scalars

A Passive Scalar node under the Passive Scalars manager node represents a passive scalar in the simulation tree.



Individual passive scalar nodes can be renamed or deleted.

Passive Scalar Properties

Number	A property that indicates the number of passive scalar materials preceding the current one. For example, for the first passive scalar material that you create, this property would equal zero. These values remain unchanged even when you delete one or more of these objects.
Maximum clip value	Sets the maximum scalar value used for clipping.
Minimum clip value	Sets the minimum scalar value used for clipping.

Clip	Clipping prevents scalar values from exceeding limits you set.	
	Off	No clipping applied.
	Below Min, Above Max	Prevents scalar values from rising above the Maximum clip value, and falling below the Minimum clip value.
	Below Min	Prevents scalar values from falling below the Minimum clip value.
	Above Max	Prevents scalar values from rising above the Maximum clip value.
Transport	Provides the option to deactivate the convection and/or diffusion terms in Eqn. (164) and Eqn. (165). The source term is always active. The transient term is only active for transient simulations.	
	Convection and Diffusion	Both diffusion and convection are activated.
	Convection only	The diffusion term is deactivated.
	Diffusion only	The convection term is deactivated.
	No Convection or Diffusion	Both convection and diffusion terms are deactivated.
Use GGDH	Uses the Generalized Gradient Diffusion Hypothesis. A Reynolds Stress Turbulence model must be activated first.	
	Activated	Use GDDH in calculating passive scalar transport.
	Deactivated	Do not use GDDH. The default.

The Lagrangian Passive Scalar Source Node

The optional Passive Scalar Source node is activated by the Passive Scalar property Source Active. The node specifies source terms for the scalars.



Properties

Method Selects the method used for the source terms of the individual passive scalars.

- **Constant** Applies the constant vector of values to all individual passive scalars as their source term.
- **Composite** Allows you to determine a different type and vector of source term for each individual passive scalar.
- **Injection Table** Set the value of the passive scalars according to an injection table. The table specifies the location of injection points, the time of injection, and the value of each passive scalar separately. Used only with table injectors.

Dimensions The dimensionality of the value, Mass/Volume-Time (read-only). The passive scalar itself is always dimensionless, but its source term has units referring to the density of the receiving phase over time.

See the formulation section.

Working with Lagrangian Phase Interactions

Lagrangian phases can interact both with other Lagrangian phases and with phases of other types using Multiphase Interaction models.

Lagrangian/Eulerian Passive Scalar Interactions

Both Lagrangian and Eulerian frameworks give the option of using passive scalars as user-defined variables of arbitrary value, assigned to fluid phases or individual particles. They are passive because they do not affect the physical properties of the simulation. An intuitive way to think of passive scalars is as marker paint on a particle or parcel, or marker dye in a fluid, but with numerical values instead of colors, and with no appreciable mass or volume.

The Passive Scalar Interaction model allows transfer of values between a Lagrangian and an Eulerian passive scalar. It becomes available when a Lagrangian and an Eulerian phase both have passive scalars defined. When you select it, the Multiphase Interaction node also appears, if it was not already present. The Passive Scalar Interaction model has no properties.

To set up an interaction between Lagrangian and Eulerian passive scalars:

- Left-click on the Phase Interactions node under Multiphase Interaction and select New. This adds a Phase Interaction node, with a Models node beneath it.
- Left-click on the Models node and select the Select models... item.

- The *Phase Interaction Model Selection* dialog opens. Select Passive Scalars Interaction and either Area Weighted or Volume Weighted. Nodes for these appear under the Models node. See the formulation section for a description of area and volume weighting.



- Open the pull-down menu of the Passive Scalar Interaction node, and select the Lagrangian phase to be involved in the interaction.



- The properties of the Area Weighted or Volume Weighted nodes list the two interacting phases. Use the pull-down menus for each phase to assign a passive scalar to the phase.



- The Area Weighted node has the nodes Coefficient 1 to 2 and Coefficient 2 to 1 under it. The Volume Weighted node as a single Coefficient node under it.



Set the Method for determining the value of each coefficient, using the pull-down menu in the coefficient properties.



Working with Primary Atomization

Primary atomization refers to the process of forcing liquid through a small orifice at a high pressure, resulting in a fine spray of liquid droplets. STAR-CCM+ currently provides a primary atomization model that simulates the behavior observed in pressure-swirl atomizers.

What Is the Pressure-Swirl Atomizer?

The pressure-swirl atomizer, also commonly known as a simplex atomizer, is widely used for liquid-fuel injection in gas turbines, oil furnaces, and direct-injection spark-ignited automobile engines. In this atomizer type, a swirling motion is imparted to the injected liquid through nozzles that are known as swirl ports. The liquid accelerates as it flows through the swirl ports into a central swirl chamber. Under the action of centrifugal force, it spreads out in the form of a conical sheet and develops a hollow air core. It

emerges from the orifice as an unstable thinning liquid sheet, then breaking up into ligaments and then droplets. A sketch of a hollow-cone pressure swirl injector is shown in the figure below.



In the illustration above, three stages in the atomization process are identified: film formation, sheet breakup, and atomization. δ_o is the initial film thickness in the nozzle; $2h_0$ is the initial thickness of the film as it exits the nozzle, and $2h_b$ is the thickness of the film at the point of breakup. θ is the spray angle, which is defined as $0.5 \times (\text{inner cone angle} + \text{outer cone angle})$.

Using the Pressure Swirl Atomizer

Two steps are required to implement the pressure swirl atomizer in STAR-CCM+. First select the Linearized Instability Sheet Atomization (LISA) model for the Lagrangian phase being injected through the atomizer. Then choose the Pressure Swirl injector type for the injector representing the atomizer.

What Is the Linearized Instability Sheet Atomization Model?

The Linearized Instability Sheet Atomization (LISA) model is a primary atomization model that represents the action of a pressure swirl atomizer. This model corresponds to the LISA model node in the Lagrangian phase

model tree, and selecting it activates the Pressure Swirl Injector type for injectors. The model generates initial size and velocity values for droplets entering through pressure swirl injectors active for the Lagrangian phase, as described in the formulation.



LISA Model Properties

None

Selecting the LISA Model

Select the LISA model from the *Primary Atomization* group box of the *Phase Model Selection* dialog.



This model becomes available when a gas model (single- or multi-component) is selected. It is not compatible with the Multiphase Mixture model.

Choosing this option makes the Pressure Swirl Injector option available in the Type property list for all injectors present in the simulation.



The LISA model uses the parameters that are set on pressure swirl injectors to predict primary atomization.

Working with Secondary Breakup

Liquid droplets sometimes break up under the action of non-uniform surface forces, a phenomenon that is known as secondary breakup. The response of droplets to non-uniform surface forces is invariably to deform, with the deformation resisted by surface tension and viscous forces inside the droplet. Characteristic measures of this behavior are the Weber and Ohnesorge numbers

$$\mathrm{We} = \frac{\rho_g |\mathbf{v}_s|^2 D_p}{\sigma} \tag{1311}$$

$$\mathrm{Oh} = \frac{\mu_l}{\sqrt{\rho_l D_p \sigma}}. \tag{1312}$$

Depending on the magnitude of these forces, various breakup regimes have been identified, each characterized by the shape of the deforming droplets. An example categorization (after Stiesch [272]) for low Oh is given below.

Category	Weber Number
• Vibrational breakup	~ 12
• Bag breakup	< 20
• Bag/ streamer breakup	< 50
• Stripping breakup	< 100
• Catastrophic breakup	> 100



Other categorizations identify more or fewer regimes; the above is representative.

Detailed modeling of even one breakup regime is difficult. At most, the goal of secondary breakup models is to predict when breakup occurs and what diameters result from it. A selection of models is provided in STAR-CCM+:

- KHRT (stripping, catastrophic)
- Reitz-Diwakar (bag, stripping)
- TAB (vibrational)
- SSD (catastrophic)

The most appropriate regime for each model is noted in parentheses. However it is not unusual to find models that are applied outside these regimes. Where possible, results from more than one breakup model can be compared to add confidence.

These models become available when a gas model (single- or multi-component) is selected. It is not compatible with the Multiphase Mixture model.

What Is the KHRT Droplet Breakup Model?

The KHRT Breakup model combines two submodels, one based on Kelvin-Helmholtz (KH) theory and one based on Rayleigh-Taylor (RT) theory. The KHRT Breakup node represents this model in the Lagrangian phase model tree.



Both breakup submodels consider the growth of instabilities on a droplet and provide expressions for their wavelength and frequency, as discussed in the formulation. Kelvin-Helmholtz instabilities are due to the slip velocity of the droplet, which eventually shears small child droplets off the parent [258], corresponding to the stripping regime. Rayleigh-Taylor instabilities are due to the acceleration of the droplet and tend to shatter the droplet completely [255], corresponding to the catastrophic regime. The KH and RT submodels compete: instabilities due to both can grow simultaneously; if they grow for long enough, breakup occurs due to the RT instabilities. Otherwise KH breakup occurs.

Because there is a tendency for the KH submodel to generate a bi-modal size distribution, the smaller child droplets are accumulated into a new parcel; the parent parcel retains the larger parent droplets. However, to prevent too many child parcels being created, these parcels are only released when their mass reaches a predefined proportion of the original mass of the parent.

Child parcels are given a lateral velocity proportional to the kinetic energy of the droplet oscillation at the instant of breakup. This velocity tends to generate a spreading effect from an injector which, to some extent, makes a cone injector unnecessary.

The KHRT model was developed as a replacement for the TAB model [255]. Although the latter can still be preferred at low Weber numbers, the KHRT model is preferred at higher Weber numbers. The principal difficulty in using the model is tuning the coefficients to the application under consideration. This difficulty reflects the complexity of the physics and is not due to a flaw in the model.

KHRT Breakup Properties

Child Parcels	Number of child parcels to create during a breakup event.

KHRT Breakup Expert Properties

Normal Velocity Coefficient	Normal velocity coefficient for child parcels, A_1 in Eqn. (1426)
B0	KH length coefficient, Eqn. (1424)
B1	KH time coefficient, Eqn. (1423)
C3	RT length coefficient, Eqn. (1429)
Ctau	RT time coefficient, Eqn. (1428)
Maximum fractional shed mass	Maximum fraction of the parcel mass that is shed before the child parcels are created.

What Is the Reitz-Diwakar Droplet Breakup Model?

The Reitz-Diwakar breakup model is based on observed length- and time-scales of droplet breakup. The Reitz-Diwakar Breakup node represents this model in the Lagrangian phase model tree.



As discussed in the formulation, the Reitz-Diwakar breakup model, [259] and [260], assumes that breakup occurs in one of two possible modes

- "bag" breakup, in which the non-uniform pressure field around the droplet causes it to expand in the low-pressure wake region and eventually disintegrate when surface tension forces are overcome.
- "stripping" breakup, in which liquid is sheared or stripped from the droplet surface.

The model is simple and does not generate child parcels, hence it is relatively cheap to run. It can be used to estimate the importance of breakup before turning to a more complex model, or to provide support to conclusions obtained from such a model.

Reitz-Diwakar Properties

None

Reitz-Diwaker Expert Properties

WeCrit	Minimum Weber number for breakup, see Eqn. (1431).
Cb2	Bag breakup time-scale coefficient, see Eqn. (1432).
Cs1	Coefficient prescribing the onset of stripping breakup, see Eqn. (1433).
Cs2	Stripping breakup time-scale coefficient, see Eqn. (1434).

What Is the TAB Droplet Breakup Model?

The TAB breakup model is an extension of the TAB distortion model which assumes that breakup occurs when the TAB distortion reaches a prescribed magnitude. The TAB Breakup node represents this model in the Lagrangian phase model tree.



The TAB breakup model begins with the assumption that breakup occurs when the TAB distortion, as calculated by the TAB distortion model, exceeds unity. Breakup replaces the parent particles with child particles whose diameter is chosen from a Rosin-Rammler distribution (see the formulation for equations). By default, the breakup event creates no new parcels—the original parcel is retained, but the particle diameter change to the new child value. New parcels can optionally be created however, in which case the droplet diameter in each is chosen independently. This

feature can be useful in modeling the generation of a distribution of droplet sizes through breakup. Parcels that are involved in a breakup event are given a lateral velocity proportional to the kinetic energy of the droplet oscillation at the instant of breakup. This velocity tends to generate a spreading effect from an injector which, to some extent, makes a cone injector unnecessary.

Despite being based on a single mode of oscillation in the vibrational regime, the TAB model reproduces the same characteristic time-scales in low and high Weber number limits as the Reitz-Diwakar model. Typically, however, the TAB model is used at low Weber numbers; hollow-cone gasoline sprays are an example of a preferred application for this model [231]. Outside its range of validity, the model tends to underpredict droplet sizes.

TAB Breakup Properties

Child Parcels	Number of child parcels to create during a breakup event.

TAB Breakup Expert Properties

Cv	Normal velocity coefficient, see Eqn. (1436)
K	Energy ratio, see Eqn. (1435)
Rosin-Rammler Exponent	Exponent in Rosin-Rammler size distribution
Cb	Critical distortion coefficient, see Eqn. (1419)

What is the SSD Breakup Model?

In the Stochastic Secondary Droplet (SSD) breakup model, breakup is a random process unrelated to original droplet size. This results in a log normal distribution of droplet diameters in the long time limit, as discussed in the formulation. The SSD Breakup node represents this model in the simulation tree.



The particles interact with the Eulerian fluid phase as they travel. Particles break up when three criteria are met:

- The particle radius is larger than the critical radius (calculated from local droplet conditions using Eqn. (1437)).
- The Weber number (calculated from local droplet conditions) is larger than a critical Weber number (set in the interface).
- The particle's local time (the accumulated residence time) is larger than the breakup time (Eqn. (1438)). The model tracks parcels as they cross the domain and compares the residence time with a locally evaluated breakup time. The residence time is only calculated on parcels with a radius larger than a critical radius and a Weber number larger than a critical Weber number.

When the particle breaks up, the local time is reset for the new parcels.

You can control the average number of particles per child parcel and the maximum number of child parcels formed in each breakup event using the Target Count and Maximum Child Parcels properties. The model uses these values when calculating the size, number of parcels, and particles per parcel (parcel count) of the child parcels. This gives you some control over the description of the spray.

Note that neither Target Count nor Maximum Child Parcels changes the physics of the breakup event, since these parameters do not change the diameter distribution functions. Rather, these parameters influence the total number of parcels in the calculation and thus the relative expense of the computation. A larger number of samples usually means an improved representation of the drop size distribution.

This model differs from other breakup models in three ways:

- The parent drop no longer exists after the breakup event. (The particle numbering system re-numbers after the breakup, so a particle before the event with the same number as a particle after the event is not the same particle.)
- Because the model works with a user-specified particle count, you seldom have very large particle counts that are difficult to handle numerically.
- The breakup distribution function, $\Phi(r)$ in Eqn. (1439), is a function of local conditions and depends only weakly on the diameter of the parent. Though the distribution is limited to the range between the current radius and half the critical radius, the function depends more on the Weber number, which is a strong function of local flow.

SSD Breakup Properties

WeCrit	The critical Weber number, We_{cr} in Eqn. (1437). The larger this value is, the larger droplets can get before breaking up. The default value is 6.
Target Count	The average number of particles per child parcel, derived from Eqn. (1439). The default value is 1000.
B1	The breakup parameter, B_1 in Eqn. (1438). The larger this value is, the later breakups occur. The default value is $\sqrt{3}$.

SSD Breakup Expert Properties

K0	Controls the mean distribution of child diameters; K_0, equal to $\langle \xi \rangle$ in Eqn. (1439). The default value is -0.1.
K1	Controls the breadth of the distribution of diameters; K_1 in Eqn. (1440). The default value is 0.1.
Maximum Child Parcels	The maximum number of parcels that can result from a breakup. Use this to limit the greater computational expense of a large number of child parcels. The default value is 10.
Normal Velocity Coefficient	Controls the radial diffusion of the spray; K_3 in Eqn. (1441). The larger it is, the more rapid the radial diffusion of the spray. Use a value of 0 if the spray angle is known and fixed at the injector. The default value is 0.1.

Working with Lagrangian Phase Boundary Conditions

The "boundary conditions" for a Lagrangian phase describe how particles behave when they impinge on a boundary. This behavior varies according to the boundary interaction mode and whatever conditions and values are required at the boundary by the interaction mode and the Lagrangian phase models. For some boundary types, this behavior follows obvious default modes with no additional user input. For other boundaries, however, some modeling is sometimes required, and so boundary conditions for each Lagrangian phase must be specified. You can choose to specify:

- Default boundary conditions which apply at each of these boundary types.
- Specific boundary conditions at a particular boundary.

Phase models that are associated with boundary interaction modes are:

- Bai-Gosman
- Satoh
- Ice Accretion

When you work with the Discrete Element Method (DEM), a special DEM mode is available.

Setting Default Lagrangian Phase Boundary Conditions

The Boundary Conditions of the Lagrangian phase are contained within the Boundary Conditions manager node for a Lagrangian phase.



It contains a sub-node for each of the boundary or interface types for which you can set the default boundary interaction mode. Boundary or interface types are only included if at least one boundary or interface of the type is present in the simulation.



The properties of a boundary type sub-node define the interaction mode of the Lagrangian phase for boundaries of that type, unless the mode is over-ridden at a specific boundary. Depending on the selected boundary interaction mode, additional sub-nodes can be added for boundary values.

Lagrangian Phase Boundary Conditions Properties

Mode	Selects the boundary interaction mode for the Lagrangian phase at this boundary type. The actual list that you see depends on the compatibility with the boundary type and the selected Lagrangian phase models.

Setting Lagrangian Phase Boundary Conditions for a Specific Boundary

Lagrangian phase boundary conditions can be set at specific boundaries as well as for the entire phase. The Lagrangian Specification node allows you to set these conditions. This node is a sub-node of the Physics Conditions node for the boundary of interest.



To set phase-specific conditions:

- Select the Lagrangian Specification node.
- Set Method to Specify for Boundary.

  The Phase Conditions manager node appears for the boundary.

Lagrangian Specification Properties

Method	Selects the method for the Lagrangian boundary condition specification.	
	Use Phase Defaults	Specifies that the default boundary conditions for all Lagrangian phases are used at this boundary.
	Specify for Boundary	Specifies that specific boundary conditions are provided for all Lagrangian phases at this boundary. A Phase Conditions sub-node is added as a child to the Physics Values node for this boundary.

What Is the Phase Conditions Manager?

The Phase Conditions manager node is added to a Physics Values node of a boundary when you choose to specify Lagrangian phase boundary conditions specifically for that boundary.



There are sub-nodes of the Phase Conditions node representing each of the Lagrangian phases.

Selecting the Mode

The Physics Conditions > Mode node within each phase, which has properties, allows you to specify the boundary interaction mode at the boundary for that phase. Depending on the selected boundary interaction mode, additional sub-nodes can be added for boundary values.



Phase Boundary Interaction Mode Properties

Active Mode	Selects the boundary interaction mode for the Lagrangian phase at this boundary. The actual list that you see depends on the compatibility with the boundary type and the selected Lagrangian phase models.

Setting a Phase Impermeable Boundary

Flow boundaries can be set to phase impermeable for Lagrangian multiphase flows. Phase impermeable means that particles do not pass through the boundary; instead they interact with the boundary as if it were a wall. Fluid in the continuous phase continues to pass through as normal.

To set a phase impermeable boundary:

- On a flow boundary, select the phase representation of the boundary node within the Phase Conditions manager.

- In the *Properties* window, set Type to Phase Impermeable.



The boundary node within the Phase Conditions manager changes to a wall icon. The wall icon indicates that Lagrangian particles see the boundary as a wall.

To set specific wall interaction options on the impermeable boundary:

- Select the Physics Conditions > Lagrangian Specification node of the parent boundary and set Method to Specify for Boundary. The Physics Conditions and Physics Values manager nodes appear beneath the impermeable boundary node for the phase.



Understanding Boundary Interaction Modes

The boundary interaction mode defines how particles behave when they impinge on a boundary. The boundary interaction mode controls what conditions and values, if any, must be set to determine the behavior.

The availability of a boundary interaction mode for a Lagrangian phase at a boundary depends on both the boundary type and the Lagrangian phase models, as summarized in the table below.

The Phase Impermeable type only appears if you have selected the Phase Impermeable option for a Lagrangian Phase within the Phase Conditions node of a flow boundary object.

Boundary/Interface Type	Massless Particle Interaction Modes	Material Particle Interaction Modes	DEM Particle Interaction Modes
Axis	Escape	Reflect	N/A
Baffle	Composite Escape* Stick	Bai-Gosman Composite Escape Rebound* Satoh Stick Vaporize	Escape DEM*
Contact	Composite Escape* Stick	Bai-Gosman Composite Escape Ice Accretion Rebound* Satoh Stick Vaporize	Escape DEM*
Fluid Film	Composite Escape* Stick	Bai-Gosman Composite Escape Ice Accretion Film Damping Fluid Film Rebound* Satoh Stick Vaporize	Escape DEM*
Flow-Split Outlet	Escape	Escape	Escape
Free-Stream	Escape	Escape	Escape
Fully Developed	Escape	Escape	Escape
Internal	Transmit	Transmit	Transmit
Mass Flow Inlet	Escape	Escape	Escape

Phase Impermeable	Composite Escape* Stick	Bai-Gosman Composite Escape Rebound* Satoh Stick Vaporize	Escape DEM*
Porous Baffle	Composite Escape Stick Transmit*	Composite Escape Rebound Stick Transmit*	Transmit DEM*
Pressure Outlet	Escape	Escape	Escape
Stagnation Inlet	Escape	Escape	Escape
Symmetry Plane	Escape	Reflect	N/A
Velocity Inlet	Escape	Escape	Escape
Wall	Composite Escape* Stick	Bai-Gosman Composite Escape Ice Accretion Rebound* Satoh Stick Vaporize	Escape DEM*

* default

Bai-Gosman

Bai-Gosman mode is available only with the Bai-Gosman Impingement model at walls, baffles, fluid film, phase impermeable, and contact interfaces. A procedure is adopted to determine some of the droplet post-impingement quantities depending on wall temperature, surface roughness and whether it is wet or dry.

Composite

Composite mode is available at baffle, contact, fluid film, phase impermeable, porous baffle, and wall boundaries. It is available with massless and material particles but not with DEM. The aim of the mode is to allow the actual boundary interaction mode for each parcel to be selected probabilistically, based on local conditions.

Selecting Composite mode adds a physics condition Composite Default Mode that allows you to select the main boundary interaction mode.

It also adds a manager of Mode Probabilities under Physics Values. Available boundary interaction modes are listed under this manager and the probability of its occurrence can be specified for each.



By default, the probability of each mode under Mode Probabilities is set to be constant with a value of `0.0`. Although you can use the usual options for a scalar profile to set the values, remember that the total sum of probability must be less than or equal to one. The probability of a certain event is 1, and

the probability of an impossible event is 0. Choosing the Field Function method allows the probability to be expressed in terms of local conditions, such as impingement velocity and wall temperature.



The probability of the default mode is automatically calculated as one minus the sum of the probabilities of all other modes. For this reason, it is recommended that the mode with the highest chance of occurrence is set as default.

DEM

DEM mode is only available with the Discrete Element Method (DEM). This mode makes the material properties of the DEM particles available at the boundary.

DEM Trap

DEM mode is only available with the Discrete Element Method (DEM) when the DEM Trap Walls model is selected. This mode allows you to specify the rate at which DEM particles leave a boundary.

Film Damping

Film Damping mode allows the motion of solid particles to be damped immediately before and after wall impingement as they pass through a fluid film. This mode is only available for solid particles on fluid film boundaries.



Selecting Film Damping mode adds a physics condition Film Damping Sub-Mode. The behavior of the particle as it impinges on the wall having passed through the film can be specified as stick, rebound, escape, or composite.



Fluid Film

In this mode, droplets which impinge on the Fluid-Film interface are converted to fluid film: their mass, momentum, and energy are added to the conservation equations maintained by the Fluid Film model, ensuring global conservation.

Ice Accretion

Ice Accretion mode is available only with the Ice Accretion model at walls and contact interfaces. It means that the droplet is added to the ice layer instantaneously upon contact with the boundary.

Escape

Escape mode means that the particle is removed from the simulation, effectively escaping the solution domain. The mass, momentum, and energy (where applicable) of the parcels are also removed.

Rebound

Rebounding particles remain active in the simulation, but with their velocity modified according to restitution relationships. For walls, baffles and contact interfaces, this mode is the default. When a two-way coupled particle rebounds off a Fluid Film, some of the energy of the particle can be transferred to the film.

Reflect

Reflect mode means that the particle rebounds perfectly from the boundary, effectively equivalent to rebound with unit restitution coefficients.

Satoh

Satoh mode is available only with the Satoh Wall Impingement model at walls, baffles, fluid film, phase impermeable, and contact interfaces. A portion of particles encountering a boundary rebounds, while the residue either escapes, sticks, or forms a fluid film. The response depends on the particle normal velocity, diameter, and the boundary properties. Satoh mode can be used to model the impingement of oil droplets in oil-mist separators.

Stick

Stick mode is available only for unsteady flows. Stuck particles remain stuck to the same boundary face with the same coordinates until either they evaporate or the face is no longer a wall. In the latter case, they become active (free) particles again. Stuck particles get the velocity of the boundary face to which they are attached, but note that this velocity is ignored for tracking purposes: if the wall has a net velocity relative to the mesh, the parcel is not tracked along the wall. No net force is allowed to act on the particle, and there is no momentum transfer. Heat and mass transfer continue as for free particles. In the current version of STAR-CCM+, there is no heat transfer directly to the wall using this mode. When a two-way coupled particle sticks to a Fluid Film, the kinetic energy of the particle is transferred to the film.

Transmit

Transmit mode means that the particle is allowed to cross the interface, with its position and velocity transformed in a way consistent with the interface topology.

Vaporize

Vaporize mode is available only with the Droplet Evaporation model at walls, baffles, and contact interfaces. It means that the droplet evaporates instantaneously upon contact with the boundary. If the Two-Way Coupling model is active, the mass, momentum, and energy of the droplet is transferred to the continuous phase; otherwise this mode is functionally identical to Escape mode.

Modeling Bai-Gosman Wall Impingement

The Bai-Gosman wall impingement model is based on literature findings and mass, momentum, and energy conservation constraints. To reflect the stochastic nature of the impingement process, the model uses a random procedure to determine some of the droplet post-impingement quantities. This randomization allows secondary droplets resulting from a primary droplet splash to have a distribution of sizes and velocities.

This section includes instructions on selecting this model.

The Bai-Gosman Wall Impingement Model

The Bai-Gosman wall impingement model ([232], [233], [234]) provides a methodology for modeling the behavior of droplets impacting on a wall. In particular, this model attempts to predict how and when droplets break up or stick to the wall. This model is used with impermeable boundaries (wall, contact, and baffle) as well as with fluid film. The Bai-Gosman Wall Impingement node represents this model in the Lagrangian phase model tree.



Bai-Gosman Wall Impingement Properties

Child Parcels	Number of child parcels to create during a splash/breakup event.
Lower Transition Temperature	Temperature for transition between ranges 1 and 2 (T_{12} in Eqns. (1444) and (1459)). This value is expected to be approximately the boiling temperature of the droplet.
Upper Transition Temperature	Temperature for transition between ranges 2 and 3 (T_{23} in Eqns. (1459) and (1464)). This value is expected to be approximately the Leidenfrost temperature of the droplet. See Using the Leidenfrost Temperature.
WeT1	Minimum Weber number for breakup We_{T1} in temperature ranges 2 (Eqns. (1460) and (1461)) and 3 (Eqns. (1465) and (1466)).
WeT2	Minimum Weber number for breakup and spread We_{T2} in temperature range 2, see Eqns. (1461) and (1462).

Wall State

This property defines whether an impermeable boundary (other than a Fluid Film boundary) is treated as wet or dry. In the case of a Fluid Film boundary, the wall is treated as wet when the film thickness is larger than zero. See Bai-Gosman Wall Impingement.

Wet

Specifies that the wall is wet; droplet impingement produces different outcomes, depending on Weber number and wall temperature:

At low Weber numbers, droplets adhere, rebound, or spread.

At high Weber numbers, droplets splash, breakup and spread, or break up and rebound.

Dry

Specifies that the wall is dry; droplet impingement produces different outcomes, depending on Weber number and wall temperature:

At low Weber numbers, droplets spread or rebound.

At high Weber numbers, droplets splash, breakup and spread, or break up and rebound.

Leidenfrost effect occurs when the temperature is high enough. See Using the Leidenfrost Temperature.

Bai-Gosman Wall Impingement Expert Properties

Rosin-Rammler Exponent	Exponent in Rosin-Rammler size distribution, q in Eqn. (1457).
Cf	Wall friction coefficient for splash-generated droplets, see Eqn. (1458).
Aw	Coefficient governing the onset of the splash regime, see Eqn. (1448).
a0	Coefficient governing the number of splash-generated droplets, see Eqn. (1456).
Minimum Ejection Angle	Minimum ejection angle from the wall of splash-generated droplets.
Maximum Ejection Angle	Maximum ejection angle from the wall of splash-generated droplets.
WeRebound	Minimum Weber number for rebound in range 1. The default value is 2.
WeSpread	Minimum Weber number for spread in range 1. The default value is 20.

Selecting the Bai-Gosman Wall Impingement Model

The Bai-Gosman wall impingement model is only used for liquid droplets (single- or multi-component). It also requires a Gas model in the physics continuum. To activate the model, select Bai-Gosman Wall Impingement

from the *Wall Impingement* group box of the *Phase Model Selection* dialog.



Choosing this option makes the Bai-Gosman boundary interaction mode available in the Mode property list for walls, baffles, contact interfaces, and fluid films. The additional Adhere mode node appears for walls, baffles, and contact interfaces only.

Selecting the Bai-Gosman Adhere Mode

When the Bai-Gosman mode is activated (either directly through the Mode node or when Composite mode is selected), an additional node Adhere Mode appears in the Physics Conditions. The Adhere Mode allows you to select which specific boundary interaction mode is used when the impingement

satisfies the Adhere conditions. The default mode for this condition is Stick.



Modeling Satoh Wall Impingement

The Satoh Wall Impingement Model is a wall impingement feature for modeling oil droplets in oil-mist separators.

This section includes instructions on selecting this model.

The Satoh Wall Impingement Model

The Satoh model node represents the Satoh Wall Impingement model. It provides a method for modeling the behavior of oil droplets in a "blow-by" flow as they encounter the walls, baffles, and porous plates of an oil-mist separator.



According to Satoh and others [263], the particle diameter and its velocity normal to the boundary determine the outcome of a collision between a particle and a boundary. The relationship between the two is assumed to follow two power law curves as depicted below. Below the lower curve,

called the Initial Spread Line, all the particles rebound. Above the upper curve, called the Full Spread Line, the particles spread over the boundary. In the area between the curves some particles rebound and the rest spread.



As per Satoh and others [263], the droplet diameter is in the range 0 m to 8E-6 m. The droplet impingement velocity normal to the wall is in the range 0.01 m/s to 13 m/s. In this range, small droplets with low normal velocity rebound from the wall whereas large droplets with high normal velocity tend to spread over the surface to form a film.

In STAR-CCM+ fluid film formation is dependent on both model selection and the boundary type. An order of precedence is used to model spreading behavior: if the Fluid Film model is not available, Stick mode is used. If neither mode is applicable on the incident boundary, the particle is allowed to Escape.

Satoh Wall Impingement Properties

Reference Diameter	Reference particle diameter, D_{ref}, used to calculate the non-dimensional diameter for the initial (1468) and full spread lines (1469).
Initial Pre-exponential	Scalar multiplier, a_{ISL}, for initial spread line.
Initial Exponent	Exponent, n_{ISL}, for the dimensionless diameter for the initial spread line.
Full Pre-exponential	Scalar multiplier, a_{FSL}, for full spread line.
Full Exponent	Exponent, n_{FSL}, for the dimensionless diameter for the full spread line.

Selecting the Satoh Wall Impingement Model

The Satoh wall impingement model is only used for liquid droplets (single- or multi-component). It also requires a Gas model in the physics continuum. To activate the model, select Satoh Wall Impingement from the

Wall Impingement group box of the *Phase Model Selection* dialog.



Choosing this option makes the boundary interaction mode available in the Mode property list for walls, baffles, and contact interfaces.



Modeling Ice Accretion

This section describes the Ice Accretion model and how to select it.

The Ice Accretion Model

The Ice Accretion model node represents the Ice Accretion model. When this model is active, particles impinging on a boundary ice layer are accounted for in the mass and energy balance equations from the Deicing model.

This model governs the transfer of particle mass to an ice layer on impact with an ice-layer boundary, to become part of the overall mass balance of the ice layer.



Ice Accretion Model Properties

None.

Selecting the Ice Accretion Model

The Ice Accretion model becomes available only if you have the Deicing model selected in the physics continuum and these models selected under the Lagrangian Phases node:

- Lagrangian Energy
- Two-Way Coupling
- The single-component droplet model or the single-component particle model

Select the Ice Accretion option from the *Optional Models* group as shown below.



The dialog shows the Ice Accretion model activated.



If the Auto-select recommended Phase models box has been checked, then selecting the Ice Accretion model causes the Energy and Two-Way Coupling models to be selected as well. If the box is not checked, and the models are

not already selected, a warning message appears beneath the *Enabled Phase Models* group box to indicate that more selections are required.



Trying to iterate without the DeIcing model selected results in a warning similar to the one shown below.



Boundary Values Reference

Depending on how the Lagrangian phase boundary conditions and models are set, you can control the interaction with boundaries using the following values:

- Normal Restitution Coefficient
- Tangential Restitution Coefficient
- Erosion Ratio

Normal Restitution Coefficient

Added to the tree when Rebound is chosen as the interaction mode at a boundary.

Tangential Restitution Coefficient

Added to the tree when Rebound is chosen as the interaction mode at a boundary.

Erosion Ratio

Added to the tree when the Erosion model is activated in the Lagrangian phase. Depending on the method that is chosen for the erosion ratio, a sub-node is added to this node.

Erosion Ratio Properties

Method	Selects the method for calculating the erosion ratio.	
	Ahlert	Specifies that the erosion ratio is calculated using the Ahlert Correlation. An Ahlert node is added as a child to this node.
	Field Function	Specifies that the erosion ratio is calculated using a field function. This method is typically user-defined and can involve only constants and particle-based functions. A field function method node is added as a child to this node.
	DNV	Specifies that the erosion ratio is calculated using the DNV Correlation. A DNV node is added as a child to this node.
	Neilson-Gilchrist	Specifies that the erosion ratio is calculated using the Neilson-Gilchrist Correlation. A Neilson-Gilchrist node is added as a child to this node.
	Oka	Specifies that the erosion ratio is calculated using the Oka Correlation. An Oka node is added as a child to this node.
Dimensions	The dimensionality of the erosion ratio (read-only).	

Ahlert

This sub-node is added to the erosion ratio node when Ahlert is chosen as the erosion ratio method. It represents the Ahlert Correlation. The default values for this method are for liquid-borne semi-rounded sand particles eroding aluminum, as given by McLaury et al. [247].

Ahlert Properties

Material Coefficient	A material-dependent constant. It is the quantity K in Eqn. (1392).
Shape Coefficient	A factor to account for the shape of the particles. It is the quantity F_s in Eqn. (1392). Suggested values are 1 for angular particles, 0.53 for semi-rounded particles and 0.2 for fully rounded particles.
Transition Angle	The angle α_0 which is the upper limit of validity of the Low Alpha Function and the lower limit of validity of Eqn. (1393).
Low Alpha Function	The polynomial for the angle function $f(\alpha)$ in Eqn. (1392) to be used below the transition angle α_0. The polynomial is in terms of the incidence angle in radians.
Coefficient W	The coefficient w in Eqn. (1393) for the angle function $f(\alpha)$ above the transition angle α_0.
Coefficient X	The coefficient x in Eqn. (1393) for the angle function $f(\alpha)$ above the transition angle α_0.
Coefficient Y	The coefficient y in Eqn. (1393) for the angle function $f(\alpha)$ above the transition angle α_0.
Coefficient Z	The coefficient z in Eqn. (1393) for the angle function $f(\alpha)$ above the transition angle α_0 (read-only).
Reference Velocity	The normalization factor for the particle velocity. It is the quantity u_{ref} in Eqn. (1392).
Velocity Exponent	The exponent for the normalized velocity. It is the quantity n in Eqn. (1392).

DNV

This sub-node is added to the erosion ratio node when DNV is chosen as the erosion ratio method. It represents the DNV Correlation. The default values for this method are for air-borne sand eroding carbon steel, and are taken from Haugen *et al.* [242].

DNV Properties

Material Coefficient	A material-dependent constant. It is the quantity K in Eqn. (1394). The default value is 2.0E-9.
Reference Velocity	The normalization factor for the particle velocity. It is the quantity u_{ref} in Eqn. (1394). The default value is 1.0 m/s.
Velocity Exponent	The exponent for the normalized velocity. It is the quantity n in Eqn. (1394). The default value is 2.6.
Angle Function	The polynomial for the angle function $f(\theta)$ in Eqn. (1394). It is evaluated in terms of the incidence angle in degrees. See Eqn. (1395).

Neilson-Gilchrist

This sub-node is added to the erosion ratio node when Neilson-Gilchrist is chosen as the erosion ratio method. It represents the Neilson-Gilchrist Correlation. The default values for this method are for liquid-borne sand eroding AISI 4130 steel, and are taken from Wallace et al. [273].

Neilson-Gilchrist Properties

Transition Angle	The angle at which cutting deformation changes behavior. It is the quantity α_0 in Eqn. (1397).
Deformation Cutoff Velocity	The normal velocity at which deformation wear is minimized. It is the quantity K in Eqn. (1398).
Cutting Coefficient	The material-dependent constant governing cutting erosion. It is the quantity ε_C in Eqn. (1397).
Deformation Coefficient	The material-dependent constant governing deformation erosion. It is the quantity ε_D in Eqn. (1398).

Oka

This sub-node is added to the erosion ratio node when Oka is chosen as the erosion ratio method. It represents the Oka Correlation. The default values for this method are for air-borne sand eroding 0.025% carbon steel, and are taken from Oka *et al.* [253] and Oka and Yoshida [254], except for the Reference Erosion Ratio, which is derived from the DNV Correlation.

Oka Properties

Reference Velocity	The normalization factor for the particle velocity. It is the quantity u_{ref} in Eqn. (1399).
Velocity Exponent	The exponent for the normalized velocity. It is the quantity k_2 in Eqn. (1399).
Reference Erosion Ratio	The erosion ratio at reference conditions and normal incidence. It is the quantity e_{90} in Eqn. (1399).
Coefficient Hv	The coefficient H_v in the angle function Eqn. (1400). It is typically, but not necessarily, equal to the Vickers hardness of the eroded material in GPa.
Coefficient n1	The coefficient n_1 in the angle function Eqn. (1400).
Coefficient n2	The coefficient n_2 in the angle function Eqn. (1400).
Reference Diameter	The normalization factor for the particle diameter. It is the quantity D_{ref} in Eqn. (1399).
Diameter Exponent	The exponent for the normalized diameter. It is the quantity k_3 in Eqn. (1399).

Working with Injectors

In STAR-CCM+, particles enter the fluid continuum through injectors. While the Lagrangian phase defines *what* particles enter the domain and how they behave, the injector defines *where,* in what direction, and with what frequency, the particles enter. Injectors are contained within a top-level injectors manager node.

Eight types of injector are currently provided in STAR-CCM+:

- The Part injector generates a cloud of injection points that are positioned according to the geometry of the parts you select.
- The Point injector injects particles from a single point in a direction you specify.

- The Hollow Cone injector injects particles from a single point, with inner, and outer angles you specify.
- The Solid Cone injector injects particles from a single point, with only the outer angle you specify.
- The Surface injector injects particles from the face centroids of a geometric surface you select.
- The Pressure Swirl injector is used with the Linear Instability Sheet Atomization (LISA) model to generate particles with properties based on injector properties you specify.
- The Lattice injector is used with the Discrete Element Method (DEM) to inject particles into the system with specified, mesh-independent spatial distribution.
- The Random injector is used with the Discrete Element Method (DEM) to inject particles into a region with a random point distribution. The Particle Amount Specification controls the number of particles that are injected.
- The Film Stripping injector is used with the Fluid Film model, to simulate droplets being formed and removed from the surface of a fluid film. Both wave stripping and edge stripping can contribute to this phenomenon.

These types are explained in greater detail in the section, Understanding Injector Types. Further documentation is provided on:

- Using injectors
- Using injector particle size distributions
- Injector Conditions Reference
- Injector Values Reference

Injectors Manager Node

Injector nodes represent injectors and are all contained within the Injectors manager node, which has properties and its own pop-up menu. The injector manager can hold any number of injectors up to the memory limits of the computer.



Injector Manager Expert Properties

Tree Threshold	Specifies the maximum estimated number of injection points per region before the algorithm changes to a tree search. There is a memory and CPU time penalty that is incurred in building the tree. If there are enough points, the algorithm repays this penalty by using the tree search to locate the points faster than the default brute-force algorithm.

Understanding Injector Types

Ten injector types are provided in STAR-CCM+:

- Part injector
- Point injector
- Hollow cone injector
- Solid cone injector
- Random injector —used with the Discrete Element Method (DEM)
- Surface injector
- Pressure Swirl injector
- Table injector
- Lattice injector—used with the Discrete Element Method (DEM)
- Film Stripping—used with the Fluid Film model

Each type has distinctive features that determine where and how particles are injected. A common property among all injector types is the injector Inputs property. As its use varies between injector types, it is important to understand how the property is used in each injector: this property is discussed in the individual descriptions below. The type of the injector can also collaborate with the Lagrangian phase models to determine which conditions and values are necessary to define the initial conditions completely.

A key role of the type of injector is to determine how the injection points are generated from the Inputs of the injector. Injection points are where the injector introduces particles—in other words the initial positions of the particles.

Part Injector

A part injector represents a collection of injection points that are extracted from the geometry of the parts of the injector. If desired, the points can be a subset of the points that are defined on the parts. The point inclusion

probability property of the injector sets the probability that a particular point is used. An example of a part injector (defined by three line probes) is shown below.



Point Injector

A point injector injects at a single user-specified point which is the position of the injector. The initial cell location for this point is found by searching the parts of the injector, which are limited to regions. Thus, if you can specify which region contains the point, the input parts can be used to accelerate the cell location process. You can also use input parts to resolve ambiguities. For example, suppose that the point lies on the interface between two regions but the intention is not to inject into one of them. An

error is thrown if a cell location for the point injector cannot be found by searching its input parts. Typically this means that the point is outside the specified region or regions.



Hollow Cone Injector

A hollow cone injector injects parcels in a cone formation about an axis you specify. An inner cone angle and an outer cone angle, define the cone, as shown in the diagram below. (The diagram shows only half the cone; the injector uses a whole cone.)



The specified number of parcel streams are injected in random directions conforming to a uniform distribution on the surface of the cone. Given the direction, the velocity is obtained using the specified velocity magnitude. In steady simulations, each iteration uses the same stochastically chosen direction and size distribution. In unsteady simulations, the direction and size distribution are randomized again and newly sampled each time step, if you activate the Cone Randomization property. For this injector, the Inputs property is used to define the region within which the origin of the axis is located.



Solid Cone Injector

A solid cone injector is similar to the hollow cone type, except that there is no inner cone angle parameter (the angle is treated as zero).

Surface Injector

A surface injector injects parcels from each face of the surface comprising the parts of the injector. Parcels are injected at the centroid of each face in the direction of the face surface normal. For this type, parts are selected from those parts that belong to the model geometry.

A unique feature of the surface injector type is that the injection flow rate can be defined with a flux, that is, as a rate per unit area. The final flow rate is the product of the face area and the specified flux, whether this flux is a particle, mass or volume flux. In the following example, the initial location

of each parcel coincides with the centroid of each face on the starting surface. The mass of each parcel is proportional to the area of the face from which it came (as a mass flux has been specified).



When the flow rate through a surface injector is defined directly, the flow rate through each face scales to the ratio of the face area to the overall surface area.

Pressure Swirl Injector

When used with the LISA model, this injector type represents a pressure swirl atomizer. A thin liquid film is formed within the injector nozzle that breaks up into droplets soon after entering the gas. A mathematical model represents the breakup process. This model derives from a stability analysis of the two-dimensional liquid film.

For this injector type, the Inputs property is used to define the region within which the origin of the axis is located. The inner cone angle and outer cone angle are used to calculate the spray angle, which is defined as half the sum of the inner and outer cone angles. An injector upstream pressure value is also required, to use in predicting the exit velocity of the thin liquid film.

Note: This injector is only available for use with the three-dimensional space model.

Table Injector

A table injector injects parcels from points specified in an input table. The table also specifies the injection times and the properties of the parcels, such as mass, diameter, and temperature.

Input parts for a table injector must be regions. The injection points can be anywhere within a specified region.

Using Injectors

Injector nodes represent injectors within the injector manager. They are created using the right-click menu on the injectors manager node. The injector node has its own properties. It contains condition and value objects that apply to the Lagrangian phase to which it is assigned.



An injector has a type, which determines how it generates injection points. Particles are introduced at these injection points and are assigned to the user-specified Lagrangian phase. The goal of an injector is to provide the initial state of particles, according to its user-specified conditions and values. The injector type and the models of the Lagrangian phases determine these conditions and values.

Injector Properties

Inputs	Selects the parts that are used to generate the injection points. How the parts are used and which parts can be selected depend on the injector type.
Type	The injector type. The possible types are listed below.
	Film Stripping — A film stripping injector.
	Lattice Injector — A lattice injector.
	Part Injector — A part injector.
	Point Injector — A point injector.
	Hollow Cone Injector — A hollow cone injector.
	Solid Cone Injector — A solid cone injector.
	Random Injector — A random injector.
	Surface Injector — A surface injector.
	Pressure Swirl Injector — A pressure swirl injector.
	Table Injector — A table injector.
Number of Points	The number of injection points this injector produces (read-only). The injection points are first generated when the solution is initialized.
Index	A unique identifier for this injector (read-only)
Lagrangian Phase	Selects the Lagrangian phase from a drop-down list.

To set up an injector, do the following:

- Create the injector
- Assign a Lagrangian phase
- Set the injector type
- Assign the injector parts
- Set conditions and values

Other operations that can be performed on an injector node are:

- Renaming an injector
- Deleting an injector

Creating an Injector

To create an injector, right-click the Injectors manager node and select New.



Assigning a Lagrangian Phase to the Injector

To assign a Lagrangian phase to an injector, select the injector node and select that phase from the drop-down list of the Lagrangian Phase property.



Setting the Injector Type

To set the injector type, select the injector node and select the required type from the drop-down lost of the Type property.



Assigning the Injector Parts

To assign the input parts of the injector, click the right half of the Inputs property of the injector node. This action activates the *Inputs* dialog.



The displayed selection contains parts appropriate to the injector type; not all parts are appropriate for every injector.

Setting Conditions and Values

Each injector contains sub-nodes that are managers of injector conditions and values, which work in ways similar to the conditions and values for a STAR-CCM+ simulation.



Renaming an Injector

Similar to other objects, an injector can be renamed within STAR-CCM+.

Deleting an Injector

Injectors can be deleted from an existing simulation.

Injector Manager Pop-Up Menu

This pop-up menu item is available when you right-click the Injectors manager node.

New	Creates an injector with a default name.

See also: common menu items.

Using Injector Particle Size Distributions

Particle size specification at an injector can be either by diameter or mass. Among the methods available, four define the particle size using a particle size distribution, each represented by a node.



The available methods are:

- Constant
- Field Function
- Log-Normal (using a particle size distribution)
- Normal
- Rosin-Rammler (using a particle size distribution)
- Table (iteration)
- Table (particle size CDF) (using a particle size distribution)
- Table (particle size CDF, time) (using a particle size distribution)
- Table (r)

- Table (time)
- Table (x,y,z)
- Table (xyz,time)
- User Code

In each case, a range of particle sizes is generated between a user-specified minimum and maximum, according to the distribution of the method.

What Is Log-Normal?

The Log-Normal size distribution is a smooth continuous distribution depending on four parameters: mean size, standard deviation, minimum size, and maximum size.

It is available for both the Particle Diameter and Particle Mass Particle Size Specification options in the Lagrangian Multiphase model and has properties.

Log-Normal Size Distribution Properties

Mean	The mean size, $\bar{D}$ in Eqn. (1474).
Standard Deviation	The standard deviation, σ in Eqn. (1474).
Minimum	The minimum size that the distribution generates.
Maximum	The maximum size that the distribution generates.

What Is Normal?

The Normal size distribution is a smooth continuous distribution depending on four parameters: mean size, standard deviation, minimum size, and maximum size.

It is available for both the Particle Diameter and Particle Mass Particle Size Specification options in the Lagrangian Multiphase model and has properties.

Normal Size Distribution Properties

Mean	The mean size, $\bar{D}$ in Eqn. (1475).
Standard Deviation	The standard deviation, σ in Eqn. (1475).
Minimum	The minimum size that the distribution generates.
Maximum	The maximum size that the distribution generates.

What Is Rosin-Rammler?

The Rosin-Rammler size distribution is a smooth continuous distribution depending on four parameters: reference size, exponent, minimum size, and maximum size.

It is available for both the Particle Diameter and Particle Mass Particle Size Specification options in the Lagrangian Multiphase model and has properties.

Rosin-Rammler Size Distribution Properties

Exponent	The exponent, q in Eqn. (1476).
Reference	The reference size, D_{ref} in Eqn. (1476).
Minimum	The minimum size that the distribution generates.
Maximum	The maximum size that the distribution generates.

What Is Table (particle size CDF)?

A table can be used to specify the particle size distribution using a cumulative distribution function (CDF). A particle size CDF table has a column of 'CDF' values and a column of 'data' values for the particle sizes.

Note: The CDF values must be strictly increasing—duplicate values are not permitted.

It is available for both the Particle Diameter and Particle Mass Particle Size Specification options in the Lagrangian Multiphase model and has properties.

Table (Particle Size CDF) Properties

Table: CDF	Looks for the column in the table with the heading CDF (read-only).
Table: Data	Identifies the selected column in the table file as the data column containing the particle sizes for the corresponding CDF values.
Table	Selects a table from among the table sub-nodes in the Tables node of the Tools node.

You must have read a table into STAR-CCM+ to specify these properties; otherwise the drop-down lists are empty and the Table: Data and Table properties show as `Null`. The other property, Table: (CDF), is read-only since it simply indicates the title of the column that the tabular method expects.

What Is Table (particle size CDF, time)?

This table can be used to specify a particle distribution using a cumulative distribution function (CDF) that varies over time.

The first column of the table lists the sizes of the particles, which must increase monotonically. Subsequent columns have headers of the form `CDF [t=####]` where `####` is a time value. These time values must increase monotonically, without duplicates, as you go down the header row. Each column of the table corresponds to the starting time for use of the next CDF.

Under each time column, each entry is a CDF value; these must increase monotonically down the column, without duplicates, though the rate of increase is independent between columns.

Example:

```
diameter CDF [t=0.01] CDF [t=0.50] CDF [t=1.2] CDF [t=1.24]
0.01     0.000        0.000       0.00        0.000
0.03     0.002        0.030       0.05        0.004
0.04     0.004        0.040       0.12        0.078
0.06     0.007        0.050       0.89        0.512
0.11     1.000        1.000       1.00        1.000
```

Table (Particle Size CDF, Time) Properties

Table: Particle Size	The header of the column listing particle sizes. In the example, this is diameter.
Table: Starting Time Column	The header of the column for the starting simulation time, for example, CDF [t=0.50]. The starting time must be less than or equal to the starting time of the simulation.
Table	Selects a table from among the table sub-nodes in the Tables node of the Tools node.
Interpolation	Selects interpolation method, Spline, Step, or Linear.

You must have read a table into STAR-CCM+ to specify these properties; otherwise the drop-down lists are empty and the properties show as `Null`.

Using Injection Tables

Injection tables are the input method for table injectors. These tables specify the location of injection points and the time of injection. They can also specify some properties of the injected parcels. There are three kinds of injection table:

- Injection Points table
- Unsteady Injection Points table
- Parcels table

Injection Points Table

Injection Points tables are available with both steady and unsteady Lagrangian Multiphase simulations. A table injector using an Injection Points table has the following features:

- Each row of the table represents an injection point.
- Any injection point that is not in the parts specified under the injector Inputs property is discarded.
- The injector uses the Parcel Streams value for the number of parcel streams per injection point.
- Parcels are injected according to the Update Frequency property of the solver.

- The injector can use any of the Flow Rate Distribution options, but only the Per Injector Point or Per Parcel Stream options allow you to specify the flow rate in the injection table.
- The table can include any properties of the parcels listed under the physics values node of the injector. Examples of such properties include particle diameter, velocity, temperature, passive scalar values, parcel mass fraction, and parcel mole fraction.

Example table:

X(cm)	Y(cm)	Z(cm)	Diameter	U	V	W	Mass Flow Rate	Temperature
0	0	6.0	1.e-05	0	0	0.001	0.001	350
0	0	2.0	2.e-05			0.002	0.001	350

Injection Points Table Properties

The table is represented by a node named Injection Points Table, located under the injector Values node. The table node has the following properties:

Table	The table name, selected by drop-down menu from the tables defined under Tools > Tables. The default is None.
X coordinate Y coordinate Z coordinate	The names specified in the table for the columns for the X, Y, and Z coordinates. For the example table in this section, the entries would be `X(cm)`, `Y(cm)`, and `Z(cm)`.
Coordinate System	The coordinate system in which the injection points are specified. This can be the laboratory system or a local coordinate system you have created. This coordinate system applies only to the injection points, not to any other related properties, such as velocity. The default is Laboratory.

Unsteady Injection Points Table

Unsteady Injection Points tables are available only with unsteady Lagrangian Multiphase simulations. A table injector using an Unsteady Injection Points table has the following features:

- Each row of the table represents an injection point at a specific time value.
- The total number of injection points is equal to the maximum index found in the table.

- Each injection point is identified in the table by a unique positive index.
- Numbering for the injection point indices starts from 1.
- For every injection point index, there is at least one row in the table.
- For every injection point index, the time values must be specified in ascending order.
- At any given time, for each injection point index, the locations are interpolated in time as follows:
  - If the current value of time is less than the minimum time value found for that injection point, parcels are injected from the location corresponding to the minimum time value.
  - If the current value of time is greater than the maximum time value found for that injection point, parcels are injected from the location corresponding to the maximum time value
- Any injection point that is not in the parts specified under the injector Inputs property is discarded.
- The injector uses the Parcel Streams value for the number of parcel streams per injection point.
- Parcels are injected according to the Update Frequency property of the solver.
- The injector can use any of the Flow Rate Distribution options, but only the Per Injector Point or Per Parcel Stream options allow you to specify the flow rate in the injection table.
- The table can include any properties of the parcels listed under the physics values node of the injector. Examples of such properties include particle diameter, velocity, temperature, passive scalar values, parcel mass fraction, and parcel mole fraction.

Example table, for a simulation with time-step of 0.001:

Index	X(cm)	Y(cm)	Z(cm)	Diameter	Time
1	0	0	1.0	1.e-05	0.0012
2	0	0	2.0	2.e-05	0.0022
1	0	0	3.0	3e-05	0.0032
3	0	0	3.0	3e-05	0.0035
2	0	0	4.0	4e-05	0.0042

There are three injection points, indexed 1, 2, and 3. At every time-step, all three injection points inject parcels. At every time-step, the position of the injection point is interpolated, based on the values given in the table.

- For the time 0.001, corresponding to the first time-step, points inject from their earliest given position. For example, the earliest time specified for point 1 is 0.0012 s, which is greater than the current physical time, so point 1 injects from position 1 cm.

Injection Point	Position
1	1 cm
2	2 cm
3	3 cm

- At time 0.002, corresponding to the second time-step, there are no positions specified for points 1 or 3.
  - Point 1 injects from a position between 1 cm and 3 cm; the exact position depends on the interpolation option specified.
  - Point 3 injects from a position of 3 cm, because that is still its earliest specified position.
- At time 0.003, corresponding to the third time-step, there are no positions specified for points 2 or 3.
  - Point 2 injects from a position between 2 cm and 4 cm; the exact position depends on the interpolation option specified.
  - Point 3 injects from a position of 3 cm, because that is still its earliest specified position.
- After the value for time passes 0.004, all three injection points inject parcels from the position for the last, highest time value:

Injection Point	Position
1	3 cm
2	4 cm
3	3 cm

Unsteady Injection Points Table Properties

The table is represented by a node named Unsteady Injection Points Table, located under the injector Values node. The table node has the following properties:

Table	The table name, selected by drop-down menu from the tables defined under Tools > Tables. The default is None.

X coordinate Y coordinate Z coordinate	The names specified in the table for the columns for the X, Y, and Z coordinates. For the example table in this section, the entries would be `X(cm)`, `Y(cm)`, and `Z(cm)`.
Coordinate System	The coordinate system in which the injection points are specified. This can be the laboratory system or a local coordinate system you have created. This coordinate system applies only to the injection points, not to any other related properties, such as velocity. The default is Laboratory.
Time	The name of the column listing time values (read-only). The column in the table must be named `time`, `Time`, or `TIME`.
Injection Point Index	The name of the column listing a unique positive index for each injection point (read-only). The column in the table must be named `index`, `Index`, or `INDEX`.

Expert Properties

Interpolation	Selects interpolation method, Spline, Step, or Linear.

Parcels Table

Parcels tables are available only with unsteady Lagrangian Multiphase simulations. A table injector using a parcels table has the following features:

- Each row of the table represents a parcel at a specific time value.
- Parcels are injected only at the times specified in the table. There is no interpolation.
- A set of rows with time values that lie between the current physical time and the current physical time plus the time-step are treated as a single block of parcels and are injected at the start of the most recent time-step.
- Any parcel that is not in the parts specified under the injector Inputs property is discarded.
- The injector uses the Particle Count value for the number of particles per parcel.
  - This value can be constant or read from the injection table.

  - This value is taken as 1 if the DEM model is active, or if the value is not set.
- The table can include any properties of the parcels listed under the physics values node of the injector. Examples of such properties include particle diameter, velocity, temperature, passive scalar values, parcel mass fraction, and parcel mole fraction.

Example table, for a simulation with time-step of 0.001:

X(cm)	Y(cm)	Z(cm)	Diameter	Time
0	0	1.0	1.e-05	0.0
0	0	2.0	2.e-05	0.0
0	0	3.0	3e-05	0.003
0	0	4.0	3e-05	0.0031
0	0	5.0	4e-05	0.0032

The table has two time blocks, corresponding to injection times 0.0 and 0.003 seconds.

- There are two parcels in the first time block, and three parcels in the second time block.
- In the first time step (corresponding to physical time = 0.001) two parcels are injected from 1.0 cm and 2.0 cm.
- In the second time step (corresponding to physical time = 0.002), no parcels are injected.
- In the third time-step (corresponding to physical time = 0.003), no parcels are injected.
- In the fourth time step (corresponding to physical time = 0.004) three parcels are injected from 3.0 cm, 4.0 cm and 5.0 cm.

Parcel Table Properties

The table is represented by a node named Parcels Table, located under the injector Values node. The table node has following properties:

Table	The table name, selected by drop-down menu from the tables defined under Tools > Tables. The default is None.
X coordinate Y coordinate Z coordinate	The names specified in the table for the columns for the X, Y, and Z coordinates. For the example table in this section, the entries would be `X(cm)`, `Y(cm)`, and `Z(cm)`.

Coordinate System	The coordinate system in which the injection points are specified. This can be the laboratory system or a local coordinate system you have created. This coordinate system applies only to the injection points, not to any other related properties, such as velocity. The default is Laboratory.
Time	The name of the column listing time values. The column in the table must be named `time`, `Time`, or `TIME`.

Injector Conditions Reference

Injectors can be set up with the following conditions:

- Angular Velocity Specification
- Coarse Graining Option Specification
- Flow Rate Distribution
- Flow Rate Specification
- Injection Time Rate Specification
- Lattice Structure
- Orientation Specification
- Overlap Specification
- Particle Amount Specification
- Particle Size Specification
- Species Specification
- Time Randomization
- Velocity Specification

Angular Velocity Specification

DEM particles can be injected with an angular velocity specified using vector notation.

Angular Velocity Specification Properties

Method	Specifies the distribution of the flow rate.	
	Components	Enter the angular velocity as a vector with magnitude. This option adds an Angular Velocity node to the Values node.

Coarse Graining Option Specification

This node is added to the tree when the Coarse Grain Particle model is activated. This condition defines the number of individual particles represented in each simulated parcel.

Coarse Graining Option Specification Properties

Method	Specifies how the number of individual particles per coarse-grained parcel is determined.	
	Constant	Indicates that the number of individual particles per coarse-grained parcel is specified directly. This is the only option for random injectors. This option adds a Particle Count node to the Values node. The default value is 1.
	Dynamic	Indicates that STAR-CCM+ dynamically calculates the number of individual particles per coarse-grained particle. The number depends on other injector conditions and values, such Flow Rate Specification, Particle Size Specification, and Mass Flow Rate. The number need not be an integer, but is never less than 1.

Flow Rate Distribution

This node is added to the tree when the injector is any type other than a point injector. There are three ways to distribute the flow rate. As an example, consider an injector consisting of 20 injection points. Each point injects 10 parcels every time step; that is to say, each injection point has 10 parcel streams. Let the time step be 1 s and the specified flow rate be 1 kg/s.

- If the flow rate distribution is Per Injector, the total mass per second is 1 kg.

- If the flow rate distribution is Per Injection Point, the total mass per second is 20 kg.
- If the flow rate distribution is Per Parcel Stream, the total mass per second is 200 kg.

Flow Rate Distribution Properties

Distribution	Specifies the distribution of the flow rate.	
	Per Parcel Stream	Indicates that the specified flow rate is applied at each parcel stream of each injection point of this injector.
	Per Injection Point	Indicates that the specified flow rate is applied at each injection point of this injector. (Not available for point injectors.)
	Per Injector	Indicates that the specified flow rate is the total for all the injection points of this injector.

Flow Rate Specification

This condition specifies how the flow rate of particles at the injector is defined.

Flow Rate Specification Properties

Method — Selects the method for entering the flow rate.

- **Mass Flow Rate** — Indicates that the rate at which the injector introduces particles is deduced from a mass flow rate. Hence this condition is not available with the Massless Particles model. A Mass Flow Rate node is added to the Values node.
- **Volume Flow Rate** — Indicates that the rate at which the injector introduces particles is deduced from a volume flow rate. Hence this condition is not available with the Massless Particles model. A Volume Flow Rate node is added to the Values node.
- **Particle Flow Rate** — Indicates that the rate at which the injector introduces particles is specified directly. A Particle Flow Rate node is added to the Values node.
- **Mass Flux** — Indicates that the rate at which the injector introduces particles is specified using a mass flux. The final mass flow rate depends on the surface area of the injector part. A Mass Flux node is added to the Values node. This method is only available for surface type injectors.
- **Volume Flux** — Indicates that the rate at which particles are introduced is specified using a volume flux. The final volume flow rate depends on the surface area of the injector part. A Volume Flux node is added to the Values node. This method is only available for surface type injectors.

Particle Flux — Indicates that the rate at which particles are introduced is specified using a particle flux. The final particle flow rate depends on the surface area of the injector part. A Particle Flux node is added to the Values node. This method is only available for surface type injectors.

Dynamic — Indicates that STAR-CCM+ dynamically calculates the rate at which particles are introduced. The particle flow rate depends on the Particle Amount Specification and number of seeds values. This method is only available for random type injectors.

Injection Time Rate Specification

Unlike other Lagrangian particles, DEM particles interact with each other. Hence it is not always possible to inject the required mass simultaneously at the beginning of an iteration. This option allows you to change this default behavior to smooth the injection over the time-step.

Injection Time Rate Specification Properties

Time Rate — Selects if the DEM particles are injected simultaneously at the start of the time-step or over a series of sub steps.

- Time-step — All particles are injected simultaneously at the start of the iteration. The default setting. Issues can occur where many particles are injected simultaneously
- Substep — The particle injection occurs throughout the solution of the time-step. Using this method avoids issues where large numbers of particles are injected simultaneously but can cause the solvers to slow.

Lattice Structure

This node is added to the tree when the injector has lattice injector type. Allows injection of particles into the system with specified, mesh-independent spatial distribution. The main advantage of the lattice

injector when compared to the other injector types is the ability to introduce higher particle loading rapidly and make the particle distribution stable with respect to remeshing.

Lattice Structure Properties

Lattice	Selects the distribution type of injected DEM particles.	
	Simple cubic	Particle injection points correspond to the corners of a cube.
	Body centered cubic	Particle injection points correspond to the corners and centroid of a cube.
	Face centered cubic	Particle injection points correspond to the corners and face centers of a cube.



Simple cubic Body centered cubic Face centered

Orientation Specification

Added to the tree when a non-spherical DEM particle shape is selected. Used to specify the initial orientation of the particle as it is injected.

Orientation Specification Properties

Method	Selects the method for entering the particle orientation.	
	Angles	The orientation depends on the three Euler angles relative to the global co-ordinate frame.

Overlap Specification

Added to the tree when DEM particles are being injected with a table injector. Activated to specify that overlapping particles can be injected.

Overlap Specification Properties

Allow Overlap	Specifies whether overlapping particles canbe injected.	
	Activated	Overlap allowed.
	Deactivated	Overlap not allowed.

Particle Amount Specification

This condition specifies the method for defining the quantity of particles injected. Added to the tree when the random injector type is chosen.

Particle Amount Specification Properties

Method	Selects the method for entering the particle amount.	
	Porosity Limit	Indicates that the quantity of particles is specified by setting the volume fraction taken up by particles. When selected, a Porosity Limit node is added to the Values node. This method is the default.
	Number of Particles	Indicates that the quantity of particles is specified directly. When selected, a Number of Particles node is added to the Values node.

Particle Size Specification

This condition specifies the method for defining the size of particles from the injector.

Particle Size Specification Properties

Method	Selects the method for entering the particle magnitude.	
	Particle Size	Indicates that the particle size is specified directly. Since the particles are assumed to be spherical, this means that a Particle Diameter node is added to the Values node.
	Particle Mass	Indicates that the particle size is deduced from its mass. A Particle Mass node is added to the Values node.

Species Specification

Similar to the species specification for a continuum, this object gives you a choice of how to specify the proportions of components in the injected particle:

- Mass fraction—the ratio of the mass of a component over the total mass of the particle
- Mole fraction—the ratio of the molar amount of a component over the molar amount of the particle

Time Randomization

This object is available for non-DEM injectors in unsteady simulations. When activated, it causes the model to choose parcel injection time randomly within the time-step, rather than the default injection of all parcels at the start. This randomization can, for example, improve the appearance of animations where otherwise groups of parcels would appear to "march forward" together.

Time Randomization Properties

Random	Facilitates injection of parcels at random times.	
	Activated	Injection time for parcels is chosen randomly within the time-step.
	Deactivated	All parcels are injected at the start of the time-step.

Velocity Specification

The initial velocity of injected particles can be entered as a vector, or as a magnitude and direction, in which case the direction would be specified as a direction vector. You can also control the reference frame.

Velocity Specification Properties

Method	Specifies the method for entering velocity	
	Components	Enter the velocity as a vector with magnitude. This option adds a Velocity node to the Values node.
	Magnitude + Direction	Enter a direction vector and a magnitude. Two nodes are added to the Values node—Injection Direction and Velocity Magnitude.
Reference Frame	The injection velocity can be specified relative to the mesh or in the absolute frame.	
	Relative to Mesh	The velocity is relative to the mesh.
	Absolute	The velocity is absolute.

Injector Values Reference

Depending on how the injector conditions are set, you can define injections using the following values:

- Angular Velocity
- Axis
- Injection Cycle Limit
- Injection Direction
- Injector Diameter

- Injector Upstream Pressure
- Inner Cone Angle
- Mass Flow Rate
- Mass Flux
- Number of Particles
- Number of Seeds
- Outer Cone Angle
- Parcel Streams
- Particle Charge
- Particle Count
- Particle Diameter
- Particle Flow Rate
- Particle Flux
- Particle Mass
- Particle Orientation
- Particle Temperature
- Passive Scalar
- Point Inclusion Probability
- Porosity Limit
- Position
- Sheet Breakup Constant
- Spacing
- Species Mass Fraction
- Species Mole Fraction
- Swirl
- Velocity
- Velocity Magnitude
- Volume Flow Rate
- Volume Flux

Angular Velocity

Added to the tree when the DEM Particle type is selected. Vector notation is used to specify the angular velocity of injected particles where the magnitude and direction of the vector are equivalent to the angular rotation rate and the axis of rotation respectively.

Angular Velocity Properties

Coordinate System	The coordinate system in which the position is defined. Select from a drop-down list of available systems.	
Method	Selects the method for entering the angular velocity.	
	Composite	Set the angular velocity of each component separately. A Composite array node is added as a child to this node.
	Constant	Specifies the angular velocity as a single, constant value. A Constant node is added as a child to this node.
	Field Function	Defines the angular velocity using a field function (typically user defined). A field function method node is added as a child to this node.
	Normally Distributed	Specifies the angular velocity randomly using a normal distribution. A Normally distributed node is added as a child to this node.
	Table (iteration)	Defines the angular velocity as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (r)	Defines the force as a function of radius. A Table (r) node is added as a child to this node.
	Table (time)	Defines the angular velocity as a function of physical time. A Table (time) node is added as a child to this node.
	Table (x,y,z)	Defines the force as a function of position. A Table (x,y,z) node is added as a child to this node.

Table (xyz, time)	Defines the force as a function of position and time. A Table (xyz, time) node is added as a child to this node.
Uniformly distributed	Specifies the angular velocity randomly using a normal distribution. A Uniformly distributed node is added as a child to this node.
User Code	Adds a User Code child node that can be used to select user coding.

Dimensions	The dimensionality of the volume flow rate (read-only).

Axis

Added to the tree when the hollow cone, solid cone, or pressure swirl injector types are chosen. This value defines the central axis of the cone within which parcels are injected.

Axis Properties

Direction	The direction of the axis. Specified as a vector in the chosen coordinate system.
Origin	The origin of the axis (and therefore the point from which parcels are injected). Specified as a position vector in the chosen coordinate system.
Coordinate System	The coordinate system in which the axis is defined. Select from a drop-down list of available systems.

Cone Randomization

Added to the tree when the hollow cone or solid cone injector types are chosen, in an unsteady simulation. This determines whether injection points change position randomly from time-step to time-step, or remain in fixed positions.

Cone Randomization Properties

Rerandomize	Specifies whether the injection points are randomized again for each time-step.	
	Actived	Randomization again at each time-step.
	Deactivated	Re-use the initial randomization. This is the default.
Value	Proportion of the distance between initial positions over which injection points can vary between one time-step and the next, expressed as a value between 0 and 1. Effective only when Rerandomize is activated. The default value is 1.	

Injector Diameter

Added to the tree when the pressure swirl injector type is chosen. This value defines the diameter of the injector hole through which the droplet fluid flows.

Injector Diameter Properties

Value	The diameter of the injector hole.

Injection Cycle Limit

Added to the tree when the random injector type is chosen. This value sets the number of times the injector can try to insert the Number of Seeds. This value and the Porosity Limit work together to limit the random injector.

Injection Direction

Added to the tree when the Magnitude + Direction option is chosen as the Velocity Specification. It is entered as a vector profile.

Injector Upstream Pressure

Added to the tree when the pressure swirl injector type is chosen. This value sets the pressure head available to force liquid through the pressure swirl nozzle, and is used to predict the liquid sheet exit velocity.

Injector Upstream Pressure Properties

Value	The upstream pressure, which is used to calculate Δp in Eqn. (1402).

Inner Cone Angle

Added to the tree when the hollow cone or pressure swirl injector types are chosen. This value defines the internal angle of the cone about the central axis. A value of zero results in a solid cone formation.

Inner Cone Angle Properties

Value	The internal angle of the cone about the central axis.

Mass Flow Rate

Added to the tree when Mass Flow Rate is chosen as the Flow Rate Specification.

In most cases, specifying mass flow rate means that each parcel in the parcel streams has the same mass (unsteady) or mass flow rate (steady). For the surface injector type, A_f/A_s gives the scale for the mass or mass flow rate of each parcel. A_f is the area of the face through which the parcel enters and A_s is the total area of the surface injector.

Mass Flow Rate Properties

The following table applies to all injector types and distribution conditions except for those that are listed in the next table:

Method	Selects the method for entering the injector mass flow rate.	
	Constant	Specifies the mass flow rate as a single, constant value. A Constant node is added as a child to this node.
	Table (iteration)	Defines the mass flow rate as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (time)	Defines the mass flow rate as a function of physical time. A Table (time) node is added as a child to this node.
	User Code	Defines the mass flow using a user subroutine. A User Code node is added as a child to this node.

Dimensions The dimensionality of the mass flow rate (read-only).

The following table applies to the surface and part injectors, for Per Injector Point and Per Parcel Stream distributions:

Method Selects the method for entering the injector mass flow rate.

Constant	Specifies the mass flow rate as a single, constant value. A Constant node is added as a child to this node.
Field Function	Defines the mass flow rate using a field function (typically user defined). A field function method node is added as a child to this node.
Table (iteration)	Defines the mass flow rate as a function of iteration number. A Table (iteration) node is added as a child to this node.
Table (r)	Defines the mass flow rate as a function of radius. A Table (r) node is added as a child to this node.
Table (time)	Defines the mass flow rate as a function of physical time. A Table (time) node is added as a child to this node.
Table (x,y,z)	Defines the mass flow rate as a function of position. A Table (x,y,z) node is added as a child to this node.
Table (xyz, time)	Defines the mass flow rate as a function of position and time. A Table (xyz, time) node is added as a child to this node.
User Code	Defines the mass flow rate using a user subroutine. A User Code node is added as a child to this node.

Dimensions The dimensionality of the mass flow rate (read-only).

Mass Flux

Added to the tree when the injector has surface injector type and the Mass Flux option has been chosen for the Flow Rate Specification. It is entered as a scalar profile. The actual mass flow rate through each individual face of the injector surface is the product of the face surface area and the specified mass flux.

Number of Particles

Added to the tree when the random injector type is chosen and the Number of Particles option has been chosen for the Particle Amount Specification.

Number of Particles Properties

Value	Number of particles to inject in a single operation.

Number of Seeds

The number of seeds is the number of particles the random injector tries to inject at the same time. It does not define the total number of particles injected—the Particle Amount Specification sets that number. The Number of Seeds node is added to the tree when the random injector type is chosen and the Porosity Limit value is chose for the Particle Amount Specification.

For most cases, set this value as one tenth of the total number of particles that can be packed inside the volume of the injector region. Calculate the optimal number of seeds for a case as follows:

- Calculate the volume V_r of the injector region.
- Calculate the volume V_p of a single particle.
- Calculate the number of seeds as the number of particles that would fill 0.1 of the region volume as $\frac{V_r}{0.1V_p}$.

Number of Seeds Properties

Value	Number of seed particles to inject in a single operation.

Outer Cone Angle

Added to the tree when the hollow cone, solid cone, or pressure swirl injector types are chosen. This value defines the outer angle of the cone within which parcels are injected.

Outer Cone Angle Properties

Value	The external angle of the cone about the central axis.

Parcel Streams

This value sets the number of parcel streams per injection point (and per time-step in unsteady simulations). This quantity has no influence on the mass or volume flow rate from the injector, but instead can influence accuracy, depending on other selections. Specifically:

- With the Turbulent Dispersion model, increasing the number of parcel streams increases accuracy by creating a larger ensemble of tracks from each injection point, each representing a different realization of the particle trajectory. Typically 10–100 parcel streams are specified for this purpose.
- With particle size distributions, increasing the number of parcel streams increases the resolution of the particle size distribution. Typically 5-20 parcel streams are specified for this purpose.

The combined effect of both of the above, is multiplicative. If neither is relevant, use the default number of parcel streams (1). CPU time is directly proportional to this quantity.

Parcel Streams Properties

Value The number of parcel streams.

In steady simulations, Value is the number of parcel tracks originating from each injection point. In unsteady simulations, this number of parcels is introduced at each injection point at the start of each time-step.

The number of parcel streams can be a constant or derived from an expression. Expressions allow the value to vary with time or other conditions. Field functions are allowed.

All values, whether from expressions or constants, are rounded to the nearest integer. Values lower than 1 generate an error.

Particle Charge

Added to the tree when the Coulomb Force model is active in the Lagrangian phase that is associated with the injector. This value is entered as a scalar profile.

Particle Count

Added to the tree only when Injector Type is set to Table Injector and Table Type is set to Parcels Table. This value is the number of particles represented by a single parcel. It can be a constant or read from the injection table. The default value is 1. This node is not available with the DEM model, since each DEM parcel always represents one particle.

Particle Diameter

Added to the tree when Particle Size is chosen as the Particle Size Specification. The standard methods for a scalar profile are available for this value, along with more methods shown in the properties table below.

Particle Diameter Properties

Method	Selects the method for specifying the injector particle diameter. Common methods for a scalar profile can be used, as well as the additional options shown below.	
	Constant	Specifies the particle diameter as a single, constant value. A Constant node is added as a child to this node.
	Field Function	Defines the particle diameter using a field function (typically user defined). A field function method node is added as a child to this node.
	Log-Normal	Specifies a range of particle diameters using a Log-normal particle size distribution. A Log-Normal node is added as a child to this node.
	Normal	Specifies a range of particle diameters using a normal particle size distribution. A Normal node is added as a child to this node.
	Rosin-Rammler	Specifies a range of particle diameters using a Rosin-Rammler particle size distribution. A Rosin-Rammler node is added as a child to this node.
	Table (iteration)	Defines the particle diameter as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (particle size CDF)	Specifies a range of particle diameters using a tabulated particle size distribution. A Table (particle size CDF) node is added as a child to this node.

Table (r)	Defines the particle diameter as a function of radius. A Table (r) node is added as a child to this node.
Table (time)	Defines the particle diameter as a function of physical time. A Table (time) node is added as a child to this node.
Table (x,y,z)	Defines the particle diameter as a function of position. A Table (x,y,z) node is added as a child to this node.
Table (xyz,time)	Defines the particle diameter as a function of position and time. A Table (xyz, time) node is added as a child to this node.
User Code	Defines the particle diameter using a user subroutine. A User Code node is added as a child to this node.

Dimensions The dimensionality of the particle diameter (read-only).

Particle Flow Rate

Added to the tree when Particle Flow Rate is chosen as the Flow Rate Specification.

For all but the surface injector type, specifying particle flow rate means that each parcel in the parcel streams has the same particle count (unsteady) or particle flow rate (steady). For the surface injector type, A_f/A_s gives the scale for the mass or mass flow rate of each parcel. A_f is the area of the face through which the parcel enters and A_s is the total area of the surface injector.

Particle Flow Rate Properties

Method	Selects the method for entering the particle flow rate.	
	Constant	Specifies the particle flow rate as a single, constant value. A Constant node is added as a child to this node.
	Table (iteration)	Defines the particle flow rate as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (time)	Defines the particle flow rate as a function of physical time. A Table (time) node is added as a child to this node.
Dimensions	The dimensionality of the particle flow rate (read-only).	

Particle Flux

Added to the tree when the injector has surface injector type and Particle Flux is chosen as the Flow Rate Specification. It is entered as a scalar profile. The actual particle flow rate through each individual face of the injector surface is the product of the face surface area and the specified particle flux.

Particle Mass

Added to the tree when Particle Mass is chosen as the Particle Size Specification.

Particle Mass Properties

Method

Selects the method for specifying the injector particle mass. Common methods for a scalar profile can be used, as well as the additional options shown below.

Constant	Specifies the particle mass as a single, constant value. A Constant node is added as a child to this node.
Field Function	Defines the particle mass using a field function (typically user defined). A field function method node is added as a child to this node.
Log-Normal	Specifies a range of particle masses using a Log-normal particle size distribution. A Log-Normal node is added as a child to this node.
Normal	Specifies a range of particle masses using a normal particle size distribution. A Normal node is added as a child to this node.
Rosin-Rammler	Specifies a range of particle masses using a Rosin-Rammler particle size distribution. A Rosin-Rammler node is added as a child to this node.
Table (iteration)	Defines the particle mass as a function of iteration number. A Table (iteration) node is added as a child to this node.
Table (particle size CDF)	Specifies a range of particle masses using a tabulated particle size distribution. A Table (particle size CDF) node is added as a child to this node.
Table (r)	Defines the particle mass as a function of radius. A Table (r) node is added as a child to this node.

Table (time)	Defines the particle mass as a function of physical time. A Table (time) node is added as a child to this node.
Table (x,y,z)	Defines the particle mass as a function of position. A Table (x,y,z) node is added as a child to this node.
Table (xyz, time)	Defines the particle mass as a function of position and time. A Table (xyz, time) node is added as a child to this node.
User Code	Defines the particle mass using a user subroutine. A User Code node is added as a child to this node.

Dimensions The dimensionality of the particle mass (read-only).

Particle Orientation

Added to the tree when a non-spherical DEM particle shape is selected. Used to specify the initial orientation of the particle as it is injected.

Particle Orientation Properties

Method	Selects the method for entering the particle orientation.	
	Constant	Specifies the particle orientation as a single, constant value. A Constant node is added as a child to this node.
	Field Function	Defines the particle orientation using a field function (typically user defined). A field function method node is added as a child to this node.
	Normally distributed	Specifies the particle orientation randomly using a normal distribution. A Normally distributed node is added as a child to this node.
	Table (iteration)	Defines the particle orientation as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (time)	Defines the particle orientation as a function of physical time. A Table (time) node is added as a child to this node.
	Uniformly distributed	Specifies the particle orientation randomly using a normal distribution. A Uniformly distributed node is added as a child to this node.
Dimensions	The dimensionality of the particle orientation (read-only).	

Particle Temperature

Added to the tree when the Lagrangian Energy model is active in the Lagrangian phase that is associated with the injector. This value is entered as a scalar profile.

Passive Scalar

Added to the tree when the Passive Scalar model is active in the Lagrangian phase that is associated with the injector.



Properties

Method — Selects the method used for setting the value of the scalar on the particle as the particle enters the simulation.

- Constant — Set the values of all individual scalars as a vector of constants.
- Composite — Set the value of each passive scalar separately (constant, field function, table or user code). If you use a table injector, individual passive scalars can be initialized by an injection table.
- Injection Table — Set the value of the passive scalar for each component according to an injection table. The table specifies the location of injection points, the time of injection, and the passive scalar value of each component separately. Used only with table injectors.

Dimensions — The dimensionality of the value, Mass/Volume-Time (read-only). The passive scalar itself is always dimensionless, but its source term has units referring to the density of the receiving phase over time.

Point Inclusion Probability

This node is added to the tree when the injector has part injector type, and sets the probability of a point being included in the set of points from which particles are injected. In the example that is shown below, a line probe derived part with a resolution of 11 is used as the input part for the injector. This configuration means that 10 points are available from which to inject particles. In the first screenshot, the point inclusion probability is set to 1.0, which means that the injector uses all points that are extracted from the part. In the second screenshot, the point inclusion probability is set to 0.4, meaning that only a subset of the points is used. In the third screenshot, the Re-randomize property is activated, regenerating the injection points at the start of each time-step.

Point Inclusion Probability = 1.0



Point Inclusion Probability = 0.4



Point Inclusion Probability = 0.4, with re-randomizing option activated



Point Inclusion Probability Properties

Value	The probability that a point in the injector is used as an injection point. The probability is applied to each part point.
Re-randomize	Controls how injection points are regenerated (available for transient analyses only).
	Activated: The injection points are regenerated at the start of each time-step.
	Deactivated: The injection points are generated once at the start of the simulation and are then fixed (as in a steady analysis).

Porosity Limit

The random injector porosity is defined as:

$$\text{Porosity} = \frac{\text{Total Volume of Injected Partices}}{\text{Volume of Regions Assigned to the Injector}}$$

The porosity limit defines the maximum packing limit for particles that are introduced into the injected region. As the injector injects particles that are composed of spheres, the porosity is always less than 1. For example, when

using single spheres the maximum value is 0.74. STAR-CCM+ stops injecting particles if the limit is exceeded, or if there is no space for further particles in the region.

This injector value is available when you select the random injector type.

Porosity Limit Properties

Value	Specifies the porosity limit for the random injector, which acts as an injection stopping condition. The porosity is the ratio of the total injected particles volume to the region volume.

Position

This node is added to the tree when the injector has point injector type, to define the position of the injector.

Position Properties

Coordinate System	The coordinate system in which the position is defined. Select from a drop-down list of available systems.
Value	The value and units of the position. Clicking the customizer executes the define value action. An expression is not accepted for this entry.

Sheet Breakup Constant

Added to the tree when the pressure swirl injector type is chosen. This value is an empirical constant that calibrates the ratio of the sheet surface disturbance at breakup to its initial value in the dispersion relation of the LISA model.

Sheet Breakup Constant Properties

Value	Value of the empirical constant ξ in Eqn. (1409). It is set to 12 by default.

Spacing

Added to the tree when the lattice injector type is chosen. This value allows you to control the dimensions of the bounding box in the x, y, and z directions.

Spacing Properties

Value	The value and units of the lattice injector position. Clicking the customizer executes the define value action. An expression is not accepted for this entry.

Species Mass Fraction

Added to the tree when the Lagrangian Species model is active in the Lagrangian phase that is associated with the injector. Also added when the Method property of the Species Specification injector condition is set to `Mass Fraction`. The mass fraction of a component in the particle is the ratio of its mass over the total mass of the particle. This object functions in a manner similar to the species mass fraction of a continuum.

Mass Fraction Properties

Method	Selects the method for specifying component mass fractions through the injector.	
	Constant	Set the mass fraction of each component in the order they appear in the mixture. A Constant array node is added as a child to this node.
	Composite	Set the mass fraction of each component separately. A Composite array node is added as a child to this node.
	Injection Table	Set the mass fraction of each component according to an injection table. The table specifies the location of injection points, the time of injection, and the mass fraction of each component separately. Used only with table injectors.
Dimensions	The dimensions of the mass fraction (read-only).	

Species Mole Fraction

Added to the tree when the Lagrangian Species model is active in the Lagrangian phase that is associated with the injector. Also added when the Method property of the Species Specification injector condition is set to `Mole Fraction`. The mole fraction of a component in the particle is the ratio of the molar amount of the component over the total molar amount of the particle. This object functions in a manner similar to the species mole fraction of a continuum.

Mole Fraction Properties

Method	Selects the method for specifying component mole fractions through the injector.	
	Constant	Set the mole fraction of each component in the order they appear in the mixture. A Constant array node is added as a child to this node.
	Composite	Set the mole fraction of each component separately. A Composite array node is added as a child to this node. If you use a table injector, individual species mass/mole fractions can be initialized by an injection table.
	Injection Table	Set the mole fraction of each component according to an injection table. The table specifies the location of injection points, the time of injection, and the mole fraction of each component separately. Used only with table injectors.
Dimensions	The dimensions of the mole fraction (read-only).	

Swirl

Added to the tree when the hollow cone injector type is chosen. This value defines additional angular velocity about the axis, added to existing velocity.

Swirl Properties

Value	Scalar value of the swirl. It is set to zero by default.

Velocity

Added to the tree when the Components option is chosen as the Velocity Specification. This value is entered as a vector profile.

If the particle injection velocity is to correspond to the fluid velocity at the injection point, take the following steps:

- In the *Properties* window for the Values > Velocity node, select Field Function for the Method property. A Field Function node is added as a sub-node to the Velocity node.

- In the *Properties* window for the Field Function node, select Velocity for the Vector Function property as shown below.



Velocity Magnitude

Added to the tree when the Magnitude + Direction option is chosen as the Velocity Specification. This value is entered as a scalar profile.

Volume Flow Rate

Added to the tree when Volume Flow Rate is chosen as the Flow Rate Specification.

For all but the surface injector type, specifying volume flow rate means that each parcel in the parcel streams has the same volume (unsteady) or volume flow rate (steady). For the surface injector type, A_f/A_s gives the scale for the volume or volume flow rate of each parcel. A_f is the area of the face through which the parcel enters and A_s is the total area of the surface injector.

Volume Flow Rate Properties

Method	Selects the method for entering the volume flow rate.	
	Constant	Specifies the volume flow rate as a single, constant value. A Constant node is added as a child to this node.
	Table (iteration)	Defines the volume flow rate as a function of iteration number. A Table (iteration) node is added as a child to this node.
	Table (time)	Defines the volume flow rate as a function of physical time. A Table (time) node is added as a child to this node.
Dimensions	The dimensionality of the volume flow rate (read-only).	

Volume Flux

Added to the tree when the injector has surface injector type and Volume Flux is chosen as the Flow Rate Specification. It is entered as a scalar profile. The actual volume flow rate through each individual face of the injector surface is the product of the face surface area and the specified volume flux.

Working with Particle Tracks

A track records the history of a parcel. This history is the locus of parcel positions which can be visualized much like a streamline. STAR-CCM+ allows other data to be recorded too, such as particle velocity and temperature, which can be used for subsequent visualization and analysis. Because there is the potential to generate a large amount of data, tracks can be recorded in a dedicated file under your control.

The particle track represents the track of one Lagrangian phase that is taken from a track file. A node of the particle track (named Continuum 1: Phase 1 tracks in the following screenshot) has its own properties.



This object is created automatically when you create a track file, and cannot be deleted directly. It is deleted only when you delete the track file that is associated with it.

You can, however, rename this object.

Particle Track Properties

Track File	The track file from which this part is derived (read-only).
Phase Id	The ID number of the Lagrangian phase with which the track file is associated (read-only).
Phase Name	The names of the continuum and Lagrangian phase with which the track file is associated when created (read-only).

See also: common menu items.

What Is a Track File?

A track file records the tracks of parcels. The Lagrangian Multiphase model (specifically in the Lagrangian Track File model) gives the option to create it. An external file is used because there is the potential to generate—and regenerate—a large amount of data during the simulation. This data is

saved for subsequent visualization and analysis. It is represented in the object tree by a track file node (intake in the following screenshot), which has properties and a pop-up menu and can be renamed or deleted.



Track File Properties

Path	The location of the external data source of the track file. This path is relative to the location of the simulation file. Edit to change to a different track file.

Track File Pop-Up Menu

Reload File...	Rereads the track file data from the external file.

See also: common menu items.

Modeling with a Track File

This section describes the Track File model and how to select it.

What Is the Track File Model?

The Track File node, which has properties, represents the Track File model. This model indicates that the tracks of parcels are recorded in a track file and, if so, what variables are recorded.

A temporary track file only is created while the simulation runs. It is essential to save the simulation so that this temporary track file is moved to a final file having the same name as the simulation. Only the final file can be brought into STAR-CCM+ for analysis, as discussed in section, Creating a Track File.



Every time a track file entry is written, the values of all the user-selected variables are written, together with the following:

- Parcel Centroid, the instantaneous parcel centroid.
- Parcel Index.
- Time (unsteady) or Particle Residence Time (steady).
- Particle Count (unsteady) or Particle Flow Rate (steady).

Hence, you do not need to select the field functions corresponding to these variables. A parcel generates a track file entry at the following events:

- Injection.
- Crossing an internal face.
- Interacting with a boundary face.

Track File Model Properties

Scalars	Selects a set of scalar field functions, with a Select Objects dialog, for recording on the track file. Three scalars that are automatically selected, Parcel Index, Time (or Particle Residence Time in steady simulations), and Particle Count (or Particle Flow Rate in steady simulations). They appear in the *Selected* list in grayed-out text, since they cannot be removed from the selector.
Vectors	Selects a set of vector field functions, with a Select Objects dialog, for recording on the track file. One vector that is automatically selected, Parcel Centroid appears in the *Selected* list in grayed-out text, since it cannot be removed from the selector.

Selecting the Track File Model

Select the Track File option from the *Track Sampling* group, as shown below.



Creating a Track File

To create an external track file (`.trk`), do the following:

- Select the Lagrangian Track File model.
- Step the solution at least one iteration or time-step. Stepping creates a temporary track file. The extension of this file is the first available in the series `.trk.tmp`, `.trk.tmp2`, `.trk.tmp3`, and so on.

- Save the simulation file. This action replaces the temporary file with a permanent track file (`.trk`).

By default, the external track file has the same base file name as the simulation file. If this name is already in use, for example by another server, then the base file name is appended with a unique number. The format of the track file is unique to STAR-CCM+.

Modeling with Boundary Sampling

This section describes the Boundary Sampling model and how to select it.

What Is the Boundary Sampling Model?

The Boundary Sampling node, which has properties, represents the Boundary Sampling model. This model records parcel states when they impact particular boundaries.



Every time a parcel impacts one of the designated boundaries, the values of all the user-selected variables are stored, together with the following:

- Parcel Centroid, the instantaneous parcel centroid.
- Parcel Index.
- Time (unsteady) or Particle Residence Time (steady).
- Particle Count (unsteady) or Particle Flow Rate (steady).

The Boundary Sampling model differs from the Track File model in that it does not have a file that is associated with it. Sampled quantities are stored in memory, and are accessed through particle track parts that are added to the Particle Tracks folder when boundaries are added to the Boundaries property of the model node. One particle track part is added for each boundary that is selected to participate in sampling.

Boundary Sampling Model Properties

Boundaries	The set of boundaries for which particle sampling is active.
Scalars	Selects a set of scalar field functions, with a Select Objects dialog, for recording on the boundary. Three scalars that are automatically selected, Parcel Index, Time (or Particle Residence Time in steady simulations), and Particle Count (or Particle Flow Rate in steady simulations) appear in the *Selected* list. They are in grayed-out text, since they cannot be removed from the selector.
Vectors	Selects a set of vector field functions, with a Select Objects dialog, for recording on the boundary. One vector that is automatically selected, Parcel Centroid appears in the *Selected* list in grayed-out text, since it cannot be removed from the selector.

Selecting the Boundary Sampling Model

Select the Boundary Sampling option from the *Track Sampling* group, as shown below.



Selecting the DEM Boundary Sampling Model

If you are using the Discrete Element Method (DEM) and the Boundary Sampling Model, the DEM Boundary Sampling Model becomes available. This model makes the Impact Velocity field function available and tracks particle velocity at the initial impact point for the selected surfaces.



Mapping Cell and Boundary Field Functions to Parcels

Both cell and boundary field function data can be mapped to parcels any time a field function is required on parcels, as shown below.



***Track files* and *boundary sampling* facilitate recording field functions of cells and incident boundaries in particle tracks. However, field function mapping can also be used:**

- to evaluate user-defined source terms,
- to evaluate composite boundary interaction,

- when displaying scenes or evaluating reports, and so on.

Mapping Cell Data to Parcels

When a cell field function is mapped to parcels, each parcel gets the value from its host cell.

When a vertex field function is mapped to parcels, each parcel gets the arithmetic average of the vertex values of the vertices belonging to its host cell.

Mapping Boundary Data to Parcels

When a boundary field function is mapped to parcels, each parcel just gets the value of its host boundary face, but only if the latter is defined. Parcels which are intersecting a boundary have a host face (the face they are intersecting), whereas parcels which are in the interior do not have a host face. For the latter, a null value is mapped to the parcels for the boundary field function.

What Is the Particle Tracks Manager?

The Particle Tracks manager node, contains all of the particle track nodes in the object tree. The particle track manager can hold any number of particle tracks up to the memory limits of the computer.



Particle tracks get their data from track files. The Track Files node, which has its own pop-up menu, is the manager object for working with track files in the simulation tree. It is added to the tree once you select the Lagrangian Multiphase model. The track files manager can hold any number of track files up to the memory limits of the computer.



Reading a Track File

To read a saved track file, right-click the Track Files manager node and select Track File...



Select an available track file from the *Open* dialog, then click Open. This action adds a track file to the object tree using the data from the specified file.



The path to the track file is relative to the location of the simulation file.

The new track file appears in the Track Files node. It is initially assigned a name corresponding to the file name (without the `.trk` extension), but can be renamed.



Rereading a Track File

If the external file has been modified since you read it into STAR-CCM+, it is easy to update the track file object. Simply right-click the node and select Reload file...



Using Particle Tracks in Analysis

It is possible to visualize a particle track in a scene.



Particle tracks function similar to any other parts in a simulation in the following ways:

- You can manage the particle tracks through the scene, or through an individual displayer as shown in the following screenshot.



- A particle track can also serve as an input part for certain derived parts: a cell surface, section, threshold, or iso surface.
- It is possible to select particle tracks when specifying the parts for an xyz internal table.

When you specify the scalar for an xyz internal table, you can also access specific field functions that correspond to the particle tracks.



Example of Using Particle Track Data

When a track file is read into a simulation, STAR-CCM+ creates field functions for each of the scalar and vector variables that was chosen in the Track File model. For example, if scalar variables Parcel Index, Particle Count,

and Particle Surface Area were selected for the Scalars property of the node of the Track File model, the corresponding field functions Track: Parcel Index, Track: Particle Count and Track: Particle Surface Area would appear when the track file is created. These functions become available in appropriate scalar and vector function selectors throughout the simulation.

As an example of using particle track data, suppose that it was necessary to plot the distribution of parcels crossing the plane section that is shown below.



The steps that are required to obtain such a plot are as follows:

- Define a derived part that is based on the particle tracks in to obtain a discrete set of points.
- Create an XY Plot and set its Part property to the newly created derived part.
- Choose appropriate particle track variables for the X and Y scalar functions.

Defining a Derived Part Based on Particle Tracks

In this example, the steps for creating a plane section are:

- Make an appropriate scene visible in the *Graphics* window. Right-click the Derived Parts node and choose New Part > Section > Plane...
- In the *Create Section* dialog, remove the default region from the selector in the *Input Parts* group box, and replace it with the particle tracks object belonging to a Lagrangian phase. More than one particle tracks object could be selected if available.
- Set the other properties of the plane section as desired. A completed dialog is shown below.



- Click Create to create the plane section.
- Rename the newly created derived part node if desired. In this example, the node is called Track-based section.

Creating an XY Plot with a Part Based on Particle Tracks

- Right-click the Plots node and choose New Plot > XY Plot from the pop-up menu.
- Select the newly created plot node (XY Plot 1 in this example) and, in the *Properties* window, click to the right of the Parts property so that the *Select Objects* dialog appears.

- In the *Select Objects* dialog, move the plane section that was created previously to the *Selected* list and click OK.

  The properties of the XY Plot 1 node shows the selection in the Parts property.



Selecting X and Y Functions Based on Particle Tracks

The X axis of the plot is set to show the X coordinates of particles crossing through the plane, and the Y axis of the plot shows the Y coordinates. In this way, a cross-sectional distribution of particle positions is obtained.

- Select the Plots > XY Plot 1 > X Type node and, in the *Properties* window, set the Type property to Scalar.

The child node is changed to Scalar instead of Position.



- Select the X Type > Scalar node and, in the *Properties* window, set the Scalar property to Track > Position > Laboratory > X.



- Select the Plots > XY Plot 1 > Y Types node and, in the *Properties* window, activate Smooth Values.

This step is necessary, since particle tracks do not store cell center values.



- Select the Y Types > Scalar node and, in the *Properties* window, set Scalar to Track > Position > Laboratory > Y.



The final plot is shown below. Chart properties have been set so that both axes have the same dimensions and scale.



Other parcel and particle quantities could be plotted in a similar way, provided they are available from the track file.

Animating Particle Tracks

Particle tracks can be animated in a scene just like streamlines.

First, create a streamline displayer in a scene. Open its node and select the Animations sub-node. Set the Animation Mode property to `Tracers`.



This action adds the Streamline Settings sub-node to the Animations node. This sub-node allows you to control the appearance of the tracers.



As you choose your settings, set the Tail Length property to zero to make this displayer a particle animation.

Once your animation settings are ready, make your streamline displayer access the particle track data. Select the Parts sub-node of the streamline displayer node.



In the *Properties* window, activate the in-place object selector by clicking the right column of the Parts property. Select the particle tracks and click OK.



Once you have selected a particle track, specify what scalar function to display. Select the Scalar Field sub-node of the scalar displayer node.



Set the Function property to `Track > Particle Residence Time`.



With the function selected, make sure that the scene display is open and active, and then click the (Play/Pause Animation) toolbar button to start the animation. The toolbar contains a dynamic set of buttons for controlling the movement.



Once you have finished, you can record the animation onto separate files.

Pop-Up Menu of the Track Files Manager

Track File	Executes the read track file action, and so adds a track file to the Track Files node.

Load Balancing for Lagrangian/Eulerian Flows

For parallel processing, the traditional approach to decomposing the volume mesh is to give each processor approximately the same number of cells. For cases where the same Eulerian flow equations are solved in each cell, this approach generally gives good performance.

For cases containing particles, this traditional form of decomposition can lead to imbalanced CPU loads across the parallel cores. If one or two partitions contain most of the particles, then those partitions require

significantly more compute time than partitions with few particles. The traditional decomposition does not account for this difference. See the example in the following figure, where Eulerian mesh-based domain decomposition leads to poor CPU load and low parallel efficiency.



Dynamic Load Balancing for Lagrangian Simulations

STAR-CCM+ provides an alternative approach to domain decomposition that accounts for the presence of particles. There are two key components involved:

- The partitioning solver:

  The main load-balancing solver. This solver rebalances the domain decomposition in response to information provided by physics solvers. It also rebalances the decomposition in response to mesh topology changes, such as those caused by interface intersections. See The Partitioning Solver.

- The Lagrangian load-balancing solver:

  This solver provides the partitioning solver with information about the computational load that is required to solve particle physics in each Eulerian cell. See The Lagrangian Load-Balancing Solver.

These two solvers operate together. The Lagrangian solver collects the information that is required to rebalance the domain and passes it to the partitioning solver. The partitioning solver coordinates the final load balancing, accounting for all load-balancing criteria as it does so. This approach optimizes the domain decomposition in response to changing particle distributions. The computational load across all cores remains balanced as particles migrate throughout the domain.

Running with Load Balancing Activated

To run with load balancing, use the following steps:

1. Load the simulation into a parallel server. See Working with Parallel Servers.
2. Activate the Lagrangian load-balancing solver. See The Lagrangian Load-Balancing Solver.
3. Start the run.

The load-balancing framework must complete at least one cycle before the load-balancing information can be collected and used to rebalance mesh. The first iteration or time step therefore runs with the original domain decomposition based on cell count.

Understanding the Lagrangian Multiphase Solver

This solver, which you can adjust with properties, controls the update of the Lagrangian Multiphase model, according to the solution methodology.



Depending on model selections, sub-nodes can be made available:

- When you select Steady as your time model, the Lagrangian Steady solver is made available as a sub-node.
- When you select Implicit Unsteady as your time model, the Lagrangian Implicit Unsteady solver is made available as a sub-node.
- When you select Explicit Unsteady as your time model, the Lagrangian Explicit Unsteady solver is made available as a sub-node.
- When you select the Two-Way Coupling model for a Lagrangian phase, the Two-Way Coupling solver is made available as a sub-node.
- When you select the Lagrangian Multiphase model, the Lagrangian load-balancing solver is made available as a sub-node.

Lagrangian Multiphase Solver Properties

Maximum Courant Number	The value of Co_{max} in Eqn. (1313) used to provide an upper bound for the local time-step. Reducing this value can increase accuracy, at the expense of CPU time.
Minimum Courant Number	The value of Co_{min} in Eqn. (1314) used to provide a lower bound for the local time-step. Reducing this value can increase accuracy, at the expense of CPU time.
Maximum Parcel Count	The maximum number of parcels per Lagrangian phase. If this limit is exceeded, the simulation throws an error. A value of 0 (the default) means unlimited.

Lagrangian Multiphase Solver Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating. Do not activate this feature as your first choice in most situations, since it is primarily a "debugging" tool and not supported for normal simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.
Temporary Storage Retained	Retains or deactivates temporary storage at the end of the iteration. Examples include the Lagrangian source terms that become active in the Two-Way Coupling model. These quantities become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.

The Lagrangian Steady Solver

This solver, which you can adjust with properties, is available as a sub-node of the Lagrangian Multiphase solver when you use the Steady time model. It provides properties relevant to the steady solution methodology of the Lagrangian Multiphase solver.



Lagrangian Steady Solver Properties

Update Frequency	Sets the number of updates of the continuous phase solution per update of the Lagrangian Multiphase solution. Increasing this value can save CPU time.
First Iteration	Specifies the first iteration in which the Lagrangian solver runs.
Maximum Residence Time	Sets the residence time beyond which particles are not tracked, to prevent particles being tracked indefinitely. Set this quantity high enough to avoid prematurely terminating tracks.
Verbosity	Controls the level of detail that the solver presents is in the *Output* window.

None	No output at all.
Low	If any tracks reached maximum residence time or maximum substeps, a single message is printed at the end of the Lagrangian solver stage. This setting is the default.
High	Same as Low, plus an update every 1000 substeps of the number of active parcels and the minimum residence time observed among those parcels.

Maximum Sub-Steps	Sets the maximum number of substeps for which particles are tracked, to prevent particles being tracked indefinitely. Set this quantity high enough to avoid prematurely terminating tracks.
Sub-Step Reporting Frequency	Sets the number of substeps between each instance of reporting solver progress.

The Lagrangian Implicit Unsteady Solver

This solver, which you can adjust with properties, is available as a sub-node of the Lagrangian Multiphase solver when you use the Implicit Unsteady time model. It provides properties relevant to the unsteady solution methodology of the Lagrangian Multiphase solver.



Lagrangian Implicit Unsteady Solver Properties

Update Frequency	Selects the frequency of updates of the continuous phase solution per update of the Lagrangian Multiphase solution. In scenes containing Lagrangian data, do not set the update policy to `Iteration` when this Lagrangian solver is set to update after the last iteration. Otherwise, the scene updates before the Lagrangian solver does. Recommended — Runs the solver every iteration when there are two-way coupled phases, or after the last iteration when there are only one-way coupled phases.

	Once per Time-Step	Runs the solver in the first iteration when there are two-way coupled phases, or after the last iteration when there are only one-way coupled phases.
	Every Iteration	Runs the solver in every iteration.
Verbosity	Controls the level of detail that this solver presents in the *Output* window.	
	None	No output at all.
	Low	A single message is printed at the end of the time-step when any tracks reached maximum substeps. Low is the default setting.
	High	Same as Low, plus an update every 100 substeps of the number of active parcels and the minimum percentage progress observed for those parcels. Progress is based on the ratio of the current tracking time to the overall time-step.
Maximum Sub-Steps	Sets the maximum number of substeps for which particles are tracked, to prevent particles being tracked indefinitely. Set this quantity high enough to avoid prematurely terminating tracks.	
Sub-Step Reporting Frequency	Sets the number of substeps between each instance of reporting solver progress.	

The Lagrangian Explicit Unsteady Solver

This solver, which you can adjust with properties, is available as a sub-node of the Lagrangian Multiphase solver when you use the Explicit Unsteady time model. It provides properties relevant to the unsteady solution methodology of the Lagrangian Multiphase solver.



Lagrangian Explicit Unsteady Solver Properties

Maximum Sub-Steps

Sets the maximum number of substeps for which particles are tracked, to prevent particles being tracked indefinitely. Set this quantity high enough to avoid prematurely terminating tracks.

Verbosity

Controls the level of detail that this solver presents in the *Output* window.

None	No output at all.
Low	If any tracks reached maximum sub-steps, a single message is printed at the end of the time-step. This setting is the default.
High	Same as Low, plus an update every 100 substeps of the number of active parcels and the minimum percentage progress observed for those parcels. Progress is based on the ratio of the current tracking time to the overall time-step.

The Two-Way Coupling Solver

This solver, which you can adjust with properties, is available as a sub-node of the Lagrangian Multiphase solver when you use the Two-Way Coupling model. It controls the under-relaxation and linearization of the sources appearing in the continuous phase equations due to the Lagrangian phases.



Two-Way Coupling Solver Properties

Under-Relaxation Factor	The fraction of the new source values which are used when updating source values. Lower values can stabilize simulations with strong coupling between the phases.
Maximum Volume Fraction	Sets the maximum volume fraction apparent to the continuous phase.

Under-Relaxation Factor Ramp Properties

Ramp Method	Activates or deactivates the ramp for calculation.	
	No Ramp	Leaves off ramp calculation.
	Linear Ramp	Activates linear ramp value calculation. The Linear Ramp node becomes available in the Under-Relaxation Factor Ramp node.

Expert Properties

Linearize Sources	Indicates when, if possible, to split the sources into explicit and implicit components. Typically, splitting stabilizes simulations with strong coupling between the phases.
Limit Sources	When the current volume fraction exceeds the maximum volume fraction and this property is activated, the two-way coupling sources are limited proportionately to (maximum volume fraction)/(current volume fraction). This limitation can improve stability. There is no effect when the current volume fraction is less than the maximum.

The Lagrangian Load-Balancing Solver

This solver maintains and controls the Lagrangian component of the load-balancing framework. See Load Balancing for Lagrangian/Eulerian Flows.

Lagrangian Load-Balancing Solver Properties

Activated	When activated, Lagrangian load-balancing is turned on for parallel servers. See Load Balancing for Lagrangian/Eulerian Flows.
Parcel Weight	Controls the relative weight of the parcel versus the cell. This parameter allows you to control the balance of Lagrangian physics versus Eulerian physics. The weight can range from 1 (the default), as Lagrangian as possible, to 0, strictly Eulerian.
Rebalance Limit	Sets the point at which the load balance is even enough and the solver stops. The limit is the ratio of the loading between the most and least loaded domains. Choose a value of the limit that avoids rebalancing efforts beyond useful levels. The lower the limit, the less frequent the rebalancing. Setting the value close to 1 results in 'always rebalance' behavior. The default value is 0.8.

The Lagrangian Passive Scalar Solver

The Lagrangian Passive Scalar Solver node is activated with the Passive Scalar model.



Properties

None.

Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver does update the solution.

Solution Methodology

The solution of equations in a Lagrangian framework implies the time-marching of ordinary differential equations (ODEs). The set of equations depends on the models that are invoked for the Lagrangian phase in question; the one common element is an ODE for the position of each parcel, Eqn. (1316).

This section describes the methodology that is used to solve the system of equations that the Lagrangian Multiphase and the Lagrangian Phase models generate. It covers:

- Unsteady Procedure
- Steady Procedure
- Time-Step Controls

Unsteady Procedure

A time-step in an unsteady simulation advances the solution from time t to time $t + \Delta t$. A single iteration of the unsteady solver for the Lagrangian Multiphase model involves:

1. Recovering the Lagrangian Multiphase solution at time t.
2. Advancing the Lagrangian Multiphase solution to time $t + \Delta t$ by time-marching. A local time-step is used for each parcel. The time-step is adjusted dynamically according to the time-step controls in force.

These two steps are executed in each iteration of the unsteady solver, before the flow solver. If the Two-Way Coupling model is active, source terms that the Lagrangian phase models compute are stored for subsequent application in their respective transport equations.

In an explicit unsteady case, this procedure means that the Lagrangian Multiphase solution is obtained using the continuous phase solution from the previous time-step.

Steady Procedure

An iteration in a steady simulation updates the steady solution. For the Lagrangian Multiphase model, the steady solution consists of

- Tracks, generated by the Track File model.
- Cell fields such as volume fraction and sources.
- Boundary fields such as incident mass flux.

A steady state is attained when these entities become invariant with further iteration. Despite this steady characteristic of the solution, it results from time-marching, where the time variable is the residence time of the parcels. The procedure in each iteration is:

1. Deactivate the Lagrangian Multiphase solution.
2. Generate the Lagrangian Multiphase solution by time-marching each parcel until it has left the computational domain, or has been removed, or until the user-specified maximum residence time is reached. A local time-step is used for each parcel, which is adjusted dynamically according to the time-step controls in force.

These two steps are executed in each iteration of the steady solver, before the flow solver. If the Two-Way Coupling model is active, source terms that the Lagrangian phase models compute are stored for subsequent application in their respective transport equations.

If the Two-Way Coupling model is not active, it is possible to freeze the Lagrangian Multiphase solver until convergence is reached, using its Solver Frozen property, potentially saving significant CPU time. The Lagrangian Multiphase solution can then be obtained by unfreezing the solver and stepping the solution. Even when the Two-Way coupling model is active, an economic route to a steady solution can be to update the Lagrangian Multiphase solution less frequently than the other solvers. Use the Update Frequency property of the Lagrangian Steady solver. However, do not use the Solver Frozen feature as your first choice in other situations, since it is primarily a "debugging" tool and not supported for normal simulation work.

Time-Step Controls

The local time-step δt_p controls the accuracy of the Lagrangian Multiphase solution. This time-step is used to time-march each parcel. The time-step is calculated dynamically using user-defined bounds and model time-scales.

You can use the Parcel Time Step field function to obtain the local time-step in field function expressions that are evaluated during substepping, for example in a passive scalar source. STAR-CCM+ can use the Parcel Time Step field function for analysis or visualization when temporary storage is turned on or when the Lagrangian Multiphase solver is running.

User-Defined Bounds

Three parameters are available for you to control the local time-step δt_p. Smaller δt_p can increase accuracy, until temporal error becomes negligible. The three parameters are:

- The physical time-step Δt in unsteady simulations. Clearly $\delta t_p \leq \Delta t$.
- The maximum Courant number Co_{max}, which implies

$$\delta t_p \leq \frac{\mathrm{Co}_{max}\Delta x}{\max(|\mathbf{v}|, |\mathbf{v}_p|)} \tag{1313}$$

  where Δx is a characteristic length-scale of the cell containing the parcel.
- The minimum Courant number Co_{min}, which implies

$$\delta t_p \geq \frac{\mathrm{Co}_{min}\Delta x}{\max(|\mathbf{v}|, |\mathbf{v}_p|)}. \tag{1314}$$

Model Time-Scales

Models can also propose constraints for δt_p. Among the model time-scales are:

- The momentum relaxation time-scale, Eqn. (1334), provided by the Drag Force model.
- The thermal relaxation time-scale, Eqn. (1359), provided by the Lagrangian Energy model.
- The residual of the eddy interaction time, Eqn. (1387), provided by the Turbulent Dispersion model.
- The mass transfer time-scale $m_p/\dot{m}_p$ provided by the Droplet Evaporation model.

Using the Lagrangian Multiphase Model with Time Models

The Lagrangian Multiphase model is compatible with all Time models except Harmonic Balance.



Unsteady Simulation

In unsteady simulations, the approach can be categorized as a *discrete element method* [237]. In this method, the solution at a given instant in time is the state (position, velocity, temperature, and so on.) of each parcel. The number of nominally identical particles that a parcel represents, or *particle count*, is part of the state of the parcel. Optionally, the tracks of parcels can also be recorded in unsteady simulations: these tracks are the history of the parcel state, which can be visualized as a trajectory.

Steady Simulation

In steady simulations, the approach can be categorized as a *trajectory method* [237]. The solution is typically the tracks of the parcels, and the variation of state (velocity, temperature, and so on) along them. Each track records the complete lifetime of one parcel, from when it is injected to when it leaves the domain. The *particle flow rate* along each track is part of its state.

Parcels are indeed supposed to leave the domain in steady simulations. This approach is unsuitable for configurations that permanently carry particles in the continuous phase, whether through flow features (for example, recirculation zones) or geometry (for example, mixing vessels). For these configurations, an unsteady simulation can be more appropriate.

Using the Lagrangian Multiphase Model with Motion Objects

- The Lagrangian Multiphase model is compatible with all motion objects. However, a simulation using the Discrete Element Method (DEM) can only use rigid motion (Rotation or Rotation and Translation) objects.

Lagrangian Multiphase Formulation

This section explains the following formulation of the Lagrangian Multiphase model:

- Basic Equations of Motion
- Massless Particle Equation of Motion
- Mass Balance for a Material Particle
- Mass Transfer
- Component Mass Balance for a Material Particle
- Component Mass Transfer
- Momentum Balance for a Material Particle
- Drag Coefficient
- Field Charging
- Film Damping
- Momentum Transfer
- Energy Balance for a Material Particle
- Heat Transfer Coefficient
- Energy Transfer
- Quasi-Steady Single-Component Droplet Evaporation
- Multi-Component Droplet Evaporation
- Sherwood Number
- Turbulent Dispersion
- Volume Fraction
- Erosion

- Boundary Interaction Modes
- Particle Size Distributions
- TAB Distortion
- KHRT Droplet Breakup
- Reitz-Diwakar Droplet Breakup
- TAB Droplet Breakup
- Bai-Gosman Wall Impingement
- Satoh Wall Impingement Model
- Impingement Heat Transfer
- Collision Detection Model
- Passive Scalar Source Terms
- Passive Scalar Interactions

A bibliography is also provided.

Basic Equations of Motion

The most basic particle description involves only its position $\mathbf{r}_p(t)$ and velocity $\mathbf{v}_p(t)$. These two quantities relate through the equation of motion

$$\frac{\mathrm{d}\mathbf{r}_p}{\mathrm{d}t} = \mathbf{v}_p - \mathbf{v}_g . \tag{1315}$$

The grid velocity $\mathbf{v}_g(\mathbf{x}, t)$ is evaluated at the particle position $\mathbf{r}_p(t)$; its appearance in Eqn. (1315) indicates that the convention in STAR-CCM+ is that $\mathbf{v}_p(t)$ is the absolute velocity of the particle, whereas $\mathbf{r}_p(t)$ is the position of the particle with respect to the frame of reference.

Individual particles are not tracked, instead a single parcel represents a set of identical particles, at some mean centroid $\mathbf{r}_\pi(t)$. The velocity of the parcel is assumed to be the same as its constituent particles, hence its equation of motion is

$$\frac{\mathrm{d}\mathbf{r}_\pi}{\mathrm{d}t} = \mathbf{v}_p - \mathbf{v}_g \tag{1316}$$

where the grid velocity is now evaluated at $\mathbf{r}_\pi(t)$.

In subsequent sections, a quantity which is identical for a parcel and its particles, such as velocity in the preceding paragraph, are described as a "particle" quantity and given the subscript p. The term "parcel" and the subscript π are reserved for quantities which are exclusively defined on parcels.

Massless Particle Equation of Motion

The continuous phase transports the massless particles, meaning that their velocity is always identical to that of the surrounding fluid

$$\mathbf{v}_p = \mathbf{v}. \tag{1317}$$

For convenience in later equations, the particle slip velocity is defined as the velocity deficit of the particle relative to the continuous phase

$$\mathbf{v}_s \equiv \mathbf{v} - \mathbf{v}_p. \tag{1318}$$

A massless particle therefore has zero slip velocity.

Mass Balance for a Material Particle

The equation of conservation of mass of a material particle is simply

$$\frac{\mathrm{d}m_p}{\mathrm{d}t} = \dot{m}_p \tag{1319}$$

where m_p is the mass of the particle and $\dot{m}_p$ the rate of mass transfer to the particle. The latter is zero unless a mass transfer model is active, such as the Quasi-Steady Single-Component Droplet Evaporation model.

If the Two-Way Coupling model is active, $\dot{m}_p$ is accumulated over all the parcels and applied in the continuous phase continuity equation.

Mass Transfer

The rate of mass transfer to a single particle from the continuous phase is $\dot{m}_p$.

In unsteady simulations, the rate of mass transfer from all particles in a cell c to the continuous phase is

$$S_m = -\frac{1}{\Delta t}\sum_{\pi}\left(\int_t^{t+\Delta t}\int_{V_c}\delta(\mathbf{r}-\mathbf{r}_\pi)n_\pi\dot{m}_p\mathrm{d}V\mathrm{d}t\right) \tag{1320}$$

where the volume integral is over the cell; the Dirac delta function filters out parcels which are not in the cell. The summation is over all parcels for which the Two-Way Coupling model is active. The discrete form of Eqn. (1320) is

$$S_m = -\frac{1}{\Delta t}\sum_{\pi}\sum_{\delta t_p} n_\pi \dot{m}_p \delta t_p \tag{1321}$$

where the second summation is over all δt_p for which parcel π is in cell c.

In steady simulations, the rate of mass transfer from all particles in a cell c to the continuous phase is

$$S_m = -\sum_{\pi}\left(\int_0^{\infty}\int_{V_c}\delta(\mathbf{r}-\mathbf{r}_\pi)\dot{n}_\pi\dot{m}_p \mathrm{d}V\mathrm{d}t\right) \tag{1322}$$

or, in discrete form,

$$S_m = -\sum_{\pi}\sum_{\delta t_p}\dot{n}_\pi\dot{m}_p\delta t_p\,. \tag{1323}$$

The field S_m is applied in the continuous phase continuity equation.

Component Mass Balance for a Material Particle

For multi-component material particles, a mass conservation equation can be written for each component

$$\frac{\mathrm{d}m_{pi}}{\mathrm{d}t} = \dot{m}_{pi} \tag{1324}$$

where m_{pi} is the mass of component i in the particle and $\dot{m}_{pi}$ the rate of mass transfer to this component. The latter is zero unless a mass transfer model is active, such as the Multi-Component Droplet Evaporation model. The overall mass balance, Eqn. (1319), is satisfied when Eqn. (1324) is summed over all the components.

If the Two-Way Coupling model is active, $\dot{m}_{pi}$ is accumulated over all the parcels and applied in the continuous phase species conservation equation. As with single-component material particles, $\dot{m}_p$ is accumulated over all the parcels and applied in the continuous phase continuity equation.

Component Mass Transfer

The rate of mass transfer to a single particle component from the continuous phase is $\dot{m}_{pi}$.

In unsteady simulations, the rate of mass transfer of component i from all particles in a cell c to the continuous phase is

$$S_{mi} = -\frac{1}{\Delta t}\sum_{\pi}\left(\int_t^{t+\Delta t}\int_{V_c}\delta(\mathbf{r}-\mathbf{r}_\pi)n_\pi\dot{m}_{pi}\mathrm{d}V\mathrm{d}t\right) \tag{1325}$$

where the volume integral is over the cell; the Dirac delta function filters out parcels which are not in the cell. The summation is over all parcels for which the Two-Way Coupling model is active. The discrete form of Eqn. (1320) is

$$S_{mi} = -\frac{1}{\Delta t}\sum_{\pi}\sum_{\delta t_p} n_\pi \dot{m}_{pi}\delta t_p \tag{1326}$$

where the second summation is over all δt_p for which parcel π is in cell c.

In steady simulations, the rate of mass transfer of component i from all particles in a cell c to the continuous phase is

$$S_{mi} = -\sum_{\pi}\left(\int_0^\infty \int_{V_c} \delta(\mathbf{r}-\mathbf{r}_\pi)\dot{n}_\pi \dot{m}_{pi}\mathrm{d}V\mathrm{d}t\right) \tag{1327}$$

or, in discrete form,

$$S_{mi} = -\sum_{\pi}\sum_{\delta t_p} \dot{n}_\pi \dot{m}_{pi}\delta t_p\,. \tag{1328}$$

The field S_{mi} is applied in the continuous phase conservation equation for component i.

Momentum Balance for a Material Particle

The generic form of the equation of conservation of momentum for a material particle is

$$m_p \frac{\mathrm{d}\mathbf{v}_p}{\mathrm{d}t} = \mathbf{F}_s + \mathbf{F}_b\,. \tag{1329}$$

Here $\mathbf{F}_s$ represents the forces acting on the surface of the particle, and $\mathbf{F}_b$ the body forces. These forces in turn are decomposed into

$$\mathbf{F}_s = \mathbf{F}_d + \mathbf{F}_p + \mathbf{F}_{vm} \tag{1330}$$

$$\mathbf{F}_b = \mathbf{F}_g + \mathbf{F}_u \tag{1331}$$

where:

- $\mathbf{F}_d$ is the drag force
- $\mathbf{F}_p$ is the pressure gradient force
- $\mathbf{F}_{vm}$ is the virtual mass force
- $\mathbf{F}_g$ is the gravity force

- $\mathbf{F}_u$ is the user-defined body force

The momentum transfer to the particle from the continuous phase is simply $\mathbf{F}_s$. If the Two-Way Coupling model is active, $\mathbf{F}_s$ is accumulated over all the parcels and applied in the continuous phase momentum equation.

Drag Force

The equation for the drag force is

$$\mathbf{F}_d = \frac{1}{2} C_d \rho A_p |\mathbf{v}_s| \mathbf{v}_s \tag{1332}$$

where C_d is the drag coefficient of the particle, ρ the density of the continuous phase, $\mathbf{v}_s$ the particle slip velocity and A_p the projected area of the particle. This force appears in the Drag Force model.

Eqn. (1332) can also be written

$$\mathbf{F}_d = \frac{m_p \mathbf{v}_s}{\tau_v} \tag{1333}$$

in which τ_v is the *momentum relaxation time-scale*

$$\tau_v = \frac{2m_p}{C_d \rho A_p |\mathbf{v}_s|}. \tag{1334}$$

Pressure Gradient Force

The equation for the pressure gradient force is

$$\mathbf{F}_p = -V_p \nabla p_{static} \tag{1335}$$

where V_p is the volume of the particle and ∇p_{static} is the gradient of the static pressure in the continuous phase. The Pressure Gradient Model activates the pressure gradient force.

Virtual Mass Force

The equation for the virtual mass force [230] is

$$\mathbf{F}_{vm} = C_{vm} \rho V_p \left(\frac{\mathrm{D}\mathbf{v}}{\mathrm{D}t} - \frac{\mathrm{d}\mathbf{v}_p}{\mathrm{d}t} \right) \tag{1336}$$

where C_{vm} is the virtual mass coefficient. The default value of 0.5 for this coefficient is for a sphere in a uniform, inviscid, incompressible flow [248]. The operator $\mathrm{D}/\mathrm{D}t$ is the material (substantive) derivative. This force appears in the Virtual Mass model.

Gravity Force

The equation for the gravity force is

$$\mathbf{F}_g = m_p \mathbf{g} \tag{1337}$$

where $\mathbf{g}$ is the gravitational acceleration vector. This force is activated when the Gravity model is active in the continuous phase.

User-Defined Body Force

The equation for the user-defined body force is

$$\mathbf{F}_u = V_p \mathbf{f}_u \tag{1338}$$

where $\mathbf{f}_u$ is the user body force (per unit volume). This force appears in the User-Defined Body Force model.

Drag Coefficient

The drag coefficient C_d is a function of the small-scale flow features around the individual particles. These features are impractical to resolve spatially, and so the usual practice is to obtain the drag coefficient from correlations, typically derived from experiment or theoretical studies. These correlations differ according to the nature of the dispersed phase. Droplets, bubbles, and solid particles have different correlations. The shape of the particle, the presence of inter-phase mass and energy transfer, and so on, all affect these correlations.

STAR-CCM+ provides two general methods for defining the drag coefficient:

- Schiller-Naumann correlation
- field function

A third method, that is based on the Liu dynamic drag coefficient, is available when the TAB Distortion model is active.

Schiller-Naumann Correlation

The Schiller-Naumann correlation [264] is suitable for spherical solid particles, liquid droplets, and small-diameter bubbles. It is formulated as

$$C_d = \begin{cases} \dfrac{24}{\mathrm{Re}_p}(1 + 0.15\mathrm{Re}_p^{0.687}) & , \quad \mathrm{Re}_p \le 10^3 \\ 0.44 & , \quad \mathrm{Re}_p > 10^3 \end{cases} \tag{1339}$$

where Re_p is the particle Reynolds number, that is defined as

$$\mathrm{Re}_p \equiv \frac{\rho \left|\mathbf{v}_s\right| D_p}{\mu} \tag{1340}$$

and D_p is the particle diameter. This correlation is available only when the continuous phase is viscous.

Field Function

The field function drag coefficient method allows C_d to be specified directly with a scalar field function, *f*:

$$C_d = f. \tag{1341}$$

There is no restriction governing the availability or applicability of this method.

Liu Dynamic Drag Coefficient

The Liu dynamic drag coefficient [246] is intended to account for the dependence of the drag of a liquid droplet on its distortion under the action of aerodynamic forces. As a basis, it uses the following expression for the drag coefficient of an undistorted sphere.

$$C_{d,\,sphere} = \begin{cases} \dfrac{24}{\mathrm{Re}_p}\left(1 + \dfrac{1}{6}\mathrm{Re}_p^{2/3}\right) & \mathrm{Re}_p \le 1000 \\ 0.424 & \mathrm{Re}_p > 1000 \end{cases} . \tag{1342}$$

As the distortion of the droplet increases, its shape is assumed to become a disk whose axis is aligned with the relative velocity. This shape increases the drag on the droplet. The Liu drag coefficient models this effect by noting that the high Reynolds number limit of the drag coefficient of a disk is 1.54. It then assumes that the disk drag is 1.54/0.424 higher than the sphere drag at all Reynolds numbers, and that the drag of intermediate shapes can be interpolated between those two extremes. So

$$C_d = C_{d,\,sphere}(1 + 2.632y). \tag{1343}$$

The interpolation factor y is 0 for a sphere and 1 for a disk. It is identified as the TAB distortion, Eqn. (1419), and hence the Liu drag coefficient requires the TAB distortion model to calculate this quantity.

Field Charging

Field charging occurs when unipolar ions, traveling along an electric field, bombard a particle and lose their charge to it, until the charge on the particle reaches a saturation level. The saturation charge of a spherical particle exposed to an electric field is:

$$Q_s = \frac{3\pi\varepsilon_c\varepsilon_r}{\varepsilon_r + 2}|E|d_p{}^2 \tag{1344}$$

where:

- ε_c is the permittivity of the medium.
- ε_r is the relative permittivity of the particle.
- E is the strength of the electric field, calculated from Maxwell's equations.
- d_p is the diameter of the particle.

The total surface charge on the particle Q_p evolves as:

$$\frac{dQ_p}{dt} = \frac{Q_s}{\tau_q}\left(1 - \frac{Q_p}{Q_s}\right)^2 \tag{1345}$$

where the time constant τ_q is:

$$\tau_q = \frac{4\varepsilon_c}{\rho K} \tag{1346}$$

- ρ is the charge density.
- K is the ion mobility.

Film Damping

Film Damping mode allows the motion of solid particles to be damped immediately before and after wall impingement as they pass through a fluid film. This mode is only available for solid particles on fluid film boundaries.

In this situation, the equation of conservation of momentum for a material particle is assumed to reduce to:

$$m_p \frac{\mathrm{d}\mathbf{v}_p}{\mathrm{d}t} = \mathbf{F}_d + \mathbf{F}_{vm} \tag{1347}$$

where:

- $\mathbf{F}_d$ is the drag force
- $\mathbf{F}_{vm}$ is the virtual mass force

Drag Force

The equation for drag force can be used in its standard form. As the particles are small, spherical, and solid the drag coefficient of the particle, C_d can be calculated using the Schiller-Naumann correlation.

Virtual Mass Force

The equation for virtual mass force is simplified using two assumptions:

- The virtual mass coefficient, C_{vm} is equal to 0.5
- The acceleration of the fluid film is negligible compared to the acceleration at the particle:

$$\left|\frac{\mathrm{D}\mathbf{v}}{\mathrm{D}t}\right| \ll \left|\frac{\mathrm{d}\mathbf{v}_p}{\mathrm{d}t}\right| \tag{1348}$$

Using these assumptions the virtual mass force becomes:

$$\mathbf{F}_{vm} = -0.5\rho V_p \frac{\mathrm{d}\mathbf{v}_p}{\mathrm{d}t} \tag{1349}$$

Momentum Transfer

The rate of momentum transfer to a single particle from the continuous phase is $\mathbf{F}_s + \dot{m}_p\mathbf{v}_p$, where $\mathbf{F}_s$ is the force acting on surface of the particle, defined in Eqn. (1330), and $\dot{m}_p$ the rate of mass transfer to the particle.

In unsteady simulations, the rate of momentum transfer from all particles in a cell c to the continuous phase is

$$\boldsymbol{S}_{\mathbf{v}} = -\frac{1}{\Delta t}\sum_{\pi}\left(\int_t^{t+\Delta t}\int_{V_c}\delta(\mathbf{r}-\mathbf{r}_\pi)n_\pi(\mathbf{F}_s + \dot{m}_p\mathbf{v}_p)\mathrm{d}V\mathrm{d}t\right) \tag{1350}$$

where the volume integral is over the cell; the Dirac delta function filters out parcels which are not in the cell. The summation is over all parcels for which the Two-Way Coupling model is active. The discrete form of Eqn. (1350) is

$$\boldsymbol{S}_{\mathbf{v}} = -\frac{1}{\Delta t}\sum_{\pi}\sum_{\delta t_p} n_\pi(\mathbf{F}_s + \dot{m}_p\mathbf{v}_p)\delta t_p \tag{1351}$$

where the second summation is over all δt_p for which parcel π is in cell c.

In steady simulations, the rate of momentum transfer from all particles in a cell c to the continuous phase is

$$\mathbf{S_v} = -\sum_{\pi}\left(\int_0^{\infty}\int_{V_c}\delta(\mathbf{r}-\mathbf{r}_\pi)\dot{n}_\pi(\mathbf{F}_s+\dot{m}_p\mathbf{v}_p)\mathrm{d}V\mathrm{d}t\right) \tag{1352}$$

or, in discrete form,

$$\mathbf{S_v} = -\sum_{\pi}\sum_{\delta t_p}\dot{n}_\pi(\mathbf{F}_s+\dot{m}_p\mathbf{v}_p)\delta t_p\,. \tag{1353}$$

The field $\mathbf{S_v}$ is applied in the continuous phase momentum equation.

Energy Balance for a Material Particle

A material particle is assumed to be internally homogeneous which, from a thermal point of view, implies a low Biot number, for example, less than ~0.1. The generic form of the equation of conservation of energy consistent with this assumption is

$$m_p c_p \frac{\mathrm{d}T_p}{\mathrm{d}t} = Q_t + Q_{rad} + Q_s \tag{1354}$$

Here, Q_t represents the rate of convective heat transfer to the particle from the continuous phase, Q_{rad} represents the rate of radiative heat transfer, and Q_s other heat sources.

Convective Heat Transfer

The convective heat transfer term is formulated as

$$Q_t = fhA_s(T-T_p) \tag{1355}$$

where A_s is the particle surface area, and h is the heat transfer coefficient. The factor f is a mass transfer correction, for which the formulation of El Wakil and others [239] is used

$$f = \frac{z}{e^z-1} \tag{1356}$$

with

$$z = \frac{-\dot{m}_p c_p}{hA_s}\,. \tag{1357}$$

In the limit $\dot{m}_p \to 0$, $f \to 1$.

The heat transfer to the particle is always active in the energy equation. It can also be written

$$Q_t = \frac{f m_p c_p (T - T_p)}{\tau_T} \tag{1358}$$

in which τ_T is the *thermal relaxation time-scale*

$$\tau_T = \frac{m_p c_p}{h A_s}. \tag{1359}$$

If the Two-Way Coupling model is active, Q_t is accumulated over all the parcels and applied in the continuous phase energy equation.

Radiative Heat Transfer

If the Particle Radiation model is active, heat transfer due to radiation is included in the energy equation. The radiative heat transfer, Q_{rad}, to the particle is defined as:

$$Q_{rad} = \frac{A_s}{4} Q_{a,p} (G - 4\sigma T_p^4) \tag{1360}$$

where A_s is the particle surface area, $Q_{a,p}$ is the absorption efficiency of the particle, G is the incident radiative heat flux, and σ is the Stefan-Boltzmann constant.

Other Sources of Heat Transfer

The other terms in the energy equation are the user-defined energy source and the latent heat transfer:

$$Q_s = V_p q_u + \dot{m}_p L_{eff} \tag{1361}$$

where q_u is the user energy source (per unit volume) and L_{eff} is the effective latent heat of transferred material, which depends on the mass transfer model and the composition of the transferred material. The user energy source is activated in the energy equation by the User-Defined Energy Source model.

Heat Transfer Coefficient

As with the drag coefficient, the heat transfer coefficient h must be defined using a correlation, generally but not necessarily empirical. It is often given in terms of the particle Nusselt number

$$\mathrm{Nu}_p \equiv \frac{h D_p}{k} \tag{1362}$$

where k is the thermal conductivity of the continuous phase.

STAR-CCM+ provides two methods for defining the heat transfer coefficient:

- Ranz-Marshall correlation
- field function

Ranz-Marshall correlation

The Ranz-Marshall correlation [257] is suitable for spherical particles up to Re_p ~5000. It is formulated as

$$\mathrm{Nu}_p = 2(1 + 0.3\mathrm{Re}_p^{1/2}\mathrm{Pr}^{1/3}) \tag{1363}$$

in which Pr is the Prandtl number of the continuous phase. This correlation is available only when the continuous phase is viscous.

Field Function

The field function heat transfer coefficient method allows h to be specified directly with a scalar field function, f:

$$h = f. \tag{1364}$$

There is no restriction governing the availability or applicability of this method.

Energy Transfer

The rate of total energy transfer to a single particle from the continuous phase is

$$Q_t + \mathbf{F}_s \cdot \mathbf{v}_p + \frac{1}{2}\dot{m}_p\mathbf{v}_p^2 + \dot{m}_p h,$$

where Q_t is the surface heat transfer defined in Eqn. (1355), $\mathbf{F}_s$ is the surface force that is defined in Eqn. (1330), $\dot{m}_p$ the rate of mass transfer to the particle and h the enthalpy of transferred material.

In unsteady simulations, the rate of energy transfer from all particles in a cell c to the continuous phase is

$$S_E = -\frac{1}{\Delta t}\sum_{\pi}\left(\int_t^{t+\Delta t}\int_{V_c}\delta(\mathbf{r}-\mathbf{r}_\pi)n_\pi\left(Q_t + \mathbf{F}_s \cdot \mathbf{v}_p + \frac{1}{2}\dot{m}_p\mathbf{v}_p^2 + \dot{m}_p h\right)\mathrm{d}V\mathrm{d}t\right) \tag{1365}$$

where the volume integral is over the cell; the Dirac delta function filters out parcels which are not in the cell. The summation is over all parcels for which the Two-Way Coupling model is active. The discrete form of Eqn. (1365) is

$$S_E = -\frac{1}{\Delta t}\sum_{\pi}\sum_{\delta t_p} n_\pi\left(Q_t + \mathbf{F}_s \cdot \mathbf{v}_p + \frac{1}{2}\dot{m}_p \mathbf{v}_p^2 + \dot{m}_p h\right)\delta t_p \tag{1366}$$

where the second summation is over all δt_p for which parcel π is in cell c.

In steady simulations, the rate of energy transfer from all particles in a cell c to the continuous phase is

$$S_E = -\sum_{\pi}\left(\int_0^\infty \int_{V_c} \delta(\mathbf{r} - \mathbf{r}_\pi) n_\pi\left(Q_t + \mathbf{F}_s \cdot \mathbf{v}_p + \frac{1}{2}\dot{m}_p \mathbf{v}_p^2 + \dot{m}_p h\right) \mathrm{d}V\mathrm{d}t\right) \tag{1367}$$

or, in discrete form,

$$S_E = -\sum_{\pi}\sum_{\delta t_p} n_\pi\left(Q_t + \mathbf{F}_s \cdot \mathbf{v}_p + \frac{1}{2}\dot{m}_p \mathbf{v}_p^2 + \dot{m}_p h\right)\delta t_p\,. \tag{1368}$$

The field S_E is applied in the continuous phase energy equation.

Quasi-Steady Single-Component Droplet Evaporation

The quasi-steady single-component droplet evaporation model assumes droplets to be internally homogeneous, consisting of a single liquid component. The rate of change of droplet mass due to quasi-steady evaporation $\dot{m}_p$ can be written [270]

$$\dot{m}_p = -g^* A_s \ln(1 + B) \tag{1369}$$

in which B is the Spalding transfer number and g^* is the mass transfer conductance (to be precise, in the limit $B \to 0$). Effectively, the transfer number represents the driving force for evaporation, which is a function of thermodynamic conditions of the liquid and vapor. Conductance, on the other hand, represents geometrical and mechanical effects, such as the size and velocity of the droplet. Expressions for these quantities depend on the mode of evaporation, of which there are three:

- Super-Critical
- Heat Transfer Limited Evaporation
- Vapor Diffusion Limited Evaporation

The conductance is positive definite by definition. The transfer number can be positive or negative, with the latter implying condensation. That is, condensation is treated as “reverse evaporation”.

Super-Critical

This mode is active when the droplet temperature exceeds the critical temperature of the material, $T_p \geq T_c$. The droplet is forced to vaporize instantaneously.

Heat Transfer Limited Evaporation

This mode is active when the vapor at the droplet surface is saturated and subcritical. The evaporation rate results from the balance between heat transfer from the vapor and latent heat transfer due to evaporation. The condition for a saturated vapor is that the surface equilibrium mole fraction of the vapor exceeds unity:

$$X_{v,s} = \frac{p_{sat}(T_p)}{p} \geq 1$$

in which $p_{sat}(T_p)$ is the vapor saturation pressure at the droplet surface temperature. Under this condition, the transfer number is

$$B = \frac{c_p(T - T_p)}{L}$$

where L is the latent heat of vaporization. The conductance is

$$g^* = \frac{k\mathrm{Nu}_p}{c_p D_p} \tag{1370}$$

where Nu_p is the Nusselt number.

Vapor Diffusion Limited Evaporation

This mode is active when the vapor at the droplet surface is subcritical and not saturated, that is, when neither of the other two modes is active. The evaporation rate then depends on the rate at which vapor can diffuse away from the droplet. Under this condition, the transfer number is [269]

$$B = \frac{Y_{v,s} - Y_v}{1 - Y_{v,s}}$$

where Y_v is the vapor mass fraction in the "free stream", that is, cell, and $Y_{v,s}$ is the surface equilibrium vapor mass fraction

$$Y_{v,s} = X_{v,s}\frac{W_v}{W_s}.$$

The molecular weights W_v and W_s are of the vapor and gas mixture at the droplet surface respectively. Finally, the conductance is

$$g^* = \frac{\rho D_v \mathrm{Sh}_p}{D_p} \tag{1371}$$

where D_v is the molecular diffusivity of the vapor and Sh_p the Sherwood number.

Multi-Component Droplet Evaporation

The multi-component droplet evaporation model assumes droplets to be internally homogeneous, consisting of an ideal mixture of liquid components, some of which are transferred. This means that they can vaporize. There can also be inert components in both the droplet and the gas. The rate of change of mass of each transferred component due to quasi-steady evaporation $\dot{m}_{pi}$ can be written [270]

$$\dot{m}_{pi} = -\varepsilon_i g^* A_s \ln(1 + B) \tag{1372}$$

in which B is the Spalding transfer number, g^* is the mass transfer conductance (to be precise, in the limit $B \to 0$) and ε_i the fractional mass transfer rate, which has the property

$$\sum_T \varepsilon_i = 1 \tag{1373}$$

with the summation over the T transferred components. Effectively, the transfer number represents the driving force for evaporation, which is a function of thermodynamic conditions of the liquid and vapor mixtures. On the other hand, the conductance represents geometrical and mechanical effects, such as the size and velocity of the droplet. Expressions for these quantities depend on the mode of evaporation, of which there are three:

- Super-Critical
- Heat Transfer Limited Evaporation
- Vapor Diffusion Limited Evaporation

The conductance is positive definite by definition. The transfer number can be positive or negative, with the latter implying condensation. That is, condensation is treated as “reverse evaporation”.

Super-Critical

This mode is active when the droplet temperature exceeds the highest critical temperature of the transferred components, $T_p \geq \max(T_{c,i})$. The transferred components are forced to vaporize instantaneously, leaving behind the non-volatile components, if any.

Heat Transfer Limited Evaporation

This mode is active when the vapor at the droplet surface is saturated and subcritical. The evaporation rate then depends on the balance between heat transfer from the vapor and latent heat transfer due to evaporation. The first condition for a saturated vapor is:

$$\sum_T X_{is} \geq 1 \tag{1374}$$

where X_{is} is the equilibrium mole fraction of transferred component *i* at the droplet surface. X_{is} is evaluated using Raoult's law:

$$X_{is} = \frac{p_{sat}^i(T_p)}{p} X_{ip}$$

in which $p_{sat}^i(T_p)$ is the component saturation pressure that is evaluated at the droplet surface temperature, and X_{ip} is the component mole fraction in the droplet.

When the condition in Eqn. (1374) is satisfied, the volatile components vaporize in proportion to their mass fractions, that is:

$$\varepsilon_i = \frac{Y_{ip}}{\sum_T Y_{ip}} \tag{1375}$$

The second condition for a saturated vapor is:

$$\sum_T Y_i = 1 \tag{1376}$$

where Y_i is the mass fraction of component *i* in the "free stream", that is the cell containing the droplet.

The transfer number is

$$B = \frac{c_p(T - T_p)}{\sum_T \varepsilon_i L_i} \tag{1377}$$

where L_i is the latent heat of vaporization.

Finally, the conductance is

$$g^* = \frac{k\mathrm{Nu}_p}{c_p D_p} \tag{1378}$$

where Nu_p is the Nusselt number.

Vapor Diffusion Limited Evaporation

This mode is active when the vapor at the droplet surface is subcritical and not saturated, that is, when neither of the other two modes is active. The evaporation rate then depends on the rate at which vapor components can diffuse away from the droplet. From [269], the fractional mass transfer rate is:

$$\varepsilon_i = \frac{Y_{is}(1+B) - Y_i}{B} \tag{1379}$$

The transfer number is:

$$B = \frac{\sum_T Y_{is} - \sum_T Y_i}{1 - \sum_T Y_{is}} \tag{1380}$$

where:

Y_i is the mass fraction of component i in the "free stream", that is the cell containing the droplet.

Y_{is} is its surface equilibrium mass fraction:

$$Y_{is} = X_{is}\frac{W_i}{W_s} \tag{1381}$$

The molecular weights W_i and W_s are of the component and gas mixture at the droplet surface respectively. Finally, the conductance is:

$$g^* = \frac{\rho D_v \mathrm{Sh}_p}{D_p} \tag{1382}$$

where D_v is the molecular diffusivity of the vapor (assumed equal for all components) and Sh_p the Sherwood number.

Sherwood Number

The Sherwood number is effectively a dimensionless mass transfer conductance, which is related to the actual mass transfer conductance through Eqn. (1371). As with the drag coefficient and heat transfer coefficient, the Sherwood number must be defined using a correlation, generally but not necessarily empirical.

STAR-CCM+ provides two methods for defining the Sherwood number:

- Ranz-Marshall correlation
- field function

Ranz-Marshall correlation

The Ranz-Marshall correlation [257] is suitable for spherical particles up to Re_p ~5000. It is formulated as

$$\mathrm{Sh}_p = 2(1 + 0.3\mathrm{Re}_p^{1/2}\mathrm{Sc}^{1/3}) \tag{1383}$$

in which Sc is the Schmidt number of the continuous phase. This correlation is available only when the continuous phase is viscous.

Field Function

The field function Sherwood number method allows Sh_p to be specified directly with a scalar field function, *f*:

$$\mathrm{Sh}_p = f. \tag{1384}$$

There is no restriction governing the availability or applicability of this method.

Turbulent Dispersion

Following Gosman and Ioannides [240], a particle is assumed to pass through a sequence of turbulent eddies as it traverses a turbulent flow field. Here, an eddy is a local disturbance to the Reynolds-Averaged velocity field. The particle remains in the eddy until either the eddy time-scale τ_e is exceeded, or the separation between the particle and the eddy exceeds the length scale of the eddy, l_e. In practice, an eddy transit time, τ_c serves in place of the latter.

Velocity Realizations

The particle experiences an instantaneous fluid velocity $\mathbf{v}$ in each eddy, which is

$$\mathbf{v} = \bar{\mathbf{v}} + \mathbf{v}' \tag{1385}$$

where $\bar{\mathbf{v}}$ is the local Reynolds-Averaged velocity and $\mathbf{v}'$ is the eddy velocity fluctuation, unique to each particle. The latter is a normal (Gaussian) deviate with zero mean value and a standard deviation that comes from the eddy velocity scale

$$u_e = \frac{l_t}{\tau_t}\sqrt{\frac{2}{3}}. \tag{1386}$$

The turbulence model provides the length and time-scales of the turbulence, l_t and τ_t. For example, their ratio is $\sqrt{k}$ with a K-Epsilon or K-Omega turbulence model.

Once generated, a single realization of $\mathbf{v}'$ continues to apply to a particle until its eddy interaction time is exceeded:

$$\tau_I = \min(\tau_e, \tau_c). \tag{1387}$$

Thus, $\mathbf{v}$ as given by Eqn. (1385) is used to evaluate the particle slip velocity, Eqn. (1318) and feeds through to other models, for example, using the particle Reynolds number, Eqn. (1340).

Eddy Time-Scale

The eddy time-scale measures the lifetime of the eddy: the maximum interval over which a single realization of $\mathbf{v}'$ remains valid. It can be related to the diffusion of a passive scalar according to the underlying turbulence model, giving

$$\tau_e = \frac{2\mu_t}{\rho u_e^2}. \tag{1388}$$

Eddy Transit Time

Particles can interact with the eddy for less than the eddy time-scale when they have non-zero slip velocity, which precludes massless particles by definition. With non-zero slip velocity, it is possible for the particle to "cross" the eddy, and hence escape it. The eddy transit time is defined only for material particles when the Drag Force model is activated, in which case it is estimated from

$$\tau_c = \begin{cases} \infty & \tau_v \leq \dfrac{l_e}{|\mathbf{v}_s|} \\ -\tau_v \ln\left(1 - \dfrac{l_e}{\tau_v |\mathbf{v}_s|}\right) & \tau_v > \dfrac{l_e}{|\mathbf{v}_s|} \end{cases} \tag{1389}$$

where τ_v is the momentum relaxation time-scale, Eqn. (1334), and l_e is the eddy length scale $u_e \tau_e$.

Volume Fraction

The volume fraction of a Lagrangian phase is the fraction of the local cell volume which that phase occupies. It is calculated for each Lagrangian phase in which either the Material Particles or DEM Particles model is active, and is made available as the "Volume Fraction of *Phase*" field function.

In addition, the volume fraction ϕ_c of all Lagrangian phases in which the Two-Way Coupling model is active is accumulated. This volume is seen as a void by the continuous phase: the volume available for the continuous phase decreases by this fraction. In practice, ϕ_c is under-relaxed to promote stability:

$$\phi_c = \alpha\left(\sum \phi\right) + (1 - \alpha)\phi_{c,\,old} \tag{1390}$$

In this expression, α is the Under-Relaxation Factor property of the Two-Way Coupling solver and $\sum \phi$ is the sum of the volume fractions of the relevant Lagrangian phases. The value of ϕ_c which is contributed to the continuous phase void fraction is limited to $\min(\phi_c, \phi_{c,\,max})$, again to promote stability. The value $\phi_{c,\,max}$ is the Maximum Volume Fraction property of the Two-Way Coupling solver.

Erosion

This section describes the formulation that the Erosion model uses to calculate the erosion rate, which is presented as a field function.

Erosion Rate

The erosion rate is defined as the mass of wall material eroded per unit area per unit time. It is calculated on wall, baffle, and contact boundary faces by accumulating the damage that each particle impact does on the face:

$$E_f = \frac{1}{A_f}\sum_{\pi(f)} \dot{m}_\pi e_r \tag{1391}$$

in which A_f is the area of the face, $\dot{m}_\pi$ is the mass flow rate of particles in parcel π impacting on the face and e_r is the erosion ratio. The summation is over all parcels which strike the face in this iteration or time-step. The erosion rate therefore depends on the flow (whether and how particles impact on the wall) and the chosen method for the erosion ratio. There are four built-in methods for the erosion ratio:

- Ahlert Correlation
- DNV Correlation
- Neilson-Gilchrist Correlation
- Oka Correlation

together with the possibility of specifying it directly with a field function.

Ahlert Correlation

The Ahlert correlation [228] for the erosion ratio is

$$e_r = KF_s f(\alpha)\left(\frac{u_{\text{rel}}}{u_{\text{ref}}}\right)^n \tag{1392}$$

in which:

- K is a material-dependent constant.
- F_s is a factor to account for the shape of the particles.
- $f(\alpha)$ is a function expressing the dependency on the particle incidence angle.
- u_{rel} is the magnitude of the relative velocity of the particle with respect to the wall, $|\mathbf{u}_{\text{rel}}|$, $(\mathbf{u}_{\text{rel}} = \mathbf{u}_{\text{particle}} - \mathbf{u}_{\text{wall}})$.
- u_{ref} is a constant reference velocity.
- n a constant exponent.

The shape coefficient F_s is reported to take the value 1 for angular particles, 0.53 for semi-rounded particles and 0.2 for fully rounded particles.

The angle function $f(\alpha)$ is split into two ranges. Below the user-specified transition angle α_0 it is a polynomial in α, the incidence angle in radians. Above the transition angle, $f(\alpha)$ follows a trigonometric relationship

$$f(\alpha) = x\cos^2\alpha\sin w\alpha + y\sin^2\alpha + z. \tag{1393}$$

The constants w, x and y are user-defined, whereas z is calculated internally by requiring that $f(\alpha)$ is continuous at α_0.

The default coefficients for the Ahlert correlation are for liquid-borne semi-rounded sand particles eroding aluminum, as given by McLaury et al [247].

DNV Correlation

The DNV correlation [242] for the erosion ratio is

$$e_r = Kf(\theta)\left(\frac{u_{\text{rel}}}{u_{\text{ref}}}\right)^n \tag{1394}$$

in which:

- K is a material-dependent constant, with a default value of 2.0E-9, for steel.
- $f(\theta)$ is a function expressing the dependency on the particle incidence angle. The default form of the function is

  $$f(\theta) = A\theta + B\theta^2 + C\theta^3 + D\theta^4 + E\theta^5 + F\theta^6 + G\theta^7 + H\theta^8 \tag{1395}$$

  where the values of the coefficients are:

  $$\begin{aligned} A &= 9.37 & E &= 170.137 \\ B &= -42.295 & F &= -98.398 \\ C &= 110.864 & G &= 31.211 \\ D &= -175.864 & H &= -4.170 \end{aligned}$$

  θ is radians and the range of $f(\theta)$ is 0 to $\frac{\pi}{2}$.
- u_{rel} is the magnitude of the relative velocity of the particle with respect to the wall, $|\mathbf{u}_{\text{rel}}|$, ($\mathbf{u}_{\text{rel}} = \mathbf{u}_{\text{particle}} - \mathbf{u}_{\text{wall}}$).
- u_{ref} is a constant reference velocity. The default value is 1 m/s.
- n a constant exponent. The default value is 2.6.

The DNV correlation is superficially similar to the Ahlert correlation but differs in that $f(\theta)$ is a polynomial in terms of the particle incidence angle θ in degrees. Coefficients are not interchangeable between the Ahlert and DNV correlations.

The default coefficients for the DNV correlation are for air-borne sand eroding carbon steel, and are taken from [242].

Neilson-Gilchrist Correlation

The Neilson-Gilchrist [249] correlation for the erosion ratio is

$$e_r = e_{rC} + e_{rD} \tag{1396}$$

in which e_{rC} and e_{rD} represent contributions from cutting and deformation respectively. The cutting erosion is modeled as a function of the incidence angle α

$$e_{rC} = \begin{cases} \dfrac{u_{\mathrm{rel}}^2 \cos^2\alpha \sin\dfrac{\pi\alpha}{2\alpha_0}}{2\varepsilon_C} & \alpha < \alpha_0 \\ \dfrac{u_{\mathrm{rel}}^2 \cos^2\alpha}{2\varepsilon_C} & \alpha \geq \alpha_0 \end{cases} \tag{1397}$$

where:

- α_0 is the user-specified transition angle.
- ε_C is the user-specified cutting coefficient.
- u_{rel} is the magnitude of the relative velocity of the particle with respect to the wall, $|\mathbf{u}_{\mathrm{rel}}|$, $(\mathbf{u}_{\mathrm{rel}} = \mathbf{u}_{\mathrm{particle}} - \mathbf{u}_{\mathrm{wall}})$.

The deformation erosion is similarly

$$e_{rD} = \frac{\max(u_{\mathrm{rel}} \sin\alpha - K, 0)^2}{2\varepsilon_D} \tag{1398}$$

with ε_D the deformation coefficient and K the cut-off velocity, below which no deformation erosion occurs.

The default coefficients for the Neilson-Gilchrist correlation are for liquid-borne sand eroding AISI 4130 steel, and are taken from Wallace and others [273].

Oka Correlation

The Oka correlation [253] and [254] for the erosion ratio is

$$e_r = e_{90} g(\alpha) \left(\frac{u_{\mathrm{rel}}}{u_{\mathrm{ref}}}\right)^{k_2} \left(\frac{D_p}{D_{\mathrm{ref}}}\right)^{k_3}. \tag{1399}$$

where:

- The angle function $g(\alpha)$ is defined as

$$g(\alpha) = (\sin\alpha)^{n_1}(1 + H_v(1 - \sin\alpha))^{n_2} \tag{1400}$$

  with n_1, n_2 and H_v user-specified constants. Oka and others identify the value of H_v as the Vickers hardness of the eroded material in units of GPa.
- u_{rel} is the magnitude of the relative velocity of the particle with respect to the wall, $|\mathbf{u}_{\mathrm{rel}}|$, $(\mathbf{u}_{\mathrm{rel}} = \mathbf{u}_{\mathrm{particle}} - \mathbf{u}_{\mathrm{wall}})$.
- u_{ref} is the user-specified reference velocity.
- D_{ref} is the user-specified reference diameter.
- k_2 and k_3 are user-specified exponents.

By inspection of Eqn. (1399) and Eqn. (1400), e_{90} is revealed to be the reference erosion ratio at $u_{\mathrm{rel}} = u_{\mathrm{ref}}$, $D_p = D_{\mathrm{ref}}$ and $\alpha = 90°$.

The purported strength of the Oka model is that the coefficients for a particular combination of eroded material and eroding material can be derived from more fundamental coefficients. These coefficients are specific to either the eroded material or the eroding material. Hence, for example, the fundamental coefficients for sand can serve as a basis for both sand-steel erosion and sand-aluminum erosion. The fundamental coefficients for the eroding material, in turn, are shown to be derivable from measurable properties of the eroding material such as its Vickers hardness.

STAR-CCM+ bases the Oka correlation on the coefficients in Eqn. (1399) and Eqn. (1400); for more information on relating these coefficients to fundamental coefficients, see [253] and [254]. The default coefficients are for air-borne sand eroding 0.25% carbon steel, and are taken from [253] and [254], except for e_{90} which is derived from the DNV correlation.

LISA Model

The LISA model that is implemented in STAR-CCM+ is based primarily on the work of Senecal and others [268] and Schmidt and others [265]. The LISA model combines the sheet breakup and atomization stage into one, so that there are two stages modeled overall.

Assumptions

Assumptions that are made in the LISA model are:

- A hollow-cone spray forms.
- The model only calculates a mean droplet diameter, and generates droplets at the injector position by assuming a Rosin-Rammler size distribution. The droplets are distributed uniformly in the spray cone.

- The primary breakup model does not operate in the spray region. The breakup length L_b is assumed to be small, and the injection flowrate is steady. The atomized particles appear at the injection point immediately.
- Most numerical models assume that the slip velocity between the liquid sheet and the surrounding air is equal to the absolute velocity of the liquid near the injector ($V_s \equiv V_f$).

Film Formation

The centrifugal motion of the liquid within the injector creates a liquid film surrounding an air core. The thickness and velocity of the liquid film are calculated based on the input parameters for the pressure swirl injector. Initial sheet thickness within the injector, δ_0, is

$$\dot{m}_f = \pi \rho u_f \delta_0 (d_0 - \delta_o) \tag{1401}$$

where $\dot{m}_f$ is the mass flowrate through the injector (specified by you), ρ is the density of the continuous phase, and d_0 is the injector exit diameter.

The axial component u_f of the film velocity $\mathbf{v}_f$ at the injector exit is difficult to calculate from first principles. Instead, the approach from Han [241] is used. In this approach, the injector exit velocity profile is assumed to be uniform, and the absolute film velocity magnitude V_f is related to the injection pressure head Δp by:

$$V_f = |\mathbf{v}_f| = k_v \sqrt{\frac{2\Delta p}{\rho}} \tag{1402}$$

where $\mathbf{v}_f = u_f \mathbf{i} + v_f \mathbf{j}$ The value of the velocity coefficient k_v is obtained [268] as

$$k_v = max\left[0.7, \frac{4\dot{m}}{\pi d_0^2 \rho_f \cos\theta} \sqrt{\frac{\rho_f}{2\Delta p}}\right] \tag{1403}$$

The constant axial component is then

$$u_f = V_f \cos\theta \tag{1404}$$

where θ is the spray angle.

Sheet Breakup and Atomization

This section is only a brief summary of the breakup and atomization stage: details can be obtained from the paper by Senecal and others [268]. The model represents the film exiting the injector as a two-dimensional, viscous, incompressible liquid sheet of thickness $2h$ moving through a quiescent gas

medium. The model also includes the effects of surrounding gas, liquid viscosity, and surface tension on the breakup of the liquid sheet. This phenomenon is depicted in the figure below.



Fluctuating velocities and pressure for both the liquid and the surrounding gas are produced by imposing a spectrum of infinitesimal wave-type disturbances on the initially steady motion. These disturbances, which are defined within a coordinate system that moves with the sheet, have the form:

$$\eta = \eta_0 \exp(ikx + \omega t) \tag{1405}$$

In this equation, η_0 is the initial wave amplitude, $k = 2\pi/\lambda$ is the wavenumber, and $\omega = \omega_r + i\omega_i$ is the complex growth rate. The most unstable disturbance has the largest value of ω_r, denoted by Ω_s, and is assumed to be responsible for breakup. It is therefore desired to obtain a dispersion relation $\omega = \omega(k)$ from which the most unstable disturbance can be deduced. It has been shown that two solutions or modes exist that satisfy the conservation equations and their associated boundary conditions at the upper and lower interfaces for the disturbed liquid sheet. For the first solution, called the sinuous mode, the waves on the upper and lower boundaries are in phase. For the second solution, called the varicose mode, the waves are π radians out of phase. Numerous authors (for example, Squire [271], Rangel and Sirignano [256], and Senecal and others [268]) have shown that the sinuous mode is dominant for low velocities and low gas to liquid density ratios. In addition, the modes become indistinguishable for high velocity flows.

As derived in Senecal and others [268], the dispersion relation for the sinuous mode is

$$\begin{aligned}&\omega^2[\tanh(kh) + \Gamma] + \omega[4\nu_f k^2 \tanh((kh) + 2i\Gamma k V_s)] + \\ &4\nu_f^2 k^4 \tanh(kh) - 4\nu_f^2 k^3 \xi \tanh(\xi h) - \Gamma V_s^2 k^2 + \frac{\sigma k^3}{\rho_f} = 0\end{aligned} \tag{1406}$$

where V_s is the magnitude of the relative velocity between the liquid film and the gas, $\Gamma = \rho/\rho_f$, $\xi = k^2 + \omega/v_f$, and σ is the surface tension.

With some simplifications, the growth rate for the sinuous mode is

$$\omega_r = -\frac{2\nu_f k^2 \tanh(kh)}{\tanh(kh)+\Gamma} + \frac{\sqrt{4\nu_f^2 k^4 \tanh^2(kh) - \Gamma^2 V_s^2 k^2 - [\tanh(kh)+\Gamma](-\Gamma V_s^2 k^2 + \sigma k^3/\rho_f)}}{\tanh(kh)+\Gamma} \tag{1407}$$

This expression can be simplified further depending on whether the disturbance waves being formed are long or short. The distinction between short or long waves is made by calculating the Weber number for the gas phase, $We_g = \rho V_s^2 h/\sigma$ (based on the relative liquid velocity, sheet half-thickness, and gas density). *Short waves* are formed for $We_g > 27/16$, corresponding to high-speed flows with $\rho/\rho_f \ll 1$; otherwise *long waves* (corresponding to low-speed flows) are formed. Most modern injectors produce short waves.

Short Waves

For short waves the simplified expression for the growth rate is

$$\omega_r = -2\nu_f k^2 + \sqrt{4\nu_f^2 k^4 + \Gamma V_s^2 k^2 - \frac{\sigma k^3}{\rho_f}} \tag{1408}$$

Here, $k = K_s$ such that ω_r has its maximum value (denoted by Ω_s). For this condition, the sheet breaks up at a length and time that are given, respectively, by

$$L_b = V_f \tau_b \qquad \tau_b = \frac{\zeta_b}{\Omega_s} \tag{1409}$$

where $\zeta_b = \ln(\eta_b/\eta_0)$ is the sheet breakup empirical constant that you define (default value equal to 12). The ratio η_b/η_0 is the ratio of the surface disturbance at sheet breakup to the initial disturbance. The sheet half-thickness h_b [231] at the breakup is

$$h_b = \frac{h_0(d_0 - \delta_0)}{2L_b \sin\theta + d_0 - \delta_0} \tag{1410}$$

where

$$h_0 = 0.5\delta_0 \cos\theta \tag{1411}$$

The diameter of the ligament at sheet breakup is

$$d_L = \sqrt{\frac{16h_b}{K_s}} \tag{1412}$$

Long Waves

Viscosity has a minor effect on wave growth in the long wave regime. In this case K_s is directly obtained from an inviscid analysis as:

$$K_s = \frac{\rho_g V_s^2}{2\sigma} \tag{1413}$$

For the inviscid case:

$$\Omega_s = \sqrt{\frac{\Gamma V_s^2 K_s^2 - \frac{\sigma K_s^3}{\rho_f}}{K_s h_b}} \tag{1414}$$

$$L_b = V_s \tau_b = \frac{V_s \zeta_b}{\Omega_s} \tag{1415}$$

Eqn. (1410), Eqn. (1414) and Eqn. (1415), determine the sheet half-thickness h_b at the sheet breakup. The diameter of the ligament at sheet breakup is:

$$d_L = \sqrt{\frac{8h_b}{K_s}} \tag{1416}$$

Droplet Diameter Distribution

After calculating d_L, the mean diameter d_D for the Rosin-Rammler distribution (with a spread parameter of 3.5 and a specified dispersion angle) is

$$d_D = 1.88 d_L (1 + 3Oh)^{1/6} \tag{1417}$$

where Oh, the Ohnesorge number, is

$$Oh = \frac{\mu_f}{(\rho_f \sigma d_L)^{1/2}} \tag{1418}$$

TAB Distortion

The goal of the TAB distortion model is to calculate the instantaneous displacement x of the droplet equator from its equilibrium position, as illustrated in the following diagram (after Baumgarten [231]).



For simplicity of analysis, this displacement normalizes to the droplet diameter

$$y = \frac{2x}{C_b D_p} \tag{1419}$$

in which C_b is an empirical constant. The Taylor analogy between a distorting droplet and a damped spring-mass system [252] gives a second-order ordinary differential equation governing y

$$\frac{\rho_l D_p^2 \ddot{y}}{4} = \frac{C_k \sigma}{D_p}\left(\frac{\mathrm{We}}{\mathrm{We}_{crit}} - 2y\right) - C_d \mu_l \dot{y} \tag{1420}$$

where σ and μ_l are respectively the surface tension and viscosity of the droplet and We its Weber number, Eqn. (1311). C_d, C_k and We_{crit} are empirical constants with default values 5, 8 and 12 respectively. C_b does not appear in Eqn. (1420): it is not a necessary parameter to determine y. Eqn. (1420) is integrated for each droplet from injection, with initial conditions $y = \dot{y} = 0$. The field functions TAB Distortion and TAB Distortion Rate correspond to y and $\dot{y}$.

KHRT Droplet Breakup

Kelvin-Helmholtz instabilities are postulated as the cause of stripping breakup. The model schematics and key equations are illustrated below.

KH Model **RT Model**



Child drop diameter $= 2B_0\Lambda_{KH}$ Child drop diameter $= C_3\Lambda_{RT}$

Breakup time constant $= \dfrac{1.894B_1D_p}{\Lambda_{KH}\Omega_{KH}}$ Breakup time constant $= \dfrac{C_\tau}{\Omega_{RT}}$

Reitz [258] proposed expressions for the wavelength and growth-rate of these instabilities

$$\Lambda_{KH} = \frac{4.51D_p(1+0.45\mathrm{Oh}_r^{0.5})(1+0.4\mathrm{Ta}_r^{0.7})}{(1+0.865\mathrm{We}_r^{1.67})^{0.6}} \tag{1421}$$

$$\Omega_{KH} = \frac{0.34+0.385\mathrm{We}_r^{1.5}}{(1+\mathrm{Oh}_r)(1+1.4\mathrm{Ta}_r^{0.6})}\sqrt{\frac{4\pi\sigma}{3m_p}} \tag{1422}$$

in which $\mathrm{We}_r = 0.5\mathrm{We}$ is the Weber number, which is based on the droplet radius,

$$\mathrm{Oh}_r = \mu_l\sqrt{\frac{2}{\rho_l D_p\sigma}}$$

is the Ohnesorge number based on the droplet radius and $\mathrm{Ta}_r = \mathrm{Oh}_r\sqrt{\mathrm{We}_r}$ is the Taylor number that is based on the droplet radius. The wavelength and growth-rate in turn are used to formulate a characteristic time-scale and length-scale for breakup caused by these instabilities

$$\tau_{KH} = \frac{1.894 B_1 D_p}{\Lambda_{KH}\Omega_{KH}} \tag{1423}$$

$$D_{KH} = 2B_0\Lambda_{KH}. \tag{1424}$$

The coefficient B_0 takes the value 0.61. The coefficient B_1 can be tuned to account for uncertainties in, for example, the droplet initial conditions. A wide range of values have been reported in the literature, from 1.73 to 60 or higher.

If $D_p > D_{KH}$, breakup due to KH instabilities is possible, in which case the parent droplet diameter decreases according to

$$\frac{\mathrm{d}D_p}{\mathrm{d}t} = \frac{D_{KH} - D_p}{\tau_{KH}}. \tag{1425}$$

However, this is allowed to occur only if breakup due to RT instabilities (see below) does not occur.

The mass shed during KH breakup is accumulated [255] until it reaches a user-specified proportion of the parent parcel's original mass, default 3%. One or more child parcels are created, with child droplets having diameter D_{KH}. The diameter of the droplets in the parent parcel remains unchanged; the number of droplets is calculated to conserve mass. Child parcels are also given a velocity normal to the original velocity of the parent parcel:

$$v_\perp = A_1\Lambda_{KH}\Omega_{KH}. \tag{1426}$$

The normal velocity coefficient A_1 can be tuned to produce a known spray angle; alternatively the spray angle may be specified directly at a cone injector, in which case A_1 should be set to zero.

Rayleigh-Taylor instabilities are postulated as the cause of "catastrophic" breakup. Following Senecal et al [267], STAR-CCM+ incorporates the effect of viscosity on these instabilities. The relationship between their wavenumber and growth-rate as a function of the droplet acceleration a

$$\omega_{RT} = -k_{RT}^2\left(\frac{\mu_l + \mu_g}{\rho_l + \rho_g}\right) + \sqrt{k_{RT}\left(\frac{\rho_l - \rho_g}{\rho_l + \rho_g}\right)a - \frac{k_{RT}^3\sigma}{\rho_l + \rho_g} + k_{RT}^4\left(\frac{\mu_l + \mu_g}{\rho_l + \rho_g}\right)^2} \tag{1427}$$

is solved numerically to give the wavelength Λ_{RT} corresponding to the maximum growth-rate Ω_{RT}. This is assumed to be the critical mode. In the limit of negligible viscosity, this method reproduces the inviscid RT formulation that is used by [255], for example. The justification for including viscosity is that its effects can be significant; implementations which neglect it, typically have to introduce improvised modifications, such

as a user-defined breakup length [261].

The wavelength and growth-rate in turn are used to formulate a characteristic time-scale and length-scale for breakup that is caused by these instabilities

$$\tau_{RT} = \frac{C_\tau}{\Omega_{RT}} \tag{1428}$$

$$D_{RT} = C_3 \Lambda_{RT}. \tag{1429}$$

The coefficient C_τ takes the default value 1. The coefficient C_3 takes the default value 0.1, with values in the range 0.1-5.33 being reported in the literature.

If $D_p > D_{RT}$, a critical Rayleigh-Taylor instability is assumed to be growing on the droplet. If the instability persists for a duration τ_{RT}, a breakup event occurs. The parent droplets shatter completely, producing child droplets with new diameter D_{RT}.

Reitz-Diwakar Droplet Breakup

The Reitz-Diwakar breakup model, [259] and [260], assumes that breakup occurs in one of two possible regimes. In either case, the droplet diameter decreases according to

$$\frac{\mathrm{d}D_p}{\mathrm{d}t} = \frac{D_s - D_p}{\tau_b} \tag{1430}$$

in which D_s is the stable diameter and τ_b the breakup time-scale. Both D_s and τ_b depend on the active breakup regime.

The first breakup regime, bag breakup, is caused by the droplet expanding into the low-pressure wake region behind it, and eventually disintegrating when surface tension forces are overcome. It is assumed to be possible if

$$\mathrm{We} > \mathrm{We}_{crit} \tag{1431}$$

with $\mathrm{We}_{crit} = 12$ the default. The stable diameter is obtained by making Eqn. (1431) an equality. The characteristic time-scale for this breakup regime is

$$\tau_b = \frac{C_{b2} D_p}{4} \sqrt{\frac{\rho_l D_p}{\sigma}}. \tag{1432}$$

The second breakup regime, stripping breakup, is caused by liquid being sheared or stripped from the droplet surface. It is assumed to be possible if Eqn. (1431) is satisfied and additionally

$$\mathrm{We} > \max(2C_{s1}\mathrm{Re}^{0.5}, \mathrm{We}_{crit}) . \tag{1433}$$

The default value for C_{s1} is 0.5. The stable diameter is obtained by making Eqn. (1433) an equality. The characteristic time-scale for this breakup regime is

$$\tau_b = \frac{C_{s2}}{2}\sqrt{\frac{\rho_l}{\rho}}\frac{D_p}{|\mathbf{v}_s|} . \tag{1434}$$

The default value for C_{s2} is 20. When both Eqn. (1431) and Eqn. (1433) are satisfied, the regime with the smallest characteristic time-scale dominates.

TAB Breakup

The TAB breakup model relies on the TAB distortion model to calculate the normalized droplet distortion y, Eqn. (1419). When y reaches unity, the droplet is assumed to be critically distorted; a breakup event is then triggered. The Sauter mean diameter D_{32} of the post-breakup droplets is related to the diameter D_p of the parent droplet through an energy balance [252], giving

$$\frac{D_p}{D_{32}} = 1 + \frac{KC_b^2C_k}{5} + \frac{3C_b^2m_p}{4\sigma\pi}\left(\frac{K}{5} - \frac{C_v^2}{6}\right)\dot{y}^2 . \tag{1435}$$

Here, C_k and C_b are inherited from the TAB distortion model, with C_b taking the default value 0.5. K is the ratio of the total energy in distortion and oscillation to the energy in the fundamental mode, default 10/3. In STAR-CCM+, the diameter of child droplets is obtained by sampling a Rosin-Rammler distribution with the Sauter mean diameter given by Eqn. (1435) and a user-specified exponent (spread). This is done separately for the parent parcel and, optionally, each child parcel.

Post-breakup parcels are also given a velocity normal to the original velocity of the parent parcel. This is proportional to their rate of distortion at the instant of breakup $\dot{y}_b$

$$v_\perp = \frac{1}{2}C_vC_bD_p\dot{y}_b . \tag{1436}$$

In both Eqn. (1435) and Eqn. (1436), C_v is the normal velocity coefficient. This coefficient can be tuned to produce a known spray angle; alternatively the spray angle may be specified directly at a cone injector, in which case C_v should be set to zero.

SSD Breakup

In the Stochastic Secondary Droplet (SSD) breakup model, breakup is a random process unrelated to original droplet size. This results in a log normal distribution of droplet diameters in the long time limit.

The particles interact with the continuous fluid phase as they travel. Particles break up when they meet three criteria:

- The particle radius is larger than the critical radius, r_c.
- The Weber number is larger than a critical Weber number, We_{cr}.
- The particle's accumulated residence time is larger than the breakup time. τ_{bu} (Eqn. (1438)). The model tracks parcels as they cross the domain and compares the residence time with a locally evaluated breakup time. The residence time is only calculated on parcels with a radius larger than a critical radius and a Weber number larger than a critical Weber number.

The first two criteria are related according to:

$$r_c = \frac{\mathrm{We}_{cr}\sigma_l}{\rho_g v_s^2} \tag{1437}$$

where:

- r_c is the critical radius
- We_{cr} is the critical Weber number specified in the interface. The default value is 6.
- σ_l is the surface tension of the liquid.
- ρ_g is the density of the continuous phase.
- v_s, slip velocity, is the relative velocity of the particle through the continuous phase.

Particles cease to exist when they break up into child particles. In turn, these child particles can also break up, if they meet the criteria above. The higher the critical Weber number, the larger the particles when breakup stops.

Particle break up only if they have a residence time greater than the breakup time. The breakup time scale is defined as:

$$\tau_{\mathrm{bu}} = B_1\sqrt{\frac{\rho_l}{\rho_g}}\frac{r_p}{v_s} \tag{1438}$$

where:

- τ_{bu} is the breakup time.

- B_1 is a parameter regulating speed of breakup, set in the interface ; larger values mean breakup occurs later. The default value is $\sqrt{3}$.
- ρ_l is the density of the liquid.
- ρ_g is the density of the continuous phase.
- r_p is the radius of the particle.
- v_s, slip velocity, is the relative velocity of the droplet through the continuous phase.

When breakup happens, the parent parcel is replaced by a number of child parcels, each containing a number of identical particles. The particle diameters in a given parcel follow a log-normal distribution:

$$\Phi(r) = \frac{1}{2}\left(1 + \mathrm{erf}\left[\frac{\log\left(\frac{r}{r_0}\right) - \langle\xi\rangle}{\sqrt{2}\langle\xi^2\rangle}\right]\right) \tag{1439}$$

where:

- $\Phi(r)$ is the breakup distribution function, the probability density as a function of r.
- r is the radius of the child droplet.
- r_0 is the radius of the parent droplet.
- $\langle\xi\rangle$ is the first moment of $\Phi(r)$, equal to a constant, K_0, set in the interface with default value -0.1. More negative values of K_0 give a smaller mean distribution of child diameters. Positive values of K_0 give larger mean child diameters.
- $\langle\xi^2\rangle$ is the second moment of $\Phi(r)$,

$$\langle\xi^2\rangle = K_1 \ln\left(\frac{\mathrm{We}}{\mathrm{We}_{\mathrm{cr}}}\right) \tag{1440}$$

  where:

  - K_1 is a constant set in the interface, default value 0.1. K_1 must be positive.
  - We is the Weber number of the particle.
  - $\mathrm{We}_{\mathrm{cr}}$ is the critical Weber number, a constant set in the interface, default value 6.

  $\langle\xi^2\rangle$ is the variance of the distribution, and K_1 is directly proportional to it. The larger they are, the broader the distribution of diameters.

Child parcels receive the velocity of the parent parcel multiplied by a factor, $|\boldsymbol{W}_{\mathrm{bu}}|$:

$$\left|W_{\mathrm{bu}}\right| = K_3 r_{\mathrm{p,c}} \nu = \frac{K_3 r_{\mathrm{p,c}}}{\tau_{\mathrm{bu}}} \tag{1441}$$

where:

- K_3 is a constant set in the interface, Normal Velocity Coefficient, default value 0.1, limited to values between 0 and 1. The larger K_3 is, the more rapid the radial diffusion of the spray. Use a value of 0 if the spray angle is known and fixed at the injector.
- $r_{\mathrm{p,c}}$ is the radius of the child particle.
- $\nu = \frac{1}{\tau_{\mathrm{bu}}}$ is the breakup frequency, the reciprocal of the breakup time.

Bai-Gosman Wall Impingement

The Bai-Gosman model [232][233][234] aims to predict the outcomes of liquid droplets impacting an impermeable boundary ("wall") or fluid film. It categorizes possible outcomes as being in one of six possible regimes, illustrated in the following diagram.



For fluid film boundaries, momentum and kinetic energy lost by the droplets are gained by the film.

The choice of regime for a given impingement event is made using four parameters:

- The incident Weber number

$$\mathrm{We}_I = \frac{\rho_p \mathbf{v}_{r,n}^2 D_p}{\sigma} \tag{1442}$$

  in which $\mathbf{v}_{r,n} = (\mathbf{v}_p - \mathbf{v}_w) \cdot \mathbf{n}_w$ is the normal component of the particle velocity relative to the wall. $\mathbf{n}_w$ is the unit vector normal to the boundary.

- The Laplace number

$$\mathrm{La} = \frac{\rho_p \sigma D_p}{\mu_p^2}. \tag{1443}$$

- The boundary temperature T_w (provided the Energy model is active in the physics continuum).
- The wall state; a boundary can be either wet or dry. For fluid film boundaries, this is determined by the presence or absence of a fluid film. For other boundaries, this is determined by the user-selected Wall State property.

  For a dry wall state:

  - At Weber numbers below We_c and wall temperatures below T_{12}, impinging droplets spread out on the wall.
  - At higher Weber numbers or wall temperatures, droplets can spread, break up and spread, break up and rebound, or rebound.
  - Above We_c, impinging droplets splash.

  For a wet wall state:

  - At Weber numbers below 2 and wall temperatures below T_{12}, impinging droplets adhere to the wall.
  - At higher Weber numbers or wall temperatures, droplets can spread, break up and spread, break up and rebound, or rebound.
  - Above We_c, impinging droplets splash.

Leidenfrost effect occurs when the temperature is high enough, at T_{12}, reducing friction and heat transfer. It peaks at T_{23}. See Using the Leidenfrost Temperature.

Regime Transition Criteria

The idealized map of impingement regimes is shown below, in which there are three temperature ranges, separated by two transition temperatures:

- T_{12}, separating range 1 from range 2, is expected to be approximately the boiling temperature of the droplet.
- T_{23}, separating range 2 from range 3, is expected to be approximately the Leidenfrost temperature of the droplet. See Using the Leidenfrost Temperature.

The transition criteria for wetted walls are given below for the three wall temperature ranges.



The Weber numbers 2 and 20, shown in the preceding diagram, can be adjusted in the expert properties of this model.

The transition criteria for dry walls are given below for the three wall temperature ranges.



Temperature Range 1

The first temperature range is expressed as:

$$T_w \leq T_{12}\,. \tag{1444}$$

It also applies when the Energy model is not active in the physics continuum. The model recognizes the following regimes in this temperature range:

- Adhere

  $$\mathrm{We}_I \leq 2$$

  in which the impinging droplet adheres to the boundary in nearly spherical form. If the boundary is a fluid film, the droplets transfer their mass, momentum, and kinetic energy to the film.

  In STAR-CCM+, it is possible to select the specific boundary interaction mode to control the outcome of a particle satisfying the adhere condition. The default in unsteady is Stick, while in steady the default is Escape.

- Rebound

  $$\mathrm{We}_I \leq 20$$

in which the impinging droplet bounces off the boundary after the impact. This is functionally identical to the Rebound boundary interaction mode with restitution coefficients

$$e_n = 0.993 - 1.76\theta + 1.56\theta^2 - 0.49\theta^3 \tag{1445}$$

$$e_t = \frac{5}{7} \tag{1446}$$

in which θ is the droplet incidence angle that is measured from the boundary.

- Spread

$$\mathrm{We}_I \leq \mathrm{We}_c \tag{1447}$$

in which the droplet merges with an existing film on a wetted boundary. If the boundary is a fluid film, the droplets transfer their mass, momentum, and kinetic energy to the film. The limiting Weber number We_c is given by

$$\mathrm{We}_c = A\ \mathrm{La}^{-0.18} \tag{1448}$$

For smooth walls or wet rough walls

$$A = A_w \tag{1449}$$

where A_w is the user-defined splash onset coefficient with a default value, 1320.

For dry, rough walls A is interpolated from the table below

Roughness μm	0	0.05	0.14	0.84	3.1	12.0	∞
A/A_w	3.9879	3.9879	3.4348	1.9955	1.5576	1.0015	1

In the current version of STAR-CCM+, the spread regime is implemented as rebound with restitution coefficients

$$e_n = 0 \tag{1450}$$

$$e_t = 1 \tag{1451}$$

- Splash

$$\mathrm{We}_I > \mathrm{We}_c \tag{1452}$$

in which the impinging droplet breaks up into smaller droplets, some of which are reflected from the boundary. Residue may remain stuck to the wall. Each impinging parcel, of mass m_π, produces a user-specified number of secondary parcels which are randomly reflected from the wall within an ejection cone. The total mass of the secondary parcels, m_s, is determined by the ratio

$$r_m = \frac{m_s}{m_\pi} = 0.2 + c_{ws} X_r \tag{1453}$$

where X_r is a random number that is distributed uniformly in the range $[0 - 1]$. c_{ws} is a constant, for wet walls

$$c_{ws} = 0.75 \tag{1454}$$

and for dry walls

$$c_{ws} = 0.6 \tag{1455}$$

The total number of secondary droplets that are created per splash is given by

$$N_s = a_0 \left(\frac{\mathrm{We}_I}{\mathrm{We}_m} - 1 \right) \tag{1456}$$

with a_0 a user-specified coefficient (default 5) and We_m a Weber number which, in this temperature range, is equal to We_c. The actual number of secondary droplets that are created by a splash event tends to N_s as the number of secondary parcels tends to infinity. This is ensured by randomly selecting the diameters and number of secondary droplets per parcel from a Rosin-Rammler distribution with user-specified exponent q and reference diameter

$$D_{ref} = \left(\frac{r_m}{N_s \Gamma(1 + 3/q)} \right)^{1/3} D_p \tag{1457}$$

in which $\Gamma(x)$ is the gamma function. The default exponent, 1, produces a Chi-squared distribution.

The velocity of the secondary droplets is calculated from two contributions; one due to the normal incident component $\mathbf{v}_n$ and another due to the tangential incident component $\mathbf{v}_t$. They are denoted $\mathbf{v}_n'$ and $\mathbf{v}_t'$ respectively.

The magnitude of $\mathbf{v}_n'$ for each secondary parcel is calculated from the overall conservation of energy as described in [234]. The direction of $\mathbf{v}_n'$ is chosen by randomly selecting an azimuthal angle in the range $[0°, 360°]$ and an ejection angle from the wall in a user-specified range, default $[5°, 50°]$.

The $\mathbf{v}_t$ component contributes directly to the secondary droplet velocity as

$$\mathbf{v}_t' = c_f \mathbf{v}_t \tag{1458}$$

where c_f is the wall friction coefficient which is estimated to be in the range 0.6 to 0.8; its default value is set to 0.7.

Finally, the parent parcel sticks to the wall. Its mass is reduced by a factor $(1 - r_m)$, while the number of droplets which it represents remains unchanged.

Temperature Range 2

The second temperature range is expressed as:

$$T_{12} < T_w \leq T_{23}. \tag{1459}$$

In this range, there is no contact or deposition due to an intervening vapor film. The regimes are determined by two characteristic Weber numbers We_{T1} and We_{T2} as follows:

- Rebound:

$$\mathrm{We}_I \leq \mathrm{We}_{T1} \tag{1460}$$

- Break-up and Rebound:

$$\mathrm{We}_{T1} < \mathrm{We}_I \leq \mathrm{We}_{T2} \tag{1461}$$

  in which the incident droplets shatter and rebound from the wall. This regime is implemented as a variant of the splash regime in the first temperature range except that no residue is left on the wall ($r_m = 1$) and $\mathrm{We}_m = \mathrm{We}_{T1}$.

- Break-up and Spread:

$$\mathrm{We}_{T2} < \mathrm{We}_I \leq \mathrm{We}_c \tag{1462}$$

  in which the incident droplets shatter and spread along the wall. This regime is implemented as a variant of the splash regime in the first temperature range except that no residue is left on the wall ($r_m = 1$), $\mathrm{We}_m = \mathrm{We}_{T2}$ and the ejection angle is zero.

- Splash:

$$\mathrm{We}_I > \mathrm{We}_c \tag{1463}$$

  which is identical to the splash regime in the first temperature range except that no residue is left on the wall ($r_m = 1$).

Following [232], default values of 30 and 80 for water are used for We_{T1} and We_{T2} respectively.

Temperature Range 3

The third temperature range is expressed as:

$$T_w > T_{23}. \tag{1464}$$

Again, wall contact is prevented by a vapor film and the droplets are elevated from and move tangentially to the surface. The regimes are as follows

- Spread:

$$\mathrm{We}_I \leq \mathrm{We}_{T1} \tag{1465}$$

- Break-up and Spread:

$$\mathrm{We}_{T1} < \mathrm{We}_I \leq \mathrm{We}_c \tag{1466}$$

  in which the incident droplets shatter and spread along the wall. This regime is implemented as a variant of the splash regime in the first temperature range except that no residue is left on the wall ($r_m = 1$), $\mathrm{We}_m = \mathrm{We}_{T1}$ and the ejection angle is zero.

- Splash:

$$\mathrm{We}_I > \mathrm{We}_c \tag{1467}$$

  which is identical to the splash regime in the first temperature range except that no residue is left on the wall ($r_m = 1$).

Satoh Wall Impingement Model

The Satoh model [263] aims to predict the outcomes of oil droplets impacting solid boundaries, particularly in oil-mist separators. Droplets impinging on boundaries are assumed to either rebound or spread depending on their diameter, D and velocity normal to the wall, Vp.

An intermediate region where a combination of rebounding and spreading occurs is defined between two lines, Initial Spread Line (ISL) and Full Spread Line (FSL). These lines both follow power law curves.

- Initial Spread Line

$$Vp_{ISL} = a_{ISL}\left(\frac{D}{D_{ref}}\right)^{n_{ISL}} \tag{1468}$$

- Full Spread Line

$$Vp_{FSL} = a_{FSL}\left(\frac{D}{D_{ref}}\right)^{n_{FSL}} \tag{1469}$$

where D_{ref} is the reference diameter, a is the pre-exponent multiplier and n is the exponent value.

Boundary Interaction Modes

This section describes the formulation of the Lagrangian phase boundary interaction modes. Currently only the rebound interaction mode is detailed.

Rebound

Rebound mode is available only at walls, baffles, and contact interfaces. Rebounding particles remain active in the simulation; the mode is distinguished by its treatment of the particle velocity. The rebound velocity relative to the wall velocity is determined by the impingement velocity and user-defined restitution coefficients:

$$(\mathbf{v}_p - \mathbf{v}_w)^R = e_t(\mathbf{v}_p - \mathbf{v}_w)_t^I - e_n(\mathbf{v}_p - \mathbf{v}_w)_n^I. \tag{1470}$$

The superscripts R and I denote rebound and impingement respectively; the subscripts n and t denote wall-normal and tangential respectively. Since the left-hand side of Eqn. (1470) can be split into orthogonal n and t components, it can be split into two equations:

$$(\mathbf{v}_p - \mathbf{v}_w)_t^R = e_t(\mathbf{v}_p - \mathbf{v}_w)_t^I \tag{1471}$$

$$(\mathbf{v}_p - \mathbf{v}_w)_n^R = -e_n(\mathbf{v}_p - \mathbf{v}_w)_n^I \tag{1472}$$

which serve to emphasize the definition of the restitution coefficients as the constants of proportionality between impingement and rebound velocities. Both coefficients may range from 0 to 1. The latter, "perfect" elastic rebound, is the default.

The tangential velocity of a wall boundary is zero unless a value is explicitly prescribed through a wall sliding option. In other words, it is non-zero only at no-slip walls.

Particle Size Distributions

A range of particle sizes generated by an injector may be represented statistically by a particle size distribution. In STAR-CCM+ a particle size distribution is quantified using a cumulative distribution function (CDF), $F(D)$. Here, "size" may mean diameter or mass, depending on the particle size specification, but we use the symbol D. The flow rate specification of the injector determines the precise definition of the cumulative distribution function:

- when the mass flow rate is specified, the cumulative distribution function $F_m(D)$ gives the fraction of the mass flow rate of the injector with a size smaller than D. Hence $F(D_M) = 0.5$ defines D_M as the **mass-median** size.

- when the volume flow rate is specified, the cumulative distribution function $F_V(D)$ gives the fraction of the volume flow rate of the injector with a size smaller than D. Hence $F(D_M) = 0.5$ defines D_M as the **volume-median** size.
- when the particle flow rate is specified, the cumulative distribution function $F_N(D)$ gives the fraction of the particle flow rate of the injector with a size smaller than D. Hence $F(D_M) = 0.5$ defines D_M as the **number-median** size.

This definition provides three essential properties of a cumulative distribution function:

$$\begin{aligned} \lim_{D \to 0} F(D) &= 0 \\ \lim_{D \to \infty} F(D) &= 1 \\ \frac{\mathrm{d}F(D)}{\mathrm{d}D} &\geq 0 \end{aligned} \tag{1473}$$

STAR-CCM+ provides three internal cumulative distribution functions.

- Log-Normal
- Normal
- Rosin-Rammler

In addition, the function may be tabulated as pairs of $(D, F(D))$ values. It is your responsibility to ensure that these values respect the properties in Eqn. (1473), with the additional constraint that CDF values must be strictly increasing—duplicate values are not permitted.

Given a cumulative distribution function, STAR-CCM+ generates particle sizes by dividing the distribution into ranges of equal mass, volume, or number (depending on how the flow rate is prescribed). The number of ranges is the number of parcel streams for the injector.

Log-Normal

The log-normal size distribution [245] is a normal (Gaussian) distribution which uses the logarithm of the particle size as the independent variable. Its cumulative distribution function is

$$F(D) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{\ln D - \ln \bar{D}}{\sigma\sqrt{2}}\right)\right] \tag{1474}$$

in which the parameters are the standard deviation σ and the mean size $\bar{D}$.

Normal

The normal size distribution is a normal (Gaussian) distribution which uses the particle size as the independent variable. Its cumulative distribution function is

$$F(D) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{D - \bar{D}}{\sigma\sqrt{2}}\right)\right] \tag{1475}$$

in which the parameters are the standard deviation σ and the mean size $\bar{D}$.

Rosin-Rammler

The Rosin-Rammler distribution [245] was developed to describe the volume distribution of particles as a function of their diameter, $F_V(D_p)$. In STAR-CCM+, it is extended to be a generic size distribution, with cumulative distribution function

$$F(D) = 1 - \exp\left[-\left(\frac{D}{D_{ref}}\right)^q\right] \tag{1476}$$

in which the parameters are the exponent q and the reference size D_{ref}. This form identifies the Rosin-Rammler distribution as a Weibull distribution.

As noted above, the Rosin-Rammler distribution may be a cumulative mass, volume, or number distribution, depending on the flow rate specification of the injector.

Impingement Heat Transfer

The Impingement Heat Transfer model is designed to predict the transfer of heat between a wall and a stream of droplets impinging on it. In a simulation with two-way coupling, the droplets can change the temperature of the wall. This temperature change can happen even in cases where the wall is above the Leidenfrost temperature. Heat can be transferred between droplets and wall through the gas phase, albeit more slowly than without the Leidenfrost effect. The droplets can then cool the wall below the Leidenfrost point, so that liquid films can form.

The model makes the following assumptions:

- The area through which energy flows from the drop to the wall face is no larger than the area of the wall face.
- For a wall boundary condition where the temperature is fixed (one-way coupling), the particle surface temperature over a dynamically calculated impingement area is fixed. The particle temperature is assumed to be at the droplet centroid a distance r from the wall.

- The distortion of the droplet is neglected and a diffuse spray is assumed.

At high mass loadings (seen with volume fractions larger than 0.3), these assumptions can break down, resulting in odd behavior in the gas phase. This behavior includes, but is not limited to, spurious temperature fluctuations, unphysical particle temperatures, and general difficulty converging.

The energy equation for the droplet or particle is:

$$mc_p \frac{dT_p}{dt} = \dot{Q}_c + \dot{Q}_r + \dot{Q}_l + \dot{Q}_s + \dot{Q}_{wp} \tag{1477}$$

where:

- $\dot{Q}_c$ is the rate of heat transfer from convection.
- $\dot{Q}_r$ is the rate of heat transfer from radiation.
- $\dot{Q}_l$ is the rate of heat transfer from mass transfer.
- $\dot{Q}_s$ is the rate of heat transfer from user-defined heat sources.
- $\dot{Q}_{wp}$ is the rate of heat transfer from wall conduction.

Heat that is transferred from the wall to the drop is (Wruck [274]):

$$\dot{Q}_{wp} = A_{cont} \frac{2\sqrt{t_{cont}}}{\sqrt{\pi}\Delta t_p} \frac{b_w b_p}{b_w + b_p} (T_w - T_p) \tag{1478}$$

where terms with a w or a p refer to the wall or particle, respectively, and:

- A_{cont} is the effective contact area:

  $$A_{cont} = \frac{\pi}{4} D_{eff}^{\ 2}$$

  where D_{eff} is an effective diameter, time-averaged over the droplet contact time (Akao *et al.* [229]):

  $$D_{eff} = D_0 0.61 We_p^{\ 0.38}$$

  where D_0 is the droplet diameter and We_p is the droplet Weber number.
- t_{cont} is the contact time between particle and wall:

$$t_{cont} = \begin{cases} \frac{\pi}{4}\sqrt{\frac{\rho_p D_0^{\ 3}}{\sigma_p}} & K < 40 \\ \sqrt{\frac{\pi}{2}}\left(\frac{\rho_p D_0^{\ 5}}{\sigma_p u_r^{\ 2}}\right)^{0.25} & K \geq 40 \end{cases}$$

where:

- ρ_p is the particle density.
- σ_p is the particle surface tension.
- u_r is the velocity of the particle relative to the wall.
- K is a dimensionless number proportional to the impact energy. Birkhold *et al.* [235] define this number as:

$$K = Ca^{5/4} La^{3/4}$$

where Ca is the Capillary number and La is the Laplace number.

- Δt_p is the particle time-step.
- b_w and b_p are the heat penetration coefficients as evaluated with the materials of the wall or the particle.

$$b_w = \sqrt{\kappa_w \rho_w (C_p)_w} \qquad b_p = \sqrt{\kappa_p \rho_p (C_p)_p}$$

where:

- κ_w and κ_p are the thermal conductivities of the wall and the particle.
- ρ_w and ρ_p are the densities of the wall and the particle.
- $(C_p)_w$ and $(C_p)_p$ are the specific heats of the wall and the particle.
- T_w and T_p are the temperatures of the wall and the particle.

Eqn. (1478) can be written as:

$$\dot{Q}_{wp} = A_p h_{wp} (T_w - T_p) \tag{1479}$$

where:

- A_p is the cross-sectional area of the particle.
- h_{wp} is:

$$h_{wp} = \frac{2A_{cont}\sqrt{t_{cont}}(b_w b_p)}{A_p \sqrt{\pi}(b_w + b_p)\Delta t_p}$$

Collision Detection Model

STAR-CCM+ uses a host cell approach to detect collisions. Two parcels can collide only if they are in the same cell. STAR-CCM+ uses the faster of the following collision algorithms to detect collisions in any particular cell:

- the NTC (No Time Counter) collision detection algorithm [266]
- the O'Rourke algorithm [251]

This section also describes collision outcomes.

NTC Algorithm

If a cell contains N droplets, the expected number of collisions in the cell over a time interval is given by summing the probability of all possible collisions

$$M_{\text{coll}} = \frac{1}{2}\sum_{i=1}^{N_p} q_i \sum_{j=1}^{N_p} q_j \frac{v_{i,j}\sigma_{i,j}\Delta t}{V} \tag{1480}$$

where:

- $v_{i,j}$ is the relative velocity between two colliding parcels
- $\sigma_{i,j}$ is the collision cross section of the two drops defined as

$$\sigma_{i,j} = \pi(r_i + r_j)^2 \tag{1481}$$

- Δt is the time-step size
- V is the cell volume
- N_p is the number of parcels in a cell
- q_i is the number of droplets in parcel i.

The factor of one-half is a result of symmetry. This summation can be modified by pulling a constant factor outside:

$$M_{\text{coll}} = \frac{(qv\sigma)_{\max}\Delta t}{2V}\sum_{i=1}^{N_p} q_i \sum_{i=1}^{N_p} \frac{q_j v_{i,j}\sigma_{i,j}}{(qv\sigma)_{\max}} \tag{1482}$$

The value of $(qv\sigma)_{\max}$ is used for scaling the selection probability of a collision. The chosen value is sufficiently large so that the following restriction holds

$$\frac{q_j v_{i,j}\sigma_{i,j}}{(qv\sigma)_{\max}} < 1 \tag{1483}$$

It is assumed that a representative subsample of parcels may be randomly selected from the set of parcels in the cell. This statistical approximation allows a constant multiplier to reduce the limits of summation as follows:

$$M_{\text{coll}} = \sum_{i=1}^{\sqrt{M_{\text{cand}}}} q_i \sum_{i=1}^{\sqrt{M_{\text{cand}}}} \frac{q_j v_{i,j} \sigma_{i,j}}{(q v \sigma)_{\max}} \tag{1484}$$

where

$$M_{\text{cand}} = \frac{N_p^2 (q v \sigma)_{\max} \Delta t}{2V} \tag{1485}$$

The overall cost is proportional to the product of the limits of the summation, namely M_{cand}. The value of M_{cand} is linearly proportional to N_p, because q goes as $1/N_p$. The double summation is evaluated using an acceptance-rejection scheme. The number of candidate pairs (M_{cand}) is selected randomly from the cell population.

After a pair has been selected, a random number from a uniform distribution (from 0 to 1) is used to determine if the candidate pair actually collides. A new random number is chosen for every pair, even for pairs within the same cell. A collision takes place between parcels i and j if the deviate r satisfies the inequality

$$r < \frac{q_g v_{i,j} \sigma_{i,j}}{(q v \sigma)_{\max}} \tag{1486}$$

The variable q_g represents the greater number of droplets between q_i and q_j. If the collision is accepted, then q_l, the lesser number of droplets, actually participates in the collision. This distinction is important in the case of droplet coalescence, where one parcel of droplets absorbs the other. Whenever employed for a given cell, the NTC algorithm is linearly proportional to N_p.

If the spray is so dense that $M_{\text{cand}} > N_p^2/2$, then direct calculation of collisions using the O'Rourke approach is more efficient than the NTC algorithm for this cell. STAR-CCM+ then switches to the O'Rourke algorithm automatically.

O'Rourke Algorithm

The O'Rourke collision algorithm is a direct technique because it considers all possible collision partners.

The probability of any droplet colliding with any other droplet is given by

$$p_{i,j} = \frac{v_{i,j}\sigma_{i,j}\Delta t}{V} \tag{1487}$$

The mean expected number of collisions between a droplet in parcel i and the droplets in parcel j is given by

$$\bar{\mu} = q_j \frac{v_{i,j}\sigma_{i,j}\Delta t}{V} \tag{1488}$$

The number of collisions is determined by sampling from a Poisson distribution with a mean of $\bar{\mu}$

$$P_n = \frac{\mu}{n!} e^{-\bar{\mu}} \tag{1489}$$

Eqn. (1489) gives the probability of n collisions. P_0 gives the probability of zero collisions. Thus, whenever employed for a given cell, the O'Rourke approach is proportional to N_p^2.

Collision Outcomes

After a collision is detected between two droplets, different collision outcomes are possible: bounce, permanent coalescence, and separation (grazing collision). However, this implementation ignores the formation of satellite droplets which may occur at high collision Weber numbers.

Collision outcomes are described by three non-dimensional parameters:

- the collision Weber number We_{coll}

  $$We_{\text{coll}} = \frac{\rho_l (v_{i,j})^2 (r_1 + r_2)}{\sigma} \tag{1490}$$

  The relative velocity $v_{i,j}$ is

  $$v_{i,j} = v_1 - v_2 \tag{1491}$$

  where v_1 and v_2 are the velocities of the colliding drops, and r_1 and r_2 are their corresponding radii.

- the impact parameter B

  $$B = \frac{b}{r_1 + r_2} \tag{1492}$$

  where b is calculated by taking the distance from the center of one drop to the relative velocity vector that is placed on the center of the other drop. When $B = 0$, the collision is head-on. When $B = 1$, the drops barely graze each other.

- the drop size ratio γ

$$\gamma = \frac{r_2}{r_1} \quad , r_2 > r_1 \tag{1493}$$

The overall collision dynamics algorithm is described in the flowchart [244].



$Y_{r,2}$ is a random number in the range [0,1].

The coalescence efficiency E_{coal} is defined as the probability of coalescence once a collision has occurred. O'Rourke and Bracco [250] employ the approach of Brazier-Smith *et al* [236] to obtain

$$E_{coal} = \min\left[1.0, \frac{2.4f(\gamma)}{We}\right] \tag{1494}$$

where

$$f(\gamma) = \gamma^3 - 2.4\gamma^2 + 2.7\gamma \tag{1495}$$

The bounce efficiency E_{boun} is

$$E_{boun} = \min\left[1.0, \left(\frac{We}{2.4f(\gamma)}\right)^{1/3}\right] \tag{1496}$$

If coalescence occurs, the number of coalescences m that each droplet undergoes is determined by finding the value of m for which

$$\sum_{n=0}^{m-1} P_n < xx < \sum_{n=0}^{m} P_n \tag{1497}$$

where P_n is the probability of n collisions given by Eqn. (1489).

When coalescence occurs, the number of droplets from the less populous parcel is subtracted from the more populous parcel. If the outcome is grazing collision or bounce, only one collision is calculated per droplet. The post-collision velocity, temperature, and mass of the droplets are calculated based on the conservation of mass, momentum, and energy.

Passive Scalar Source Terms

When a Passive Scalar Source node is activated for a Lagrangian phase, the change of value of the passive scalar over time is described by:

$$\varphi_{t+\delta t} = \varphi_t + \frac{S_\varphi}{\rho}\delta t \tag{1498}$$

where:

- φ_t is the value of the scalar at a time t.
- $\varphi_{t+\delta t}$ is the value of the scalar at a later time $t+\delta t$.
- S_φ is the user-defined source term, with units of density/time.
- ρ is the phase density.

If the Passive Scalar Source node is inactive, the scalar value remains constant unless by interaction with a Eulerian passive scalar.

Passive Scalar Interactions

There are two forms of interaction between passive scalar values: volume weighted and area weighted.

Volume Weighted Interaction

Volume weighted interactions model diffusion, with distinct forward and backward rates. This interaction conserves the density weighted passive scalar.

The volume weighted interaction models the mass flow rate as:

$$J = k(\rho_1\varphi_1 - \rho_2\varphi_2)A_p \tag{1499}$$

where:

- J is the mass flow rate.
- k is the interaction coefficient set under the Phase Interaction node.
- ρ_1 is the density of the Lagrangian particle.
- ρ_2 is the density of the Eulerian fluid.
- φ_1 is the value of the passive scalar for the Lagrangian particle.
- φ_2 is the value of the passive scalar for the Eulerian fluid.
- A_p is the surface area of the particle.

The source term for a typical Lagrangian particle in the parcel is:

$$S_{\varphi_1} = (-J)/V_p \tag{1500}$$

where:

- S_{φ_1} is the value of the passive scalar source for the particle.
- V_p is the volume of the particle.

The source term for the Eulerian cell is:

$$S_{\varphi_2} = n_p \cdot J/V_c \tag{1501}$$

where:

- n_p is the number of particles in the parcel.
- S_{φ_2} is the value of the passive scalar source for the fluid.
- V_c is the volume of the cell.

Area Weighted Interaction

Area weighted interactions model adsorption/desorption.

The area weighted interaction models the mass flow rate as:

$$J = (k_{1\rightarrow 2}\rho_1\varphi_1 - k_{2\rightarrow 1}\rho_2\varphi_2)A_p \tag{1502}$$

where:

- J is the mass flow rate.
- $k_{1\rightarrow 2}$ and $k_{2\rightarrow 1}$ are the interaction coefficients set under the Phase Interaction node.
- ρ_1 is the density of the Lagrangian particle.
- ρ_2 is the density of the Eulerian fluid.
- φ_1 is the value of the passive scalar for the Lagrangian particle.
- φ_2 is the value of the passive scalar for the Eulerian fluid.

- A_p is the surface area of the particle.

The source term for a typical Lagrangian particle in the parcel is:

$$S_{\varphi_1} = (-J)/A_p \tag{1503}$$

where:

- S_{φ_1} is the value of the passive scalar source for the particle.
- A_p is the surface area of the particle.

The source term for the Eulerian cell is:

$$S_{\varphi_2} = n_p \cdot J/V_c \tag{1504}$$

where:

- n_p is the number of particles in the parcel.
- S_{φ_2} is the value of the passive scalar source for the fluid.
- V_c is the volume of the cell.

Generalized Gradient Diffusion Hypothesis

The Generalized Gradient Diffusion Hypothesis model (GGDH) provides improved prediction for passive scalar diffusion in highly anisotropic turbulent flows. This hypothesis models the passive scalar diffusion coefficient as a tensor quantity proportional to the Reynolds stresses. The diffusion coefficient is:

$$\mathbf{D} = \frac{\mu}{\sigma}\nabla\varphi + \frac{C_s}{\sigma_t}\tau\,\mathbf{R}\cdot\nabla\varphi \tag{1505}$$

where:

- μ is the viscosity.
- σ is the molecular Schmidt number, a material property.
- φ is the value of the passive scalar.
- $C_s = 0.2$, a model constant.
- σ_t is the turbulent Schmidt number, assumed to have a value of 0.9, consistent with the turbulent Prandtl number used for energy.
- $\tau = k/\varepsilon$, where k is kinetic energy and ε is the dissipation rate.
- $\mathbf{R}$ is the Reynolds stress tensor.

The diffusion coefficient $\mathbf{D}$ is then used in the passive scalar transport equation, Eqn. (164).

Lagrangian Multiphase Bibliography

[228] Ahlert, K. 1994. “Effects of particle impingement angle and surface wetting on solid particle erosion of AISI 1018 steel”, MS Thesis,where: University of Tulsa, USA.

[229] Akao, J., Akaki, K., Mori, S., and Moriyama, A. 1980. “Deformation behaviors of a liquid droplet impinging on to a hot metal surface”, *Transactions of the Iron and Steel Institute of Japan*, **20**, pp. 737-743.

[230] Auton, T.R., Hunt, J.C.R., and Prud’Homme, M. 1988. “The force exerted on a body in inviscid unsteady non-uniform rotational flow”, *J. Fluid Mech.*, **197**, pp. 241–257.

[231] Baumgarten, C. 2006. “Mixture Formation in Internal Combustion Engines”, Springer-Verlag.

[232] Bai, C. 1996. “Modeling of spray impingement processes”, Ph.D Thesis, University of London.

[233] Bai, C., and Gosman, A.D. 1995. “Development of methodology for spray impingement simulation”, SAE Technical Paper Series 950283.

[234] Bai, C.X., Rusche, H., and Gosman, A.D. 2002. “Modeling of gasoline spray impingement”, *Atomization & Sprays*, **12** (1-3), pp. 1–27.

[235] Birkhold, F., Meinast, U., Wassermann, P., and Deutschmann, O. 2006. "Analysis of the Injection of Urea-water-solution for Automotive SCR DeNOx-Systems: Modelling of Two-phase Flow and Spray-Wall Interaction", SAE-2006-01-0643.

[236] Brazier-Smith, P., Jennings, S., and Latham, J. 1972. “The Interaction of Falling Rain Drops: Coalescence” Proceedings of the Royal Society of London A, **326**, pp. 393–408.

[237] Crowe, C.T., Sommerfeld, M., and Tsuji, Y. 1998. *Multiphase Flows with Droplets and Particles*. CRC Press, Boca Raton.

[238] Eckhause, J.E., Reitz, R.D. 1995. “Modeling heat transfer to impinging fuel sprays in direct-injection engines”, *Atomization and Sprays*, **5**, pp. 213–242.

[239] El Wakil, M.M., Ueyhara, O.A., and Myers, P.S. 1954. “A theoretical investigation of the heating-up period of injected fuel droplets vaporizing in air”, NACA Tech. Note 3179.

[240] Gosman, A.D., and Ioannides, E. 1983. “Aspects of computer simulation of liquid-fueled combustors”, AIAA, *J. Energy*, **7**(6), pp. 482–490.

[241] Han, Z., Parish, S., Farrell, P.V., and Reitz, R.D. 1997. “Modeling Atomization Processes of Pressure-Swirl Hollow-Cone Fuel Sprays”, *Atomization and Sprays*, **7**(6) pp. 663–684.

[242] Haugen, K., Kvernvold, O., Ronold, A., and Sandberg, R. 1995. “Sand erosion of wear-resistant materials: Erosion in choke valves”, Wear, **186-187**, pp. 179–188.

[243] Incopera, F.P., and DeWitt, D.P. 1996. *Fundamentals of Heat and Mass Transfer*. 4th Edition, John Wiley & Sons

[244] Kralj, C. 1995. “Numerical Simulation of Diesel Flow Processes”, Ph.D. Thesis, Imperial College, London.

[245] Lefebvre, A.H. 1989. *Atomization and Sprays*, Taylor & Francis, Bristol PA

[246] Liu, A.B., Mather, D., and Reitz, R.D. 1993. “Modeling the Effects of Drop Drag and Breakup on Fuel Sprays”, SAE Paper 930072.

[247] McLaury, B.S., Shirazi, S.A., Shadley, J.R., and Rybicki, E.F. 1996. "Modeling erosion in chokes", ASME FED conference, **236**(1), pp. 773–781.

[248] Milne-Thompson, L.M. 1968. *Theoretical Hydrodynamics,* 5th edition, McMillan & Co., New York.

[249] Neilson, J.H., and Gilchrist, A. 1968. "Erosion by a stream of solid particles", Wear, **11**, pp. 111–122.

[250] O'Rourke, P.J., and Bracco, F. 1980. "Modeling of Drop Interactions in Thick Sprays and a Comparison with Experiments" Proceedings of the Institution of Mechanical Engineers, **9**, pp. 101–106.

[251] O'Rourke, P.J. 1981. "Collective Drop Effects on Vaporizing Liquid Sprays", Department of Mechanical and Aerospace Engineering, Princeton University.

[252] O'Rourke, P.J., and Amsden, A.A. 1987. "The TAB Method for Numerical Calculation of Spray Droplet Breakup", SAE Paper 872089.

[253] Oka, Y.I, Okamura, K., and Yoshida, T. 2005. "Practical estimation of erosion damage caused by solid particle impact. Part 1: Effect of impact parameters on a predictive equation", Wear, **259**, pp. 95–101.

[254] Oka, Y.I., and Yoshida, T. 2005. "Practical estimation of erosion damage caused by solid particle impact. Part 2: Mechanical properties of materials directly associated with erosion damage", Wear, **259**, pp. 102–109.

[255] Patterson, M.A., and Reitz, R.D. 1998. "Modeling the Effects of Fuel Spray Characteristics on Diesel Engine Combustion and Emission", SAE Paper 980131.

[256] Rangel, R.H., and Sirignano, W.A. 1991. "The linear and nonlinear shear instability of a fluid sheet", *Phys. Fluids A.*, **3** p. 2392.

[257] Ranz,W.E., and Marshall,W.R. 1952. "Evaporation from drops Parts I and II", *Chem. Eng. Prog.*, **48**(3), p. 141.

[258] Reitz, R.D. 1987. "Modeling Atomization Processes in High-Pressure Vaporizing Sprays", Atomisation and Spray Technology, **3**, pp. 309–337.

[259] Reitz, R.D., and Diwakar, R. 1986. "Effect of Drop Breakup on Fuel Sprays", SAE Paper 860469.

[260] Reitz, R.D., and Diwakar, R. 1987. "Structure of High-Pressure Fuel Sprays", SAE Paper 870598.

[261] Ricart, L.M., Reitz, R.D., Dec, J.M. 1997. "Comparisons of Predicted Diesel Spray Liquid Penetration and Vapor Fuel Distributions with In-Cylinder Optical Measurements", Paper No. 97-ICE-52, ICE-Vol. 29-2, 1997 Fall Technical Conference, ASME.

[262] Rohsenow, W.M. 1952. 'A method of correlating heat transfer data for surface boiling liquids', Transactions of the ASME, **74**, p. 969.

[263] Satoh, K., Kawai, T., Ishikawa, M., and Matsuoka, T., 2000. "Development of Method for Predicting Efficiency of Oil Mist Separators", *SAE Technical Paper Series*, 2000-01-1234

[264] Schiller, L., and Naumann, A. 1933 "Ueber die grundlegenden Berechnungen bei der Schwerkraftaufbereitung", *VDI Zeits.*, **77**(12), pp. 318–320.

[265] Schmidt, D.P., Nouar, I., Senecal, P.K., Rutland, C.J., Martin, J.K., Reitz, R.D., and Hoffman, J.A. 1999. "Pressure-Swirl Atomization in the Near Field", *SAE Technical Paper Series*, 1999-01-0496.

[266] Schmidt, D.P., and Rutland, C.J. 2000. "A New Droplet Collision Algorithm", *J Comput Phys*, **164**, pp. 62–80.

[267] Senecal, P.K., Richards, K.J., Pomraning, E., Yang, T., Dai, M.Z., McDavid, R.M., Patterson, M.A., Hou, S., and Shethaji, T. 2007. "A New Parallel Cut-Cell Cartesian CFD Code for Rapid Grid Generation Applied to In-Cylinder Diesel Engine Simulations", SAE Paper 2007-01-0159.

[268] Senecal, P.K., Schmidt, D.P., Nouar, I., Rutland, C.J., and Reitz, R.D. 1999."Modeling High Speed Viscous Liquid Sheet Atomization", *Int. J. Multiphase Flow*, **25**, pp. 1073–97.

[269] Sirignano, W.A. 1999. *Fluid dynamics and transport of droplets and sprays*, Cambridge University Press, Cambridge.

[270] Spalding, D.B. 1960. "A standard formulation of the steady convective mass transfer problem", *Int. J. Heat Mass Transfer*, **1**, pp. 192-207.

[271] Squire, H.B. 1953. "Investigation of the instability of a moving liquid film", *Brit. J. Appl. Phys.*, **4** p. 167.

[272] Stiesch, G. 2003. "Modeling Engine Spray and Combustion Processes", Springer (Berlin).

[273] Wallace, M.S., Peters, J.S., Scanlon, T.J., Dempster, W.M., McCulloch, S., and Ogilvie, J.B. 2000. "CFD-based erosion modeling of multi-orifice choke valves", ASME, Paper FEDSM2000-11244.

[274] Wruck, N. 1998. "Transientes Seiden von Tropfen beim Wandaufprall", Ph.D. Thesis, RWTH Aachen.

Lagrangian Multiphase Field Functions Reference

The following field functions are made available to the simulation when the Lagrangian Multiphase model is used. Where applicable, the Lagrangian phase model which makes the field function available is given in brackets after the description.

A number of the field functions are only available when the Lagrangian Multiphase solver's Temporary Storage Retained property is ticked. These are indicated by the word "temporary" after the description.

Accretion Ratio	The ratio of the rate at which ice melts to the rate at which droplets impinge, according to $-\dot{m}_{melt}/\dot{m}_d$.
Aborted Parcel Count	The number of parcels that are aborted in each cell during tracking. (temporary)
Droplet Dynamic Viscosity	The dynamic viscosity of the droplet material, μ_l. (TAB Distortion, KHRT breakup, Primary Atomization)
Droplet Latent Heat of Vaporization	The latent heat of vaporization of the droplet material, L. (Droplet Evaporation)
Droplet Saturation Pressure	The absolute saturation pressure of the droplet material, p_{sat}. (Droplet Evaporation)
Droplet Surface Tension	The surface tension of the droplet material, σ. (TAB Distortion, Secondary breakup, Primary Atomization)
Droplet Weber Number	The droplet Weber number, Eqn. (1311). (Secondary breakup)

Eddy Velocity Fluctuation	The difference between the instantaneous and Reynolds-Averaged velocity in eddies of the particles. It is the quantity $\mathbf{v}'$ in Eqn. (1385). (Turbulent Dispersion, temporary)
Erosion Rate	The rate at which material is being eroded from a boundary. It is the quantity E_f in Eqn. (1391). (Erosion)
Incident Mass Flux of *Phase*	The mass flux of material particles of *phase* impinging on the boundary. (Material Particles)
Incident Particle Flux of *Phase*	The rate of massless particles of *phase* impinging on the boundary. (Massless Particles)
Injector Index	The index of the injector which created each parcel.
Lagrangian Charge Density	The total particle charge per unit cell volume. (Coulomb Force, Two-Way Coupling.)
Lagrangian Energy Source	The rate of energy transfer to the continuous phase from all Lagrangian phases with a Two-Way Coupling model. It is the quantity S_E in Eqn. (1366) and Eqn. (1368). (Material Particles, Two-Way Coupling, Energy, temporary)
Lagrangian Mass Source	The rate of mass transfer to the continuous phase from all Lagrangian phases with a Two-Way Coupling model. It is the quantity S_m in Eqn. (1321) and Eqn. (1323). (Material Particles, Two-Way Coupling, Mass Transfer, temporary)

Lagrangian Momentum Source	The rate of momentum transfer to the continuous phase from all Lagrangian phases with a Two-Way Coupling model. It is the quantity $\boldsymbol{S}_v$ in Eqn. (1351) and Eqn. (1353). (Material Particles, Two-Way Coupling, temporary)
Lagrangian Species Source (*Species*)	The rate of mass transfer of *species* to the continuous phase from all Lagrangian phases with a Two-Way Coupling model. (Material Particles, Two-Way Coupling, Mass Transfer, Species, temporary)
Number Density of *Phase*	The number of massless particles of *phase* per unit volume of space. (Massless Particles)
Number of Collisions	The number of collisions in a cell for a particular time-step. (No Time Counter, temporary)
Parcel Centroid	The position of parcels, $\mathbf{r}_\pi$.
Parcel Index	A unique identifier for each parcel. The index is unique within a Lagrangian phase.
Parcel Mass	The mass of parcels, $m_\pi = n_\pi m_p$. (Material Particles). (Unsteady)
Parcel Mass Flow Rate	The mass of parcels, $\dot{m}_\pi = \dot{n}_\pi m_p$. (Material Particles). (Steady)
Parcel Time Step	The local time-step, δt_p in Eqn. (1313) and Eqn. (1314).
Parcel Volume	The volume of parcels, $V_\pi = n_\pi V_p$. (read-only). (Unsteady)

Parcel Volume Flow Rate	The volume of parcels, $\dot{V}_\pi = \dot{n}_\pi V_p$. (Material Particles). (Steady)
Particle Charge	The particle charge q. (Coulomb Force).
Particle Connected Component Size	The number of spheres in the particle cluster to which the given particle belongs. See the field function Particle Connected Components.
Particle Connected Components	The labels (integers) of particle clusters. These clusters are groups of particles that form a single connected network through particle-particle contact.
Particle Count	The number of particles represented by a parcel, n_π. (Unsteady)
Particle Density	The density of the particle material, ρ_p. (Material Particles)
Particle Diameter	The diameter of particles, D_p. (Material Particles)
Particle Flow Rate	The rate at which particles flow along parcel tracks, $\dot{n}_\pi$. (Steady)
Particle Incidence Angle	The angle between the particle velocity vector and the plane of face on which it is incident, if any, $\theta = \mathrm{asin}(\mathbf{v}_p \cdot \mathbf{n} / \lvert\mathbf{v}_p\rvert)$. This function is provided for use in boundary condition expressions, it is not available for visualization, reporting, or recording on the track file. (Material Particles)
Particle Mass	The mass of single particles, m_p. Note, not the mass of parcels. (Material Particles)

Particle Mass Fraction of *species*	The mass fraction, $Y_{pi} = m_{pi}/m_p$, of species *i* in a single particle.
Particle Molar Concentration of *Species*	The molar concentration, $C_{pi} = \rho_p Y_{pi}/W_{pi}$, of species *i* in a single particle, where ρ_p is the particle mixture density, Y_{pi} is the species mass fraction, and W_{pi} is the species molecular weight.
Particle Mole Fraction of *Species*	The mole fraction, $X_{pi} = Y_{pi} W_p / W_i$, of species *i* in a single particle, where Y_{pi} is the species mass fraction, W_p is the molecular weight of the particle mixture, and W_{pi} is the molecular weight of the species.
Particle Nusselt Number	The non-dimensional heat transfer coefficient of particles, Nu_p, defined in Eqn. (1362). (Lagrangian Energy)
Particle Residence Time	The "age" of particles as measured from time of injection.
Particle Reynolds Number	The quantity Re_p in Eqn. (1340). (Material Particles)
Particle Sherwood Number	The non-dimensional mass transfer conductance, Sh_p, defined in Eqn. (1371). (Mass Transfer)
Particle Slip Velocity	The velocity deficit of particles relative to the continuous phase, $\mathbf{v}_s = \mathbf{v} - \mathbf{v}_p$. (Material Particles)
Particle Specific Heat	The specific heat of particles. (Lagrangian Energy)

Particle Surface Area	The surface area of single particles, A_s. Note, not the surface area of parcels. (Material Particles)
Particle Temperature	The temperature of particles, T_p. (Lagrangian Energy)
Particle Velocity	The absolute velocity of particles, $\mathbf{v}_p$.
Particle Volume	The volume of single particles, $V_p = m_p/\rho_p$. Note, not the volume of parcels. (Material Particles)
TAB Distortion	The normalized distortion of droplets, y in Eqn. (1419). (TAB Distortion, temporary)
TAB Distortion Rate	The rate of change of TAB Distortion, $\dot{y}$. (TAB Distortion, temporary)
Void Fraction	The fraction of the cell volume available to the fluid, unoccupied by a solid or dispersed phase.
Volume Fraction of *Phase*	The volume of material particles of *phase* per unit volume of space. (Material Particles)

Using the Fluid Film Model

Problems where a thin film of fluid exists on solid boundaries are common in engineering practice. Examples include fluid film on the internal surfaces of automotive engines, gas turbines, spray-cooling systems and ink-jet printers. The Fluid Film Model in STAR-CCM+ provides a mathematical description of the behavior of such films, subject to certain assumptions, as described in the Formulation.

Depending on the conditions, complex interactions can occur between the fluid film and the surrounding environment. These interactions require further mathematical modeling to capture the details of the developed flows. Examples of physical processes that can be modeled with fluid films include gravity driven flow, shear forces on the film interface, external forcing through user-defined momentum source terms and evaporative cooling.



There are different methods by which a film can be formed and removed:

- You can specify the properties of an initial layer of fluid film on a boundary surface.
- A fluid film can flow from a fluid film inlet, and can be removed at an outlet.
- You can specify both a fluid film mass source and a fluid film mass sink to form or remove a fluid film.
- A film can be formed as droplets impact boundary surfaces. For this method, you define injectors for fluid particles with the Lagrangian Multiphase model and a Film-Lagrangian Phase Interaction model with impingement. You can also model the impingement of a dispersed Eulerian phase on a surface to form a fluid film.

  It is also possible to model the stripping of droplets from the surface of a fluid film due to wave-induced instabilities or sharp-angled edges. Droplets that are formed in this way are later tracked using the Lagrangian Multiphase model.
- A film can be formed as a gas, or gas component, condenses onto a surface. The film can be removed as liquid evaporates from the surface.

For this method, define a Film-Eulerian Phase Interaction together with the Film Evaporation and Condensation model. The Film Boiling model is available as an extension of the Film Evaporation and Condensation model.

You can model fluid film melting and solidification. STAR-CCM+ also supports multi-component fluid film models: you can model a mixture of different species components within the film.

Fluid film is represented in the STAR-CCM+ continuum by the Fluid Film node, which is subject to common menu items.



To begin working with the fluid film model:

- Set up your continuum
- Define the fluid film phase with phase models
- Create the shell region and initialize the film

The Fluid Film model activates a solver so that you can customize calculations.

Fluid film thickness, surface velocity values, and interfacial shear stresses are field functions that you can post-process, along with other results.

Selecting the Fluid Film Model

To use the Fluid Film model:

1. In the *Physics Model Selection* dialog, select Three Dimensional from the *Space* group box.
2. A fluid film simulation is a transient analysis, so select Implicit Unsteady from the *Time* group box.
3. Specify the material model as Gas or Multi-Component Gas.
4. Select Coupled Flow or Segregated Flow from the *Flow* group box.
5. Select Fluid Film from the *Optional Models* group box.

The *Physics Model Selection* dialog contains your selections similar to the example shown in the following screenshot. In this example, the option Segregated Flow was selected.



6. Select the appropriate optional models:
   - If you intend to model fluid film stripping or impingement, select the Multiphase Interaction and the Lagrangian Multiphase models.
   - If you intend to model heat transfer or energy through the fluid film, select Coupled Energy or one of the Segregated Fluid Energy models.
   - If you intend to model film evaporation and condensation, select the Multiphase Interaction model.

Defining Fluid Film Phases

The process of selecting the Fluid Film model in the *Physics Model Selection* dialog automatically activates the object that controls the definition of fluid film phases.



The model contains no fluid film phases in its initial state. The process for creating and defining each required phase involves the following steps:

1. Create a fluid film phase
2. Select the models
3. Select the material
4. Customize material properties

Repeat these steps until all fluid film phases have been specified.

Create a New Fluid Film Phase

The Fluid Film Phases manager node (which resides within the Fluid Film model node) is initially unpopulated.

- Right-click this node and select New from the pop-up menu.



A new fluid film phase node appears beneath the manager node, with its own Models sub-node.



Once a phase has been created, it can be renamed so that it can be clearly identified in the simulation.

Select the Models

Each fluid film phase has models that are associated with it. These models are selected using the *Phase Model Selection* dialog.

1. Right-click the Models node and then click Select models...



2. Select the recommended models.

   An example phase model selection dialog is shown below.



3. Click Close to close the dialog.

Nodes corresponding to the chosen models appear within the node tree for the fluid film phase.



Select the Material

The default material model appears within the Models node for the phase. To replace this material:

1. Right-click the material node and select Replace with...



A dialog appears with a set of available replacement materials.



2. Select the desired replacement material in the tree and click OK.

Customize Material Properties

Further changes to material properties can be made by accessing the individual material property nodes within the material node as shown.



Fluid Film Phase Manager Pop-Up Menu

This pop-up menu item is available when you right-click the Fluid Film Phases manager node.

New	Creates a fluid film phase with a default name and a Models sub-node.

Working with Fluid Film Phase Models

Some fluid film phase models are variants of STAR-CCM+ continuum models that are customized for fluid film conditions. The following is a summary of these fluid film phase models:

- Shell Three-Dimensional Space

- Liquid Material
- Multi-Component Liquid Material
- Fluid Film Flow
- Segregated Temperature
- Coupled Temperature
- Laminar Viscous Regime
- Constant Density Equation of State

Modeling Shell Three-Dimensional Space

The shell three-dimensional space model is similar to the three-dimensional space model for a physics continuum. Activating this model lets you create a shell region from a region boundary.

The Shell Three Dimensional node represents this model for the phase.



Selecting the Shell Three-Dimensional Space Model

- In the *Phase Model Selection* dialog, select the Shell Three Dimensional option from the *Space* group box.



Modeling Liquid Materials

The fluid film Liquid model is similar to the Liquid model for a physics continuum.

The Liquid node represents this model for the phase. Its child node is a single-component liquid material to which the conventions for managing a single-component liquid apply.



Selecting the Liquid Material Model

- In the *Phase Model Selection* dialog, select the Liquid option from the *Material* group box.



Modeling Multi-Component Liquid Materials

The fluid film Multi-Component Liquid model is similar to the Multi-Component Liquid model for a physics continuum.

The Multi-Component Liquid node represents this model for the phase. Its child node is a liquid mixture material to which the conventions for managing a multi-component liquid apply.



The fluid film evaporation and condensation model is available only when the Multi-Component Liquid material model is selected. You model the evaporation or condensation of a single liquid component by selecting only one liquid component in a liquid mixture.

Selecting the Multi-Component Liquid Material Model

- In the *Phase Model Selection* dialog, select Multi-Component Liquid from the *Material* group box.



Modeling Fluid Film Flow

The Fluid Film Flow model solves conservation equations for mass and momentum. Film thickness comes from the mass equation. These equations are solved in a segregated manner, that is, one after the other. The energy equation, however, can either be solved segregated or as part of the coupled energy equation. A parabolic velocity profile is assumed in the direction normal to the wall.

The Flow Model node represents this model for the phase.



Flow Model Properties

Maximum Volume Fraction	Sets the maximum volume fraction ($\phi_{film,\,max}$ in Eqn. (1516)) for the fluid film phase. Setting this property allows you to include the effect of the liquid film on the void fraction. The volume that the film takes up is no longer available to the gas, in the cells that neighbor the shell region. To ensure stability, set the maximum volume fraction lower than 1 (for example, 0.75). To deactivate the volume fraction effect, set this value to 0.

Flow Model Expert Properties

Maximum Film Thickness	Sets the maximum film thickness (h_f in Eqn. (1507)) for the fluid film phase.

Selecting the Fluid Film Flow Model

- In the *Phase Model Selection* dialog, select Flow Model from the *Fluid Film Flow* group box.



Modeling Segregated Temperature Energy

The Segregated Temperature model solves the total energy equation for the fluid film phase, which is similar to the Segregated Temperature model applied to the physics continuum. It also accounts for energy transferred due to droplets being added or removed.

The Segregated Temperature node represents this model for the phase.



Segregated Temperature Energy Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Selecting the Segregated Temperature Energy Model

To select the Segregated Temperature film phase model:

1. Select the Segregated Temperature model for the continuum.

2. In the *Phase Model Selection* dialog, select Segregated Temperature from the *FluidFilmSegregatedEnergy* group box.



Modeling Coupled Temperature Energy

The Coupled Temperature model solves the total energy equation for the fluid film phase. It also accounts for energy transferred due to droplets being added or removed.



- To access the expert properties of the Coupled Temperature model, select the Coupled Temperature node.

Coupled Temperature Energy Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Selecting the Coupled Temperature Energy Model

To select the Coupled Temperature film phase model:

1. Select the Coupled Energy model for the continuum.
2. In the *Phase Model Selection* dialog, select Coupled Temperature from the *FluidFilmCoupledEnergy* group box.



Modeling the Laminar Viscous Regime

The Laminar Viscous Regime model defines the flow within the film phase as being laminar, as opposed to being inviscid, or turbulent. In this respect, it is similar to the standard Laminar model.

The Laminar node represents this model for the phase.



Selecting the Laminar Viscous Regime Model

To select the Laminar Viscous Regime model:

1. Select the Fluid Film Flow model for the phase.

   If the Auto-select recommended models checkbox is activated, the Laminar *Viscous Regime* is selected automatically when selecting the Fluid Film Flow model.

2. If the Auto-select recommended models checkbox is deactivated, select Laminar from the *Viscous Regime* group box.



Modeling the Constant Density Equation of State

The Constant Density model defines the fluid film phase as incompressible, with a single value of density throughout the phase. In this respect, it is similar to the standard constant density model.

The Constant Density node represents this model for the phase.



Selecting the Constant Density Equation of State Model

To select the Constant Density model:

1. Select the Liquid or Multi-Component Liquid model for the phase.
2. In the *Phase Model Selection* dialog, select Constant Density from the *Equation of State* group box.



Using Shell Regions

A fluid film is not modeled directly within a regular STAR-CCM+ region. Instead, STAR-CCM+ uses a shell region that you create as the space within which the fluid film flows. It is represented in the object tree by a shell region node (Film_Wall-Y shell and ReceivingFilm_Wall-Y shell in the following screenshot), which has properties and a pop-up menu.



The shell region is effectively two-dimensional—only one cell thick. Its boundaries are *edges*. Working with a shell region requires creating the region, and then setting up its interfaces and conditions for movement of a fluid along a surface.

Creating Shell Regions

A shell region is a surface domain in space. Shell regions can be created from existing region boundaries.

To set up a shell region:

1. Decide which boundary is the location where the fluid film forms.
2. Right-click the node of that boundary and select Create Shell Region.



A new shell region is created.



Each shell region is given a unique name that is based on the name of the boundary which is used to create the shell region, but it can be renamed.

3. Select the shell region node and select the fluid film phase that is associated with the shell region.



This action adds sub-nodes to the shell region and boundary objects, allowing you to define appropriate conditions for the fluid film.



4. Select the individual boundary nodes in the newly created shell region and specify the boundary type.



The boundary type that you select applies conditions for that type to the fluid film.

Multiple Boundaries

You can select multiple boundaries using the `<Ctrl>` key at the same time, and convert them to one or more shell regions in a single operation. Boundaries within the same fluid region belong to the same shell region. If film passes from one shell region to another, join the two regions using an interface.

Joining Shell Regions

One shell region is joined to a region or another shell region using an interface. When you create a shell region, an interface between the original fluid region and the shell region is created automatically at the boundary that you selected for creating the shell region.



When the fluid film is expected to flow over multiple shell regions, an interface must exist between them. These different shell regions are created from boundaries belonging to different regions. Similarly, if shell regions have been created separately within a region, create an interface manually between them.



- To create this interface, select adjacent edge boundaries within the shell regions, and create an in-place interface.



The resulting interface node has its own icon to indicate a shell-shell interface.



Setting Film Boundary Conditions

Once the shell region is connected to a fluid film phase, it becomes an active part of the simulation. Therefore, the shell region and the boundary objects acquire typical physics condition and physics value sub-nodes.



With these objects, you can work with boundary conditions and values for the fluid film flow just as you would for any other flow in STAR-CCM+.

Using Motions and Moving Reference Frames

STAR-CCM+ supports motions and moving reference frames for fluid film models.

To model a moving shell region, assign the appropriate motion or moving reference frame to the volume region from which the shell was created. You specify the motion and the reference frame in the Physics Values > Motion Specification node in the volume region.



The specified motion and reference frame are set in the Physics Values > Motion Specification node in the associated shell region. These settings are read-only in the shell region.



The Lead Region property indicates the volume region that the shell region follows. This region is the volume region from which the shell region was created.

Defining Initial Film Conditions

The shell region includes the Initial Condition Option node so that you can specify initial conditions that are unique to that region. You can use this node to set initial film conditions on a particular shell region.



When you set the Option property to `Specify Region Values`, you are excluding this shell region from the global settings of the continuum. Instead, the shell region has its own set of initial conditions. As you can see from the following screenshot, initial conditions include the Fluid Film Thickness and its Velocity.



Initial temperature is also available when you activate the model

Accepting Impinging Droplets

For simulations where the Lagrangian Multiphase model is selected and fluid droplets are present, fluid films can be formed. These films form when the fluid droplets impinge the interface between a fluid film shell region and the fluid region.

For consistency, select the same materials for the fluid film and Lagrangian phases. Although differences in material properties are allowed between the phases, avoid these differences. They would lead to improper conservation of mass or enthalpy.

To model impinging droplets:

1. Select the Lagrangian phase models for modeling impingement.
2. Select the Film-Lagrangian Phase Interaction model.
3. Select the impingement model.

4. Set the fluid-film boundary mode.

Selecting Lagrangian Phase Models for Impingement

Select the correct Lagrangian phase models so that the Lagrangian droplets can exchange mass, momentum, and energy with the fluid film. If the densities of the phases differ, volume is not conserved.

1. For the Lagrangian phase, select Two-Way Coupling from the *Optional Models* group box.



2. Additionally, you can also select the Bai-Gosman or Satoh Wall Impingement models to model interactions between droplets and walls in greater detail.

Selecting the Film-Lagrangian Phase Interaction Model

Select the Film-Lagrangian Phase Interaction model before modeling droplet impingement or film stripping:

1. Create a phase interaction.
2. Select Multiphase Interaction > Phase Interactions > [Phase Interaction] > Models and choose *Select Models...*

3. In the *Phase Interaction Model Selection* dialog, select Film-Lagrangian Phase Interaction from the *Phase Interaction Topology* group-box.



4. Click Close.

   The Film-Lagrangian Phase Interaction node appears in the simulation tree.



5. In the Properties window, set Fluid Film Phase to the fluid film continuum and Lagrangian Phase to the Lagrangian phase.
6. If modeling a multi-component fluid film and a multi-component droplet, map the components between the two phases.

Film-Lagrangian Phase Interaction Properties

Fluid Film Phase	The fluid film continuum interacting with the Lagrangian phase.
Lagrangian Phase	The Lagrangian phase interacting with the fluid film phase.

Selecting the Impingement Model

The Impingement model can be selected from the Film-Lagrangian Phase Interaction once the Film-Lagrangian Phase Interaction model has been defined:

1. Select Multiphase Interaction > Phase Interactions > [Phase Interaction] > Models and choose *Select Models...*
2. In the *Phase Interaction Model Selection* dialog, select Impingement from the *Optional Models* group box.



3. Click Close.

The Impingement node appears in the simulation tree.



Setting the Fluid-Film Lagrangian Phase Boundary Condition

When you create a shell region for the fluid film, an additional boundary type, Fluid-Film Boundary, is added within the Boundary Conditions node of each Lagrangian phase.



Within each participating Lagrangian droplet phase, set the boundary interaction mode for Fluid-Film boundaries to Fluid Film. This option transfers all Lagrangian droplets to the fluid film. Alternatively, you can the Bai-Gosman or Satoh Wall Impingement models and set the boundary interaction mode accordingly. This method models the impingement in more detail by calculating which droplets rebound or stick.

To set the boundary interaction mode globally for all fluid film boundaries:

1. Select the Fluid-Film Boundary > Boundary Conditions > Mode node.
2. In the *Properties* window, set Active Mode to Fluid Film.



3. If you are using the Bai-Gosman or Satoh Wall Impingement models, select the Bai-Gosman or Satoh options respectively.

To set the boundary interaction mode locally for specific boundaries:

1. For the boundary in question, select the Physics Conditions > Lagrangian Specification node.

2. In the *Properties* window, set Method to Specify for Boundary.



This selection activates a Phase Conditions manager node in that boundary node.

3. Select the Phase Conditions > Phase > Physics Conditions > Mode node.

4. In the *Properties* window, set Active Mode to Fluid Film (or Bai-Gosman / Satoh).



The Fluid Film, Bai-Gosman, or Satoh options are available only when the Film-Lagrangian Phase Interaction is valid. Removing or adding models to the film phase or Lagrangian phase after setting a Lagrangian boundary condition can cause the Film-Lagrangian Phase Interaction to become temporarily invalid. In that case, the Active Mode defaults to Rebound without warning. Always check if the boundary mode is still set Fluid Film when changing the phase models.

Activating Film Stripping

In STAR-CCM+, you can model fluid being removed from a film due to wave stripping across the surface, or edge stripping on sharp corners. Waves generate instabilities on the surface, and these instabilities form droplets which are stripped off the film. These droplets are entrained in the bulk fluid flow. Similarly, droplets are also stripped from the film when the fluid film flows over a sharp edge. Droplets that enter the bulk flow are tracked using the Lagrangian Multiphase model in STAR-CCM+.

Setting up a film stripping simulation includes the same basic technique of setting up a fluid film shell region. Lagrangian features are added with the selection of Lagrangian Multiphase in the continuum, activating wave and/or edge stripping, and using film stripping injectors.

It is possible for droplets to form on the surface of a fluid film, and to be immediately forced back into the film on the subsequent time step.

To activate film stripping:

1. In the *Lagrangian Phase Model Selection* dialog, select Two-Way Coupling from the *Optional Models* group box.



2. Set up a Film-Lagrangian Phase Interaction.
3. Set its phase component properties before continuing.

4. In the *Phase Interaction Model Selection* dialog, select Wave Stripping and/or Edge Stripping from the *Optional Models* group box.



5. If modeling a multi-component fluid film and a multi-component droplet, map the components between the two phases.

Film Stripping Injector

A special type of injector is used to signify a film surface from which droplets leave the film and are injected into the bulk flow.

1. Create a Lagrangian injector, and set the injector Type property to Film Stripping.



2. Set the Parts property to Fluid-Film Interface boundary on the original fluid region.



3. Set the Lagrangian Phase property to the previously created Lagrangian Phase.



4. Set injector conditions and values, namely the Time Randomization condition and the Parcel Streams value.

Modeling Wave Film Stripping

Wave film stripping models the ejection of droplets from the fluid film as a result of wave-induced instabilities. A guide to using wave film stripping is provided in the Activating Film Stripping section. The theory and

methodology for wave film stripping is provided in detail in the Formulation section. The Wave Stripping node represents this model for the phase interaction.



The Wave Stripping model assumes that sinusoidal perturbations disturb the flat film surface, and that the most unstable wavelength grows linearly. In some cases, these assumptions are not satisfied: when a film drips from underneath a horizontal surface or ceiling, for example. To obtain correct results, you can tune the model with a validation study. The following parameters, available under Properties, are provided:

- c_H to control the thickness threshold at which stripping occurs
- c_D to control the size of the droplets that result from stripping
- c_B to control the rate at which film mass is stripped

Consider the following example: a case with a “hanging film” building up on a ceiling. One would intuitively expect the stripping speed to increase as the film grows thicker, eventually reaching some equilibrium thickness. The parameters make it possible (in a simulation-dependent manner) to retrieve this more physical behavior.



Wave Stripping Properties

cH	The coefficient (c_H in Eqn. (1518)) for minimum film thickness to initiate stripping. The default setting is `1.0`. Lowering this value makes wave stripping occur at a lower film thickness threshold.
cD	The droplet diameter coefficient (c_D in Eqn. (1520)). The default setting is `3.78`. Increasing this number leads to bigger particles (but the total amount of mass that is stripped remains the same).
cB	The breakup rate scaling coefficient (c_B in Eqn. (1522)). The default setting is `1.0`. Increasing this number increases the rate of wave stripping (more particles per time-step).
Gas Velocity Scaling Factor	The constant that empirically relates the shear velocity to the free stream velocity ($K_{*,\infty}$ in Eqn. (1524)). The default setting is `10.0`.

Selecting the Wave Film Stripping Model

To select the Wave Stripping film phase model:

1. Select the Lagrangian Multiphase model for the continuum.
2. Select the Film-Lagrangian Phase Interaction model.
3. In the *Phase Interaction Model Selection* dialog, select Wave Stripping from the *Optional Models* group box.



Modeling Edge Film Stripping

Edge film stripping models the ejection of droplets from the fluid film when it flows over a sharp edge. A guide to using edge film stripping is provided in the Activating Film Stripping section. The theory and methodology for edge film stripping is provided in the Formulation section.

The Edge Stripping node represents this model for the phase interaction.



Edge Film Stripping Properties

Minimum Corner Angle	Defines the minimum sharp edge angle to which the edge film stripping model is applied. This angle allows you to define which edges cause stripping, and which edges are due to mesh discretization.

Edge Film Stripping Expert Properties

q	A shape factor that is used in the Rosin-Rammler droplet size distribution (see Eqn. (1535)). Reducing the value of q generates statistically smaller droplets.
Critical FR	The critical force ratio that must be exceeded before edge stripping occurs (see Eqn. (1525)). This property also influences the stripping fraction (see Eqn. (1531)). Reducing the critical force ratio increases the occurrence of edge stripping.
c1	The droplet diameter scale factor, which is used to scale the parent droplet size (see Eqn. (1532)). Reducing this value reduces the droplet size.
Limit parent droplet diameter	Specifies whether the maximum parent droplet diameter is limited. Activated — Activate the upper limit for parent droplet diameter. Deactivated — Deactivate the upper limit for the parent droplet diameter.
Maximum diameter	Sets the maximum parent droplet diameter. This limit is only applied if the Limit parent droplet diameter property is activated.
Randomized droplet diameter	Specifies whether the droplet diameter is calculated or randomized.

Activated	Generate randomized droplet diameters using the Rosin-Rammler size distribution (see Eqn. (1535)).
Deactivated	Generate droplets whose diameters are the average of the size distribution (see Eqn. (1537)).

Selecting the Edge Film Stripping Model

To select the Edge Stripping film phase model:

1. Select the Lagrangian Multiphase model for the continuum.
2. Select the Film-Lagrangian Phase Interaction model.
3. In the *Phase Interaction Model Selection* dialog, select Edge Stripping from the *Optional Models* group box.



Mapping the Components Between a Multi-Component Fluid Film and a Multi-Component Droplet

You can model impingement and stripping for multi-component fluid films in STAR-CCM+. Map every component in the fluid film continuum to its corresponding component in the Lagrangian phase for multi-component simulations:

1. Select the Multiphase Interaction > Phase Interactions > Phase Interaction Name > Models > Component Mapping node.



2. In the *Properties* windows, click the ellipsis of the Connectivity property.

3. In the *Component Mapping Connectivity* dialog that appears, match the fluid film components to their corresponding Lagrangian components using the drop-down menus provided.



4. Click OK.

If you have not mapped all components correctly, the following warning dialog appears during runtime.



Modeling Impingement of a Eulerian Phase to a Fluid Film

You can use a Film-Eulerian Phase Interaction to model the impingement of a dispersed Eulerian phase on a surface to form a fluid film.

The Film-Eulerian Phase Interaction model supports mass transfer from a dispersed Eulerian phase into the fluid film domain. However, the reverse (film stripping, converting fluid film mass to the dispersed Eulerian phase) is not supported.

This model supports single-component phases only: the transfer of a multi-component Eulerian phase to fluid film is not supported. The same material properties must be specified for both the fluid film phase and the dispersed Eulerian phase.

You can specify the impingement efficiency, that is, the fraction of the total mass impinging on the fluid film boundary that is transferred from the Eulerian phase into the fluid film.

For more information on the Inertial Impingement model, see the Inertial Impingement Model Formulation.

To model mass transfer due to inertial impingement on a surface:

1. Select the Film-Eulerian Phase Interaction model.
2. Select the Inertial Impingement model.
3. Specify the Fluid-Film Impingement Efficiency.

Selecting the Film-Eulerian Phase Interaction Model

Before you can model mass transfer due to inertial impingement, select the Film-Eulerian Phase Interaction model:

1. Create a phase interaction.
2. In the Multiphase Interaction > Phase Interactions > [Phase Interaction] node, right-click Models and choose *Select Models...*
3. In the *Phase Interaction Model Selection* dialog, select Film-Eulerian Phase Interaction from the *Phase Interaction Topology* group box.



4. Click Close.

The Film-Eulerian Phase Interaction node appears in the simulation tree.



5. In the Properties window, set Fluid Film Phase to the fluid film continuum and Eulerian Phase to the Eulerian phase.
6. If modeling a multi-component fluid film and a multi-component droplet, map the components between the two phases.

Film-Eulerian Phase Interaction Properties

Fluid Film Phase	The fluid film continuum interacting with the Eulerian phase.
Eulerian Phase	The Eulerian phase interacting with the fluid film phase.

Selecting the Inertial Impingement Model

The Inertial Impingement model can be selected from the Film-Eulerian Phase Interaction when the Film-Eulerian Phase Interaction model is selected:

1. In the Multiphase Interaction > Phase Interactions > [Phase Interaction] node, right-click Models and choose *Select Models...*
2. In the *Phase Interaction Model Selection* dialog, select Inertial Impingement from the *Optional Models* group box.



3. Click Close.

The Inertial Impingement node appears in the simulation tree.



The Inertial Impingement model has no properties to set.

Setting the Fluid-Film Impingement Efficiency

When you select the Inertial Impingement model in a Film-Eulerian phase interaction, you specify the impingement efficiency of the liquid Eulerian phase at the fluid-film interface.



The impingement efficiency is the fraction of the total mass impinging on the fluid film boundary that is transferred from the Eulerian phase into the Fluid Film domain. The default is a constant value of 1.0.

You can use this setting to model different impingement conditions without changing the physics setup.

In the *Properties* window, you set the method for defining the impingement efficiency.

Method	Specifies the method for defining the impingement efficiency.	
	Constant	Specify a constant value for the impingement efficiency.
	Field Function	Specify a field function that represents the impingement efficiency.

Modeling Fluid Film Evaporation and Condensation

STAR-CCM+ provides an evaporation and condensation model for fluid films. This model is a Film-Eulerian Phase Interaction model between a multi-component liquid film and a multi-component gas. Other than being a Film-Eulerian Phase Interaction model, the fluid film evaporation, and condensation model is similar to the VOF evaporation and condensation model.

This section describes the following:

- Setting up the fluid film evaporation and condensation model
- Defining a Film-Eulerian Phase Interaction model
- Selecting the fluid film evaporation and condensation model
- Properties of the fluid film evaporation and condensation model
- Defining evaporation and condensation mixture components

The mathematical formulation and field function references for the fluid film evaporation and condensation model are also available.

Setting Up the Fluid Film Evaporation and Condensation Model

To model fluid film evaporation and condensation in STAR-CCM+:

1. Select the physics models that are required for a fluid film simulation. Include these models:
   - Select Multi-Component Gas material model in the physics continuum
   - Select the Multiphase Interaction model in the physics continuum.

- If you want to model the effect of evaporation and condensation on heat balance, select either Coupled Energy or one of the Segregated Fluid Energy models.

2. Set up the fluid film phase and choose Multi-Component Liquid as the material model.
3. Define a Film-Eulerian Phase Interaction.
4. Select the fluid film evaporation and condensation model.
5. Define the evaporation and condensation mixture components.

Selecting the Film-Eulerian Phase Interaction Model

Select the Film-Eulerian Phase Interaction model before modeling fluid film boiling or evaporation and condensation:

1. Create a phase interaction.
2. Select Multiphase Interaction > Phase Interactions > [Phase Interaction] > Models and choose *Select Models...*
3. In the *Phase Interaction Model Selection* dialog, select Film-Eulerian Phase Interaction from the *Phase Interaction Topology* group-box.



4. Click Close.

The Film-Eulerian Phase Interaction node appears in the simulation tree.



5. In the Properties window, set the Fluid Film Phase property to the fluid film continuum. The Continuous Phase property is automatically set to the physics continuum.

Film-Eulerian Phase Interaction Properties

Fluid Film Phase	The fluid film continuum interacting with the Eulerian phase.
Continuous Phase	The physics continuum for the Eulerian phase. (Read-only)

Selecting the Fluid Film Evaporation and Condensation Model

Once you have set up a Film-Eulerian Phase Interaction, select the Evaporation and Condensation model:

1. Right-click Multiphase Interaction > Phase Interactions > [Phase Interaction] > Models and choose Select Models...
2. In the *Phase Interaction Model Selection* dialog, select Component Mapping from the *Optional Models* group box.

   The Boiling and Evaporation/Condensation models become available in the *Optional Models* group box.
3. Select Evaporation/Condensation from the *Optional Models* group box.



4. Click Close.

Fluid Film Evaporation and Condensation Model Properties

The Evaporation/Condensation node appears in the simulation tree to indicate that the fluid film evaporation and condensation model has been selected.



The default fluid film evaporation and condensation model is implemented as a diffusion limited model. For that reason, it is not directly applicable to regimes where energy transport dominates the phase exchange. For example, during condensation of pure vapor or when evaporation takes place close to boiling point. In such saturated regimes, you can switch to using the thermally limited evaporation and condensation model.

For more information on the evaporation and condensation model implementation, see the formulation section.

Evaporation/Condensation Properties

Implicit Coupling

Switches between the default hydrodynamically limited evaporation and condensation model and the thermally limited evaporation and condensation model.
For more information, see Evaporation and Condensation Solution Procedure in the formulation section.

Activated

Activates the thermally limited implementation of the evaporation and condensation model.
This model is used for saturated regimes, for example, during condensation of pure vapor or when evaporation takes place close to boiling point.

Deactivated

STAR-CCM+ uses the hydrodynamically limited evaporation and condensation model by default. This model is used for non-saturated regimes.

Under_Relaxation Factor

At each iteration, this property governs the extent to which the newly computed solution supplants the old solution.

Evaporation/Condensation Expert Properties

Saturation Threshold

The saturation ratio that is used in Eqn. (1562) to determine whether saturation occurs.
This setting lets you adjust when the fluid film evaporation and condensation model switches into saturation mode.

N_seeds

The nucleation density for condensation droplets.
This value is N in Eqn. (1566).
For more information, see Steam Condensation at a Dry Wall in the formulation section.

R_seeds

The minimum radius of condensation droplets.
This value is R_{min} in Eqn. (1566).
For more information, see Steam Condensation at a Dry Wall in the formulation section.

Linearize Film Energy — Linearizes the effect of phase exchange in the film energy equation.

- Activated — Activates the expert property.
- Deactivated — Default option.

Linearize Gas Energy — Linearizes the effect of phase exchange in the gas energy equation.

- Activated — Activates the expert property.
- Deactivated — Default option.

Condensation — Includes condensation in the phase interaction.

- Activated — Default option.
- Deactivated — Deactivates the expert property.

Evaporation — Includes evaporation in the phase interaction.

- Activated — Default option.
- Deactivated — Deactivates the expert property.

Setting the Liquid Film and Vapor Phase Mixture Components

This task applies to both the Boiling and the Evaporation/Condensation models.

Set the liquid film and vapor phase component pairs using the Component Mapping model node:

1. Select the Component Mapping model node.

2. Edit the Connectivity property and match the liquid film and vapor phase components using the drop-down menus provided.



Modeling Fluid Film Boiling

The multi-component liquid film boiling model models vaporization above boiling temperature. This model assumes that boiling takes place either on hot film walls or inside film cells. Boiling at hot walls increases the heat flux between the film and the wall. The mechanisms causing this effect are:

- Evaporation (latent heat) transport

  Heat is supplied near the wall to bubbles which then move away into the liquid film. In subcooled boiling, simultaneous evaporation at the base of a bubble and condensation at its tip can occur.

- Micro-convection

  Growth and collapse of bubbles cause random liquid motion.

- Vapor-liquid exchange

  A Reynolds analogy model in which bubble growth causes an exchange of liquid between the wall and bulk regions.

- Surface quenching

  A variation of the vapor-liquid exchange, assuming transient conduction to the cold liquid in contact with the wall after bubble departure.

- Wake flow

Liquid motion behind a departing bubble causes convection from the wall.

- Thermo-capillary flow

  Small variations in surface tension due to temperature differences between the base and tip of a bubble cause a jet of hot liquid to flow away from the wall.

Note: Boiling takes place at the film wall. However, no boiling occurs at edges.

Caution: Film boiling requires a conformal mesh at film-solid interfaces.

The mathematical formulation and field function references for the fluid film boiling model are also available.

To activate the film boiling model:

1. Set up a Film-Eulerian Phase Interaction.
2. Right-click Multiphase Interaction > Phase Interactions > [Phase Interaction] > Models and choose Select Models...
3. In the *Phase Interaction Model Selection* dialog, select Component Mapping from the *Optional Models* group box.

   The Boiling and Evaporation/Condensation models become available in the *Optional Models* group box.
4. Select Boiling from the *Optional Models* group box.



The Boiling node appears in the simulation tree to indicate that the multicomponent film boiling model has been selected.



5. Set the boiling and HTCxArea properties.
6. Define the boiling mixture components.

Boiling Properties

Boiling Model Parameter c_sf	Characterizes the fluid wall interface. c_{sf} in Eqn. (1570).
n	The Prandtl Number Exponent. n in Eqn. (1570).
Leidenfrost Temperature	The temperature $\Delta T_L = T_L - T_{sat}$ at which the Leidenfrost effect occurs. See Eqn. (1576).
Boiling Model Parameter c_s	Parameter that is used to compute the wall heat flux. c_s in Eqn. (1576).
Constant of Critical Heat Flux	c_{max} in Eqn. (1571).
Constant of Minimal Heat Flux	c_{min} in Eqn. (1572).
Wall Boiling Evaporation Factor c_e	The wall boiling evaporation c_e factor that determines how much of the boiling heat flux generates vapor. See Eqn. (1578).

Troubleshooting Fluid Film Boiling

Error: FilmEvaporationModel: Detect negative latent heat of boiling

The latent heat is the difference between the total enthalpy of the liquid (film) and gas phase. When this error message appears, you can evaluate the enthalpy of the film and gas phase with a contour plot to check for errors.

In addition, this error message can occur when you have incorrectly defined the Heat of Formation or Specific Heat material properties for the film or gas phases.

Modeling Fluid Film Melting and Solidification

STAR-CCM+ provides a melting and solidification model for fluid films. The fluid film melting and solidification model is similar to the VOF melting and solidification model.

This section describes the following:

- Selecting the Fluid Film Melting and Solidification Model
- Setting Fluid Film Melting-Solidification Model Parameters
- Setting up Solidified Mass Removal

The mathematical formulation and field function references for the fluid film melting and solidification model are also available.

Selecting the Fluid Film Melting and Solidification Model

Before you can select the fluid film melting and solidification model, you activate the appropriate continuum models. You then create the liquid film phase and select the appropriate models.

Activating the Continuum Models

To activate the continuum models:

1. In the continuum, open the *Physics Model Selection* dialog.
2. Select Three Dimensional from the *Space* group box.
3. A fluid film simulation is a transient analysis, so select Implicit Unsteady from the *Time* group box.
4. Specify the material model as Gas or Multi-Component Gas.
5. Select Non-reacting from the *Reaction Regime* group box.
6. Select the appropriate option from the *Flow* group box:
   - Coupled Flow
   - Segregated Flow
7. Select Fluid Film from the *Optional Models* group box.
8. Select the appropriate energy option from the *Optional Models* group box:
   - For a Coupled Flow model, select Coupled Energy.
   - For a Segregated Flow model, select either Segregated Fluid Temperature, Segregated Fluid Isothermal, or Segregated Fluid Enthalpy.
9. Select the models that are appropriate for the continuum from the following group boxes:

- *Equation of State*
- *Viscous Regime*

10. Click Close.

Creating the Liquid Film Phase

You create the liquid film phase and select the appropriate models. The liquid film can be a single-component liquid or a multi-component liquid.

To create the liquid film phase:

1. In the Models > Fluid Film > Fluid Film Phases node, create a new phase.
2. In the new phase, open the *Phase Model Selection* dialog.
3. From the *Material* group box, select Liquid or Multi-Component Liquid.
4. In the *Equation of State* group box, select the model that is appropriate for the liquid film phase in your simulation.

   The fluid film melting and solidification model is available for selection.
5. Select Melting-Solidification from the *Optional Models* group box.



6. Click Close.

Setting Fluid Film Melting-Solidification Model Parameters

The Melting-Solidification node represents the fluid film melting and solidification model in the simulation tree.



- To set the appropriate properties, such as the under-relaxation factor for melting and solidification, select the Melting-Solidification node.

Melting-Solidification Properties

Solidified Film Option	Specifies how the solidified mass is treated in the simulation. For more information, see Setting up Solidified Mass Removal. None: The fluid film is treated as a homogeneous mixture of solidified mass and liquid. For more information, see Melting-Solidification Enthalpy Formulation. Remove Solidified Film: The solidified mass is removed. All of the mass that remains in the fluid film is liquid. For more information, see Removal of Solidified Mass.

Melting-Solidification Expert Properties

Smoothing Steps	If the Morpher is enabled, this property affects the smoothing of the morphed surface that is used to model the surface of the solidified mass. At each iteration, the smoothing process is repeated the specified number of times.
Solid Time-step Factor	If the Morpher is enabled, this property scales the morphing speed of the morphed surface that is used to model the surface of the solidified mass. In practical solidification simulations, the morphing of the solidified mass front is typically so small that it does not noticeably influence the solution at each time step. This setting lets you speed up the calculation by the specified factor. For example, by setting this factor to 10, and reducing the Maximum Physical Time by the same factor, you can obtain (to good approximation) the same solidified mass shape, but approximately 10 times faster.
Linearize Melting	Linearizes the influence of latent heat. Activated: Activates the expert property. Deactivated: Default option.
Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution.

Additional material properties are added when the Melting-Solidification

model is selected.

The Melting-Solidification model requires the fraction solid curve, the solidus, and liquidus temperatures and the latent heat of fusion.

For a single-component fluid film model, you set these properties as the material properties of the liquid.



For a multi-component fluid film model, you set these properties as the mixture properties of the multi-component liquid.



It is good practice to use a field function for the viscosity and thermal conductivity of the film, since the values of these properties differ for liquid and solid states. Make these properties a function of either relative solid volume fraction or temperature.

Solidified Mass Removal

A major application of the Melting-Solidification model for fluid films is airfoil icing. An important aspect of airfoil icing is run-back: the flow of liquid water over the layer of recently solidified ice.

You can configure the Melting-Solidification model so that the solidified mass is removed from the fluid film. All of the mass that remains in the fluid film is liquid.

For more information, see Setting up Solidified Mass Removal.

In airfoil icing simulations, it is important to represent the geometry of the airfoil as a combination of the airfoil and the ice layer, as the geometry has a significant effect on the flow field. Accurately modeling the flow field is essential for predicting aerodynamic characteristics such as lift and drag, and also for modeling the self-reinforcing impact on the ice accretion rate itself.

To ensure that the geometry is accurately represented, you can use the Morpher to adjust the surface of the solidified volume when the solidified mass is removed.

If the Morpher is activated and configured, the wall morphs to model the surface of the solidified volume. You can specify the density of the solid; STAR-CCM+ uses the specified density to compute the morphing that is applied to the surface of the solidified mass.

If the Morpher is not activated, or is not correctly configured, the solidified mass is removed with no effect on the geometry of the wall.

Setting up Solidified Mass Removal

By default, the Melting-Solidification model treats the fluid film as a homogeneous mixture of solidified mass and liquid. A Relative Solid Volume Fraction is used to represent the amount of solid.

You can configure the Melting-Solidification model so that the solidified mass is removed from the fluid film. All of the mass that remains in the fluid film is liquid.

For more information, see Solidified Mass Removal.

To set up solidified mass removal:

1. In the *Melting-Solidification Properties* window, set the Solidified Film Option to Remove Solidified Film.

   The Solid Density node appears under the Melting-Solidification node. If you are using the Morpher to model the surface of the solidified mass, you can specify the density of the solid.

2. Select the Solid Density node and select the Method to use for specifying the density of the solid.

   For more information, see Solid Density Parameters.

3. Specify the appropriate constant value, or select the appropriate field function.

To activate and configure the Morpher:

1. Activate the Morpher.

   To set up a Morphing motion, click Tools > Motions > New > Morphing.

2. Apply the Morphing motion to the fluid region.

   Under the Regions > [Fluid] > Physics Values > Motion Specification node, set the Motion property to Morphing.

   For more information on setting up a morphing motion, see Working with Morphing.

3. At the interface between the fluid film and the gas (on the gas side), configure the Morpher to allow STAR-CCM+ internal models to add to the displacement.

   Under the Regions > [Fluid] > Boundaries > [Boundary Name] > Physics Conditions > Morpher node, set the Method property to Internal.

4. In the *Melting-Solidification Properties* window, set the Smoothing Steps property to the appropriate value.

   This property affects the smoothing of the morphed surface that is used to model the surface of the solidified mass. At each iteration, the smoothing process is repeated the specified number of times.

Solid Density Parameters

The Solid Density node lets you set the Method to use for specifying the density of the solid.

Solid Density Properties

Method	Liquid	The density of the solid is identical to the density of the liquid. This is the default setting.
	Constant	The density of the solid is specified as a constant. The default value is 916.7 kg/m^3.
	Field Function	The density of the solid is specified as a field function.

Setting Fluid Film Solver Parameters

The Fluid Film node represents the Fluid Film solver in the object tree. This node is subject to common menu items.



The Fluid Film solver node contains sub-nodes (instead of properties) that let you customize solver settings for the following:

- Segregated Flow (as it applies to the fluid film)
- Film Velocity
- Film Thickness

The fluid film solvers generally use the same techniques as the Segregated Flow solvers. Differences arise from the fact that these solvers work with film equations on surfaces, and film control volumes depend on film distribution.

Setting Segregated Flow Solver Parameters for Fluid Film

The Fluid Film Segregated Flow solver is represented in the object tree by the Segregated Flow sub-node of the Fluid Film solver. You can use its properties to adjust Segregated Flow solver parameters specifically for the fluid film calculations.



Fluid Film Segregated Flow Solver Properties

Temporary Storage Retained	Retains or removes data from temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations. Activated — Retain temporary storage. Deactivated — Do not retain temporary storage.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. In general, to freeze flow or turbulence, activate Temporary Storage Retained for both flow and turbulence for one iteration after freezing any solver. Activated — The solver does not update the solution. Deactivated — The solver updates the solution.

Setting Film Velocity Solver Parameters

The Film Velocity node represents the Film Velocity solver in the object tree. This node has properties and is subject to common menu items.



This solver uses a form of the Segregated Flow Velocity solver for the fluid film. It solves the momentum conservation equation, Eqn. (1510).

Film Velocity Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed velocity supplants the old one. The default value is 0.7.

Setting Film Thickness Solver Parameters

The Film Thickness node represents the Film Thickness solver in the object tree. This node has properties and is subject to common menu items.



This solver calculates the mass conservation equation (Eqn. (1507)) which is expressed in terms of film thickness corrections. The Film Thickness solver is equivalent to the Segregated Flow Pressure solver which calculates the mass conservation equation that is expressed in terms of pressure corrections.

Film Thickness Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed film thickness field supplants the old one. The default value is 0.7.

Post-Processing Fluid Film Results

You can show the movement of droplets (particles) in a scalar scene for particle impingement and/or film stripping. The example in the following screenshot shows particles being stripped from one wall in a channel and landing on a lower wall, forming a film.



This type of display involves making selections as follows:

- The Lagrangian phase as the part in the scalar displayer



- Particle Mass as the scalar value to display



Fluid Film Formulation

The fluid film model accounts for transport of conserved quantities within the film and its interaction with surroundings. STAR-CCM+ uses calculations that include:

- Mass Conservation
- Momentum Conservation
- Grid Flux
- Energy Conservation
- Species Conversation

- Volume Fraction
- Wave Stripping
- Edge Stripping
- Evaporation and Condensation
- Boiling
- Melting and Solidification

This section also includes nomenclature and a bibliography.

Major assumptions that are made in the basic model formulation are:

- The film is thin enough for laminar boundary layer approximation to apply.
- The film stays attached to the boundary unless an internal or separate model predicts film separation.
- There is a parabolic velocity profile across the film.
- The analysis is transient.

Using these assumptions, the film conservation equations that are given in the following sections are integrated and discretized for each film cell.

If wall velocity is $\mathbf{v}_w$, average film velocity is $\bar{\mathbf{v}}_f$, shear force at the film free surface is τ_{fs}, and ξ is the local coordinate normal to the wall, the assumed velocity profile across the film is:

$$\mathbf{v}(\xi) = \mathbf{v}_w + \left(3\frac{\bar{\mathbf{v}}_f - \mathbf{v}_w}{h_f} - \frac{\tau_{fs}}{2\mu}\right)\xi - \left(3\frac{\bar{\mathbf{v}}_f - \mathbf{v}_w}{2h_f^2} - \frac{3\tau_{fs}}{4\mu h_f}\right)\xi^2 \tag{1506}$$

Numerical Implementation

The film solution domain is made of shell regions. The mesh of shell regions is created using the boundary faces of the volume mesh. A film cell is made from two-dimensional boundary faces. The edge length and the film thickness define the area of an edge face. The volume of the film cell is defined with the boundary area and the film thickness.

The film governing equations are integrated over the volume of fluid film in each film cell to obtain a set of algebraic equations.

The discretized equations of mass and momentum are solved in a segregated manner. The energy equation can either be solved segregated or as part of the coupled energy equation.

Mass Conservation Equation

The equation that governs the film mass conservation is as follows:

$$\frac{\partial}{\partial t}\int_V \rho_f \, dV + \int_A \rho_f (\mathbf{v}_f - \mathbf{v}_g) \cdot d\mathbf{a} = \int_V \frac{s_m}{h_f} dV \tag{1507}$$

This equation is used to calculate the film thickness h_f. It is specific in the sense that the volume V and the surface A are functions of the film thickness and its spatial distribution.

Here, ρ_f is the film density, $\mathbf{v}_f$ is the film velocity, $\mathbf{v}_g$ is the grid velocity, and the subscript f denotes the fluid film values.

See also Grid Flux.

The quantity s_m is the mass source/sink per unit area. Contributions to the mass source or sink come from droplet impingement, film stripping, mass transfer, or user-defined sources.

Grid Flux

The grid flux at a face, as used in Eqn. (2538) to compute the face mass flux, is calculated as follows.

For a moving reference frame:

$$G_f = \mathbf{v}_g \cdot \mathbf{a}_f \tag{1508}$$

where $\mathbf{v}_g$ is the grid velocity and $\mathbf{a}$ is the face area.

For a moving mesh:

$$G_f = \frac{\mathbf{a}_s \cdot \boldsymbol{h}}{\Delta t} \tag{1509}$$

where $\mathbf{a}_s$ is the area that is swept by the movement of the two vertices that make up the edge of the shell face, $\boldsymbol{h}$ is a vector with the magnitude of the film thickness in the wall surface-normal direction, and Δt is the time step.

Momentum Conservation Equation

The momentum conservation equation is:

$$\frac{\partial}{\partial t}\int_V \rho_f \mathbf{v}_f \, dV + \int_A \rho_f \mathbf{v}_f (\mathbf{v}_f - \mathbf{v}_g) \cdot d\mathbf{a} = \int_A \mathbf{T}_f \cdot d\mathbf{a} - \int_A p_f \, d\mathbf{a} + \int_V \left(\rho_f \mathbf{f}_b + \frac{\mathbf{s}_m}{h_f}\right) dV \tag{1510}$$

where $\mathbf{v}_g$ is the grid velocity, $\mathbf{s}_m$ is the momentum source corresponding to the mass source s_m , p_f is the pressure, $\mathbf{f}_b$ is the body force (for example, gravity, or the pseudo-force in a moving reference frame), and $\mathbf{T}_f$ is the viscous stress tensor within the film. The kinematic and dynamic conditions at the interface between the film and the surrounding fluid (film free surface) are satisfied:

$$(\mathbf{v}_f)_{\text{int}} = (\mathbf{v})_{\text{int}} \tag{1511}$$

$$(\mathbf{T}_f \cdot d\mathbf{a} + p_f\, d\mathbf{a})_{\text{int}} = (\mathbf{T} \cdot d\mathbf{a} + p d\mathbf{a})_{\text{int}} \tag{1512}$$

The quantities without the subscript f are associated with the surrounding fluid. Assuming that the normal components of the viscous and convective terms are negligible, the pressure distribution within the fluid film is obtained from Eqn. (1510) as:

$$p_f(\xi) = p_{\text{int}} - \mathbf{s}_m \cdot \mathbf{n} - \rho_f \mathbf{f}_b \cdot \mathbf{n}(h_f - \xi) + \int_{\xi}^{h_f} \frac{\partial}{\partial t}(\rho_f \mathbf{v}_f \cdot \mathbf{n}) d\xi \tag{1513}$$

where $\mathbf{n}$ is the wall surface unit vector pointing towards the film, and ξ is the local coordinate normal to the wall. Eqn. (1513) assumes that the force $\mathbf{s}_m$ is applied at the film free surface.

Energy Conservation Equation

The energy conservation equation is:

$$\frac{\partial}{\partial t}\int_V \rho_f E_f dV + \int_A [\rho_f H_f(\mathbf{v}_f - \mathbf{v}_g) + \mathbf{v}_g p_f] \cdot d\mathbf{a} = \int_A \mathbf{q}''_f \cdot d\mathbf{a} + \int_A \mathbf{T}_f \cdot \mathbf{v}_f d\mathbf{a} + \int_V \mathbf{f}_b \cdot \mathbf{v}_f dV + \int_V \frac{s_e}{h_f} dV \tag{1514}$$

where E_f is the film total energy, H_f is the film total enthalpy, $\mathbf{q}''_f$ is the film heat flux, $\mathbf{T}_f$ is the fluid viscous stress tensor, and $\mathbf{f}_b$ is the body force. s_e is the energy source/sink term per unit film area. Contributions to the energy source or sink come from droplet impingement, film stripping, or user-defined sources.

Species Conservation Equation

The species mass conservation equation of species i is:

$$\frac{\partial}{\partial t}\int_V \rho_f Y_{i,f} dV + \int_A \rho_f (\mathbf{v}_f - \mathbf{v}_g) Y_{i,f} \cdot d\mathbf{a} = \int_A \left(\frac{\mu}{\sigma}\right) \nabla Y_{i,f} \cdot d\mathbf{a} + \int_V \frac{s_{m,i}}{h_f} dV \tag{1515}$$

where $Y_{i,f}$ is the mass fraction of species i, and σ is the molecular Schmidt number. The quantity $s_{m,i}$ is the mass source or since for species i.

Volume Fraction

The volume fraction for the film, in the neighboring gas cell, is computed as:

$$\phi_{film} = min(\frac{V_{film}}{V_{gas}}, \phi_{film,max}) \tag{1516}$$

where V_{film} is the film volume (the area multiplied by the film height) and V_{gas} is the volume of the neighboring cell.

You can set $\phi_{film,max}$ in the Maximum Volume Fraction property of the fluid film phase.

Wave Stripping

Depending on flow conditions and the geometry of the wall on which the film resides, the film can form droplets and this process is called film stripping. Two stripping mechanisms are modeled:

- Stripping due to wave instability generated by an adjacent fluid flow
- Stripping due to body-force induced instability (gravity, film acceleration)

Adjacent Flow and Body-Force Induced Instability

Due to external forces such as gravity or shearing at the interface between a fluid film and its surrounding fluid, the film surface can develop instabilities and start shedding droplets. The model described here is based on the original documentation of Hervé Foucart [278]. It postulates that stripping at the surface occurs in a three-stage process, as represented schematically in the following illustration of the liquid film stripping mechanism.



First, waves develop at the liquid-gas interface. As these waves grow, they become unstable and form cylinders of liquid which are ejected into the surrounding gas. The cylinders then break up further forming spherical droplets. The model is closed by computing the radius of the child droplets as a function of the most unstable wavelength characterizing the surface instability.

It is accepted that for surface stripping:

- Ejection occurs when critical conditions are satisfied, that is a minimum height is required for droplet ejection.
- The ejected fluid volume forms a cylinder, as a result of the Kelvin-Helmholtz instability.
- The cylinder then breaks into droplets due to Rayleigh-Taylor instabilities.

Having characterized surface stripping as a function of wavelength, the stripping model relationships are now closed by computing the wavelength for which surface break-up is most likely to occur. This wavelength is the most unstable, and is calculated as the resonance wavelength using the dispersion equation. Solving the dispersion equation for surface waves, the resonance wavelength is obtained as:

$$\lambda_r = \frac{2\pi}{\rho_f \mathbf{f}_b \cdot \mathbf{n}}\left(\frac{1}{3}\rho\tilde{\mathbf{v}}_r^2 - \sqrt{\left(\frac{1}{3}\rho\tilde{\mathbf{v}}_r^2\right)^2 - \rho_f \mathbf{f}_b \cdot \mathbf{n}\sigma}\right) \tag{1517}$$

where $\tilde{\mathbf{v}}_r^2$ is the square of relative velocity between the film and surrounding fluid $(\mathbf{v}_f - \mathbf{v}) \cdot (\mathbf{v}_f - \mathbf{v})$, σ is the surface tension, and $\mathbf{f}_b$ is the body force acting on film which also includes the inertial force. The minimum height necessary for droplet ejection is then:

$$h_{\min} = c_H \frac{\lambda_r}{2\pi} \tag{1518}$$

The film height to strip, that is the surface wave amplitude, is:

$$h_a = \left[\frac{3}{4}\left(\frac{2}{3.78}\right)^3\right]^2 \pi\lambda_r \tag{1519}$$

The resulting droplet diameter is:

$$D_d = c_D \sqrt{\frac{\lambda_r h_a}{\pi}} \tag{1520}$$

The default setting for c_D is 3.78. The number 3.78 in Eqn. (1519) arises from Rayleigh theory. It means that the diameter of each droplet that is generated from the cylinder is 3.78 times the diameter of that cylinder.

All droplets initially have the same velocity as the film and are placed at a point between the cell center and the boundary face center, depending on the radius of the ejected droplet.

Although no consideration of time scale is included in the original stripping model, the time scale can be defined in terms of a linear growth rate as:

$$t_b = \frac{\lambda_r(\rho_f + \rho)}{|\mathbf{v}_r|\sqrt{(1-\mathrm{We})\rho_f\rho}} \tag{1521}$$

where We is the Weber number and the most unstable mode occurs when $\mathrm{We} = 2/3$.

Using this time scale, stripping is expected to occur n_b times during a time-step Δt:

$$n_b = c_B \frac{\Delta t}{t_b} \tag{1522}$$

Free-Stream Velocity

The Wave Stripping model requires a measure of the free-stream velocity ($\mathbf{v}_r$ as used in Eqn. (1517)).

STAR-CCM+ calculates two measures of the free-stream velocity and uses the highest.

- One measure is the gas velocity in the nearest gas cell. The disadvantage with this measure is that the value depends on the mesh size, and the approximation becomes less valid as the mesh size decreases.
- A second measure is a mesh-independent estimate of the free-stream velocity based on the shear velocity. This second measure is derived below.

At the free surface, on the gas side, the shear (or frictional) velocity is defined as:

$$\mathbf{u}_* = -\sqrt{\frac{\tau}{\rho}}\,\mathbf{n}_\tau \tag{1523}$$

where $\mathbf{n}_\tau$ is the unit vector in the direction of the shear stress τ.

It is assumed that the free-stream gas velocity can be reasonably approximated as:

$$\mathbf{v}_r \approx \mathbf{v}_{l,g} + K_{*,\infty}\mathbf{u}_* = \mathbf{u}_{*,\infty} \tag{1524}$$

where $\mathbf{v}_{l,g}$ is the interfacial velocity and $K_{*,\infty}$ is a constant that relates the shear velocity to the free-stream velocity. This constant is specified as the Gas Velocity Scaling Factor in the Wave Stripping model properties.

Edge Stripping

Maroteaux et al. [280] proposed a model for the break-up of liquid films over a sharp edge. This model is based on the Rayleigh-Taylor instability and contains the break-up criterion and resulting droplet diameter distribution. However, Gubaidullin [281] has raised questions about the validity of the breakup criterion. Therefore, an alternative criterion is adopted, proposed by Friedrich et al. [282], that is based on a force balance. The droplet diameter distribution still follows Maroteaux.



The model is only applied at edges where the corner angle (θ in the diagram above) is greater than the user-defined minimum corner angle, θ_{min}. If a boundary cell has multiple sharp edges, stripping occurs over all of them in a weighted manner.

The force ratio FR, which is essentially the ratio of the film momentum flux to the surface tension and gravity force, is:

$$FR = \frac{We_f}{1 + \frac{1}{\sin\theta} + Bo_f \frac{L_b}{h_f \sin\theta}} \tag{1525}$$

where We_f is the film Weber number (see Eqn. (1526)), Bo_f is the Bond number (see Eqn. (1527)), L_b is the break-up length (see Eqn. (1528)) and h_f is the film thickness.

The film Weber number is:

$$We_f = \frac{\rho_f u_f^2 h_f}{\sigma} \tag{1526}$$

where ρ_f is film density, u_f is film velocity projected orthogonal to the stripping edge, h_f is film thickness and σ is film surface tension.

The film Bond number is:

$$Bo_f = \frac{\rho_f g_\theta h_f^2}{\sigma} \tag{1527}$$

where g_θ is the component of acceleration normal to the downstream wall.

Arai and Hashimoto [283] give the break-up length as:

$$L_b = 0.0388 h_f^{0.5} Re_f^{0.6} We_{rel}^{-0.5} \tag{1528}$$

where Re_f is the film Reynolds number (see Eqn. (1529)) and We_{rel} is the relative Weber number (see Eqn. (1530)).

The film Reynolds number is:

$$Re_f = \frac{\rho_f u_f h_f}{\mu_f} \tag{1529}$$

where μ_f is the film viscosity.

The relative Weber number is:

$$We_{rel} = \frac{h_f \rho_g (u_g - u_f)^2}{2\sigma} \tag{1530}$$

where ρ_g is the gas density and u_g is the gas velocity projected orthogonally to the stripping edge.

The break-up is deemed to occur when $FR > FR_C$, where FR_C is the user-defined critical force ratio, which is proposed in [282]. The critical force ratio default value is 1. Of the fluid crossing the edge, only a fraction x_S separates from the film. This fraction is approximated using the formula below:

$$x_S = \begin{cases} 0 & FR \le FR_C \\ 0.44(FR - FR_C) & FR_c < FR \le (FR_C + 1.6) \\ 0.057(FR - FR_C - 1.6) + 0.704 & (FR_C + 1.6) < FR \le (FR_C + 6.792) \\ 1 & (FR_C + 6.792) < FR \end{cases} \tag{1531}$$

which is based on the experimental data in [282].

For droplet formation, the model that is presented in [280] is adopted, which estimates the parent droplet diameter D_d as:

$$D_d = c_1\sqrt{\frac{\lambda_m h_f}{\pi}} \tag{1532}$$

where c_1 is a user-defined droplet diameter scale factor. This factor is set to 3.78 by default. The wave length λ_m is calculated from the most unstable wavenumber k_m, which maximizes the growth rate:

$$\omega = -\left(\frac{\sigma - (\rho_f - \rho_g)a/k^2}{2\mu_f h_f}\right)\left(\frac{(kh_f)\sinh(kh_f)\cosh(kh_f) - k^2h_f^2}{\cosh^2(kh_f) + k^2h_f^2}\right) \tag{1533}$$

where the acceleration a is computed from:

$$a = \frac{u_f^2\theta}{h_f(\pi + \theta)} \tag{1534}$$

and λ_m is related to k_m by $\lambda_m = 2\pi/k_m$. You can choose to provide an upper limit for the parent droplet diameter D_d by activating the Limit parent droplet diameter property and supplying a value for the Maximum diameter.

According to the model, the cumulative droplet size distribution $F(D)$ is a Rosin-Rammler distribution:

$$F(D) = 1 - e^{-(D/X)^q} \tag{1535}$$

where q is a user-defined parameter, which is set to 1.5 by default, and:

$$X = \frac{D_d}{(3\ln 10)^{1/q}} \tag{1536}$$

There are two options for determining the droplet diameter:

- Generate droplets with a deterministic diameter, namely the average of the size distribution that is:

  $$D = X\Gamma\left(1 + \frac{1}{q}\right) \tag{1537}$$

  where $\Gamma(x)$ represents the Gamma function.
- Generate droplets with a randomized diameter D according to Eqn. (1535).

You can select these options by either activating or deactivating the Randomize droplet diameter property.

The droplets are ejected at the film velocity speeds in terms of both magnitude and direction.

Edge Stripping with Flow from Both Sides

For a case where flow comes toward an edge from both sides, the flows are combined into one stream of Lagrangian droplets that conserves mass, momentum, species, and energy. For example, if both flows have a different temperature, the droplets get a temperature somewhere in between. For opposing flows, the stripping fraction x_S (Eqn. (1531)) is also set to 1 for both sides (which is full stripping). Droplet diameter is computed based on the film properties of the side with the biggest flux. The following diagram illustrates opposing-flow edge stripping.



Inertial Impingement

The Inertial Impingement model uses the inertia of the dispersed droplets that approach the wall to compute the mass flux of the droplets that impinge on the wall to form a fluid film. Impingement occurs on the walls at which a shell interface has been created. The mass flux due to impingement is:

$$m_i = \gamma \rho_i \mathbf{U}_i \cdot \mathbf{A} \alpha_i \tag{1538}$$

where:

- m_i is the mass flux leaving the system.
- γ is the impingement efficiency.
- ρ_i is the cell-centered density.
- $\mathbf{U}_i$ is the cell-centered velocity.
- $\mathbf{A}$ is the cell wall area vector (normal to the wall in the outward facing direction).
- α_i is the cell-centered volume fraction.

The impingement efficiency specifies the fraction of the flux incident on the surface that impinges on the fluid film. The impingement efficiency can be a value from 0 through 1.

The kinematic and dynamic conditions at the interface between the fluid film and the surrounding multiphase fluid (film free surface) are satisfied. The impinging mass flux that is calculated at the interface cell is used in all of the other transport equations.

Limitations

The Inertial Impingement model has the following limitations:

- Only a single-component liquid phase can impinge. Multi-component liquid phases are not supported.
- The liquid phase impinging on the fluid-film interface must have the exact same material properties as the fluid-film phase.
- The Eulerian multiphase domain does not experience any reduction in its domain due to the formation of liquid film.

Evaporation and Condensation

The species mass flux for every component is conserved at the interface between the gas and the fluid film, such that:

$$\rho_{g_s} Y_{g_s,i}(v_{g_s} - \dot{h}) - \rho_{g_s} D_{g,i} \frac{dY_i}{dy}\bigg|_{g_s} = \rho_{l_s} Y_{l_s,i}(v_{l_s} - \dot{h}) - \rho_{l_s} D_{l,i} \frac{dY_i}{dy}\bigg|_{l_s} \tag{1539}$$

where, evaluated at surface s:

- ρ_{g_s} and ρ_{l_s} are the gas and liquid film densities
- $Y_{g_s,i}$ and $Y_{l_s,i}$ are the mass fractions for the gas and liquid film
- v_{g_s} and v_{l_s} are the normal velocity components for the gas and liquid film
- $D_{g,i}$ and $D_{l,i}$ are the gas and liquid film molecular diffusion coefficients
- $\dot{h}$ is the rate of change of film thickness

The total mass flux is also conserved:

$$\rho_{g_s}(v_{g_s} - \dot{h}) = \rho_{l_s}(v_{l_s} - \dot{h}) \tag{1540}$$

Combining Eqn. (1539) and Eqn. (1540):

$$\rho_{l_s}(v_{l_s} - \dot{h})(Y_{g_s,i} - Y_{l_s,i}) + \rho_{l_s} D_{l,i} \frac{dY_i}{dy}\bigg|_{l_s} - \rho_{g_s} D_{g,i} \frac{dY_i}{dy}\bigg|_{g_s} = 0 \tag{1541}$$

The evaporation rate is defined as:

$$\dot{m}_v = -\rho_{l_s} \dot{h} \tag{1542}$$

with v_{l_s} considered to be zero from here on. The summation of Eqn. (1541) over all liquid film component N_L becomes:

$$\left[1 - \sum_j^{N_L} Y_{g_s,j}\right] \dot{m}_v = -\sum_j^{N_L} \rho_{g_s} D_{g,i} \frac{dY_i}{dy}\bigg|_{g_s} \tag{1543}$$

and since

$$-Y_{g_s,j} \dot{m}_v = -\rho_{g_s} D_{g,i} \frac{dY_i}{dy}\bigg|_{g_s} \tag{1544}$$

for all $N_{P,L}$ species i in the film that are inert, the summation can be limited to the N_V interacting components. The total evaporation rate can therefore be expressed as:

$$\dot{m}_v = -\frac{\sum_j^{N_v} \rho_{g_s} D_{g,i} \left.\frac{dY_i}{dy}\right|_{g_s}}{1 - \sum_j^{N_v} Y_{g_s,j}} \tag{1545}$$

Note that Eqn. (1545) is valid only below saturation conditions. That is:

$$\sum_j^{N_v} Y_{g_s,j} < 1$$

At saturation conditions, the algebra that is used to derive Eqn. (1545) is not correct.

The normal derivative is treated through the species transfer coefficients $s_{t,i}$ and the Spalding transfer number B such that:

$$\dot{m}_v = -\frac{\sum_j^{N_V} s_{t,j}(Y_{g_c,j} - Y_{g_s,j})}{1 - \sum_j^{N_V} Y_{g_s,j}} \cdot \frac{\ln(1+B)}{B} \tag{1546}$$

where the subscript c indicates a cell value and B is defined as:

$$B = \frac{\sum_j^{N_V} Y_{g_s,j} - \sum_j^{N_V} Y_{g_s,\infty}}{1 - \sum_j^{N_V} Y_{g_s,j}} \tag{1547}$$

where it is assumed that $Y_{g_s,\infty} \approx Y_{g_s,c}$

The interfacial gas mass fraction $Y_{g_s,j}$ is required to evaluate Eqn. (1546). $Y_{g_s,j}$ is related to $Y_{l_s,j}$ through Raoult's law, as follows:

Consider the vapor pressure $\tilde{p}_i$ of a component:

$$\tilde{p}_i = a_i p^*_i \tag{1548}$$

where p^*_i is the vapor pressure of the pure component i and a_i is the activity, which is dependent on the liquid mixture. This expression can be written as:

$$\tilde{p}_i = \gamma_i X_{l_s,i} p^*_i(T_s) \tag{1549}$$

Raoult's law states that $\gamma_i \approx 1$, an approximation that is followed here. Given the components vapor pressures, their molar fractions at the interface can be derived from:

$$X_{g_s,\,i} = \frac{\tilde{p}_i}{p} \tag{1550}$$

where p is the ambient pressure. The conversion to interface mass fraction is obtained using:

$$Y_{g_s,\,i} = \frac{X_{g_s,\,i} W_i}{\sum_j^{N_V} X_{g_s,\,j} W_j + \sum_j^{N_{P,\,G}} X_{g_s,\,j} W_j} \tag{1551}$$

While $X_{g_s,\,i}$ is known for the evaporation components, it is unknown for the inert species. A good approximation is the introduction of a background molar weight, which is calculated at a small distance away from the interface:

$$W_{bg} = \frac{\sum_j^{N_{P,\,G}} X_j W_j}{\sum_j^{N_{P,\,G}} X_j} \tag{1552}$$

The interfacial background molar fraction is:

$$X_{g_s,\,bg} = 1 - \sum_j^{N_V} X_{g_s,\,j} \tag{1553}$$

The interfacial vapor mass fraction can then be approximated as:

$$Y_{g_s,\,i} \approx \frac{X_{g_s,\,i} W_i}{X_{g_s,\,bg} W_{bg} + \sum_j^{N_v} X_{g_s,\,j} W_j} \tag{1554}$$

Eqn. (1549) assumes knowledge of $X_{l_s,\,i}$. Presently, the normal species profile in the film is approximated as a constant, such that the interfacial molar fractions follow from the associated cell center values:

$$X_{l_s,\,i} = X_{l_c,\,i} \tag{1555}$$

The component evaporation rate can be expressed as:

$$\dot{m}_{v,\,i} = Y_{g_s,\,i} \dot{m}_v - \rho_{g_s} D_{g,\,i} \left. \frac{dY_i}{dy} \right|_{g_s} \tag{1556}$$

which is valid under all conditions.

The interfacial heat flux conservation is expressed as:

$$k_g \frac{dT}{dy}\bigg|_{g_s} - k_l \frac{dT}{dy}\bigg|_{l_s} - \dot{Q}_v = 0 \tag{1557}$$

where k denotes thermal conductivity and:

$$\dot{Q}_v = \sum_j^{N_V} \Delta H_i^{vap} \dot{m}_{v,i} \tag{1558}$$

which is a condition for the interfacial surface temperature T_s.

Note that a combination of Eqn. (1556) and Eqn. (1558) leads to an expression for the total evaporation rate which is valid at all conditions:

$$\dot{m}_v = \frac{\dot{Q}_v + \sum_i^{N_V} \Delta H_i^{vap} - \rho_{g_s} D_{g,i} \frac{dY_i}{dy}\bigg|_{g_s}}{\sum_i^{N_V} \Delta H_i^{vap} Y_{g_s,i}} \tag{1559}$$

Evaporation and Condensation Solution Procedure

This section describes how STAR-CCM+ calculates evaporation and condensation in both non-saturated and saturated regimes.

Below Saturation

Below the saturation temperature (also referred to as boiling temperature), evaporation and condensation rates are controlled by the following:

- Hydrodynamic effects - vapor needs to be delivered to or carried away from the surface, which happens at a finite rate.
- Thermal effects - heat needs to be conducted towards or away from the interface to balance the latent heat of evaporation.

STAR-CCM+ attempts to satisfy Eqn. (1560), in which $\dot{Q}_v$ is given by Eqn. (1558), which in turn relies on Eqn. (1556), which requires Eqn. (1545). Despite the apparent complexity, this condition is simply finding the interfacial temperature T_s such that Eqn. (1560) is satisfied.

This problem can be approached in many ways. The present implementation treats it by iteratively looking for the root of the function:

$$f(T_s) = k_g \frac{dT}{dy}\bigg|_{g_s} - k_l \frac{dT}{dy}\bigg|_{l_s} - \dot{Q}_v(T_s) \tag{1560}$$

through the secant method.

STAR-CCM+ uses the following process:

- Start with two initial values, $T_{s,0}$ and $T_{s,1}$, and evaluate f_0 (shorthand for $f(T_{s,0})$) and f_1.
- Compute $T_{s,n} = T_{s,n-1} - f_{n-1}\left(\frac{T_{s,n-1} - T_{s,n-2}}{f_{n-1} - f_{n-2}}\right)$ for $n = 3..N_{it}$

At present, the first initial value is taken to be the actual interfacial temperature upon entering the optimizer, and the second initial value is taken 5 percent lower than the first. Note that the solution of T_s is limited from below by T_{sat}, at which point it is assumed that the exclusively thermally limited rate (essentially stating that vapor is infinitely available) is valid.

At Saturation

Saturation can be detected by either, or both of, the following conditions:

$$\sum_{i}^{N_V} Y_{g_s,j} = 1 \tag{1561}$$

$$\sum_{i}^{N_V} Y_{g_\infty,j} = 1 \tag{1562}$$

The first condition (Eqn. (1561)) is obvious, as it is essentially stating that the temperature at the interface has reached the boiling temperature. The second condition (Eqn. (1562)) states that a pure vapor (without the presence of non-condensable gases) can only be in quasi-steady equilibrium at the interface if its temperature is the saturation temperature. If either of these conditions are detected, the hydrodynamic procedure described above is cut short and the interfacial temperature is set to T_{sat}.

Final Application of Thermal Equilibrium

Once T_s is found (either below the boiling temperature for hydrodynamically limited conditions or exactly at the boiling temperature for thermally limited conditions), STAR-CCM+ finalizes the calculation by using the following steps:

1. Compute $\dot{Q}_v$ via enforcing the thermal equilibrium of Eqn. (1560).
2. Compute $\dot{m}_v$ using Eqn. (1559).
3. Compute $\dot{m}_{v,i}$ using Eqn. (1556).

Note that if $T_s < T_{sat}$, these final steps are not necessary: they are fully consistent with the quantities that are computed during the iterative process, and would therefore be already available. However, these final steps are presented to illustrate that, once T_s has been found, evaporation rates can be calculated in an identical fashion for both the non-saturated and the saturated regimes.

Steam Condensation at a Dry Wall

The method that is outlined above runs into issues at perfectly dry cold walls that are in contact with a pure vapor (for example, at the start of the simulation).

The heat transfer coefficient in the film is generally given by:

$$h_f = \frac{2k_f}{h} \tag{1563}$$

so the following approximation (based on the assumption of piecewise linear temperature profile) holds:

$$k_l \frac{dT}{dy}\bigg|_{l_s} \approx \frac{2k_f}{h}(T_{l,c} - T_s) \tag{1564}$$

For vanishing film thickness and $T_s > T_{l,c}$, Eqn. (1564) tends to infinity, expressing that at zero film thickness the heat is carried towards the wall infinitely fast. At saturation conditions, the thermally limited model (which is assumed here) effectively predicts infinite condensation rates, which causes the simulation to diverge. (The heat transfer coefficient that is calculated by STAR-CCM+ does not actually go to infinity, as a certain minimal thickness (currently 10^{-8}m) is used. However, the calculated value is still extremely high.)

Aside from molecular considerations (such as kinetic theory and interfacial non-equilibrium), the main physics that is missing is in how the condensation is modeled. Condensation at a dry wall does not happen in

the form of a flat film, but rather as droplets growing at nuclei. The growth of these droplets is restricted by surface tension effects, and the interfacial contact area between vapor and liquid is very small, so it is likely that these become limiting factors at the early stages of condensation.

To implement a model that behaves in the same way that the purely hydrodynamically limited model behaves (that is, models the condensation as a fully-wetting film), but remains stable, a simplified drop-wise approach is used. The impact of the model is only on the interfacial area.

For the derivation, consider a film cell with area A and thickness h. Assume that there is a nucleation density of N (m^{-2}) and that the nuclei have a minimum radius of R_{min}. It is assumed that the droplets form with a 90 degree contact angle (a temporary simplification), and that they are all identical, with a radius R. For a given film thickness, the radius of the drops (assuming hemispheres) can be expressed as:

$$R = \left(\frac{3h}{2\pi N}\right)^{1/3} \tag{1565}$$

It is assumed that the evaporation rates (as calculated before) still hold and that the only effect is due to the smaller total interfacial area between the vapor and the film. This ratio can be described as:

$$f_A = 2\pi N R_{eff}^2 \tag{1566}$$

which behaves as $N^{1/3}h^{2/3}$.

When this factor f_A is smaller than 1, it is used as a multiplier for the evaporation rates. The model requires two additional parameters: the nucleation density N and the minimum radius R_{min}. These parameters are specified as properties of the evaporation and condensation model. N is specified as N_seeds and R_{min} is specified as R_seeds. For more information, see Evaporation/Condensation Expert Properties.

This model stabilizes the solution without affecting the longer-term result (that is, once a fully-wetting film has formed and grown for a while).

Boiling

Boiling is considered when the temperature of the fluid film is above the boiling temperature and the boiling property is activated. Below the boiling temperature, the same equilibrium conditions as for evaporation are assumed, but the vapor mass rates are determined differently:

Bulk Boiling

The heat transfer to the interface between the phases is used to compute the mass rate of evaporation when film cells are boiling. It is assumed that:

- The vapor bubbles are at the saturation temperature.
- The temperature of the liquid is approximately the temperature of the mixture.
- All the heat flux from the liquid to the interface is used in mass transfer due to evaporation.

$$\dot{m}_v = \frac{\mathrm{HTCxArea}(T - T_{sat})}{\Delta H^{vap}} \tag{1567}$$

$\mathrm{HTCxArea}$ is the product of the heat transfer coefficient between the vapor bubbles and the surrounding liquid, and the specific contact area (contact area per unit volume) between the two. ΔH^{vap} is the latent heat of vaporization of the mixture and computes as:

$$\Delta H^{vap} = \sum_i Y_{g,i} \Delta H_i^{vap} \tag{1568}$$

where $Y_{g,i}$ is the mass fraction and ΔH_i^{vap} is the latent heat of pure component i. This latent heat is the difference between the total enthalpy of its liquid phase $H_{i,l}$ and its gas phase $H_{i,g}$.

$$\Delta H_i^{vap} = H_{i,g} - H_{i,l} \tag{1569}$$

Wall Boiling

The empirical standard pool boiling correlation [284] is used to model liquid film boiling at a hot wall. The film starts to boil when the wall temperature exceeds the boiling temperature of the liquid. The heat flux passing from the wall to the film increases in proportion to the cubic power of $\Delta T_W = T_W - T_{sat}$ as:

$$q_{boil} = \frac{\mu_l \Delta H^{vap}}{\mathrm{Pr}_l^n} \sqrt{\frac{g(\rho_l - \rho_g)}{\sigma}} \left(\frac{C_{p,l} \Delta T_W}{c_{sf} \Delta H^{vap}} \right)^3 \tag{1570}$$

where:

- the viscosity μ_l, the Prandtl number Pr_l^n, and the specific heat $C_{p,l}$ are evaluated for the liquid phase.
- ρ_l and ρ_g denote the densities of the liquid and gas phases.
- g and σ denote the gravity acceleration and surface tension coefficient respectively.

- c_{sf} and n characterize the fluid wall interface.

  For example, for water-brass contact, the values are $c_{sf} = 0.006$ and $n = 3$.

- T_{sat} represents the saturation temperature of the mixture and is computed from the temperature-dependent saturation pressure of its components by applying Dalton's law.

The heat flux increases as T_W increases but does not exceed a critical heat flux that is:

$$q_{max} = c_{max}\rho_g^{1/2}\Delta H^{vap}[g(\rho_l - \rho_g)\sigma]^{1/4} \tag{1571}$$

where $c_{max} = 0.15$. After q_{max} has been attained, the heat flux decreases to its minimum at the Leidenfrost temperature T_L:

$$q_{min} = c_{min}\rho_g\Delta H^{vap}[g(\rho_l - \rho_g)\sigma]^{1/4}(\rho_l + \rho_g)^{-1/2} \tag{1572}$$

where $c_{min} = 0.09$.

The wall temperature of maximal heat flux $T_S = T_{sat} + \Delta T_S$ can be obtained from equations for q_{max} and q_{boil} resulting in

$$\Delta T_S = \frac{c_{sf}c_{max}^{1/3}\mathrm{Pr}_l^{n/3}\sigma^{1/4}\rho_g^{1/6}\Delta H^{vap}}{C_{p,l}\mu_l^{1/3}g^{1/12}(\rho_l - \rho_g)^{1/12}} \tag{1573}$$

Based on this correlation, the heat flux from the wall is computed as follows:

If $\Delta T_W < \Delta T_S$, then

$$q_W = q_{boil} \tag{1574}$$

Else, if $\Delta T_W < c_s\Delta T_S$, then

$$q_W = q_{max} \tag{1575}$$

Else, if $\Delta T_W < \Delta T_L$, then

$$q_W = q_{min} + \frac{\Delta T_L - \Delta T_W}{\Delta T_L - c_s\Delta T_S}(q_{max} - q_{min}) \tag{1576}$$

Else,

$$q_W = q_{min}\frac{\Delta T_W}{\Delta T_L} \tag{1577}$$

$\Delta T_L = T_L - T_{sat}$ and c_s is assumed to be 1.2.

The component mass rates of vaporization are then computed as

$$\dot{m}_{v,i} = c_e Y_{g,i} \frac{q_W}{\Delta H^{vap}} \tag{1578}$$

where c_e is the wall boiling evaporation factor that determines how much of the boiling heat flux generates vapor. A value smaller than one means that a portion of the boiling heat flux does not generate vapor but heats up the film, for example:

- when heat is supplied near the wall to bubbles which then move away into the liquid film and condensate again inside the film
- when evaporation takes place at the base of a bubble and condensation at its tip.

Melting and Solidification

Melting-Solidification Enthalpy Formulation

The Melting-Solidification model does not track the liquid-solid interface explicitly. Instead, the model uses an enthalpy formulation to determine the distribution of the solidified portion of the liquid-solid film.

For basic melting and solidification, the enthalpy of the liquid-solid film h_{ls}^* includes the latent heat of fusion h_{fusion}:

$$h_{ls}^* = h_{ls} + (1 - \alpha_s^*) h_{fusion} \tag{1579}$$

where h_{ls} is the sensible enthalpy. The relative solid volume fraction α_s^* is defined as the portion of the volume of the liquid-solid film which the solid state occupies. In the enthalpy model, the relative solid volume fraction α_s^* is a function of temperature:

$$\alpha_s^* = \begin{cases} 1 & \text{if} \quad T^* < 0 \\ f(T^*) & \text{if} \quad 0 < T^* < 1 \\ 0 & \text{if} \quad 1 < T^* \end{cases} \tag{1580}$$

where T^* is the normalized temperature that is defined as:

$$T^* = \frac{T - T_{solidus}}{T_{liquidus} - T_{solidus}} \tag{1581}$$

The function $f(T^*)$ is called the fraction solid curve. For a linear dependence between α_s^* and T^*, the solidification path is defined as:

$$f(T^*) = 1 - T^* \tag{1582}$$

The solidification path can also be specified as a table, in which case linear dependencies are assumed between the points in the table. The values of temperature in the table have to increase (decrease) monotonically as the values of the relative solid fraction decrease (increase) monotonically.

If melting and solidification take place at one temperature (that is, $T_{\text{solidus}} = T_{\text{liquidus}}$), a linear solidification path is assumed and a small temperature interval of 0.002 K is automatically introduced.

The physical properties of the liquid-solid phase can be constant, or can be a function of temperature and the relative solid volume fraction.

Removal of Solidified Mass

When the Solidified Film Option property is set to Remove Solidified Film, and the Morpher is activated, the Melting-Solidification model tracks the liquid-solid interface.

For more information, see Setting up Solidified Mass Removal.

During every time step, a *Solid Film Thickness Increment* Δh_s is computed, updated at every inner iteration during the segregated process. This value represents the increase in solidified thickness during the time step. Δh_s is then used to construct the sources and sinks into the various transport equations.

The updates (at every inner iteration) are computed as:

$$\Delta h_s' = \begin{cases} -\dfrac{\Delta H(T) - \Delta H_{liq}}{\Delta H_{fus}}(h + \Delta h_s) & \text{if} \quad T \geq T_{liq} \\ h \cdot \alpha s & \text{otherwise} \end{cases} \tag{1583}$$

with $\Delta h_s'$ the correction to Δh_s. This expression implies that (within each time step) the amount of solidified mass increases when the temperature is below the liquidus temperature, but is reduced if the temperature is above the liquidus temperature. This behavior is in accordance with intuition. However, the present implementation has a lower limit:

$$\Delta h_s \geq 0$$

and therefore this feature only allows mass to solidify; it does not allow the reverse process of melting. Once the solidified mass has been removed, it leaves the system and cannot be retrieved.

The solid film thickness increment can be related to a mass sink $\frac{kg}{m^2 s}$:

$$\dot{m} = \frac{\rho_{liq}\Delta h_s}{\Delta t} \tag{1584}$$

which is used as a base for corresponding sources in the continuity, momentum, energy, and species transport equations in the film.

Morphing - Vertex Displacement

As described above, when the Solidified Film Option property is set to Remove Solidified Film, the solidified mass is removed from the film, leaving only liquid behind.

For each time step, the solidification is computed as follows:

- A mass sink is applied to the liquid film, corresponding to how much liquid is turning into solid mass.
- This mass sink is used to calculate the velocity of the solid front (that is, the velocity that the interface between solid and liquid is moving), also taking into account the density of the solid mass.
- This velocity is used to calculate the solid thickness increment Δh_s for that time step; simply $\Delta h_s = V \times dt$.

Typically, this solidified mass accumulates on the wall: it remains in place (it does not flow any further) and deforms the surface as seen by the gas and the remaining film. This effect is modeled by feeding the displacement Δh_s into the Morpher.

For each time step, the displacement of every face is:

$$\bar{d}_{face} = \Delta h_s \hat{A} \tag{1585}$$

where $\hat{A}$ is a unit vector, normal to the film surface and pointing into the gas (that is, the growth is assumed to happen in the direction normal to the surface). The displacement $\bar{d}_{face}$ is then interpolated to displacements $\bar{d}_{vert}$ on the vertices.

The Solid Time-step Factor value (K) acts as a multiplier for the solid front velocity (that is, $V = K \times V$). Consequently, within each time step, the solid thickness increment is increased by a factor K. This can be interpreted as a form of solid multiplier (that is, each 1kg/s of solidifying liquid is turned into K kg/s of solid). This could be considered as a mass source, although there are no transport equations solved for the solid: the solid is removed from the simulation and its effect is included only through the morphing of the boundary. The total physical time should be reduced by a factor K to get the same physical result.

Morphing - Smoothing

When the vertices are displaced according to the algorithm that is described above, the displacement can lead to roughness in the deformation, forming a zig-zag pattern. This roughness can cause instability issues in the fluid film transport.

To resolve this issue, an optional smoother is available. The smoother recomputes the displacement as a weighted average, including the displacements at the nearest neighbors. This smoothing can be repeated multiple times: set the Smoothing Steps property to the appropriate value.

Fluid Film Nomenclature

The following is a guide to the terms used in this section. The subscript f denotes the fluid film values. The subscript g denotes a multi-component gas value and the superscript s denotes a surface value. The quantities without the subscript f are associated with the surrounding fluid.

ΔH^{vap}	enthalpy of vaporization
α_s^*	relative solid volume fraction
ξ	local coordinate normal to the wall
λ_r	resonance wavelength
μ_f	film viscosity
σ	surface tension coefficient
ρ_f	film density
τ_{fs}	shear force at the film free surface
$v_{g,n}^s$	normal component of gas velocity at surface s
$v_{f,n}^s$	normal component of film velocity at surface s
B	Spalding transfer number
Bo_f	film Bond number
c_H	Minimum film thickness coefficient
c_D	Droplet diameter coefficient
c_B	Breakup rate scaling coefficient
c_e	wall boiling evaporation factor
D_g	gas molecular diffusion coefficient
D_f	film molecular diffusion coefficient
FR	force ratio
$\mathbf{f}_b$	body force
h_a	film height that is stripped (surface wave amplitude)
h_f	film thickness
h_{fusion}	latent heat of fusion
k_f	film thermal conductivity
k_g	gas thermal conductivity
L_b	break-up length
$\dot{m}_v$	evaporation rate

N_v	number of interacting components
N_f	number of components in fluid film
N_g	number of components in gas
$N_{P,f}$	number of inert components in fluid
$N_{P,g}$	number of inert components in gas
$\mathbf{n}$	wall surface unit vector
p_f	pressure
p^*	saturation pressure
Pr	Prandtl number
Re_f	film Reynolds number
r_d	droplet radius
s_m	mass source/sink per unit area
$\mathbf{s}_m$	momentum source
s_t	species transfer coefficient
$\mathbf{T}_f$	viscous stress tensor
T	temperature
T_L	Leidenfrost temperature
T_{sat}	saturation temperature
t_b	time scale
$\mathbf{v}_f$	film velocity
$\mathbf{v}_w$	wall velocity
$\bar{\mathbf{v}}_f$	average film velocity
We_f	film Weber number
We_{rel}	relative Weber number
X_g^s	gas molar fraction
X_f^s	fluid film molar fraction
x_S	stripping fraction
Y_g^s	gas mass fraction
Y_f^s	fluid film mass fraction

Fluid Film Bibliography

[275] Bai, C. and Gosman, A.D. 1996. "Mathematical modeling of wall films formed by impinging sprays", SAE Technical Paper Series 960626.

[276] Sirignano, W.A. 1999. *Fluid Dynamics and Transport of Droplets and Sprays*, Cambridge University Press, New York.

[277] Torres, D.J., O'Rourke, P.J., and Amsden, A.A. 2003. "Efficient multicomponent fuel algorithm", *Combust. Theory Modelling*, **7**, p. 67.

[278] Foucart, H. Private Communication.

[279] Ledoux, M. 1993. *Modelisation des Sprays*, Cours de DEA.

[280] Maroteaux, F., Llory, D., Le Coz, J-F., Habchi, C. 2002. *Liquid film atomization on wall edges separation criterion and droplets formation model*, J. Fluids Engineering, **124**, pp. 565-575.

[281] Gubaidullin, A. 2007. *Comments on Liquid film atomization on wall edges separation criterion and droplets formation model*, J. Fluids Engineering, **129**, pp. 665-666.

[282] Friedrich, M. A., Lan, H., Drallmeier, J. A., Armaly, B. F. 2007. *A separation criterion with experimental validation for shear-driven films in separated flows*, Accepted for publication by ASME J. Fluids Engineering.
[283] Arai, T., Hashimoto, H., 1985. *Disintegration of a thin liquid sheet in a concurrent gas stream*, Proceedings of the Third International Conference on Liquid atomization and Spray Systems, London
[284] White, F.M. 1988. "Heat and Mass Transfer". Addison Wesley.

Fluid Film Field Functions Reference

The following primitive field function is made available to the simulation when you activate the Fluid Film model. This object has typical base properties.

Fluid Film Thickness	h_f in Eqn. (1507).

The following primitive field functions are made available when the Temporary Storage Retained expert property of the Fluid Film Segregated Flow solver is activated. They allow you to see whether any stripping occurs and how fast.

Film Wave Stripping Minimum Thickness	The minimum thickness that allows stripping, $h_{\min}$ in Eqn. (1518).
Film Wave Stripping Time Scale	The time scale (t_b in Eqn. (1521)) of wave stripping, where the film thickness exceeds $h_{\min}$.

The following primitive field functions are made available when the Fluid Film Evaporation and Condensation model is activated. They allow you to monitor how fast evaporation or condensation occurs.

Film Evaporation Latent Heat Flux	The latent heat due to evaporation (W/m^2) that is injected at the interface.
Film Evaporation Rate	The rate of evaporation in kg/m^2-s for every species.
Film Evaporation Mass Fraction	The saturated vapor mass fraction on the gas side of the interface, for every interacting component.

The following primitive field functions are made available when the Fluid Film Boiling model is activated.

Film Boiling Latent Heat	The latent heat of phase transfer due to boiling (J/kg).
Film Boiling Latent Heat Flux	The total power (W/m^2) that is spent during the phase change due to wall boiling. This field function is available on both gas and liquid sides.
Film Boiling Rate	The evaporation rate due to wall boiling (kg/m^2-s) for every interacting component. This field function is available on both gas and liquid sides.
Film Boiling Temperature	The boiling temperature (K) of the liquid mixture.
Film Bulk Boiling Latent Heat Flux	The total power (W/m^2) that is spent during the phase change due to boiling in film cells. This field function is available on both gas and liquid sides.
Film Bulk Boiling Rate	The evaporation rate due to boiling inside the film cell (kg/m^2-s) for every interacting component. This field function is available on both gas and liquid sides.
Film Boiling Flag	Indicates if the conditions for boiling are met in the film $T > T_{sat}$

The following primitive field function is made available when the Fluid Film Melting and Solidification model is activated.

Relative Solid Volume Fraction	The volume fraction of the fluid film that is in the solid state.

Using the Discrete Element Method

The Discrete Element Method (DEM) is an engineering numerical method to simulate motion of many interacting discrete objects that are typically solid particles. Although DEM modeling demands significant computing power, it provides detailed resolution that other methods cannot achieve.

Established by Cundall and Strack [287], the DEM model is an extension of the Lagrangian modeling methodology to include dense particle flows. The distinct characteristic of DEM is that inter-particle contact forces are included in the equations of motion. These forces cannot be ignored for highly loaded flows, that is, many interacting particles.

STAR-CCM+ uses a classical mechanics method to model DEM, and is based on soft-particle formulation where particles are allowed to develop an overlap. The calculated contact force is proportional to the overlap, as well as to the particle material and geometric properties.



Applications include modeling the granular flow of:

- Sand
- Food particles
- Powders
- Capsules
- Slurries

DEM in STAR-CCM+ is implemented within the Lagrangian framework and uses Lagrangian phase models to define the particles. In addition, the DEM model uses its own phase and phase interaction models for specific DEM particle properties and their behavior during contact with each other or wall boundaries.

The DEM model also uses Lagrangian injectors and its own lattice injector and random injector to introduce particles into the fluid region.

After running the solver, you can clear DEM particles from the region by activating the Lagrangian Dem Solution Fields option in the Clear Solution dialog.

Regarding mesh motion, the DEM model is only compatible with rigid body motion (Rotation or Rotation and Translation) objects. Hence, you cannot use the DEM model with DFBI.

Considerations for the Discrete Element Method

When setting up a simulation that uses DEM, take the following points into consideration:

- Use a coarser mesh that typical CFD applications. If a particle is in a small cell and occupies most of its volume, it can cause high void fractions and thus solver instability.
- Typically, the most robust approach to solving the continuous phase with DEM is to use the coupled solver with incompressible, ideal gas.
- To reduce computation times, reduce the Young's Modulus of the particles. However, this change in physical properties compromise the accuracy and validity of your results.

The DEM model is represented in the STAR-CCM+ physics continuum by the Discrete Element Model (DEM) node.



Discrete Element Model Properties

None.

Selecting the Discrete Element Model

As a prerequisite, select the following continuum models:

- In the *Physics Model Selection* dialog, select the Lagrangian Multiphase model.
- Select the Implicit Unsteady model.

Then, to select the DEM model:

- In the *Optional Models* group box, select the Discrete Element Model (DEM) option.



Note: If you set up Lagrangian phases with DEM particles, clearing the DEM model automatically deletes the Lagrangian phases, because they no longer serve any purpose in the simulation. To prevent the loss of those phases, clear the DEM Particles phase model in each Lagrangian phase before clearing the DEM continuum model.

This section describes the following:

- Modeling DEM particles
- Modeling drag force in DEM

- Modeling drag torque in DEM
- Modeling phase interactions in DEM
- Defining injectors in DEM
- What is the DEM solver?
- DEM formulation
- DEM field functions reference

Modeling Discrete Element Method Particles

The DEM Particles model specifies that the Lagrangian phase represents solid particles and the Discrete Element Method (DEM) models the particle behavior within the simulation. The momentum balance formulation determines how the particles interact with each other and with solid boundaries. The image below shows the effect of particle-particle collisions and particle-wall collisions.



Caution: If a particle spans 3 or more mesh partitions (due to small cells and a high number of computational processes), the solution can miss particle-particle or particle-wall collisions, or record them late. To avoid this error, either use a coarser mesh, or increase the size of mesh partitions by reducing the number of processes.

As the DEM model simulates solid particles, the only available Lagrangian phase material models are solid, multi-component coal, and multi-component solid. Similarly, the only available Lagrangian phase equation of state models are constant density and polynomial density.

The DEM Particles model is represented in the Lagrangian phase by the DEM Particles node.



DEM Lagrangian phases include phase models that allow you to define the shape of the particles as:

- Composite particles—Non-spherical particles with fixed bonds between the particle components.
- Particle clumps—Non-spherical particles with breakable bonds between the particle components.
- Spherical Particles—Single spherical particles.

DEM Particles Properties

None.

Selecting the DEM Particles Model

As a prerequisite, select the following continuum models:

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM) model.

Then, to select the DEM Particles phase model for a newly created Lagrangian phase:

- In the *Phase Model Selection* dialog, select the DEM Particles model from the *Particle Type* group box.



What Is the Composite Particles Model?

The Composite Particles model uses multiple shapes (spheres) to represent non-spherical particles. The shapes are fixed together and do not separate during the simulation. If you wish to model particle bonds that can break, select the Particle Clumps model.

You define the arrangement of spheres within a composite particle using the particle shapes editor to compose the shape yourself, or using the Update from Part option on the Composite Particles node to have STAR-CCM+ do it.

The image below shows an example of a composite particle that uses five spheres to represent its shape.



The Composite Particles model is represented in the STAR-CCM+ Lagrangian phase by the Composite Particles node. Select the node to access the properties of the model. Right-click the node to access the pop-up menu of the Composite Properties model.



Selecting the Composite Particles Model

As a prerequisite, select the following continuum and phase models:

- In the *Physics Model Selection* dialog, select the Discrete Element Method (DEM) model.

- In the *Phase Model Selection* dialog, select the DEM Particles model.

Then, to select the Composite Particles phase model:

- In the *Phase Model Selection* dialog, select the Composite Particles model from the *Particle Shape* group box.



Thereafter, you can use the particle shape editor or the Composite Particles Part property to create shapes.

What Is the Particle Clumps Model?

The Particle Clumps model uses multiple shapes (spheres) to represent non-spherical particles. You can use the particle shape editor to define the initial particle assembly. The join between these spheres is modeled using the Bonded Particles phase interaction model. This model allows the parts

of the clump to separate if the bond between particles exceeds the maximum tensile or shear stress. If you wish to model non-spherical, fixed particles, select the composite particles model.

In contrast to the composite particles model, the particle clumps model defines particles as assemblies of spheres that are connected through parallel bonds. The composite particles model only treats the individual spheres of a particle as separate entities for contact detection, where contacts between particles must be resolved separately. In the remainder of the calculation, the particle is a single object. With the particle clumps model, each sphere in the assembly is treated as an individual particle, with a separate position, velocity, angular velocity, and orientation that are maintained by the system.

When the injector injects a particle, bonds immediately form between particles in contact. When a bond breaks, the clumps break up into their individual shapes. The initial configuration of the particle clump is considered to be in a state of equilibrium, that is, all the contact forces are balanced.

The equations controlling the particle clumps model specify that when the bond breaks, the particles interact according to the Hertz-Mindlin law. The bonding and grain forces do not act in parallel in this case. This fact allows STAR-CCM+ to define the initial particle shape without the excessive limitations that the small contact overlap of the soft particle model imposes.

The image below shows an example of particle clumps breaking up as they strike a wall boundary.



The Particle Clumps model is represented in the STAR-CCM+ Lagrangian phase by the Particle Clumps node. Select the node to access the properties of the model. Right-click the node to access the pop-up menu of the Particle Clumps model.



Selecting the Particle Clumps Model

As a prerequisite, select the following continuum and phase models:

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM) model.
- In the *Phase Model Selection* dialog, select the DEM Particles model.

Then, to select the Particle Clumps phase model:

- In the *Phase Model Selection* dialog, select the Particle Clumps model from the *Particle Shape* group box.



What Is the Spherical Particles Model?

The Spherical Particles model uses a single sphere to represent the particle. If you wish to model non-spherical particles, select the Composite Particles or Particle Clumps models

The Spherical Particles model is represented in the Lagrangian phase by the Spherical Particles node.



Spherical Particle Model Properties

None.

Selecting the Spherical Particles Model

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM) model.
- In the *Phase Model Selection* dialog, select the DEM Particles model.

Then, to select the Spherical Particles phase model:

- In the *Phase Model Selection* dialog, select the Spherical Particles model from the *Particle Shape* group box.



What Is the Coarse Grain Particle Model?

A common problem with using the Discrete Element Method is the high computational cost of simulating a realistic number of particles. One way to reduce this cost is to use the Coarse Grain Particle model.

In this model, each simulated particle is a parcel that represents multiple smaller particles. STAR-CCM+ uses the coarse-grained particle for calculating fluid-particle interactions, while it continues to calculate contact dynamics such as collisions on the fine scale particles. The number of individual particles per coarse-grained particle depends on injector conditions and values, such Flow Rate Specification, Particle Size Specification, and Mass Flow Rate.

The Coarse Grain Particle model is best used for fluidized beds applications, where accuracy of particle-fluid interactions is more important than accuracy of particle-particle interactions.

Do not use the Coarse Grain Particle model when the mesh size is smaller than the size of the coarse-grained particle.

The Coarse Grain Particle model is represented in the Lagrangian phase by the Coarse Grain Particle node.



Coarse Grain Particle Model Properties

None. See Coarse Graining Option Specification in the Injector Conditions Reference for the Lagrangian model.

Selecting the Coarse Grain Particle Model

As a prerequisite, select the Lagrangian Multiphase and Discrete Element Model (DEM) models.

- Under the Lagrangian Multiphase node, right-click the Lagrangian Phases node and create at least one phase.
- Under a phase, right-click the Models node and select, in order, the DEM Particles model, then the Coarse Grain Particle model.

Working with Individually Composed Shapes

STAR-CCM+ allows you to define composite particles or particle clumps as an arrangement of spheres. You define the relative dimensions and positions of the spheres in an arbitrary unit system, as STAR-CCM+ scales the final particle that is injected into the system according to the injector properties. The scale factor is:

$$\left(\frac{V_{injector}}{V_{particle}}\right)^{\frac{1}{3}}$$

where:

- $V_{particle}$ is the particle volume computed using Monte Carlo integration
- $V_{injector}$ is the injector volume

The injector volume is the volume of a sphere with the diameter specified at the injector:

$$V_{injector} = \pi\left(\frac{D_{injector}^{3}}{6}\right)$$

where:

- $D_{injector}$ is the injector diameter

Each sphere is represented in the STAR-CCM+ Lagrangian phase by a Sphere node, under the Shapes manager. Select a node to access the properties of the corresponding sphere.



Composite Particles and Particle Clump Properties

Principal Moment of Inertia	Read-only property reporting the principal moment of inertia of the particle, which is composed of user-defined shapes. This property is used when calculating the momentum balance (see Eqn. (1588)).
Volume	Read-only property reporting the volume of the particle.
Area	Read-only property reporting the area of the particle.
Sphericity	Read-only property reporting the sphericity of the particle. This property is used when calculating the Haider and Levenspiel drag coefficient (see Eqn. (1644)).

Composite Particles and Particle Clumps Pop-Up Menu

Update Properties	Updates the particle properties that are based upon the current geometry, that is, the user-defined shapes.
Show model	Displays a representation of the particle in the *Graphics* window.

Sphere Properties

Position	Specifies the center-point of the sphere using Cartesian coordinates that are relative to an arbitrary coordinate system. This property is used to locate the spheres relative to one another.
Radius	Specifies the radius of the sphere with an arbitrary unit. The actual size of the sphere is not specified here, as the entire composite particle or clump is scaled to match the injector.

Example of Creating a Particle Using Shapes

The following example uses shapes to create a composite particle with a non-spherical shape. A similar procedure can also be used to create a

particle clump.

First, right-click the Phase > Models > Composite Particles > Shapes manager and select New to create a sphere. In the *Properties* window, set the Radius and Position properties. As the position and radius are arbitrary, with the spheres defined relative to one another, you can typically define the first sphere with the default values.



You can continue to create numerous spheres to represent the particle. Define each particle relative to the first.



If two particles have similar properties, you can copy and paste the node to create an identical sphere and then change its properties.



To visualize the particle, right-click the Composite Particles node and select Show model.



The *Graphics* window then displays the particle as a *Composite particle* scene.



Working with Sphere-Filled Parts

The composite particle method simulates particles with complex shapes by arrangements of spheres that are kept in fixed relationships to each other. You can create a composite particle by selecting a previously defined part and having STAR-CCM+ fill the part with a specified number of spheres.



Original part. Composite particle, small number of spheres. Composite particle, large number of spheres.

This method works well for largely convex objects such as rocks or grains, but is more likely to fail for complex concave objects.

Simulating sharp edges and corners requires small spheres, but small spheres produce a small time-step in a DEM simulation—and hence higher computation time. To avoid small time-steps, you can edit individual

spheres to the largest size acceptable to you, after STAR-CCM+ creates the initial composite particle.

Every sphere that is created by this method has some part of the surface exposed. As a result, the interior of the particle is hollow in some cases. Spheres that are completely embedded in the particle have no effect on simulation.

The method does not always produce a symmetrical composite particle when filling a symmetrical part. If symmetry considerations are important in the simulation, try a different number of spheres and see if the result is more symmetrical, or use the particle shapes editor instead.

A good practice is to start with a small number of spheres and see what the resulting representation looks like. If the representation is not satisfactory, then increase the number of spheres.

Creating Sphere-Filled Parts

To create a composite particle by filling a part with spheres:

1. Select the Composite Particles node under the Models node for the Lagrangian Phases.
2. In the *Composite Particles - Properties* window, edit Part.
3. In the *Edit* dialog, select the part that you want to use as a composite particle, then click OK.
4. In the *Composite Particles - Properties* window, set the Number of Spheres property.
5. Right-click the Composite Particles node and select Update from Part. STAR-CCM+ creates a composite particle that is based on the selected shape, using the specified number of spheres.



6. To visualize the particle, right-click the Composite Particles node and select Show model.

Setting Phase Boundary Conditions

The Lagrangian phase boundary conditions specify how the particles interact with wall boundaries. As the Discrete Element Method (DEM) is implemented within the Lagrangian framework, you can find further reading on the interaction of particles with boundaries in the Understanding Boundary Interaction Modes section.

Note: For contacts between particle and boundary, the model ignores the tangential component of the contact force when the viscous regime is set to Inviscid.

The Lagrangian phase boundary conditions are represented in the STAR-CCM+ Lagrangian phase by the Lagrangian Phase > Boundary Conditions node.



Mode

The mode defines how the particle behaves when it interacts with the boundary. Depending on the boundary type, the following boundary modes are available when using DEM:

- Escape—Particles can leave the region through the boundary.
- Transmit—Particles can cross an interface, from one region to another.
- DEM mode—Particles contact and rebound off the boundary. Selecting this mode activates the Boundary Properties node, which defines the boundary surface material properties.

- DEM trap mode—Particles leave the region through the boundary at a rate you specify.

Boundary Properties

The boundary properties specify the surface material properties at the boundary and are only available with the DEM mode and DEM trap mode. These properties influence how the particles behave when they contact the boundary. For example, a wood boundary surface would absorb more of the particle contact than an aluminum surface.

You can specify the material properties by either setting the material properties manually, or using a materials database.

Flow Rate Specification

The flow rate specification is available with the DEM trap mode only. This condition specifies how the flow-rate of particles at the boundary is defined.

Flow Rate Specification Properties

Method	Selects the method for entering the flow rate.	
	Mass Flow Rate	Indicates that the rate at which particles leave the boundary as a mass flow rate. Hence this condition is not available with the Massless Particles model. A Mass Flow Rate node is added to the Values node.
	Volume Flow Rate	Indicates that the rate at which particles leave the boundary as a volume flow rate. Hence this condition is not available with the Massless Particles model. A Volume Flow Rate node is added to the Values node.
	Particle Flow Rate	Indicates that the rate at which particles leave the boundary is specified directly. A Particle Flow Rate node is added to the Values node.

Physics Values

You can specify the rate at which DEM particles leave a boundary in the Physics Values node. The Mass Flow Rate, Volume Flow Rate, or Particle Flow Rate are specified as scalar profiles.



Modeling Particle Discharge on Wall Boundaries

A fluidized bed is a long geometry with a small orifice through which particles are discharged at one end. To save simulation time and avoid meshing the small features of the orifice, use the DEM Trap Walls model to approximate a discharge boundary condition through the wall of such geometries.

The DEM Trap Walls model allows you to specify the rate at which DEM particles leave a boundary. The DEM Trap Walls node represents this model in the simulation tree.



To specify a discharge boundary condition:

- Select the DEM Trap Walls model.
- Set the Mode of the boundary to DEM trap mode.

  The Mode node is contained within the Phase Conditions node of wall boundaries when DEM is active in the region. Setting the Mode to DEM trap mode adds a Flow Rate Specification node to the Physics Conditions for the boundary, and a corresponding value flow rate node depending on the selection for Flow Rate Specification.
- Set the Flow Rate Specification for the boundary.
- Set the value of the flow rate.

Unlike a fluid, particles do not form a continuum. Therefore, the outflow boundary conditions cannot be guaranteed. The DEM Trap Walls model only ensures that the specified maximum outflow is not exceeded. If the flow rate is set to an unrealistically high value, this model does not force particles to exit the domain.

Selecting the DEM Trap Walls Model

As a pre-requisite, select the DEM Particles model from the *Particle Type* group-box in the *Phase Model Selection* dialog. Then select DEM Trap Walls from the *Optional Models* group box.



The DEM Traps Walls node appears in the simulation tree.



Modeling Specific Material Properties for Specific Walls

You can specify different materials for different wall boundaries of a region while using DEM.

1. When you create the region, use the option Create a Boundary for Each Part Surface. See Boundary Options.
2. For each boundary where you want to customize the material properties:
   a. In the Regions node, select one of the relevant boundaries, go to Physics Conditions > Lagrangian Specification, and change the Method property to Specify for Boundary.
   b. Open up the newly created Phase Conditions node, right-click the Solid > Al node, and select Replace With...
   c. Under Physics Models > Multiphase Interaction > Phase Interactions, define a new Phase Interaction pair between the DEM particle phase and the newly customized boundary. See Modeling Phase Interactions in the Discrete Element Method.

   Repeat these steps for each customized wall boundary.

Modeling Drag Force in the Discrete Element Method

The drag force model calculates the force acting on a particle due to viscosity-induced fluid resistance. As DEM is implemented within the Lagrangian framework, Eqn. (1332) in the Lagrangian Formulation section is used to calculate the drag force.

STAR-CCM+ calculates the drag force using drag coefficients. It is impractical to determine these drag coefficients spatially, so STAR-CCM+ uses experimental and theoretical correlations to obtain suitable values. For DEM, the following methods are available to determine the drag coefficient:

- Di Felice
- Gidaspow
- Haider and Levenspiel

These methods are provided in addition to the Schiller-Naumann Lagrangian drag coefficient methods.

Field functions for these drag coefficients are always available, regardless of whether you have selected the drag coefficient method or not. As a result, you can combine the STAR-CCM+ built-in drag coefficients with your own calibration factors in a user-defined field function. For example:

0.5*$ParticleGidaspowDragCoefficient

The Drag Force model and Drag Coefficient method are represented in the STAR-CCM+ Lagrangian phase by the Drag Force and Drag Coefficient nodes respectively. No properties are provided on the Drag Forces model node; select the Drag Coefficient node to access the properties the appropriate method.



Drag Coefficient Method Properties

Dimensions	The dimensionality of the drag coefficient (read-only).	
Method	Selects the method for specifying the drag coefficient.	
	Di Felice	Defines the drag coefficient using the Di Felice correlation. Use this method for dense particulate flow. This option is only available with the Two-Way Coupling Lagrangian phase model.
	Field Function	Defines the drag coefficient using a field function, which is typically user-defined.
	Gidaspow	Defines the drag coefficient using the Gidaspow correlation. Use this method for simulating the flow through fluidized beds where the initial packing of particles leads to low-volume fractions. This option is only available with the Two-Way Coupling Lagrangian phase model.
	Haider & Levenspiel	Defines the drag coefficient using the Haider & Levenspiel correlation. Use this method for non-spherical particles with Re < 2.6×10^5. This option is only available with the Two-Way Coupling Lagrangian phase model.
	Schiller-Naumann	Defines the drag coefficient using the Schiller-Naumann correlation, which is one of the standard Lagrangian drag coefficient methods.

Selecting the Drag Force Model

As a prerequisite, select the following physics and phase models:

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM).
- In the *Phase Model Selection* dialog, select the DEM Particles model.

Then, to select Drag Force for DEM:

- In the *Phase Model Selection* dialog, select Drag Force from the *Optional Particle Forces* group box.



What Is the Di Felice Drag Coefficient Method?

The Di Felice drag coefficient method [288] is suitable for dense particulate material flow. Eqn. (1639) in the Formulation section is the method for calculating the Di Felice drag coefficient.

The Di Felice drag force method is represented in the Lagrangian phase by the Di Felice node.



Di Felice Drag Coefficient Method Properties

None.

What Is the Gidaspow Drag Coefficient Method?

The Gidaspow drag coefficient method [290] is suitable for the flow of spherical particles in both low- and high-concentrations. It can be used to simulate flow through fluidized beds where the initial packing of particles transitions to low-volume fractions. The regular drag models, when used for diluted flows, can underpredict the drag in these flow regimes.

The Gidaspow drag coefficient method is a combination of the Wen Yu and Ergun methods, where a cutoff void fraction determines the point at which one method switches to the other. Eqn. (1641) (Wen Yu) and Eqn. (1642) (Ergun) in the Formulation section are the relevant method equations.

The Gidaspow drag force method is represented in the Lagrangian phase by the Gidaspow node. Select the node to access the Gidaspow properties.



Gidaspow Drag Coefficient Method Properties

Exponent	The Wen-Yu power exponent (Eqn. (1642)).
Cutoff Void Fraction	The void fraction cutoff, ν_{min} in Eqn. (1641), for the transition between the Wen-Yu model and the Ergun equation.

What Is the Haider and Levenspiel Drag Coefficient Method?

The Haider and Levenspiel drag coefficient method [291] is suitable for non-spherical particles with Re < 2.6 x 10^5. Eqn. (1643) in the Formulation section is the relevant method equation.

The Haider and Levenspiel drag force method is represented in the Lagrangian phase by the Haider and Levenspiel node.



Haider & Levenspiel Drag Coefficient Method Properties

None.

Modeling Drag Torque in the Discrete Element Method

The Drag Torque model calculates the rotational drag exerted on a particle as it spins in a viscous fluid, according to Eqn. (1645). This model is available in the *Optional Particle Forces* group box for a phase defined as DEM Particles in the Lagrangian Multiphase model.

The Drag Torque model is represented in the Lagrangian phase by the Drag Torque node. Use the Rotational Drag Coefficient node to set the method that STAR-CCM+ uses to predict drag torque



Note: The Drag Torque model applies only to spherical particles.

Rotational Drag Coefficient Properties

Dimensions	The dimensionality of the drag coefficient (read-only).	
Method	Selects the method for specifying the drag coefficient.	
	Sommerfeld	Use the definition supplied in Eqn. (1647)
	Field Function	Define your own field function, to specify the rotational drag coefficient.

Modeling Lift Force in the Discrete Element Method

STAR-CCM+ DEM provides two lift force models, one for spin lift force, and another for shear lift force. In this context, lift forces are forces normal to the particle velocity; they are not necessarily forces in the upward direction.



Modeling the Spin Lift Force

The Spin Lift Force model calculates the force acting on a spinning particle moving in a fluid. STAR-CCM+ calculates the spin lift force using a particle lift coefficient. You can define the lift coefficient using a field function or the Sommerfeld method. See Particle Spin Lift Force, Eqn. (1650).

The Spin Lift Force model and the Particle Lift Coefficient method are represented in the Lagrangian phase by the Spin Lift Force node and the Particle Lift Coefficient Method node respectively. No properties are provided on the Spin Lift Force node; select the Particle Lift Coefficient Method node to access the properties of the appropriate method.



Spin Lift Force Coefficient Properties

Method — Selects the method for specifying the spin lift coefficient.

- Field function — Define your own field function to specify the spin lift force coefficient.
- Sommerfeld — Calculates the spin lift force using the model derived by Sommerfeld [299], given as Eqn. (1650) in the formulation.

Selecting the Spin Lift Force Model

This model is available in the *Optional Particle Forces* group box for a phase defined as DEM Particles in the Lagrangian Multiphase model.



Modeling the Shear Lift Force

The Shear Lift Force model calculates the force acting on a particle moving relative to a fluid where there is a velocity gradient in the fluid orthogonal to the relative motion. STAR-CCM+ calculates the shear lift force using a shear lift coefficient. You can define the lift coefficient using a field function, the Saffman method, or the Sommerfeld method. See Particle Spin Lift Force, Eqn. (1655) and Eqn. (1656).

The Shear Lift Force model and the Shear Lift Coefficient method are represented in the Lagrangian phase by the Shear Lift Force node and the Shear Lift Coefficient Method node respectively. No properties are provided on the Shear Lift Force node; select the Shear Lift Coefficient Method node to access the properties of the appropriate method.

Note: The Shear Lift Force model is available within the Lagrangian Multiphase model as well as within the Discrete Element Model.

Shear Lift Force Coefficient Properties

Method — Selects the method for specifying the shear lift coefficient.

- Field function — Define your own field function to specify the spin lift force coefficient.
- Saffman — Use the definition of Saffman [298] provided as Eqn. (1655). This method recovers the original Saffman asymptotic solution for low Reynolds numbers.
- Sommerfeld — Use the definition of Sommerfeld [299] provided as Eqn. (1656). This definition covers a broader range of Reynolds numbers than the Saffman definition.

See Particle Shear Lift Force.



Selecting the Shear Lift Force Model

This model is available in the *Optional Particle Forces* group box for a phase defined as DEM Particles in the Lagrangian Multiphase model.



Modeling Phase Interactions in the Discrete Element Method

The DEM Phase Interaction model determines how particles behave when they come into contact with each other or solid boundaries. For further information on creating phase interactions and selecting phase interaction models, refer to the section on defining phase interactions. The DEM Phase Interaction model provides access to other models that define particle interaction behavior.

- Hertz Mindlin model
- Walton Braun model
- Conduction Heat Transfer model
- Parallel Bonds model
- Linear Cohesion model

- Bonded Particles model
- Rolling Resistance model

The DEM Phase Interaction model is represented in the STAR-CCM+ phase interaction continuum by the DEM Phase Interaction node. Select the node to access the properties of the model



DEM Phase Interaction Model Properties

First phase	Specifies the first interacting DEM phase. Only DEM phases that are defined under the Lagrangian Multiphase model are available.
Second phase	Specifies the second interacting DEM phase. Only DEM phases that are defined under the Lagrangian Multiphase model are available.

Selecting the DEM Phase Interaction Model

As a prerequisite, select the following physics model:

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM).

Then, to select the DEM Phase Interaction model:

- In the *Phase Interaction Model Selection* dialog, select DEM Phase Interaction from the *Phase Interaction Topology* group box.



What Is the Hertz Mindlin Model?

The Hertz Mindlin no-slip contact model is the standard model that is used for modeling contact between particles.

The Hertz Mindlin model is represented in the STAR-CCM+ phase interaction continuum by the Hertz Mindlin node. Select the node to access the properties of the model.



Hertz Mindlin Properties

Static Friction Coefficient	C_{fs} (see Eqn. (1597)).
Rolling Friction Coefficient	$c_{\text{rolling friction}}$ (see Eqn. (1589)).
Normal Restitution Coefficient	See Eqn. (1600).
Tangential Restitution Coefficient	See Eqn. (1600).

Selecting the Hertz Mindlin Model

To select the Hertz Mindlin model:

- In the *Phase Interaction Model Selection* dialog, select Hertz Mindlin from the *Dem Contact Model* group box.



What Is the Walton Braun Model?

The Walton Braun model is an elastic-plastic contact model, which is useful in situations where a collision between two bodies includes plastic deformation. This situation is typical for materials like grain, soybeans, or coal where the collisions lead to material deformation and the impact energy is dissipated during the collision.

In order to simulate inelastic material, STAR-CCM+ provides a hysteretic plastic-elastic contact model that is based on the formulation of Walton and Braun ([303], [304]).

The Walton Braun model is represented in the STAR-CCM+ phase interaction continuum by the Walton Braun node. Select the node to access the properties and expert properties of the model.



Walton Braun Properties

These property settings are available for you to calibrate when comparing results in STAR-CCM+ to your experimental results.

Static Friction Coefficient	C_{fs} (see Eqn. (1597)).
Rolling Friction Coefficient	$c_{\text{rolling friction}}$ (see Eqn. (1589)).
Normal Restitution Coefficient	See Eqn. (1600).
Tangential Restitution Coefficient	See Eqn. (1600).
Stiffness Ratio	Controls the ratio of release and accumulation of energies during loading and/or unloading, therefore this property must be a non-negative, non-zero value. Setting this property to 0 means the material is perfectly plastic and it only deforms. The Walton Braun model is not suitable for simulating such physics.
Yield Strength Fraction	Also known as yield stress fraction (see Eqn. (1612)).

Walton Braun Expert Properties

Freeze Tangential Force	Controls whether to include tangential force in the calculations.	
	Activated	Does not include tangential force.
	Deactivated	Includes tangential force.
Freeze Torque	Controls whether to include torque in the calculations.	
	Activated	Does not include torque.
	Deactivated	Includes torque.

Selecting the Walton Braun Model

To select the Walton Braun model:

- In the *Phase Interaction Model Selection* dialog, select Walton Braun from the *Dem Contact Model* group box.



What is the Linear Spring Model?

The Linear Spring model calculates contact between particles or between particles and a wall. The relationship between particle compression and contact forces is based on the relationship found in linear springs. The model uses constants for normal and tangential spring stiffness that are either set by you directly or are based on properties of the particle: the radius, maximum overlap, and Young's modulus (see Eqn. (1605)).

This model is represented under the phase interactions node by the Linear Spring node. Select the node to set the properties.



Linear Spring Properties

Static Friction Coefficient	The friction coefficient μ. See Eqn. (1605).
Normal Restitution Coefficient	See Eqn. (1600).
Tangential Restitution Coefficient	See Eqn. (1600).
Normal Spring Stiffness Method	Specifies the method for assigning normal spring stiffness K_n to the bodies in the interaction. See Eqn. (1605). Constant: Stiffness is a constant set by the user. Particle Material Based: Stiffness is calculated based on the material composing the particle.
Tangential Spring Stiffness Method	Specifies the method for tangential normal spring stiffness K_t to the bodies in the interaction. See Eqn. (1605). Constant: Stiffness is a constant set by the user. Particle Material Based: Stiffness is calculated based on the material composing the particle.

Particle Material Based Properties

Maximum Overlap/Radius Ratio	A constant set by the user, based on the overlap distance and the particle radius. The default is 0.1.

What Is the Conduction Heat Transfer Model?

The Conduction Heat Transfer model calculates conduction heat transfer between two particles or a particle and a wall. This model is represented in the STAR-CCM+ phase interaction continuum by the Conduction Heat

Transfer node. This node has no properties.



The energy state of the continuous phase is not affected by particle-particle or particle-wall heat conduction unless you also activate the Two-Way Coupling Model.

When the Two-Way Coupling model is active in conjunction with Conduction Heat Transfer, the wall heat flux coefficients are modified according to Eqn's (1618) and (1619). See Modeling Two-Way Coupling and the Conduction Heat Transfer Model formulation section.

Selecting the Conduction Heat Transfer Model

The Conduction Heat Transfer Model is available for phase interactions involving two particles or a particle and a wall. Therefore, set the First phase property of the DEM Phase Interaction model to the DEM Lagrangian phase and the Second phase property to either the DEM Lagrangian phase or a wall.

As a prerequisite, select the following phase interaction model:

- In the *Phase Interaction Model Selection* dialog, select either the Hertz Mindlin or Walton Braun models.

Then, to select the Conduction Heat Transfer model:

- In the *Phase Interaction Model Selection* dialog, select Conduction Heat Transfer from the *Optional Models* group box.



What Is the Parallel Bonds Model?

The Parallel Bonds model introduces attractive inter-particle forces to the particle system. The bonded particle model uses the concept of a massless bar connecting a pair of bonded particles. The bar can transmit force and torque between particles and it is also subject to cracking under load.



The inter-particle force consists of the traditional grain-grain interaction (represented by the specified contact model, Hertz-Mindlin or Walton-Braun) and the parallel bond forces.

Potential applications of this particular model include:

- Pharmaceutical and biomedical—the tableting process, the breakage of particles, and deforming cells in blood flow
- Nuclear—checking the effects of certain obstructions to the flow of nuclear waste material; heat transfer studies
- Chemical and metallurgical—particle breakage and dust formation during handling of solid raw materials; also phase change
- Shipping—ice breakup, specifically ice breakup upon collision with dams or shorelines, as well as upon collision with shipping vessels
- Petroleum—hydrates and asphaltenes. This application can include specifying how particles bond with each other based on various parameters.

The Parallel Bonds model is represented in the STAR-CCM+ phase interaction continuum by the Parallel Bonds node. Select the node to access the properties and expert properties of the model.



Controlling Bond Formation

Parallel bond formation depends on the expert properties in the model, Bond Start Time and Bond End Time. When the simulation time is in the interval between these times, any particles that collide establishes a parallel bond in addition to the regular contact force being calculated.

Parallel Bonds Properties

Bond Start Time

Specifies the time at which the bonds begin to be created. The default is 0.

Bond End Time

Specifies the time at which the bond creation process stops. The default is 0.

Tensile Strength Method

Specifies the method for assigning tensile strength to the bond. The bond fails if tensile stress exceeds δ_{max} in Eqn. (1633).

- **No Failure Limit** — The method does not allow the bond to break.
- **Constant** — Bond strength is a constant set by the user.
- **Uniform Random** — The method assigns a strength to each bond, selected from a uniform distribution with upper and lower limits set by the user.
- **Normal Random** — The method assigns a strength to each bond, selected from a normal distribution with mean, standard deviation, and upper and lower limits set by the user.

Shear Strength Method

Specifies the method for assigning shear strength to the bond. The bond fails if shear stress exceeds τ_{max} in Eqn. (1634).

- **No Failure Limit** — The method does not allow the bond to break.
- **Constant** — Bond strength is a constant set by the user.

Uniform Random	The method assigns a strength to each bond, selected from a uniform distribution with upper and lower limits set by the user.
Normal Random	The method assigns a strength to each bond, selected from a normal distribution with mean, standard deviation, and upper and lower limits set by the user.

Parallel Bonds Expert Properties

Bond Radius Multiplier	Scales the radius of the particle bond. The default multiplier is 1.

Selecting the Parallel Bonds Model

The Parallel Bonds Model is only available for a phase interaction involving two particles. It is not available for use on a phase already using the Bonded Particles model, such as Particle Clumps. To select the Parallel Bonds model, set the First phase and Second phase properties of the DEM Phase Interaction model to DEM Lagrangian phases.

As a prerequisite, select the following physics, phase, and phase interaction models:

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM).
- In the *Phase Model Selection* dialog, select the DEM Particles model.
- In the *Phase Interaction Model Selection* dialog, select the DEM Phase Interaction and either the Hertz Mindlin or Walton Braun models.

Then, to select the Parallel Bonds model:

- Close the model selection window, select DEM Phase Interaction node, and in the properties window set the first and second phase to the DEM phase to be bonded.

- In the *Phase Interaction Model Selection* dialog, select Parallel Bonds from the *Optional Models* group box.



What Is the Linear Cohesion Model?

Cohesion modeling facilitates simulation of inter-molecular attraction forces (Van der Waals forces) between particle surfaces. For some classes of simulations, such as dry powders, these forces significantly affect the outcome and cannot be ignored. STAR-CCM+ implements a variant of the cohesion model that is known as the Linear Cohesion model, which allows you to choose between two major models by making an entry for F in the properties of the Linear Cohesion node (see formulation). The cohesion model acts in tandem with the contact model, so when particles attach both contact and cohesion work simultaneously until the particles detach again.

The Linear Cohesion model is represented in the STAR-CCM+ phase interaction continuum by the Linear Particles node. Select the node to access the properties of the model.



Linear Cohesion Properties

Work of Cohesion	W in Eqn. (1635).
Factor	The multiplication model blending factor F, with values varying between JKR and DMT models (see Eqn. (1635)): • For the JKR model [293], enter 1.5. • For the DMT model [285], enter 2.0.

Selecting the Linear Cohesion Model

As a prerequisite, select the following physics and phase interaction models:

- In the *Physics Model Selection* dialog, select the Discrete Element Model (DEM).
- In the *Phase Interaction Model Selection* dialog, select the DEM Phase Interaction and either the Hertz Mindlin or Walton Braun models.

Then, to select the Linear Cohesion model:

- In the *Phase Interaction Model Selection* dialog, select Linear Cohesion from the *Optional Models* group box.



What Is the Bonded Particles Model?

The bonded particle model is a combination of the Hertz Mindlin and Parallel Bonds model, and is activated by default with the Particle Clumps model. When the particles are attached, the Parallel Bonds model determines the connection and inter-particle forces. If the tensile or shear stresses between particles exceeds the maximum limits, the bond breaks and the particles separate. Thereafter, the Hertz Mindlin model determines the particle-particle interaction upon collision.

The Bonded Particles model is represented in the STAR-CCM+ phase interaction continuum by the Bonded Particles node. Select the node to access the properties of the model.



Bonded Particles Properties

Static Friction Coefficient	C_{fs} (see Eqn. (1597)).
Rolling Friction Coefficient	$c_{\text{rolling friction}}$ (see Eqn. (1589)).
Normal Restitution Coefficient	See Eqn. (1600).
Tangential Restitution Coefficient	See Eqn. (1600).
Maximum Tensile Stress	The maximum tensile stress the bond between particles can withstand before they detach. See Eqn. (1633).
Maximum Shear Stress	The maximum shear stress the bond between particles can withstand before they detach. See Eqn. (1634).

Bonded Particles Expert Properties

Bond Radius Multiplier	Scales the radius of the particle bond.

Selecting the Bonded Particles Model

Selecting the Particle Clumps phase model automatically selects the Bonded Particles model. This model is not available by any other method.

What Is the Rolling Resistance Model?

With the DEM model active in a physics continuum, rolling resistance is an Optional Model under Models in a Phase Interaction.



To use the rolling resistance model:

- Expand the Rolling Resistance node.
- Right-click on the Rolling Resistance Methods node and select Edit... to access the dialog for selecting the method.
- Select one of these methods from the Method drop-down menu: Force Proportional, Constant Torque, or Displacement Damping.



Force Proportional Properties

This method is the default version of Rolling Resistance. The Force Proportional method is represented in the STAR-CCM+ phase interaction continuum by the Force Proportional node.



Coefficient of rolling resistance	μ_r (see Force Proportional Method). It can take on any non-negative real value. The default value is 0.0010.

Constant Torque Properties

The Constant Torque method is represented in the STAR-CCM+ phase interaction continuum by the Constant Torque node.

Coefficient of rolling resistance	μ_r (see Eqn. (1636)). It can take on any non-negative real value. The default value is 0.1.
Moment limiter factor	Sets the limit used by the Use Limiter property. It can take on any non-negative real value. It operates only when Use Limiter is active. The default value is 0.08.
Use Limiter	Sets an upper limit on the angular momentum of the particles around the axes of moment. The value of the limit is set in the Moment Limiter Fraction property.

Activated	Activates moment limiting.
Deactivated	Leaves moment limiting inactive (the default).

Displacement Damping Properties

The Displacement Damping method is represented in the STAR-CCM+ phase interaction continuum by the Displacement Damping node.

Coefficient of rolling resistance	μ_r (see Eqn. (1637)). It can take on any non-negative real value. The default value is 0.1.
Theta Maximum	θ_r^{max}, the maximum angular displacement before rolling begins. It is available only when Use Damper is active. The default value is 0.035 radian.
Viscous Damping Factor	ν, sets the factor used by the Use Damper property. It can take on any value from 0 to 1. It is available only when Use Damper is active. The default value is 0.1.
Viscous Damping f multiplier	f, provides a multiplier for the Viscous Damping Factor property. It can take on any value from 0 to 1. It is available only when Use Damper is active. The default value is 1.
Use Damper	Sets an upper limit on the angular momenta of particles around their axes of moment. The value of the limit is set in the Moment Limiter Fraction property. Activated — Activates viscous damping. Deactivated — Leaves viscous damping inactive (the default).

Visualizing and Reporting Contact

In general, collisions involving DEM particles are not the same as those involving Lagrangian parcels. With DEM particles, there can be more than one collision between particles or between a particle and a wall, if particles are modeled as spherical composites. In the following illustration, arrows mark multiple collisions between a composite particle and the wall.



STAR-CCM+ provides discretized collision information for visualization and reporting, through the following field functions:

- Contact Damping Force
- Contact Elastic Force
- Contact Normal
- Contact Normal Overlap
- Contact Relative Velocity
- Contact Tangential Overlap

Modeling Artificial Viscosity

When two particles collide, or when a particle collides with a wall, the DEM module uses the impact velocity to predict the outcome. Impact velocities can increase to the same order of magnitude as the speed of sound in the particle material. When this occurs, the numerical models can over-predict the overlap between the colliding objects. The large overlap results in

inaccurate results or unphysical ones, such as particles passing through walls. You can avoid or reduce these problems by applying the artificial viscosity model.

With the DEM model active in a physics continuum, artificial viscosity is an Optional Model under Models in a Phase Interaction.

Note: The name refers only to the similarity between the formulation of the model and formulations for viscosity. It does not refer to any physical effect resembling real viscosity.

This model introduces a force during compression (while the particles are overlapping and their centers are moving nearer) based on the von Neumann-Richtmyer artificial viscosity [302]:

$$F_v = -f(Q_2|v_n| - Q_1 c)\frac{m_{eq}}{R_{eq}}v_n$$

where:

- F_v is the artificial viscosity force.
- v_n is the component of the velocity in the contact normal direction.
- m_{eq} is the equivalent (harmonic mean) mass of the particles.
- R_{eq} is the equivalent (harmonic mean) radius of the particles.
- c is the speed of sound in the material. It is estimated from the equivalent mass and radius, and from the contact stiffness for the specific contact model in use.
- Q_1 is the linear coefficient.
- Q_2 is the quadratic coefficient.

The factor f is a linear ramp function that is scaled to the particle sizes as:

$$f = \begin{Bmatrix} 0, & r/R_{eq} < p_{\text{lower}} \\ \dfrac{r/R_{eq} - p_{\text{lower}}}{p_{\text{upper}} - p_{\text{lower}}}, & p_{\text{upper}} \geq r/R_{eq} \geq p_{\text{lower}} \\ 1, & r/R_{eq} > p_{\text{upper}} \end{Bmatrix}$$

where:

- r is the penetration, the overlap distance between particles or between particle and wall.
- p_{upper} and p_{lower} are the upper and lower ramp limits. They are defined in terms of the ratio of the penetration r to the equivalent particle radius R_{eq}.

This ramp function f allows you to use artificial viscosity without affecting typical low penetration contacts.

Caution: The energy that is dissipated by artificial viscosity represents unresolved processes of damage, deformation, and entropy. Do not use the artificial viscosity model if you want to study these or similar collision effects in detail. Instead, use models that address those effects; for example, to study particle breakage, use the Bonded Particles model.

Artificial Viscosity Properties

Linear Coefficient	The linear coefficient Q_1, which must be nonnegative. The default value is 0.3.
Quadratic Coefficient	The quadratic coefficient Q_2, which must be nonnegative. The default value is 1.
Upper ramp limit	The upper ramp limit, p_{upper}. The value must lie between zero and one, and p_{upper} must be larger than p_{lower}. The default value is 0.2.
Lower ramp limit	The lower ramp limit, p_{lower}. The value must lie between zero and one, and p_{lower} must be smaller than p_{upper}. The default value is 0.1.

All these values are dimensionless.

The default value of Q_1, 0.3, is high, to enhance stopping at bounding surfaces. A value of 0.06 is more typical of shock problems.

The default value of Q_2, 1, is the value for an elastic fluid. Values typically range from 1.5 to 0.6.

The default values for p_{upper} and p_{lower}, 0.2 and 0.1, are high. Contact force models are valid only for values lower than 0.05, but the high values create a default situation where artificial viscosity comes into play even in borderline cases. Bring values down for stiff particles and up for soft ones.



Defining Injectors in the Discrete Element Method

An injector setup for DEM particles follows the Lagrangian methodology. Some of the Lagrangian injector types are not allowed for DEM particles. The main difference between Lagrangian and DEM particles with respect to the injector is the concept of a particle as opposed to a parcel. For a Lagrangian phase (Material and Massless particle types) the parcel can represent many particles at once or a fraction of a particle. For DEM particles, the parcel always represents a single particle, except in the case of the Coarse Grain Particle model.

When setting up the injector, ensure that the injector injects particles into the system stress-free and does not violate the basic small overlap assumptions of the contact model: particles cannot overlap with other particles or with the wall. Particles that overlap are eliminated during injection. This condition imposes a limitation for injector flow rates when injecting particles too close to each other in space or time. As a result, the injector flow rate has a finite maximum value that depends on particle diameter, injection velocity, and distribution, as well as the unsteady iteration time.

The following special injector types are available for DEM particles only:

- Lattice injector
- Random injector

Lattice Injector

When used with the Discrete Element Method (DEM), the lattice injector allows injection of particles into the system with specified, mesh-independent spatial distributions. You can use other injector types as usual. However, the main advantage when compared to the other injector types is the ability to introduce higher particle loading rapidly and make the particle distribution stable with respect to remeshing.

This injector places injection points at nodes of a regularly spaced lattice structure. Currently, the supported lattice types, which are selected in the Lattice property of the Lattice Structure injector condition, are:



Simple cubic Body-centered cubic Face-centered cubic

The Spacing injector value allows you to control the dimensions of the bounding box in the x, y, and z directions. The absolute placement of the lattice structure depends on a bounding box of all generating regions.

When you create a lattice injector, you can assign parts from the Geometry > Parts node or from the Regions node. If you assign from the Parts node, assign the part containing the lattice to the fluid region in question. The part does not need a mesh.

The only parts available from the Parts node are those from the *New Shape Part* menu: block, cone, cylinder, and sphere, shown combined in this example:



The following image displays particles in a body-centered cubic structure.



Random Injector

A random injector injects particles into a region with a random point distribution. This injector is suitable if you wish to inject particles in a less uniform, more arbitrary manner. The Particle Amount Specification

controls the amount of injected particles, using either the Number of Particles or the Porosity Limit method. Porosity Limit is the default.

If the Porosity Limit method is chosen, the injector works as follows:

1. Check if the Injection cycle limit has been reached. If not, continue; otherwise stop.
2. Randomly inject a number of particles. The Number of Seeds property sets this quantity.
3. Eliminate any overlapping particles.
4. Calculate the quantity of particles per unit volume -- also known as the porosity.
5. If the calculated porosity is less than the Porosity Limit, repeat steps 1 - 4. STAR-CCM+ continues injecting particles until the Porosity Limit is reached. The porosity limit for a fully packed volume is about 0.6.

If the Number of Particles method is chosen, the injector injects the specified number of non-overlapping particles, up to the limits of the volume.

To speed up the injection process, assign non-zero initial velocities. Particles with actual velocities pack faster and make room for new particles to be injected.

As the random injector injects particles anywhere within its assigned region, create a dedicated region surrounding the volume in which particles enter the simulation:

- Prepare the simulation geometry by defining a separate part within which the injected particles are randomly distributed.
- Assign this part to a separate region that is connected to the main fluid region using internal interfaces.
- Ensure that the injection region and the main fluid region belong to the same physics continuum.

When you create a random injector, you can assign parts from the Parts node under the Geometry node, as well as from the Regions node. However, the only parts available from the Parts node are those from the *New Shape Part* menu: block, cone, cylinder, and sphere.

In the example below, the injector injects particles into the cylindrical region, which is connected to the main volume using internal interfaces.



What Is the Discrete Element Method Solver?

The DEM solver controls the solution update for the Discrete Element Method (DEM). The DEM solver is represented in STAR-CCM+ by the DEM Solver node. Select the node to access the expert properties of the model.



Controlling the DEM Time-Step

By default, the duration of the Rayleigh wave (Eqn. (1590)) across the surface of the sphere limits the DEM time-step. Furthermore, a minimum of 10 time-steps are necessary in order to resolve particle collisions correctly (Eqn. (1591)). The distance traveled by particles within a single time-step is

constrained (Eqn. (1592)) to account for contact among particles and between particles and walls.

To account for the particle movements with precision, STAR-CCM+ uses substeps to assess physics such as flow and gravity (if applicable) as well as DEM. This results in dynamic time-step calculations by default.

To control the time-step, you can control the dynamic calculations of the solver in one of two ways:

- Change the Time Scale property to reduce the allowed fraction of the Rayleigh wave limit.



- Alternatively, enter the desired time-step in the Static Timestep property. This causes STAR-CCM+ to use the value of the time-step that you enter. Do not make this option your first choice, as it can cause a failure to converge.



To restore dynamic calculations, simply enter the default value of `0.0` in this property.

Discrete Element Method Solver Expert Properties

Grid Scale	Sets the scale of the grid cell to the particle diameter.
Reordering Frequency	Sets the number of iterations until STAR-CCM+ reorders the sequence in which DEM particles are accessed in memory. The default value is 0 (no reordering).
Time Scale	Sets the allowed fraction of the Rayleigh wave limit (Eqn. (1590)).
Verbose	Provides more output for this value while the simulation is running. This property can be useful for debugging problems when they occur. Activated — Activates verbose output. Deactivated — Leaves output at its default.
Smooth Large Particles	Smooths large particles on the Eulerian mesh for two-way coupling.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. Do not make activating this feature your first choice in most situations, as it is primarily a "debugging" tool and not supported for normal simulation work. Activated — The solver does not update the solution. Deactivated — The solver updates the solution.
Void Fraction Limit	Sets the void fraction limit for particle smoothing.
Smoothing Radius	Sets the smoothing radius (normalized to the particle radius).

Smoothing Sigma	Sets the smoothing distribution sigma (normalized to the particle radius).
Static Timestep	Instructs STAR-CCM+ to use the value of the time-step that you enter. This option should not be your first choice, as it can cause a failure to converge. The default setting of `0.0` leaves the dynamic calculation uninterrupted.
Skin	Sets a distance around each particle, the "skin." STAR-CCM+ runs contact detection only when the particle moves through a distance equal to the skin thickness. The default value is 0.

Best Practices

DEM gives high resolution, but at a high cost in computational power. Best practices for DEM mostly consist of finding ways to reduce the computational cost. Use the following guidelines:

- Maximize the number of processors. However, do not exceed about 1000 particles per processor. If you continue to add processors, speed goes down as processors spend more time communicating with each other. (So, for three million particles, use about three thousand processors, not more.)
- In composite particles, minimize the number of spheres in a particle.
- Maximize particle size or use the Coarse Grain Model to maximize parcel size.

  One DEM particle can represent several objective particles. Using larger particles reduces particle count and increases the DEM time-step, which reduces both the time that is needed and the number of processors needed. Consider the example of a simulation with particles in a rotating drum. Enlarging the particles from 2 mm to 3 mm produced the same results in half the time, using 5 processors instead of 18.

  Using the Coarse Grain model reduces time that is spent on contact physics while maintaining full accuracy for particle-fluid interactions. Coarse Grain is best used for fluidized beds applications, where accuracy of particle-fluid interactions is more important than accuracy of particle-particle interactions.

Note: The change in size in the example simulation also results in slightly different bulk properties (such as angle of repose) that depend on contact physics. You can often compensate for similar changes in bulk properties by slightly adjusting the parameters of the DEM Contact Model, such as friction.

Note: To maintain the accuracy of the DEM model, scaled-up particles or coarse-grain parcels must still be smaller than the smallest significant geometrical feature of the surroundings. In the example simulation, the parcels had to be smaller than the smallest important length scale in the geometry.

- Minimize Young's modulus. This practice increases the DEM time-step, resulting in faster computation. For applications with low-confinement shear flow regimes, you can often reduce Young's modulus by several orders of magnitude and still get accurate results. For example, in one simulation, going from 1000 MPa to 100 MPa reduced time to one third, using 5 processors instead of 18.
- Make the mesh as coarse as practical.
- Reduce the modeled physical time as much as practical:
  - When possible, optimize injector settings to introduce the maximum number of particles in the minimum physical time.
  - Sometimes, you can reduce time by increasing physical speeds of moving geometrical boundaries, such as using higher rpm on a rotating boundary.
  - If you only need to establish the direction of a trend, you can use a short physical time.
- To speed memory access, use the DEM solver property Reordering Frequency. On reordering, STAR-CCM+ resets the sequence in which DEM particles are accessed in memory. For quasi-stationary cases, it is enough to reorder once, after the particles are injected and settle. For rapidly moving particles, reorder at least once every 20 iterations. The benefit of reordering decreases with repetition.
- To reduce time that is spent running contact detection, use the DEM solver property Skin. The Skin property defines a volume around each DEM particle, a "skin." The solver runs contact detection only when the particle moves through a distance equal to the skin thickness. As a rule of thumb, set skin thickness to $10v_p t$ where v_p is a characteristic particle velocity and t is length of time-step.

Discrete Element Model Formulation

This section provides the following formulation of the DEM model:

- Momentum balance for DEM particles
- DEM particle times scale
- Contact force modeling
- Conduction heat transfer
- Particle bonding
- Linear cohesion
- DEM drag force
- DEM drag torque
- DEM lift force

A bibliography is also included.

Momentum Balance for Discrete Element Method Particle

The momentum balance for a DEM particle is derived from the momentum balance of the material particle (Eqn. (1331)). The DEM modeling introduces extra body force representing inter-particle interaction due to particle contacts with other particles and with mesh boundaries:

$$\mathbf{F}_b = \mathbf{F}_g + \mathbf{F}_u + \mathbf{F}_c \tag{1586}$$

where:

$$\mathbf{F}_c = \sum_{\text{neighbor particles}} \mathbf{F}_{\text{contact}} + \sum_{\text{neighbor boundaries}} \mathbf{F}_{\text{contact}} \tag{1587}$$

A user-specified contact model controls the exact formulation of contact force.

Besides the standard Lagrangian linear momentum equation, the DEM particle equations of motion incorporate angular momentum conservation equations:

$$\frac{d}{dt}\mathbf{L}_p = \frac{d}{dt}(\mathbf{I}_p\omega_p) = \sum_{\text{neighbor particles}} \mathbf{T}_{\text{contact}} + \sum_{\text{neighbor boundaries}} \mathbf{T}_{\text{contact}} \tag{1588}$$

where contact torque is computed as:

$$\mathbf{T}_c = \mathbf{r}_c \times \mathbf{F}_c - \mu_r |\mathbf{r}_c| |\mathbf{F}_c| \frac{\boldsymbol{\omega}_p}{|\boldsymbol{\omega}_p|} \tag{1589}$$

and where:

- $\mathbf{L}_p$ is the particle angular momentum.
- $\mathbf{I}_p$ is the particle moment of inertia.
- $\mathbf{T}_c$ represents the torque on an individual particle due to contact force acting on the particle at a point other than the particle center of gravity.
- $\mathbf{r}_c$ is a vector from the particle center of gravity to the contact point.
- $\mathbf{F}_c$ is the contact force acting on the particle.
- μ_r is the coefficient of rolling friction.
- $\boldsymbol{\omega}_p$ is the particle angular velocity.

The detailed torque formulation is supplied by the inter-particle contact model.

Discrete Element Method Particle Times Scale

The maximum time-step that is allowed for a DEM particle is constrained by the assumption that the force acting on a particle is only affected by the particle's immediate neighbors during a single time-step duration. The time-step is therefore limited by the time it takes the Rayleigh wave to propagate across the surface of the sphere to the opposite pole [292]:

$$\tau_1 = \pi \frac{R_{min}}{V_{Rayleigh}} \tag{1590}$$

where R_{min} is the minimal sphere radius.

The Rayleigh wave velocity depends on material properties, and the exact value is obtained as a solution to a secular equation [297]. STAR-CCM+ uses an approximation of the solution that provides sufficient precision without incurring high computational cost [294], [295].

Besides the wave propagation, more time-step limiting criteria are applied to moving particles. The duration of impact of two perfectly elastic spheres assuming the Hertz contact theory was derived by Timoshenko [300] to be:

$$\tau_2 = 2.94\left(\frac{5\sqrt{2}\pi\rho}{4}\frac{1-\nu^2}{E}\right)^{\frac{2}{5}} \frac{R}{\sqrt[5]{v_{impact}}} \tag{1591}$$

In order to resolve the collisions a minimum of 10 time-steps is required for collision resolution.

The final restriction on a DEM particle time-step is geometric. It is based on the assumption that particles must not move too far within the time-step. This condition prevents missing contacts between DEM particles, as well as particles and walls. Therefore each particle is constrained such that it takes at least 10 time-steps for the particle to move the full length of the radius.

$$\tau_3 = \left(\frac{R}{v_{particle}}\right) \tag{1592}$$

The final particle time-step is determined as a minimum of τ_1, τ_2, τ_3. In practice τ_1 is typically the limiting factor while τ_2 and τ_3 only constrain particles that are moving fast, or the Young's modulus of the material is set low in order to accelerate the simulation.

Contact Force Modeling

Contact force formulation in DEM is typically a variant of the spring-dashpot model. The spring generates repulsive force pushing particles apart and the dashpot represents viscous damping and allows simulation of collision types other than perfectly elastic. The forces at the point of contact are modeled as a pair of spring-dashpot oscillators. One oscillator represents the normal direction and the other the tangential direction of force with respect to the contact plane normal vector.

STAR-CCM+ provides two contact models:

- Hertz Mindlin
- Walton Braun

Hertz-Mindlin No-Slip Contact Model

The Hertz-Mindlin contact model is a variant of the non-linear spring-dashpot contact model that is based on the Hertz-Mindlin contact theory [292], [289]. The forces between two spheres, A and B, are described by the following set of equations.

$$\mathbf{F}_{\text{contact}} = \mathbf{F}_n + \mathbf{F}_t \tag{1593}$$

where $\mathbf{F}_n$ is the normal and $\mathbf{F}_t$ is a tangential force component.

The normal direction is defined by the following:

- Normal force:

$$\mathbf{F}_n = -K_n d_n - N_n v_n \tag{1594}$$

- Normal spring stiffness:

$$K_n = \frac{4}{3} E_{eq} \sqrt{d_n R_{eq}} \tag{1595}$$

- Normal damping:

$$N_n = \sqrt{(5 K_n M_{eq})}\, N_{n\ \mathrm{damp}} \tag{1596}$$

The tangential direction is defined by the following:

- The tangential force is defined by $-K_t d_t - N_t v_t$ if $|K_t d_t| < |K_n d_n| C_{fs}$, where C_{fs} is a static friction coefficient. Otherwise:

$$\mathbf{F}_t = \frac{|K_n d_n| C_{fs} d_t}{|d_t|} \tag{1597}$$

- Tangential spring stiffness:

$$K_t = 8 G_{eq} \sqrt{d_t R_{eq}} \tag{1598}$$

- Tangential damping:

$$N_t = \sqrt{(5 K_t M_{eq})}\, N_{t\ \mathrm{damp}} \tag{1599}$$

where:

The damping coefficient equals 1 if C_{rest} equals 0. Otherwise:

$$N_{\mathrm{damp}} = \frac{-\ln(C_{n\ \mathrm{rest}})}{\sqrt{\pi^2 + \ln(C_{n\ \mathrm{rest}})^2}} \tag{1600}$$

The equivalent radius is defined as:

$$R_{eq} = \frac{1}{\frac{1}{R_A} + \frac{1}{R_B}} \tag{1601}$$

The equivalent particle mass is:

$$M_{eq} = \frac{1}{\frac{1}{M_A} + \frac{1}{M_B}} \tag{1602}$$

The equivalent Young's modulus is expressed as:

$$E_{eq} = \frac{1}{\frac{1-\nu_A^2}{E_A} + \frac{1-\nu_B^2}{E_B}} \tag{1603}$$

The equivalent shear modulus is:

$$G_{eq} = \frac{1}{\frac{2(2-\nu_A)(1+\nu_A)}{E_A} + \frac{2(2-\nu_B)(1+\nu_B)}{E_B}} \tag{1604}$$

where:

- M_A and M_B are masses of sphere A and sphere B.
- d_n and d_t are overlaps in the normal and tangential directions at the contact point.
- R_A and R_B are the radii of the spheres.
- E_A and E_B are the Young's modulus of the spheres.
- ν_A and ν_B are the Poisson coefficients.
- ν_n and ν_t are the normal and tangential velocity components of the relative sphere surface velocity at the contact point.

For particle-wall collisions, the formulas stay the same, but the wall radius and mass are assumed to be $R_{wall} = \infty$ and $M_{wall} = \infty$, so the equivalent radius is reduced to $R_{eq} = R_{particle}$ and equivalent mass to $M_{wall} = M_{particle}$.

Di Renzo and Di Maio [289]proposed several formulations for detailed treatment of tangential micro-slip and tangential force computation. STAR-CCM+ uses the formulation that Tsuji proposed [301], where the tangential force is assumed to be non-linear but the detail micro-slip tracking is replaced by analytical expression. The resulting code is computationally efficient while still matching experimental data.

Linear Spring Contact Model

This model is based on the work of Cundall and Strack [287], using a parallel linear spring-dashpot model for the normal force and a parallel linear spring-dashpot in series with a slider for the tangential force. In both, the spring accounts for the elastic part of the response and the dashpot accounts for dissipation from plastic deformation.



The normal and tangential forces f_n and f_t are:

$$f_n = -K_n \cdot \delta_n - N_n v_n$$

$$f_t = \begin{cases} -K_t \cdot \delta_t - N_t v_t & \text{if } (|K_t \cdot \delta_t| < |K_n \cdot \delta_n| C_{fs}) \\ \dfrac{|K_n \cdot \delta_n| C_{fs} \delta_t}{|\delta_t|} & \text{otherwise} \end{cases} \tag{1605}$$

where:

- K_n is the normal spring constant.
- δ_n is the normal deformation distance.
- $N_n = 2N_{n\ \mathrm{damp}}\sqrt{K_n M_{eq}}$ where:
  - $N_{n\ \mathrm{damp}}$ is the normal damping coefficient as in Eqn. (1600).
  - M_{eq} is the equivalent particle mass as in Eqn. (1602).
- C_{fs} is the coefficient of static friction
- K_t is the tangential spring constant.
- δ_t is the tangential deformation distance.

- $N_t = 2N_{t\,\mathrm{damp}}\sqrt{K_t M_{eq}}$ where $N_{t\,\mathrm{damp}}$ is the tangential damping coefficient.

When the Linear Spring model uses the Particle Material Based method, the spring stiffness is estimated based on the assumption that the linear model is a linearized Hertz model, in which the maximum normal contact force is:

$$K_n \delta_{max} = \frac{4}{3} E_{eq} \sqrt{R_{eq} \delta_{max}} \delta_{max}$$

where:

- δ_{max} is the maximum overlap.
- E_{eq} is the equivalent Young's modulus defined in Eqn. (1603).
- R_{eq} is the equivalent radius defined in Eqn. (1601).

To calculate the spring constants, let:

$$\lambda_{max} = \frac{\delta_{max}}{R_{eq}}$$

This gives the normal spring constant K_n as:

$$K_n = \frac{4}{3} \sqrt{\lambda_{max}} E_{eq} R_{eq} \tag{1606}$$

The tangential spring constant K_t is:

$$K_t = 8 \sqrt{\lambda_{max}} G_{eq} R_{eq} \tag{1607}$$

where G_{eq} is the equivalent shear modulus defined in Eqn. (1604).

Walton Braun Hysteretic Contact Model

The Walton Braun hysteretic plastic-elastic contact system is characterized by the following relations:

$$\mathbf{F}_n = -K_1 \delta \mathbf{n} \tag{1608}$$

when the contact is loading, and:

$$\mathbf{F}_n = -K_2 (\delta - \delta_{\mathrm{deformation}}) \mathbf{n} \tag{1609}$$

during unloading, where:

$$K_1 = 1.6 \pi R Y_0 \tag{1610}$$

$$K_2 = \frac{1}{E_f} K_1 \tag{1611}$$

$$Y_0 = E \cdot \text{YieldStressFraction} \tag{1612}$$

$$\delta_{\text{deformation}} = \delta_{\max}(1 - E_f) \tag{1613}$$

E is the material Young modulus, δ is the contact normal overlap, $\mathbf{n}$ is the contact normal vector, and Yield Stress Fraction is a user-controllable model property defining the onset of plastic deformation. E_f is a user-controllable energy fraction that defines the amount of energy that is recovered during the unloading and $\delta_{\max}$ is the maximum overlap that is reached during the contact loading.

Contact Cycle

When particle contact is formed, the state of the contact is tracked at each time-step as follows:

- Contact is considered loading when $\mathbf{v}_r \cdot \mathbf{n} \geq 0$ and $\delta > \delta_{\max}$, where $\mathbf{v}_r$ is relative surface velocity.

  When contact is loading, the force is calculated using K_1 stiffness and $\delta_{\max}$ is updated to a value of δ.
- If the contact is not loading, it is either unloading ($\mathbf{v}_r \cdot \mathbf{n} < 0$) or reloading ($\delta \leq \delta_{\max}$). During unloading, the K_2 stiffness is used to compute the normal force and the value of $\delta_{\max}$ is not updated.
- If the particles separate completely, the contact tracking information is removed and subsequent contact is treated as new.

Geometric Interpretation

The loading/unloading cycle can be geometrically represented as follows:



Conduction Heat Transfer Model

Two bodies are assumed to exchange heat through conduction when they make physical contact. The two bodies can be two particles or a particle and a wall. There are four cases:

- Particle-particle (see [286]):

  The heat exchange q_{ij} from particle i to particle j is:

  $$q_{ij} = 4r_c k(T_j - T_i) \tag{1614}$$

  where:

  - r_c is the contact area radius.
  - k is the equivalent thermal conductivity of the two particles, where

    $$\frac{1}{k} = \frac{1}{k_i} + \frac{1}{k_j}$$

    where k_i and k_j are the thermal conductivities for particles i and j.
  - T_i and T_j are the temperatures of particle i and particle j.
- Particle-wall, adiabatic wall: heat exchange is zero.

- Particle-wall, fixed wall temperature:

  The heat exchange q_{iw} from the wall to particle i is:

$$q_{iw} = k(T_w - T_i) \tag{1615}$$

  where:

  - k is the equivalent thermal conductivity, equal to k_i.
  - T_i and T_w are the temperatures of particle i and the wall.
- Particle-wall, fixed wall heat flux:

  The heat exchange q_{iw} from the wall to particle i is:

$$q_{iw} = A_c q_w \tag{1616}$$

  where:

  - A_c is the contact area between the particle and the wall.
  - q_w is the heat flux density of the wall.

Two-Way Coupling Model

If the simulation includes a continuous phase, its energy state is unaffected by particle-particle and particle-wall heat conduction unless you activate the Two-Way Coupling model. With this model active, the Lagrangian energy source term is added to the energy equation of the continuous phase.

The linearized wall heat flux, q_w, is expressed as:

$$q_w = a + bT_c + cT_w + dT_w{}^4 \tag{1617}$$

where:

- T_c and T_w are the cell and wall temperatures, respectively.
- a, b, c, and d are wall heat flux coefficients.

The contribution of DEM particle-wall conductive heat flux to the wall heat flux can be made by modifying wall heat flux coefficients a and c as follows:

For fixed wall heat flux conditions, only a needs to be modified:

$$a = \sum \frac{A_c}{A} q_w \tag{1618}$$

where:

- A_c is the particle-wall contact area.
- A the boundary face area.

For fixed wall temperature conditions, a and c are modified as:

$$a = \sum 4 r_c k T_p$$
$$c = \sum 4 r_c k \tag{1619}$$

where:

- k is the particle thermal conductivity.
- r_c the contact radius.
- T_p the particle temperature.

Particle Bonding Model

The parallel bond forces are represented in the following governing equations [296].

Force and torque on a particle due to parallel bonds is:

$$\mathbf{F}_i = F_n \mathbf{n}_i + F_s \mathbf{t}_i \tag{1620}$$

$$\mathbf{M}_i = M_n \mathbf{n}_i + M_s \mathbf{t}_i \tag{1621}$$

where the terms F_n, M_n denote the normal and F_s, M_s denote the shear components of forces and torques with respect to the contact plane, and:

$$\Delta F_n = k_n A \Delta U_n \tag{1622}$$

$$\Delta F_s = -k_s A \Delta U_s \tag{1623}$$

$$\Delta M_n = -k_s J \Delta \Omega_n \tag{1624}$$

$$\Delta M_s = -k_n I \Delta \Omega_s \tag{1625}$$

where:

$$A = \pi R^2 \tag{1626}$$

$$I = \frac{1}{4}\pi R^4 \tag{1627}$$

$$J = \frac{1}{2}\pi R^4 \tag{1628}$$

$$k_n = \frac{E}{R_a + R_b} \tag{1629}$$

$$k_s = 6\left(\frac{G}{R_a + R_b}\right) \tag{1630}$$

$$\Delta U = V_{\text{relative}} \Delta t \tag{1631}$$

where ΔU_n is the normal component of ΔU, and ΔU_s is the tangential component of ΔU.

$$\Delta \Omega = \omega_{\text{relative}} \Delta t \tag{1632}$$

where $\Delta \Omega_n$ is the normal component of $\Delta \Omega$, and $\Delta \Omega_s$ is the tangential component of $\Delta \Omega$.

The parallel bond can be treated as brittle and disintegrative when maximum shear and tensile stresses exceed prescribed limit values. The stress limit values are calculated based on beam theory to be:

$$\delta_{\max} = \frac{-F_n}{A} + \frac{|M_s| R}{I} \tag{1633}$$

$$\tau_{\max} = \frac{|F_s|}{A} + \frac{|M_n| R}{J} \tag{1634}$$

If either value exceeds the predefined maximum, the bond is removed and the forces are further calculated only based on the granular interaction of particles.

Linear Cohesion Model

This variant of cohesion modeling can use either of two models that are described in published literature: the Johnson-Kendall-Roberts (JKR) model [293] or the Derjaguin-Muller-Toporov (DMT) model [285].

Cohesion force between two spherical particles is expressed in both cases as:

$$F_{cohesion} = R_{min} W \pi F \tag{1635}$$

where R_{min} is the minimal radius of surfaces in contact, W is the work of cohesion $[J/m^2]$, and F is a multiplication model blending factor with values as follows:

- $\frac{3}{2}$ for the JKR model
- 2 for the DMT model

The principal difference between JKR and DMT model formulation is the surface area where the cohesion force is acting. In the JKR model, this area is limited to direct contact, while in the DMT model it also incorporates the "neck" area in immediate contact proximity.

Rolling Resistance Model

This section covers the rolling resistance methods applicable to DEM particles.

- Force Proportional Method
- Constant Torque Method
- Displacement Damping Method

If you are using the Hertz-Mindlin, Walton-Braun, or Bonded Particle contact models together with the Rolling Resistance model, the Rolling Resistance model applies a rolling friction coefficient to the contact models.

Force Proportional Method

This method only supplies a coefficient of rolling resistance, μ_r, also called the coefficient of rolling friction. It can take on any positive real value. The default value is 0.001.

Constant Torque Method

The Constant Torque method applies a torque to resist rolling, calculated as:

$$M_r = \frac{-\omega_{rel}}{|\omega_{rel}|}\mu_r R_{eq} F_n \tag{1636}$$

where:

- ω_{rel} is the relative rotation, given as $\omega_{rel} = \omega_i - \omega_j$, where:
  - ω_i is the rotation of the particle for which the resistance moment M_r is being calculated.
  - ω_j is the rotation of the other particle in the contact.
- μ_r is the coefficient of rolling friction, also called coefficient of rolling resistance. The default value is 0.1.
- R_{eq} is the particle equivalent radius, that is, the harmonic mean of the radii of the two particles.
- F_n is the magnitude of the contact force.

This method also uses a "Moment Limiter." This function limits the size of the angular momentum of the particle, measured around the axis of moment. This maximum is 0.08 of the moment that negates the lesser of the two rotations ω_i and ω_j in the next time-step.

This limit can control specious rotations in particle compacts. These rotations can arise from a large number of contacts or from high compact pressure.

Displacement Damping Method

The Displacement Damping method combines a viscous damping resistance with a spring-like resisting torque generated by elastic energy stored in the particle as it tries to turn around a finite contact area.

The spring resistance torque at time t is:

$$M^k{}_{r,\,t+\delta t} = M^k{}_{r,\,t} - k_r \omega_{rel} \Delta t \tag{1637}$$

where $\left|M^k{}_{r,\,t}\right| \le k_r \theta_r{}^{max}$ and $k_r = (\mu_r R_{eq} F_n)/\theta_r{}^{max}$, and where:

- $\theta_r{}^{max}$ is the maximum angular displacement before rolling.
- ω_{rel} is the relative rotation, given as $\omega_{rel} = \omega_i - \omega_j$, where:
  - ω_i is the rotation of the particle for which the resistance moment M_r is being calculated.
  - ω_j is the rotation of the other particle in the contact.
- μ_r is the coefficient of rolling friction, also called coefficient of rolling resistance. The default value is 0.1.
- R_{eq} is the particle equivalent radius, that is, the harmonic mean of the radii of the two particles.
- F_n is the magnitude of the contact force.

You can activate the viscous damping torque or leave it off. When active, it is:

$$M_r{}^d = \begin{Bmatrix} -\nu C^{\text{crit}} \omega_{\text{rel}} & \text{if } (\left|M^k{}_{r,\,t}\right| \le k_r \theta_r{}^{max}) \\ -f \nu C^{\text{crit}} \omega_{\text{rel}} & \text{if } (\left|M^k{}_{r,\,t}\right| > k_r \theta_r{}^{max}) \end{Bmatrix} \tag{1638}$$

where:

- ν is the damping factor parameter.
- f is a damping factor multiplier
- C^{crit} is there critical viscous damping constant, computed as:

  $$C_r{}^{\text{crit}} = 2\sqrt{I_r k_r}$$

  where I_r is the equivalent moment of inertia for the relative rotational vibration mode about the contact point.

Discrete Element Method Drag Coefficient

The following section covers the drag force coefficients available with the DEM drag force model.

- Di Felice
- Gidaspow
- Haider and Levenspiel

Di Felice Drag Coefficient Method

This drag method introduces an extra term in the fluid drag force expression (Eqn. (1332)) to take into account the effect of the presence of other particles around a particle. The Di Felice drag coefficient is given as:

$$C_d = \left[0.63 + \frac{4.80}{\sqrt{\mathrm{Re}_p}}\right]^2 \varepsilon_i^{-\xi} \tag{1639}$$

where Re_p defines the particle Reynolds Number, ε_i is void fraction around a particle, and:

$$\xi = 3.7 - 0.65\exp[-0.5(1.5 - \log_{10}\mathrm{Re}_p)^2] \tag{1640}$$

The term $\varepsilon_i^{-\xi}$ takes into account the effect of enhanced drag on a particle, due to the presence of other particles around it.

Gidaspow Drag Coefficient Method

STAR-CCM+ incorporates the equations in drag coefficient form as follows:

$$C_d = \frac{4}{3}\left(150\frac{1-\nu_f}{\nu_f \Re_p} + 1.75\right) \qquad \text{if } \nu_f < \nu_{min} \tag{1641}$$

Otherwise:

$$C_d = \frac{(24 + 3.6 \bullet \Re_p^{0.687})}{\Re_p}\nu_f^{-3.65} \tag{1642}$$

where ν_f is the void fraction, ν_{min} is the cutoff void fraction (user-defined, with a default setting of 0.8), and $\Re_p$ is the particle Reynolds number. The exponent of -3.65 can also be user-defined in the Exponent property of the Gidaspow node.

Haider and Levenspiel Drag Coefficient Method

The drag coefficient C_d depends on the particle Reynolds number and the particle sphericity:

$$C_d = \frac{24}{\mathrm{Re}_p}(1 + A\mathrm{Re}_p^B) + \frac{C}{\left(1 + \frac{D}{\mathrm{Re}_p}\right)} \tag{1643}$$

and:

$$\begin{aligned} A &= 8.1716e^{-4.0665\phi} \\ B &= 0.0964 + 0.5565\phi \\ C &= 73.690e^{-5.0746\phi} \\ D &= 5.3780e^{6.2122\phi} \end{aligned} \tag{1644}$$

where ϕ is the particle sphericity and Re_p is the particle Reynolds number.

Drag Torque Model

Drag torque reduces the difference in the rotational velocities between a particle and the fluid in which it is immersed. The drag is a torque $\boldsymbol{T}$ applied to a DEM particle (Sommerfeld [299]):

$$\mathbf{T} = \frac{\rho}{2}\left(\frac{D}{2}\right)^5 C_R|\boldsymbol{\Omega}|\boldsymbol{\Omega} \tag{1645}$$

where:

- ρ is the fluid density
- D is the particle diameter.
- C_R is the rotational drag coefficient.
- $\boldsymbol{\Omega}$ is the relative angular velocity of the particle to the fluid, or slip-rotation, given by:

$$\boldsymbol{\Omega} = \frac{1}{2}\nabla \times \mathbf{u} - \boldsymbol{\omega}_p \tag{1646}$$

  where:

  - $\mathbf{u}$ is the fluid velocity
  - $\boldsymbol{\omega}_p$ is the angular velocity of the particle.

The rotational drag coefficient C_R used here is defined as [299]:

$$C_R = \begin{pmatrix} \dfrac{12.9}{\mathrm{Re}_R^{0.5}} + \dfrac{128.4}{\mathrm{Re}_R}, & 32 \le \mathrm{Re}_R < 1000 \\ \dfrac{64\pi}{\mathrm{Re}_R}, & \mathrm{Re}_R < 32 \end{pmatrix} \tag{1647}$$

Re_R is the rotational Reynolds number, defined as:

$$\mathrm{Re}_R = \frac{\rho D^2 |\mathbf{\Omega}|}{\mu} \tag{1648}$$

where μ is the dynamic viscosity.

Although Sommerfeld restricts the rotational Reynolds number to values below 1000, the model places no upper limit on it.

Discrete Element Method Lift Forces

Lift forces in DEM simulations can arise from particle spin, particle shear, or both. In this context, lift forces are taken to mean forces normal to the particle velocity; they are not necessarily forces in the upward direction.

Particle Spin Lift Force

This force applies to a spinning particle moving relative to a fluid. The force is given by:

$$\mathbf{F}_{LR} = \frac{\rho\pi}{8} D^2 C_{LR} |\mathbf{u}_s| \frac{\mathbf{\Omega} \times \mathbf{u}_s}{|\mathbf{\Omega}|} \tag{1649}$$

where:

- ρ is the fluid density.
- D is the particle diameter.
- C_{LR} is the coefficient of rotational lift.
- $\mathbf{u}_S$ is the fluid velocity relative to the moving particle, known as the slip velocity.
- $\mathbf{\Omega}$ is the relative angular velocity of the particle to the fluid, known as the slip rotation.

Sommerfeld [299] gives the coefficient of rotational lift C_{LR} as:

$$C_{LR} = 0.45 + \left(\frac{\mathrm{Re}_R}{\mathrm{Re}_p} - 0.45 \right) e^{-0.5684 \mathrm{Re}_R^{0.4} \mathrm{Re}_p^{0.3}} \tag{1650}$$

where:

- Re_R is the rotational Reynolds number.
- Re_p is the particle Reynolds number.

Sommerfeld reports this formula for C_{LR} as valid for particle Reynolds numbers below 140.

You can also define C_{LR} for yourself as a field function.

The rotational Reynolds number Re_R is given as:

$$\mathrm{Re}_R = \frac{\rho D^2 |\mathbf{\Omega}|}{\mu} \tag{1651}$$

where:

- ρ is the fluid density.
- D is the particle diameter.
- $\mathbf{\Omega}$ is the relative angular velocity of the particle to the fluid, known as the slip rotation.
- μ is the dynamic viscosity of the fluid.

The slip rotation $\mathbf{\Omega}$ is given by:

$$\mathbf{\Omega} = \frac{1}{2}\nabla \times \mathbf{u} - \boldsymbol{\omega}_p \tag{1652}$$

where:

- $\mathbf{u}$ is the fluid velocity.
- $\boldsymbol{\omega}_p$ is the angular velocity of the particle.

Particle Shear Lift Force

This force applies to a particle moving relative to a fluid where there is a velocity gradient in the fluid orthogonal to the relative motion.

Saffman [298] gives the shear lift force as:

$$\mathbf{F}_{LS} = 1.615 D^2 (\rho\mu)^{0.5} \left|\frac{\delta \mathbf{u}}{\delta y}\right| \mathbf{u}_S \tag{1653}$$

where:

- D is the particle diameter.
- ρ is the fluid density.
- μ is the dynamic viscosity of the fluid.
- $\mathbf{u}$ is the fluid velocity.

- $\mathbf{u}_S$ is the fluid velocity relative to the moving particle, known as the slip velocity.
- The y direction is the direction of the velocity gradient.

The three-dimensional version of this equation is:

$$\mathbf{F}_{LS} = C_{LS}\frac{\rho\pi}{8}D^3(\mathbf{u}_S \times \boldsymbol{\omega}) \tag{1654}$$

where:

- $\boldsymbol{\omega} = \nabla \times \mathbf{u}$, the curl of the fluid velocity.
- C_{LS} is the shear lift coefficient.

You can define C_{LS} as a field function or use either of two published definitions:

- Saffman [298] provides a definition that recovers the original Saffman asymptotic solution for low Reynolds numbers:

$$C_{LS} = \frac{4.1126}{\mathrm{Re}_S^{0.5}} \tag{1655}$$

- Sommerfeld [299] provides a definition for a broader range of Reynolds numbers:

$$C_{LS} = \frac{4.1126}{\mathrm{Re}_S^{0.5}} f(\mathrm{Re}_p, \mathrm{Re}_S) \tag{1656}$$

where:

$$f(\mathrm{Re}_p, \mathrm{Re}_S) = \begin{cases} (1 - 0.3314\beta^{0.5})e^{-0.1\mathrm{Re}_p} + 0.3314\beta^{0.5} & , (\mathrm{Re}_p \le 40) \\ 0.0524(\beta\mathrm{Re}_p)^{0.5} & , (\mathrm{Re}_p > 40) \end{cases} \tag{1657}$$

and

$$\beta = 0.5\frac{\mathrm{Re}_S}{\mathrm{Re}_p} \tag{1658}$$

where:

- Re_S is the Reynolds number for shear flow, $\mathrm{Re}_S = \frac{\rho D^2|\boldsymbol{\omega}|}{\mu}$.
- Re_p is the particle Reynolds number.

Discrete Element Method Bibliography

[285] Butt, H-J., Graf, K. and Kappl, M. 2006. *Physics and Chemistry of Interfaces*, Wiley-VCH.

[286] Cheng, G. J., Yu, A. B. and Zulli, P. 1999. “Evaluation of effective thermal conductivity from the structure of packed bed”, *Chemical Engineering Science*, 54, pp. 4199-4209.

[287] Cundall, P. A., Strack, O. D. L. 1979. “A discrete numerical model for granular assemblies”, *Geotechnique*, **29**, pp. 47-65.

[288] Di Felice, R. 1994. “The voidage function for fluid-particle interaction systems”, *Int. J. Multiphase Flow*, **20**, pp. 153-159.

[289] Di Renzo, Alberto and Di Maio, Francesco Paolo. 2004. “Comparison of contact-force models for the simulation of collisions in DEM-based granular flow codes”, *Chemical Engineering Science* **59** pp. 525-541.

[290] Gidaspow, D. 1994. *Multiphase Flow and Fluidization-Continuum and Kinetic Theory Descriptions*, Academic Press.

[291] Haider, A. and Levenspiel, O. 1989. “Drag coefficient and terminal velocity of spherical and non-spherical particles”, *Powder Technology*, **77**, pp. 143-152.

[292] Johnson, K.L. 1985. *Contact Mechanics*, Cambridge University Press, Cambridge.

[293] Johnson, K.L. 1987. *Contact Mechanics*, Cambridge University Press, Cambridge.

[294] Pham, Chi Vinh, Ogden, R.W. 2004. “On formulas for the Rayleigh wave speed”, *Wave Motion*, **39** pp. 191-197.

[295] Pichugin, A.V. 2008. “Approximation of the Rayleigh wave speed”.

[296] Potyondy, D.O, and Cundall, P.A. 2004. “A bonded-particle model for rock”, *Int. J. Rock Mechanics & Mining Sciences* **41** pp. 1329–1364.

[297] Rayleigh, L. 1885. “On waves propagated along the plane surface of an elastic solid”, *Proc. Lond. Math. Soc.* **17**.

[298] Saffman, P.G. 1965. "The lift on a small sphere in a slow shear flow," *J. Fluid Mech.* **22** 385-400. (See also Saffman, P.G. 1968. Corrigendum to: “The lift on a small sphere in a slow shear flow,“ *J. Fluid Mech.* **31** 624.)

[299] Sommerfeld, M. 2000. “Theoretical and Experimental Modelling of Particulate Flows,” *Technical Report Lecture Series 2000-06*, von Karman Institute for Fluid Dynamics pp 20-23.

[300] Timoshenko, S. 1951. *Theory of Elasticity*, McGraw-Hill Book Company.

[301] Tsuji, Y., Tanaka, T., Ishida, T. 1992. “Lagrangian numerical simulation of plug flow of cohesionless particles in horizontal pipe.” *Powder Technology* **71**:3 pp. 239-250.

[302] Von Neumann, J., and Richtmyer, R. D. 1950. "A method for the numerical calculation of hydrodynamic shocks", *J. Appl. Phys.* 21, 232–237.

[303] Walton, O.R. 1993. “Numerical simulation of inclined chute flows of monodisperse inelastic, frictional spheres”, *Mech. Mater.* **16** 239–247.

[304] Walton, O.R. 1993. “Numerical simulation of inelastic, frictional particle-particle interactions”, in *Particulate Two-Phase Flow*, M.C. Roco, Ed., Butterworth–Heinemann, Stoneham, MA, pp. 884–911.

[305] Walton, O.R. and Braun, R.L. 1986. “Viscosity, granular-temperature, and stress calculations for shearing assemblies of inelastic, frictional disks”, *J. Rheology* **30**:5, pp. 949–980.

Discrete Element Method Field Functions Reference

The following primitive field functions are available in the simulation when using the DEM model. These objects all have the same base properties.

Contact Damping Force	The non-elastic component of the contact force between two particles, that is, the contact force minus the elastic component. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.
Contact Elastic Force	The elastic component of the contact force between two particles, that is, the contact force minus the damping component. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.
Contact Force	The total contact force between two particles, that is, the damping force plus the elastic force. See $\mathbf{F}_c$ in Eqn. (1587).
Contact Normal	A unit length vector perpendicular (orthogonal) to the contact plane. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.
Contact Normal Overlap	The contact normal vector that is multiplied by the overlap, giving the direction and magnitude of the particle overlap, normal to the contact plane. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.

Contact Overlap	The amount of overlap between two particles or a particle and wall boundary upon contact. This overlap occurs as the particles are modeled using soft-particle formulation, which permits a small overlap. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.
Contact Relative Velocity	The relative overlap between two contacting particles. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.
Contact Tangential Overlap	The magnitude and direction of the particle overlap, tangential to the contact. To access this field function, select the DEM Particles, DEM Phase Interaction, and either the Hertz Mindlin or Walton Braun models.
Contact Torque	$\mathbf{T}_{\text{contact}}$ in Eqn. (1589).
Compression Force	Sum of values for $\mathbf{F}_c$.
Effective Particle Volume	The volume of a particle minus the overlapping volume beyond the plane(s) of intersection with any other particle(s).



Parcel DEM Timestep	See Eqns. (1590)-(1592).
Parcel Halo Count	Used in debugging.
Parcel Halo Marker	Used in debugging.
Particle Angular Momentum	$\mathbf{L}_p$ in Eqn. (1588).
Particle Angular Velocity	ω_p in Eqn. (1588).
Particle Moment Inertia	$\mathbf{I}_p$ in Eqn. (1588). This quantity can be used for validation.
Particle Rotation Kinetic Energy	Allows post-processing of the kinetic energy of the rotating particle.
Particle Shape Index	Allows post-processing by shape.
Particle Size	Sets the particle display size in the scene. This scalar would typically be used with the Point Size Scale Field on a scalar displayer to provide plots of particles that are scaled according to their physical size.
Particle Volume Overlap	How much of the particle is obscured by the contact area.
Void Fraction	The fraction of the cell volume available to the fluid, unoccupied by a solid or dispersed phase.

Migrating from EDEM

To convert an EDEM simulation to run under STAR-CCM+, remember:

- Plots of time-dependent quantities should be set before simulation in STAR-CCM+ and after simulation in EDEM.
- Geometry meshing requirements are different in STAR-CCM+ compared to EDEM (concepts of volume mesh, regions).
- Cell size should be at least four times particle size, and as much larger as practical.
- Optimal number of CPUs: STAR-CCM+ readily scales up to at least 64 CPUs. These CPUs can be distributed through a cluster.
- The Hertz-Mindlin model in STAR-CCM+ uses two coefficients of restitution, one for energy dissipation due to damping of normal component of contact force, another for energy dissipation due to damping of tangential component of contact force. The Hertz-Mindlin model in EDEM uses only one parameter for coefficient of restitution.
- Differences between EDEM and STAR-CCM+ terminology include the following:

EDEM	**STAR-CCM+**
Write-out interval	Auto-save time
Factory	Injector
Rayleigh timestep	Time scaling factor
Simulated time	Time stopping criteria parameter

Using the Eulerian Multiphase Model

For those multiphase models that require multiple Eulerian phases, the Eulerian Multiphase model is used to create and manage those phases. The corresponding node (Eulerian Multiphase in the following screenshot), which is subject to common menu items, is activated as soon as the model is selected. Right-click this node and choose New from the pop-up menu to create a Eulerian phase.



The Multiphase Segregated Flow, Eulerian Multiphase Mixture, and Volume of Fluid (VOF) models require you to select the Eulerian Multiphase model before they become available. The process of defining new phases is common to all three models, as is the selection of phase materials, viscous regime and phase equations of state. The relevant field functions are made available for each phase, such as Mass Flux, Mass Imbalance, Velocity, and Vorticity.

While the concept of volume fraction is different for these models, how you set the volume fraction for boundaries and initial conditions is the same.

Eulerian Phase Manager Pop-Up Menu

This pop-up menu item is available when you right-click the Eulerian Phases manager node.

New	Creates a Eulerian phase with a default name with a Models sub-node.

Selecting the Eulerian Multiphase Model

To solve problems using the Eulerian Multiphase model, first select the Eulerian Multiphase material model in the *Physics Model Selection* dialog. The Multiphase Interaction model is selected automatically.

Choose Eulerian Multiphase Mixture, Multiphase Segregated Flow, or Volume of Fluid (VOF) from the *Eulerian Multiphase Model* group box. The Multiphase Equation of State model is selected automatically, along with any other related models.

An example of selections for Multiphase Segregated Flow is shown below.



The remaining selections in this dialog depend on problem characteristics that are not related to Eulerian phases.

Defining Eulerian Phases

A Eulerian phase in STAR-CCM+ is a phase that is modeled in a Eulerian framework. The definition of each Eulerian phase includes the set of models applicable to its material; two materials constitute two different Eulerian phases.

Eulerian phases are stored under the Eulerian Phases node, which has a pop-up menu. Each Eulerian phase node contains a Models node, on which you select phase models, and an Initial Conditions node, on which you set the initial volume fraction and other appropriate initial conditions. As Eulerian phase nodes are subject to common menu items, you can copy and paste the properties of one phase onto another.



Initially, the Eulerian Phases node is empty. The process for creating and defining Eulerian phases is outlined below:

1. Creating a Eulerian Phase
2. Renaming the Phase
3. Selecting Phase Models
4. Replacing the Phase Material
5. Customizing Material Properties
6. Defining Volume Fraction in Eulerian Multiphase Flows

Repeat Steps 1 - 6 for all phases

Repeat these steps until all Eulerian phases have been specified. When all Eulerian phases have been defined, define phase interactions for pairs of phases that have significant contact with each other.

Creating a Eulerian Phase

The Eulerian Phases manager node (which is under the Eulerian Multiphase model node) is initially unpopulated.

- Right-click the Eulerian Phases node and select New from the pop-up menu.



A new Eulerian phase node appears beneath the Eulerian Phases node, with its own Models and Initial Conditions sub-nodes.



Renaming the Phase

The newly created phase can be renamed if desired.

1. Right-click the phase node and select Rename as shown below.



2. Enter the new name into the inline editor and press `<Enter>`.

Selecting Phase Models

Each Eulerian phase has phase models that are associated with it. A Eulerian phase model plays the same role in a Eulerian phase as a conventional model does in a continuum. It defines how a physical phenomenon in a phase is represented, what the primary variables of the simulation are and what mathematical formulation is used to generate the solution. Eulerian phase model selection is not identical to model selection for a continuum, since Eulerian phases implicitly inherit certain model choices such as those for space, time, and motion.

Eulerian phase models are stored under the Models manager node for the phase. They are selected using the *Phase Model Selection* dialog, which is accessed by right-clicking the Models node and choosing Select models...



The *Phase Model Selection* dialog is used in the same way as the *Physics Model Selection* dialog. That is, it allows you to select and review phase models, change which models have been selected, and supports auto-selection of recommended models. The choice of material and equation of state models

available in the dialog depends on the physics models you selected at the continuum level (either Volume of Fluid (VOF) or Multiphase Segregated Flow). An example *Phase Model Selection* dialog is given below:



Once you have selected the phase models, click Close.

Nodes corresponding to the chosen models appear under the Models node for the Eulerian phase.



If you try to run a simulation when you have not completed the phase model selection, the error below appears:



Eulerian Phase Models Manager Pop-Up Menu

Select Models...	Lets you select the Eulerian phase models.

See also: common menu items.

Replacing the Phase Material

The default material model appears within the Models node for the phase. To replace the default material with something more appropriate:

1. Right-click the material node and selecting Replace with.



The *Replace Material* dialog appears.

2. Navigate to the required material, select its radio button, and click OK to confirm the selection.



Customizing Material Properties

Further changes to material properties can be made by accessing the individual material property nodes within the material node as shown.



If the flow is laminar, the dynamic viscosity of a liquid can be set to use the Non-Newtonian Generalized Power Law.

Defining Volume Fraction in Eulerian Multiphase Flows

The volume fraction of a phase is the ratio of the volume that the phase occupies to the computational cell volume. Define the initial and boundary conditions of this scalar before running the simulation. You specify the volume fraction for each phase under the corresponding phase node.

When you specify initial and boundary values, there is no restriction on the values that you can set for the volume fractions of the phases. However, STAR-CCM+ uses normalized values for each phase volume fraction to ensure that the total volume fraction is 1.0.

If you are using the Volume of Fluid (VOF) model with only two phases, the volume fraction of the second-defined phase is calculated by subtracting the volume fraction value of the first-defined phase from 1.0. Any value that you specify is ignored.

- Setting Initial Volume Fractions
- Setting Volume Fraction at Flow Boundaries

This section applies to flows using either the Multiphase Segregated Flow model or the Volume of Fluid (VOF) model.

Setting Initial Volume Fraction

Following the phase specification, the Initial Conditions node of each phase contains a Volume Fraction sub-node.



Use this node to define the initial spatial distribution of the phase. The volume fraction of a phase is entered as a scalar profile.

Setting Volume Fraction at Flow Boundaries

Following the Eulerian phase specification, the node of every **open** boundary in the simulation contains a Phase Conditions node that lets you specify the physics conditions and values for each phase.

For each phase, use the Physics Values > Volume Fraction node to define the boundary condition on the phase volume fractions. The volume fraction of a phase is entered as a scalar profile.

When you define an outlet pressure boundary, you must set a volume fraction to cover the case of backflow. Otherwise, this volume fraction does not influence the solution domain.



A zero-flux condition is imposed at **closed** (wall) boundaries. Therefore, no user specification is needed there.

Modeling the Eulerian Phase Material

The material of each Eulerian phase is defined by selecting one of the material models in the *Phase Model Selection* dialog. For the Multiphase Segregated Flow model, you can choose from:

- Gas
- Particle
- Liquid
- Multi-component gas
- Multi-component liquid

For the VOF Multiphase flow model, the Particle option is deactivated.

Additional documentation on working with single-component materials can be found in the section, Managing single-component materials. Documentation on working with multi-component materials can be found in the section, Managing multi-component materials. This latter section includes details on adding individual mixture components.

Upon selecting a material model, the corresponding node is displayed under the Models manager node for the phase. The material model node (Gas and Liquid in the following screenshot), which is subject to common menu items, represents the Eulerian phase material model. Each Eulerian phase has its own material model, since the phase represents a distinct physical substance.



In Multiphase Segregated Flow cases, the options for the phase material include *Particle*. Selecting this option allows you to define a dispersed particle phase that is modeled using the Eulerian Multiphase approach. To model the effect of particle-particle interaction within the phase, use the Solid Pressure Force model.

To define the phase material, select the appropriate material model from the *Material* group box, as shown below.



Modeling the Eulerian Phase Equation of State

A Eulerian phase requires its own Equation of State model to define how the density of the material is computed. This section explains the equation of state models available for Eulerian phases. The particular models available depend on the underlying Multiphase Model for the continuum:

- What is the Constant Density Eulerian Phase Model?
- What is the Ideal Gas Eulerian Phase Model?
- What is the Polynomial Density Eulerian Phase Model?
- What is the User Defined EOS Eulerian Phase Model?

What Is the Constant Density Eulerian Phase Model?

The Constant Density node, which is subject to common menu items, represents the Constant Density Eulerian phase model. The model is the same as the Constant Density model for the STAR-CCM+ simulation in general, except that it applies to the Eulerian phase rather than throughout

the continuum.



Select the Constant Density option from the *Equation of State* group, as shown below.



Properties

None

What Is the Ideal Gas Eulerian Phase Model?

The Ideal Gas node, which has properties and is subject to common menu items, represents the Ideal Gas Eulerian phase model. The model is the same as the Ideal Gas model for the STAR-CCM+ simulation in general, except

that it applies to the Eulerian phase rather than throughout the continuum. In the context of Eulerian Multiphase, this model is intended for use in low Mach number flows (typically for Mach numbers less than 0.3).



Activate the Gas Eulerian phase model to make the Ideal Gas option available. Select it from the *Equation of State* group, as shown below.



Ideal Gas Properties

Incompressible	Controls whether the ideal gas is compressible (that is, dependent on pressure).	
	Activated	Removes dependency on pressure.
	Deactivated	Allows dependency on pressure.

What Is the Polynomial Density Eulerian Phase Model?

The Polynomial Density node, which is subject to common menu items, represents the Polynomial Density Eulerian phase model. The model is the same as the Polynomial Density model for the STAR-CCM+ simulation in general, except that it applies to the Eulerian phase rather than throughout the continuum.



Select the Polynomial Density option from the *Equation of State* group, as shown below.



Properties

None

What Is the User Defined EOS Eulerian Phase Model?

The User-Defined EOS node, which has properties and is subject to common menu items, represents the User-Defined Equation of State Eulerian phase model. The model is the same as the User Defined EOS model for the STAR-CCM+ simulation in general, except that it applies to the Eulerian phase rather than throughout the continuum.



Select the User-Defined EOS option from the *Equation of State* group, as shown below.



User Defined EOS Properties

Compressible	Controls whether the fluid is compressible (that is, dependent on pressure).	
	Activated	Allows dependency on pressure.
	Deactivated	Removes dependency on pressure.

Modeling the Eulerian Phase Viscous Regime

A Eulerian phase requires its own Viscous Regime model. The model depends on the viscous regime for the continuum. If you select the Laminar or Inviscid models, all Eulerian phases are automatically set to match the continuum and cannot be changed. If you select the Turbulent model in the continuum, then phases can be set to use either the Laminar or Turbulent models (although Turbulent is activated by default).

To activate a laminar viscous regime in a turbulent continuum:

1. Right-click on the Models node of the phase to be changed.
2. In the *Model Selection* dialog, do the following:
   a. Deselect Turbulent in the *Enabled Models* box.

The *Viscous Regime* box appears as below:



b. Select Laminar from the *Viscous Regime* box.

3. Click Close.

The selected models are appears as below:



Using the Eulerian Multiphase Mixture Model

The Eulerian Multiphase Mixture model is a simplified multiphase model that can be used to model suspension-like multiphase flows. In this model, the computational efforts are reduced by assuming the suspension to be a homogeneous single-phase system.

Industrial applications of the Eulerian Multiphase Mixture model include steam generators and engine cooling systems.

Carry out the following steps when setting up a Eulerian multiphase mixture simulation:

1. Select the Eulerian Multiphase Mixture model and other physics models.
2. Define Eulerian phases and select phase models.
3. Define phase interactions and select phase interaction models.

What is the Eulerian Multiphase Mixture Model?

Using multiphase models in numerical flow modeling of solid/fluid suspension-like flows is highly complex, particularly for flows that have wide particle distributions.

To simplify this complexity, the Eulerian Multiphase Mixture model covers three different levels of approximation:

- Homogeneous Multiphase model

  The multiphase mixture is treated as a homogeneous single-phase system, with the following assumptions ([306], [310]):

  - The mixture of phases behaves as a single fluid.

    A single set of transport equations for mass, momentum, and energy is solved for the mixture as a whole.
  - The phases are in thermal equilibrium.

    The two phases have the same temperature and the phase distribution is computed based on heat balance.
  - There is no relative/slip velocity between the phases.
  - The phases flow through porous medium with porosity χ.
- Algebraic Slip model

This model uses empirical correlations for the slip ratio between phase velocities to calculate the void fraction. However, the kinematic effect of relative motion between the phases is not accounted for.

- Drift Flux model.

  This model accounts for the kinematic effects of relative motion between the phases.

Selecting the Eulerian Multiphase Mixture Model

To select the Eulerian Multiphase Mixture model:

1. In the *Physics Model Selection* dialog, select the Eulerian Multiphase material model.

   The Multiphase Interaction model is selected automatically.

2. In the *Eulerian Multiphase Model* group box, select Eulerian Multiphase Mixture.



The Multiphase Equation of State model and the Two-Phase Thermodynamic Equilibrium model are selected automatically.

An example of model selections for Eulerian Multiphase Mixture is shown below.



Selecting the Two-Phase Thermodynamic Equilibrium Model

By default, the Two-Phase Thermodynamic Equilibrium model is selected automatically when you select the Eulerian Multiphase Mixture model.

If you deactivate Auto-Select Recommended Models, the Two-Phase Thermodynamic Equilibrium model becomes available for selection when you select the Eulerian Multiphase Mixture model.



Setting Up Phase Interactions

When the Eulerian Multiphase Mixture model is activated in the physics continuum, the following models are selected automatically in each phase interaction:

- Multiphase Material
- Two-Phase Equilibrium Interaction

The Multiphase Material node under the phase interaction Models node lets you specify the method that is used to determine the latent heat.

The following interphase transfer models are available for Eulerian multiphase mixtures:

- Algebraic Slip
- Drift Flux
- Boiling

What is the Algebraic Slip Model?

The Algebraic Slip model uses empirical correlations for the slip ratio between phase velocities to calculate the void fraction. However, the kinematic effect of relative motion between the phases is not introduced into the governing transport equations. This approach provides a better estimate of the distribution of void fraction and mixture density than the homogenous model (the Eulerian Multiphase Mixture model).

Selecting the Algebraic Slip Model

To select the Algebraic Slip model:

1. Create the necessary phase interaction.
2. Select the phase interaction models.

   By default, the Two-Phase Equilibrium Interaction model and the Multiphase Material model are selected automatically.

3. In the *Optional Models* group box, select Algebraic Slip.



4. Click Close.

The Algebraic Slip node is added under the Phase Interaction > Models node.



You specify the Velocity Ratio value.

Velocity Ratio Properties

Method	Specifies the method for defining the slip velocity ratio. This value is S in Eqn. (1672).	
	Constant	Defines the slip velocity ratio as a constant value.
	Field Function	Defines the slip velocity ratio using a field function (typically user defined).

What is the Drift Flux Model?

The Drift Flux model accounts for the kinematic effects of relative motion between the phases. The relative motion is defined using kinematic constitutive equations for the particular flow regime and thermal-hydraulic conditions.

Selecting the Drift Flux Model

To select the Drift Flux model:

1. Create the necessary phase interaction.
2. Select the phase interaction models.

   By default, the Two-Phase Equilibrium Interaction model and the Multiphase Material model are selected automatically.
3. In the *Optional Models* group box, select Drift Flux.



4. Click Close.

The Drift Flux node is added under the Phase Interaction > Models node.



You specify the Distribution Parameter and Drift Velocity values.

Distribution Parameter Properties

Method — Specifies the method for defining the distribution factor that reflects non-uniformity in the void fraction and velocity distribution. This value is C_0 in Eqn. (1680).

- Constant — Defines the distribution parameter as a single three-part (or two-part for two-dimensional), comma-separated, number.
- Field Function — Defines the distribution parameter using a field function (typically user defined).

Drift Velocity Properties

Method — Specifies the method for defining the volume-weighted mean drift velocity. This value is $\mathbf{v}_{vj}$ in Eqn. (1680).

- Constant — Defines the drift velocity as a single three-part (or two-part for two-dimensional), comma-separated, number.
- Field Function — Defines the drift velocity using a field function (typically user defined).
- Composite — Defines the drift velocity as three components (for example, X, Y, Z) for three-dimensional problems, or two components (for example, X, Y) for two-dimensional problems.

Selecting a Boiling Model

To include the wall boiling heat flux in your simulation, you select the appropriate boiling model in the phase interaction.

When the Boiling model is activated in the *Phase Interaction Model Selection* dialog, the Rohsenow Boiling and Transition Boiling models are available for selection in the *Boiling Models* group.



The selected boiling model appears in the Multiphase Interaction > Phase Interactions > Phase Interaction 1 > Models node. The Transition Boiling model is shown below.



Set the appropriate Rohsenow Boiling properties or Transition Boiling properties.

Multiphase Mixture Field Functions Reference

The following field functions are available when the Multiphase Mixture model is activated. These objects all have the same base properties.

Saturation Pressure of <phase>	The saturation pressure of the specified phase.
Saturation Temperature of <phase>	The saturation temperature of the specified phase.
Static Enthalpy	The static enthalpy.
Static Enthalpy at Saturation Temperature of <phase>	The static enthalpy at the saturation temperature of the specified phase.

Eulerian Multiphase Mixture Model Formulation

The Eulerian Multiphase Mixture model formulation includes the following sections:

- Homogeneous (Eulerian Multiphase Mixture) Model
- Algebraic Slip Model
- Drift Flux Model

Homogeneous (Eulerian Multiphase Mixture) Model

Governing Equations

The following equations describe mass, momentum, and energy conservation of the mixture of phases in porous media:

$$\frac{d}{dt}\int_V \chi\rho_m dV + \int_A \chi\rho_m(\mathbf{v}_m - \mathbf{v}_g)\cdot da = \int_V \chi s_m dV \tag{1659}$$

$$\frac{d}{dt}\int_V \chi\rho_m\mathbf{v}_m dV + \int_A \chi\rho_m\mathbf{v}_m(\mathbf{v}_m - \mathbf{v}_g)\cdot da = \tag{1660}$$
$$-\int_A \chi p\mathbf{I}\cdot da + \int_A \chi\mathbf{S}\cdot da + \int_V \chi(\mathbf{f}_g + \mathbf{f}_p + \mathbf{s}_m)dV$$

$$\frac{d}{dt}\int_V \chi((\rho_m E_m)_{\text{fluid}} + (1-\chi)(\rho C_p)_{\text{solid}}T)dV + \int_A \chi[\rho_m H_m(\mathbf{v}_m - \mathbf{v}_g) + \mathbf{v}_g p]\cdot da = \tag{1661}$$
$$-\int_A \dot{\mathbf{q}}\cdot da + \int_A \chi(\mathbf{S}\cdot\mathbf{v}_m)\cdot da + \int_V \chi(\mathbf{f}_g + \mathbf{v}_m + s_h)dV$$

The mixture density ρ_m and mass-weighted mixture velocity $\mathbf{v}_m$ are defined as:

$$\rho_m = \sum_{i=1}^{n}\alpha_i\rho_i \qquad \mathbf{v}_m = \frac{\sum_{i=1}\alpha_i\rho_i\mathbf{v}_i}{\rho_m} \tag{1662}$$

where s_m, $\mathbf{s}_m$, and s_h are user-defined sources of mass, momentum, and energy, and $\mathbf{f}_g$ and $\mathbf{f}_p$ are the forces due to gravity and porous resistance.

The porous resistance force is approximated as:

$$\mathbf{f}_p = \mathbf{P} \cdot \mathbf{v}_{m_{sup}} \tag{1663}$$

where $\mathbf{P}$ is the porous resistance tensor that is composed of the viscous $\mathbf{P}_\mu$ and inertial resistance tensor $\mathbf{P}_i$:

$$\mathbf{P} = \mathbf{P}_\mu + \mathbf{P}_i \left|\mathbf{v}_{m_{sup}}\right| \tag{1664}$$

and $\mathbf{v}_{m_{sup}}$ is the mixture superficial velocity, which is defined as $\mathbf{v}_{m_{sup}} = \chi \mathbf{v}_m$.

The stress tensor, $\mathbf{S}$, is estimated as:

$$\mathbf{S} = \mu_{eff}\left[(\nabla \mathbf{v}_m + \nabla \mathbf{v}_m^T) - \frac{2}{3}(\nabla \bullet \mathbf{v}_m)\mathbf{I}\right] \tag{1665}$$

where μ_{eff} is the effective viscosity.

Turbulence quantities in the porous medium are prescribed either directly or by specifying the turbulent intensity and length scale or turbulent intensity and viscosity ratio.

Total energy E_m and total enthalpy H_m of the mixture are defined as:

$$E_m = H_m - \frac{p}{\rho_m} \tag{1666}$$

$$H_m = h_m + \frac{\left|\mathbf{v}_m\right|^2}{2} \tag{1667}$$

where h_m is the mixture-specific static enthalpy.

The heat flux $\dot{\mathbf{q}}$ in a porous medium is approximated as:

$$\dot{\mathbf{q}} = -[\chi \mathbf{K}_{eff,\, fluid} + (1-\chi)\mathbf{K}_{eff,\, solid}] \cdot \nabla T \tag{1668}$$

where $\mathbf{K}_{eff}$ are the effective thermal conductivities in fluid and solid.

Modeling Phase Change

Assuming that the liquid and vapor phase are in thermodynamic equilibrium, steam mass quality (vapor mass fraction) x is a function of the static thermal enthalpy h ([306], [310]):

$$x = \max\left[0, \min\left(1, \frac{h_m - h_{ls}}{h_{vs} - h_{ls}}\right)\right] \tag{1669}$$

where h_{ls} and h_{vs} are the enthalpies of liquid and vapor at saturation temperature T_{sat}.

The vapor volume fraction α is calculated using the expression:

$$\alpha = \frac{x}{x + (1 - x)\frac{\rho_{vs}}{\rho_{ls}}} \tag{1670}$$

where ρ_{vs} and ρ_{ls} are densities of vapor and liquid at saturation temperature.

Algebraic Slip Model

The Algebraic Slip model uses empirical correlations for the slip velocity ratio to obtain a better estimate of the distribution of void fraction and mixture density than the homogenous model (the Eulerian Multiphase Mixture model).

The Algebraic Slip model uses the same set of transport equations as the homogenous model (Eqn. (1659), Eqn. (1660), and Eqn. (1661)). No kinematic effects of relative motion between the two phases are considered in the governing equations.

The Algebraic Slip (or void fraction) model uses empirical correlations for the slip velocity ratio S:

$$S = \frac{\mathbf{v}_v \cdot \mathbf{g}}{\mathbf{v}_l \cdot \mathbf{g}} \tag{1671}$$

where $\mathbf{g}$ is the gravity vector.

In the Algebraic Slip model, the expression for the vapor volume fraction (Eqn. (1670)) is replaced with:

$$\alpha = \frac{x}{x + (1 - x)\frac{\rho_{vs}}{\rho_{ls}}S} \tag{1672}$$

where the slip velocity ratio S is calculated using an appropriate correlation.

Drift Flux Model

The Drift Flux model accounts for the kinematic effects of relative motion between the phases. The Drift Flux model has the advantages of the homogenous model (the Eulerian Multiphase Mixture model) of being easy to formulate and apply, and the influences of relative motion between the phases are included in the governing equations.

Relative motion for the Drift Flux model is defined using kinematic constitutive equations for the particular flow regime and thermal-hydraulic conditions.

The main assumptions of the Drift Flux model are ([306], [307], [310]):

- The mixture of phases behaves as a single fluid.
- The phases are in thermal equilibrium.
- The phases can move relative to each other.
- The phases flow through porous medium with porosity χ.

Governing Equations

The mass, momentum, and energy conservation of a mixture of phases are described with the following set of transport equations ([306], [307], [310]):

$$\frac{d}{dt}\int_V \chi\rho_m dV + \int_A \chi\rho_m(\mathbf{v}_m - \mathbf{v}_g)\cdot da = \int_V \chi s_m dV \tag{1673}$$

$$\frac{d}{dt}\int_V \chi\rho_m\mathbf{v}_m dV + \int_A \chi\rho_m\mathbf{v}_m(\mathbf{v}_m - \mathbf{v}_g)\cdot da = \tag{1674}$$
$$-\int_A \chi p\mathbf{I}\cdot da + \int_A \chi\mathbf{S}\cdot da + \int_V \chi(\mathbf{f}_g + \mathbf{f}_p + \mathbf{s}_m)dV - \int_A \chi\frac{\alpha}{1-\alpha}\frac{\rho_v\rho_l}{\rho_m}\bar{\mathbf{v}}_{vj}\bar{\mathbf{v}}_{vj}\cdot da$$

$$\frac{d}{dt}\int_V \chi((\rho_m E_m)_{\text{fluid}} + (1-\chi)(\rho C_p)_{\text{solid}}T)dV + \tag{1675}$$
$$\int_A \chi[\rho_m H_m(\mathbf{v}_m - \mathbf{v}_g) + \mathbf{v}_g p]\cdot da =$$
$$-\int_A \dot{\mathbf{q}}\cdot da + \int_A \chi(\mathbf{S}\cdot\mathbf{v}_m)\cdot da - \int_A \chi\alpha\frac{\rho_v\rho_l}{\rho_m}(h_{vs} - h_{ls})\bar{\mathbf{v}}_{vj}\cdot da +$$
$$\int \chi\left(\mathbf{f}_g\cdot\mathbf{v}_m + s_h + \frac{\alpha(\rho_l - \rho_v)}{\rho_m}\bar{\mathbf{v}}_{vj}\cdot\nabla p\right)dV$$

Here $\bar{\mathbf{v}}_{vj}$ is the mean drift velocity of the vapor phase, which is defined as:

$$\bar{\mathbf{v}}_{vj} = \mathbf{v}_v - \mathbf{j} \tag{1676}$$

where $\mathbf{v}_v$ is the velocity of the vapor phase and $\mathbf{j}$ is the volume-weighted mixture velocity:

$$\mathbf{j} = \alpha\mathbf{v}_v + (1-\alpha)\mathbf{v}_l \tag{1677}$$

The drift velocity can be specified using kinematic constitutive equations [309], which can be written as:

$$\bar{\mathbf{v}}_{vj} = \bar{\mathbf{v}}_{vj}(\alpha, \sigma, g, \mathbf{v}_m, p_m, \text{etc.}) \tag{1678}$$

Knowing the drift velocity, the void-fraction-weighted mean velocity of individual phases can be estimated:

$$\mathbf{v}_v = \mathbf{v}_m + \frac{\rho_l}{\rho_m}\bar{\mathbf{v}}_{vj} \qquad \mathbf{v}_l = \mathbf{v}_m - \frac{\alpha\rho_v}{(1-\alpha)\rho_m}\bar{\mathbf{v}}_{vj} \tag{1679}$$

The mean drift velocity $\bar{\mathbf{v}}_{vj}$ can also be expressed as:

$$\bar{\mathbf{v}}_{vj} = \frac{[(C_0-1)\mathbf{v}_m + \mathbf{v}_{vj}]\rho_m}{C_0\alpha\rho_v + (1-C_0\alpha)\rho_l} \tag{1680}$$

where:

- C_0 is a distribution factor that reflects non-uniformity in the void fraction and velocity distribution
- $\mathbf{v}_{vj}$ is the volume-weighted mean drift velocity which could be estimated using appropriate constitutive equations.

Auxiliary Variables

The auxiliary variables that are associated with the Drift Flux model are:

- Mean drift velocity $\bar{\mathbf{v}}_{vj}$.
- Volume-weighted mean drift velocity $\mathbf{v}_{vj}$ and the distribution parameter C_0, as defined by appropriate correlations.

  See Lellouche-Zolotar Correlation.
- Volume-weighted mean velocities of vapor $\mathbf{v}_v$ and liquid $\mathbf{v}_l$ phase (Eqn. (1679)).
- Velocity or slip ratio S between the phases, which can be calculated as:

$$S = \frac{\mathbf{v}_m\rho_m + \bar{\mathbf{v}}_{vj}\rho_l}{\mathbf{v}_m\rho_m + \bar{\mathbf{v}}_{vj}\rho_v\dfrac{\alpha}{1-\alpha}} \tag{1681}$$

Lellouche-Zolotar Correlation

In a bubbly or a churn turbulent bubbly flow regime, Lellouche and Zolotar [308] recommended the following correlations.

For $\mathbf{v}_{vj}$:

$$\mathbf{v}_{vj} = C_2\left[\frac{(\rho_l - \rho_g)}{\rho_l^2} \cdot \sigma_g\right]^{1/4} \frac{(1 - r_g)^{1/2}}{(1 + r_g)} \tag{1682}$$

where:

- $C_2 = 1.41$ in the tube bundle region.
- $C_2 = 4.5$ above the tube bundle.

For C_0:

$$C_0 = \frac{L}{(\kappa_l + (1 - \kappa_l) r_g^{\gamma})} \tag{1683}$$

where:

$$\kappa_l = \kappa_0 + (1 - \kappa_0)(\rho_g / \rho_l)^{1/4} \tag{1684}$$

$$\kappa_0 = [1 + e^{-Re/10^5}]^{-1} \tag{1685}$$

$$L = \frac{1 - e^{-c_l r_g}}{1 - e^{-c_l}} \tag{1686}$$

$$c_l = \frac{4.0 p_{cr}^2}{p(p_{cr} - p)} \tag{1687}$$

$$\gamma = \left[\frac{1 + 1.57\frac{\rho_g}{\rho_l}}{1 - \kappa_0}\right] \tag{1688}$$

and:

- p_{cr} is the critical pressure, which has a value of $221.14 \times 10^5 \mathrm{N/m}^2$ for steam and water.

Also, as recommended in [308], the value of κ_0 is trapped such that:

$$\kappa_0 = \min(0.8, \kappa_0) \tag{1689}$$

or

$$\mathrm{Re} = \min(\mathrm{Re}, 1.3863 \times 10^5) \tag{1690}$$

The Reynolds number in Eqn. (1690) is calculated as:

$$\mathrm{Re} = \frac{\mathrm{Gd}}{\mu l} \tag{1691}$$

where:

- d is either the equivalent hydraulic diameter inside the tube bundle or the inner shell diameter above the tubes.
- G is the total mass flow rate entering the secondary-side of the generator.

Multiphase Mixture Bibliography

[306] Collier, John G., and Thome, John R. "Convective Boiling and Condensation" (3rd ed. in paperback), Clarendon Press, Oxford, 1996.

[307] Hibiki, T. and Ishii, M. 2003. "One-dimensional drift-flux model and constitutive equations for relative motion between phases in various two-phase flow regimes", *Int. Journal of Heat and Mass Transfer*, **46**, pp. 4935-4948.

[308] Lellouche, G.S. and Zolotar, B.A., "Mechanistic Model for Predicting Two-Phase Void Fraction for Water in Vertical Tubes, Channels, and Rod Bundles", Special Report, *ELECTRIC POWER RESEARCH INSTITUTE*, February 1982.

[309] "One-dimensional Drift-Flux Model and Constitutive Equations for Relative Motion Between Phases in Various Two-Phase Flow Regimes", ANL-77-47, USA, 1977.

[310] Thome, John. R., "Engineering Data Book III", Wolverine Tube, Inc., 2010.

Using the Multiphase Segregated Flow Model

The Multiphase Segregated Flow model in STAR-CCM+ is based on an Eulerian-Eulerian formulation where each distinct phase has its own set of conservation equations. Phases are considered to be mixed on length scales smaller than the length scales to resolve, and coexist everywhere in the flow domain. The volume fraction is the portion of a volume that a phase occupies.

This concept of coexisting phases is called "interpenetrating continua". The concept assumes that you are interested in the time averaged behavior of the flow, rather than the instantaneous behavior. This assumption has important implications for understanding volume fraction in the context of the Multiphase Segregated Flow model.

Carry out the following steps when setting up a multiphase segregated flow simulation:

1. Selecting the Multiphase Segregated Flow model and other physics models.
2. Defining Eulerian phases and selecting phase models.
3. Defining phase interactions and selecting phase interaction models.
4. Defining volume fractions for inflow boundaries and initial conditions.
5. Specifying solver settings.

A formulation section for multiphase segregated flows is included. This section also includes information on the following features:

- Recommended solver settings
- Turbulence in multiphase flows
- Energy in multiphase flows
- Combustion in multiphase flows
- Modeling bulk and wall boiling
- Modeling interphase mass transfer
- Modeling dissolution mass transfer rate
- Modeling particle-phase flows
- Modeling boundary permeability
- Modeling particle size distribution (S-Gamma)
- Using mass source terms

What Is the Multiphase Segregated Flow Model?

The Multiphase Segregated Flow model is used to model Eulerian Multiphase cases (also referred to as Eulerian-Eulerian, Euler-Euler, EMP). The Multiphase Segregated Flow model solves a set of conservation equations for each Eulerian phase present in the simulation. The pressure is assumed to be the same in all phases. The volume fraction gives the share of the flow domain that each phase occupies. Each phase has its own velocity and physical properties. The term "segregated" refers to the fact that the solution algorithm uses a SIMPLE-type approach, which has separate pressure and velocity solvers.

The Multiphase Segregated Flow node, which has properties and is subject to common menu items, is activated as soon as the model is selected.



Interfaces separate the multiple phases present in the simulation. The motions of the phases influence the interface between each pair of phases. The phases can be any kind of fluid in the sense of moving gas, liquid or

solid particles. In general, the phases are not in equilibrium and each phase has its own velocity, energy, and other variables and its own physical properties.

The conservation equations for each phase variable require closure by the definition of phase interactions at each phase interface. This definition consists of suitable models for the interfacial area, and for the rates of interphase transfer of mass, momentum and energy. These closures characterize and vary between different multiphase flow patterns, and are often not known.

Properties

Velocity Convection	Selects the convection scheme for flow velocity.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme.
Volume Fraction Convection	Selects the convection scheme for volume fraction.	
	1st-order	Selects the first-order upwind convection scheme.
	2nd-order	Selects the second-order upwind convection scheme. If you use this option, set the Volume Fraction Minimum expert property to have a value greater than the default setting of `0.0`.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model and the STAR-CCM+ discretization techniques, do not change the *Expert* settings. The values in that category have been optimized for accuracy and performance.

Flow Boundary Diffusion	Activate/deactivate the flow-boundary diffusion fluxes (or viscous fluxes for flow models)—see Eqn. (2558). This property is activated by default.	
	Activated	Activate flow-boundary diffusion.
	Deactivated	Deactivate flow-boundary diffusion.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Volume Fraction Minimum	Sets a lower limit on the value of volume fraction in a cell for any phase. The default value is `0.0`. If you have the Volume Fraction Convection set to `2nd-order`, enter a value greater than zero. When instability issues arise, set to `1.0E-25`.	

Selecting the Multiphase Segregated Flow Model

The following steps describe how to select and use the Multiphase Segregated Flow Model:

- In the *Physics Model Selection* dialog, select option Eulerian Multiphase in the *Material* group box.

The *Eulerian Multiphase Model* group box appears as shown below.



- Select option Multiphase Segregated Flow from the *Eulerian Multiphase Model* group box. The *Viscous Regime* group box appears.
- Select Laminar or Turbulent from the *Viscous Regime* group box as desired.

The Laminar option has been selected in the dialog that is shown below.



For Time models, use either Steady or Implicit Unsteady.

Modeling Phase Interactions in Multiphase Segregated Flow

In Multiphase Segregated Flow simulations, the phase interaction model defines the topology and physics of the interface that separates any two interacting phases. Information on the process of creating phase interactions and selecting phase interaction models is provided in the section Defining Phase Interactions. For multiphase segregated flows, there are two types of phases interactions:

- Eulerian Continuous-Dispersed
- Granular Dispersed-Dispersed

Either model can be selected from the *Phase Interaction Topology* group box in the *Phase Interaction Model Selection* group box.



What Is the Eulerian Continuous-Dispersed Model?

For each Eulerian multiphase phase interaction that is not between particles in the Granular Pressure model, it is necessary to define which phase is continuous and which is dispersed. This setting is defined on the properties of the Eulerian Continuous-Dispersed model node.



Eulerian Continuous-Dispersed Properties

Continuous Phase	Choose the continuous phase of the interacting pair from the drop-down list of previously defined Eulerian phases.
Dispersed Phase	Choose the dispersed phase of the interacting pair from the drop-down list of previously defined Eulerian phases.

When you choose the Eulerian Continuous Dispersed model, the following models are selected automatically:

- Drag Force
- Multiphase Material
- Interaction Length Scale
- Interaction Area Density

The Multiphase Material node in the phase interaction Models manager lets you set any additional material properties that are required by some of the phase interaction models.

The following interphase transfer models are available for multiphase segregated flows:

- Lift force model
- Wall Lubrication force model
- Turbulent dispersion force model
- Virtual mass force model
- Interphase Energy Transfer

If the continuous phase is a liquid and the dispersed phase is a particle, you can also select the Suspension Rheology model.

Modeling Interaction Length Scale

The interaction length scale is used to define non-dimensional parameters such as the Reynolds number for a phase interaction, and also to compute an interaction area density. The interaction length scale for a continuous-dispersed interaction is chosen to be the mean particle size.

The Interaction Length Scale model is selected automatically when you select the Drag Force model.



The Interaction Length Scale node is added to the phase interaction models.



In the *Properties* window, you set the method for defining the interaction length scale.

Interaction Length Scale Model Properties

Method — Specifies the method for defining the interaction length scale.

- **Constant** — Specifies a constant value that represents the mean particle size for the phase interaction.
- **Field Function** — Specifies a field function that represents the mean particle size for the phase interaction. For example, in a flow with a constant number density of particles, particle size could be specified to be proportional to the phase volume fraction.
- **Particle Diameter** — Uses the particle diameter of the particle phase, which is set as the dispersed phase. This method is available when the Granular Pressure model is selected in the physics continuum.
  You can specify the particle diameter in the Material Properties node in the particle phase definition.
- **Sauter Mean Diameter** — Uses the Sauter Mean Diameter of the dispersed phase. This method is made available when the S-gamma model is activated for the dispersed phase.
- **Kurul Podowski** — Specifies bubble size as a function of local subcooling.
  The Kurul Podowski interaction length scale is an empirical relation that is used in boiling flows. It is available only when the boiling mass transfer rate has been selected with interphase mass transfer.

Kurul Podowski Interaction Length Scale

The parameters that are required when calculating this length scale are set as properties of the Kurul Podowski node.



Kurul Podowski Interaction Length Scale Properties

Minimum Diameter	Minimum bubble diameter
Maximum Diameter	Maximum bubble diameter
Minimum Diameter Subcooling	Liquid subcooling temperature corresponding to minimum bubble diameter
Maximum Diameter Subcooling	Liquid subcooling temperature corresponding to maximum bubble diameter

Modeling Interaction Area Density

The interaction area density specifies the interfacial area available for drag, heat, and mass transfer between each pair of phases in an interaction.

The Interaction Area Density model is selected automatically when you select the Drag Force model.



Use the Method property of the Interaction Area Density node to choose how you want to define the value.



Interaction Area Density Properties

Method — Specifies the method to use for modeling the interaction area density.

- **Spherical Particle** — Defines the interaction area density using the surface area of a spherical particle, from Eqn. (1696).
  This model is not applicable at high particle volume fractions.
- **Symmetric** — Uses the availability of the continuous phase to modify the interaction area density, from Eqn. (1697).
  This model can be used for a wide range of volume fractions.
- **Field Function** — Choose a field function in the Field Function child node.

Using the Suspension Rheology Model

The presence of solid particles suspended in a Newtonian liquid is known to lead to non-Newtonian behavior, as the mixture viscosity becomes dependant on the volume fraction of the particle phase. For a Eulerian multiphase flow, you can model this behavior using the Suspension Rheology phase interaction model.

The Suspension Rheology model is available only when the Eulerian Continuous-Dispersed phase interaction model is selected, and the continuous phase is specified as a liquid and the dispersed phase is specified as a particle.



The Multiphase Interaction node with the Suspension Rheology model selected is shown below.



When you select the Suspension Rheology model, two new properties are added under the Multiphase Material > Material Properties node: Relative Viscosity and Normal Relative Viscosity.

For more information, see Setting the Suspension and Emulsion Rheology Model Properties.

Using the Emulsion Rheology Model

The presence of particles of liquid drops suspended in a Newtonian liquid is known to lead to non-Newtonian behavior, as the mixture viscosity becomes dependant on the volume fraction of the particle phase. For a Eulerian multiphase flow, you can model this behavior using the Emulsion Rheology phase interaction model.

The Emulsion Rheology model is available only when the Eulerian Continuous-Dispersed phase interaction model is selected, and both the continuous phase and the dispersed phase are specified as a liquid phase.



The Multiphase Interaction node with the Emulsion Rheology model selected is shown below.



When you select the Emulsion Rheology model, three new properties are added under the Multiphase Material > Material Properties node: Relative Viscosity, Normal Relative Viscosity, and Surface Tension.

For more information, see Setting the Suspension and Emulsion Rheology Model Properties.

Setting the Suspension and Emulsion Rheology Model Properties

This section includes the following:

- Setting the Material Properties
- Setting the Anisotropy Tensor
- Selecting the Drag Force Model for Emulsions

Setting the Material Properties

To use the Suspension Rheology and Emulsion Rheology models, you set the appropriate material properties.

Suspension and Emulsion Rheology Multiphase Material Properties

Normal Relative Viscosity	This property can be defined by a Constant, a Field Function, or the Morris and Boulay Model. If you set the Relative Viscosity property to use the Krieger-Dougherty model [342], this value defaults to zero.
Relative Viscosity	The dimensionless relative viscosity (see Eqn. (1698)) is used to describe the numerics of the mixture viscosity. Mathematically, this value tends to infinity as the dispersed phase reaches the limit of maximum packing. This property can be defined by a Field Function, the Krieger-Dougherty Model (see Eqn. (1699)), or the Morris and Boulay Model (see Eqn. (1700)).
Surface Tension	This property can be defined by a Constant or a Field Function.

The relative viscosity is a strong function of volume fraction, so it is sensitive to small variations of volume fraction in the simulation. Therefore, it is recommended that you run these models as unsteady with a small time-step, and use an initial volume fraction that is below the maximum packing.

The Krieger-Dougherty Model for Relative Viscosity

The Krieger-Dougherty model describes non-Newtonian flow behavior in rigid sphere suspensions. The model takes explicit notice of the interactions between neighboring spherical particles.

For more information, see The Krieger-Dougherty Model Formulation.

The Krieger-Dougherty Model node is shown below:



Krieger-Dougherty Model Relative Viscosity Properties

Maximum Packing	The maximum critical packing fraction (assuming random close packing); ϕ_m in Eqn. (1699). This value is 0.645 for hard spheres. Note that this value is a theoretical limit and will be lower in most practical cases. The value is lower again for non-spherical particles.
Maximum Viscosity	Specifies the maximum value of the relative viscosity that is allowed in the simulation. In Eqn. (1699), the relative viscosity η_r tends to infinity as the dispersed phase reaches the limit of maximum packing.
Intrinsic Viscosity Exponent	The intrinsic viscosity; $[\eta]$ in Eqn. (1699). This value is 2.5 for spherical particles.

The Morris and Boulay Model for Relative Viscosity

The Morris and Boulay Model node for Relative Viscosity is shown below:



Morris and Boulay Model Relative Viscosity Properties

Maximum Packing	The maximum packing fraction (assuming random close packing); ϕ_m in Eqn. (1700). This value is 0.645 for hard spheres. Note that this value is a theoretical limit and will be lower in most practical cases. The value is lower again for non-spherical particles.
Maximum Viscosity	Specifies the maximum value of the relative viscosity that is allowed in the simulation. In Eqn. (1700), the relative viscosity η_r tends to infinity as the dispersed phase reaches the limit of maximum packing.
Contact Contribution	The contact contribution, K_s in Eqn. (1700).

The Morris and Boulay Model node for Normal Relative Viscosity is shown below:



Morris and Boulay Model Normal Relative Viscosity Properties

Maximum Packing	The maximum packing fraction (assuming random close packing); ϕ_m in Eqn. (1701). This value is 0.645 for hard spheres. Note that this value is a theoretical limit and will be lower in most practical cases. The value is lower again for non-spherical particles.
Maximum Viscosity	Specifies the maximum value of the relative viscosity that is allowed in the simulation. In Eqn. (1701), the normal relative viscosity η_n tends to infinity as the dispersed phase reaches the limit of maximum packing.
Contact Contribution	The contact contribution, K_n in Eqn. (1701).

Setting the Anisotropy Tensor

Both the Suspension Rheology and Emulsion Rheology models have an Anisotropy Tensor option.

The Morris and Boulay model was derived from experiments in pipes and Couette devices where the flow becomes aligned in a particular direction governed by the geometry. The normal stress was shown to be anisotropic and, at different strengths, lead to migration of the particles in different directions (see Eqn. (1702)).

The Emulsion Rheology > Anisotropy Tensor node is shown below:



By default, the Anisotropy Tensor is an Isotropic Normal Stress Tensor.

For more information, see Anisotropy Tensor Formulation.

Selecting the Drag Force Model for Emulsions

The Emulsion Rheology model can use either the Suspension Drag Force model or the Emulsion Topology Inversion Drag Force model, depending on whether the phases are allowed to invert.

You select these models on the Drag Force > Linearized Drag Coefficient node.



Modeling Drag Force

This section explains the following about the Drag Force model that is used in phase interactions:

- Selecting the Drag Force Model
- What Is the Drag Force?
- What Is the Drag Coefficient?
- What Is the Linearized Drag Coefficient?

Selecting the Drag Force Model

The Drag Force model is selected by default on opening the *Phase Interaction Model Selection* dialog for the first time. If it has been deselected, however, it can be selected from the *Optional Phase Interaction Models* group box as shown.



The *Phase Interaction Model Selection* dialog shows the Drag Force model activated.



What Is the Drag Force?

The Drag Force node, which has properties and is subject to common menu items, represents the Drag Force model. This model calculates the force on a particle in the dispersed phase due to its velocity relative to the continuous phase, as from Eqn. (1710).



In STAR-CCM+, the Drag Force is modeled as a linear multiplier, A_D in Eqn. (1710), of the relative velocity between two phases. There are two approaches that can be taken to obtain this multiplier in the phase interaction model: the standard drag force coefficient and the linearized drag force coefficient type.

Drag Force Properties

Drag Force Coefficient Type	The type of coefficient to use.	
	Standard	Specifies the standard drag force coefficient.
	Linearized	Specifies the linearized drag force coefficient.

What Is the Drag Coefficient?

The Drag Coefficient node, which has properties and is subject to common menu items, is a sub-node of the Drag Force node.



Drag Coefficient Properties

Dimensions — The dimensionality of the drag coefficient (read-only).

Method — Selects the method to use for specifying the drag coefficient.

- **Constant** — Defines the standard drag coefficient using a constant value. A constant method node is added as a child to this node.
- **Field Function** — Defines the standard drag coefficient using a field function. A field function method node is added as a child to this node.
- **Schiller-Naumann** — Defines the standard drag coefficient using the Schiller-Naumann correlation. This correlation is available only when the continuous phase is viscous. However, the continuous phase can be either Newtonian or non-Newtonian. The model is suitable for spherical solid particles, liquid droplets, and small-diameter (spherical) bubbles.
- **Hamard and Rybczynski** — Defines the standard drag coefficient using the Hamard and Rybczynski correlation. This correlation is available only when the continuous phase is viscous. It is suitable for spherical liquid droplets.
- **Wang Curve Fit** — Defines the standard drag coefficient using the Wang Curve Fit correlation. This correlation is available only when the continuous phase is viscous, and when the dispersed phase is a fluid (it is not available for solid particles).
  This correlation is suitable for water-air bubble systems or similar systems at normal pressure and temperature.

Bozz-ano-Dente	Defines the standard drag coefficient using the Bozzano and Dente correlation. It is available when the dispersed phase is a fluid (it is not available for solid particles). Additionally, the correlation requires that gravity is activated for the physics continuum. Specify the surface tension value in the Multiphase Material > Material Properties > Surface Tension node for the parent phase interaction. The Surface Tension node appears automatically when the correlation is selected. The Bozzano-Dente correlation is suitable for a wide range of liquid-gas bubble systems including water at low and high pressure, and highly viscous liquids such as some oils, glycerol, or molten glass. For small bubble diameters in low viscosity systems (for example, water), Bozzano-Dente gives a different asymptote to other standard drag correlations.
Tomiyama	Defines the standard drag coefficient using the Tomiyama correlation. It is only available when the continuous phase is viscous, and when the dispersed phase is a fluid (it is not available for solid particles). Additionally, the correlation requires that gravity for the physics continuum. Specify the surface tension value in the Multiphase Material > Material Properties > Surface Tension node for the parent phase interaction. The Surface Tension node appears automatically when the correlation is selected. The Tomiyama correlation is suitable for liquid-gas bubble systems with a low Reynolds number and low Morton number flow, for example, small bubbles in an air-water system. In such cases, the Bozzano-Dente model over-predicts the drag coefficient. The Tomiyama drag coefficient has its own properties that determine the drag force, depending on how contaminated the bulk flow is by impurities, such as surfactants. Surfactants reduce the liquid surface tension and prevent internal circulation in bubbles, hence increasing drag.

What Is the Linearized Drag Coefficient?

The Linearized Drag Coefficient node, which has properties and is subject to common menu items, is a sub-node of the Drag Force node.



The linearized drag coefficient is used as a direct multiplier of the relative velocity between phases in Eqn. (1710).

Linearized Drag Coefficient Properties

Dimensions — The dimensionality of the drag coefficient (read-only).

Method — Selects the method to use for specifying the drag coefficient.

- Arastoopour — Suitable for solid dispersed-phase applications such as fluidized beds. The Arastoopour drag coefficient method node is added as a child. See Eqn. (1727). Only available when the dispersed phase is a solid.
- Constant — Defines the drag coefficient using a constant value. A constant method node is added as a child to this node.
- Emulsion Topology Inversion — Only available when the Emulsion Rheology model is selected.
  Allows the phases to invert when the volume fraction of the dispersed phase reaches the specified Inversion Volume Fraction.
  When the dispersed volume fraction becomes larger than the specified value, the drag, the relative viscosity, and the normal relative viscosities are swapped between the phases. After inversion, the relative viscosities are calculated using a new value for the maximum packing in the relative viscosity models.
- Field Function — Defines the drag coefficient using a field function. A field function method node is added as a child to this node.
- Gidaspow — Suitable for solid dispersed-phase applications such as fluidized beds. The Gidaspow drag coefficient method node is added as a child. See Eqn. (1728). Only available when the dispersed phase is a solid.

Suspension	Only available when the Suspension Rheology or Emulsion Rheology model is selected. A linearized form of the Schiller-Neumann drag model which allows A_D to become infinitely large when maximum packing is reached. The maximum-packed structure prevents slip from occurring between the phases.
Syamlal-O'Brien	Uses a drag formula that is based on terminal velocities in settling beds. The Syamlal O'Brien drag coefficient method node is added as a child. See Eqn. (1729). Only available when the dispersed phase is a solid.
Symmetric	Recommended for flows where phase topology inversion occurs. The Symmetric drag coefficient node is added as a child.

What Is the Arastoopour Drag Coefficient Method?

The Arastoopour drag coefficient method is provided for cases of high particle loading. Arastoopour node and its properties represent it.



Arastoopour Drag Coefficient Properties

Exponent	The exponent n in Eqn. (1727).

What Is the Gidaspow Drag Coefficient Method?

The Gidaspow drag coefficient method is provided for cases of high particle loading. The Gidaspow node and its properties represent it.



Gidaspow Drag Coefficient Properties

Exponent	The hindered settling exponent n in Eqn. (1728).
Transition Volume Fraction	When the particle volume fraction is below this value a modified form of Stoke's law is used to compute the drag, as shown in Eqn. (1728). The Ergun equation is used for loadings above this value.

What Is the Syamlal O'Brien Drag Coefficient Method?

The Syamlal O'Brien drag coefficient method is based on measurements of terminal velocity that is made in settling beds. The Syamlal O'Brien node and its properties represent it.



Syamlal O'Brien Drag Coefficient Properties

Transition Volume Fraction	The cut-off volume fraction α_{tr} in Eqn. (1732).

Syamlal O'Brien Drag Expert Properties

C1	Value that is used as the exponent of α_c above the cut-off volume fraction in Eqn. (1732).
C2	Value that is used as the multiplier of α_c below the cut-off volume fraction in Eqn. (1732).

What Is the Symmetric Drag Force Method?

The symmetric drag method is recommended for flows where phase topology inversion occurs, that is, the dispersed phase in one region of the domain becomes the continuous phase in another. The Symmetric node represents it. No properties are required as the model equations can access all the quantities that they require.



Drag Correction Properties

Dimensions	The dimensionality of the drag correction (read-only).
Method	Selects the method to use for specifying the drag correction.
	Field Function — Defines the drag correction using a field function. A field function method node is added as a child to this node.
	Volume Fraction Exponent — Defines the drag correction using the Volume Fraction Exponent. This correlation is available only when the continuous phase is viscous. It is suitable for spherical solid particles, liquid droplets, and small-diameter (spherical) bubbles.

Volume Fraction Exponents Properties

Volume Fraction Exponent (continuous)	Sets the value of n_v. It is included to make sure that the drag tends to zero as the amount of the continuous phase material tends to zero, given by Eqn. (1777)
Maximum Packing	This setting limits the maximum packing, $\alpha_{d,max}$, given by Eqn. (1778)

Tomiyama Drag Coefficient Properties

Contamination State	Selects the method to use for specifying contamination of the bulk flow by impurities, such as surfactants. Surfactants reduce the liquid surface tension and prevent internal circulation in bubbles, hence increasing drag. The Contaminated water option is the default option since it applies to most industrial cases.

Pure	Uses the drag coefficient that is described by Eqn. (1749).
Moderate	Specifies a higher drag coefficient than the Pure option, before bubble deformation effects become dominant. See Eqn. (1750).
Contaminated	Specifies a higher drag coefficient than the Moderate and Pure options to account for the presence of surfactants. See Eqn. (1751).

Modeling Lift Force

This section explains the following about the Lift Force phase interaction model:

- Selecting the Lift Force Model
- What Is the Lift Force?
- What Is the Lift Coefficient?
- What Is the Lift Correction?

Selecting the Lift Force Model

Select Lift Force from the *Optional Phase Interaction Models* group box in the *Phase Interaction Model Selection* dialog as shown.



What Is the Lift Force?

The Lift Force node, which is subject to common menu items, represents the Lift Force model for calculating interphase momentum transfer. This models the perpendicular lift force that a particle can experience in a non-uniform or swirling flow field.



The Lift Force model requires a Lift Force Coefficient.

Lift Force Properties

None.

What Is the Lift Coefficient?

The Lift Coefficient node, which has properties and is subject to common menu items, is a sub-node of the Lift Force node. This value supplies the lift coefficient in Eqn. (1810).



The lift coefficient can be set as a constant value, or by a field function, or by Tomiyama method.

Lift Coefficient Properties

Dimensions	Dimensionless.	
Method	Selects the method to use for specifying the lift coefficient.	
	Constant	Defines the lift coefficient using a constant value. A constant method node is added as a child to this node.
	Field Function	Defines the lift coefficient using a field function. A field function method node is added as a child to this node.
	Tomiyama	Defines the standard lift coefficient using the Tomiyama correlation. It is only available when the continuous phase is viscous, and when the dispersed phase is a fluid (it is not available for solid particles). Additionally, the correlation requires that gravity is activated for the physics continuum. Specify the surface tension value in the Multiphase Material > Material Properties > Surface Tension node for the parent phase interaction. The Surface Tension node appears automatically when the correlation is selected.

Lift Coefficient Expert Properties

Lift Force Coefficient Correction	On	Activates the lift correction. For more information, see What Is the Lift Correction?.
	Off	Deactivates the lift correction.

If the Drag Force Coefficient Type is set to Standard (the default setting), the Lift Force Coefficient Correction option is activated by default when you select the Lift Force model. However, if the Drag Force Coefficient Type is set to Linearized, the Lift Force Coefficient Correction property is deactivated and cannot be set.

What Is the Lift Correction?

The Lift Correction is used in the Lift Force model to improve numerical stability in the simulation.

When you use drag correction in a simulation where the Lift Force model is also used, the reduction of drag force (see Eqn. (1710) and Eqn. (1711)) can render the simulation numerically unstable. To resolve this, the lift correction is used to account for multi-particle effects by modifying the single-particle Lift Coefficients.

The lift correction can be approximated by the drag correction, or specified by a field function. For more information, see Lift Force.

The Lift Correction node, which has properties and is subject to common menu items, is a sub-node of the Lift Force node.



Lift Correction Properties

Dimensions	Dimensionless.
Method	Selects the method to use for specifying the lift correction.
Drag Correlated Lift Correction	Uses the drag coefficient correction to approximate the lift coefficient correction. This approximation is useful for cases when swarming is modeled.
Field Function	Defines the lift correction using a field function. A field function method node is added as a child to this node.

Modeling Wall Lubrication Force

This section explains the following about the Lift Force phase interaction model:

- What Is the Wall Lubrication Force?
- What Is the Wall Lubrication Coefficient?
- Selecting the Wall Lubrication Force

What Is the Wall Lubrication Force?

The wall Lubrication Force node, which is subject to common menu items, represents the Wall Lubrication model for calculating interphase momentum transfer. This models the force that a rising bubble experiences that prevents it from touching the wall.



Wall Lubrication Force Properties

None.

What Is the Wall Lubrication Coefficient?

The Wall Lubrication Coefficient node, which has properties and is subject to common menu items, is a sub-node of the Wall Lubrication Force node. This supplies the wall lubrication coefficient in Eqn. (1814).



The wall lubrication coefficient can be set using the Antal equation or by using a field function.

Wall Lubrication Properties

Dimensions — The dimensionality of the wall lubrication coefficient (read-only).

Method — Selects the method to use for specifying the wall lubrication coefficient.

- Antal — Defines the standard lift coefficient using the Antal correlation. An Antal node is added as a child to this node where you can specify the coefficients for C_{w1} and C_{w2} in Eqn. (1816). This is the default setting.
- Field Function — Defines the lift coefficient using a field function. A field function method node is added as a child to this node where you can add a function for C_{WL} in Eqn. (1814).

Selecting the Wall Lubrication Force

To select the Wall Lubrication Force model:

- Select Lift Force from the *Optional Phase Interaction Models* group box in the *PhaseInteraction Model Selection* dialog.

The Wall Lubrication Force becomes available in the *Optional Phase Interaction Models* group box as shown.



Modeling Turbulent Dispersion Force

This section explains the following about the Turbulent Dispersion Force model:

- What Is the Turbulent Dispersion Force?
- Selecting the Turbulent Dispersion Force

Also see the formulation section.

What Is the Turbulent Dispersion Force?

The Turbulent Dispersion Force node represents the Turbulent Dispersion Force model that accounts for the interaction between the dispersed phase and the surrounding turbulent eddies. This node has properties and is subject to common menu items. A turbulence mode must be selected before you can apply this option.



Turbulent Dispersion Force Properties

Calibration Constant C_0 in Eqn. (1788), with a default value of 1.

Turbulent Dispersion Force Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model, do not change the *Expert* settings.

Use Logarithmic Form

Analytically, the logarithmic form (Eqn. (1787)) is equivalent to the basic form (Eqn. (1786)), but it has a much simpler discretization on a collocated grid. This simplicity makes it more reliable whenever extreme volume fractions occur in neighboring cells -- for example, due to condensation or separation processes. In such applications, this reliability in turn improves convergence history.

Activated: Use Eqn. (1787).

Deactivated: Use Eqn. (1786), the default.

Use Volume Fraction Smoothing

This option reduces numerical checkerboarding in the presence of dominant turbulent dispersion forces and hence improves convergence. It does not affect the converged solution.

Activated: Use smoothing.

Deactivated: No smoothing, the default.

Use Tensorial Diffusivity

Where a Reynolds Stress model has been selected for the continuous phase, activating this option makes the simulations use an anisotropic model for turbulent dispersion (Eqn. (1789)). If this option is not active, or if a turbulence model based on the eddy viscosity hypothesis is selected for the continuous phase, then the simulation uses the isotropic model (Eqn. (1788)) instead.

Activated: Use Eqn. (1789).

Deactivated: Use Eqn. (1788), the default.

Turbulent Dispersion Force Prandtl Number

Dimensions — The dimensionality of the Prandtl number (read only).

Method — Selects the method to use for specifying the Prandtl number, σ_α in Eqn. (1790).

- Constant — Takes a default value of 1, as for a passive scalar.
- Field Function — Defines the Prandtl number using a field function. A field function method node is added as a child to this node.
- Inertial Correction — Recommended for large particles or bubbles.
- Tchen Correction — Recommended for small particles or bubbles.

Selecting the Turbulent Dispersion Force

Select Turbulent Dispersion Force from the *Optional Phase-Pair Interaction Models* group box in the *Phase Interaction Model Selection* dialog as shown.



The dialog shows the Turbulent Dispersion Force model activated.



Modeling Virtual Mass Force

This section explains the following about the Virtual Mass Force model:

- What Is the Virtual Mass Force?
- Selecting the Virtual Mass Force

What Is the Virtual Mass Force?

The Virtual Mass Force node, which has properties and is subject to common menu items, represents the Virtual Mass Force model for calculating interphase momentum transfer. The virtual mass force accounts for the additional resistance that is experienced by a particle undergoing acceleration, and can therefore be significant in accelerating flows. Examples include short bubble columns or applications with swirl-promoting devices.



Virtual Mass Coefficient Properties

Virtual mass coefficient	Defines the virtual mass coefficient that appears in Eqn. (1794).
Dimensions	The dimensionality of the virtual mass coefficient (read-only).
Method	Selects the method to use for specifying the lift coefficient.

Constant	Defines the virtual mass coefficient using a constant value. A constant method node is added as a child to this node.
Field Function	Defines the virtual mass coefficient using a field function. A field function method node is added as a child to this node.
Spherical Particle	Defines the virtual mass coefficient using the Spherical Particle method. A Spherical Particle node is added as a child to this node. This option is the default method.
Zuber	Defines the virtual mass coefficient using the Zuber method. A Zuber node is added as a child to this node.

Spherical Particle Properties

Free Stream Minimum Fraction	Defines the freestream fraction that appears in Eqn. (1805). The default value is `0.0`

Zuber Properties

Free Stream Minimum Fraction	Defines the freestream fraction that appears in Eqn. (1809). The default value is `-10.0`

Selecting the Virtual Mass Force

Select Virtual Mass Force from the *Optional Phase Interaction Models* group box in the *Phase Interaction Model Selection* dialog as shown.



Providing Additional Material Properties

The Multiphase Material node is used to hold any additional material property nodes that are required by the phase interaction models.



Currently, two more property nodes can be added, as follows:

- Surface tension -- required when the Bozzano-Dente drag law, Tomiyama drag law, or Tomiyama lift law have been selected.
- Latent heat -- required when the bulk boiling or wall boiling models have been selected.

What Is the Granular Dispersed-Dispersed Model?

If you have activated the Granular Pressure model and have created more than one particle, it is necessary to define an interaction between the two particle phases. This interaction is defined on the properties of the Granular Dispersed-Dispersed model node.



Granular Dispersed-Dispersed Properties

Phase 1	Choose a particle phase of the interacting pair from the drop-down list of previously defined Eulerian phases. The list contains the particle phases only: you cannot select another type of phase.
Phase 2	Choose a different particle phase of the interacting pair than you defined in Phase 1.

Granular Dispersed-Dispersed Reference Value

Coefficient of restitution	Defines the coefficient of restitution for collisions between particles in the phase interaction. This value, e, is used in the kinetic regime formulation. The default value of e = 0.9.

Once the Granular Dispersed-Dispersed model has been chosen, select the models for:

- Granular Drag Force

Modeling Granular Drag Force

This section explains the following about the Granular Drag Force model that is used in particle-particle phase interactions:

- What is the Granular Drag Force?
- What is the Linearized Drag Coefficient?
- Selecting the Granular Drag Force

What is the Granular Drag Force?

The Granular Drag Force node, which has properties and is subject to common menu items, represents the Granular Drag Force model. It is related to the Drag Force model for continuous-dispersed phases.



In STAR-CCM+, the Granular Drag Force is modeled as a linear multiplier, A_D in Eqn. (1710), of the relative velocity between two phases.

Granular Drag Force Properties

Drag Force Coefficient Type	The type of coefficient to use.	
	Linearized	Specifies the linearized drag force coefficient.

What Is the Linearized Drag Coefficient?

The linearized drag coefficient is used as a direct multiplier of the relative velocity between phases in Eqn. (1710). The Linearized Drag Coefficient node, which has properties and is subject to common menu items, is a sub-node of the Granular Drag Force node.



Linearized Drag Coefficient Properties

Dimensions — The dimensionality of the drag coefficient (read-only).

Method — Selects the method to use for specifying the drag coefficient.

- Constant — Defines the standard drag coefficient using a constant value.
  A constant method node is added as a child to this node.
- Field Function — Defines the standard drag coefficient using a field function.
  A field function method node is added as a child to this node.
- Gera-Syamlal — Defines the standard drag coefficient using the Gera-Syamlal granular drag force model.
  The Gera-Syamlal node is added as a child to this node. This node lets you define the coefficient of friction for the phase interaction.

Gera-Symlal Properties

Coefficient of friction — Defines the coefficient of friction between particles for the phase interaction.
This value, c_f, is used in the Gera-Syamlal granular drag force formulation.
The default value for c_f is 0.

Selecting the Granular Drag Force

The Granular Drag Force becomes available when the Granular Dispersed-Dispersed model is selected. If the Auto-select recommended models option is activated, it is selected automatically. Otherwise, it can be selected from the *Optional Phase Interaction Models* group box as shown.



Modeling Interphase Mass Transfer

Four options are provided for modeling mass transfer between Eulerian phases:

- Bulk boiling, which is defined by the boiling mass transfer rate model.
- Diffusion mass transfer, which is defined by the dissolution mass transfer rate model.
- Interphase reactions, which are defined by the interphase reaction model.
- User-defined mass transfer, which is defined by the user-defined mass flux model.

Documentation for bulk boiling is provided within the section, Modeling the boiling process. This section contains information on the dissolution mass transfer rate, as follows:

- What Is the Dissolution Mass Transfer Rate?
- What is the Continuous Phase Sherwood Number?
- What is the Dispersed Phase Sherwood Number?
- What is the Equilibrium Coefficient?
- Selecting the Dissolution Mass Transfer Rate
- What is the User-Defined Mass Flux

What Is the Dissolution Mass Transfer Rate?

The Dissolution Mass Transfer Rate node, which has properties and is subject to common menu items, represents the Dissolution Mass Transfer Rate model. This model calculates the mass of a component that is transferred from one multi-component phase to another.



In STAR-CCM+, the rate of transfer of species from one phase to another is governed by Eqn. (2030).

Dissolution Mass Transfer Rate Properties

Phase 0 Components	Lets you specify the relationship between the components being transferred between phases.

The *Dissolution Mass Transfer Rate - Phase 0 Components* dialog lets you specify the components to transfer. The components of the continuous phase are activated on the left-hand column. The components of the dispersed phase can be matched in the right-hand column as shown:



For each transfer activated, the continuous phase component name appears under the Equilibrium Coefficient > Composite node:



What is the Continuous Phase Sherwood Number?

The Continuous Phase Sherwood Number node, which has properties and is subject to common menu items, is a sub-node of the Dissolution Mass Transfer Rate node that is used to calculate the mass transfer coefficient, which is given by Eqn. (2036).



Continuous Phase Sherwood Number Properties

Dimensions	The dimensionality of the Sherwood Number (read-only).
Method	Selects the method.
	Constant — A constant method node is added as a child to this node.
	Field Function — A field function method node is added as a child to this node.
	Ranz Marshall — A Ranz Marshall method node that represents use of the Ranz-Marshall correlation (Eqn. (1937)) is added as a child to this node.

What is the Dispersed Phase Sherwood Number?

The Dispersed Phase Sherwood Number node, which has properties and is subject to common menu items, is a sub-node of the Dissolution Mass Transfer Rate node that is used to calculate the mass transfer coefficient, which is given by Eqn. (2036).



Dispersed Phase Sherwood Number Properties

Dimensions	The dimensionality of the Sherwood Number (read-only).
Method	Selects the method.
	Constant — A constant method node is added as a child to this node.
	Field Function — A field function method node is added as a child to this node.

What is the Equilibrium Coefficient?

The Equilibrium Coefficient node, which has properties and is subject to common menu items, is a sub-node of the Dissolution Mass Transfer Rate node.



The Equilibrium Coefficient node lets you specify the appropriate equilibrium coefficient:

- For gas-liquid interactions, the ratio of gas to liquid mole fractions, which are given by Eqn. (2028).
- For liquid-liquid interactions, the ratio of continuous to dispersed mole fractions, which are given by Eqn. (2032).

Equilibrium Coefficient Properties

Dimensions — The dimensionality of the Equilibrium Coefficient (read-only).

Method

- Composite — Read-only manager for the Equilibrium Coefficients that are activated for each component transfer. The Composite node does not have properties.

Equilibrium Coefficient for Individual Component Transfer Properties

Dimensions — The dimensionality of the Equilibrium Coefficient method (read-only).

Method — Selects the method.

For liquid-liquid interactions, the equilibrium coefficient specification is limited to Constant or Field Function.

- Constant — A constant method node is added as a child to this node. This node can be used to set a constant value.
- Field Function — A field function method node is added as a child to this node. This node can be used to select an appropriate function.
- Henry's Law — A Henry's Law method node is added as a child to this node. This node can be used to sea value for Henry's Law Coefficient, which are given by Eqn. (2039).
- Raoult's Law — A Raoult's Law method node is added as a child to this node. This child node does not have properties. A Saturation Pressure node is added to the liquid phase. This node can be used as an input to Eqn. (2038).

Selecting the Dissolution Mass Transfer Rate

The Dissolution Mass Transfer Rate model can be selected from the *Interphase Mass Transfer Rate* group box. To make it available:

- Define both of the phases in the interaction to have more than one component--usually one multi-component liquid and one multi-component gas.
- Specify the continuous and dispersed phases.
- Select the Interphase Mass Transfer model.



Note: If you select the Dissolution Mass Transfer Rate model, and the Phase Coupled Fluid Energy model is already activated, then select the Interphase Energy Transfer phase-interaction model to make sure that the phase temperatures reach sensible values. The Dissolution Mass Transfer Rate rate model assumes that the effect of interphase mass transfers on the phase temperatures is negligible.

What Is the User-Defined Interphase Mass Flux?

This is a basic model for interphase mass transfer that is not governed by energy or species transfer. It is the only model available when the Phase Coupled Fluid Energy model is not active. An example of where this model can be useful is a simulation of bubble dissolution into de-aerated water.

When selected, the User-Defined Interphase Mass Flux node represents the model, which has the Interphase Mass Flux as its child node. The Method property of this node is used to specify how the transfer rate is supplied.



Note: If you select the User-Defined Interphase Mass Flux phase-interaction model, and the Phase Coupled Fluid Energy model is already activated, then select the Interphase Energy Transfer phase-interaction model to make sure that the phase temperatures reach sensible values. The User-Defined Interphase Mass Flux model assumes that the effect of interphase mass transfers on the phase temperatures is negligible.

Properties of the User-Defined Interphase Mass Flux

Method	Select the method to use in providing the value of mass flux.	
	Constant	Adds a node on which to set the constant value.
	Field Function	Adds a node on which to select a field function.

Modeling Interphase Reactions

You can model mass transfer between Eulerian phases as a reaction between the phases.

To make the Interphase Reaction model available:

- Define both of the phases in the interaction.

  The phases must have more than one component--usually one multi-component liquid and one multi-component gas.
- Specify the continuous and dispersed phases.
- Select the Interphase Mass Transfer model.

The Interphase Reaction model can be selected from the *Interphase Mass Transfer Rate* group box.



Setting Up Interphase Reactions

The Interphase Reaction node contains a Reactions sub-node. The Reactions sub-node for the Interphase Reaction model is shown below.



Rather than using properties, this sub-node has a pop-up menu with which you can create objects for setting up interphase reactions.

Reactions Pop-Up Menu

New Reaction	Creates an interphase reaction. This new object contains sub-nodes for reactants, products, and the reaction coefficient.
Import Reactions	Lets you import a reaction from the Material Database.

See also: common menu items.

You use the Reactions node to define the reactions that occur between the phase components.

To set up an interphase reaction:

1. Right-click the Reactions node and select New Reaction.

A new reaction definition is added.



2. In the appropriate nodes, specify the reactants and products.

   This procedure is similar to the corresponding procedure for Lagrangian particle reactions. For more information, see Setting Up Particle Reactants and Products.

3. Specify the properties of the reaction components.

4. In the Properties > Reaction Coefficient node, specify the appropriate reaction rate properties.

   This procedure is similar to the corresponding procedure for Lagrangian particle reactions. For more information, see Defining Reaction Rates.

Specifying the Reaction Component Parameters

A reaction component is any constituent of a multi-component mixture that has already been defined as a reactant or product in an interphase reaction equation.



To specify the required parameters for a particular reaction component:

1. Select the component.
2. In the *Properties* window, specify the appropriate values:

Stoich. Coeff.	Set the stoichiometric coefficient that is appropriate to the current reaction.
Rate Exponent	Set the rate exponent that is appropriate to the current reaction.

Modeling Interphase Energy Transfer

When the Phase Coupled Fluid Energy model is active, you can choose to model energy transfer between Eulerian phases. Activating the Interphase Energy Transfer model for a phase interaction permits this energy transfer between the two phases that are selected on the interaction. The Interphase Energy Transfer node represents this model. It has a single sub-node for the Nusselt number.



What Is the Nusselt Number?

The Method property of the Nusselt Number node is used to provide the necessary input for Eqn. (1936).

Properties of the Nusselt Number

Method	Selects the method to use in specifying the Nusselt number that is required by Eqn. (1936).	
	Constant	Adds a Constant child node on which to set the constant value.
	Field Function	Adds a Field Function child node from which to select an appropriate function.
	Ranz Marshall	Adds the Ranz Marshall child node that represents use of the Ranz-Marshall correlation that is given in Eqn. (1937).

Working with Segregated Multiphase Flow Solvers

Two solvers are invoked when the Multiphase Segregated Flow model is selected:

- Multiphase segregated flow solver
- Volume fraction solver

When setting up a multiphase flow case it is often necessary to adjust the solver settings. Recommended settings are discussed in the section Setting Up Solvers for Multiphase Flow Cases.

What Is the Multiphase Segregated Flow Solver?

This solver, which has properties and is subject to common menu items, controls the solution update for the Multiphase Segregated Flow model. It is responsible for solving the velocity field for each individual phase and the overall pressure solution. It provides two more solvers:

- Phase Coupled Velocity
- Pressure



Solver Expert Properties

The properties in the *Expert* category are for a temporary debug situation, at the expense of simulation accuracy and higher memory usage. Do not change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.
Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Update reconstruction gradients with each iteration.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.	
	Activated	Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you deactivate the checkbox again.
	Deactivated	Do not reset reconstruction gradients. Gradients are recomputed at the next iteration if they had been zeroed.
Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.

Interaction Force Storage Retained

Allows you to retain the information that is stored for each phase interaction force model at the end of the iteration. These quantities become available as field functions during subsequent iterations.
You can use this option when you want to visualize the phase interaction forces, such as drag, turbulent drag, lift, virtual mass, and wall lubrication.

Activated — Retain the phase interaction force storage.

Deactivated — Do not retain the phase interaction force storage.

What Is the Phase Coupled Velocity Solver for Multiphase Flow?

The phase coupled velocity solver solves the discrete conservation equations to obtain the intermediate velocity field for all the phases. The corresponding Phase Coupled Velocity Solver node controls the under-relaxation factor and algebraic multigrid parameters for the Eulerian phase momentum equations.



The equations for velocity component correction are solved for all phases simultaneously. One advantage of this solver is that it permits greater robustness in solving multiphase flows with high levels of phase interaction.

Phase Velocity Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7. For the theoretical background, see Implicit Iteration.

What Is the Pressure Correction Solver for Multiphase Flow?

Controlled by the multiphase segregated flow solver, the pressure solver controls the under-relaxation factor and algebraic multigrid parameters for the pressure correction equation.



Phase Velocity Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.3.

What Is the Volume Fraction Solver for Multiphase?

This solver, which has properties and is subject to common menu items, controls the solution update for the volume fractions of all the phases. The discretized form of the continuity equation (Eqn. (1692)), is solved for each phase.



To make sure that Eqn. (1693) is satisfied after solving the volume fraction equations for all phases, the volume fraction of each phase in a cell is normalized by the sum of the volume fractions of all the phases in that cell.

Volume Fraction Solver Properties

Explicit Relaxation Factor	This property specifies the multiplier that is applied to the provisional increment of the solution. The default value is 1.0.
Under-Relaxation Factor	At each iteration, this value is supplied to the AMG Linear Solver component of the Volume Fraction solver. This value enhances diagonal dominance of the matrix and restricts the provisional increment of the solution. The default value is 0.5.

The default settings use only the implicit under-relaxation factor for robustness. In some scenarios, such as a fine mesh, it can take a while for the Volume Fraction to propagate through the system. In these cases, you can split the under-relaxation property value between the two factors. The product of the two factors (that is, Under-Relaxation Factor times Explicit Relaxation Factor) defines the under-relaxation setting. For example, instead of using Under-Relaxation Factor = `0.15`, you can specify the combination Under-Relaxation Factor = `0.6`, Explicit Relaxation Factor = `0.25`. This combination can help the volume fraction to converge faster without upsetting the overall convergence of the simulation.

Note: The same approach can be used for S-Gamma under-relaxation. In general, it is good practice to use the same under-relaxation factors for the Volume Fraction and S-Gamma equations since their ratio is a key parameter in interphase interactions.

Volume Fraction Solver Expert Properties

The properties in the *Expert* category are for a temporary debug situation, at the expense of simulation accuracy and higher memory usage. Do not change these properties unless you are thoroughly familiar with the STAR-CCM+ discretization techniques.

Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Update reconstruction gradients with each iteration.

Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) is first-order estimates.	
	Activated	Reconstruction gradients are maintained at zero.
	Deactivated	Reconstruction gradients are not maintained at zero.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Setting Up Solvers for Multiphase Flow Cases

For multiphase flow cases, it is recommended that the default values for the solver under-relaxation factors are changed as follows:

- For velocity, turbulence, turbulent viscosity, and energy solvers, decrease the under-relaxation factor to 0.3
- For the pressure solver, decrease the under-relaxation factor to 0.1

The under-relaxation factor of each solver can be adjusted by selecting the solver node, then changing the value of the Under-Relaxation property in the *Properties* window.

- For the volume fraction and S-Gamma solvers, adjust the Explicit Relaxation Factor and the Under-Relaxation Factor property settings to decrease the product of the two values to 0.1

In some cases, decrease the under-relaxation values further to reach a converged solution. For example, decrease the under-relaxation to 0.1 for the velocity, energy and turbulence solvers, and 0.01 for the pressure and volume fraction solvers.

Additionally, it can be necessary to decrease the solver convergence tolerance. This setting is a property of the AMG Linear Solver component of each solver (the following example shows where this property is located for the Pressure solver). The default value for the Convergence Tolerance is 0.1, but this value can be reduced to between 1.0E-3 and 1.0E-4 in most cases.



Modeling Turbulence in Multiphase Segregated Flow

The presence of a dispersed phase can modify the turbulence of a continuous phase in a number of ways:

- Shear production is modified through the effect of buoyancy, drag, and other forces on velocity profiles.
- The passage of the dispersed phase particles generates velocity fluctuations in the continuous phase. The wake of the dispersed phase particles is a source of turbulence.
- The interaction of velocity fluctuations and volume fraction fluctuations can generate further turbulence enhancement or damping terms.

STAR-CCM+ allows you to model the effects of turbulence on continuous and dispersed phases in several ways:

- Continuous and dispersed phases can each be modeled with their own set of equations for the turbulence energy and dissipation. You can model phases independently as laminar, turbulent or with different turbulence models.

  The turbulence equations are identical to the single phase formulation, except each phase is scaled with a factor of the volume fraction of that phase. The same phase volume fraction also modifies convection and diffusion fluxes of phase-turbulence quantities.
- The turbulence of the dispersed phase can be calculated from the turbulence of the continuous phase using the Turbulence Response Model. This algebraic model reduces the computations that are required as only the continuous phase equations are solved. However, the model couples the turbulence of dispersed phases to the continuous ones.
- The phases can also be coupled through bubble/particle induced turbulence of the dispersed phase which adds more source terms to the continuous phase turbulence energy and dissipation equations.

Continuous Phase Turbulence Models

The models available for modeling turbulence in the continuous phase are:

- Standard K-Epsilon Model, with High-y+ Wall Treatment
- Standard Two-Layer K-Epsilon Model, with Two-Layer Wall Treatment
- Realizable K-Epsilon Model, with Two-Layer Wall Treatment
- Linear Pressure Strain Model, with High-y+ Wall Treatment

- Linear Pressure Strain Two-Layer Model, with Two-Layer Wall Treatment
- Quadratic Pressure Strain Model, with High-y+ Wall Treatment



You can expand the phase-based turbulence models to use extra bubble/particle induced turbulence models that exist at the phase interaction level:

- Gosman
- Troshko Hassan
- Tchen



Disperse Phase Turbulence Models

The models available for modeling turbulence in the dispersed phase are:

- Laminar
- Turbulence Response
- Standard K-Epsilon Model, with High-y+ Wall Treatment
- Standard Two-Layer K-Epsilon Model, with Two-Layer Wall Treatment
- Realizable K-Epsilon Model, with Two-Layer Wall Treatment
- Linear Pressure Strain Model, with High-y+ Wall Treatment
- Linear Pressure Strain Two-Layer Model, with Two-Layer Wall Treatment
- Quadratic Pressure Strain Model, with High-y+ Wall Treatment



If you expect the dispersed phase to carry a significant proportion of the wall stress or wall heat transfer, do not use the Laminar model.

Note: These models were derived for two-phase flow, but are also applicable for dilute mixtures of a number of dispersed phases.

Turbulence Response Model

The Turbulence Response Model predicts velocity fluctuations in the dispersed phase using algebraic correlations to the velocity fluctuations in the continuous phase (see Eqn. (1781)).



Selecting Turbulence Response

The Particle Induced Turbulence option becomes available within the *Reynolds-Averaged Turbulence Models* box of the *Phase Model Selection* dialog after you have chosen the Reynolds-Averaged Navier-Stokes model. Click this option as shown below.



The *Turbulence Response Models* box appears and you can select one of the following:

- Issa Turbulence Response Model
- User-Defined Turbulence Response Model
- Tchen Turbulence Response Model

Once you have selected one of the turbulence response models, a Response Coefficient Turbulent Viscosity node is added to the Solvers node.

Issa Turbulence Response Model

The Issa Turbulence Response model is defined as a correlation for the turbulence response coefficient with a volume fraction correction.



Issa Turbulence Response Model Properties

Reference Continuous Phase	Select the continuous phase of the multiphase flow.

Issa Turbulence Response Model Expert Properties

Normal Stress Term	An explicit term that directly incorporates divergence and turbulent kinetic energy into the stress tensor according to the full Boussinesq approximation. This term often has a negligible effect in incompressible single-phase flows, but use it with care in Eulerian multiphase flows where it can cause a strong turbophoretic effect. Activated — Use the term. Deactivated — Do not use the term (the default).

Selecting Issa Turbulence Response Model

The Issa Turbulence Response option becomes available within the *Turbulence Response Models* box of the *Phase Model Selection* dialog after you have chosen the Turbulence Response model. Click this option as shown below.



User-Defined Turbulence Response Model

The User-Defined Turbulence Response model allows you to define the turbulence response coefficient. When the turbulence response coefficient is equal to 1, the dispersed phase fluctuates in tandem with the continuous phase.



User-Defined Turbulence Response Model Properties

Reference Continuous Phase	Select the continuous phase of the multiphase flow.

Uawe-Defined Turbulence Response Model Expert Properties

Normal Stress Term	An explicit term that directly incorporates divergence and turbulent kinetic energy into the stress tensor according to the full Boussinesq approximation. This term often has a negligible effect in incompressible single-phase flows, but use it with care in Eulerian multiphase flows where it can cause a strong turbophoretic effect. Activated — Use the term. Deactivated — Do not use the term (the default).

Turbulence Response Model Properties

Method	Selects the method to use in specifying the turbulence response coefficient.

Selecting User-Defined Turbulence Response Model

The User-Defined Turbulence Response option becomes available within the *Turbulence Response Models* box of the *Phase Model Selection* dialog after you have chosen the Turbulence Response model. Click this option as shown below.



Response Coefficient Turbulent Viscosity Solver

The Response Coefficient Turbulent Viscosity node is added to the Solvers node when a turbulence response model is selected. It governs the update of the turbulent viscosity resulting from the presence of a dispersed phase.



Response Coefficient Turbulent Viscosity Solver Properties

Under-Relaxation Factor	Governs the portion of the newly computed turbulent viscosity field that is used to update the turbulent viscosity. The default value of 1 implies that the turbulent viscosity is entirely computed from the new field.
Maximum Ratio	Sets the maximum allowable ratio of turbulent to laminar viscosity. By default it is set to 100,000.

Solver Expert Properties

The properties in the *Expert* category are for a temporary debug situation, at the expense of simulation accuracy and higher memory usage. Do not change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Tchen Turbulence Response Model

The Tchen Turbulence Response model provides a dispersed phase turbulent diffusivity for gas flows laden with heavy particles (see Eqn. (1909)).



Tchen Turbulence Response Model Properties

Turbulent Prandtl Number	The coefficient σ_0 in Eqn. (1790), with a default value of 1.0, representing basic passive diffusivity.
Crossing Trajectories Coefficient	The calibration coefficient C_β in Eqn. (1869), with a default value of 1.8.
Reference Continuous Phase	Select the continuous phase of the multiphase flow.

Tchen Turbulence Response Model Expert Properties

Normal Stress Term

An explicit term that directly incorporates divergence and turbulent kinetic energy into the stress tensor according to the full Boussinesq approximation. This term often has a negligible effect in incompressible single-phase flows, but use it with care in Eulerian multiphase flows where it can cause a strong turbophoretic effect.

Activated — Use the term.

Deactivated — Do not use the term (the default).

Selecting the Tchen Turbulence Response Model

The Tchen Turbulence Response option becomes available within the *Turbulence Response Models* box of the *Phase Model Selection* dialog after you have chosen the Turbulence Response model. Click this option as shown below.



Particle Induced Turbulence Models

The particle induced turbulence models add more source terms to the continuous phase turbulence kinetic energy and dissipation equations. The models represent modification of turbulence due to the presence of a dispersed phase.

The models are:

- Gosman model
- Troshko and Hassan model
- Tchen model



Selecting Particle Induced Turbulence

The Particle Induced Turbulence option becomes available within the *Optional Models* box of the *Phase Interaction Model Selection* dialog after you have chosen the Eulerian Continuous-Dispersed model. Click this option as shown below.



Gosman Particle Induced Turbulence Model

The Gosman Particle Induced Turbulence model adds source terms to the turbulent kinetic energy equation (see Eqn. (1917)) and to the turbulent dissipation equation (see Eqn. (1921)).



Selecting the Gosman Particle Induced Turbulence Model

To select the Gosman Particle Induced Turbulence model, do the following:

1. In the *Physics Model Selection*, select the Turbulent viscous regime at the continuum level.
2. Create a continuous phase and select a turbulence model.
3. Create a dispersed phase and select a turbulence model.
4. Set the Reference Continuous Phase property to the continuous phase in the Issa Turbulence Response node or in the User-Defined Turbulence Response node that you chose.
5. In the *Phase Interaction Model Selection* dialog, select the Turbulence Response model, the Turbulent Dispersion Force model and the Particle Induced Turbulence from the *Optional Models* box.

6. The Gosman Particle Induced Turbulence option becomes available within the *Particle Induced Turbulence Models* box of the *Phase Interaction Model Selection* box.



Troshko and Hassan Particle Induced Turbulence Model

The Troshko and Hassan Particle Induced Turbulence model describes bubble induced turbulence effects. It adds a source term to the turbulent kinetic energy equation (see Eqn. (1911)) and another to the dissipation equation (see Eqn. (1913)). The model uses the Virtual Mass phase interaction model but if the model is not selected the virtual mass coefficients default to the value of a spherical bubble, which is given by Eqn. (1910).



Properties of the Troshko and Hassan Particle Induced Turbulence Model

Troshko and Hassan C3 coefficient	Selects the calibration constant C_3 for the dissipation term in Eqn. (1913). By default this constant is set to 0.45.
Damping coefficient	This parameter is not part of the model derivation, but defaults to 1 and can be used by the user to regulate the proportion of the source term in Eqn. (1912) that is used in the calculation.

Selecting the Troshko and Hassan Particle Induce Turbulence Model

To select the Troshko and Hassan Particle Induced Turbulence model, do the following:

1. In the *Physics Model Selection*, select the Turbulent viscous regime at the continuum level.
2. Create a continuous phase and select a turbulence model.
3. Create a dispersed phase and select a turbulence model.
4. In the *Phase Interaction Model Selection*, select the Particle Induced Turbulence from the *Optional Models* box.

The Troshko and Hassan Particle Induced Turbulence option becomes available within the *Particle Induced Turbulence Models* box of the *Phase Interaction Model Selection* box.



Tchen Particle Induced Turbulence Model

The Tchen Particle Induced Turbulence model provides sources terms to the continuous turbulent kinetic energy equation due to the interaction with particles, based on Tchen's theory. See the formulation section.



Properties of the Tchen Particle Induced Turbulence Model

Crossing Trajectories Coefficient	The calibration coefficient C_β in Eqn. (1869), with a default value of 1.8.
Production Calibration Coefficient	The coefficient C_0 in Eqn. (1922) defining the source term for turbulence energy or for Reynolds stress equations. This parameter is not part of the model derivation, but defaults to 1 and you can adjust it to test the sensitivity of the solution to this term.
Dissipation Calibration Coefficient	The coefficient $C_{\varepsilon 3}$ in Eqn. (1924) defining the corresponding source term for the tubulent dissipation equation. Parameter $C_{\varepsilon 3}$ has a default value of 1.44, and you can adjust it independently of constants for other terms.
Turbulent Prandtl Number	The coefficient σ_0 in Eqn. (1790), with a default value of 1.0, representing basic passive diffusivity.

Selecting the Tchen Particle Induced Turbulence Model

To select the Tchen Particle Induced Turbulence model, do the following:

1. In the *Physics Model Selection*, select the Turbulent viscous regime at the continuum level.
2. Create a continuous phase and select a turbulence model.
3. Create a dispersed phase and select a turbulence model.
4. In the *Phase Interaction Model Selection* dialog, select the Turbulence Response model, the Turbulent Dispersion Force model and the Particle Induced Turbulence Models from the *Optional Models* box.

The Tchen Particle Induced Turbulence option becomes available within the *Particle Induced Turbulence Models* box of the *Phase Interaction Model Selection* dialog box.



Modeling Energy in Multiphase Segregated Flow

Energy conservation within a multiphase flow can be modeled using the Phase Coupled Fluid Energy model. This model solves the energy conservation equation for each Eulerian phase within the continuum.



When you select this model for the physics continuum, also choose the particular form of energy equation that is required in each Eulerian phase.

Solving for Enthalpy

The enthalpy field can be of interest in some applications, especially those involving phase change. In such applications, it is recommended that you also set the Heat of Formation and Standard State Temperature properties for each phase material. These two parameters, along with the specific heat of the phase, allow the enthalpy for a particular phase to be fully defined. Consistent definitions of phase enthalpies (that is with phase reference enthalpies defined against a reference temperature that is common for all phases) also correctly defines the magnitude and sign of enthalpy of phase.

These reference condition inputs have been named Heat of Formation and Standard State Temperature as commonly used in reacting flows. However, any other consistent set of reference data can be used instead.

For phase change calculations, a typical usage is to specify phase enthalpies at the saturation temperature corresponding to the system pressure. Example values for two-phase boiling water systems operating at 101,325 Pa are:

- Enthalpy reference temperature: 373.12 K
- Enthalpy reference value (liquid): 417436 J/kg
- Enthalpy reference value (vapor): 2674950 J/kg

The enthalpy results that are computed for this system can then be compared directly with steam table values.

Selecting the Phase Coupled Fluid Energy Model

The Phase Coupled Fluid Energy option becomes available within the *Optional Models* group box of the *Physics Model Selection* dialog after you have chosen the Multiphase Segregated Flow model. Click this option as shown below.



An additional *Energy* option appears within the *Phase Model Selection* dialog for each Eulerian phase in the continuum. An example of this dialog is shown below.



Choose between Segregated Fluid Enthalpy and Segregated Fluid Temperature.

Modeling Combustion in Multiphase Segregated Flow

Combustion in a multiphase flow has applications that include:

- Fluid catalytic cracking
- Fluidized beds
- Riser flow

To use the STAR-CCM+ facilities for solving combustion problems within a multiphase flow:

- Before you select phase models, activate the Phase Coupled Fluid Energy model in the single-phase continuum.



- For the phase, in the *Phase 1 Model Selection* dialog, select either a Multi-Component Gas or a Multi-Component Liquid material model from the *Material* group.
- Select Segregated Fluid Enthalpy from the *Energy* group. This model is required for each phase in which you want to use combustion.
- Specify that this phase requires chemical reaction simulation by selecting Reacting in the *Reaction Regime* group.

- Specify whether your combustion is premixed, non-premixed or partially premixed by selecting the corresponding radio button in the *Reactive Flow* group.



The availability of specific combustion modeling options depends on this choice and on the material model that you chose earlier.

Accessing Combustion Models

After selecting a combustion flame type, you are presented with a choice of combustion models valid for use with the flame type. If the combustion model you want is not available for the flame type, deactivate the flame type from the *Enabled Models* section and choose the correct alternative. You can switch between combustion models that are presented for a particular flame type without this step. For example, when you choose Homogeneous Reactor, you can switch among the three flame types even if you do not

deactivate the combustion model. The relationships among material models, combustion flame types, and specific combustion modeling options are outlined in the following table.

		Material	
		Multi-Component Gas	Multi-Component Liquid
Combustion	Non-Premixed	• EBU, including Ebu coefficients, Arrhenius coefficients, and the user reaction coefficient • Homogeneous Reactor	• EBU, including Ebu coefficients and the user reaction coefficient
	Partially Premixed	• EBU, including Ebu coefficients, Arrhenius coefficients, and the user reaction coefficient • Homogeneous Reactor	• EBU, including Ebu coefficients and the user reaction coefficient
	Premixed	• Homogeneous Reactor	

Most of the other features available with gas phase reactions, such as the species source and specific heat with NASA polynomials, are available with liquid phase reactions as well.

Modeling the Boiling Process

The Wall Boiling and Bulk Boiling models within the Eulerian Multiphase framework of STAR-CCM+ are aimed at forced-flow, subcooled, boiling applications. In subcooled boiling, local boiling occurs on a heated surface despite the mean temperature of the liquid being less than the saturation temperature. The degree of liquid subcooling is expressed either as a temperature difference or an enthalpy difference below the saturation value.

Provided the flow regime near the wall remains close to bubbly flow, a standard subcooled boiling model can be pushed towards zero subcooling conditions (that is, saturated boiling). The model can also be pushed towards higher heat fluxes approaching Departure from Nucleate Boiling (DNB) conditions.

A successful wall boiling model predicts the amount of wall superheat, which is the amount the wall temperature exceeds the saturation temperature, as well as the axial and radial distribution of vapor. The Eulerian Multiphase framework is appropriate for modeling this kind of wall boiling process for several reasons:

- By definition, the phases in subcooled boiling are not in thermal equilibrium and each phase requires its own energy equation.
- The phases in subcooled boiling require their own equation of motion, as the bubble moves away from the boiling surface while cool liquid moves to take its place.
- The condensation rate of vapor in the bulk subcooled liquid is a part of the wall boiling process and decides:
  - when net vapor generation first begins on the wall, and
  - how the vapor is distributed downstream of this point.

  Both of these non-equilibrium flow features can be different in reality from predictions using equilibrium models.

The Bulk Boiling model in STAR-CCM+ is an essential part of subcooled boiling. It is able to model boiling and condensation in the interior of the flow, providing the rate of interphase mass transfer only depends on the following:

- The heat transfer rate from the liquid and vapor to a saturated interface between these two phases.
- The interfacial area density that is provided by a well-defined characterization of bubbles by their size and number density.

This means that the bulk boiling model is suitable for modeling vapor bubbles whose initial size and concentration:

- are already known as an inlet condition (or initial condition in the case of transient simulations)
- are estimated by a size distribution model (such as S-gamma)
- are generated by the Wall Boiling model.

This thermally driven model is not suitable for describing boiling without bubbles (homogeneous nucleation), evaporation (concentration-driven multi-component mass transfer), nor cavitation (inertially limited mass

transfer). It is also not suitable when the resulting interphase mass transfer rate occurs on time scales that are much shorter than flow time scales. This scenario is unlikely to be solvable with a segregated solver algorithm.

What Is Bulk Boiling?

The Bulk Boiling model assumes that the phase temperatures are not in equilibrium, except at the interface where they are both equal to the saturation temperature. Bulk boiling is activated by selecting Interphase Mass Transfer phase interaction models.



STAR-CCM+ can solve for bulk boiling or condensation at a rate comparable to flow timescales. To solve these problems, use the interphase mass transfer models. The following models are provided:

- Boiling Mass Transfer Rate
- Multi-component Boiling Mass Transfer Rate
- User-defined Interphase Mass Flux

What Is the Boiling Mass Transfer Rate?

This model is only available when the Phase Coupled Fluid Energy model is active in the physics continuum. It is used to model the rate of bulk boiling or condensation between phases. The properties of three child nodes, the

Continuous Phase Nusselt Number, the Disperse Phase Nusselt Number and the Interface Temperature, are used to define the parameters that are required by the model equation.



Boiling Mass Transfer Rate Properties

Relaxation Factor	Adjusts the model sensitivity. Reduce it to improve convergence.

Properties of the Continuous Phase Nusselt Number

This property controls the heat transfer rate to the boiling interface from the liquid side (see Eqn. (1938)).

Method	Selects the method to use in specifying the continuous phase Nusselt number Nu_c in Eqn. (1940).	
	Constant	Adds a Constant child node on which to set the constant value.
	Field Function	Adds a Field Function child node from which to select an appropriate function.
	Ranz Marshall	Adds the Ranz Marshall child node that represents use of the Ranz-Marshall correlation that is given in Eqn. (1937).

Properties of the Disperse Phase Nusselt Number

This property controls the heat transfer rate to the boiling interface from the vapor side (see Eqn. (1939)).

Method — Select the method to use in specifying the dispersed phase Nusselt number Nu_d in Eqn. (1941).

- Constant — Adds a Constant child node on which to set the constant value.
- Field Function — Adds a Field Function child node from which to select an appropriate function.

Properties of the Interface Temperature

The interface temperature is often defined as a constant saturation temperature for a particular system pressure. In STAR-CCM+, a field function can optionally be used for the interface temperature, so that depth can be accounted for, or overpressure due to surface tension in micron-sized bubbles.

Method — Select the method to use in specifying the interface temperature T_{cd} in Eqn. (1942).

- Constant — Adds a node on which to set the constant value.
- Field Function — Adds a node on which to select a field function.

What is the Multi-component Boiling Mass Transfer Rate Model?

The Multi-component Boiling Mass Transfer Rate model is used to model the rate of bulk boiling or condensation between multi-component phases. The model calculates the mass, energy, and momentum transfer between a continuous and a dispersed multi-component phase. The two phases can be either a gas and a liquid, or two liquids that are divided by an interface.

Note: STAR-CCM+ implements the full multi-component boiling model when the Maxwell-Stefan diffusion model is active. However, when the Maxwell-Stefan diffusion model is not active, the dispersion is assumed to be dilute and a reduced multi-component boiling model is used.

For more information, see Reduced Model for Dilute Dispersions in the Multi-component Boiling Mass Transfer Rate formulation.

This model is only available when the Phase Coupled Fluid Energy model is active in the physics continuum.

To select the Multi-component Boiling Mass Transfer Rate model:

- Define both of the phases in the interaction to have more than one component--usually one multi-component liquid and one multi-component gas.
- Specify the continuous and dispersed phases.
- Select the Interphase Mass Transfer model.

  The *Interphase Mass Transfer Rate* options become available.

- Select the Multi-component Boiling Mass Transfer Rate model from the *Interphase Mass Transfer Rate* group box.



- Click Close.

The Multi-component Boiling Mass Transfer Rate node is shown below:



Multi-component Boiling Mass Transfer Rate Properties

Relaxation Factor	Adjusts the model sensitivity. Reduce it to improve convergence. The default value is 0.5, but lower values may be required for high mass transfer rates.
Phase 0 Components	Lets you specify the relationship between the components that are transferred between phases.

Multi-component Boiling Mass Transfer Rate Expert Properties

High Flux Correction Option	When activated, the high mass transfer rate coefficients are calculated. You can deactivate this option to skip the high mass transfer rate coefficients calculation: the coefficients are set to unity. This can save computing time for low mass transfer rates, when the value of these coefficients is expected to be close to unity.
Max Point Solver Iteration	Specifies the maximum number of iterations performed by the Point Solver. The default value is 100.
Point Solver Tolerance	Specifies the maximum value of the Point Solver residual. The default value is 0.001.
Point Solver Underrelaxation Factor	Adjusts the sensitivity of the Point Solver. Reducing this value, or the Point Solver Tolerance value (above), can increase the computing time significantly.

Specifying the relationship between the phase components

The *Multi-component Boiling Mass Transfer Rate - Phase 0 Components* dialog lets you specify the components to transfer. You match the components of the continuous phase (in the left-hand column) with the components of the

dispersed phase in the right-hand column:



For each transfer activated, the continuous phase component name appears under the Equilibrium Coefficient > Composite node:



Multi-component Boiling Mass Transfer Rate Model Nodes

The properties of child nodes are used to define the parameters that are required by the model equation.

- Continuous Phase Nusselt Number
- Disperse Phase Nusselt Number
- Continuous Phase Sherwood Number

  See What is the Continuous Phase Sherwood Number?
- Dispersed Phase Sherwood Number

  See What is the Dispersed Phase Sherwood Number?
- Equilibrium Coefficient

  See What is the Equilibrium Coefficient?

Selecting the Bulk Boiling Model

As a prerequisite, the following models must be activated for the continuum and Eulerian phases:

- In the *Physics Model Selection* dialog for the continuum, select the Phase Coupled Fluid Energy physics model.
- In the *Phase Model Selection* dialog for all applicable phases, select either the Segregated Fluid Temperature or Segregated Fluid Enthalpy phase model.

Then, to select the Bulk Boiling model:

- In the *Phase Interaction Model Selection* dialog, select the Interphase Mass Transfer phase interaction model.



- If the Auto-select recommended models option is activated, the Boiling Mass Transfer Rate model is automatically selected.



What Is Wall Boiling?

In a heated channel, where the flow at the entrance is subcooled, boiling begins when the wall temperature exceeds the saturation temperature of the liquid. Small bubbles start to appear on the wall at the point of Inception of Boiling (IB). Downstream of the point of IB when the degree of subcooling becomes smaller, the void fraction starts to grow from the point of Net Vapor Generation (NVG). Downstream of the point of NVG, the void fraction grows rapidly. Bubbles detach from the wall and move into the subcooled liquid. Some of these bubbles condense in the subcooled liquid, enhancing heat transfer to the liquid. Further downstream, the bulk liquid reaches the saturation point and saturated boiling begins.



The wall boiling process can be broken down into three stages that are based upon the different types of heat transfer, as illustrated in the diagram below.



The method for calculating each type of heat transfer is described below:

- **Evaporation** – Modeling of heat transfer at the wall begins with a prediction of the nucleation site number density. This value determines the number of bubbles per unit area, which is based on the amount of superheating on the wall. This number density is used with the bubble departure frequency and bubble departure diameter to model the evaporative heat flux across the nucleation sites.
- **Quenching** – When buoyancy causes a bubble to leave the surface, liquid fills the space that the bubble occupied, which enhances heat transfer. This process is called quenching. The quenching heat flux is modeled using the quenching heat transfer coefficient. The area over which quenching takes place is the bubble influence wall area fraction.
- **Convective Heating** – The remaining area is available for convective heat transfer. The wall dryout area fraction how much convective heat transfers to the liquid phase and how much transfers to the vapor phase.

The evaporation and quenching mechanisms can support much higher heat transfer rates than single-phase, forced flow.

The wall boiling model is comprised of several submodels, which calculate various aspects of the boiling process. Do not change the default properties for these submodels, unless you have supporting experimental data. The Setting Up a Wall Boiling Simulation section provides guidelines on the wall boiling model for both new and expert users.



Wall Boiling Properties

Relaxation Factor	Adjusts the model sensitivity. Reduce it to improve convergence.

Selecting the Wall Boiling Model

As a prerequisite, the following models must be activated for the continuum and Eulerian phases:

- In the *Physics Model Selection* dialog for the continuum, select the Phase Coupled Fluid Energy physics model.
- In the *Phase Model Selection* dialog for all applicable phases, select the Segregated Fluid Temperature phase model.

Then, to select the Wall Boiling model:

- In the *Phase Interaction Model Selection* dialog, select the Interphase Mass Transfer phase interaction model. The Boiling Mass Transfer Rate model is automatically selected.



- Deselect the Constant Interaction Length Scale model and instead, select the Kurul Podowski Interaction Length Scale model from the *Interaction Length Scale* group box.



- Alternatively, you can select the User-Defined Interaction Length Scale to use the S-Gamma model for predicting bubble size.

- Finally, select the Wall Boiling model from the *Optional Models* group box.



Setting Up a Wall Boiling Boundary

After you select the Wall Boiling phase interaction model, a Wall Interphase Mass Transfer Option appears in the Physics Conditions of each wall boundary. A User Wall Heat Flux Coefficient Specification appears for each phase in the Phase Conditions.



Wall Interphase Mass Transfer Option Properties

Method	Active	Allow wall boiling to occur for this boundary.
	None	Do not allow wall boiling to occur on this boundary. For example, use this option when a boundary is adiabatic.

Setting Up a Wall Boiling Simulation

This section provides recommendations and considerations for setting up a wall boiling simulation for both new and expert users.

Basic Setup

To set up a basic simulation:

- First, verify that your application is suitable for modeling as a subcooled boiling process. Base your model on real experimental or industrial conditions, or at least on an energy balance confirming that subcooled boiling is expected.
- Then, start from the default model selections and default model calibrations.

Condensation is a part of subcooled boiling, so you have to activate the Bulk Boiling model before activating the Wall Boiling model.

Selection of the Wall Boiling model requires that you set valid and accurate data for the following simulation properties:

- For Reference Values > Gravity, make sure that acceleration due to gravity is acting in the correct direction.
- For Reference Values > Reference Pressure, set the value to the expected outlet pressure, if setting zero relative pressure conditions at the outlet. Otherwise, set the value to the design operating pressure.
- For Phase Interaction > Multiphase Material > Surface Tension, set the surface tension coefficient, at saturation temperature, for the liquid-vapor interface.
- For Eulerian Phases > Gas > Molecular Weight, set an appropriate material property for the vapor phase.

Advanced Setup

The wall boiling model is intended to cover forced-flow, subcooled boiling up to pressures of around 155 bar, with as little calibration of model constants as possible. Therefore, it is built up mechanistically from many simpler submodels:

- Nucleation Site Number Density
- Bubble Departure Diameter
- Bubble Departure Frequency
- Bubble Influence Wall Area Fraction
- Bubble Induced Quenching Heat Transfer
- Wall Dryout Area Fraction

The output from each of these submodels can be inspected graphically at run time.

If suitable experimental data is available, these submodels can also be calibrated or replaced.

Working Fluid

The wall boiling submodels are implemented to use all of the phase and interface properties you specify (such as density, saturation enthalpy, surface tension).

Note: Some of the original calibrations are based upon water as the working fluid. Therefore, some adjustments are required when using a working fluid other than water.

The wall contact angle for the Hibiki Ishii Nucleation Site Number Density and Kocamustafaogullari Bubble Departure Diameter models is specific to the combination of working fluid and boiling surface.

Note: This angle is a nominal value that is based on room temperature rather than a value measured under boiling conditions.

Adjust the following options, from older correlations for use with other working fluids:

- Kurul Podowski Interaction Length Scale
- Lemmert Chawla Nucleation Site Number Density
- Tolubinsky Kostanchuk Bubble Departure Diameter

Convergence

High heat flux wall boiling cases can sometimes be difficult to converge for two reasons:

- Interphase mass transfer rates due to bulk condensation are determined by instantaneous values of interfacial area and interfacial energy balance. However, in a segregated solver algorithm, these mass transfer rates are not directly coupled to the hydrodynamic equations, so they are only limited by mass balances after full convergence.
- The most sensitive part of the wall boiling model is the nucleation site number density. The number of active sites depends on wall superheat, perhaps to the power of 1.8 to 5 or even beyond. In a segregated solution algorithm, the resulting evaporation rate can overshoot the liquid supply to the wall until full convergence is achieved.

Rather than turning on full power immediately in a new computation and then under-relaxing all equations heavily, a more natural strategy is given below:

- Obtain a converged flow and heat transfer solution at inlet temperature, but without any boiling.
- Restart the simulation with bulk condensation activated and obtain a converged flow.
- Activate wall boiling and gradually increase power from just below the value to induce boiling at the wall, to full power over a number of iterations, time-steps, or restarts.

If you experience difficulty in achieving convergence, adjust the relaxation factor that is provided for the Boiling Mass Transfer Rate (the terms that cover bulk condensation). Typical values are 0.5 but this value can be reduced to make the residuals behave more monotonically.

Another relaxation factor is provided for the Wall Boiling model. It is applied to the phase energy equations at the wall and is intended to account for the sensitivity of evaporation rate to wall temperature. The recommended value is 0.5 when the default Lemmert Chawla model is selected for Nucleation Site Number Density. Reduce the relaxation factor to about 0.1 when the Hibiki Ishii model is used.

Sometimes unnecessary difficulties are caused by inappropriate turbulence profiles at inlets. The symptoms are that the turbulent viscosity levels collapse on entry with possible consequences for both nearby wall heat transfer coefficients and for global residual behavior. For long pipe and channel geometries, you can use the same model, but with boiling deactivated, to compute and save fully developed outlet velocity and turbulent profiles in a file. These boundary conditions can then be used as inlet conditions for a boiling simulation.

What Is the Nucleation Site Number Density?

The nucleation site number density determines the number of locations on the heated surface where bubbles form, per unit area. This is the leading factor determining the evaporation rate in a mechanistic model of subcooled boiling (see Eqn. (1991)). The diagram below illustrates this

parameter:



The options for calculating the nucleation site number density are:

- Lemmert Chawla Nucleation Site Number Density
- Hibiki Ishii Nucleation Site Number Density
- User-Defined Nucleation Site Number Density

What Is the Lemmert Chawla Nucleation Site Number Density Model?

The Lemmert Chawla model [346] is a correlation against wall superheat and is the default model that is recommended for initial studies on the grounds of robustness. The default calibration assumes that the number of active nucleation sites varies with the wall superheat to the power of 1.805. This calibration is used with the Tolubinsky and Kostanchuk model for calculating bubble departure diameter, for example for modeling forced convection subcooled boiling of water at 45 bar. An alternative calibration is discussed in the formulation section.

Properties of the Lemmert Chawla Nucleation Site Number Density Model

Calibration Constant A	A in Eqn. (2007). The default value of 1.805 is chosen according to Kurul and Podowski [343].
Calibration Constant B	B in Eqn. (2007). The default value is 0.0
Maximum Superheat	ΔT_{max} in Eqn. (2006). Sets the upper limit for the wall superheat value that is passed to this submodel, which prevents overflow during initial convergence. Make sure that this limit does not affect the value of wall superheat predicted in the converged solution. The default value is 25.0 delta K.
Reference Nucleation Site Number Density	n_0'' in Eqn. (2007). The default value is 12366.44783 / m^2.
Reference Wall Superheat	ΔT_0 in Eqn. (2007). The default value is 1.0 delta K.

Selecting the Lemmert Chawla Model

To select the Lemmert Chawla model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- If the Auto-select recommended models checkbox is activated, the Lemmert Chawla model is automatically selected.

- If the Auto-select recommended models checkbox is not activated, select the Lemmert Chawla model from the *Nucleation Site Number Density* group box.



What Is the Hibiki Ishii Nucleation Site Number Density Model?

The Hibiki Ishii model is a more advanced nucleation site number density model. It is applicable for pressures up to 198 bar. Use the Hibiki Ishii model with the Kocamustafaogullari model for bubble departure diameter.

The model has the following characteristics:

- It considers a boundary condition for wall superheating.
- It is validated against numerous sets of experimental data.
- It has a wide range of applicability in terms of mass flow, pressure, and contact angle.

Properties of the Hibiki Ishii Nucleation Site Number Density Model

If you use a working fluid other than water, adjust the Wall Contact Angle. If you know average cavity density and the cavity length scale values for the working surface, adjust these also.

Maximum Superheat	Sets the upper limit for the wall superheat value that is passed to this submodel (see Eqn. (2008)), which prevents overflow during initial convergence. Make sure that this limit does not affect the value of wall superheat predicted in the converged solution.
Wall Contact Angle	Data that you define for your working fluid. A nominal value at room temperature is required. The default of 0.722 radians (41.37^{o}) is an approximate value for water systems (see Eqn. (2011)).
Average Cavity Density	The average number of nucleation sites per unit area, which is used when calculating the nucleation site number density (see Eqn. (2011)). This can be tuned to suit a particular surface finish.
Wall Contact Angle Scale	Reference value that is fixed by model authors to calibrate model for dependence on Wall Contact Angle (see Eqn. (2011)). The standard value for most applications is 0.722 radians (41.37^{o}).
Cavity Length Scale	Value that is used when calculating the nucleation site number density (see Eqn. (2011)). This can be tuned to suit a particular surface finish.
Density Function Constant C0, C1, C2, and C3	Values that are fixed by the model authors to calibrate the pressure effect function, which is used when calculating nucleation site number density (see Eqn. (2013)). The standard values for most applications are -0.1064, 0.48246, -0.22712, and 0.05468.

Selecting the Hibiki Ishii Model

To select the Hibiki Ishii model:

- In the *Phase Interaction Model Selection* dialog, deactivate the Auto-select recommended models checkbox.
- Select the Wall Boiling model from the *Optional Models* group box.
- Select the Hibiki Ishii model from the *Nucleation Site Number Density* group box.



What Is the Bubble Departure Diameter?

The bubble departure diameter determines the diameter of the bubble at the instant it leaves the nucleation site. This is the second of the three factors determining the evaporation rate in subcooled boiling (see Eqn. (1991)). The diagram below illustrates this parameter:



The options for calculating the bubble departure diameter are:

- Tolubinsky Kostanchuk Bubble Departure Diameter
- Kocamustafaogullari Bubble Departure Diameter
- User-Defined Bubble Departure Diameter

What Is the Tolubinsky Kostanchuk Bubble Departure Diameter Model?

The Tolubinsky Kostanchuk model is a correlation against liquid subcooling and the default model that is recommended for initial studies on the grounds of robustness. It matches the default Lemmert Chawla model for nucleation site number density.

Properties of the Tolubinsky Kostanchuk Bubble Departure Diameter Model

Reference Diameter	Reference diameter that is used when calculating the bubble departure diameter (see Eqn. (2017)).
Reference Diameter Subcooling	Reference diameter subcooling (that is, the difference between the saturation temperature and liquid temperature for the reference diameter) used when calculating the bubble departure diameter (see Eqn. (2017)).
Minimum Diameter	Sets the lower limit for the bubble departure diameter. This setting makes sure that the result is within plausible limits.
Maximum Diameter	Sets the upper limit for the bubble departure diameter. This setting makes sure that the result is within plausible limits.

Selecting the Tolubinsky Kostanchuk Model

To select the Tolubinsky Kostanchuk model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- If the Auto-select recommended models checkbox is activated, the Tolubinsky Kostanchuk model is automatically selected.

- If the Auto-select recommended models checkbox is not activated, select the Tolubinsky Kostanchuk model from the *Bubble Departure Diameter* group box.



What Is the Kocamustafaogullari Bubble Departure Diameter Model?

The Kocamustafaogullari model is more recent, more general and based on force balance with an adjustment for pressure dependence. This model must be used whenever the Hibiki Ishii model is selected for nucleation site number density.

Properties of the Kocamustafaogullari Bubble Departure Diameter Model

If you use a working fluid other than water, adjust the Wall Contact Angle.

Calibration Constant	Calibration constant (see Eqn. (2019)). The standard value is 1.5126×10^{-3} m/radian (2.64×10^{-5} m/degree).

Wall Contact Angle — Data that you specify for your working fluid. A nominal value at room temperature is required. The default is 0.722 radians (41.37°) is an approximate value for water systems (see Eqn. (2019)).

Selecting the Kocamustafaogullari Model

To select the Kocamustafaogullari model:

- In the *Phase Interaction Model Selection* dialog, deactivate the Auto-select recommended models checkbox.
- Select the Wall Boiling model from the *Optional Models* group box.
- Select the Kocamustafaogullari model from the B*ubble Departure Diameter* group box.



What Is the Bubble Departure Frequency?

The bubble departure frequency determines how many bubbles leave a nucleation site per second. This is last of three key factors determining the evaporation rate in subcooled boiling (see Eqn. (1991)). The diagram below illustrates this parameter:



The options for calculating the bubble departure frequency are:

- Cole Bubble Departure Frequency
- User-Defined Bubble Departure Frequency

What Is the Cole Bubble Departure Frequency Model?

This frequency is equivalent to a terminal velocity scale over bubble departure diameter, and the overall evaporation rate is calibrated around this assumption (see Eqn. (2020)). There are no adjustable parameters, but you can define an alternative model.

Properties of the Cole Model

None.

Selecting the Cole Model

To select the Cole model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- If the Auto-select recommended models checkbox is activated, the Cole model is automatically selected.

- If the Auto-select recommended models checkbox is not activated, select the Cole model from the *Bubble Departure Frequency* group box.



What Is the Bubble Influence Wall Area Fraction?

The bubble influence wall area fraction estimates the fraction of the wall area that is affected by the sweep of liquid inflow beneath a departing bubble.

The options for calculating the bubble influence wall area fraction are:

- Kurul Podowski Bubble Influence Wall Area Fraction
- User-Defined Bubble Influence Wall Area Fraction

What Is the Kurul Podowski Bubble Influence Wall Area Fraction Model?

The Kurul Podowski model assumes that the wall area influenced by bubble-induced quenching is larger than the nucleation site area, by a specified factor (see Eqn. (2026)).

Properties of the Kurul Podowski Bubble Influence Wall Area Fraction Model

Area Coefficient	Ratio of nucleation site area density and bubble-influenced wall area density, which is given by (Eqn. (2026)). The default value is 2.0.

Selecting the Kurul Podowski Model

To select the Kurul Podowski model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- If the Auto-select recommended models checkbox is activated, the Kurul Podowski model is automatically selected.
- If the Auto-select recommended models checkbox is not activated, select the Kurul Podowski model from the *Bubble Influence Wall Area Fraction* group box.



What Is the Quenching Heat Transfer?

When a bubble leaves the heated surface, cooler liquid fills the space that it occupied. The heat transfer during this process is known as quenching heat transfer. The Quenching Heat Transfer Coefficient is used to calculate the quenching heat flux.

The options for calculating the quenching heat transfer coefficient are:

- Del Valle Kenning Bubble Induced Quenching Heat Transfer Coefficient
- User-Defined Bubble Induced Quenching Heat Transfer Coefficient

When you activate Wall Boiling, a model is automatically selected to correct the bubble induced quenching heat flux:

- Bubble Induced Quenching Temperature

What Is the Del Valle Kenning Bubble Induced Quenching Heat Transfer Coefficient Model?

This model calculates the bubble induced quenching heat transfer coefficient using Eqn. (2022).

Properties of the Del Valle Kenning Bubble Induced Quenching Heat Transfer Coefficient Model

Wait Coefficient	Proportion of bubble departure cycle period between departure of one bubble and nucleation of the next (see Eqn. (2023)).

Selecting the Del Valle Kenning Model

To select the Del Valle Kenning model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- If the Auto-select recommended models checkbox is activated, the Del Valle Kenning model is automatically selected.

- If the Auto-select recommended models checkbox is not activated, select the Del Valle Kenning model from the *Bubble Induced Quenching Heat Transfer Coefficient* group box.



What Is the Bubble Induced Quenching Temperature?

This model corrects the Bubble Induced Quenching Heat Flux so that it uses the temperature of the liquid brought to the wall by the action of the departing bubble. This temperature can be different to the liquid temperature computed at the cell center next to the wall. The correction is taken from a location on the undisturbed liquid phase temperature profile away from the departure site.

Properties of the Bubble Induced Quenching Temperature Model

Quenching Temperature Option	Defines the location for sampling the liquid temperature profile.	
	Wall Cell	The liquid temperature at the cell center next to the wall is used as the quenching temperature. This is the default setting.
	Fixed Number of Diameters	The liquid temperature is estimated at a location that is specified as a fixed number of the local bubble departure diameter size. You set this in the Fixed Number of Diameters sub-node.
	Fixed Yplus	The liquid temperature is estimated at a location that is specified as a fixed number of wall turbulence length scales for the liquid phase. You set this in the Fixed Y+ sub-node.

Selecting the Bubble Induced Quenching Temperature Model

To select the Bubble Induced Quenching Temperature Model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- The Bubble Induced Quenching Temperature model is automatically selected.

What Is the Wall Dryout Area Fraction?

The wall dryout area fraction specifies how much of the heat flux that is applied at the wall is directed towards the vapor convection, as opposed to liquid convection and evaporation (see Eqn. (1992)). It can be used to improve robustness during initial convergence, or for indicating approach to Departure from Nucleate Boiling (DNB) conditions in converged solutions.

The options are:

- Basic Wall Dryout Area Fraction
- User-Defined Wall Dryout Area Fraction

What Is the Basic Wall Dryout Area Fraction Model?

The Basic model defines the area fraction as zero until the volume fraction of the vapor near the wall exceeds a specified value. It then transitions the area fraction to unity as the vapor volume fraction approaches unity.

To reduce grid-dependency, the near-wall volume fraction that is used for the area fraction calculation can optimally be defined as an average over a prescribed thickness of notional bubbly layer next to the wall.

Properties of the Basic Wall Dryout Area Fraction Model

Bubbly Layer Relaxation Factor	Controls the update of the Bubbly Layer Volume Fraction that is used by the dryout model. It helps convergence behavior once the dryout criterion has been exceeded. The default value is 0.5	
Bubbly Layer Option	Defines the thickness of the layer that is used for averaging volume fraction in the dryout criterion.	
	Wall Cell	The value that is computed at the cell center next to the wall represents the entire bubbly layer. This means the cell-center volume fraction is used in the dryout criterion. This is the default setting.
	Fixed Number of Diameters	The bubbly layer thickness is defined by the local bubble departure diameter that is multiplied by the value you set in the Fixed Number of Diameters sub-node.
	Fixed Yplus	The bubbly layer thickness is defined by the wall turbulence length scale for the vapor phase that is multiplied by the value you set in the Fixed Y+ sub-node.
Wall Dryout Breakpoint	The vapor volume fraction at which point heat transfer to the vapor phase begins (see Eqn. (2000)). The default value is 0.9	

Selecting the Basic Wall Dryout Area Fraction Model

To select the Basic Wall Dryout Area Fraction model:

- In the *Phase Interaction Model Selection* dialog, select the Wall Boiling model from the *Optional Models* group box.
- If the Auto-select recommended models checkbox is activated, the Basic Wall Dryout Area Fraction model is automatically selected.
- If the Auto-select recommended models checkbox is not activated, select the Basic Wall Dryout Area Fraction model from the *Wall Dryout Area Fraction* group box.



Specifying User-Defined Wall Boiling Parameters

Each boiling parameter has an option that allows you to specify your own values, as opposed to STAR-CCM+ calculating the values internally. The properties and usage of models you define are the same, so they are all documented in this common section.

Use these options to specify experimental data in the simulation, or you are implementing your own physics model.



Properties of User-Defined Models

The properties of all models you define are identical to those used for scalar profiles.

Selecting a User-Defined Model

To select a user-defined model:

- Deactivate the Auto-select recommended models checkbox.
- Select the Wall Boiling model.

- Select the model for the parameters you wish to specify manually. For example, the User-Defined Bubble Departure Frequency model is selected below.



Wall Boiling Field Functions Reference

The following field functions are made available to the simulation when the Wall Boiling model is used. These objects all have the same base properties.

BubblyLayerVolumeFraction	α_δ in Eqn. (1999).
BubblyLayerVolumeFraction-Gradient	α_g' in Eqn. (1999).
BubbleDepartureDiameter	d_w in Eqn. (2017) for the Tolubinsky Kostanchuk model or Eqn. (2019) for the Kocamustafaogullari model.
BubbleDepartureFrequency	f in Eqn. (2020).

BubbleInfluenceWallAreaFraction	K_{quench} in Eqn. (2026).
BubbleInducedQuenchingHeatTransferCoefficient	h_{quench} in Eqn. (2022).
BubbleInducedQuenchingTemperatureFactor	s_{quench} in Eqn. (2025).
BubbleInducedQuenchingTemperatureDistance	y_{quench} in Eqn. (2025).
NucleationSiteNumberDensity	n'' in Eqn. (2007) for the Lemmert Chawla model or Eqn. (2011) for the Hibiki-Ishii model.
WallInterfaceMassFlux	$\dot{m}''_{evap}$ in Eqn. (2004).
BoundaryHeatFlux<Liquid Phase Name>	q''_{wl} in Eqn. (1994).
BoundaryHeatFlux<Phase Interaction Name>	q''_{wlg} in Eqn. (1995).
BoundaryHeatFlux<Vapor Phase Name>	q''_{wg} in Eqn. (1996).

Modeling Particle-Phase Flows

Particle-particle interactions can be significant in solid particle flows and therefore require their own set of models.

STAR-CCM+ provides the Granular Pressure model for gas-solid flows, such as fluidized beds. The model introduces the concept of the "packing limit." Below this limit, particles are loosely packed and inter-particle collisions determine their motion; these collisions are handled as part of a Granular Temperature model that becomes available when Granular Pressure is selected. Above the packing limit, friction between particles dominates their motion and a frictional stress model is used.

The following sections describe the Granular Pressure model:

- What is the Granular Pressure Model?
- Selecting the Granular Pressure Model
- What is the Granular Temperature Model?
- Selecting the Granular Temperature Model
- Setting Granular Pressure at a Particle Phase
- Setting Granular Pressure at Boundaries
- Formulation for the Granular Pressure

In STAR-CCM+, the Solid Pressure Force model is provided for dispersed particle phases to account for the forces that arise during particle-particle interaction when the particles are packed closely together. It is a way of limiting the maximum particle volume fractions. The following sections are provided:

- What Is the Solid Pressure Force Model?
- Selecting the Solid Pressure Force Model
- Formulation for the Solid Pressure Force

What Is the Granular Pressure Model?

The Granular Pressure model estimates the stresses of particle phases in a gas-solid fluidized bed. Fluidized bed reactors are widely used in combustion, catalytic cracking and various other chemical and metallurgical processes. Granular pressures are also encountered in many industrial processes.

The closeness of particles to one another, or "packing," determines the type of particle motion. Where the particles are closely packed, the motion is dominated by contact and a frictional stress model is used. When the particles are below the packing limit, then collisions and kinetic energy determine the motion, handled as part of the Granular Temperature model.



STAR-CCM+ supports multiple particle phases in the frictional model.

Granular Pressure Properties

Determines the regimes that are modeled.

Radial Distribution Function	Ding-Gidaspow	Use the Ding-Gidaspow radial distribution function. Eqn. (1821) is used for packing below the packing limit and Eqn. (1822) is used close to the packing limit.
Frictional Stress Formulation	Specifies the motion behavior of the particles when they approach the maximum packing limit. Select an appropriate frictional stress formulation.	
	Schaeffer	Use the Schaeffer [359] frictional regime equations. You set the appropriate properties in the Schaeffer sub-node.

Modified Johnson	Use the Modified Johnson [340] frictional regime equations. You set the appropriate properties in the Modified Johnson sub-node.
Kinetic-Based	Use the collisional solid pressure to limit the particle volume fraction near maximum packing. The collisional solid pressure relies on the correction to the particle radial distribution function near maximum packing (Eqn. (1822)).

Schaeffer Properties

Specifies the motion behavior of the particles. This behavior is based on the work by Schaeffer [359]. Only available when the Frictional Stress Formulation property is set to Schaeffer.

Angle of Internal Friction	Used to determine the effective granular viscosity, which is given by Eqn. (1854). It is set to 25° by default.

Modified Johnson Properties

Specifies the motion behavior of the particles. This behavior is based on the work by Johnson and Jackson [340]. Only available when the Frictional Stress Formulation property is set to Modified Johnson.

Critical Volume Fraction	Solid volume fraction at which the frictional stress is activated. Set to 0.5 by default.
Angle of Internal Friction	Used to determine the effective granular viscosity, which is given by Eqn. (1854). Set to 28.5° by default.

Modified Johnson Expert Properties

Multiplier Fr	Multiplication factor in Eqn. (1848). Set to 0.05 by default.
Exponent r	Numerator exponent in Eqn. (1848). Set to 2.0 by default.
Exponent s	Denominator exponent in Eqn. (1848). Set to 5.0 by default.

Selecting the Granular Pressure Model

Select the Granular Pressure Model through the Granular Pressure option of the *Optional Models* section of the *Physics Model Selection* dialog, as shown below. This option becomes available after selecting the Multiphase Segregated Flow model.



Once Granular Pressure has been activated for the continuum, you can specify relevant settings in the particle phase.

What is the Granular Temperature Model?

The Granular Temperature model calculates the collisions and motions of particles below the packing limit. Select this model only if you wish to model the collision-dominated regime of flow. Do not select this model if you only wish to model the frictional regime.

Selecting the Granular Temperature Model

The Granular Temperature model becomes available when you select the Granular Pressure model, in the *Optional Models* section of the *Physics Model Selection* dialog.



When you select the Granular Temperature model, the Granular Temperature Transport model is selected automatically. However, if you deactivate Auto-select recommended models before you select the Granular Temperature model, three models are available in the *Granular Temperature* section of the *Physics Model Selection* dialog.



The Granular Temperature models are:

- Algebraic Granular Temperature.

  This model applies to cases where the local equilibrium assumption holds. These cases include bubbling fluidized beds and applications involving packed beds of particles. This model is less computationally expensive than the Granular Temperature Transport model.

- Granular Temperature Transport.

  This model applies to any granular flow case. It takes into account convection, production, diffusion, and dissipation of granular energy. It also accounts for contributions to granular energy due to inter-phase drag. If you select this model, you must select the Granular Energy Transfer model for phase interactions involving the particle dispersed phase.

- User-Defined Granular Temperature.

  This model lets you specify the minimum granular temperature directly.

Algebraic Granular Temperature Properties

This model uses Eqn. (1828) to solve for the granular temperature.

Granular Temperature Minimum	Sets a minimum value for the continuum. Set to 1.0E-10 m^2/s^2 by default.
Particle Kinetic Viscosity	Specifies the motion behavior of the particles below the maximum packing limit. Use the corresponding Linearized Drag Law with the chosen Kinetic viscosity. Gidaspow: Use the Gidaspow Kinetic Viscosity model (see Eqn. (1840)). Selected by default. Syamlal: Use the Syamlal Kinetic Viscosity model (see Eqn. (1841)).

Granular Temperature Transport Properties

When you select this model, STAR-CCM+ adds the Granular Temperature solver to the list of solvers.

Granular Temperature Minimum	Sets a minimum value for the continuum. Set to 1.0E-10 m^2/s^2 by default.
Particle Kinetic Viscosity	Specifies the motion behavior of the particles below the maximum packing limit. Use the corresponding Linearized Drag Law with the chosen Kinetic viscosity. Gidaspow: Use the Gidaspow Kinetic Viscosity model (see Eqn. (1840)). Selected by default. Syamlal: Use the Syamlal Kinetic Viscosity model (see Eqn. (1841)).
Granular Diffusion Coefficient	Specifies the method for calculating the diffusion coefficient in the Granular transport equation. Gidaspow: Use the Gidaspow method (see Eqn. (1825)). Selected by default. Syamlal: Use the Syamlal method (see Eqn. (1826)).

User-Defined Granular Temperature Properties

Granular Temperature Minimum	Sets a minimum value for the continuum. Set to 1.0E-10 m^2/s^2 by default.
Particle Kinetic Viscosity	Specifies the motion behavior of the particles below the maximum packing limit. Use the corresponding Linearized Drag Law with the chosen Kinetic viscosity.
	Gidaspow: Use the Gidaspow Kinetic Viscosity model (see Eqn. (1840)). Selected by default.
	Syamlal: Use the Syamlal Kinetic Viscosity model (see Eqn. (1841)).

Setting Granular Pressure in a Particle Phase

If you choose the Granular Pressure and Granular Temperature models for the continuum, the Granular Pressure model and one of the Granular Temperature models are automatically selected when you specify the phase type as Particle.



Note: When using the Granular Pressure model, it is important to use a Laminar viscous regime. Granular Pressure models the turbulence and is incompatible with a turbulent particle phase.

Setting Phase Model Properties

This section describes the properties of granular pressure particle phase models.

Granular Pressure Properties

Sets the maximum packing limit for that phase.



Maximum Solid Fraction	Set the upper packing limit of the particles. It is set to 0.624, the limit for rigid spherical particles, by default.

Algebraic Granular Temperature Properties

This option is available only when the Algebraic Granular Temperature model is selected in the physics continuum.

Uses Eqn. (1828) to solve for the granular temperature.



User-Defined Granular Temperature Properties

This option is available only when the User-Defined Granular Temperature model is selected in the physics continuum.

You can input your own values for granular temperature using a method that is specified in the Granular Temperature sub-node. A Constant value of 1.0E-6 m^2/s^2 is set by default.



Granular Temperature Transport Properties

This option is available only when the Granular Temperature Transport model is selected in the physics continuum.



Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Convection	Select the convection scheme for flow velocity.	

1st-order	Selects the first-order upwind convection scheme.
2nd-order	Selects the second-order upwind convection scheme.

Setting the Particle Diameter

You set the particle diameter in the Material Properties node in the particle phase definition.



Particle Diameter Properties

Method	Specifies the method for defining the particle diameter.	
	Constant	Specify a constant value that represents the mean particle diameter.
	Field Function	Specify a field function that represents the mean particle diameter.
	Sauter Mean Diameter	This method is made available when the S-gamma model is activated for the dispersed phase.

Setting Phase Reference Values

This section describes the properties of granular pressure particle phase reference values.

The phase Reference Values have the following properties:

Coefficient of Restitution Properties

Value	Represents the ratio of speeds of the particles before and after an impact. By default it is set to 0.9. A value of 1 indicates a perfectly elastic collision.

Setting Phase Initial Conditions

This section describes the properties of granular pressure particle phase initial conditions.

The phase Initial Conditions nodes are available only when any of the Granular Temperature submodels have been selected when selecting the Physics models.

Granular Temperature

You can input your own values for the initial granular temperature in a sub-node that is based on the method you select.

A Constant value of 1.0E-6 m^2/s^2 is set by default.

Velocity

You can specify the initial velocity as a vector. It is entered as a vector profile.

Volume Fraction

You can input your own values for the initial volume fraction in a sub-node that is based on the method you select.

Setting Granular Pressure at Boundaries

To set the granular pressure at boundaries:

1. Choose the wall shear stress specification method.
2. Choose the granular temperature specification method.
3. Specify the appropriate values under the Physics Values node:
   - Granular Temperature
   - Wall Coefficient of Restitution
   - Specularity Coefficient

Setting the Wall Shear Stress Specification Method

The Physics Conditions > Shear Stress Specification node lets you choose the wall shear stress specification method.

It is recommended that you set the Shear Stress Specification for the particle phase to Slip at wall boundaries. If an estimate of the Specularity Coefficient is available, set this option to Partial-Slip and specify the appropriate value for the specularity coefficient.



Shear Stress Specification Properties

Method	Defines the wall shear stress specification method.	
	No-Slip	The tangential velocity is set to 0.
	Partial-Slip	Adds the Specularity Coefficient in a Physics Values node.
	Slip	No shear stress at the boundary.

Setting the Granular Temperature Specification Method

The Physics Conditions > Granular Temperature Specification node lets you choose the granular temperature specification method.

To estimate the granular temperature at boundaries, either extrapolate the granular temperature at the face cell, use the Johnson-Jackson method, or use a user-defined value.

The granular temperature value that is specified is used to calculate the solid pressure at the boundaries, which has a significant influence on results.



Granular Temperature Specification Properties

The Option property determines the method of inputting the granular temperature for the particle phase at your chosen boundary.

Extrapolation	Extrapolate from a face cell to the boundary. This option is not available if the wall shear stress specification method is set to Partial-Slip.

User	Specify your own boundary values for the Granular Temperature in a Physics Values node.
Johnson-Jackson	Use the Johnson-Jackson model for wall shear stress for particle phases in granular flows. This model is a more physically sound method of estimating granular temperature near the walls. This model can be used with both the Slip and Partial-Slip wall shear stress specification method. This option is available only when the Granular Temperature Transport model is selected. Specify the Wall Coefficient of Restitution and the Specularity Coefficient in a Physics Values node.

Granular Temperature Properties

Method	Sets the method to use for specifying your own values for granular temperature. A Constant value of 1.0E-6 m^2/s^2 is set by default.

Wall Coefficient of Restitution Properties

Value	The coefficient of restitution for collisions between particles and the wall. This value represents the ratio of speeds of the particles before and after an impact. By default the coefficient of restitution is set to 0.9. A value of 1 indicates a perfectly elastic collision.

Specularity Coefficient Properties

Method	Sets the method to use for specifying your own values for the specularity coefficient. A value of 0 corresponds to "Slip", while a value of 1 corresponds to "No-Slip".

You need to set the specularity coefficient only when on the wall shear stress specification method is set to Partial-Slip. The specularity coefficient does not depend on the granular temperature specification.

What Is the Solid Pressure Force Model?

The Solid Pressure Force model adds a compaction or 'solid pressure' force to the momentum equations of all the dispersed particle phases. This force becomes significant as the specified maximum packing limit is reached, and acts to repel particles and prevent them from exceeding the maximum volume fraction. The Solid Pressure Force node and its properties represent it.



Solid Pressure Force Properties

Maximum Solid Fraction	The value of the maximum packing limit $\alpha_{p,\,max}$ in Eqn. (1856), which is set to 0.624 by default (representative of rigid spherical particles).

Expert Properties

Solid Pressure Force Constant	The model parameter A in Eqn. (1856), which is set to -600 by default.

Selecting the Solid Pressure Force Model

The Solid Pressure Force Model is selected through the Solid Pressure Force option of the *Optional Models* section of the *Physics Model Selection* dialog, as shown below. This option becomes available after selecting the Multiphase Segregated Flow model.



Modeling Particle Size Distribution

Particle size is a vital parameter for calculating particle hydrodynamics in dispersed multiphase flows. In flows where the particles in the dispersed phase are bubbles, for example, the size of the bubbles can change continuously due to breakup and coalescence. As interfacial transfer terms are related to the surface area of the dispersed phase, it is essential to account for the particle size and its distribution to correctly simulate the multiphase flow behavior.

In STAR-CCM+, the S-Gamma model that was developed by Lo and Rao [347], and Lo and Zhang [348] has been implemented to account for the particle size and its distribution. This model, which assumes a log-normal distribution of particle sizes, is based on predicting the transport of the moments of the particle size distribution. These moments are:

- zeroth moment: the particle number density, n.

- second moment: related to the interfacial area density.
- third moment: related to the dispersed phase volume fraction.

One transport equation is solved for each moment. As the third moment is based on the dispersed phase volume fraction, this can be derived from the volume fraction equation that is solved by the Multiphase Segregated Flow solver. You can choose whether to solve for just the second moment, or for the zeroth and second moments, by selecting the required option in the properties of the S-Gamma model node. More details are provided in the formulation section for this model.

Particles change size due to breakup and coalescence as they interact with each other. Additional models are provided to account for this interaction between phases. These models are selected on the multiphase interaction node for a pair of Eulerian phases.

The log-normal size distribution is defined by a mean particle diameter and its variance. The mean particle diameter is always updated in the calculation, but the variance is only updated when the zeroth moment is solved in addition to the second moment, and breakup and/or coalescence is active. If the zeroth moment is not solved, the variance of the distribution remains at the initial value you provide. This initial value could be zero.

After selecting the S-Gamma model, define the log-normal size distribution of particles in a dispersed phase entering the simulation. Also provide the initial distribution at the start of the simulation. In both cases, the particle mean size and its statistical variance are used to define the distribution.

One key benefit of predicting the particle size distribution is that the mean particle size can be coupled back into the phase transport equations through the interaction length scale. The Sauter Mean diameter is one of the additional field functions that are exposed by the S-Gamma model.

Extensions to the S-Gamma model can be supplied using the region source terms for the zeroth and second moments.

What Is the S-Gamma Model?

The S-Gamma node, which has properties and is subject to common menu items, represents the S-Gamma phase model for particle size distribution.



As discussed above, you can choose to solve only one extra transport equation, for the second moment, or two extra transport equations for the zeroth and second moments. Use the Number of S-Gamma Equations property to make this selection.

S-Gamma Model Properties

We Crit	Specifies the critical Weber number We_{cr} for breakup (Eqn. (2060)) and coalescence (Eqn. (2076)).
Number of S-Gamma Equations	Selects the number of extra equations to solve in addition to the volume fraction. **One Equation** Solves for the transport of interfacial area density (the second moment). **Two Equation** Solves for both the transport of particle number density and the interfacial area density (zeroth and second moments) in Eqn. (2041).

Selecting the S-Gamma Model

Select the S-Gamma model for dispersed phases in a Eulerian Multiphase simulation. It is not intended to apply to continuous phases. Select the S-Gamma option from the *Optional Models* group of the *Phase Model Selection* dialog of the Eulerian phase, as shown below.



Following selection of the S-Gamma model, provide inputs for the particle mean diameter and its variance for flow boundaries and initial conditions.

If you wish to include the effects of breakup and coalescence on the predicted size distribution, select the S-Gamma Breakup and S-Gamma Coalescence models in a multiphase interaction that includes the dispersed phase.

What Is the S-Gamma Solver?

This solver, which has properties and is subject to common menu items, controls the solution update for the S-Gamma model.



S-Gamma Solver Properties

Post-Solver Relaxation Factor	This property is applied to the Under-Relaxation Factor to restrict the solution increment. The default value is `1.0`.
Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is `0.15`.

It is recommended that you use the same relaxation factor settings for this solver and the Volume Fraction solver.

Expert Properties

The properties in the *Expert* category are for a temporary debug situation, at the expense of simulation accuracy and higher memory usage. Do not change these properties unless you are thoroughly familiar with the STAR-CCM+ discretization techniques.

Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Uses reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Updates reconstruction gradients with each iteration.
Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration.	
	Activated	Retains temporary storage.
	Deactivated	Does not retain temporary storage.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.	
	Activated	Reconstruction gradients are maintained at zero.
	Deactivated	Reconstruction gradients are not maintained at zero.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Setting Mean Size and Variance (Boundaries and Initial Conditions)

Set the particle size distribution for flow boundaries and initial conditions. As the S-Gamma model assumes a log-normal particle size distribution, set the mean particle size and its variance to define the distribution.

Setting variance to zero results in a mono-dispersed phase (single particle size). If the one-equation option has been chosen, the variance at flow boundaries is ignored, and the initial value that is set for the continuum is used. For both initial conditions and boundary conditions, the process of setting the mean size and variance is the same:

- Choose the method for setting the size distribution
- Set the mean particle diameter
- Set the variance of particle diameter

Choosing the Size Distribution Method

When the S-Gamma model is selected for a Eulerian phase, a Size Distribution Specification node is added in two places:

- Within the Initial Conditions manager node for the Eulerian phase.



- Within the Physics Conditions folder for the Eulerian phase that is contained within the Phase Conditions folder for a boundary.



Child nodes are added to the companion Physics Values folder (in the case of a boundary) or Initial Conditions folder (in the case of initial conditions) depending on the chosen method. These nodes are:

- Sauter Mean Diameter
- Particle Size Variance

Size Distribution Specification Properties

Method	Controls whether to include particle size variance in the calculations.	
	Sauter Mean Diameter	Set the mean diameter only (see Eqn. (2043)).
	Sauter Mean Diameter and Variance of Diameter	Set the mean diameter and particle size variance (see Eqns. (2043) and (2045)). Adds a Particle Size Variance profile node to the object tree.

Setting the Mean Diameter

Use the Sauter Mean Diameter node to set the mean particle diameter of the log-normal distribution for boundary or initial conditions. This node functions as a scalar profile, and is shown for an initial condition below.



Setting the Particle Size Variance

Use the Particle Size Variance node to set the variance of the log-normal particle size distribution for boundaries and initial conditions. The node is added when the Size Distribution Specification node is set to include it. This node functions as a scalar profile, and is shown for a boundary condition below.



Coupling the Mean Size with the Interaction Length Scale

You can couple the mean size with the interaction length scale. Within the Multiphase Interaction model node, open the sub-nodes for the appropriate phase interaction. Open the User-Defined Interaction Length Scale node, and select the Interaction Length Scale node. In the properties, select `Field Function` as the method.



In the Scalar Function property of the Field Function method node, select the `Sauter Mean Diameter` field function for the appropriate phase. The example in the following screenshot shows this field function for the gaseous phase, Air.



Setting Region Sources for the Particle Size Distribution

The S-Gamma Source Option node is a phase condition at the region level. Its properties give you access to the source terms of the size distribution equations.



Setting the S-Gamma Source Terms

When you activate the S-Gamma Source Term property of the S-Gamma Source Option node (by activating its checkbox), scalar profile nodes appear in the object tree that let you specify sources in the size distribution equations.



These source terms are:

- S_0 and the source derivative (see Eqns. (2040), (2041), (2044) and (2046))
- S_2 and the source derivative (see Eqns. (2041), (2043) and (2049))

S-Gamma Source Option Properties

S-Gamma Source Term	Provides access to the S_0 and S_2 source terms.	
	Activated	Allows you to specify source terms in the object tree as scalar profiles.
	Deactivated	Leaves the source term calculations at their internal defaults.

Modeling Breakup with S-Gamma

The phase interaction within the Multiphase Interaction model node contains the objects for selecting and using the S-Gamma Breakup model, represented in the object tree by a model node with that name which has properties, and is subject to common menu items.



For breakup, the S-Gamma model considers the balance between disruptive forces (due to shear and turbulence) and restoring forces (due to surface tension) on the particle (such as a droplet). In laminar flows, the viscous effects dominate hence this regime is named "viscous breakup". In turbulent flows, the interactions with turbulence eddies dominate and this regime is named "inertial breakup". More details are provided in the formulation section for this model. It is possible to post-process certain results through a field function that this model activates.

S-Gamma Breakup Model Properties

Breakup Factor	A multiplier for breakup terms.

Selecting the S-Gamma Breakup Model

This model is selected on a Eulerian phase interaction. It becomes available when the continuous and dispersed phases have been set on the Eulerian Continuous-Dispersed model. The S-Gamma model is required as a phase model for the dispersed phase of the interacting pair.

In the *Phase Interaction Model Selection* dialog, select the S-Gamma Breakup option from the *Optional Models* group, as shown below.



The phase interaction model selection dialog shows the S-Gamma Breakup model activated.



Working with Viscous Breakup

In the S-Gamma Breakup multiphase interaction model node, the Viscous Breakup manager node lets you control the specific characteristics of the viscous breakup process with its properties.



Viscous Breakup Properties

Viscous Breakup Timescale	Sets how the breakup timescale is calculated.	
	Standard	Uses the default timescale from Eqn. (2058).
	User Defined	Activates the User-Defined Viscous Breakup Timescale method node.
Viscous Breakup Fragments	Sets the number of fragments that form during breakup.	
	Standard	Uses the default fragmentation number in Eqn. (2057).
	User Defined	Activates the User-Defined Viscous Breakup Fragments method node.
Viscous Critical Diameter	Sets how the critical diameter is calculated.	
	Standard	Uses the default critical diameter from Eqn. (2056).
	User Defined	Activates the User-Defined Viscous Critical Diameter method node.

Setting the User-Defined Viscous Breakup Timescale

Set the exponent of the bubble diameter in the timescale expression (Eqn. (2054)) as a property of the User-Defined Viscous Breakup Timescale node, and set the multiplier τ' using the Viscous Breakup Timescale sub-node. The Viscous Breakup Timescale sub-node functions as a scalar profile.



User-Defined Viscous Breakup Timescale Properties

Diameter Exponent	Sets the exponent a for the particle diameter in the timescale expression (Eqn. (2054)). The default setting is `1.0`.

Setting a User-Defined Viscous Breakup Fragment Number

Set the exponent b of the bubble diameter in the number of fragments expression (Eqn. (2053) as a property of the User-Defined Viscous Breakup Fragments node, and set the multiplier N' for the number of fragments using the Viscous Breakup Fragments sub-node. The Viscous Breakup Fragments sub-node functions as a scalar profile.



User-Defined Viscous Breakup Fragments Properties

Diameter Exponent	Sets the exponent b for the particle diameter in the fragment number expression (Eqn. (2053)). The default setting is `3.0`.

Setting the User-Defined Viscous Critical Diameter

Set the viscous critical diameter, d_{cr}, using the Viscous Breakup Dcrit sub-node of the User-Defined Viscous Critical Diameter node. The Viscous Breakup Dcrit sub-node functions as a scalar profile.



Working with Inertial Breakup

In the S-Gamma Breakup multiphase interaction model node, the Inertial Breakup manager node lets you control the specific characteristics of the inertial breakup process with its properties.



Inertial Breakup Properties

Concentration Correction Factor	C_α in Eqn. (2060). This is used in the calculation of the critical droplet diameter, since the presence of nearby droplets dampens the disruptive power of the inertial forces.	
Inertial Breakup Timescale Coefficient	Time coefficient k_{br} in Eqn. (2061).	
Inertial Breakup Timescale	Sets how the breakup timescale is calculated.	
	Standard	Uses the default timescale calculation from Eqn. (2061).
	User Defined	Activates the User-Defined Inertial Breakup Timescale method node.
Inertial Breakup Fragments	Sets how the breakup fragmentation is calculated.	
	Standard	Uses the default number of fragments from Eqn. (2062).
	User Defined	Activates the User-Defined Inertial Breakup Fragments method node.
Inertial Critical Diameter	Sets how the critical diameter is calculated.	
	Standard	Uses the default calculation from Eqn. (2060).
	User Defined	Allows user-defined parameters. Activates the User-Defined Inertial Critical Diameter method node.

For the critical diameter User Defined option, you can specify a constant or field function. For timescale and fragments use the function of the form $N(d) = N'd^b$, in which case N' is set as a constant or field function and b is set as a constant in the Diameter Exponent property (see Eqn. (2053)).

Setting the User-Defined Inertial Breakup Timescale

Set the exponent of the bubble diameter in the timescale expression (Eqn. (2054)) as a property of the User-Defined Inertial Breakup Timescale node, and set the multiplier τ' using the Inertial Breakup Timescale sub-node. The Inertial Breakup Timescale sub-node functions as a scalar profile.



User-Defined Inertial Breakup Timescale Properties

Diameter Exponent	Controls the exponent for the particle diameter in the timescale expression. The default setting is `1.5`.

Setting the User-Defined Inertial Breakup Fragment Number

Set the exponent b of the bubble diameter in the number of fragments expression (Eqn. (2053) as a property of the User-Defined Inertial Breakup Fragments node, and set the multiplier N' for the number of fragments using

the Inertial Breakup Fragments sub-node. The Inertial Breakup Fragments sub-node functions as a scalar profile.



User-Defined Inertial Breakup Fragments Properties

Diameter Exponent	Controls the exponent for the particle diameter in the fragment number expression. The default setting is `0.0`.

Setting the User-Defined Inertial Critical Diameter

Set the inertial critical diameter, d_{cr}, using the Inertial Breakup Dcrit sub-node of the User-Defined Inertial Critical Diameter node. The Inertial Breakup Dcrit sub-node functions as a scalar profile.



Modeling Coalescence with S-Gamma

The phase interaction within the Multiphase Interaction model node contains the objects for selecting and using the S-Gamma Coalescence model. The S-Gamma Coalescence model node that appears in the object tree has properties, and is subject to common menu items.



For coalescence, the S-Gamma model considers the probability of collisions of the particles (such as droplets), the contact time of two colliding particles and the drainage time of liquid film between the particles. Similar to the breakup model, there is the "viscous collision" (or viscous coalescence) regime and the "inertial collision" (or inertial coalescence) regime. The drainage time is a function of the state of the particle surface, whether it is fully or partially mobile or immobile. The model therefore considers the breakup and coalescence processes in great detail, as discussed in the formulation section for this model. It is possible to post-process certain results through field functions that this model activates.

S-Gamma Coalescence Model Properties

Coalescence Factor	Calibration coefficient, F_{cl} in Eqn. (2063).
Hamaker Constant	A_H, used in the calculation of critical film thickness (Eqn. (2072)).

Selecting the S-Gamma Coalescence Model

This model is selected on a Eulerian phase interaction. It becomes available when the continuous and dispersed phases have been set on the Eulerian Continuous-Dispersed model. The S-Gamma model is required as a phase model for the dispersed phase of the interacting pair.

In the *Phase Interaction Model Selection* dialog, select the S-Gamma Coalescence option from the *Optional Models* group, as shown below.



The phase interaction model selection dialog shows the S-Gamma Coalescence model activated.



Working with Viscous Coalescence

In the S-Gamma Coalescence multiphase interaction model node, the Viscous Coalescence manager node lets you control the specific characteristics of the viscous coalescence process with its properties.



Viscous Coalescence Properties

Viscous Collision Rate	Sets how the collision rate is calculated.	
	Standard	Calculates the collision rate using Eqn. (1984).
	User Defined	Activates the User-Defined Viscous Coalescence Collision Rate method node.
Viscous Coalescence Probability	Sets the mode of calculating the probability of a collision resulting in coalescence. The choice of this mode can affect the particle size, since increased mobility enhances coalescence, possibly leading to predictions of excessively large particle sizes (typically for bubbles or drops).	
	Fully Immobile Interface	Coalescence is blocked by contaminants in the fluid (Eqn. (2068)).
	Partially Mobile (Short Collision Time)	Coalescence is partially blocked by contaminants in the fluid (Eqn. (2069)).
	Partially Mobile (Quasi-Steady Flow in Film)	Coalescence takes place faster due to less blockage (Eqn. (2070)).
	Fully Mobile Interface	Coalescence takes place as quickly as possible--no material is blocking it (Eqn. (2071)).
	User Defined	Allows you to specify your own parameters. Activates the User-Defined Viscous Coalescence Probability method node.

Setting the User-Defined Viscous Coalescence Collision Rate

The default expression for K_{coll} in Eqn. (2066) can be replaced with a collision rate that you specify using the Viscous Coalescence Collision Rate node. This node, which functions as a scalar profile, is a child of the

User-Defined Viscous Coalescence Collision Rate node.



Setting the User-Defined Viscous Coalescence Probability

The default expression for coalescence probability P_{cl} in Eqn. (2067) can be replaced with a value you specify using the Viscous Coalescence Probability node. This node, which functions as a scalar profile, is a child of the User-Defined Viscous Coalescence Probability method node.



Working with Inertial Coalescence

In the S-Gamma Coalescence multiphase interaction model node, the Inertial Coalescence manager node lets you control the specific characteristics of the inertial coalescence process with its properties.



Inertial Coalescence Properties

Probability Factor	$k_{cl,2}$, used in the calculations for inertial collision probability (Eqn. (2074)).	
Inertial Collision Rate	Sets how the collision rate is calculated.	
	Standard	Uses the default collision rate calculation from Eqn. (2073).
	User Defined	Activates the User-Defined Inertial Coalescence Collision Rate method node.
Inertial Coalescence Probability	Sets the method of calculating the probability of a collision resulting in coalescence.	
	Standard	Uses the default coalescence probability calculation from Eqn. (2074).
	User Defined	Activates the User-Defined Inertial Coalescence Probability method node.

Setting the User-Defined Inertial Coalescence Collision Rate

The default expression (Eqn. (2073)) for inertial coalescence collision rate, $K_{coll,i}$, can be replaced with a value you specify using the Inertial Coalescence Collision Rate node. This node, which functions as a scalar

profile, is a child of the User-Defined Inertial Coalescence Collision Rate node.



Setting the User-Defined Inertial Coalescence Probability

The default expression (Eqn. (2074)) for the inertial coalescence probability can be replaced with a value you specify using the Inertial Coalescence Probability node. This node, which functions as a scalar profile, is a child of the User-Defined Inertial Coalescence Probability node.



S-Gamma Field Functions Reference

The following field functions are made available to the simulation when the S-Gamma model is used. These objects all have the same base properties. The phase name is appended, to the function name, such as SauterMeanDiameter of Air.

ParticleSizeVariance	σ in Eqn. (2045).
SauterMeanDiameter	d_{32} in Eqns. (2043) and (2045).
SecondSizeDistributionMoment	S_2 in Eqns. (2041), (2043) and (2049).
ThreeZeroDiameter	d_{30} in Eqns. (2044) and (2045).

ZerothSizeDistributionMoment	S_0 in Eqns. (2040), (2041), (2044) and (2046).

S-Gamma Breakup Field Functions Reference

The following extra field function is made available to the simulation when the S-Gamma Breakup model is used. Its name usually has the phase interaction appended, such as Inertial Breakup Dcrit of Water-Air.

Inertial Breakup Dcrit	The critical diameter d_{cr} in Eqn. (2060).

S-Gamma Coalescence Field Functions

The following field functions are made available to the simulation when the S-Gamma Coalescence model is used. The names of these functions usually have the name of the phase interaction appended, such as Inertial Coalescence Collision Rate of Water-Air.

Inertial Coalescence Collision Rate	The inertial collision rate $K_{coll,\,i}$ in Eqn. (2073).
Inertial Coalescence Probability	The inertial probability $P(d_{eq})$ in Eqn. (2074).
Viscous Coalescence Collision Rate	The viscous collision rate $K_{coll,\,v}$ in Eqn. (2066).
Viscous Coalescence Probability	The viscous probability $P_v(d_{eq})$ in Eqn. (2067).

Modeling Boundary Permeability

You can define a boundary as being permeable or impermeable for a particular phase. This setting can be applied to the following boundary types:

- Flow-split outlet

  You can specify that for a particular phase this boundary behaves like a wall (it is *Impermeable*).

- Wall

  You can specify that for a particular phase this boundary behaves like an outlet (it is *Permeable*). You can also set this boundary to be semi-permeable by specifying the proportion of the phase that is removed at the boundary.

These options are specified in the properties of the phases under the Phase Conditions node for the boundary.

An example where boundary permeability settings would be advantageous is a case modeling a bubble column containing air and water. In this case, you could specify a degassing boundary at the outlet, which would allow air to exit the fluid domain but not water. This would be achieved by selecting the Phase Impermeable option for the water phase at the outlet boundary. With such an approach, the mass of water within the column is conserved, while allowing the air bubbles to escape. The code ensures that the mass flow of air at the inlet and outlet of the domain is balanced. The volume fractions of each phase throughout the domain are then computed according to continuity.

An example involving a permeable wall phase condition is a case that involves droplet deposition on a wall. The permeable wall approach assumes that when droplets impinge on a surface they stick there and hence are removed from the flow field. The permeable wall phase condition says that the phase leaves the flow domain at the rate that it approaches the surface. It is also assumed that the deposition layer is sufficiently thin so as not to affect the flow field. Where not all the droplets stick to the wall, the wall permeability factor could be adjusted to limit the mass flux that is removed at the wall.

Phase Condition Properties

Type	The phase-specific boundary option.	
	Flow-Split Outlet	Normal outlet behavior (default option for flow-split outlets).
	Phase Impermeable	Specifies that the outlet behaves like a wall for the selected phase. This option is available on flow-split outlets.

Phase Permeable	Specifies that the wall behaves like an outlet for the selected phase. This option is available on wall boundaries. Also allows you to define a semi-permeable wall using the Wall Permeability setting.
Wall	Normal wall behavior (default option for wall boundaries).

See also: common menu items.

Defining an Impermeable Outlet

When a flow-split outlet is defined as impermeable for a particular phase, it behaves like a wall for that phase.

To define an impermeable outlet phase condition, first set up your physics continuum, phases, regions, and boundaries. Make sure that the outlet boundary type is set to Flow-Split Outlet. Next, expand the Phase Conditions

node for the outlet boundary. Select the node for the phase that is associated with the Impermeable condition. In the *Properties* window, change the Type to Phase Impermeable.



The icon for the phase changes from to , indicating that the outlet behaves like a wall for that phase.



Defining a Permeable Wall

The permeable wall phase condition allows the selected phase to permeate though the wall, thus the wall behaves like an outlet for that phase. Other phases in the domain still treat the boundary as a wall. This phase condition can be used to model the deposition of a phase on a surface, or to model a gas leaving a liquid at a free surface (degassing), for example. At a phase permeable wall, any component of phase velocity parallel to the wall is treated according to the wall boundary conditions. For example, zero for a non-slip wall in a stationary geometry case.

You can define the proportion of mass flux that is removed from the phase at the permeable wall by defining the Wall Permeability factor. When the Wall Permeability factor is set to 1.0, the wall behaves as a pure outlet for that phase. Setting the Wall Permeability to a value from 0 through 1 results in a semi-permeable wall.

To define a permeable wall phase condition, first set up your physics continuum, phases, regions, and boundaries. Next, expand the Phase Conditions node for the wall boundary. Select the node for the phase that is associated with the Permeable condition. In the *Properties* window, change the Type to Phase Permeable.



When a phase is associated with a Permeable wall condition, a Physics Values > Wall Permeability node becomes available for the phase. You can adjust the proportion of mass flux that is removed from the phase at the wall boundary by entering a value for the Wall Permeability factor. A value of 1.0 is a pure outlet, and a value of 0.0 is a wall.



Using Mass Source Terms in Eulerian Multiphase Flows

Mass source terms can be used for Eulerian multiphase flows. To activate mass source terms in a region select the Mass Source Option in the region *Physics Conditions*, and tick the checkbox in the *Properties* window. This reveals three nodes in the *Physics Values* for the region:

- Mass Source
- Mass Source Pressure Derivative

- Mass Source Volume Fraction Derivative



Mass Source for Eulerian Multiphase Flows

This node is added to the region physics values when the Mass Source Option region condition is activated. The values for the mass source can be entered as constant or composite values. The dimensions of the mass source term are mass per volume per time.

Note: When you create a mass source for a particular phase, create source terms for other transport equations. The source terms describe the inflow or outflow of other variables that are carried on the mass source, for example, momentum, turbulence and energy.

Mass Source Pressure Derivative for Eulerian Multiphase Flows

This node represents linearization of the phase Mass Source with respect to pressure. The values for the mass source pressure derivative can be entered as constant or composite values. The derivative of the source term is used to linearize the overall volume correction equation. The default value is zero. Providing a value for the derivative helps to stabilize the solution.

Mass Source Volume Fraction Derivative for Eulerian Multiphase Flows

This node represents linearization of the phase Mass Source with respect to volume fraction. The values for the mass source volume fraction derivative can be entered as constant or composite values. The derivative of the source term is used to linearize the phase mass balance equations. The default value is zero. Providing a value for the derivative helps to stabilize the solution.

Defining Mass Source Terms Using a Constant or Composite Method

In Eulerian multiphase simulations, two techniques are available for defining mass source terms:

- Constant
- Composite

The constant method allows you to define constant values for all phases within a single property (set using an array profile). The composite method provides separate nodes for each phase, allowing you to define the source terms using different techniques. For example, the composite method allows you to define the source term for each phase using a different field function.

To select either the constant or composite method, first activate the mass source option for the corresponding region. Next, select the Mass Source, Mass Source Pressure Derivative, or Mass Source Volume Fraction physics value node. In the *Properties* window, set the Method property to the desired option.

In the screenshot below, the composite method is selected for the Mass Source, and the constant method has been selected for the two mass source derivatives.



Defining Constant Values for All Phases

The Constant method allows you to specify constant mass source values for all phases within a single property. It is also possible to define constant values using the composite technique, but it is necessary to set these values under separate nodes in the simulation tree.

To define values using the constant method, make sure that the method for the corresponding mass source node is set to Constant. Select the Constant node and enter values in the profile array.

To check which array value corresponds to which phase, click the ... for the Value property. This opens the *Constant - Value* dialog, which lists each phase separately.

Defining Values Using a Composite Method

The Composite method allows you to specify mass source values for each phases independently. It provides a separate node for each phase, which can be used to enter values using field functions, for example.

To define values using the constant method, make sure that the method for the corresponding mass source node is set to Composite. Expand the nodes for the appropriate phase, and define the value or field function in the usual way.

Multiphase Segregated Flow Formulation

This section contains the formulation for the models and methods that are related to the Multiphase Segregated Flow solver.

- Conservation equations of mass and momentum
- Phase interactions
- Interaction length scales
- Interaction area density
- Suspension and emulsion rheology
- Interphase momentum transfer
- Drag force: overview
- Drag force: models
- Drag force: linearized drag coefficient methods
- Drag force: standard drag coefficient methods
- Drag force: drag correction methods
- Turbulent dispersion force
- Virtual mass force
- Lift force
- Wall lubrication force
- Granular pressure
- Solid pressure force
- Multiphase turbulence modeling
- Dispersed Phase Turbulence models
- Particle Induced Turbulence models
- Wall permeability
- Energy conservation equation
- Interphase energy transfer
- Interphase boiling mass transfer
- Multi-component boiling mass transfer
- Thermal boundary conditions on walls
- Wall boiling
- Multi-component mass transfer
- Dissolution mass transfer
- S-Gamma model for particle size distribution

- S-Gamma breakup model formulation
- S-Gamma coalescence model formulation
- Bibliography

Conservation Equations of Mass and Momentum

The conservation equations for a generic phase i are:

Continuity

$$\frac{\partial}{\partial t}(\alpha_i \rho_i) + \nabla \cdot (\alpha_i \rho_i \mathbf{v}_i) = \sum_{j \neq i} (m_{ij} - m_{ji}) + S_i^{\alpha} \tag{1692}$$

where:

α_i is the volume fraction of phase i

ρ_i is the density of phase i

$\mathbf{v}_i$ is the velocity of phase i

S_i^{α} is the phase mass source term

m_{ij} is the mass transfer rate to phase i, from phase j $(m_{ij} \geq 0)$

m_{ji} is the mass transfer rate to phase j, from phase i $(m_{ji} \geq 0)$

In addition, the volume fractions must satisfy the equation:

$$\sum_i \alpha_i = 1 \tag{1693}$$

Momentum

$$\frac{\partial}{\partial t}(\alpha_i \rho_i \mathbf{v}_i) + \nabla \cdot (\alpha_i \rho_i \mathbf{v}_i \mathbf{v}_i) = \tag{1694}$$

$$-\alpha_i \nabla p + \alpha_i \rho_i \mathbf{g} + \nabla \cdot [\alpha_i (\tau_i + \tau^t{}_i)] + \mathbf{M}_i + (\mathbf{F}_{int})_i + \mathbf{S}_i^v + \sum (m_{ij} \mathbf{v}_j - m_{ji} \mathbf{v}_i)$$

where:

p is the pressure, assumed to be equal in both phases

$\mathbf{g}$ is the gravity vector

τ_i and τ_i^t are the molecular and turbulent stresses, respectively

$\mathbf{M}_i$ is the interphase momentum transfer per unit volume

$\mathbf{S}_i^v$ is the phase momentum source term

$(\mathbf{F}_{int})_i$ represents internal forces (such as solid pressure force between particles)

m_{ij} is the mass transfer rate from phase j to phase i

m_{ji} is the mass transfer rate from phase i to phase j

The interphase momentum transfer represents the sum of all the forces the phases exert on one another and satisfies the equation:

$$\sum M_i = 0$$

Phase Interactions

Different kinds of multiphase flow are characterized by their Phase Interactions. Each interaction identifies:

- which phases interact (commonly a phase pair that is separated by a distinct interface)
- the interface topology (continuous-disperse is a standard model)
- models for the interface length scale and area density
- models for interphase transfers across the interface.

Phase Interactions are typically used to define a single flow regime for example 'bubbly flow' or 'droplet flow'.

It is also possible to allow more than one Phase Interaction between each pair of phases, for example where the interface topology changes in different parts of the flow. Models that you define for interphase scales and transfers can then be included in which Phase Interaction is dominant under different local conditions.

Interaction Length Scales

For continuous-dispersed phase interactions, the interface length scale is taken to be an effective mean diameter of the dispersed phase particles. Where particles are not spherical, this is absorbed as a correction factor into interphase transfer models such as for bubble drag and lift force. In general, the Sauter mean diameter is preferred since this gives the correct mean interfacial area where there is a distribution of particle sizes. If the S-gamma model for predicting particle size distribution is used, the result for Sauter mean diameter can be selected as a field function for the interaction length scale.

The interaction length scale, l_{cd} in a continuous-dispersed phase interaction, can be provided as a constant or a field function. For boiling flows, when boiling mass transfer is active, the Kurul Podowski length scale is provided as an alternative.

Kurul Podowski Interaction Length Scale

This is an empirical relation that is used in boiling flows, which specifies bubble size as a function of local subcooling. The length scale is predicted as

$$l_{cd} = \frac{D_{min}(\Delta T - \Delta T_{D,\,min}) + D_{max}(\Delta T_{max} - \Delta T_{D,\,min})}{\Delta T_{D,\,max} - \Delta T_{D,\,min}} \tag{1695}$$

where:

D_{min} and D_{max} are the minimum and maximum bubble diameters, respectively (with default values of 0.15 mm and 2.0 mm)

$\Delta T_{D,\,min}$ is the liquid subcooling corresponding to the minimum bubble diameter (with a default value of 13.5 K)

$\Delta T_{D,\,max}$ is the liquid subcooling corresponding to the maximum bubble diameter (with a default value of -5.0 K)

$\Delta T = T_{cd} - T_c$ where T_{cd} is the interface temperature for the phase interaction--usually the saturation temperature. T_c is the temperature of the continuous phase in the interaction pair.

Interaction Area Density

The interaction area density specifies the interfacial area available for drag, heat, and mass transfer between each pair of phases in an interaction. Heat and mass transfer models use the interaction area density directly, while drag models use one quarter of the interfacial area as an estimate of the projected area. Any correction factors, such as for non-spherical particle shapes or particle crowding, are assumed to be covered in correlations for the drag, heat, and mass transfer coefficients. There are two built-in models for interaction area density in STAR-CCM+: spherical particle and symmetric. Additionally, an interaction area density can also be supplied.

Spherical Particle Interaction Area Density

This is based on the surface area of a spherical particle, and is not applicable at high particle volume fractions. For this model, the interaction area density a_{cd} between a continuous phase (subscript c) and a dispersed phase (subscript d) is:

$$a_{cd} = \frac{6\alpha_d}{l_{cd}} \tag{1696}$$

where:

α_d is the volume fraction for the dispersed phase in the interacting phase pair

l_{cd} is the interaction length scale for the pair

Symmetric Particle Interaction Area Density

The availability of the continuous phase modifies the Symmetric Particle Interaction Area Density, and so can be used for a wide range of volume fractions:

$$a_{cd} = \frac{6\alpha_c\alpha_d}{l_{cd}} \tag{1697}$$

where:

α_d is the volume fraction for the dispersed phase in the interacting phase pair

α_c is the volume fraction for the continuous phase in the pair

l_{cd} is the interaction length scale for the pair

Suspension and Emulsion Rheology

The presence of particles suspended in a Newtonian liquid is known to lead to non-Newtonian behavior, as the mixture viscosity becomes dependant on the volume fraction of the particle phase. For a Eulerian multiphase flow, you can model this behavior using the Suspension Rheology and the Emulsion Rheology phase interaction models.

Relative Viscosity Material Property

In suspension rheology, the dimensionless relative viscosity is used to describe the numerics of the mixture viscosity. The viscosity tends to infinity as the dispersed phase reaches the limit of maximum packing.

The relative viscosity is:

$$\eta_r = \frac{\eta}{\mu_c} \tag{1698}$$

where:

- η is the mixture viscosity.
- μ_c is the continuous phase (Newtonian) viscosity.

There are many models with a similar form. One of the earliest from the study of suspensions is the Krieger-Dougherty model [342].

The Krieger-Dougherty Model Formulation

The Krieger-Dougherty model was based on the assumption that crowding plays a vital role in the origin of non-Newtonian flow behavior in rigid sphere suspensions. Measurements on latex systems did not detect non-Newtonian behavior at concentrations below 20% by volume of suspended polymer. The Krieger-Dougherty model takes explicit notice of the interactions between neighboring spherical particles. The resulting flow equation compares well with experimental viscometric data on synthetic latexes and solutions of high polymers.

In STAR-CCM+, the Krieger-Dougherty model has the form:

$$\eta_r = \left(1 - \frac{\phi}{\phi_m}\right)^{-[\eta]\phi_m} \tag{1699}$$

where:

- $[\eta]$ is the intrinsic viscosity. $[\eta]$ = 2.5 for spherical particles.
- ϕ is the volume fraction.

- ϕ_m is the maximum critical packing fraction. For hard spheres, ϕ_m = 0.645.

Morris and Boulay Model Formulation

In STAR-CCM+, the relative shear viscosity of the Morris and Boulay model [352] is defined as:

$$\eta_r(\phi) = 1 + 2.5\phi\left(1 - \frac{\phi}{\phi_m}\right)^{-1} + K_s\left(\frac{\phi}{\phi_m}\right)^2\left(1 - \frac{\phi}{\phi_m}\right)^{-2} \tag{1700}$$

The free parameters in this model are as follows:

- K_s is the contact contribution
- ϕ_m is the maximum packing fraction.

Normal Relative Viscosity Material Property

Krieger-Dougherty Model

The Krieger-Dougherty model has a constant zero normal relative viscosity.

Morris and Boulay Model

The normal relative shear viscosity of the Morris and Boulay model is defined as:

$$\eta_r(\phi) = K_n\left(\frac{\phi}{\phi_m}\right)^2\left(1 - \frac{\phi}{\phi_m}\right)^{-2} \tag{1701}$$

The free parameters in this model are as follows:

- K_n is the contact contribution
- ϕ_m is the maximum packing fraction.

Anisotropy Tensor Formulation

The Morris and Boulay model was derived from experiments in pipes and Couette devices where the flow becomes aligned in a particular direction governed by the geometry. The normal stress was shown to be anisotropic and, when varying in strength, lead to migration of the particles in different directions.

The normal part of the stress is defined as:

$$\tau_{p,\,NS} = -\eta_n(\phi)\eta_f\dot{\gamma}Q \tag{1702}$$

where

- ϕ is the particle volume fraction
- η_f is the continuous liquid viscosity
- $\dot{\gamma}$ is the shear rate of the liquid.
- Q is the anisotropy tensor.

X-direction Flow Aligned Anisotropy Tensor

For shear flow in the x direction:

$$Q = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} \tag{1703}$$

where the anisotropy parameters are:

- $\lambda_1 = 1.0$
- $\lambda_2 = 0.8$
- $\lambda_3 = 0.5$

Isotropic Form of Anisotropy Tensor

An isotropic version gives equal weighting to all directions by setting the Anisotropy tensor equal to the identity tensor [351].

$$Q = I \tag{1704}$$

This form is valid for 3D simulations, but the contributions are no longer anisotropic.

Interphase Momentum Transfer

In STAR-CCM+, the momentum transfer term can include contributions from the drag, virtual mass, lift, and turbulent dispersion forces:

$$\mathbf{M}_i = \sum_{j \neq i} (\mathbf{F}_{ij}^{D} + \mathbf{F}_{ij}^{VM} + \mathbf{F}_{ij}^{L} + \mathbf{F}^{TD}{}_{ij}) \tag{1705}$$

where $\mathbf{F}_{ij}$ is the force per cell volume, on phase i, due to phase j. It follows that

$$\mathbf{F}_{ji} = -\mathbf{F}_{ij}$$

The individual forces are described below along with the submodels available to calculate them.

Drag Force: Overview

Overview of the Drag Force Model

The drag force options available in STAR-CCM+ are summarized in the table

Drag Force	Linearized Drag Coefficient	Drag Coefficient	Drag Correction
Linearized	Constant Field Function Gidaspow (P) Syamlal O'Brien (P) Symmetric Drag Coefficient (M)		
Standard		Constant Field Function Schiller-Naumann (S) Hamard and Rybczynski (S) Tomiyama (B) Bozzano-Dente (B) Wang Curve Fit (A)	Field Function Volume Fraction Exponent Richardson Zaki (S) Lockett Kirkpatrick (B) Simonnet (B)

The options are qualified by the main application areas:

- (A) air bubbles in water systems only.
- (B) bubbles.
- (M) fluid-fluid mixtures in separation applications.

- (P) solid particles at high concentration.
- (S) spherical particles at moderate concentration - including small droplets or bubbles.

The Drag Force model has two top-level options depending on whether the model is based on a Linearized Drag Coefficient or a Standard definition of Drag Coefficient:

- Linearized – used for integrated models for the interphase drag or friction force.
- Standard – used for drag force models that are composed from a Standard Drag Coefficient for a single particle together with a Drag Correction factor to account for multi-particle and other effects.

The Linearized Drag Coefficient can be used for specifying complete models for particle drag that already include multi-particle effects.

The options available for Linearized Drag Coefficient are:

- Constant
- Field Function
- Gidaspow (Solid Particles)
- Syamlal O'Brien (Solid Particles)
- Symmetric Drag Coefficient (Fluid-Fluid Mixtures)

The Gidaspow and Syamlal O'Brien methods are suitable for fluidized beds and settling beds.

The Symmetric Drag Coefficient is suitable for combinations of gas and liquid phases where different phases are regarded as being the continuous phase in different locations of the same flow.

The Field Function option can be used to develop new models for effects that are not covered in the built-in models, such as corrections for particle shape. It could also be useful for modeling interphase friction force for interface topologies other than continuous-disperse.

The Standard Drag Coefficient model describes the standard drag coefficient for a single-particle, and the following options are available:

- Constant
- Field Function
- Schiller-Naumann (Spherical particles, droplets, or bubbles)
- Tomiyama (Bubbles)
- Bozzano-Dente (Bubbles)

- Wang Curve Fit (Air/Water Bubbles)

Schiller-Naumann is the recommended option for undeformable spherical particles. This covers solid particles, as well as small spherical fluid droplets or bubbles with negligible internal circulation due to their internal viscosity or their surface contamination.

Tomiyama is the recommended option for all bubbly flows as it is applicable for a wide range of bubble sizes from small spherical bubbles to the surface tension dominated regime. This option is also applicable for a wide range of liquid viscosity and for three levels of contamination by surfactants.

The Drag Correction model describes how concentration modifies the single-particle Drag Coefficient model in a multiparticle system. The following options are available:

- Field Function
- Volume Fraction Exponent
- Richardson Zaki (Spherical particles, droplets, or bubbles)
- Lockett Kirkpatrick (Bubbles)
- Simonnet (Bubbles)

The Richardson Zaki method is the recommended option for undeformable spherical particles.

For bubbly flows, four drag correction methods are suggested in the table below together with the approximate bubble size or the expected bubble behavior as a guideline. This is because the trends, conditions, and boundaries of the various bubble behavior regimes are the subject of continued research. Further possibilities can be investigated with the help of the field functions for drag coefficient and for drag correction factor.

Bubble Regime	**Air / Water Bubble Size (d)**	**Non-dimensional Size**	**Bubble Behaviors**	**Suggested Drag Correction Method**
Small spherical	< 2.75 mm	$Eo < 1$	Hindering	Richardson Zaki
Small ellipsoidal	~ 5 mm	$Eo \sim 3.3$	Hindering Deforming	Lockett Kirkpatrick

Bubble Regime	Air / Water Bubble Size (d)	Non-dimensional Size	Bubble Behaviors	Suggested Drag Correction Method
Intermediate size	~ 7-10 mm	Eo ~ 6.6-13.4	Hindering: 0-15% void fraction Swarming: 15-30% void fraction	Simonnet
Large spherical-cap in churn-turbulent flow	~ 11-14 mm	We (drift velocity) ~ 8	Breakup Coalescence Swarming	Volume Fraction Exponent

Drag Force: Models

Non-Dimensional Groups

The particle Reynolds number is used to characterize the drag of particles of equivalent diameter d, and traveling at a general slip velocity V, relative to a continuous phase c.

$$Re = \frac{\rho_c V d}{\mu_c} \tag{1706}$$

The particle Eotvos number Eo is used to characterize shape and drag of deformed bubbles traveling at terminal velocity, by comparing buoyancy force over surface tension force.

$$Eo = \frac{|\rho_c - \rho_d| g d^2}{\sigma} \tag{1707}$$

The particle Weber number We is used to characterize the shape of a deformed bubble traveling at a general slip velocity V, and its susceptibility to breakup. The Weber number compares dynamic pressure distribution around the bubble over surface tension force.

$$We = \frac{\rho_c V^2 d}{\sigma} \tag{1708}$$

The Morton number is a function of fluid properties only. It has an important role in bubble drag laws in locating the transition point between Reynolds dependent behavior and surface tension dominated behavior, mainly as a function of fluid viscosity.

$$Mo = \frac{g\mu_c^4 \left|\rho_c - \rho_d\right|}{\rho_c^2 \sigma^3} \tag{1709}$$

Linearized Drag Coefficient

The force acting on phase i due to the drag of phase j is implemented as

$$\mathbf{F}_{ij}^D = A_{ij}^D (\mathbf{v}_j - \mathbf{v}_i) \tag{1710}$$

where A_{ij}^D is the linearized drag coefficient.

For continuous-disperse flows, the linearized drag coefficient is related to a standard engineering definition of drag coefficient C_{ij}^D, for particles by:

$$A_{ij}^D = C_{ij}^D \frac{1}{2} \rho_c \left|\mathbf{v}_j - \mathbf{v}_i\right| \left(\frac{a_{ij}}{4}\right) \tag{1711}$$

where a_{ij} is the interfacial area density. The factor $\frac{a_{ij}}{4}$ represents the projected area of the equivalent spherical particle, and any corrections for the shape of non-spherical particles can be absorbed into the model for the drag coefficient C_{ij}^D.

Standard Drag Coefficient

The standard drag coefficient C_{ij}^D is often constructed from two factors.

$$C_{ij}^D = f_{ij}^D C_{ij\infty}^D \tag{1712}$$

where:

$C_{ij\infty}^D$ is the single-particle drag coefficient and is determined from the measurement of terminal velocity of a particle in an unbounded stagnant continuous phase

f_{ij}^D is a drag correction factor and is determined from measurement of terminal velocity for multiple particles of the same size traveling together at a given concentration. This correction can cover a number of

physical effects which can be difficult to separate in practice, such as hindering due to increased turbulence interaction, or co-operative motion due to swarming or coalescence

Force Balance Under Terminal Velocity Conditions

Under terminal velocity conditions, the pressure buoyancy and drag forces are in equilibrium and the two-phase momentum conservation equations for the continuous and dispersed phases reduce to the result:

$$V^2 = \frac{8\alpha_c\alpha_d}{a}\frac{\Delta\rho}{\rho_c}\frac{g}{C_D} \tag{1713}$$

where:

$\Delta\rho = |\rho_c - \rho_d|$

a is the interfacial area density for the two-phase interaction

There are two alternative standard models for interfacial area density in STAR-CCM+, which leads to:

$$V^2 = \begin{cases} \frac{4}{3}\frac{\Delta\rho}{\rho_c}\frac{gd}{C_D}\alpha_c & \text{for Spherical Particle Interfacial Area} \\ \frac{4}{3}\frac{\Delta\rho}{\rho_c}\frac{gd}{C_D} & \text{for Symmetric Interfacial Area} \end{cases} \tag{1714}$$

Eqn. (1714) is a key equation for the calibration of single-particle drag models and multi-particle correction factors from experimental measurements of terminal velocity V as a function of particle size d, and continuous phase volume fraction α_c.

The results of such a two-phase calibration are applied in the multiphase context, and in flows where particles are not traveling at terminal velocity. The terminal velocity is replaced with the computed slip velocity $|\mathbf{v}_j - \mathbf{v}_i|$ when determining the drag coefficient C_{ij}^D and drag correction factor f_{ij}^D for each phase interaction ij.

Role of Interfacial Area Density Method in Force Balance

The Spherical Particle Interfacial Area method is geometrically exact for spherical particles of a given diameter. Any corrections for non-spherical shape can be absorbed into the drag coefficient. Regardless of particle shape, the term $\alpha_c \Delta\rho g d$ in the force balance correctly reflects the change in effective buoyancy force with concentration due to the reducing difference in density between the equivalent-volume spherical particle and the two-phase mixture as α_c falls.

The Symmetric Interfacial Area method is used where robustness is required across a wide range of volume fraction, so that the maximum packing of spherical particles can be exceeded. If an additional factor of $1/\alpha_c$ is absorbed into the drag correction factor that is used with the Symmetric Interfacial Area, both models give the same result, up to the maximum packing point. This adjustment is done automatically for the standard drag correction factor models that are implemented in STAR-CCM+. However, include the adjustment explicitly when using the Symmetric Interfacial Area Density for drag correction models from field functions or exponents.

Special Considerations for Bubble Drag Coefficients

The drag laws for bubbles which allow for reduced drag due to internal circulation includes Tomiyama/Pure, Tomiyama/Moderately Contaminated, Bozzano-Dente, and the Wang drag law. All of these are implemented in STAR-CCM+ as a function of (Re, Eo, Mo).

These laws have been formulated in the literature to reproduce the correct drag coefficient at terminal velocity. However, the laws, can result in unphysically low estimates of drag coefficient for small bubbles traveling at high slip velocity.

In STAR-CCM+, this problem is avoided by imposing a lower limit on the bubble drag coefficient. This limit is based on the lowest value of drag coefficient observed for a bubble of any size when it is traveling at its terminal velocity. This point corresponds to the slip velocity just before the bubble starts to distort from its spherical shape. It corresponds to a Weber number of about 2, so this is a physically meaningful constraint to apply at all bubble speeds.

Drag Correction Factors

The objective of drag correction methods is to reproduce experimental terminal velocities correctly as a function of particle concentration.

The traditional approach of Richardson and Zaki [356] is widely applicable for spherical particles of a given size. Drag correction for bubble columns is more complex because the bubble size can change as a function of concentration. In traditional approaches (for example, a model for the churn turbulent regime [338]), bubble drag correction has implicitly included the effects of turbulence interactions and of coalescence and breakup. A more recent trend is to isolate the effect of concentration alone on the terminal velocity of bubbles of a specified size (for example Simmonet *et al.* [361]).

To consider all these possibilities, define correction factors for particle size, terminal velocity, and drag coefficient. Let:

$$f_d = \frac{d}{d_\infty} \tag{1715}$$

$$f_V = \frac{V}{V_\infty} \tag{1716}$$

$$f_D = \frac{C_D}{C_{D\infty}} \tag{1717}$$

where the subscript ∞ indicates conditions for a single particle in an unbounded stagnant fluid.

Setting $\alpha_c = 1$ in the force balance Eqn. (1714), gives the force balance for a single particle traveling at terminal velocity in an unbounded medium:

$$V_\infty{}^2 = \frac{4}{3}\frac{\Delta\rho}{\rho_c}\frac{gd_\infty}{C_{D\infty}} \tag{1718}$$

This result applies for either of two interaction area density models. Then taking the ratio of Eqn. (1714) over Eqn. (1718) leads to the following relation between the three correction factors:

$$f_V^2 = \begin{cases} \dfrac{\alpha_c f_d}{f_D} & \text{for Spherical Particle Interfacial Area} \\[2ex] \dfrac{f_d}{f_D} & \text{for Symmetric Interfacial Area} \end{cases} \tag{1719}$$

The STAR-CCM+ model does not include the size correction factor f_d when computing f_V for two reasons:

- Modern computational methods focus on drag and drag correction for particles of a specified size, and particle size distribution is considered as a separate model.
- Size change is only relevant for drag correction of large bubbles whose drag coefficient is independent of Re.

Working with drag correction for particles of a specified size, the above relationship reduces to:

$$f_V^2 = \begin{cases} \dfrac{\alpha_c}{f_D} & \text{for Spherical Particle Interfacial Area} \\ \dfrac{1}{f_D} & \text{for Symmetric Interfacial Area} \end{cases} \tag{1720}$$

In STAR-CCM+ the relationship is used to:

- implement a drag correction factor f_D for models that are specified in literature as using a velocity correction factor f_V.
- obtain a velocity correction factor f_V for models that are specified in the literature as using drag correction factor f_D. The velocity correction factor is required to evaluate the drag coefficient for a single particle.

This last point is important in reproducing terminal velocity where the drag coefficient $C_{D\infty}$ depends on the particle Reynolds number. Evaluate the single-particle drag coefficient at the single particle terminal velocity V_∞ not at the terminal velocity V for which you are trying to estimate the drag coefficient in a multi-particle system. Apply the velocity correction factor f_V to the slip Reynolds number:

$$C_{D\infty} \equiv C_D(Re_\infty) = C_D\left(\frac{Re}{f_V}\right) \tag{1721}$$

Output Variables

To help you follow these drag computations, the following fields are available for inspection for each phase pair in a STAR-CCM+ simulation:

- The particle Eotvos number, Eo. This field only applies where the drag model depends on surface tension.
- The particle slip velocity Reynolds number, Re

- The single particle Reynolds number, Re_∞
- The single particle drag coefficient, $C_{D\infty}$
- The drag correction factor, f_D

Fixed Exponent Regimes

Where correction factors can be modeled as the continuous phase volume fraction raised to some constant power, that is:

$$f_V = \frac{V}{V_\infty} = \alpha_c^n \tag{1722}$$

$$f_D = \frac{C_D(Re,\alpha_c)}{C_D(Re_\infty,1)} = \alpha_c^{n_D} \tag{1723}$$

Then equation [1713] leads to a simple relation between these exponents:

$$n_D = \begin{cases} 1-2n & \text{for Spherical Particle Interfacial Area} \\ -2n & \text{for Symmetric Interfacial Area} \end{cases} \tag{1724}$$

For example, under Stokesian conditions ($Re_\infty < 0.1$), the constant Richardson Zaki [356] exponent of $n = 4.65$ for velocity correction leads to a drag correction exponent of:

$$n_D = -8.3 \tag{1725}$$

when using the Spherical Particle Interfacial Area, or:

$$n_D = -9.3 \tag{1726}$$

when using the Symmetric Interfacial Area.

Drag Force: Linearized Drag Coefficient Methods

The options available for Linearized Drag Coefficient are:

- Arastoopour (Solid Particles)
- Constant

- Field Function
- Gidaspow (Solid Particles)
- Syamlal O'Brien (Solid Particles)
- Symmetric Drag Coefficient (Fluid-Fluid Mixtures)

Arastoopour Drag Model

The Arastoopour Drag Model is appropriate for solid dispersed-phase applications such as fluidized beds, where particle loadings are high.

The Arastroopour drag coefficient [312] is:

$$A_{ij}^{D} = \left(0.336 + \frac{17.3}{Re_d}\right)\frac{\alpha_d \rho_d |\mathbf{v}_r|}{l_{cd}}\alpha_c^n \tag{1727}$$

Where the exponent n is set to -2.8 by default.

Gidaspow Drag Model

The Gidaspow Drag Model is appropriate for solid dispersed-phase applications such as fluidized beds where particle loadings are high. The most widely used formulae are the Ergun equation for regions of high solid particle concentrations and a modified Stokes law for regions of lower concentration [318], [373], [358].

A transition volume fraction, α_{tr}, is used as the basis for switching between these formulae as follows:

$$A_{ij}^{D} = \begin{cases} \dfrac{150\alpha_d^2 \mu_c}{\alpha_c l_{cd}^2} + \dfrac{1.75\alpha_d \rho_c |\mathbf{v}_r|}{l_{cd}} & \alpha_d \geq \alpha_{tr} \\ \dfrac{3}{4} C_D \dfrac{\alpha_d \rho_c}{l_{cd}} |\mathbf{v}_r| \alpha_c^n & \alpha_d < \alpha_{tr} \end{cases} \tag{1728}$$

The drag coefficient C_D in Eqn. (1728) is from the correlation of Schiller and Naumann Eqn. (1737).

Even below the transition volume fraction, $\alpha_d \leq \alpha_{tr}$, the inter-particle force affects the particle velocity. This hindered settling is accounted for by the correction factor α_c^n in Eqn. (1728).

Specify both the value of the transition volume fraction (by default α_{tr} = 0.2) and the hindered-settling exponent (by default n = -1.65).

Note: The Gidaspow drag coefficient method reduces to the Wen and Yu correlation when the transition volume fraction is set to 1.0.

Syamlal O'Brien Drag Model

The Syamlal O'Brien drag coefficient is based on measurements of terminal velocity in settling beds. The drag coefficient defines fluid-solid momentum exchange as a function of the particle phase fraction, particle Reynolds number, and a terminal velocity expression:

$$A_{cd}^{D} = A_{dc}^{D} = \frac{3}{4}\frac{\alpha_d \alpha_c \rho_c C_D}{V_{rs}^2 l_{cd}}\left|\mathbf{v}_r\right| \tag{1729}$$

where the bulk of the terms are the same as for the linearized drag coefficient, with the addition of the terminal velocity correlation V_{rs} for the dispersed phase. V_{rs} is:

$$V_{rs} = \frac{1}{2}\left[A - 0.06Re_d + \sqrt{(0.06Re_d)^2 + 0.12Re_d(2B - A) + A^2}\right] \tag{1730}$$

in which

$$A = \alpha_c^{4.14} \tag{1731}$$

and

$$B = \begin{cases} C_2\alpha_c^{1.28} & \alpha_c < \alpha_{tr} \\ \alpha_c^{C_1} & \alpha_c \geq \alpha_{tr} \end{cases} \tag{1732}$$

where:

α_{tr} is the transition volume fraction, which is set to 0.85 by default.

C_1 and C_2 are model constants, which are set to 2.65 and 0.8 respectively [329].

The drag coefficient C_D used in Eqn. (1729) is:

$$C_D = \left(0.63 + \frac{4.8}{\sqrt{Re_d / V_{rs}}}\right)^2 \tag{1733}$$

Symmetric Drag Model

The symmetric drag model provides the linearized coefficient as:

$$A_{cd}^D = \frac{3}{4} C_D \frac{\alpha_c \alpha_d (\alpha_c \rho_c + \alpha_d \rho_d)}{(\alpha_c l_c + \alpha_d l_d)} |\mathbf{v}_r| \tag{1734}$$

where:

l_c is the interaction length scale

l_d is the inverted topology length scale

$$C_D = \begin{cases} \frac{24}{Re_{cd}}(1 + 0.15 Re_{cd}^{0.687}) & 0 < Re_{cd} \le 1000 \\ 0.44 & Re_{cd} > 1000 \end{cases} \tag{1735}$$

with,

$$Re_{cd} = \frac{(\alpha_c \rho_c + \alpha_d \rho_d)(\alpha_c l_c + \alpha_d l_d)|\mathbf{v}_r|}{(\alpha_c \mu_c + \alpha_d \mu_d)(\alpha_c + \alpha_d)} \tag{1736}$$

Drag Force: Standard Drag Coefficient Methods

The following options are available with Newtonian continuous fluids:

- Constant
- Field Function
- Schiller-Naumann (, droplets, or bubbles)
- Hamard and Rybczynski (, droplets, or bubbles)
- Tomiyama (Bubbles)
- Bozzano-Dente (Bubbles)
- Wang Curve Fit (Air/Water Bubbles)

The following options are available with non-Newtonian continuous fluids:

- Schiller-Naumann Generalized Continuous Phase
- Schiller-Naumann Generalized Carreau-Yasuda Continuous Phase

Schiller-Naumann Drag Coefficient for Spherical Particles

The drag coefficient of spherical, rigid particles is computed based on the correlation of Schiller and Naumann [360]:

$$C_D = \begin{cases} \frac{24}{\mathrm{Re}_d}(1 + 0.15\mathrm{Re}_d^{0.687}) & 0 < \mathrm{Re}_d \leq 1000 \\ 0.44 & \mathrm{Re}_d > 1000 \end{cases} \tag{1737}$$

The dispersed phase Reynolds number, Re_d is defined as:

$$\mathrm{Re}_d = \frac{\rho_c |\mathbf{v}_r| l}{\mu_c} \tag{1738}$$

where:

μ_c is the dynamic viscosity of the continuous phase

l is interaction length scale or bubble size

This formulation is useful for spherical solid particles, droplets, and small-diameter bubbles.

Schiller-Naumann with Non-Newtonian Generalized Continuous Phase

From, [313], the drag coefficient of spherical, rigid particles in a Power Law continuous phase is:

$$C_D = \frac{24}{\mathrm{Re}_{HB}}(1 + 0.15\mathrm{Re}_{HB}^{0.687}) \tag{1739}$$

The dispersed phase Reynolds number, Re_{HB} is defined as:

$$\mathrm{Re}_{HB} = \frac{\mathrm{Re}_{PL}}{1 + \frac{7\pi}{24}\mathrm{Bi}_{HB}} \tag{1740}$$

The particle Bingham number, Bi_{HB}, is defined as:

$$\mathrm{Bi}_{HB} = \frac{\tau_0}{K}\left(\frac{d}{|\mathbf{v}_r|}\right)^n \tag{1741}$$

The Power-Law Reynolds number is:

$$\mathrm{Re}_{PL} = \frac{\rho_c |\mathbf{v}_r|^{2-n} d^n}{K} \tag{1742}$$

where:

τ_0 is the yield stress

K is the consistency factor

n is the power-law component of the model

$|\mathbf{v}_r|$ is the relative slip velocity between phases

d is the particle diameter

This formulation is useful for spherical solid particles, droplets, and small-diameter bubbles in a Power Law continuous phase.

Schiller-Naumann with Non-Newtonian Generalized Carreau-Yasuda Continuous Phase

The drag coefficient of spherical, rigid particles in a Carreau-Yasuda continuous phase is:

$$C_D = \frac{24}{\mathrm{Re}_{HB}}(1 + 0.15\mathrm{Re}_0^{0.687})(1 + 0.65(n-1)\lambda^{0.20}) \tag{1743}$$

Where n and λ are parameters in the Carreau model [321].

Re_0 is the zero-shear particle Reynolds number that uses the zero-shear viscosity of the fluid:

$$\mathrm{Re}_0 = \frac{\rho_c |\mathbf{v}_r| d}{\mu_0} \tag{1744}$$

where:

μ_0 is the zero shear viscosity of the continuous phase

d is the interaction length scale or bubble size

ρ_c is the density of the continuous phase

$|\mathbf{v}_r|$ is the magnitude of the relative velocity between the phases

This formulation is useful for spherical solid particles, droplets, and small-diameter bubbles.

Hamard and Rybczynski Drag Coefficient for Droplets

The Hamard and Rybczynski drag model is used for viscous Newtonian fluid droplets dispersed in a second immiscible viscous Newtonian fluid. From [321], the drag coefficient is:

$$C_D = \frac{24}{\mathrm{Re}} Y \tag{1745}$$

where:

Re is the particle Reynolds number

where:

$$Y = \frac{2 + 3X_E}{3 + 3X_E} \tag{1746}$$

$$X_E = \frac{\mu_d}{\mu_c} \tag{1747}$$

where:

μ_c is the dynamic viscosity of the continuous phase

μ_d is the dynamic viscosity of the dispersed phase

For gas bubbles, $X_E \ll 1$ and $Y \rightarrow 2/3$. For solid particles $Y \rightarrow 1$.

Note: The Hamard and Rybczynski drag model is not valid for high Reynolds numbers, where the Schiller-Naumann model is expected to be more accurate.

Tomiyama Drag Coefficient for Bubbles

The Tomiyama correlation is suited to systems for a wide range of conditions, as shown by the range of Eotvos (Eo), Morton (Mo), and Reynolds (Re) numbers for which it is valid:

$$\begin{aligned} 10^{-2} < Eo < 10^{3} \\ 10^{-14} < Mo < 10^{7} \\ 10^{-3} < Re < 10^{5} \end{aligned} \tag{1748}$$

The drag coefficient that is evaluated for a single bubble, which is based upon three contamination states, is from [365]:

- Pure:

$$C_D = \max\left[\min\left(\frac{16}{\mathrm{Re}}(1+0.15\mathrm{Re}^{0.687}), \frac{48}{\mathrm{Re}}\right), \frac{8\mathrm{Eo}}{3(\mathrm{Eo}+4)}\right] \tag{1749}$$

- Moderate:

$$C_D = \max\left[\min\left(\frac{24}{\mathrm{Re}}(1+0.15\mathrm{Re}^{0.687}), \frac{72}{\mathrm{Re}}\right), \frac{8\mathrm{Eo}}{3(\mathrm{Eo}+4)}\right] \tag{1750}$$

- Contaminated:

$$C_D = \max\left[\left(\frac{24}{\mathrm{Re}}(1+0.15\mathrm{Re}^{0.687})\right), \frac{8\mathrm{Eo}}{3(\mathrm{Eo}+4)}\right] \tag{1751}$$

where Re is the Reynolds number and Eo is the Eotvos number.

$$\mathrm{Re} = \frac{\rho_c |\mathbf{v}_r| l}{\mu_c} \tag{1752}$$

$$\mathrm{Eo} = \frac{|\rho_c - \rho_d| g l^2}{\sigma} \tag{1753}$$

ρ_c is the density of the continuous phase

μ_c is the dynamic viscosity of the continuous phase

ρ_d is the density of the dispersed phase

$|\mathbf{v}_r|$ is the magnitude of the relative velocity between the phases

l is interaction length scale or bubble size

g is gravitational acceleration

σ is the surface tension

Note: The Tomiyama drag coefficient was used by Tomiyama *et al.* [366] while evaluating their lift coefficients. Therefore, use the Tomiyama lift coefficient with the Tomiyama drag coefficient.

Bozzano-Dente Drag Coefficient for Bubbles

This correlation covers a wide range of bubbly flow regimes and accounts for different bubble shapes with a single expression. It has been developed for the case of a single bubble rising in a liquid column under the action of gravity. It takes both Eotvos number and Morton number into account, so can be applied to highly viscous systems such as glycerol, molten glass and some oils.

It can also be used for low pressure and high pressure water systems.

Note: For small bubbles, this correlation gives a different asymptote (Re/48) for the drag coefficient when compared to the values normally associated with contaminated or pure water systems (Re/24 or Re/16).

For this model, from Bozzano and Dente [319], the drag coefficient is:

$$C_D = f\left(\frac{a}{R_0}\right)^2 \tag{1754}$$

In which f is a friction factor and $(a/R_0)^2$ is a deformation factor. These quantities are derived as

$$f = \left(\frac{48}{\mathrm{Re}}\right)\frac{1 + 12\mathrm{Mo}^{1/3}}{1 + 36\mathrm{Mo}^{1/3}} + 0.9\frac{\mathrm{Eo}^{3/2}}{1.4(1 + 30\mathrm{Mo}^{1/6}) + \mathrm{Eo}^{3/2}} \tag{1755}$$

and

$$\left(\frac{a}{R_0}\right)^2 \cong \frac{10(1 + 1.3\mathrm{Mo}^{1/6}) + 3.1\mathrm{Eo}}{10(1 + 1.3\mathrm{Mo}^{1/6}) + \mathrm{Eo}} \tag{1756}$$

Where Re is the Reynolds number, Mo is the Morton number, and Eo is the Eotvos number.

$$\mathrm{Re} = \frac{\rho_c|\mathbf{v}_r|l}{\mu_c} \tag{1757}$$

$$\mathrm{Eo} = \frac{|\rho_c - \rho_d| g l^2}{\sigma} \tag{1758}$$

$$\mathrm{Mo} = \frac{g \mu_c^4}{\rho_c \sigma^3} \tag{1759}$$

ρ_c, μ_c are the density and dynamic viscosity of the continuous phase

ρ_d is the density of the dispersed phase

$|\mathbf{v}_r|$ is the magnitude of the relative velocity between the phases

l is interaction length scale or bubble size

g is gravitational acceleration

σ is the surface tension

Wang Drag Coefficient for Bubbles

This model is for concentrations of air bubbles in water at near atmospheric pressure. Use with care for other conditions.

Bubbles deform and depart from the spherical shape. Use the empirical correlation of Wang *et al.* [370], to obtain more realistic values. The correlation was derived using curve-fitting measurements for a single bubble rising in water:

$$C_D = \exp[a + b \ln \mathrm{Re}_d + c(\ln \mathrm{Re}_d)^2] \tag{1760}$$

Where Re_d is the dispersed phase Reynolds number.

The original coefficients are given in the table below.

$\mathbf{Re}_d$	a	b	c
$\mathbf{Re}_d \leq 1$	ln 24	–1	0
$1 < \mathbf{Re}_d \leq 450$	2.699467	––0.33581596	–0.07135617
$450 < \mathbf{Re}_d \leq 4000$	–51.77171	13.1670725	–0.8235592
$\mathbf{Re}_d > 4000$	$\ln(8/3)$	0	0

The STAR-CCM+ implementation has been modified from Wang's original method [370] so that the drag coefficient is continuous as particle Reynolds number changes. This is important to minimize the residual achievable in applications with variable particle diameter, and the modified method gives reliable convergence with normal values of relaxation factors.

The modifications are:

- The breakpoint between the first two Re regimes in the above table has been moved from Re=1 to Re=2.197115.
- The last two Re regimes in the above table have been replaced using:

$$C_D = \frac{8}{3}\frac{\mathrm{Eo}_{AW}}{\mathrm{Eo}_{AW} + 4} \tag{1761}$$

For close equivalence with the original Wang curve fit, the Eotvos number is here computed from bubble diameter d using water and air properties that are fixed at a temperature of 293K.

$$\mathrm{Eo}_{AW} = \frac{(\rho_W - \rho_A)gd^2}{\sigma_{AW}} \tag{1762}$$

ρ_W is 998.20 kg/m^3.

ρ_A is 1.2069 kg/m^3.

g is 9.81 m/s^2.

σ_{AW} is 0.072736 N/m.

Drag Force: Drag Correction Methods

The following options are available:

- Field Function
- Volume Fraction Exponent

- Richardson Zaki (Spherical particles, droplets, or bubbles)
- Lockett Kirkpatrick (Bubbles)
- Simonnet (Bubbles)

Richardson Zaki Drag Correction

This is a drag correction method for spherical particles with a wide range of applicability. According to this model, the effect of increased concentration on spherical particles is that of hindering. The terminal velocity decreases faster than would be expected purely from the reduction of buoyancy force of the two-phase mixture on the particles.

This model can also cover droplets that are small enough to retain their spherical shape and whose internal viscosity is high enough for internal circulation to be negligible. This model also applies to bubbles whose surfaces are immobile due to surface contamination, as in the case of air bubbles in regular tap water.

As originally specified, the Richardson Zaki model [356] for terminal velocity correction for spherical particles is:

$$\frac{V}{V_\infty} = \alpha_c^n \tag{1763}$$

Range	n for small tubes	n for large tubes
$\mathrm{Re}_\infty < 0.2$	4.65	4.65
$0.2 < \mathrm{Re}_\infty < 1$	$\left(4.35 + 17.5\frac{d}{D}\right)Re_\infty^{-0.03}$	$4.35Re_\infty^{-0.03}$
$1 < \mathrm{Re}_\infty < 200$	$\left(4.45 + 18\frac{d}{D}\right)Re_\infty^{-0.1}$	$4.45Re_\infty^{-0.1}$
$200 < \mathrm{Re}_\infty < 500$	$4.45Re_\infty^{-0.1}$	$4.45Re_\infty^{-0.1}$
$\mathrm{Re}_\infty > 500$	2.39	2.39

n is the Richardson Zaki exponent, D is tube diameter, and d is particle diameter.

The two figures below illustrate the wide range of applicability of the Richardson Zaki method. The first shows the terminal velocity correction factor that is computed in a Stokesian particle settling test ($\mathrm{Re}_\infty < 0.2$) against the data of Fricke and Thompson [330]. A suspension of 44-micron

resin particles in a glycerine-water mixture settles out at a rate that particle concentration affects. The rate is well-predicted either by the Richardson Zaki method, or by the Volume Fraction Exponent method using exponent n_D=-8.7.



The second figure shows the terminal velocity of 3-5 mm air bubbles rising in water against the data of Simonnet [361] under conditions $\mathrm{Re}_{\infty} > 500$. Even though the bubbles are ellipsoidal in this size range, agreement is good.



The original Richardson Zaki model is adapted for computation by:

- Omitting the small tube term as this would generate an unacceptable discontinuity between Reynolds number ranges
- Using min() and max() functions to make sure exact continuity between Reynolds number ranges
- Allowing you to update or adapt the original calibration constants.

- Preventing the drag correction reaching implausibly high or low levels as the continuous phase volume fraction vanishes. Set a maximum packing $\alpha_{d,\,max}$ for the sum of the dispersed phase volume fractions, so that the strength of the drag correction does not change further beyond this concentration.

$$f_V = max(\alpha_c, 1 - \alpha_{d,\,max}) \tag{1764}$$

$$n = max\left\{ \begin{array}{c} min\left\{ \begin{array}{c} A \\ min\left(BRe_\infty^{n_1}, CRe_\infty^{n_2}\right) \\ D \end{array} \right. \end{array} \right. \tag{1765}$$

$$Re_\infty = \frac{Re}{f_V} \tag{1766}$$

$$f_D = \begin{cases} \dfrac{\alpha_c}{f_V^2} & \text{Spherical Particle Interfacial area} \\ \dfrac{1}{f_V^2} & \text{Symmetric Interfacial Area} \end{cases} \tag{1767}$$

The single-particle Reynolds number Re_∞ is applied both to the single-particle drag law as well as to the Richardson Zaki correction itself. This is essential to reproduce terminal velocity correctly in regimes where drag coefficient is a strong function of Reynolds number.

Note: This implementation gives identical results for drag coefficient and terminal velocity regardless of which of the two interfacial area models is used, at least up to the maximum packing point.

The default calibration is applicable for spherical particles, but you can modify it:

A	B	C	D	n_1	n_2	$\alpha_{d,\,max}$
4.65	4.35	4.45	2.39	-0.03	-0.1	0.66

Lockett Kirkpatrick Drag Correction

This is a drag correction method for small ellipsoidal bubbles. It applies to air/water systems with bubble size of order 5 mm (Eotvos number around 3.3). The effect of increasing concentration is hindering.

Lockett and Kirkpatrick [349] investigated the effect of bubble concentration on the terminal velocity of 5-mm air bubbles in water. They achieved stable bubbly flow of up to 66% void fraction by using a down flow. They applied the Richardson Zaki model to spherical particles to terminal velocity data for 5-mm air bubbles in water. The model assumed that both the drag coefficient and drag correction were in the Reynolds number independent range. For this range, where $\mathrm{Re} > 500$, the Richardson Zaki exponent, n, is 2.39. They found that a better fit to the data could be achieved by applying a further empirical factor to take bubble deformation into account.

In a figure further below, the Lockett Kirkpatrick method is compared with the Simonnet *et al.* [361] method for predicting Simonnet's terminal velocity data for 7-10 mm bubbles. The agreement is good up to 15% void fraction, after which Simonnet found a swarming behavior not observed in the Lockett Kirkpatrick data from 5-mm bubbles.

The original Lockett Kirkpatrick model is:

$$\frac{V}{V_\infty} = \alpha_c^{2.39}[1 + 2.55(1 - \alpha_c)^3] \tag{1768}$$

This Lockett Kirkpatrick model is adapted for computation by limiting the correction for extreme values of α_c through a maximum packing criterion $\alpha_{d,\,max}$. $\alpha_{d,\,max}$ has a default value of 0.66.

$$\tilde{\alpha}_c = max(\alpha_c, 1 - \alpha_{d,\,max}) \tag{1769}$$

$$f_V = (\tilde{\alpha}_c)^{2.39}\left[1 + 2.55(1 - \tilde{\alpha}_c)^3\right] \tag{1770}$$

$$f_D = \begin{cases} \dfrac{\alpha}{f_V^2} & \text{Spherical Particle Interfacial area} \\[2ex] \dfrac{1}{f_V^2} & \text{Symmetric Interfacial Area} \end{cases} \tag{1771}$$

Note: This method gives identical results for drag coefficient and terminal velocity regardless of the interfacial area model that is used, at least up to the maximum packing point.

Simonnet Drag Correction

This is a drag correction for air / water bubbles in the 7-10 mm size range (Eotvos number 6.6-13.4) and 0-30% void fraction range. Effects are hindering in the range 0-15% and swarming in the range 15-30%.

Simonnet *et al.* [361] measured terminal velocities, taking great care to work with uniform size distributions. Their recommended model for drag correction factor in the 7-10 mm range is:

$$\frac{C_D}{C_{D\infty}} = \alpha_c \left[\alpha_c^{25} + \left(4.8 \frac{1-\alpha_c}{\alpha_c} \right)^{25} \right]^{-2/25} \tag{1772}$$

However, they also found that for smaller bubbles (2-6 mm) only hindering behavior occurs.

Terminal velocity computed using Simonnet's hindering/swarming method and Lockett and Kirkpatrick's [349] hindering-only methods are compared below against Simonnet's own data for 7-10 mm bubbles in air-water [361]. Both methods are a good fit to the data up to 15% void fraction when swarming behavior begins.



The Simonnet model is adapted for computation by making the correction go to zero smoothly and linearly as the continuous phase volume fraction vanishes, so that it can be used safely beyond the void fraction range of Simonnet's data.

$$f_D = sf_{D1} + (1-s)f_{D2} \tag{1773}$$

$$f_{D1} = \alpha_c \left[\alpha_c^{25} + \left(4.8 \frac{1 - \alpha_c}{\alpha_c} \right)^{25} \right]^{-2/25} \tag{1774}$$

$$f_{D2} = 0.25 \alpha_c \tag{1775}$$

$$s = 0.5 \left\{ 1 + \tanh \left(\frac{\alpha_c - 0.7058}{0.05} \right) \right\} \tag{1776}$$

Spherical Particle Interfacial Area Density and Symmetric Interfacial Area Density models give the same results for drag coefficient and terminal velocity when the Simonnet drag correction model is used in its range of applicability of 0-30% void fraction.

Volume Fraction Exponent Drag Correction

This drag correction model assumes that the drag correction factor is of the form:

$$f_D = \alpha_c^{n_D} \tag{1777}$$

where n_D is some constant power.

This method complements the other drag correction methods by covering several regimes of practical interest, such as

- Small spherical particles settling under Stokesian drag condition.
- Large spherical particles under turbulent wake drag conditions.

- Spherical cap bubbles swarming under a churn-turbulent flow regime.

Regime	**Reference**	**n**	n_D **(Spherical particle Interfacial Area)**	n_D **(Symmetric Interfacial Area)**
small spherical particles $\mathrm{Re}_\infty < 0.2$	[356]	+4.65	-8.3	-9.3
large spherical particles $\mathrm{Re}_\infty > 500$	[356]	+2.39	-3.78	-4.78
spherical cap bubbles, churn-turbulent conditions	[338]	-0.5	+2	+1

The first of these three regimes has already been shown in the figure in the Richardson Zaki section on the dilute particle settling under Stokesian conditions. The settling velocity is a strong function of concentration. Using a drag correction exponent of $n_D = -8.7$ is equivalent to using the velocity correction exponent recommended by Fricke and Thompson, $n = 4.85$, which provides a good fit to their own data [330].

The drag correction method in the churn-turbulent regime is taken from Ishii and Zuber [338] and is relevant for large spherical cap bubbles in bubble columns. In such columns, the bubble size is controlled by coalescence and breakup, and is unrelated to the initial size of the bubbles that are introduced at the inlet.

Ishii and Zuber suggested that bubbles in this regime are characterized by a Weber number of 8 and a drag coefficient of 8/3. Both of these numbers are defined in terms of their drift flux model velocity. This leads to a model predicting pure swarming behavior. Bubble equivalent diameter, terminal velocity, and drag coefficient vary as continuous phase volume fraction raised to the power of -1/2, -3/4 and +2 respectively.

Ishii and Zuber found that this model under-predicts Akita and Yoshida data for superficial velocity for air bubbles in dilute sodium sulfite solution that is presented in Figure 16 of Ishii and Zuber [338]. However it does show the correct trend of superficial velocity against bubble concentration and is a reasonable approximation for the data from the bubble column with the largest diameter of 60 cm.

In the verification calculation reproduced below, the equivalent bubble diameter varies from 11 mm to 13 mm as the void fraction increases from 1% to 30%. The terminal velocity increases from 23 cm/sec to 30 cm/sec, as expected from the relevant exponents.



The constant exponent method is adapted for computation by limiting the correction for extreme values of α_c. This is done through a maximum packing criterion $\alpha_{d,max}$, which has a default value of 0.66:

$$\tilde{\alpha}_c = max(\alpha_c, 1 - \alpha_{d,\,max}) \tag{1778}$$

$$f_D = (\tilde{\alpha}_c)^{n_D} \tag{1779}$$

$$f_V = \begin{cases} \sqrt{\dfrac{\alpha}{f_D}} & \text{Spherical Particle Interfacial Area} \\ \sqrt{\dfrac{1}{f_D}} & \text{Symmetric Interfacial Area} \end{cases} \tag{1780}$$

This velocity correction factor f_V is used to compute the single-particle terminal velocity correctly where the drag coefficient depends on Reynolds number. However, this Reynolds number adjustment is not relevant in cases where the single particle drag coefficient is not a function of Reynolds number. For example, in the Ishii Zuber model for large bubbles in the churn turbulent regime.

Turbulence Response Model

The turbulence response model [339] treats the turbulence of the dispersed phase as an empirical relation to the solution of the turbulence of the continuous phase. The turbulence response coefficient, C_t, defines the correlation between the phases as the ratio of the dispersed phase velocity fluctuation to that of the continuous phase.

$$C_t = \frac{\mu_d'}{\mu_c'} \tag{1781}$$

where:

μ_d' is the dispersed phase velocity fluctuation.

μ_c' is the continuous phase velocity fluctuation.

The dispersed phase turbulent eddy viscosity is:

$$\mu_d^t = \frac{\rho_d}{\rho_c} C_t^2 \mu_c^t \tag{1782}$$

where:

ρ_d is the dispersed phase density.

ρ_c is the continuous phase density.

If the continuous phase uses the Standard K-Epsilon Model, then the dispersed-phase turbulent kinetic energy, k_d, is:

$$k_d = C_t^2 k_c \tag{1783}$$

The K-Epsilon definition of turbulent viscosity for a single phase is:

$$\mu^t = \rho C_\mu \frac{k^2}{\varepsilon} \tag{1784}$$

where:

C_μ is a model coefficient in the Standard K-Epsilon Model.

Combine Eqn. (1782) and Eqn. (1784) to get the turbulent dissipation rate, ε:

$$\varepsilon_d = C_t^2 \varepsilon_c \tag{1785}$$

Turbulent Dispersion Force

The effect of turbulence in redistribution of non-uniformities in phase concentration is modeled by an additional turbulent dispersion force in the phase momentum equations. This term arises naturally when Reynolds averaging is applied to the instantaneous drag force. For a simple derivation and further references, see Contribution of Drag to Turbulent Dispersion. With suitable closures, the term has the following form:

$$\begin{aligned} \boldsymbol{F}_{cd}^{TD} &= A_{cd}^{D}\ \mathbf{v}_{TD} \\ \mathbf{v}_{TD} &= \boldsymbol{D}_{cd}^{TD} \cdot \left\{ \frac{\nabla \alpha_d}{\alpha_d} - \frac{\nabla \alpha_c}{\alpha_c} \right\} \end{aligned} \tag{1786}$$

where:

- $\boldsymbol{F}_{cd}^{TD}$ is the force per volume applied to the continuous phase c momentum equation due to dispersed phase d.
- $\mathbf{v}_{TD}$ is a relative drift velocity due to the use of volume-fraction weighted definitions of phase velocity.

  Alternatively, $\mathbf{v}_{TD}$ can be written:

$$\mathbf{v}_{TD} = \boldsymbol{D}_{cd}^{TD} \cdot \{\nabla \ln(\alpha_d) - \nabla \ln(\alpha_c)\} \tag{1787}$$

The logarithmic form is easier to discretize on a collocated grid. Use it for simulating flows where separation or boiling/condensation processes can produce volume fractions differing by many orders of magnitude in adjacent cells.

The tensor diffusivity coefficient $\boldsymbol{D}_{cd}^{TD}$ is usually approximated isotropically from the continuous phase turbulent diffusivity:

$$\boldsymbol{D}_{cd}^{TD} = C_0 \frac{\nu_c^t}{\sigma_\alpha} \boldsymbol{I} \tag{1788}$$

where:

- ν_c^t is the continuous phase turbulent kinematic viscosity.
- σ_α is the turbulent Prandtl number for volume fraction. This is discussed in more detail in the following section.
- The constant C_0 is not part of the model derivation, but has a default value 1 and you can vary it to test the sensitivity of the solution to this term.

When a Reynolds stress tensor $\boldsymbol{R}_c$ is modeled for the continuous phase, then an anisotropic approximation is also available:

$$\boldsymbol{D}_{cd}^{TD} = C_0 \frac{3}{2} \frac{C_\mu}{\sigma_\alpha} \frac{k_c}{\varepsilon_c} \boldsymbol{R}_c \tag{1789}$$

Turbulent Dispersion Coefficient

Instead of choosing a model for the tensor diffusivity, $\boldsymbol{D}_{cd}^{TD}$, the essential difference between alternative dispersion coefficient models can be encapsulated as an effective Turbulent Prandtl Number, σ_α. This number signifies the ratio of momentum diffusivity over the volume fraction diffusivity due to continuous phase velocity fluctuations.

There are four modeling options for the Turbulent Prandtl number:

- Constant, with default value 1 as for a passive scalar
- User-specified expression
- Tchen turbulent dispersion coefficient, recommended for small particles or bubbles
- Inertial turbulent dispersion coefficient, recommended for larger particles or bubbles

Tchen Turbulent Dispersion Coefficient

Thai-Van *et al.* [363] describe a model for effective turbulent diffusivity of particles. This is based on Simonin's extension of Tchen analysis of particle behavior in steady homogeneous turbulence.

Expressed as an effective Turbulent Prandtl number for the Turbulent Dispersion Force term, this model becomes:

$$\sigma_\alpha = \sigma_0 \sqrt{1 + C_\beta \xi^2} \, \frac{1 + \eta}{b + \eta} \tag{1790}$$

where:

- σ_0 is the unmodified Turbulent Prandtl number, assumed to be unity to represent basic passive diffusivity.
- C_β is a calibration coefficient for the crossing-trajectories effect.
- ξ is the particle slip velocity scaled by turbulent fluctuation velocity.
- η is the particle-eddy interaction time, scaled by particle relaxation time.
- b is the ratio of the coefficients of the continuous/disperse acceleration terms in the equation of motion for a particle.

For further details, see Tchen Closures.

The general effect of high crossing speed ($\xi \gg 1$) in this model is to reduce correlation and dispersion through an increased effective Turbulent Prandtl number.

For bubbly flows (phase densities $\rho_d \ll \rho_c$) with suitable interaction times, the effect of virtual mass in this model could be to increase dispersion by up to a factor of 1.5.

Inertial Turbulent Dispersion Coefficient

If the particles are larger than the inner turbulent scale, they hold a substantial share of the fluctuating kinetic energy and particle-particle interaction is an important energy redistribution mechanism. Gas-particle flows with high particle concentrations are best modeled as granular flows, but for more dilute particle concentrations (α_d between about 1% to 10%), it is often possible to obtain useful results with an inertial correction to the turbulent dispersion.

The effective Turbulent Prandtl number for inertial dispersion of heavy particles is modeled as:

$$\sigma_\alpha = \frac{\sigma_0}{\left(1 + \frac{\tau_R}{\tau_T}\right)\left(1 + \frac{\rho_c C_{VM}}{\rho_d + \rho_c C_{VM}}\right)} \tag{1791}$$

where:

- σ_0 is the unmodified Turbulent Prandtl, assumed to be unity to represent basic passive diffusivity.
- τ_R is the particle relaxation time:

$$\tau_R = \frac{\rho_d \alpha_d}{A_{cd}^D}\left(1 + \frac{\rho_c}{\rho_d} C_{VM}\right) \tag{1792}$$

- τ_T is an eddy-turnover time:

$$\tau_T = C_\mu \frac{k_c}{\varepsilon_c} \tag{1793}$$

- C_{VM} is a virtual mass coefficient for particle response, which is taken from the Virtual Mass Force model if that has been activated for simulation of an accelerating mean flow; otherwise it is assumed to have a constant value of 0.5.

The first factor in the denominator of Eqn. (1791) is accounting for the increase of the longer particle relaxation times ($\tau_R \gg \tau_T$), while the second one originates from the virtual mass force (see Contribution of Virtual Mass to Turbulent Dispersion for details). The general effect of the model is to increase dispersion through a reduced effective Turbulent Prandtl number.

Virtual Mass Force

The inertia of surrounding fluid influences the acceleration of a particle that is submerged in the flow. Inviscid flow theory allows this effect to be represented as a “virtual mass” or “added mass” equal to a constant multiplied by the mass of fluid that the particle displaces [344].

Including this “virtual mass force” can make accelerating flows more realistic. For example, the virtual mass term can influence the trajectory of bubbles in a steady but swirling flow.

Because it narrows the range of response timescales across different phases, the virtual mass term can also help non-accelerating flows to converge, by making them less sensitive to momentum or pressure relaxation factors.

A two-phase formulation for the virtual mass force is from Auton *et al.* [314]. This formulation can also be used in the multiphase context where there is one continuous phase and several dilute, dispersed phases. For non-dilute systems, it can still give a useful first approximation, especially if the virtual mass coefficient includes a correction for the concentration of particles.

The virtual mass force acting on phase i due to acceleration relative to phase j is:

$$\mathbf{F}_{ij}^{VM} = C_{ij}^{VM}\rho_c\alpha_d[\mathbf{a}_j - \mathbf{a}_i] \tag{1794}$$

where:

c is the continuous phase in the phase interaction ij.

d is the dispersed phase in the phase interaction ij.

C_{ij}^{VM} is the virtual mass coefficient for interaction ij.

Vector $\mathbf{a}_k$ is the acceleration of phase k.

The acceleration term for virtual mass is based on the rate of change of velocity of **constant mass particles**, which in the stationary coordinate frame becomes the material derivative:

$$\mathbf{a}_k = \left(\frac{Du}{Dt}\right)_k \tag{1795}$$

However, for applications where the particle mass can increase or decrease due to interphase mass transfer, $\mathbf{a}_k$ must be generalized. This is accomplished by starting with the momentum balance for the phase that is contained within a **constant control volume**. The momentum fluxes across phase-to-phase interfaces can be identified as similar terms to the momentum fluxes across intersections between the phase and the outer surfaces of the control volume. This leads to the following definition of the total inertial term for phase k in reaction to all other applied forces:

$$\rho_k\alpha_k\mathbf{a}_k = \frac{\partial}{\partial t}(\rho_k\alpha_k\mathbf{u}_k) + \nabla\cdot(\rho_k\alpha_k\mathbf{u}_k\mathbf{u}_k) - \sum_l m_{kl}\mathbf{u}_{kl} \tag{1796}$$

where:

m_{kl} is the interphase mass transfer rate per volume from phase l to phase k.

$\mathbf{u}_{kl}$ is the interphase value of velocity.

The following closure satisfies conservation of momentum in the sense that the momentum flux leaving one phase balances the momentum flux entering the other phase:

$$\mathbf{u}_{kl} = \begin{cases} \mathbf{u}_l & m_{kl} > 0 \\ \mathbf{u}_k & m_{kl} < 0 \end{cases} \tag{1797}$$

The full acceleration term, including momentum fluxes due to interphase mass transfer, then becomes:

$$\mathbf{a}_k = \left(\frac{Du}{Dt}\right)_k + \frac{1}{\rho_k \alpha_k} \sum_l m_{kl}(\mathbf{u}_k - \mathbf{u}_{kl}) \tag{1798}$$

Virtual Mass Terms in Rotating Frame

The Virtual Mass Force term that is defined above is modified as follows when used in multiple reference frame (MRF) applications. This modification is based on an absolute velocity formulation, which solves directly for velocities in the inertial frame, I, using momentum equations that are derived in the rotating frame, R.

$$\mathbf{F}_{ij}^{VM} = C_{ij}^{VM} \rho_c \alpha_d [\mathbf{a}_j^{IR} - \mathbf{a}_i^{IR}] + C_{ij}^{VM} \rho_c \alpha_d \boldsymbol{\omega} \times [\mathbf{u}_j^I - \mathbf{u}_i^I] \tag{1799}$$

where $\boldsymbol{\omega}$ is the angular velocity of the rotating frame, and the *phase-k* acceleration in the rotating frame is defined in terms of the absolute velocity by:

$$\mathbf{a}_k^{IR} = \left[\frac{D}{Dt}\mathbf{u}_k^I\right]_k^R + \frac{1}{\rho_k \alpha_k} \sum_l m_{kl}(\mathbf{u}_k^I - \mathbf{u}_{kl}^I) \tag{1800}$$

$$\left[\frac{D}{Dt}\phi\right]_k^R = \frac{\partial \phi}{\partial t} + \mathbf{u}_k^R \cdot \nabla \phi \tag{1801}$$

Virtual Mass Coefficient Limit

In the pursuit of treating the virtual mass force as an objective interphase force, Cook and Harlow [323] formulated an alternative virtual mass term by collapsing a three-field (dispersed, entrained continuous and bulk continuous phases) model back into two-fields.

The Cook and Harlow formulation becomes singular when the virtual mass associated with the dispersed phase exceeds the total mass of the continuous phase. This is not necessarily a physical constraint and it does not apply to other virtual mass formulations, such as those based on Auton *et al.* [314]. Nevertheless the Cook and Harlow concept of an entrained fluid fraction and a free stream fluid fraction can be used as a starting point when considering the maximum particle concentrations for which the virtual mass force is a relevant correction to the multiphase momentum equations. It is assumed that there is one continuous phase and that the same model for virtual mass coefficient C^{VM} is used for all dispersed phases. The fraction of "free stream" fluid, that is the fraction that is not entrained on the dispersed phases as "added mass" is:

$$\gamma = \frac{\alpha_c - C^{VM}\sum \alpha_d}{\alpha_c} \tag{1802}$$

Then an upper bound on the virtual mass coefficient is provided by a plausible assumption that the fraction of free fluid γ is not less than some minimum γ_{min}:

$$C^{VM}_{max}(\alpha_c, \gamma_{min}) = \frac{\alpha_c}{1-\alpha_c}(1-\gamma_{min}) \tag{1803}$$

Choosing $\gamma_{min} = 0$ gives a useful operating range of continuous phase volume fraction $1 \leq \alpha_c \leq 0.5$ for which a typical virtual mass coefficient of $C^{VM} = 0.5$ can be applied before the constraint starts to limit the coefficient. Beyond this range, the constraint makes the virtual mass force go to zero as the continuous phase volume fraction goes to zero, which helps maintain stability.

The virtual mass coefficient is based on the kinetic energy content in the fluid field, and it is not always helpful to consider it only in terms of entrained mass. Strong virtual mass effects, exceeding the Cook and Harlow limit $\gamma_{min} = 0$, can indeed be realistic in close-packed particles [375] or in slug flows [337]. The physical interpretation of such high virtual mass coefficients is that a relative acceleration difference of the particle can cause a fluid backflow through confined gaps between particles, or between particles and walls, at local peak velocities many times higher than the mean

relative slip velocity between phases. Under these conditions, relax the constraint by using negative values of γ_{min} to see how far accuracy can be improved without the solution becoming unstable.

Virtual Mass Coefficient Methods

There are four options for Virtual Mass Coefficient methods:

- Constant
- Field Function
- Spherical Particle
- Zuber

In a simulation, the selected virtual mass coefficient can be plotted as a field.

The default option is the Spherical Particle method which is based on a standard coefficient of 0.5 for spherical particles. For stability, the coefficient is reduced at high dispersed phase concentrations.

The Zuber option is suitable for improved accuracy in modeling the effect of increasing particle concentration, or multiple dispersed phases, in accelerating flows.

The virtual mass coefficients that are produced by these two methods are compared in the figure below, each using their own default value for free stream fraction γ_{min} as a robustness limit.



Spherical Particle Virtual Mass Coefficient

From inviscid flow theory [344], the virtual mass coefficient for a spherical particle accelerating in an unbounded three-dimensional fluid is:

$$C_{sphere}^{VM} = 0.5 \tag{1804}$$

This function is limited for stable solutions across a wide range of volume fraction by specifying a minimum free stream fraction γ_{min} with a default value of zero.

$$C_{ij}^{VM} = [(C_{sphere}^{VM})^{-2} + C_{max}^{VM}(\alpha_c, \gamma_{min})^{-2}]^{-1/2} \tag{1805}$$

Zuber Virtual Mass Coefficient

Lamb [344] provided an inviscid flow solution for a case of an inner sphere accelerating across the center of an outer stationary sphere. Zuber [375] used this solution to estimate the virtual mass coefficient for an infinite array of particles. The cube of the sphere diameter ratio represents the volume fraction.

The result in terms of the volume fraction of a single dispersed phase is:

$$C_{Zuber}^{VM} = 0.5\frac{1 + 2\alpha_d}{1 - \alpha_d} \tag{1806}$$

This function is adapted for multiphase computation by replacing α_d by the sum of dispersed phase volume fractions:

$$\sum \alpha_d = 1 - \alpha_c \tag{1807}$$

and then applying a limiter for robustness as $\alpha_c \rightarrow 0$:

$$C_{Zuber,ij}^{VM} = 0.5\frac{3 - 2\alpha_c}{\alpha_c} \tag{1808}$$

$$C_{ij}^{VM} = [C_{Zuber,ij}^{VM}{}^{-2} + C_{max}^{VM}(\alpha_c, \gamma_{min})^{-2}]^{-(1/2)} \tag{1809}$$

The default value of the notional "free stream" fraction is $\gamma_{min} = -10$. This allows the Zuber virtual mass coefficient to represent the effect of total particle concentration on the virtual mass force in the useful range $0 < \sum \alpha_d < 0.66$ before limiting is applied.

Lift Force

Lift Force Model

When the continuous phase flow field is non-uniform or swirling, particles experience a lift force perpendicular to the relative velocity. This force, $\mathbf{F}_L$, is derived by Auton *et al.* [314] as:

$$\mathbf{F}_L = C_{L,\,effective}\alpha_d\rho_c[\mathbf{v}_r \times (\nabla \times \mathbf{v}_c)] \tag{1810}$$

where $C_{L,\,effective}$ is calculated from the Lift Coefficient (C_L) and the Lift-Correction (f_l) as:

$f_l C_L$ for Spherical Particle Interaction Area Density model

$f_l \alpha_c C_L$ for Symmetric Interaction Area Density model

Lift Coefficient Methods

You can set the lift coefficient, C_L, but it is set to $C_L = 0.25$ by default, following Lance *et al.* [345]. Other models that take account of bubble size and distortion can be found in the literature. You can implement these models through field functions for the lift coefficient.

Lift Coefficient: Tomiyama

The Tomiyama lift coefficient has been evaluated based on experimental trajectories of single bubbles in a high-viscosity system. However, it is worth noting that Tomiyama *et al.* [366] found that the Tomiyama lift coefficient proposed also yields similar values to experimental data for a small bubble in a low viscosity, air-water system.

The lift coefficient that is proposed by Tomiyama *et al.* [366] is:

$$C_L = \begin{cases} 0.288\tanh(0.121\max[\mathrm{Re}, 7.374]) & \mathrm{Eo}_d < 4 \\ 0.00105\mathrm{Eo}_d^3 - 0.0159\mathrm{Eo}_d^2 - 0.0204\mathrm{Eo}_d + 0.474 & 4 \le \mathrm{Eo}_d \le 10 \\ -0.27 & 10 < \mathrm{Eo}_d \end{cases} \tag{1811}$$

Where Re is the Reynolds number (see Eqn. (1752)) and Eo_d is the modified Eotvos number that is based on the maximum horizontal dimension of a bubble.

The empirical relations between Eo_d and Eo is:

$$\mathrm{Eo}_d = \mathrm{Eo} \times \mathrm{E}^{-2/3} \tag{1812}$$

Where E is the empirical correlation for the bubble aspect ratio of spheroidal bubbles in a fully contaminated system and is:

$$\mathrm{E} = \frac{1}{1 + 0.163\mathrm{Eo}^{0.757}} \tag{1813}$$

Note: The Tomiyama [365] drag coefficient was used by Tomiyama *et al.* [366] while evaluating their lift coefficients. Therefore, use the Tomiyama lift coefficient with the Tomiyama drag coefficient.

Lift Correction Methods

A strong correlation between the lift correction and drag correction has been observed and, in light of limited literature available for lift correction, a perfect correlation of lift correction with drag correction is approximated.

The following options are available:

- Drag Correlated Lift Coefficient Correction method

  This method uses the Drag Coefficient Correction as an approximation for the Lift Coefficient Correction.

  This approximation is useful for cases when swarming is modeled.
- Field Function

Wall Lubrication Force

Antal *et al.* [311] proposed that bubbles rising close to a wall experience a force from the fluid flowing asymmetrically around the bubble. This force prevents the bubbles from touching the wall. The force per unit volume, $\mathbf{F}_{WL}$, experienced by a dispersed phase at a distance, y_w, from a wall is:

$$\mathbf{F}_{WL} = -C_{WL} y_w \alpha_d \rho_c \frac{\left|\mathbf{v}_{rll}\right|^2}{d} \mathbf{n} \tag{1814}$$

where:

α_d is the bubble volume fraction.

d is the diameter of the bubble.

C_{WL} is a function of inverse length. Its value decreases rapidly with increasing distance, y_w.

$\mathbf{n}$ is the outward facing unit normal at the nearest point on the wall, so that the force is pointing inwards to prevent the bubbles from contacting the wall.

The velocity scale in this model is based on the slip velocity component that is parallel to the wall:

$$\mathbf{v}_{rll} = \mathbf{v}_r - (\mathbf{v}_r \cdot \mathbf{n})\mathbf{n} \tag{1815}$$

The Antal model is implemented in STAR-CCM+ using two calibration coefficients, C_{w1} and C_{w2}:

$$C_{WL}(y_w) = max\left\{C_{w1} + \left(\frac{C_{w2}}{y_w}\right)d, 0\right\} \tag{1816}$$

The coefficients are chosen such that C_{w1} becomes zero a few bubble diameters away from the wall. That is, when:

$$\frac{y_w}{d} > \frac{-C_{w2}}{C_{w1}} \tag{1817}$$

In practice, this means that the force is zero at five bubble diameters. The default value for C_{w1} is -0.01 and the default value for C_{w2} is 0.05.

STAR-CCM+ also allows you to implement alternative models for the wall lubrication force by defining field functions for C_{WL}.

Granular Pressure

The granular pressure model provides a way to estimate the stresses in the granular medium. The model is applicable to gas-solid flows and is typically used for fluidized bed applications. Fluidized bed reactors are widely used in combustion, catalytic cracking and various other chemical and metallurgical processes. The particle motion can be classified into two regimes - kinetic and frictional. In the kinetic regime, the particle motion is governed by collisions between particles. By analogy with gas kinetic theory, and after assuming distribution and collision properties for the particles, a kinetic theory of granular pressures can be derived for single particle flows [332], [350].

In the frictional regime, where the particle motion is contact-dominated, empirical formulations that are borrowed from soil mechanics are employed. It is possible to work with multiple particle phases (that is, particles of different sizes) when using the Granular Pressure Model. Schaeffer, [359] described the stresses in the frictional regime that is based

on a simplified version of the critical state theory. Johnson and Jackson, [340] also developed constitutive relations for stresses in the frictional regime.

A full stress tensor is then obtained that describes the particle-particle interaction terms in the solid phase momentum equation:

$$(F_{int})_p = \nabla \cdot \bar{S}_p \tag{1818}$$

where the granular stress tensor is defined as:

$$\bar{S}_p = \left[-P_p + \left(\xi_p - \frac{2}{3}\mu_p \right) \nabla \cdot u_p \right] I \tag{1819}$$

where:

P_p is the solid pressure.

μ_p is the effective granular viscosity.

I is the identity tensor.

The subscript, p, is used to denote the particle phase.

ξ_p is the bulk viscosity. If you activate only the frictional regime ξ_p is set to 0. If you activate only the kinetic regime, ξ_p is computed using Eqn. (1837). If you activate both regimes, ξ_p is computed based on the frictional model you select.

Kinetic Regime

The kinetic regime is characterized by particle volume fractions below the maximum packing limit, $\alpha_{p,\,max}$. In the kinetic regime, the particle motion is governed by collisions between particles. The granular temperature, θ_p, is a measure of the fluctuating kinetic energy of particles:

$$\theta_p = \frac{1}{3} \langle v'_{pi} \cdot v'_{pi} \rangle \tag{1820}$$

where:

v'_{pi} is the random fluctuating velocity of the particle phase.

$\langle\ \rangle$ denotes ensemble averaging.

The probability of particle collisions is characterized by the particle radial distribution function, g_0. The radial distribution function represents the spatial distribution of the particles, and therefore their proximity, and can be related to the dispersed-phase volume fraction. The form in [326] has been implemented:

$$g_0 = \frac{3}{5}\left[1 - \left(\frac{\alpha_p}{\alpha_{p,\,max}}\right)^{1/3}\right]^{-1} \tag{1821}$$

This function becomes undefined when α_p approaches the maximum packing limit, $\alpha_{p,\,max}$. In this limit, Eqn. (1821) is supplemented by the equation that is proposed in [368] for $\alpha_p > \alpha_{crit}$:

$$\begin{aligned} g_0 = {} & 1.08\times10^3 + 1.08\times10^6(\alpha_p - \alpha_{crit}) + \\ & 1.08\times10^9(\alpha_p - \alpha_{crit})^2 + 1.08\times10^{12}(\alpha_p - \alpha_{crit})^3 \end{aligned} \tag{1822}$$

where:

$$\alpha_{crit} = \alpha_{p,\,max} - 0.001 \tag{1823}$$

The energy that is dissipated in particle collisions is characterized by the coefficient of restitution, e. You can specify e (by default e = 0.9).

Granular Temperature Transport Model

In the kinetic theory of granular pressure, θ_p is obtained by solving the conservation equation for fluctuating energy:

$$\frac{3}{2}\left[\frac{\partial}{\partial t}(\alpha_p\rho_p\theta_p) + \nabla\cdot(\alpha_p\rho_p\theta_p\mathbf{u}_p)\right] = \tau_p{:}\nabla\mathbf{u}_p + \nabla\cdot\kappa\nabla\theta_p - 3A_D\theta_p - \gamma \tag{1824}$$

The granular diffusion coefficient, κ, comes in two forms:

- The Gidaspow form:

$$\kappa = \frac{150\rho_p d_p\sqrt{\pi\theta_p}}{384(1+e)g_0}\left[1 + \frac{6}{5}\alpha_p(1+e)g_0\right]^2 + 2\alpha_p{}^2\rho_p d_p(1+e)g_0\sqrt{\frac{\theta_p}{\pi}} \tag{1825}$$

- The Syamlal form:

$$\kappa = \frac{15\alpha_p\rho_p d_p\sqrt{\pi\theta_p}}{4(41-33\eta)}\left[1 + \frac{12}{5}\eta^2(4\eta-3)\alpha_p g_0 + \frac{16}{15\pi}(41-33\eta)\eta\alpha_p g_0\right] \tag{1826}$$

where:

- ρ_p is the particle density.
- d_p is the particle diameter. Set the interaction length scale to be equal to the particle diameter.
- $\eta = \frac{1}{2}(1+e)$

The dissipation of granular energy, γ, is given by:

$$\gamma = \frac{12(1-e^2)g_0}{d_p\sqrt{\pi}}\alpha_p\alpha_{p,\,\mathrm{sum}}\rho_p\theta_p^{\;3/2} \tag{1827}$$

In the granular temperature transport model, if the particle phase is the dispersed phase, an additional model is required, the Granular Energy Transfer model. This model contributes the term with the linearized drag coefficient A_D to the granular energy equation.

Algebraic Model for Granular Temperature

By assuming local equilibrium, where energy dissipation and production balance each other, the energy equation can be reduced to a quadratic equation which is solved directly for θ_p. The equilibrium form has been implemented and the granular temperature is computed as:

$$\sqrt{\theta_p} = \frac{-K_{1p}\alpha_p D_{pii} + \sqrt{K_{1p}^2\alpha_p^2 D_{pii}^2 + 4K_{4p}\alpha_p[K_{2p}D_{pii}^2 + 2K_{3p}(D_{pij}D_{pij}))]}}{2\alpha_p K_{4p}} \tag{1828}$$

where:

$$K_{1p} = 2(1+e)\rho_p g_0 \tag{1829}$$

$$K_{2p} = \frac{4d_p\rho_p(1+e)\alpha_p g_0}{3\sqrt{\pi}} - \frac{2}{3}K_{3p} \tag{1830}$$

$$K_{3p} = \frac{d_p\rho_p}{2}\left\{\frac{\sqrt{\pi}}{3(3-e)}[0.5(3e+1) + 0.4(1+e)(3e-1)\alpha_p g_0] + \frac{8\alpha_p g_0(1+e)}{5\sqrt{\pi}}\right\} \tag{1831}$$

$$K_{4p} = \frac{12(1-e^2)\rho_p g_0}{d_p\sqrt{\pi}} \tag{1832}$$

D_{pij} is the strain rate tensor:

$$D_{pij} = \frac{1}{2}\left(\frac{\partial u_{pi}}{\partial x_j} + \frac{\partial u_{pj}}{\partial x_i}\right) \tag{1833}$$

where:

ρ_p is the particle density.

d_p is the particle diameter. Set the interaction length scale to be equal to the particle diameter.

Expressions for the solid pressure, bulk viscosity, and effective granular viscosity are obtained as functions of the particle granular temperature, θ_p.

The minimum granular temperature in the domain is set to $1\times10^{-10}\ \mathrm{m}^2/s^2$ by default but you can change the value if necessary. To estimate the granular temperature at boundaries, either extrapolate the granular temperature at the face cell or use a user-defined value. The specified granular temperature value is used to calculate the solid pressure at the boundaries, which has a significant influence on results.

The solid pressure is composed of collisional and kinetic contributions, [369]:

$$P_p = P_p^C + P_p^K \tag{1834}$$

$$P_p^C = 2g_0\rho_p\alpha_p^2\theta_p(1+e) \tag{1835}$$

$$P_p^K = \rho_p\alpha_p\theta_p \tag{1836}$$

The bulk viscosity is given by [359]:

$$\xi_p = \frac{4}{3}\alpha_p^2\rho_p d_p g_0(1+e)\sqrt{\frac{\theta_p}{\pi}} \tag{1837}$$

The effective granular viscosity is composed of collisional and kinetic contributions:

$$\mu_p = \mu_p^C + \mu_p^K \tag{1838}$$

$$\mu_p^C = \frac{4}{5}\alpha_p^2\rho_p d_p g_0(1+e)\sqrt{\frac{\theta_p}{\pi}} \tag{1839}$$

You can select between two forms of the particle kinetic viscosity:

- The form from Gidaspow [332] is:

$$\mu_p^K = \frac{10\rho_p d_p\sqrt{\pi\theta_p}}{96(1+e)g_0}\left[1+\frac{4}{5}g_0\alpha_p(1+e)\right]^2 \tag{1840}$$

- The form from Syamlal [317] is:

$$\mu_p^K = \frac{\alpha_p\rho_p d_p\sqrt{\pi\theta_p}}{6(3-e)}\left[0.5(3e+1)+\frac{2}{5}(1+e)(3e-1)g_0\alpha_p\right] \tag{1841}$$

The Syamlal version predicts lower viscosity at particle volume fractions below 0.3 [369]. It is recommended that the corresponding linearized drag law (Gidaspow or Syamlal) is used in accordance with the version of particle kinetic viscosity selected.

Mixture Maximum Packing Fraction

Yu and Standish [374] proposed an empirical correlation to estimate a maximum packing limit for a mixture of multiple particles. There is no limit to the number of particles in the mixture for this correlation.

The cumulative particle volume fraction, $\alpha_{p,\,sum}$, for a mixture of M particles is:

$$\alpha_{p,\,sum} = \sum_{n=1}^{M} \alpha_n \tag{1842}$$

The mixture maximum packing fraction is:

$$\alpha_{p,\,max}^{m} = min\left\{ \frac{\alpha_{i,\,max}}{1 - \sum_{i=1}^{i-1}\left(1 - \frac{\alpha_{i,\,max}}{p_{ij}}\right)\left(\frac{cx_i}{X_{ij}}\right) - \sum_{i=i+1}^{M}\left(1 - \frac{\alpha_{i,\,max}}{p_{ij}}\right)\left(\frac{cx_i}{X_{ij}}\right)} \right\} \tag{1843}$$

where i =1, 2, ... , M are the particle phases in the simulation arranged in order from coarsest to finest.

$$cx_i = \frac{\alpha_i}{\alpha_{p,\,sum}} \tag{1844}$$

$$X_{ij} = \begin{cases} \dfrac{1 - r_{ij}^2}{2 - \alpha_{i,\,max}} & j < i \\ 1 - \dfrac{1 - r_{ij}^2}{2 - \alpha_{i,\,max}} & j \geq i \end{cases} \tag{1845}$$

$$r_{ij} = \begin{cases} \dfrac{d_{p,\,i}}{d_{p,\,j}} & j \leq i \\ \dfrac{d_{p,\,j}}{d_{p,\,i}} & j > i \end{cases} \tag{1846}$$

$$p_{ij} = \begin{cases} \alpha_{i,\,max} + \alpha_{i,\,max}(1 - \alpha_{i,\,max})(1 - 2.35 r_{ij} + 1.35 r_{ij}^2) & r_{ij} \leq 0.741 \\ \alpha_{i,\,max} & r_{ij} > 0.741 \end{cases} \tag{1847}$$

Frictional regime

The following models are available in the frictional regime:

- Modified Johnson
- Schaeffer

Modified Johnson Frictional Model

The Modified Johnson [340] model gets activated in regions where the cumulative particle volume fraction, $\alpha_{p,\,sum}$, exceeds the specified minimum frictional volume fraction packing limit, $\alpha^f_{p,\,min}$. The solid pressure in this regime is given by the frictional solid pressure, P^f_p:

$$P^f_p = \begin{cases} \dfrac{\alpha_p}{\alpha_{p,\,sum}} Fr \left(\dfrac{(\alpha_{p,\,sum} - \alpha^f_{p,\,min})^r}{(\alpha^m_{p,\,max} - \alpha_{p,\,sum})^s} \right) & \alpha_{p,\,sum} > \alpha^f_{p,\,min} \\ 0 & \alpha_{p,\,sum} \le \alpha^f_{p,\,min} \end{cases} \tag{1848}$$

where:

$\alpha^m_{p,\,max}$ is the mixture maximum packing fraction

Fr, r, and s are user-specified empirical constants

The effective granular viscosity in this regime, μ^f_p, is given by the frictional viscosity:

$$\mu^f_p = \begin{cases} min\left(\dfrac{P^f_p \sin\phi}{\sqrt{2\boldsymbol{D_p}\boldsymbol{D_p} + \dfrac{\theta_p}{d^2_p}}}, \mu^f_{p,\,max} \right) & \alpha_{p,\,sum} > \alpha^f_{p,\,min} \\ 0 & \alpha_{p,\,sum} \le \alpha^f_{p,\,min} \end{cases} \tag{1849}$$

where:

- ϕ is the angle of internal friction. The recommended value is 28.5° when using the Johnson model.
- $\boldsymbol{D_p}$ is the strain rate tensor for the particle phase.

The term involving granular temperature, θ_p is ignored when the kinetic regime is not activated.

In the frictional regime, the bulk viscosity ξ_p is set to 0 where the cumulative particle volume fraction exceeds the minimum frictional volume fraction $\alpha^f_{p,\,min}$, and is computed using Eqn. (1837) everywhere else.

Granular Boundary Conditions

The granular boundary conditions have been implemented in STAR-CCM+ such that the two boundary conditions, shear stress specification and granular temperature specification, can be considered independent.

For the shear-stress (of the particle phase), the following condition is enforced when the Partial-Slip option is used [340]:

$$\frac{u_{sl} \cdot \tau_p \cdot \vec{n}}{\left|u_{sl}\right|} + \frac{\sqrt{3\theta_p}\pi\phi\rho_p\alpha_p g_0\left|u_{sl}\right|}{6\alpha_{p,\,max}} = 0 \tag{1850}$$

Where:

- The first term on the LHS is the component of particle stress at the wall in the direction of the slip velocity, u_{sl}.
- $\vec{n}$ is the unit normal at the boundary.
- ϕ is the specularity coefficient. This value varies between 0 for perfectly specular collisions and 1 for perfectly diffuse collisions.

For the granular temperature specification, the following condition is enforced when the Johnson-Jackson option is used [340]:

$$\kappa\nabla\theta_p \cdot \vec{n} + \frac{\pi(1-e_w^2)\rho_p\alpha_p g_0\theta_p\sqrt{3\theta_p}}{4\alpha_{p,\,max}} - \frac{\sqrt{3\theta_p}\pi\phi\rho_p\alpha_p g_0\left|u_{sl}\right|^2}{6\alpha_{p,\,max}} = 0 \tag{1851}$$

Where:

- κ is the granular diffusion coefficient.
- e_w is the coefficient of restitution for collisions between particles and the wall.

The third term on the LHS appears only when the Partial-Slip option is used for shear stress specification.

Schaeffer Model

The Schaeffer [359] model gets activated in regions where the cumulative particle volume fraction, $\alpha_{p,\,sum}$, exceeds the specified mixture maximum packing fraction, $\alpha^m_{p,\,max}$. The solid pressure in this regime is given by the frictional solid pressure, P^f_p:

$$P_p^f = \begin{cases} \frac{\alpha_p}{\alpha_{p,\,sum}}(10^{25}(\alpha_{p,\,sum} - \alpha_{p,\,max}^m)^{10}) & \alpha_{p,\,sum} > \alpha_{p,\,max}^m \\ 0 & \alpha_{p,\,sum} \le \alpha_{p,\,max}^m \end{cases} \tag{1852}$$

The effective granular viscosity in this regime, μ_p^f, is given by the frictional viscosity:

$$\mu_p^f = \begin{cases} min\left(\frac{P_p^f \sin\phi}{\sqrt{4I_{2D}}}, \mu_{p,\,max}^f\right) & \alpha_{p,\,sum} > \alpha_{p,\,max}^m \\ 0 & \alpha_{p,\,sum} \le \alpha_{p,\,max}^m \end{cases} \tag{1853}$$

where:

ϕ is the angle of internal friction and is set to $25°$ by default.

I_{2D} is the second invariant of the deviator of the strain rate tensor:

$$I_{2D} = \frac{1}{6}[(D_{p11} - D_{p22})^2 + (D_{p22} - D_{p33})^2 + (D_{p33} - D_{p11})^2] + D_{p12}^2 + D_{p23}^2 + D_{p31}^2 \tag{1854}$$

For numerical stability, the maximum value of frictional viscosity, $\mu_{p,\,max}^f$, is limited to 1000 Pa-s.

In the frictional regime, where both regime and frictional are specified, the bulk viscosity, ξ_p, is set to 0 above the mixture minimum packing fraction, $\alpha_{p,\,max}^m$, and is computed using Eqn. (1837) everywhere else.

Phase Interactions

The following models are available for phase interactions:

- Gera-Syamlal

Gera-Syamlal Granular Drag Force Model

The formulation for the Gera-Syamlal [317], [331] drag force model is:

$$A_{ij}^D = 3(1+e)\left(\frac{\pi}{2} + c_f\frac{\pi^2}{8}\right)\frac{(d_i + d_j)^2}{2\pi(\rho_i d_i^3 + \rho_j d_j^3)}\rho_i\rho_j g_0|\mathbf{v}_r| \tag{1855}$$

where e is the coefficient of restitution for the phase interaction, c_f is the coefficient of friction for this interaction. Default value for c_f is 0.

Solid Pressure Force

A compaction, or 'solid pressure', force is provided as one model of particle-particle interactions within a solid particle flow. It is implemented as the following internal force term in the momentum equation for a solid particle phase:

$$(\mathbf{F}_{int})_i = -\frac{\alpha_i}{\sum_{j=1}^{M}\alpha_{p,j}}\left[e^{A(\alpha_{p,max}-\sum_{j=1}^{M}\alpha_{p,j})}\right]\nabla\left(\sum_{j=1}^{M}\alpha_{p,j}\right) \tag{1856}$$

where:

$\alpha_{p,max}$ is the maximum packing limit

$\alpha_{p,j}$ is the volume fraction of the jth particle phase

M is the number of particle phases

A is the model constant, set to -600.0 by default

When activated during a multiphase flow simulation, the solid pressure force comes into action only when the solid volume fraction in any part of the domain approaches $\alpha_{p,max}$. The force acts to resist the formation of unphysical solid packing fractions by increasing exponentially as the solid volume fraction approaches the maximum packing limit. This model is useful to study a multiphase flow containing solids without needing to consider the effects of solid stress due to friction and/or collision.

The maximum packing limit is set to 0.624 by default, representative of rigid spherical solid particles.

Multiphase Turbulence Modeling

This section discusses the following elements for modeling turbulent dispersed phase flow:

- Turbulent dispersion -- turbulent velocity fluctuations of the continuous phase can redistribute non-uniform concentrations of particles.

  See Contribution of Drag to Turbulent Dispersion and Contribution of Virtual Mass to Turbulent Dispersion.

- Dispersed phase turbulent diffusion -- the dispersed phase needs a turbulence model for estimating the effective diffusivities for particle momentum and temperature variations during turbulent dispersion.

  See Dispersed Phase Turbulence Models.

- Particle Induced turbulence -- particles can extract or contribute energy to the continuous phase turbulent kinetic energy.

  See Particle Induced Turbulence.

The S-Gamma model for size distribution of particles in response to mean shear or turbulence is described in its own section.

The section on Granular Pressure describes the case where particles are large, heavy, or in high concentration, so that the kinetic energy of the particles and the particle-particle interactions need their own models.

Phase Averaging

The use of phase averaging generally reduces the number of closures that are needed to model turbulent multiphase flow. In particular, using phase-averaged velocity for convection fluxes eliminates the need to include a turbulent diffusion term in phase mass conservation equations.

$$\phi = \bar{\phi} + \phi' = \langle \phi \rangle_k + \phi'' \tag{1857}$$

where:

- $\bar{\phi}$ is the Reynolds average of the instantaneous variable ϕ.
- $\langle \phi \rangle_k$ is the phase average of the instantaneous variable ϕ formed by a Reynolds average that is weighted with the volume fraction for phase k:

$$\langle \phi \rangle_k = \frac{\overline{\alpha_k \phi}}{\overline{\alpha_k}} \tag{1858}$$

  where α_k is the instantaneous velocity of phase k.

Note: Elsewhere in the User Guide, α_k is used for the Reynolds average volume fraction, $\overline{\alpha_k}$.

The following Eulerian multiphase turbulence equations are derived in terms of phase-averaged turbulence stresses, kinetic energy, and dissipation [333]:

$$\boldsymbol{R}_k = \langle \mathbf{v}_k'' \mathbf{v}_k'' \rangle_k \tag{1859}$$

$$k_k = \frac{1}{2} \langle \mathbf{v}_k'' \cdot \mathbf{v}_k'' \rangle_k \tag{1860}$$

$$\varepsilon_k = \frac{\mu_k}{\rho_k} \langle 2 S_k'' \cdot S_k'' \rangle_k \tag{1861}$$

$$S_k'' = \frac{1}{2}\left\{\nabla \mathbf{v}_k'' + (\nabla \mathbf{v}_k'')^T\right\} \tag{1862}$$

where $\mathbf{v}_k$ is the instantaneous velocity of phase k.

Note: Elsewhere in the User Guide, $\mathbf{v}_k$ is used for the phase average velocity, $\langle \mathbf{v}_k \rangle_k$.

Boussinesq Closure

Using the eddy-diffusivity or Boussinesq hypothesis, the Reynolds flux of any scalar ϕ due to turbulent velocity fluctuation $\mathbf{v}'_c$ in the continuous phase can be modeled as:

$$\overline{\phi' \mathbf{v}'_c} = \overline{\phi' \mathbf{v}_c} = \overline{\phi \mathbf{v}'_c} \approx -\frac{\nu_c^t}{\sigma_\phi} \nabla \overline{\phi} \tag{1863}$$

where σ_ϕ is the ratio of turbulent diffusivities of momentum and ϕ.

Tchen Closures

Thai-Van *et al.* [363] summarises closures for in-phase and cross-phase turbulent diffusivities that are based on the work of by Simonin and others (for example see Deutch and Simonin [325]) in extending the results of Tchen theory for small particles in steady homogeneous turbulence (for example see Hinze [336]).

These closures are used in STAR-CCM+ for the Tchen options for turbulent dispersion force, for dispersed-phase turbulence, and for particle induced turbulence. They provide models for the following turbulent diffusivities:

- Continuous-phase turbulent diffusion:

$$\nu_c^t = \frac{2}{3} q_c^2 \tau_T \tag{1864}$$

- Interphase turbulent dispersion:

$$\boldsymbol{D}_{cd} = \tau_I \frac{1}{3} q_{cd} \boldsymbol{I} \tag{1865}$$

- Dispersed phase turbulent diffusion:

$$\nu_d^t = \tau_I \frac{1}{3} q_{cd} + \frac{1}{2} \tau_R \frac{2}{3} q_d^2 \tag{1866}$$

constructed using these correlations:

- $q_c^2 = \frac{1}{2}\langle \mathbf{v}_c'' \cdot \mathbf{v}_c'' \rangle_c$
- $q_{cd} = \langle \mathbf{v}_c'' \cdot \mathbf{v}_d'' \rangle_d$
- $q_d^2 = \frac{1}{2}\langle \mathbf{v}_d'' \cdot \mathbf{v}_d'' \rangle_d$

and these time scales:

- τ_T is the eddy turnover time.
- τ_I is the particle-eddy interaction time.
- τ_R is the particle relaxation time.

According to Thai-Van *et al.*, Eqn. (1866) was originally derived for turbulent gas flows laden with heavy particles. They commented that an eddy-viscosity closure for a bubbly phase could also be derived but had a negligible effect on predictions.

If a turbulence model is used for the continuous phase, then q_c^2 is readily available as the continuous-phase turbulent kinetic energy k_c, and ν_c^t corresponds to the turbulent diffusivity provided by a Standard K-Epsilon turbulence model.

From Thai-Van *et al.*, the other correlations are estimated using:

$$q_{cd} = 2q_c^2 \left[\frac{b+\eta}{1+\eta}\right] \tag{1867}$$

$$q_d^2 = q_c^2 \left[\frac{b^2+\eta}{1+\eta}\right] \tag{1868}$$

$$\tau_I = \frac{\tau_T}{\sigma_0} \frac{1}{\sqrt{1 + C_\beta \xi^2}} \tag{1869}$$

$$\tau_T = \frac{3}{2} C_\mu \frac{k_c}{\varepsilon_c} \tag{1870}$$

$$\tau_R = \tau_D \left(1 + \frac{\rho_c}{\rho_d} C_{VM}\right) \tag{1871}$$

where:

- τ_D is a basic particle timescale. It is used here more generally to represent the inverse of the linearized drag coefficient. For more information on this coefficient, see Contribution of Drag to Turbulent Dispersion.

- b is the ratio of the coefficients of the continuous/dispersed phase acceleration terms in the particle equation of motion:

$$b = \frac{1 + C_{VM}}{\frac{\rho_d}{\rho_c} + C_{VM}} \tag{1872}$$

- C_{VM} is a virtual mass coefficient for particle response, with a default value of 0.5 if a virtual mass force model has not been specified for the mean flow.
- η is the particle-eddy interaction time, which is scaled by particle relaxation time:

$$\eta = \frac{\tau_I}{\tau_R} \tag{1873}$$

- C_β is a calibration coefficient for the cross-trajectories effect. Following Deutsch and Simonin [325], choose it as 0.45 for the direction parallel to the mean relative velocity, and as 1.8 for orthogonal directions. The default value is 1.8.
- ξ is the particle slip velocity, scaled by turbulent fluctuation velocity:

$$\xi = \frac{|\mathbf{v}_r|}{\sqrt{\frac{2}{3}k_c}} \tag{1874}$$

Although Thai-Van *et al.* include a correction for fluctuations in this slip velocity estimate, this correction is neglected in STAR-CCM+.

Contribution of Drag to Turbulent Dispersion

Under Stokesian (viscous drag) conditions, the instantaneous drag force on a solid spherical particle of fixed diameter d and density ρ_d, due to velocity relative to a fluid of viscosity μ_c, is:

$$\frac{\rho_d \alpha_d}{\tau_D}(\mathbf{v}_d - \mathbf{v}_c) \tag{1875}$$

$$\tau_D = \frac{\rho_d d^2}{18\mu_c} = \frac{d^2}{18\nu_c}\left(\frac{\rho_d}{\rho_c}\right) \tag{1876}$$

The Reynolds average of the instantaneous drag force can be written as the sum of a mean drag force term and a turbulent dispersion term:

$$\overline{\frac{\rho_d \alpha_d}{\tau_D}(\mathbf{v}_d - \mathbf{v}_c)} = \frac{\rho_d}{\tau_D}\overline{\alpha_d}(\mathbf{v}_r + \mathbf{v}_{TD}) \tag{1877}$$

$$\mathbf{v}_r = \langle \mathbf{v}_d \rangle_d - \langle \mathbf{v}_c \rangle_c \tag{1878}$$

$$\mathbf{v}_{TD} = \langle \mathbf{v}_c \rangle_c - \langle \mathbf{v}_c \rangle_d \tag{1879}$$

Using the Boussinesq closure to model the last term:

$$\mathbf{v}_{TD} = \frac{\overline{\alpha_c \mathbf{v}_c}}{\overline{\alpha_c}} - \frac{\overline{\alpha_d \mathbf{v}_c}}{\overline{\alpha_d}} = \frac{\overline{\alpha_c \mathbf{v}_c'}}{\overline{\alpha_c}} - \frac{\overline{\alpha_d \mathbf{v}_c'}}{\overline{\alpha_d}} \approx -\frac{\nu_c^t}{\sigma_\alpha}\left(\frac{\nabla \overline{\alpha_c}}{\overline{\alpha_c}} - \frac{\nabla \overline{\alpha_d}}{\overline{\alpha_d}}\right) \tag{1880}$$

For non-Stokesian drag laws, and for other interfacial area models, the result of Reynolds averaging no longer has a straightforward closure, but the above result is still used as an approximation.

For small particles, it is possible to compensate partly for the non-Stokesian drag laws. This compensation is done by including a fluctuating contribution to the mean particle relative velocity used in the drag coefficient calculation (for example see Thai-Van *et al.*). However, this practice is not followed here, since more generally it is difficult to say whether particle-eddy interaction reduces or enhances mean slip [315].

For the general case, replace the Stokesian coefficient

$$\frac{\rho_d}{\tau_D}\overline{\alpha_d} \tag{1881}$$

with the mean linearized drag coefficient A_{cd}^{D} evaluated using a mean slip velocity $\mathbf{v}_r$ and mean interfacial area a_{cd}''':

$$A_{cd}^{D} = \frac{a_{cd}'''}{8}\rho_c C_D |\mathbf{v}_r| \tag{1882}$$

to give these mean drag and turbulent dispersion forces:

$$\boldsymbol{F}_{cd}^{D} = A_{cd}^{D}\,\mathbf{v}_r \tag{1883}$$

$$\boldsymbol{F}_{cd}^{TD} = A_{cd}^{D}\,\mathbf{v}_{TD} \tag{1884}$$

This form of turbulent dispersion force and its closure has been known for some time. For example, see Deutsch and Simonin [325] or Gosman *et al.* [333]. It has also recently been reconciled with alternative approaches in Drew[327] and Burns *et al.* [320].

Contribution of Virtual Mass to Turbulent Dispersion

The instantaneous virtual mass force reads:

$$\rho_c C_{VM} \alpha_d (\mathbf{a}_d - \mathbf{a}_c) \tag{1885}$$

where $\mathbf{a}$ is the acceleration of the corresponding phase. Averaging the above equation yields:

$$\rho_c C_{VM} \overline{\alpha_d}(\langle \mathbf{a}_d \rangle_d - \langle \mathbf{a}_c \rangle_d) = \rho_c C_{VM} \overline{\alpha_d}(\langle \mathbf{a}_d \rangle_d - \langle \mathbf{a}_c \rangle_c) + \rho_c C_{VM} \overline{\alpha_d}\left(\frac{\langle \alpha_c' \mathbf{a}_c \rangle}{\overline{\alpha_c}} - \frac{\langle \alpha_d' \mathbf{a}_c \rangle}{\overline{\alpha_d}}\right) \tag{1886}$$

The first term is for the inertia of the mean flow. The second one contributes to the turbulent dispersion force. It can be modeled using the Boussinesq closure, assuming particle relaxation time τ_R as an estimate of the acceleration time scale:

$$\frac{\langle \alpha_c' \mathbf{a}_c \rangle}{\overline{\alpha_c}} - \frac{\langle \alpha_d' \mathbf{a}_c \rangle}{\overline{\alpha_d}} \approx -\frac{\nu_c^t}{\tau_R \sigma_\alpha}\left(\frac{\nabla \overline{\alpha_c}}{\overline{\alpha_c}} - \frac{\nabla \overline{\alpha_d}}{\overline{\alpha_d}}\right) \tag{1887}$$

$$\tau_R = \frac{\rho_d \overline{\alpha_d}}{A_{cd}^D}\left(1 + \frac{\rho_c}{\rho_d} C_{VM}\right) \tag{1888}$$

Combining the contributions from drag and virtual mass terms:

$$A_{cd}^D \mathbf{v}_{TD} - \rho_c C_{VM} \overline{\alpha_d} \frac{\nu_c^t}{\tau_R \sigma_\alpha}\left(\frac{\nabla \overline{\alpha_c}}{\overline{\alpha_c}} - \frac{\nabla \overline{\alpha_d}}{\overline{\alpha_d}}\right) = \left(1 + \frac{\rho_c C_{VM}}{\rho_d + \rho_c C_{VM}}\right) A_{cd}^D \mathbf{v}_{TD} \tag{1889}$$

For bubbly flows ($\rho_c \ll \rho_d$), including a contribution from the virtual mass can double the overall turbulent dispersion coefficient.

While virtual mass is always important in modeling particle response in a turbulent flow, the modeling of the mean virtual mass force is optional, as not all applications include an accelerating mean flow.

Particle Induced Turbulence

This section outlines the derivation of the source term for the continuous-phase fluctuation kinetic energy equation. Out of all possible interphase forces, the drag force is assumed to make the dominant contribution to this term.

Phase averaging of the kinetic energy of the instantaneous flow gives the following composition for the kinetic energy of the velocity fluctuations in the continuous-phase:

$$k_c = \frac{1}{2}\langle \mathbf{v}_c'' \cdot \mathbf{v}_c'' \rangle_c = \frac{1}{2}\{\langle \mathbf{v}_c \cdot \mathbf{v}_c \rangle_c - \langle \mathbf{v}_c \rangle_c \cdot \langle \mathbf{v}_c \rangle_c\} \tag{1890}$$

where:

- $\frac{1}{2}\langle \mathbf{v}_c \cdot \mathbf{v}_c \rangle_c$ is the mean kinetic energy held in the instantaneous flow. The balance equation can be derived by Reynolds averaging the scalar product of the instantaneous phase velocity $\mathbf{v}_c$ with the instantaneous phase momentum equation.
- $\frac{1}{2}\langle \mathbf{v}_c \rangle_c \cdot \langle \mathbf{v}_c \rangle_c$ is the kinetic energy that is held in the mean flow. The equation can be derived from the scalar product of the phase mean velocity $\langle \mathbf{v}_c \rangle_c$ with the Reynolds averaged phase momentum equation.

The equation for k_c can be derived by subtracting these two equations. Then an instantaneous force term, assumed proportional to slip velocity $\mathbf{v}_d - \mathbf{v}_c$, and to disperse phase volume fraction α_d,

$$\mathbf{f}_c = \frac{\rho_d \alpha_d}{\tau_D}(\mathbf{v}_d - \mathbf{v}_c) \tag{1891}$$

makes the following contribution to the equation for k_c:

$$S_{kc} = \overline{\mathbf{v}_c \cdot \mathbf{f}_c} - \langle \mathbf{v}_c \rangle \cdot \overline{\mathbf{f}_c} = \overline{\mathbf{v}_c'' \cdot \mathbf{f}_c} \tag{1892}$$

$$= \frac{\rho_d}{\tau_D}\overline{\alpha_d}\{\langle \mathbf{v}_c'' \cdot \mathbf{v}_d'' \rangle_d - \langle \mathbf{v}_c'' \cdot \mathbf{v}_c'' \rangle_d + \langle \mathbf{v}_c'' \rangle_d \cdot [\langle \mathbf{v}_d \rangle_d - \langle \mathbf{v}_c \rangle_c]\} \tag{1893}$$

$$= A_{cd}^{D}\{\langle \mathbf{v}_c'' \cdot \mathbf{v}_d'' \rangle_d - \langle \mathbf{v}_c'' \cdot \mathbf{v}_c'' \rangle_d - \mathbf{v}_{TD} \cdot \mathbf{v}_r\} \tag{1894}$$

The final results are obtained by using the drag linearization coefficient A_{cd}^{D} and dispersion velocity $\mathbf{v}_{TD}$ from the earlier section Contribution of Drag to Turbulent Dispersion, together with this identity:

$$\mathbf{v}_{TD} = \langle \mathbf{v}_c \rangle_c - \langle \mathbf{v}_c \rangle_d = \frac{\overline{[\langle \mathbf{v}_c \rangle_c - \mathbf{v}_c]\alpha_d}}{\overline{\alpha_d}} = -\langle \mathbf{v}_c'' \rangle_d \tag{1895}$$

S_{kc} can now be modeled either with the help of a turbulent response model or with the help of the Tchen closures, when taken together with this further approximation:

$$\langle \mathbf{v}_c'' \cdot \mathbf{v}_c'' \rangle_d \approx \langle \mathbf{v}_c'' \cdot \mathbf{v}_c'' \rangle_c = 2k_c \tag{1896}$$

The first two terms in Eqn. (1896) represent an equilibrium or transfer between the covariance of dispersed-continuous velocity fluctuation and the turbulent kinetic energy of the continuous phase.

The last term can be viewed as the dot product of turbulent dispersion force with the mean slip velocity (or alternatively, of the mean drag force with the turbulent dispersion velocity). Its role depends on the alignment of these two vectors.

The dot product can often be small or zero in a fully developed and converged pipe flow. In a horizontal pipe flow case, this small value is because the magnitude of the slip velocity is likely to be small under fully developed conditions. In a vertical pipe flow case, it is because although the slip is not small, it is orthogonal to any gradients of volume fraction under fully developed conditions.

However, the term can be positive or negative, and large, under developing conditions or during convergence. It makes a positive contribution to continuous phase turbulent kinetic energy when mean drag force and turbulent dispersion force are acting in opposite directions. For example, it contributes when particles are falling, or when bubbles are rising, against their own concentration gradients.

Dispersed Phase Turbulence Models

The following groups of models can be selected for representing the effective diffusivity of the dispersed phase in a turbulent flow:

- Laminar
- Turbulent Response models, such as Issa or Tchen
- Reynolds-Averaged models, such as K-Epsilon or Reynolds Stress

The Laminar option is intended for compatibility with external models. Although it is true that the effective diffusivity of the dispersed phase often has little influence, the Laminar option misrepresents wall effects in a turbulent flow. Especially avoid this option especially if wall heat transfer is important.

Where there is a clearly-defined dispersed phase, then Turbulent Response provides a well-defined model for dispersed phase effective turbulent diffusivity and standard wall functions. It is also simple to use.

Full turbulence models, such as K-Epsilon or Reynolds Stress, can also be used for the dispersed phase. Without more user modeling, running full turbulence models for the dispersed phase can often give a comparable solution to turbulence response models. However, this result can happen for no better reason than similarity of flow conditions with the continuous phase together with a weak dependence of the solution on Reynolds number. So this capability is aimed at further development of user models.

Use it either for mixed regimes (continuous and dispersed topology is inverted in different parts of the same flow) or for heavy particles (where a large share of the kinetic energy is held in the dispersed phase).

There are three alternative model options within the Turbulence Response group:

- User-Defined Turbulence Response Model
- Issa Turbulence Response Model
- Tchen Turbulence Response Model

The first two provide a turbulence response coefficient C_t defined as the ratio of dispersed phase velocity fluctuations over continuous phase velocity fluctuations.

$$C_t = \frac{|\mathbf{v}_d'|}{|\mathbf{v}_c'|} \tag{1897}$$

You can provide an expression for C_t, or use the default value of one.

This Issa model provides C_t for a bubbly flow. It estimates C_t to be up to 3 for very small volume fractions of dispersed phase, but rapidly decreasing to unity for dispersed phase volume fractions greater than about 5%.

The Tchen model provides a dispersed phase turbulent diffusivity for gas flows laden with heavy particles, as specified by Thai-Van *et al.* [363]. However, this reference remarks that this model can also be sufficient for bubbly flows because of the weak influence of the eddy diffusivity of the bubbly phase.

User-Defined Turbulence Response Model

The User-Defined Turbulence Response model allows you to define the turbulence response coefficient, C_t, as a user-defined function or a constant. If C_t is equal to 1, the dispersed phase fluctuates with the continuous phase.

The dispersed phase turbulent eddy viscosity is:

$$\mu_d^t = \frac{\rho_d}{\rho_c} C_t^2 \mu_c^t \tag{1898}$$

where:

- ρ_d is the dispersed phase density.
- ρ_c is the continuous phase density.

If the continuous phase uses the Standard K-Epsilon Model, then the dispersed-phase turbulent kinetic energy, k_d, is:

$$k_d = C_t^2 k_c \tag{1899}$$

The K-Epsilon definition of turbulent viscosity for a single phase is:

$$\mu^t = \rho C_\mu \frac{k^2}{\varepsilon} \tag{1900}$$

where:

C_μ is a model coefficient in the Standard K-Epsilon Model.

Combine Eqn. (1782) and Eqn. (1784) to get the turbulent dissipation rate, ε:

$$\varepsilon_d = C_t^2 \varepsilon_c \tag{1901}$$

Issa Turbulence Response Model

The Issa Turbulence Response Model is defined as a correlation for the turbulence response coefficient, C_t, with a volume fraction correction [357]:

$$C_t = 1 + (C_t^* - 1)^{-f(\alpha_d)} \tag{1902}$$

where:

$$f(\alpha_d) = 180\alpha_d - 4.71\times10^3\alpha_d^2 + 4.26\times10^4\alpha_d^3 \tag{1903}$$

C_t^* is computed from the Issa model [370]:

$$C_t^* = \frac{3+\beta}{1+\beta+2(\rho_d/\rho_c)} \tag{1904}$$

β is defined as:

$$\beta = \frac{2A_{ij}^D l_e^2}{\alpha_d \mu_c Re_t} \tag{1905}$$

and:

$$l_e = C_\mu \frac{k_c^{3/2}}{\varepsilon_c} \tag{1906}$$

$$Re_t = \frac{|\mathbf{u}_c'| l_e}{\mathbf{v}_c} \tag{1907}$$

$$|\mathbf{u}_c'| = \sqrt{\frac{2}{3}k_c} \tag{1908}$$

The turbulence diffusivity, kinetic energy, and dissipation for the dispersed phase are then estimated as for the User Defined Response Model.

Tchen Turbulence Response Model

From Thai-Van *et al.* [363], the turbulence diffusivity for turbulent gas flows laden with heavy particles can be modeled as:

$$\nu_d^{\ t} = \tau_I \frac{1}{3} q_{cd} + \frac{1}{2} \tau_R \frac{2}{3} q_d^{\ 2} \tag{1909}$$

where:

- q_{cd} is the dispersed-phase average of the product of dispersed and continuous phase velocity fluctuations.
- q_d is half the dispersed-phase average of the product of dispersed and dispersed phase velocity fluctuations.

For more details on the closure of these correlations, see Tchen Closures.

Particle Induced Turbulence Models

STAR-CCM+ has three optional models for modeling the effect of particles on the continuous-phase turbulent kinetic energy:

- Troskho Hassan particle induced turbulence
- Gosman particle induced turbulence
- Tchen particle induced turbulence

Any of these three models can be used in combination with a K-Epsilon model for the continuous phase turbulence. You can also use a Reynolds Stress Model for the continuous phase turbulence. The Reynolds Stress Model assumes that the induced turbulence models make an isotropic contribution to each of the normal stress equations.

The Troshko-Hassan model begins by assuming that the whole of the work that is done by the drag force is an unconditionally positive production of continuous-phase pseudo-turbulence. Another feature of this model is that it assumes that this strong turbulence source is dissipated locally using a particle relaxation time scale.

The Gosman and Tchen models instead focus only on the turbulent contribution of the drag force. The dissipation time scale for the resulting source terms is assumed to be eddy-turnover time, just as for shear-induced turbulence.

Both Gosman and Tchen models result in two kinds of term in the continuous phase turbulence, either of which can make positive or negative contributions to the continuous-phase turbulence kinetic energy. (See section Particle Induced Turbulence for the meaning of these terms).

Troshko Hassan Particle Induced Turbulence Model

The Troshko and Hassan model describes bubble induced turbulence effects [367] by providing source terms to the continuous K-Epsilon model. It uses the Virtual Mass phase interaction model. If the Virtual Mass phase interaction model is not selected, then the virtual mass coefficient defaults to the value of a spherical bubble:

$$C_{ij}^{VM} = \frac{1}{2} \tag{1910}$$

The source term is derived by considering that the correlation between velocity fluctuation at the interface and interfacial force density is exactly equal to the work of interfacial force density per unit time. For bubbly flows, the drag force dominates the interfacial force. The original form from Troshko and Hassan [367] is:

$$S_{k_c} = (\mathbf{u}_j - \mathbf{u}_i) \cdot F_{ij}^D = \frac{3}{4}\frac{C_{ij}^D}{d}\alpha_d \rho_c \left|\mathbf{u}_j - \mathbf{u}_i\right|^3 \tag{1911}$$

where:

C_{ij}^D is the bubble drag coefficient.

In STAR-CCM+ Eqn. (1908) is rewritten in terms of the linearized drag coefficient, A_{IJ}^D , to account for multi-particle effects:

$$S_{k_c} = A_{ij}^D \left|\mathbf{u}_j - \mathbf{u}_i\right|^2 \tag{1912}$$

The dissipation term uses the energy source term that decays with a characteristic time, the Bubble Pseudo-Turbulence Dissipation Relaxation (BPTDR) time, t_b, and is scaled with a calibration constant C_3:

$$S_{\varepsilon_c} = \frac{C_3 S_{k_c}}{t_b} \tag{1913}$$

For the data that was analyzed, Troshko and Hassan found $C_3 = 0.45$. The bubble pseudo-turbulence dissipation relaxation time is:

$$t_b = \frac{2C_{ij}^{VM} d}{3C_{ij}^D \left|\mathbf{u}_j - \mathbf{u}_i\right|} \tag{1914}$$

Assuming that $\rho_d \ll \rho_c$ for bubbly flow and again substituting C_{ij}^D with A_{IJ}^D gives:

$$t_b = \frac{\rho_c C_{ij}^{VM} \alpha_d}{2A_{ij}^D} \tag{1915}$$

where C_{ij}^{VM} is the virtual mass coefficient.

If this model is used with the Reynolds Stress equations, its contribution is assumed to be isotropic:

$$\mathbf{S}_{Rc} = \frac{2}{3} S_{kc} \mathbf{I} \tag{1916}$$

Gosman Particle Induced Turbulence Model

Reynolds averaging of continuous-phase velocity fluctuation times the drag source term from the fluctuating part of the continuous-phase momentum equation results in the following source term for the continuous-phase turbulent kinetic energy:

$$S_{kc} = A_{cd}^D \left(q_{cd} - 2q_c^{\ 2} - \mathbf{v}_{TD} \cdot \mathbf{v}_r\right) \tag{1917}$$

where:

- q_{cd} is the dispersed-phase average of the product of dispersed and continuous phase velocity fluctuations.
- q_c is half the dispersed-phase average of the product of continuous and continuous phase velocity fluctuations.

See Particle Induced Turbulence for the origin of this term, but also keep in mind that:

- Only the contribution of the drag force has been considered in this derivation.
- The closure that is used for the drag force is good under Stokesian conditions but only approximate if A_{cd}^D represents a non-Stokesian drag law.

Gosman *et al.* [333] describe one model for this drag-related source term, which is essentially:

$$S_{kc} = A_{cd}(C_t - 1)2K_c - \mathbf{F}_{cd}^{TD} \cdot \mathbf{v}_r \tag{1918}$$

and can be arrived at simply by approximating:

$$\frac{q_{cd}}{2q_c^{\,2}} \approx \frac{|\mathbf{v}_d'|}{|\mathbf{v}_c'|} = C_t \tag{1919}$$

so the $(C_t - 1)$ contribution represents an exchange between cross-phase covariance and continuous-phase turbulent kinetic energy.

Although the final term has the appearance of a work term, in fact it has the opposite of the expected sign. If the relative motion of the phases is in the same direction as the turbulent dispersion force between the phases, the final term can become a sink of continuous phase turbulent kinetic energy.

The corresponding contribution to the Reynolds Stress equations is assumed to be isotropic:

$$\mathbf{S}_{Rc} = \frac{2}{3} S_{kc} \mathbf{I} \tag{1920}$$

The dominant dissipation time scale is assumed to be k_c/ε_c as for shear-induced turbulence, but only the $(C_t - 1)$ exchange term is allowed to contribute.

$$S_{\varepsilon c} = A_{cd}^{D} \; (C_t - 1) 2\varepsilon_c \tag{1921}$$

Tchen Particle Induced Turbulence Model

Reynolds averaging of continuous-phase velocity fluctuation times the drag source term from the fluctuating part of the continuous-phase momentum equation results in the following source term for the continuous-phase turbulent kinetic energy:

$$S_{kc} = C_0 A_{cd}^{D} \; (q_{cd} - 2q_c^{\,2} - \mathbf{v}_{TD} \cdot \mathbf{v}_r) \tag{1922}$$

where:

- q_{cd} is the dispersed-phase average of the product of dispersed and continuous phase velocity fluctuations.
- q_c is half the dispersed-phase average of the product of continuous and continuous phase velocity fluctuations.
- Constant C_0 is not part of the model derivation, but defaults to 1. You can adjust it to test the sensitivity of the solution to this term.

See Particle Induced Turbulence for the origin of this term, but also keep in mind that:

- Only the contribution of the drag force has been considered in this derivation.

- The closure that is used for the drag force is good under Stokesian conditions but only approximate if A_{cd}^{D} represents a non-Stokesian drag law.

Taken together, $q_{cd} - 2q_c^{\ 2}$ represents an exchange between cross-phase covariance and continuous-phase turbulent kinetic energy.

$A_{cd}^{D}\ (\mathbf{v}_{TD} \cdot \mathbf{v}_r)$ represents the work that is done by the turbulent dispersion force against mean slip velocity. See Particle Induced Turbulence.

The corresponding contribution to the Reynolds Stress equations is assumed to be isotropic:

$$\mathbf{S}_{Rc} = \frac{2}{3} S_{kc} \mathbf{I} \tag{1923}$$

Following Elghobashi and Abou-Arab [328], the dominant dissipation time scale is assumed to be the same as for shear-induced turbulence:

$$S_{\varepsilon c} = C_{\varepsilon 3} \frac{\varepsilon_c}{k_c} \max(S_{kc}, 0) \tag{1924}$$

A negative turbulent energy source term is not allowed to contribute to the turbulent dissipation equation.

Parameter $C_{\varepsilon 3}$ has a default value of 1.44, and you can adjust it independently of constants for other terms.

The correlations q_{cd} and $q_c^{\ 2}$ are closed using Tchen theory as described in Tchen Closures:

$$\begin{aligned} q_{cd} &= 2\left[\frac{b+\eta}{1+\eta}\right] k_c \\ q_c^{\ 2} &= k_c \end{aligned} \tag{1925}$$

where:

- η is the particle-eddy interaction time, which is scaled by particle relaxation time.
- b is the ratio of the coefficients of the continuous/disperse phase acceleration terms in the particle equation of motion:

$$b = \frac{1 + C_{VM}}{\frac{\rho_d}{\rho_c} + C_{VM}} \tag{1926}$$

After closure, the transfer term $q_{cd} - 2q_c^{\ 2}$ becomes proportional to $(b - 1)$.

So for gas-particle flow, $\rho_d \gg \rho_c$, where coefficient b is small, the transfer term becomes a sink of continuous phase turbulence kinetic energy.

On the other hand, for bubbly flow, $\rho_d \ll \rho_c$, where coefficient b tends to $3/2$ if C_{VM} is about $1/2$, the transfer term becomes a source of continuous phase turbulence kinetic energy.

Wall Permeability

A wall boundary that has been set to Phase Permeable for a particular phase allows a phase mass flux approaching the boundary to leave the system at this boundary. The size of the mass flux for phase *i* is equal to:

$$\dot{m}_i = \rho_i \mathbf{u}_i \cdot \mathbf{A} \alpha_i \gamma_i \tag{1927}$$

where:

- m_i is the mass flux leaving the system.
- ρ_i is the cell-centered density.
- $\mathbf{u}_i$ is the cell-centred velocity.
- α_i is the cell-centered volume fraction.
- $\mathbf{A}$ is the cell wall area vector (normal to the wall in the outward facing direction).
- γ_i is the wall permeability factor.

The wall permeability takes a value from 0 through 1 to control the amount of the flux incident on the surface that is removed from the flow. Any component of the wall velocity parallel to the wall is the same as for a standard wall (zero in this case for a no-slip wall). When a permeable wall phase condition is activated, the mass flux that is calculated at the wall cell is used in all the other transport equations.

Energy Conservation Equation

Energy is conserved for a Eulerian phase according to the following equation:

$$\frac{\partial}{\partial t}(\alpha_i \rho_i E_i) + \nabla \cdot [\alpha_i \rho_i H_i (\mathbf{v}_i - \mathbf{v}_g)] + \nabla \cdot \alpha_i \mathbf{v}_g p = \tag{1928}$$

$$\nabla \cdot (\alpha_i k_{eff,\,i} \nabla T_i) + \nabla \cdot (\mathbf{T}_i \cdot \mathbf{v}_i) + \mathbf{f}_i \cdot \mathbf{v}_i +$$

$$\sum Q_{;i} + \sum Q_i^{(ij)} + S_{,,\;;} + \sum (m_{;i} - m_{i;}) h_i (T_{;i})$$

where:

- E_i is the total energy.
- H_i is the total enthalpy.
- $\mathbf{T}_i$ is the viscous stress tensor.
- α_i, ρ_i and $\mathbf{v}_i$ have the same meaning as in Eqn. (1692).
- $\mathbf{v}_g$ is the grid velocity.
- T_i is the temperature.
- $k_{eff,\,i}$ is the effective thermal conductivity, from Eqn. (1929).
- $\mathbf{f}_i$ is the body force vector.
- Q_{ij} is the interphase heat transfer rate to phase i from phase j.
- $Q_i^{(ij)}$ is the heat transfer rate to phase *i* from phase pair interface *(ij)*.
- $S_{u,\,i}$ is the energy source.
- $h_i(T_{ij})$ is the phase *i* enthalpy that is evaluated at the interface temperature T_{ij}.

The effective thermal conductivity, k_{eff}, is:

$$k_{eff,\,i} = k_i + \frac{\mu_{t,\,i} C_{p,\,i}}{\sigma_{t,\,i}} \tag{1929}$$

where:

- k_i is the thermal conductivity.
- $\mu_{t,\,i}$ is the turbulent viscosity.
- $C_{p,\,i}$ is the specific heat.
- $\sigma_{t,\,i}$ is the turbulent thermal diffusion Prandtl number.

Now, total energy, E, is related to the total enthalpy, H, by

$$E_i = H_i - \frac{p}{\rho_i} \tag{1930}$$

where:

$$H_i = h_i + \frac{|\mathbf{v}_i|^2}{2} \tag{1931}$$

and

$$h_i(T_i) = h_i^{REF} + \int_{T_i^{REF}}^{T_i} C_{p,\,i}(T')dT' \tag{1932}$$

When heat is transferred only by thermal diffusion between phases, the interphase heat transfer rates satisfy $Q_{ij} = -Q_{ji}$. Where heat and mass are only transferred by phase change (bulk boiling or condensation), the heat transfer rates between each phase *i,j,* and a phase pair interface *(ij)* satisfy a heat balance:

$$Q_i^{(ij)} + Q_j^{(ij)} + (m_{ij} - m_{ji})\Delta h_{ij} = 0 \tag{1933}$$

where Δh_{ij} is the heat required to produce phase i from phase j. This is a signed quantity so that $\Delta h_{ij} = -\Delta h_{ji}$. The heat of phase change is defined from the phase enthalpy definitions that are evaluated at the interface temperature T_{ij}:

$$\Delta h_{ij} = h_j(T_{ij}) - h_i(T_{ij}) \tag{1934}$$

Interphase Energy Transfer

The heat transfer by thermal diffusion between phase *i* and phase *j* is modeled as:

$$Q_{ij} = -Q_{ji} = h^{(ij)} a_{ij}(T_j - T_i) \tag{1935}$$

where:

- Q_{ij} is the interphase heat transfer rate to phase i from phase *j*
- $h^{(ij)}$ is the mean surface average heat transfer coefficient
- a_{ij} is the interfacial area per unit volume

If *c* and *d* represent the continuous and dispersed phases in the phase interaction pair (*ij*), the mean heat transfer coefficient is computed as a function of the continuous phase thermal conductivity, the Nusselt number, and the interaction length scale:

$$h^{(ij)} = \frac{k_c Nu}{l_{cd}} \tag{1936}$$

In STAR-CCM+ the Nusselt number is:

1. A constant value (default = 2.0).
2. Supplied by a field function.
3. Computed using the Ranz-Marshall correlation [355]:

$$Nu = 2 + 0.6 Re_d^{0.5} Pr_c^{0.3} \tag{1937}$$

where:

- Pr_c is the continuous phase Prandtl number

- Re_d is the dispersed phase Reynolds number

Interphase Boiling Mass Transfer

The heat transfer from the phase-change interface to each of the two phases on either side of the interface is modeled as:

$$Q_i^{(ij)} = h_i^{(ij)} a_{ij}(T_{ij} - T_i) \tag{1938}$$

$$Q_j^{(ij)} = h_j^{(ij)} a_{ij}(T_{ij} - T_j) \tag{1939}$$

where:

- T_i and T_j are the bulk temperatures of the continuous and dispersed phases, respectively
- T_{ij} is the interface temperature

The interface temperature is often defined as a constant saturation temperature for a particular system pressure. In STAR-CCM+, a field function can optionally be used for the interface temperature, so that depth can be accounted for, or overpressure due to surface tension in micron-sized bubbles.

The heat transfer coefficient for energy supply from the continuous phase is:

$$h_i^{(ij)} = Nu_i k_i / l_{ij} \tag{1940}$$

where Nu_i is the continuous phase Nusselt number. Similarly, for the dispersed phase, the heat transfer coefficient for energy supply is:

$$h_j^{(ij)} = Nu_j k_j / l_{ij} \tag{1941}$$

where Nu_j is the dispersed phase Nusselt number. In both equations, l_{ij} is the interaction length scale for the phase pair--usually the bubble size.

The Ranz-Marshall correlation that is used for the continuous phase Nusselt number has the same definition as used for the interphase heat transfer. This correlation is not available for the dispersed phase Nusselt number since convective processes are often poorly defined inside the particle.

The interface mass flux is computed from the heat balance:

$$m^{(ij)} = -\frac{Q_i^{(ij)} + Q_i^{(ij)}}{\Delta h_{ij}} \tag{1942}$$

where:

- $m^{(ij)}$ is the signed mass transfer rate per unit interfacial area for the interaction from phase j to phase i
- Δh_{ij} is the signed enthalpy input to create phase *j* from phase *i*

Distinguish interphase mass transfer source and sink terms appearing in various phase transport equations according to sign. This technique allows you to import or export different interface values of the transported variable:

$$m_{ij} = max(m^{(ij)}, 0) \tag{1943}$$

$$m_{ji} = max(-m^{(ij)}, 0) \tag{1944}$$

so that the signed flux $m^{(ij)}$ is equivalent to $m_{ij} - m_{ji}$.

Multi-component Boiling Mass Transfer

The purpose of the Multi-component Boiling Mass Transfer model is to calculate the molar flux and the temperature at the interface. These terms are needed for the calculation of mass transfer, momentum transfer, and energy transfer between the phases. The governing equations of the Multi-component Boiling Mass Transfer model are presented in the following sections.

Species Fluxes

Before introducing the governing equations, this section defines some quantities that are used in the description of the model.

Consider a stationary coordinate reference frame and define $\mathbf{u}_i$ as the velocity of component *i* in this frame. The molar flux of species *i* is expressed as:

$$\mathbf{N}_i = c_i \mathbf{u}_i \quad \left[\frac{\mathrm{mol}}{\mathrm{m^2 s}}\right] \tag{1945}$$

Sum over all the species to get the total molar flux:

$$\mathbf{N}_t = \sum_{i=1}^{n} \mathbf{N}_i = c_t \mathbf{u} \tag{1946}$$

where the molar average velocity $\mathbf{u}$ is defined as:

$$\mathbf{u} = \sum_{i=1}^{n} x_i \mathbf{u}_i \tag{1947}$$

The diffusion flux, which is the flux of component i relative to the flux of the mixture, can now be defined. Choose $\mathbf{u}$ as the mixture reference velocity and therefore define the molar diffusion flux as:

$$\mathbf{J}_i = c_i(\mathbf{u}_i - \mathbf{u}) \tag{1948}$$

with:

$$\sum_{i=1}^{n} \mathbf{J}_i = 0 \tag{1949}$$

The molar flux $\mathbf{N}_i$ is related to the molar diffusion flux by:

$$\mathbf{N}_i = \mathbf{J}_i + c_i\mathbf{u} = \mathbf{J}_i + x_i\mathbf{N}_t \tag{1950}$$

Film Model

The assumptions of the Film Model can be summarized as follows:

- All the resistance to mass transfer is localized in a thin film that is adjacent to the phase boundary.
- Mass transfer occurs inside this film by steady-state molecular diffusion alone.
- Mass transfer through the film occurs in the direction normal to the interface (that is, composition gradients along the interface are neglected).
- There are no composition gradients outside the film.

Diffusion Flux Equation

With the previous assumptions, the species concentration gradients across the film can be calculated using finite differences and the equation for the diffusion flux becomes:

$$J_i = N_i - x_iN_t = c_t\sum_{i=1}^{n-1} k_{ij}{}^*\Delta x_j \tag{1951}$$

where we express the finite flux mass transfer coefficient matrix as

$$[k^*] = [k][\Xi] \tag{1952}$$

with $[k]$ and $[\Xi]$ being the low-flux mass transfer coefficient matrix and the finite flux correction coefficient matrix respectively.

Mass Transfer Coefficient in the Zero Flux Limit

The low-flux mass transfer coefficient matrix $[k]$ is defined as:

$$[k] = \frac{1}{l}[B]^{-1} \tag{1953}$$

$$B_{ii} = \frac{x_i}{D_{in}} + \sum_{\substack{k=1 \\ k \neq i}}^{n} \frac{x_k}{D_{ij}} \tag{1954}$$

$$B_{ij} = -x_i\left(\frac{1}{D_{ij}} - \frac{1}{D_{in}}\right) \tag{1955}$$

The length scale l that appears in Eqn. (1953) is derived from the film theory and is the thickness of the film layer. This quantity can be calculated from the Sherwood number, since:

$$\mathrm{Sh} = \frac{kd}{D} = \frac{d}{l} \tag{1956}$$

from which $l = \frac{d}{\mathrm{Sh}}$ can be obtained.

The resulting film thickness is a function of flow conditions, geometry, and fluid properties. In STAR-CCM+, only one characteristic length d and one Sherwood number Sh can be specified for each phase interaction. Therefore, the same value of l is calculated and applied for all of the transferred components.

Mass Transfer Correction Coefficient Matrix

The form of the correction coefficients depends on the mass transfer model that is adopted. For the two fluid film model, following the convention of positive flux from the vapor to the liquid, the correction matrices are:

$$[\Xi]^V = [\Phi]^V[\exp[\Phi]^V - [I]]^{-1} \tag{1957}$$

and

$$[\Xi]^L = [\Phi]^L\exp[\Phi]^L[\exp[\Phi]^L - [I]]^{-1} \tag{1958}$$

for the vapor and the liquid respectively, where:

$$\Phi_{ii} = \frac{N_i}{c_t\kappa_{in}} + \sum_{\substack{k=1 \\ k \neq i}}^{n} \frac{N_k}{c_t\kappa_{ik}} \tag{1959}$$

$$\Phi_{ij} = -N_i\left(\frac{1}{c_t\kappa_{ij}} - \frac{1}{c_t\kappa_{in}}\right) \tag{1960}$$

Interface Energy Balance

Continuity of the energy flux across the vapor - liquid interface is imposed:

$$E^V = E^L = E \tag{1961}$$

where E^V and E^L are the normal components of the energy flux at the interface. Eqn. (1961) can be rewritten as:

$$q^V + \sum_{i=1}^{''} N_i^V \overline{H}_i^V(T^I) = q^L + \sum_{i=1}^{''} N_i^L \overline{H}_i^L(T^I) \tag{1962}$$

with the heat fluxes given by:

$$q^V = h_V{}^*(T^V - T^I) \tag{1963}$$

$$q^L = h_L{}^*(T^I - T^L) \tag{1964}$$

The finite flux heat transfer coefficient in the vapor phase is calculated according to the film model as:

$$h_V{}^* = h^V(\Phi_H^V)/(\exp\Phi_H^V - 1) = \Xi_H h^V \tag{1965}$$

where h^V is the zero-flux heat transfer coefficient, $\Xi_H = (\Phi_H^V)/(\exp\Phi_H^V - 1)$ is the Ackermann correction factor and

$$\Phi_H = \sum_{i=1}^{''} N_i W_i c_{pi}/h \tag{1966}$$

Similarly, for the liquid:

$$h_L{}^* = h^L(\Phi_H^L)\exp\Phi_H^L/(\exp\Phi_H^L - 1) = \Xi_H^L h^L \tag{1967}$$

with $\Xi_H^L = (\Phi_H^L)\exp\Phi_H^L/(\exp\Phi_H^L - 1)$.

The high flux correction for h^L is expected to be close to unity due to the high value of the liquid phase heat transfer coefficient.

Species Equilibrium at the Interface

The composition of the phases on either side of the interface between the two phases are assumed to be in instant local equilibrium with each other and with the interface temperature and the system pressure:

$$y_i^I = K_i^0(P,T^I)x_i^I \tag{1968}$$

where ideal behavior of the mixtures is assumed. The mole fractions also obey the constraints:

$$\sum x_i^I = 1 \tag{1969}$$

$$\sum y_i^I = 1 \tag{1970}$$

Note that a solution of these interface state relations is possible only if at least one of the vapor-liquid equilibrium coefficients K_i^0 is less than unity while at least one of the K_i^0 is greater than unity.

Equations of the Point Solver

Consider two mixtures x and y divided by an interface. The superscripts L and V are used for bulk quantities in the liquid and in the vapor respectively, while I indicates interface quantities. The interface is considered as a surface with no resistance to mass transfer and where equilibrium conditions prevail at any time. It is assumed that no chemical reactions occur at the interface or in the film.

When the interface is not stationary, the fluxes $\mathbf{N}_i$ refer to the interface and not to the laboratory fixed frame of reference (see [362] for more details). It is assumed that the mass transfer takes place in a direction normal to the interface:

$$N_i^x = N_i^y = N_i \quad i = 1, 2, ..., n \tag{1971}$$

where N_i^x is the normal component of $\mathbf{N}_i$ and is directed from the bulk phase x to the interface I while N_i^y is directed from the interface I to the bulk phase y.

In the general case, there can be m species in the liquid and n in the gas, not all of them volatile or condensable. Since there are k couples between the volatile and condensable species, then $m-k$ species are non-volatile and $n-k$ species are non-condensable.

Thus, the variables to calculate are m interface mole fractions x_i^I, n interface mole fractions y_i^I, k interface molar fluxes N_i, and the interface temperature T^I.

Therefore, the total number of variables is $m + n + k + 1$ and the equations of the model are:

- $m - 1$ equations from the interface molar fluxes (liquid side) (see Diffusion Flux Equation for details)

$$N_i - x_i^I N_t = c_t^L \sum_{i=1}^{m-1} k_{ij}{}^{*L}(x_j^I - x_j) \quad i = 1, \text{k+1}, \ldots, \text{m-1} \tag{1972}$$

- $n - 1$ equations from the interface molar fluxes (vapor side) (see Diffusion Flux Equation for details)

$$N_i - y_i^I N_t = c_t^V \sum_{i=1}^{n-1} k_{ij}{}^{*V}(y_j - y_j^I) \quad i = 1, \text{k+1}, \ldots, \text{n-1} \tag{1973}$$

- 1 equation from the energy balance at the interface (see Interface Energy Balance for details)

$$h_V{}^*(T^V - T^I) + \sum_{i=1}^{n} N_i^V \overline{H}_i^V W_i(T^I) = h_L{}^*(T^I - T^L) + \sum_{i=1}^{m} N_i^L \overline{H}_i^L W_i(T^I) \tag{1974}$$

- k equations from the equilibrium relations at the interface (see Species Equilibrium at the Interface for details)

$$y_i^I = K_i x_i^I \quad i = 1, 2, \ldots, k \tag{1975}$$

- 1 equation from the constraint on mole fractions at the interface (liquid side)

$$\sum_{i=1}^{m} x_i^I = 1 \tag{1976}$$

- 1 equation from the constraint on mole fractions at the interface (vapor side)

$$\sum_{i=1}^{n} y_i^I = 1 \tag{1977}$$

The number of equations is, therefore, $m + n + k + 1$ as the number of variables. If species i is not volatile or non-condensable, then $N_i = 0$ in Eqn. (1972) and Eqn. (1973). If the liquid is a pure substance ($m = 1$), there are no equations of the form Eqn. (1972) and Eqn. (1976) is always satisfied.

Likewise, if the vapor is a pure substance ($n = 1$) then there are no equations of the form Eqn. (1973) and Eqn. (1977) is always satisfied.

Eqn. (1972) to Eqn. (1977) are a system of non-linear algebraic equations that are solved for each cell of the computational domain with a Newton solver.

Reduced Model for Dilute Dispersions

When the Maxwell-Stefan diffusion model is not active, the assumption of dilute dispersion is made, that is $x_i \approx 0$ for $i = 1, 2, ..., n\text{-}1$ and $x_n \approx 1$. With this assumption, only the interactions between each disperse species and the last species in the list (the solvent) are accounted for. The interactions between the disperse species are neglected. Consequently, Eqn. (1972) and Eqn. (1973) reduce to:

$$N_i - x_i^I N_t = c_t^L \kappa_{in}^L \frac{\Phi_{ii}^L e^{\Phi_{ii}^L}}{e^{\Phi_{ii}^L} - 1}(x_i^I - x_i) \tag{1978}$$

and

$$N_i - y_i^I N_t = c_t^V \kappa_{in}^V \frac{\Phi_{ii}^V}{e^{\Phi_{ii}^L} - 1}(y_i - y_i^I) \tag{1979}$$

respectively, with $\kappa_{in} = \frac{D_{in}}{l}$ and $\Phi_{ii} = \frac{N_t}{c_t \kappa_{in}}$

The other equations are not affected by the assumption of dilute dispersion.

Nomenclature

The terms that are used in the equations of the Multi-component Boiling Mass Transfer model are defined in the following table.

c_i	Molar density of $i = \rho_i / M_i$	$\left[\frac{\text{mol}}{\text{m}^3}\right]$
c_t	Mixture molar density $= \sum_{i=1}^{n} c_i$	$\left[\frac{\text{mol}}{\text{m}^3}\right]$
c_p	Specific heat	$\left[\frac{\text{J}}{\text{kgK}}\right]$
d	Characteristic length scale of the system	m
$\lvert D \rvert$	Matrix of Fick diffusion coefficients	$\left[\frac{\text{m}^2}{\text{s}}\right]$

$\lvert D \rvert$	Matrix of Maxwell-Stefan diffusion coefficients	$\left[\frac{m^2}{s}\right]$
h	Heat transfer coefficient	$\left[\frac{W}{m^2 K}\right]$
H	Specific enthalpy	$\left[\frac{J}{kg}\right]$
$\lvert I \rvert$	Unit tensor	
$\mathbf{J}_i$	Molar diffusion flux $= c_i(\mathbf{u}_i - \mathbf{u})$	$\left[\frac{mol}{m^2 s}\right]$
l	Film thickness	m
M_i	Molar mass of component i	$\left[\frac{kg}{mol}\right]$
$\mathbf{N}_i$	Molar flux of component i	$\left[\frac{mol}{m^2 s}\right]$
Sh	Sherwood number	
$\mathbf{u}$	Molar averaged referenced velocity $= \sum_{i=1}^{''} x_i \mathbf{u}_i$	$\left[\frac{m}{s}\right]$
x_i, y_i	Bulk mole fraction of component i (liquid side, vapor side)	
x_i^I, y_i^I	Interface mole fraction of component i (liquid side, vapor side)	
W_i	Molecular weight of i	$\left[\frac{g}{mol}\right]$
μ_i	Molar chemical potential of i	$\left[\frac{J}{mol}\right]$
ρ_i	Mass density of $i = c_i M_i$	$\left[\frac{kg}{m^3}\right]$
ρ_t	Mixture mass density $= \sum_{i=1}^{''} \rho_i$	$\left[\frac{kg}{m^3}\right]$

ω_i	Mass fraction of $i = \rho_i / \rho_t$
$[\Phi]$	Matrix of mass transfer rate factors
$[\Xi]$	Matrix of high flux correction factors

Thermal Boundary Conditions on Walls

Wall Energy Balance

For Eulerian multiphase flows, it is assumed that the same local wall temperature T_w is experienced by all phases in contact with the wall. This useful closure, together with an energy balance at the wall, allows the most common thermal boundary conditions to be specified for the wall, rather than for each individual phase.

The energy balance for the wall is:

$$\dot{q}''_w = \sum_i \dot{q}''_{wi} + \sum_{(ij)} \dot{q}''_{w(ij)} \tag{1980}$$

where:

- $\dot{q}''_w$ is the external heat flux from the environment to the wall
- $\dot{q}''_{wi}$ is the heat flux from the wall to the phase i fluid next to the wall
- $\dot{q}''_{w(ij)}$ is the heat flux from the wall to a phase interaction (ij) interface next to the wall

External Heat Flux

The four specifications available for modeling the external heat flux q''_w are as follows:

- Adiabatic: $\dot{q}''_w$ is zero and T_w is computed.
- Heat Flux: $\dot{q}''_w$ is specified and T_w is computed.
- Temperature: T_w is specified and $\dot{q}''_w$ is computed.
- Convection: Ambient temperature T_∞ and external heat transfer coefficient h_∞ are specified, with both $\dot{q}''_w$ and T_w computed.

  The external heat flux for the Convection option is:

$$\dot{q}''_w = h_\infty (T_\infty - T_w) \tag{1981}$$

Phase Heat Flux

The standard model for internal heat flux from each phase i to the wall, $\dot{q}''_{wi}$, is based on single-phase wall treatments within each phase, together with an assumption that each phase has a local wall contact area proportional to the volume fraction next to the wall.

For turbulent flow, these two models combine into:

$$\dot{q}''_{wi} = \alpha_i \frac{\rho_i c_{pi} u_i^*}{t_i^+}(T_w - T_i) \tag{1982}$$

where:

- α_i is the volume fraction of phase i
- u_i^* is the turbulent wall velocity scale of phase i
- t_i^+ is defined by the thermal wall law for phase i

For laminar flow, this relationship is:

$$\dot{q}''_{wi} = \frac{-\dot{\mathbf{q}}''_{wi} \cdot \mathbf{a}}{|\mathbf{a}|} = \frac{\alpha_i k_i \nabla T_i \cdot \mathbf{a}}{|\mathbf{a}|} \tag{1983}$$

where $\mathbf{a}$ is the outward facing wall face area vector.

You can also override the standard relationship for phase heat flux, by specifying three coefficients A_i, B_i and C_i:

$$-\dot{q}''_{wi} = A_i + B_i T_i + C_i T_w \tag{1984}$$

This allows various opportunities for alternative heat flux relationships, wall contact models, or linearization methods.

Note: The sign convention, which is based on the discretization of fluxes using outward facing areas. Since either side of this formula is defined to be positive when heat flows from the fluid to the wall, the coefficient of the fluid temperature, B_i, is always positive or zero. The coefficient of the wall temperature, C_i, is always negative or zero.

The internal coefficients from single-phase wall treatments are available as field functions for constructing multiphase coefficients. For example, the standard turbulent or laminar wall contact model is equivalent to specifying the coefficients:

$$A_i = \alpha_i A_{INTERNAL,\, i}$$
$$B_i = \alpha_i B_{INTERNAL,\, i} \tag{1985}$$
$$C_i = \alpha_i C_{INTERNAL,\, i}$$

Phase Interaction Heat Flux

This term is absent from the wall heat balance, unless the phase interaction (*ij*) includes phase change, in which case this heat flux drives the rate of mass transfer at the wall by:

$$\dot{q}''_{w(ij)} = -\dot{m}''_{w(ij)} h_{(ij)} \tag{1986}$$

where:

- $\dot{m}''_{w(ij)}$ is the rate of interphase mass transfer from phase *j* to phase *i* per unit wall area, and
- $h_{(ij)}$ is the enthalpy of phase change from phase *i* to phase *j*.

Standard relationships for phase interaction heat fluxes are described in the Wall Boiling formulation section.

Override a standard relationship for phase interaction heat flux, by specifying three coefficients $A_{(ij)}$, $B_{(ij)}$, and $C_{(ij)}$:

$$-\dot{q}''_{w(ij)} = A_{(ij)} + B_{(ij)} T_{(ij)} + C_{(ij)} T_w \tag{1987}$$

where:

- $T_{(ij)}$ is the temperature of the interface between phases *i* and *j*, which is typically specified as the saturation temperature.

For example, an evaporation heat flux model that depends on some power of the wall superheat, such as:

$$\dot{q}''_{w(ij)} = K \max(T_w - T_{(ij)}, 0)^n \tag{1988}$$

would be linearized by differentiation with respect to each temperature:

$$\dot{q}''_{w(ij)}(T_{(ij)}, T_w) = \dot{q}''_{w(ij)}(T_{(ij)}{}^*, T_w{}^*) + (T_{(ij)} - T_{(ij)}{}^*)\left[\frac{\partial}{\partial T_{(ij)}} \dot{q}''_{w(ij)}\right]^* + (T_w - T_w{}^*)\left[\frac{\partial}{\partial T_w} \dot{q}''_{w(ij)}\right]^* \tag{1989}$$

where an asterisk signifies a value at the previous iteration.

Collecting terms in $T_{(ij)}$ and T_w leads to the implementation:

$$
\begin{aligned}
A_{ij} &= (n-1)K\,\max(T_w{}^* - T_{(ij)}{}^*, 0)^n \\
B_{ij} &= -nK\,\max(T_w{}^* - T_{(ij)}{}^*, 0)^{n-1} \\
C_{ij} &= -B_{(ij)}
\end{aligned}
\tag{1990}
$$

Wall Boiling

Within this formulation, some model properties are referred to as "standard values". These values are calibration parameters that were carefully chosen by the original authors. Do not change unless you are an expert with experimental support. Some models are implemented specifically for water systems, so adjust the highlighted model properties if a working fluid other than water is used:

Subcooled Boiling Model

The subcooled boiling model of Kurul and Podowski [343], with contributions by many other workers, assumes that all the heat flux from the wall is transferred to the liquid next to the wall, and is then partitioned into three components:

$$
\dot{q}''_w = \dot{q}''_{conv} + \dot{q}''_{evap} + \dot{q}''_{quench} \tag{1991}
$$

These components of heat flux are described below:

- $\dot{q}''_{conv}$ is the convective heat flux (see Eqn. (1997)), which describes the removal of heat by single-phase turbulent convection on those parts of the wall that is not affected by boiling. In applications with fixed wall heat flux, this term defines the point at which the wall first exceeds saturation temperature.
- $\dot{q}''_{evap}$ is the evaporative heat flux (see Eqn. (2003)), which describes the power that is used to produce bubbles from nucleation to departure. This term is a strong function of wall superheat $T_w - T_{sat}$. In fixed heat flux applications, once boiling has started, this term is responsible maintaining a wall temperature that is slightly higher than the saturation temperature.
- $\dot{q}''_{quench}$ is the quenching heat flux (see Eqn. (2021)), which describes the enhancement of heat transfer, due to the replacement of a departing bubble by an influx of cooler liquid farther away from the wall. Bubble-induced quenching is also known in literature as liquid agitation or pumping. This term is less important when the liquid is close to saturation temperature, but in highly subcooled flows, $T_l \ll T_{sat}$, it makes an important contribution to the enhanced heat transfer due to boiling.

Vapor Contribution

Subcooled boiling models start with the assumption that only the liquid phase is in contact with the wall. However, when the vapor volume fraction next to the wall becomes high enough, liquid access to the wall and its capability to remove heat from the wall are both restricted. This can occur unintentionally during a subcooled boiling simulation, as it passes through various intermediate states from initial conditions towards a converged solution.

Alternatively, the vapor contribution is relevant when studying the approach to DNB (Departure from Nucleate Boiling) conditions, as wall heat flux is increased sufficiently.

When liquid access is restricted, vapor heat transfer removes some fraction of the wall heat flux and causes an increase in wall temperature. The heat fluxes can be represented as:

$$\dot{q}''_w = (\dot{q}''_{conv} + \dot{q}''_{evap} + \dot{q}''_{quench})(1 - K_{dry}) + K_{dry}\dot{q}''_{dry} \tag{1992}$$

where:

- $\dot{q}''_{dry}$ is the vapor contribution to convective heat flux (see Eqn. (1998)), based on single-phase turbulent convection by the vapor.
- K_{dry} is the wall contact area fraction for the vapor (see Eqn. (2000)), based on either an expression for K_{dry}, or a transition volume fraction representing the start of wall dryout.

Multiphase Framework

If other phases are present in the flow, they are assumed to be not participating in the wall contact and the wall boiling process.

Only the liquid phase l and possibly vapor phase g are assumed to be in contact with the wall, as the wall contact area fraction K_{dry} indicates.

So the combined models for subcooled boiling with possible dryout (see Eqn. (1992)) fits into the multiphase energy balance at the wall (see Eqn. (1980)) as follows:

$$\dot{q}''_w = \dot{q}''_{wl} + \dot{q}''_{wlg} + \dot{q}''_{wg} \tag{1993}$$

where:

- $\dot{q}''_{wl}$ is the heat flux between the wall and the liquid phase:

$$\dot{q}''_{wl} = (1 - K_{dry})[(1 - K_{quench})\dot{q}''_{conv} + \dot{q}''_{quench}] \tag{1994}$$

- $\dot{q}''_{wlg}$ is the heat flux between the wall and liquid-gas interface:

$$\dot{q}''_{wlg} = (1 - K_{dry})\dot{q}''_{evap} \tag{1995}$$

- $\dot{q}''_{wg}$ is the heat flux between the wall and gas phase:

$$\dot{q}''_{wg} = K_{dry}\dot{q}''_{dry} \tag{1996}$$

As outlined in the sections on phase heat flux and phase interaction heat flux, you can override any of the standard models for $\dot{q}''_{wl}$, $\dot{q}''_{wlg}$, and $\dot{q}''_{wg}$ with your own models.

Convective Heat Flux

The turbulent convection fluxes are based on single-phase wall treatments. For liquid contact with the wall, this flux is:

$$\dot{q}''_{conv} = \frac{\rho_l c_{pl} u_l^*}{t_l^+}(T_w - T_l) \tag{1997}$$

For vapor contact with the wall, this flux is:

$$\dot{q}''_{dry} = \frac{\rho_g c_{pg} u_g^*}{t_g^+}(T_w - T_g) \tag{1998}$$

Note: The volume fraction does not modify the convection contributions as the contributions are in non-boiling heat transfer, where volume fraction represents the wall contact area fraction for each phase (compare this with Eqn. (1982)).

In the subcooled boiling model, it is assumed that the liquid is in full contact with the wall and is superheated first before vapor can be generated. However, if the wall begins to dryout, the factor, K_{dry}, can be specified to represent the wall contact area fraction for the vapor.

Wall Dryout Area Fraction

Weisman and Pei [372] postulated that in wall boiling, a bubbly layer exists near the wall and, at DNB (Departure from Nucleate Boiling), this layer reaches a critical bubble concentration. At this concentration, liquid transport is no longer sufficient to remove the heat flux that is applied at the wall. Weisman-Pei found that the maximum value of average volume fraction was 0.82 in a bubbly layer 5.5 departure diameters thick.

A basic application of this DNB criterion to a CFD calculation, is that:

- The CFD computation replaces Weisman-Pei's analysis of radial transport.

- The computed distribution of bubble volume fraction, averaged over some prescribed bubbly layer thickness, reaches some critical value.

Such a computation would not expect to use the same calibration constants as the original Weisman-Pei model. Also, recalibrate the computation if alternative models are selected for the bubble departure diameter.

In STAR-CCM+, there are three options for specifying the bubbly layer thickness for the dryout criterion:

- **Wall Cell** -- the layer thickness is the width of the grid cell next to the wall
- **Fixed Number of Diameters** -- the layer thickness is a factor of the number of bubble departure diameters
- **Fixed Yplus** -- the layer thickness is a factor of the wall turbulence scales for the vapor phase

An output field function, Bubbly Layer Volume Fraction is computed on the wall. This can be easily plotted for comparison on different grids and also used in user expressions for modeling wall dryout.

There is no justification for assuming a logarithmic profile exists for bubble volume fraction, $\alpha_g(y)$, near the wall when approaching DNB conditions. Instead the distribution is approximated by a one-term expansion about the first cell center at distance y_c from the wall, using the wall-normal component of the volume fraction derivative, α_g'. This distribution can then be averaged over a layer thickness δ.

$$\begin{aligned}\alpha_\delta &= \frac{1}{\delta}\int_0^\delta dy\{\alpha_g(y_c) + \alpha_g'(y_c)(y - y_c)\} \\ &= \alpha_g(y_c) + \alpha_g'(y_c)(\delta/2 - y_c)\end{aligned} \tag{1999}$$

The simplest model for local wall dryout is that heat transfer from the wall to the vapor begins when the vapor volume fraction exceeds a specified value. At the same time, heat transfer from the wall to the liquid convection and liquid evaporation declines. This is expressed as an effective wall contact area fraction for the vapor:

$$K_{dry} = \begin{cases} 0 & \alpha_\delta \le \alpha_{dry} \\ f(\beta) & \alpha_\delta > \alpha_{dry} \end{cases} \tag{2000}$$

where:

- α_{dry} is the wall dryout break-point or the volume fraction at which dryout begins

- α_δ is the vapor volume fraction averaged over the bubbly layer thickness.
- $f(\beta)$ is the function:

$$f(\beta) = \beta^2(3-2\beta) \tag{2001}$$

Which transitions symmetrically and smoothly from 0 through 1 over $0 \le \beta \le 1$, where β is the ratio:

$$\beta = \frac{\alpha_\delta - \alpha_{dry}}{1-\alpha_{dry}} \tag{2002}$$

Which varies from 0 through 1 for $\alpha_{dry} \le \alpha_\delta \le 1$,

α_δ is the vapor volume fraction averaged over the bubbly layer thickness.

The default value for α_{dry} is 0.9. This value prevents instabilities in the computation due to unintentional dryout, as it reduces the heat flux that is directed to the liquid as the amount of liquid reduces.

Weisman and Pei [372] geometrically derived a maximum vapor volume fraction of 0.82 for use in their critical heat flux model. Adjust the wall dryout break-point, α_{dry}, when studying the approach to DNB conditions.

Evaporative Heat Flux

An evaporative heat flux is constructed for the following nucleate boiling mechanism:

- Bubbles nucleate and grow on suitably sized cavities in the wall.
- Bubbles reach a critical size, at which point the forces holding the bubble to the wall are no longer balanced.

The evaporative heat flux is:

$$\dot{q}''_{evap} = n''f\left(\frac{\pi d_w^3}{6}\right)\rho_g h_{lg} \tag{2003}$$

where:

- n'' is the nucleation site number density
- f is the bubble departure frequency
- d_w is the bubble departure diameter
- ρ_g is the vapor phase density
- h_{lg} is the latent heat

The evaporative heat flux, $\dot{m}''_{evap}$, is the rate of conversion of liquid to vapor per unit wall area:

$$\dot{m}''_{evap} = \frac{\dot{q}''_{evap}}{h_{lg}} \tag{2004}$$

Nucleation Site Number Density

A nucleation site is a gas filled cavity in the wall. In practice, there is a range of cavity sizes. A larger wall superheat allows bubble growth in more cavities of smaller size.

Two standard submodels have been implemented for the nucleation site number density:

- Lemmert Chawla – A correlation against wall superheat, and is often used together with the Tolubinsky Kostanchuk model for bubble departure diameter.
- Hibiki Ishii – A mechanistic model for active nucleation site number density as a function of critical cavity size and contact angle. It is designed for used with the Kocamustafaogullari model for bubble departure diameter, and is applicable to up 198 bar.

The Lemmert Chawla model is the default option, as its fixed power for superheat dependence means that it is easier to obtain initial solutions. The more recent Hibiki Ishii model has a wider range of applicability, but it can be difficult to reach a converged solution in high flux cases.

Under the 45 bar, 0.57 MW/m^2 conditions of the first test in Bartolomei and Chanturiya [316], these two methods (each used with its recommended departure diameter model) can produce comparable results for axial void and for wall superheat. However, the departure diameter that is associated with the Hibiki Ishii and Kocamustafaogullari solution is much smaller.

At higher pressures and higher heat fluxes, the Hibiki Ishii model has an increasingly stronger dependence on wall superheat. This model often predicts much a smaller wall superheat than the Lemmert Chawla model for the same heat flux.

Lemmert Chawla Method

The original Lemmert Chawla model [353][346] for nucleation site number density is:

$$\frac{n''}{[m^{-2}]} = (m\Delta T_{sup})^p \tag{2005}$$

where:

- n'' is the nucleation site density
- m is a calibration constant, with a default value of 185.0/K
- p is a superheat exponent, with a default value of 1.805
- ΔT_{sup} is the wall superheat, which is calculated as:

$$\Delta T_{sup} = \min\left(\max\left(T_{wall} - T_{sat}, 0\right), \Delta T_{max}\right) \tag{2006}$$

where:

- ΔT_{max} is the maximum superheat applied to the Lemmert Chawla model

STAR-CCM+ uses a modified form of the Lemmert Chawla equations. This modification allows that the superheat exponent does not have to be a constant but can itself increase significantly with wall superheat.

$$\frac{n''}{n_0''} = \left(\frac{\Delta T_{sup}}{\Delta T_0}\right)^{A + B(\Delta T/\Delta T_0)} \tag{2007}$$

where:

- A and B are calibration constants
- n_0'' is the reference nucleation site number density, this value is from $(185.0/\mathrm{K})^{1.805}$.
- ΔT_0 is the reference wall superheat

You can calibrate the Lemmert Chawla Method that is suitable for use with the Kocamustafaogullari bubble departure diameters by comparing it with the Hibiki-Ishii Method, at a given operating pressure, and for a useful range of wall superheats. For example, the table below shows calibrations of the Lemmert Chawla method for the properties of water assuming a contact angle of 0.722 radians:

P (bar)	Superheat range (K)	n_0'' (m^{-2})	ΔT_0 (K)	A	B	Fit Accuracy
45	1 - 10	138,100	1	1.348	0.35	7%
70	1 - 7	302,300	1	1.566	0.58	10%
150	1 - 3	5,665,000	1	3.095	1.7	0.1%

Using a fitted Lemmert Chawla Method instead of using the Hibiki Ishii Method gives similar results. The main advantage of such a fit is that it indicates the sensitivity of either method to wall superheat as an explicit

superheat exponent. For example, using the calibrations above, the table below shows that even at 45 bar, the nucleation site number density can vary as strongly as ΔT^5:

P (bar)	Superheat range (K)	Exponent range
45	1 - 10	1.7 - 4.8
70	1 - 7	2.2 - 5.6
150	1 - 3	4.8 - 8.2

Hibiki Ishii Method

For high wall superheats the Hibiki Ishii model can produce extreme values of wall heat fluxes. To improve robustness, a limit is applied to the amount of superheat that can be applied in this submodel:

$$\Delta T_{sup} = \min\left(\max\left(T_{wall} - T_{sat}, 0\right), \Delta T_{max}\right) \tag{2008}$$

The default value for the maximum superheat ΔT_{max} is 25K. Do not allow this limit to affect the converged solution. If necessary, adjust the limit during the calculation.

Also, use system reference pressure in place of local liquid pressure to improve robustness.

$$P_l \approx P_{REF} \tag{2009}$$

Following Hibiki and Ishii [334], the temperature of the vapor next to the wall is identified with the wall temperature.

$$T_g \approx T_{sat} + \Delta T_{sup} \tag{2010}$$

For the contact angle between the working fluid and the surface during boiling, θ, Hibiki and Ishii used a value at room temperature. They suggest that a temperature dependence of contact angle is implicitly included in the calibration of the pressure effects function $f(\rho^+)$.

The nucleation site number density is then modeled as:

$$n'' = \overline{n''}\left(1 - \exp\left[-\frac{\theta^2}{8\mu^2}\right]\right)\left(\exp\left[f(\rho^+)\frac{\lambda'}{R_c}\right] - 1\right) \tag{2011}$$

where:

- $\overline{n''}$ is the average cavity density, with a default value of 4.74 E5 sites/m^{-2}

- θ is the wall contact angle, with a default value of 0.722 radians (41.37^o), which is suitable for water systems
- μ is the wall contact angle scale, with a standard value of 0.722 radians (41.37^o)
- λ' is the cavity length scale, with default value 2.50E-6 m
- $f(\rho^+)$ is the logarithmic, non-dimensional density function representing the effect of pressure (see Eqn. (2013))
- R_c is the critical cavity radius (see Eqn. (2012))

The critical cavity radius is:

$$R_c = \frac{2\sigma(1 + (\rho_g/\rho_l)/P_l)}{\exp[h_{lg}(T_g - T_{sat})/RT_gT_{sat}] - 1} \tag{2012}$$

where:

- σ is the surface tension
- ρ_g is the vapor phase density
- ρ_l is the liquid phase density
- P_l is the liquid pressure
- h_{lg} is the latent heat
- R is the gas constant for the given vapor molecular weight.

The logarithmic, non-dimensional density function is:

$$f(\rho^+) = C_0 + C_1\rho^+ + C_2\rho^{+2} + C_3\rho^{+3} \tag{2013}$$

where:

$$\rho^+ = \log_{10}(\rho^*) \tag{2014}$$

and:

$$\rho^* = \frac{\rho_l - \rho_g}{\rho_g} \tag{2015}$$

The standard values for the density function constants are:

$$\begin{aligned} C_0 &= -0.01064 \\ C_1 &= 0.48246 \\ C_2 &= -0.22712 \\ C_3 &= 0.05468 \end{aligned} \tag{2016}$$

Bubble Departure Diameter

There are two standard methods that are implemented for calculating the bubble departure diameter:

- Tolubinsky Kostanchuk – Correlation against liquid subcooling.
- Kocamustafaogullari – Correlation that is based on bubble force balance at departure, and calibrated using steam-water data between 0.067 and 141.87 bar.

You can investigate various combinations of standard and your own submodels, but the following two combinations can both give reasonable estimates for evaporation rates:

- The Tolubinsky Kostanchuk Departure Diameter model that is used with the Lemmert Chawla Nucleation Site Number Density model.
- The Kocamustafaogullari Departure Diameter model that is used with the Hibiki Ishii Nucleation Site Number Density model.

The default option is to use Tolubinsky Kostanchuk for departure diameter, as appropriate for the default submodel for nucleation site number density.

Tolubinsky Kostanchuk Method

The Tolubinsky Kostanchuk model [364] correlates bubble departure diameter against liquid subcooling:

$$d_w = d_0 \exp\left[-\frac{\Delta T_{sub}}{\Delta T_0}\right] \tag{2017}$$

where:

- d_0 is the reference diameter with default value 0.0006 m
- ΔT_0 is the reference subcooling with default value 45K
- ΔT_{sub} is the subcooling of the liquid next to the wall:

$$\Delta T_{sub} = T_{sat} - T_l \tag{2018}$$

The minimum and maximum bubble departure diameters can be set with the Minimum Diameter and Maximum Diameter properties respectively. The defaults for these limits are 0.000025 m and 0.0014 m (25 micron and 1.4 mm).

Kocamustafaogullari Method

As noted in Hibiki and Ishii [334], it is usual to use a room temperature value for contact angle between the working fluid and surface. The actual temperature dependence is covered elsewhere in the overall model calibration.

The Kocamustafaogullari model [341] provides a correlation for departure diameter that is based on a force balance between gravity and surface tension forces, but modified by a density ratio, which makes it useful for a wide range of pressures. This model is calculated as follows:

$$d_w = d_1\theta\left(\frac{\sigma}{g\Delta\rho}\right)^{0.5}\left(\frac{\Delta\rho}{\rho_g}\right)^{0.9} \tag{2019}$$

where:

- d_1 is a calibration constant, with a default value of $1.5126\text{x}10^{-3}$ m/radian ($2.64\text{x}10^{-5}$ m/degree)
- θ is a wall contact angle for the working fluid, with a default value of 0.722 radians (41.36^o), which is suitable for water systems
- σ is the surface tension
- g is the acceleration due to gravity
- $\Delta\rho$ is the difference in density between the liquid and gas phase
- ρ_g is the vapor phase density

Bubble Departure Frequency

The standard model that is implemented for bubble departure frequency is the Cole model [322]. This is equivalent to taking a typical bubble rise velocity (estimated using unit drag coefficient) as the velocity scale, over bubble diameter d_w as the length scale:

$$f = \sqrt{\frac{4}{3}\frac{g(\rho_l - \rho_g)}{d_w\rho_l}} \tag{2020}$$

where:

- g is acceleration due to gravity
- ρ_l is the liquid phase density

Bubble Induced Quenching Heat Flux

When a bubble detaches from the wall, the space it occupied is replaced by cooler liquid. Quenching heat transfer is the component of heat flux that is used in heating this replacement liquid, as modeled by Del Valle and

Kenning [324].

The quenching heat flux is:

$$\dot{q}''_{quench} = h_{quench}(T_{wall} - T_{quench}) \tag{2021}$$

where:

- h_{quench} is the quenching heat transfer coefficient (see Eqn. (2022))
- T_{wall} is the wall temperature
- T_{quench} is the temperature at which liquid is brought to the wall by the departure of a bubble

When using the Del Valle Kenning model, the quenching heat transfer coefficient is:

$$h_{quench} = 2K_{quench} f \sqrt{\frac{\rho_l Cp_l k_l t_w}{\pi}} \tag{2022}$$

where:

- K_{quench} is the bubble influence wall area fraction (see Eqn. (2026))
- f is the bubble departure frequency (see Eqn. (2020))
- ρ_l is the liquid density
- Cp_l liquid specific heat
- k_l is the liquid conductivity
- t_w is the waiting time between bubble departure and the nucleation of the next bubble:

$$t_w = \frac{C_w}{f} \tag{2023}$$

where:

- C_w is the wait coefficient. The default value is 0.8. This value comes from an assumption by Kurul and Podowski [343] that quenching occurs between the departure of one bubble and the nucleation of next. This period is 80% of the departure cycle.

The Bubble Induced Quenching Heat Flux is implemented using the temperature difference between the wall and the liquid next to the wall, T_l, together with a slowly varying correction factor, s_{quench}:

$$q''_{quench} = h_{quench} s_{quench}(T_{wall} - T_l) \tag{2024}$$

The correction factor is defined as

$$s_{quench} = \frac{T_{wall} - T_{quench}}{T_{wall} - T_l} = \frac{T^+(y_{quench})}{T^+(y_c)} \tag{2025}$$

where:

- $T^+(y)$ is the dimensionless temperature profile that is formed by scaling wall heat flux to the liquid using liquid wall turbulence scales.
- y_c is the distance from the wall to the nearest cell center.
- y_{quench} is the representative distance from the wall at which cold liquid is drawn by the departure of a bubble. You specify this distance in terms of bubble departure diameters or of liquid wall turbulence scales.

Bubble Influence Wall Area Fraction

When an individual bubble departs from the wall with a diameter d_w, subcooled liquid flows in to fill the space underneath the detached bubble. The wall area influenced by this quenching flow is larger than the basic footprint of the bubble $\pi d_w^2/4$.

The standard model that is implemented for the bubble influence wall area fraction follows Kurul Podowski [343]:

$$K_{quench} = F_A \frac{\pi d_w^2}{4} n'' \tag{2026}$$

where:

- F_A is an area coefficient for scaling between the nucleation site area density and the wall area fraction the bubble-induced quenching influences.
- n'' is the nucleation site number density.

The default value for F_A is 2.0 and is chosen from experience with the main 45-bar test case in Bartolomei and Chanturiya [316].

Multi-Component Mass Transfer

If a phase is chosen to be a multi-component gas, liquid, or solid (particle), then the segregated species solver can be selected for that phase. This solver allows multiple species transport equations to be solved:

$$\frac{\delta \rho_j \alpha_j Y_{j,i}}{\delta t} + \nabla \cdot (\rho_j \alpha_j Y_{j,i} u_j) = \nabla \cdot (\rho_j \alpha_j D_j \nabla Y_{j,i}) + S_{j,i} + \dot{m}_{j,i} \tag{2027}$$

where:

- ρ_j, α_j and u_j are the density, volume fraction and velocity of phase j
- $Y_{j,i}$ is the mass fraction of species i in phase j, that is, the fraction of the total mass of that phase
- D_j is the mass diffusivity; combining molecular and turbulent diffusivity
- $S_{j,i}$ is a general mass source
- $\dot{m}_{j,i}$ is used here to refer specifically to the transfer of species from one phase to another.

Note: If there is a single phase with α_j equal to 1 and no interphase mass transfer, this equation reverts to the species transport equation for a single phase.

Similar to a single phase, boundary, and initial conditions are required for all species for all multi-component phases.

Dissolution Mass Transfer

For Eulerian multi-component multiphase flows, transfer between the components of each phase can occur. The Dissolution Mass Transfer model includes the case where one phase is liquid and one phase is gaseous, and the case where both phases are liquid.

The rate of this mass transfer is assumed to be slow and independent of the rate of transfers between any other species. Diffusion driven by the concentration gradient is assumed to be the mechanism that moves species between the phase bulk region and the phase boundary. This movement can occur in either direction.

Gas-Liquid Interactions

For gas-liquid interactions, the equilibrium between species at the interface is defined as:

$$X^*_{gas,i} = K_i^0 X^*_{liq,i} \tag{2028}$$

Where $X^*_{gas,i}$ and $X^*_{liq,i}$ are the equilibrium mole fractions of component i at the gas side and liquid side of the interface respectively, and K_i^0 is the equilibrium constant. With this equilibrium definition, the value of K_i^0 is independent of the choice of the continuous and dispersed phase.

The equilibrium constant can be specified as using the following techniques

- Constant value
- Field function
- Raoult's Law
- Henry's Law

The rate of transfer of the species from one phase to the other is governed by the rate at which the species can diffuse between the phase boundary and the interior of the phase.

For gas-liquid interactions, the species mass balance at the interface is:

$$\dot{m}_{l,i} + \dot{m}_{g,i} = 0 \tag{2029}$$

$$\dot{m}_{l,i} = k_l a_{lg} \rho_l (Y^*_{l,i} - Y_{l,i}) \tag{2030}$$

$$\dot{m}_{g,i} = k_g a_{lg} \rho_g (Y^*_{g,i} - Y_{g,i}) \tag{2031}$$

where:

The subscript l refers to the liquid phase and g to the gas

k is a mass transfer coefficient for either phase

a_{lg} is the interfacial area density

ρ is the phase density

$Y_{l,i}$ and $Y_{g,i}$ are the mass fractions of the two phases

Y^* refers to the equilibrium mass fraction at the interface of the two phases.

Liquid-Liquid Interactions

For liquid-liquid interactions, the equilibrium between species at the interface is defined as:

$$X^*_{cont,i} = K_i^0 X^*_{disp,i} \tag{2032}$$

Where $X^*_{cont,i}$ and $X^*_{disp,i}$ are the continuous phase and dispersed phase equilibrium mole fractions at the interface respectively. In this case, the value of the equilibrium constant K_i^0 depends on the choice of the

continuous and dispersed phase. In particular, if the continuous and disperse phases are swapped, the value of K_i^0 (either constant or field function) becomes $1/K_i^0$.

The equilibrium constant can be specified as using the following techniques

- Constant value
- Field function

For liquid-liquid interactions, the species mass balance at the interface is:

$$\dot{m}_{c,i} + \dot{m}_{d,i} = 0 \tag{2033}$$

$$\dot{m}_{c,i} = k_c a_{cd} \rho_c (Y^*_{c,i} - Y_{c,i}) \tag{2034}$$

$$\dot{m}_{d,i} = k_d a_{cd} \rho_d (Y^*_{d,i} - Y_{d,i}) \tag{2035}$$

where the subscript c refers to the continuous phase and d to the dispersed phase. The meaning of the symbols is the same as in the liquid-gas interaction case above.

Mass Transfer Coefficient

The mass transfer coefficient, k, is found using Sherwood numbers for each phase:

$$k = Sh\frac{D}{l} \tag{2036}$$

where:

D is the species diffusivity within a phase and l is the phase interaction length scale.

A single Sherwood number is assumed to apply for the whole phase. The Sherwood number for a dispersed phase or continuous phase can be specified by a constant or by a field function. The Sherwood number for a continuous phase can also be calculated using a Ranz-Marshall method:

$$Sh = 2 + 0.6Re^{1/2}Sc^{1/3} \tag{2037}$$

where:

Re is the phase pair Reynolds Number

Sc, the Schmidt number, is $\mu/(\rho D)$.

Raoult's Law

Raoult's law states that the equilibrium partial pressure of a component above the surface of a liquid is equal to the vapor pressure of a pure sample of the component multiplied by the mole fraction of the component in the liquid. Hence Raoult's law gives an equilibrium constant:

$$K_i^0 = \frac{p_{vap}}{p_{abs}} \tag{2038}$$

Raoult's Law is appropriate to volatile liquid species evaporating into or out of a gas.

Henry's Law

For cases where only small amounts of a species are able to dissolve into a liquid, for instance when a gas such as oxygen dissolves in trace amounts into water, Raoult's law no longer holds. In this case Henry's law can be used. Henry's law states that the equilibrium amount of a gas species that dissolves in a liquid is directly proportional to the partial pressure of that species in the gas:

$$K_i^0 = \frac{H_i^0}{p_{abs}} \tag{2039}$$

where:

H_i^0 is Henry's coefficient and is unique to any gas and liquid solvent pair. The default value of 4.2E9 Pa is appropriate for oxygen dissolving in water at room temperature.

S-Gamma Model for Particle Size Distribution

In the S-Gamma model transport equations are solved (in addition to volume fraction) for moments of the diameter size distribution:

$$S_\gamma = n\int_0^\infty d^\gamma P(d)d(d) \tag{2040}$$

where:

γ is the order of the moment

n is the number of particles per unit volume

d is the particle diameter

$P(d)$ is the probability density function of particle diameter [347], [348].

Moments and Mean Diameters

The zeroth-moment of the distribution, S_0, is the particle number density n, that is, the number of particles per unit volume.

The second-moment of the distribution, S_2, is related to the interfacial area density a_i:

$$a_i = \int_0^\infty n\pi d^2 P(d)d(d) = \pi S_2 \tag{2041}$$

The third-moment of the distribution, S_3, is related to the volume fraction of the particles α [348]:

$$\alpha = \int_0^\infty n\frac{\pi d^3}{6} P(d)d(d) = \frac{\pi S_3}{6} \tag{2042}$$

From the three equations that include α (volume fraction), S_2 and S_0, it is possible to derive two mean diameters [348]:

- Sauter mean diameter:

$$d_{32} = \frac{6\alpha}{\pi S_2} \tag{2043}$$

- Three-zero diameter:

$$d_{30} = \left(\frac{6\alpha}{\pi S_0}\right)^{1/3} \tag{2044}$$

If the particle size distribution is locally monodisperse, that is all the same size within a cell, then these two diameters are equal. If, however, there is a range of sizes, then the variance of particle size in log space σ^2 can be found as follows (assuming a log-normal distribution, [335]):

$$\sigma^2 = \ln\left(\frac{d_{32}}{d_{30}}\right) \tag{2045}$$

This variance, and the Sauter mean diameter, must be set for flow boundaries and initial conditions.

Transport Equations

Transport equations are solved for each of the three moments of the particle diameter size distribution. The transport equation for volume fraction (and hence S_3) is solved by the Multiphase Segregated Flow model. In One-Equation mode (given as a choice on the S-Gamma model node), only

S_2 is solved for in addition to the volume fraction. In Two-Equation mode, the equation for S_0 is also solved. Variable density does not play a role in the conservation of the number density of particles.

The transport equation for S_0 is [347], [348]:

$$\frac{\partial S_0}{\partial t} + \nabla \cdot (S_0 \mathbf{v}_d) = s_{br} + s_{cl} + s_{wb} \tag{2046}$$

$\mathbf{v}_d$ is the dispersed phase velocity that is produced by the multiphase velocity solver. s_{br} and s_{cl} are source terms representing the effect of breakup and coalescence of particles whose details are described elsewhere. s_{wb} only applies at wall boundaries when the wall boiling model is active, and accounts for the bubbles that are created on the wall. It is given as:

$$s_{wb} = \frac{dS_0}{dt} = \dot{n}A \tag{2047}$$

where:

A is the wall area of a cell

$\dot{n}$ is the local bubble creation rate

For S_2, Wei and Morel [371] give an equation for the interfacial area a_d of the dispersed phase in the form:

$$\frac{\partial a_d}{\partial t} + \nabla \cdot (a_d \mathbf{v}_d) = s_{a,\,br} + s_{a,\,cl} + s_{a,\,m} - \frac{2a_d}{3\rho}\left(\frac{\partial \rho}{\partial t} + \mathbf{v}_d \nabla \rho\right) \tag{2048}$$

$s_{a,\,br}$ and $s_{a,\,cl}$ represent the source of area due to breakup and coalescence respectively, and $s_{a,\,m}$ represents the source of area due to mass transfer between phases. The final term on the right-hand side represents the effect of variable density on interfacial area. By multiplying through by $\rho^{2/3}$, performing some algebraic manipulations, and finally dividing by π to convert to S_2, it can be shown that this equation is equivalent to:

$$\frac{\partial \rho^{2/3} S_2}{\partial t} + \nabla \cdot (\rho^{2/3} S_2 \mathbf{v}_d) = \rho^{2/3}(s_{a,\,br} + s_{a,\,cl} + s_{a,\,m}) \tag{2049}$$

This formulation is the equation that is solved in STAR-CCM+ for S_2. The mass transfer source term, $s_{a,\,m}$, is given for a mass transfer rate m_{cd} as:

$$s_{a,\,m} = \frac{dS_2}{dt} = \frac{2S_2 m_{cd}}{3\rho_d \alpha_d} + s_{a,\,wb} \tag{2050}$$

Where subscripts c and d represent the continuous and dispersed phases, respectively. When wall boiling is active, an additional source $s_{a,\,wb}$ is included, given as:

$$s_{a,\,wb} = \dot{n} A d_{dep}^2 \tag{2051}$$

where:

A is the wall area of a cell

$\dot{n}$ is the local bubble creation rate

d_{dep} is the local bubble departure diameter

S-Gamma Breakup Model Formulation

The S-Gamma equations (for S_0 and S_2) are solved with source terms added to model coalescence and breakup of bubbles and other particle types. Both these processes conserve the total volume of the dispersed phase and so no extra term is required in the volume fraction equation which is used together with S_0 and S_2 to give mean diameter.

Breakup

The source terms for breakup are found by integrating the effect of a bubble or droplet of size d breaking up over all diameters. The general form for the breakup is from Lo and Zhang [348]:

$$\frac{dS_\gamma}{dt} = \int_{d_{lo}}^{d_{hi}} nP(d)\frac{(N(d)^{1-\gamma/3}-1)}{\tau(d)} d^\gamma d(d) \tag{2052}$$

In the above N is the number of fragments, or child bubbles, produced when a bubble breaks down, and τ is the timescale for breakup. d_{lo} and d_{hi} are the upper and lower limits of a particular breakup regime. The implementation in STAR-CCM+ considers both viscous and inertial breakup regimes. For laminar flow, the lower limit for viscous breakup is the critical diameter d_{cr} and the upper limit is infinite. For turbulent flow, the upper limit is the Kolmogorov length scale L_k; if d_{cr} is greater than L_k then viscous breakup plays no role. For inertial breakup the integral is performed from the greater of L_k and d_{cr} to infinity. By writing $N(d)$ and $\tau(d)$ in the form:

$$N(d) = N' d^b \tag{2053}$$

$$\tau(d) = \tau' d^a \tag{2054}$$

The integral can be written in the form:

$$\frac{dS_\gamma}{dt} = \frac{N'^{1-\gamma/3}}{\tau'}\int_{d_{lo}}^{d_{hi}} nP(d)d^{\delta_1}d(d) - \frac{1}{\tau'}\int_{d_{lo}}^{d_{hi}} nP(d)d^{\delta_2}d(d) \tag{2055}$$

Where $\delta_1 = \gamma + b - b\gamma/3 - a$ and $\delta_2 = \gamma - a$. This form of the equation is from Hill [335].

Viscous Breakup

For viscous breakup, the critical diameter is given as [348], [335]:

$$d_{cr} = \frac{2\sigma\Omega_{cr}}{\mu_c\dot{\gamma}} \tag{2056}$$

where:

μ_c is the viscosity of the continuous phase

σ is the surface tension

$\dot{\gamma}$ is the local shear rate

Ω_{cr} is the critical capillary number [335]

The fragmentation number is assumed to be binary [348], that is:

$$N(d) = 2 \tag{2057}$$

The breakup timescale is from [348]:

$$\tau(d) = \frac{\mu_c d}{\sigma} f(\lambda) \tag{2058}$$

μ_c is the continuous phase dynamic viscosity, σ is the surface tension, and:

$$\log f(\lambda) = p_0 + p_1 \log(\lambda) + p_2 (\log(\lambda))^2 \tag{2059}$$

and $\lambda = mu_d/mu_c$ and the constants are given in [335].

Inertial Breakup

For inertial breakup, the critical bubble diameter and related values are calculated as follows [348], [335]:

$$d_{cr} = (1 + C_\alpha)\left(\frac{2\sigma We_{cr}}{\rho_c}\right)^{3/5} \varepsilon^{-2/5} \tag{2060}$$

$$\tau(d) = 2\pi k_{br}\left(\frac{3\rho_d + 2\rho_c}{192\sigma}\right)^{1/2} d^{3/2} \tag{2061}$$

$$N(d) = 2 \tag{2062}$$

where:

ρ_d and ρ_c are the dispersed and continuous phase densities

ε is the continuous phase turbulent dissipation rate

We_{cr} is the critical Weber number which you define with a default value of 0.25

C_α and k_{br} are constants with default values of 4.6 and 0.2, respectively.

S-Gamma Coalescence Model Formulation

Performing the double integral required for an exact evaluation of the coalescence term [348] is not practicable, so a simplified source term is developed based on an equivalent diameter d_{eq} for each moment, which here is taken to be equal to $d_{3\gamma}$. The source term that is given in [348] is:

$$\frac{dS_\gamma}{dt} = F_{cl}(2^{\gamma/3} - 2)\left(\frac{6\alpha}{\pi}\right)^2 k_{coll} u_{rel} P_{cl}(d_{eq}) d_{eq}^{\gamma-4} \tag{2063}$$

where:

F_{cl} is the calibration coefficient

k_{coll} is a collision rate

u_{rel} is the typical phase velocity difference

$P_{cl}(d_{eq})$ is the probability of collision leading to coalescence

To split the source term into a collision rate (in collisions per unit volume per unit time) and a coalescence probability, either of which can be replaced with functions you define. This formulation has been rearranged slightly to give:

$$\frac{dS_\gamma}{dt} = F_{cl}(2^{\gamma/3} - 2) K_{coll} P_{cl}(d_{eq}) d_{eq}^{\gamma} \tag{2064}$$

Where K_{coll} is the collision rate and is:

$$K_{coll} = k_{coll} u_{rel} d_{eq}^2 \left(\frac{6\alpha}{\pi d_{eq}^3}\right)^2 \tag{2065}$$

For a turbulent flow, viscous and inertial effects lead to coalescence and a source term of the form that is shown above is evaluated for each. By successively setting γ to 0 and 2, the source terms for the zeroth and second moment S-gamma transport equations are obtained.

Viscous Coalescence

For viscous coalescence, the collision rate is [348]:

$$K_{coll,\,v} = \left(\frac{8\pi}{3}\right)^{1/2}(\dot{\gamma}d_{eq})d_{eq}^2\left(\frac{6\alpha}{\pi d_{eq}^3}\right)^2 \tag{2066}$$

With $\dot{\gamma}$ being the shear rate. For a laminar flow, the shear rate is calculated using the velocity field and for a turbulent field it is $\dot{\gamma} = \sqrt{\varepsilon\rho_c/\mu_c}$ [347] which uses the turbulent dissipation rate, density and viscosity from the continuous phase. The probability is from [348]:

$$P_v(d_{eq}) = \exp(-t_d\dot{\gamma}) \tag{2067}$$

Physically, this is comparing the time that the bubbles interact with the time required for the film initially to separate and then drain away. Four models for drainage modes are from [347]. Starting from the model that gives the slowest drainage time (and hence lowest coalescence rate) toward the fastest, these models are called:

- Fully Immobile:

$$t_d = \frac{3\mu_c d_{eq}^2 F_i}{64\pi\sigma^2 h_{cr}^2} \tag{2068}$$

- Partially Mobile Interface (Short Collision time):

$$t_d = \frac{3}{2}\left(\frac{F_i d_{eq}^2\sqrt{\mu_d\rho_d}}{32\pi\sigma^2 h_{cr}}\right) \tag{2069}$$

- Partially Mobile Interface (Quasi-Steady state):

$$t_d = \frac{\pi\mu_d\sqrt{F_i}}{2h_{cr}}\left(\frac{d_{eq}}{4\pi\sigma}\right)^{3/2} \tag{2070}$$

- Fully Mobile:

$$t_d = \frac{3\mu_c d_{eq}}{4\sigma}\ln\left(\frac{d_{eq}}{8h_{cr}}\right) \tag{2071}$$

Where $F_i = (3/2)\pi\mu_c\dot{\gamma}d_{eq}^2$.

The critical film thickness is:

$$h_{cr} = \left(\frac{A_H d_{eq}}{24\pi\sigma}\right)^{1/3} \tag{2072}$$

where:

A_H is the Hamaker constant and its default value is 5×10^{-21}

σ is the surface tension

Reference [348] provides all equations except the fully immobile expression.

Inertial Coalescence

For inertial coalescence, the collision rate and probability are from [348]:

$$K_{coll,\,i} = \left(\frac{2\pi}{15}\right)^{1/2} (\varepsilon d_{eq})^{1/3} d_{eq}^2 \left(\frac{6\alpha}{\pi d_{eq}^3}\right)^2 \tag{2073}$$

$$P(d_{eq}) = \frac{\Phi_{max}}{\pi}\left(1 - \frac{k_{cl,\,2}^2 (We - We_0)^2}{\Phi_{max}^2}\right)^{1/2} \tag{2074}$$

Φ_{max}, the maximum phase difference is:

$$\Phi_{max} = (2h_0^2 \rho_c \sigma)/(We_0 \mu_d^2 d_{eq}) \tag{2075}$$

Also:

$$We_0 = 0.8 We_{cr} \tag{2076}$$

$$h_0 = 8.3 h_{cr} \tag{2077}$$

You can set We_{cr} and $k_{cl,\,2}$, the latter having a default value of 12.7.

Multiphase Segregated Flow Bibliography

[311] Antal, S. P., Lahey, R. T., and Flaherty, J. E. 1991. "Analysis of phase distribution in fully developed laminar bubbly two-phase flow", *Int. J. Multiphase Flow*, **7**, pp. 653–652.

[312] Arastoopour, H., Pakdel, P., and Adewumi, M. 1990. "Hydrodynamic analysis of dilute gas-solids flow in a vertical pipe", *Powder Technology*, **62**, pp. 163–170.

[313] Atapattu, D.D., Chhabra, R.P., and Uhlherr, P.H.T. 1995. "Creeping sphere motion in Herschel-Bulkley fluids: Flow field and drag", *Journal of Non-Newtonian Fluid Mechanics*, **59**, pp. 245–265.

[314] Auton, T.R., Hunt, J.C.R., and Prud'homme, M. 1988. "The force exerted on a body in inviscid unsteady non-uniform rotational flow", *J. Fluid Mech.*, **197**, pp. 241-257.

[315] Balachandar S. and Eaton J. K. 2010. "Turbulent dispersed multiphase flow", *Annu. Rev. Fluid Mech.*, **42**, pp 111-133.

[316] Bartolomei G.C. and Chanturiya V.M. 1967. "Experimental study of true void fraction when boiling subcooled water in vertical tubes", *Thermal Engineering*, **14**, pp 123-128.

[317] Benyahia, S., Syamlal, M. and O'Brien, T.J. Accesses Jan 2012. "Summary of mfix equations 2012-1", from URL https://mfix.netl.doe.gov/documentation/MFIXEquations2012-1.pdf

[318] Bouillard, J., Lyczkowski, R., and Gidaspow, D. 1989. "Porosity Distributions in a Fluidized Bed with an Immersed Obstacle", *AIChE J* **35**(6), pp. 908-922.

[319] Bozzano, G., and Dente, M. 2001. "Shape and terminal velocity of single bubble motion: a novel approach", *Computers and Chemical Engineering*, **25**, pp. 571-576.

[320] Burns, A. D., Frank, T., Hamill, I., and Shi, J-M. 2004. "The favre averaged drag model for turbulent dispersion in eulerian multi-phase flows", technical report, *ICMF*.

[321] Chhabra, R.P., 1992. "Bubbles, Drops and Particles in Non-Newtonian Fluids", *CRC Press*.

[322] Cole R. 1960. "A photographic study of pool boiling in the region of the critical heat flux", *AIChE J.*, **6**, pp. 533-542.

[323] Cook, T.L. and Harlow, F.L. 1984. "Virtual mass in multiphase flow", *Int. J. Multiphase Flow*, **10**, pp. 691-696.

[324] Del Valle M.V.H. and Kenning D.B.R. 1985. "Subcooled flow boiling at high heat flux", *Int. J. Heat Mass Transfer*, **28**, pp. 1907-1920.

[325] Deutsch,E., and Simonin, O. 1991. "Large eddy simulation applied to the motion of particles in stationary homogeneous fluid turbulence", *ASME FED*, **110**:35–42.

[326] Ding, J. and Gidaspow, D. 1990. "A bubbling fluidisation model using kinetic theory of granular flow", *AIChE Journal*, **36**(4), pp. 523–538.

[327] Drew, D. A., 2001. "A turbulent dispersion model for particles or bubbles", *J. Engineering Mathematics*, **41**:259–274, 2001.

[328] Elghobashi, S. E., and Abou-Arab, T. W., 1983. "A two-equation turbulence model for two-phase flows", *Phys. Fluids*, **26**(4):931–938.

[329] Esmaili, E. and Mahinpey, N. 2011. "Adjustment of drag coefficient correlations in three dimensional cfd simulation of gas-solid bubbling fluidized bed", *Advances in Engineering Software*, **42**, pp. 375-386.

[330] Frick, A.M. and Thompson, B.E. 1982. "Settling Velocities in Batch Sedimentation", *Oak Ridge National Laboratory*, ORNL/MIT358

[331] D. Gera, M. Syamlal, and T.J. O'Brien, 2002. *Int. J. Multiphase flow*, 30(4):419–428.

[332] Gidaspow, G. 1994. "Multiphase Flow and Fluidization - Continuum and Kinetic Theory Descriptions", *Academic Press*.

[333] Gosman, A.D., Issa, R.I., Lekakou, C., Looney, M.K., and Politis, S. 1992. "Multidimensional modelling of turbulent two-phase flows in stirred vessels", *AIChE Journal*, **38**(12), pp. 1946-1956.

[334] Hibiki T. and Ishii M., 2003. "Active nucleation site number density in boiling systems", *Int. J. Heat Mass Transfer*, **46**, pp. 2587-2601.

[335] Hill, D.P. 1998. "The Computer Simulation of Dispersed Two Phase Flows", *Ph.D. dissertation*, Imperial College, London.

[336] Hinze, J. O., 1959. *Turbulence*.

[337] Ishii, M. and Mishima, K. 1984. "Two-fluid model and hydrodynamic constitutive relations", *Nuclear Engineering Design*, **82**, pp. 107-126.

[338] Ishii, M. and Zuber, N. 1979. "Drag Coefficient and Relative Velocity in Bubbly, Droplet and Particulate Flows," *AiChe. Journal,* **25**, pp. 843-855.

[339] Issa, R. I., and Oliveira, P. J. 1994. "Numerical prediction of phase separation in two-phase flow through t-junctions", *Computers Fluids*, **23**(2), pp. 347-372.

[340] Johnson, P.C., and Jackson, R. 1987. "Frictional-collisional constitutive relations for granular materials, with application to plane shearing", *J. Fluid Mech.*, **176**, pp. 67-93.

[341] Kocamustafaogullari G., 1983. "Pressure dependence of bubble departure diameter for water", *Int. Comm. Heat Mass Transfer*, **10**, pp. 501-509.

[342] Krieger, I.M., and Dougherty, T.J. 1959. "A mechanism for non-Newtonian flow in suspensions of rigid spheres" *Transactions of the Society of Rheology*, **3**, pp. 137-152.

[343] Kurul, N. and Podowski, M.Z. 1990. "Multidimensional effects in sub-cooled boiling", *Proceedings of the Ninth Heat Transfer Conference*, Jerusalem.

[344] Lamb, H., 1945. *Hydrodynamics*, Dover Publications, New York, NY, USA.

[345] Lance, M., and Bateille, J., 1991. "Turbulence in liquid phase of a uniform bubbly air-water flow", *J. Fluid Mech.*, **222**, pp. 95-118.

[346] Lemmert, M. and Chawla, J.M. 1977. "Influence of flow velocity on surface boiling heat transfer coefficient", *Heat Transfer in Boiling*, E. Hahne and U. Grigull, Eds., Academic Press and Hemisphere, New York, NY, USA.

[347] Lo, S. and Rao, P. 2007. "Modelling of droplet breakup and coalescence in an oil-water pipeline", *6th Int. Conf. Multiphase Flow*, ICMF 2007, Leipzig, Germany.

[348] Lo, S. and Zhang, D. 2009. "Modelling of Break-up and Coalescence in Bubbly Two-Phase Flows", *J. Comp. Multiphase Flows*, **1**(1), pp. 23-38.

[349] Lockett, M.J. and Kirkpatrick R.D. 1975. "Ideal Bubbly Flow and Actual Flow in Bubble Columns", *Trans. Inst. Chem. Engrs*, **53**, pp. 267-273.

[350] Lun, C.K.K. Savage, D.J. Jefferey, S.B. and Chepurniy, N. 1984. “Kinetic theories for granular flow: inelastic particles in Couette flow and slightly inelastic particles in a general flowfield”, *Journal of Fluid Mechanics*, **140**, pp. 223–256.

[351] Miller, R.M., Singh, J.P., and Morris, J.F. 2009. "Suspension flow modeling for general geometries", *Chemical Engineering Science*, **64** (22) pp. 4597-4610

[352] Morris, J.F. and Boulay, F. 1999. "Curvilinear flows of noncolloidal suspensions: The role of normal stresses", *Journal of Rheology*, 43, p. 1213.

[353] Podowski, M.Z. and Podowski, R.M. 2009. “Mechanistic Multidimensional Modeling of Forced Convection Boiling Heat Transfer”, *Science and Technology of Nuclear Installations*, Article ID 387020, 10 pages.

[354] Politis, S., Issa, R. I., Looney, M. K., Gosman, A. D., Lekakou, C. 1992. “Multidimensional modeling of turbulent two-phase flows in stirred vessels“, *AlChE Journal*, **38**(12), pp. 1946-1956.

[355] Ranz, W.E., and Marshall, W.R. 1952. “Evaporation from drops--Part I and II”, *Chemical Engineering Progress*, **48**(3), p. 141

[356] Richardson, J, and Zaki, W. 1954. “Sedimentation and Fluidisation: Part 1”, *Trans. Inst. Chem. Engrs*, **32**, pp. 35-53.

[357] Rusche, H., Behzadi, A. and Issa, R. 2004. “Modelling of dispersed bubble and droplet flow at high phase fractions”, *Chemical Engineering Science*, **59**, pp. 759-770.

[358] Samuelsberg, A. and Hjertager, B.H. 1996. “Computational Modeling of Gas/Particle Flow in a Riser”, *AIChE J*, **42**(6), p. 1536.

[359] Schaeffer, G. 1987. “Instability in the evolution equations describing incompressible granular flow”, *J. Diff. Equ.*, **66**, pp. 19–50.

[360] Schiller, L., and Naumann, A. 1933. “Uber die grundlegenden Berechnungen bei der Schwerkraftaufbereitung”, *VDI Zeits*, **77**(12), pp. 318-320.

[361] Simonnet M., Gentric, G., Olms, E. and Midoux, N. 2007. “Experimentation determination of the drag coefficient in a swarm of bubbles”, *Chem. Eng. Sci.*, **62**, pp. 858-866.

[362] Taylor R., Krishna R., “Multicomponent Mass Transfer”, Wiley, 1993

[363] Thai-Van, D., Minier, J. P., Simonin, O., Freydier, P., and Olive, J., 1994. "Multidimensional two-fluid model computation of turbulent dispersed two-phase flows", *ASME FED*, **185**:277–291.

[364] Tolubinsky V.I. and Kostanchuk D.M. 1970. “Vapour bubbles growth rate and heat transfer intensity at subcooled water boiling”, Heat Transfer 1970, *Preprints of papers presented at the 4th International Heat Transfer Conference*, Paris, **5**, Paper No. B-2.8

[365] Tomiyama, A., Kataoka, I., Zun, I., and Sakaguchi, T. 1998. “Drag Coefficients of single bubbles under normal and micro gravity conditions”, *JSME International Journal*, Series B, **41**(2), pp. 472-479.

[366] Tomiyama, A., Tamai, H., Zun, I., and Hosokawa, S. 2002. “Transverse migration of single bubbles in simple shear flows”, *Chemical Engineering Science*, **57**, pp. 1849–1858.

[367] Troshko, A. A. and Hassan, Y. A. 2001. “A two-equation turbulence model of turbulent bubbly flows”, *International Journal of Multiphase Flow*, **27** pp. 1965-2000.

[368] van Wachem, B., Schouten, J.C., Krishna, R. and van den Bleek, C.M. 1998. "Eulerian simulations of bubbling behaviour in gas-solid fluidised beds", Computers Chem. Eng., **22**, Suppl. pp. S299–S306.

[369] van Wachem, B., Schouten, J.C., van den Bleek, C.M., Krishna, R. and Sinclair, J. L. 2001. "Comparative analysis of CFD models of dense gas-solid systems", *AIChE Journal*, **47**(5) pp. 1035–1051.

[370] Wang, D.M. 1994. "Modelling of bubbly flow in a sudden pipe expansion", Technical Report II-34, BRITE/EuRam Project BE-4098.

[371] Wei, Y. and Morel, C. 2004. "Volumetric interfacial area prediction in upward bubbly two-phase flow", *Int. J. Heat and Mass Transfer*, **47**, pp. 307–328.

[372] Weisman J. and Pei B.S. 1983. "Prediction of critical heat flux in flow boiling at low qualities", *Int. J. Heat Mass Transfer,* **26**, No. 10, pp. 1463-1477

[373] Witt, P.J. and Perry, J.H. 1996. "A Study in Multiphase Modelling of Fluidised Beds", *Computational Techniques and Applications: CTAC95*, R. May & A. Easton Eds., World Scientific, pp. 787-794.

[374] Yu, A. B., and Standish, N. 1987. "Porosity calculations of multi-component mixtures of spherical particles", *Powder Technology*, **52**, pp. 233-241.

[375] Zuber, N. 1964. "On the dispersed two-phase flow in the laminar flow regime", *Chemical Engineering Science*, **19**, pp. 897-917.

Using the Dispersed Multiphase Model

The Dispersed Multiphase (DMP) model simulates dispersed phases in a Eulerian manner. The Dispersed Multiphase model combines aspects of both the Lagrangian Multiphase (LMP) model and the Segregated (Eulerian) Multiphase (EMP) models:

- Like the EMP model, Dispersed Multiphase solves for the dispersed phases in a Eulerian manner.
- Like the LMP model, Dispersed Multiphase does not calculate the flow in the background medium (that is, the Physics Continuum).

The main difference between DMP and EMP is that the Eulerian phases in the DMP model represent only the phases that are dispersed. The continuous background phase is solved with typical single-phase models.

Dispersed Multiphase is strictly one-way coupled: the background flow can influence the dispersed phases (for example, through drag) but not the reverse. The only way in which the dispersed phases can influence the physics continuum is through impingement onto a liquid film.

For more information on modeling liquid film, see Using the Fluid Film Model.

The DMP model is particularly useful for airfoil icing applications. For an example of an airfoil icing simulation, see the Dispersed Multiphase tutorial in the Tutorial Guide.

The DMP model and the VOF model can be activated in the same simulation. The DMP model uses the VOF mixture as its background flow, and has no knowledge of the individual phases. The DMP model treats the mixture as a single phase and uses only the properties of the mixture.

For more information on the VOF model, see Using the VOF Multiphase Model.

This section covers the following:

- When to Use the Dispersed Multiphase Model
- Setting Up a Dispersed Multiphase Simulation
- Selecting the Dispersed Multiphase model
- Setting Up the Dispersed Phases
- Setting Up Phase Interactions
- Setting Boundary and Region Conditions
- Setting the Dispersed Multiphase Solver Parameters

- Dispersed Multiphase Formulation

When to Use the Dispersed Multiphase Model

Dispersed Multiphase (DMP) is an alternative to Segregated (Eulerian) Multiphase (EMP), rather than a replacement. Due to its superposing behavior, DMP is compatible with many existing single-phase models. DMP can be activated at the end of a simulation without the need to deactivate and reactivate the appropriate set of models. As DMP does not solve for the background flow, it is automatically compatible with both segregated and coupled solution of the background flow and energy.

The dispersed phase particles can exist in a background fluid that is a single- or multi-component gas or liquid. The particles themselves can be single-component gases or liquids.

Advantages of Using the DMP Model

The DMP model is useful for airfoil icing applications, for which it is more appropriate than either the EMP model or the LMP model. Using the LMP model for these applications would be time-consuming due to the large number of parcels that are required. The EMP model does not support impingement onto a liquid film, or a coupled solution of background flow and energy, which are often required in these applications.

The DMP model can also be useful in soiling applications.

Disadvantages of Using the DMP Model

The DMP model cannot use many of the complex physics models that are available with the EMP model or the LMP model, such as turbulent dispersion. The DMP model cannot model evaporation or other interactions with the background fluid.

The DMP model does not model interactions between droplets of the dispersed phase, or complicated wall interactions (such as splashing). The LMP framework provides models for these effects.

Setting Up a Dispersed Multiphase Simulation

To set up a simulation that uses the Dispersed Multiphase model, perform the following tasks:

1. Select the Dispersed Multiphase model

   See Selecting the Dispersed Multiphase model

2. Define the dispersed phases

   See Setting Up the Dispersed Phases
3. Set up phase interactions

   See Setting Up Phase Interactions
4. Set boundary and region conditions

   See Setting Boundary and Region Conditions
5. Set the Dispersed Multiphase solver parameters

   See Setting the Dispersed Multiphase Solver Parameters

The Dispersed Multiphase model can be used as a post-processing tool: you can clear the Dispersed Multiphase solution, leaving the background flow intact, and then recalculate the transport of the dispersed phases only. You can obtain a steady solution even when the entire simulation is Unsteady.

This feature can help you to perform a quick analysis of particle behavior, such as collection efficiencies. For example, if you are using a set of dispersed phases to model a range of particle diameters, you can repeat the calculation for each dispersed phase and sum the results.

Selecting the Dispersed Multiphase model

The Dispersed Multiphase model is activated from the *Physics Model Selection* dialog of the relevant physics continuum.

To select the Dispersed Multiphase model:

1. In the physics continuum, right-click the Models node and click Select Models.
2. In the *Material* group box, select a Liquid or Gas Material model.
3. In the *Flow* group box, select either the Coupled Flow model or the Segregated Flow model.

The Dispersed Multiphase option then appears in the *Optional Models* group box.



The remaining selections in this dialog depend on problem characteristics that are not related to dispersed phases.

When the Dispersed Multiphase option has been selected in the dialog, the corresponding Dispersed Multiphase model node becomes visible in the object tree.



Setting Up the Dispersed Phases

The Dispersed Multiphase model node contains a Eulerian Phase manager which allows you to define a number of dispersed phases. Each dispersed phase is a Eulerian phase, similar to the phases that are used in the Segregated Multiphase model, for which you specify the relevant models (such as a flow model, a heat transport model, and a material model).

You can use a number of dispersed phases to model particles of different materials, or particles that are of a single material but of different sizes. Each dispersed phase is independent: the phases do not interact.

To specify a dispersed phase:

1. Under the Dispersed Multiphase > Eulerian Phases node, create a phase.
2. In the new Phase node, open the *Phase Model Selection* dialog and select the appropriate material model: Gas or Liquid.
3. Select the appropriate optional models.

If the background fluid has an energy model, the dispersed phase also requires one. The only energy model option currently available is Segregated Fluid Temperature.



4. Click Close.

   The selected models are added under the Phase node.

The particle diameter is a constant for each phase and is specified as a material property.

To specify the particle diameter for a dispersed phase:

1. Under the Liquid or Gas node (the phase material model), select the Material Properties > Particle Diameter > Constant node.

2. In the *Properties* window, set the Value to the appropriate value.



Setting Up Phase Interactions

The interactions between the dispersed phase and the background (continuous phase), and between the dispersed phase and the fluid film, are set up as phase interactions. You set up these phase interactions under the Multiphase Interaction node.

In a typical scenario, the continuous phase interacts with the dispersed phase through drag force and heat transfer. The dispersed phase interacts with the fluid film phase through impingement.

For more information on setting up the fluid film model in your simulation, see Using the Fluid Film Model.

Dispersed Phase - Continuous Phase Interaction

To specify the interaction between the dispersed phase and the continuous phase:

1. Under the Multiphase Interaction > Phase Interactions node, create a phase interaction.
2. In the new Phase Interaction node, open the *Phase Interaction Model Selection* dialog and activate the Dispersed-Physics Continuum Interaction model.
3. Select the appropriate optional models.



4. Click Close.

   The selected models are added under the Phase Interaction node.
5. Select the Dispersed-Physics Continuum Interaction node and set the Dispersed Phase property to specify the appropriate phase.



For more information, see Setting Up the Dispersed Phases.

Dispersed Phase - Fluid Film Phase Interaction

To specify the interaction between the dispersed phase and the fluid film phase:

1. Under the Multiphase Interaction > Phase Interactions node, create a phase interaction.
2. In the new Phase Interaction node, open the *Phase Interaction Model Selection* dialog and activate the Film-Dispersed Phase Interaction model.
3. Select the appropriate optional models.



4. Click Close.

   The selected models are added under the Phase Interaction node.

5. Select the Film-Dispersed Phase Interaction node and set the Fluid Film Phase and the Dispersed Phase properties to specify the appropriate phases.



Setting Boundary and Region Conditions

The boundary conditions are typically set as phase conditions of the dispersed phase. The available methods to specify the velocity, volume fraction, and temperature depend on the type of boundary.

Velocity Specification

You can specify the velocity with the usual options (Components or Magnitude + Direction), or you can use the Background method, which is unique to the Dispersed Multiphase model.

The example below shows an Inlet boundary where the Velocity Specification has been set to Background.



The Background method allows the dispersed phase to take its velocity from the background flow; STAR-CCM+ does not add a Velocity node under the Physics Values. This method is useful in cases where the velocity is not known exactly, such as on Free Stream boundaries.

Region Conditions

You can set the mass, momentum, and energy source options for the dispersed phase in a region.



For more information, see the following sections:

- Mass Source Option
- Momentum Source Option
- Energy Source Option

Setting the Dispersed Multiphase Solver Parameters

When the Dispersed Multiphase model is selected, the Dispersed Multiphase solvers become available under the Dispersed Multiphase node.



The Dispersed Multiphase solver node contains sub-nodes (instead of properties) that let you customize solver settings for the following:

- Dispersed Phase Flow
- Dispersed Phase Energy

The dispersed multiphase solvers generally use the same techniques as the Segregated Flow solvers.

The dispersed multiphase solution is updated on every inner iteration.

Dispersed Phase Flow Solver

The Dispersed Phase Flow solver solves for the momentum and volume fraction.



For complex simulations, it can be useful to first let the flow converge before you activate the Dispersed Multiphase model. If the background flow is erratic and unconverged, the dispersed multiphase phases are also, and the simulation can diverge. To resolve this problem, you can freeze the Dispersed Phase Flow solver until the flow is stable.

Dispersed Phase Flow Solver Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.
Temporary Storage Retained	Retains or deletes the data that is in temporary storage at the end of the iteration.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.

Dispersed Phase Velocity Solver

The Dispersed Phase Velocity solver uses a form of the Segregated Flow Velocity solver for the fluid film. It solves the momentum conservation equation, Eqn. (2079). This node has properties and is subject to common menu items.



Dispersed Phase Velocity Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7.

Dispersed Phase Volume Fraction Solver

This solver, which has properties and is subject to common menu items, controls the solution update for the volume fractions of all the phases. The discretized form of the continuity equation Eqn. (2078), is solved for each dispersed phase.



Dispersed Phase Volume Fraction Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7.
Minimum Volume Fraction	Sets a lower limit on the value of volume fraction in a cell for each dispersed phase.
Maximum Volume Fraction	Sets an upper limit on the value of volume fraction in a cell for each dispersed phase.

Dispersed Phase Energy

The Dispersed Phase Energy solver, which has properties and is subject to common menu items, calculates the energy transport.



Dispersed Phase Energy Solver Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7.

The expert properties are the same as for the Dispersed Phase Flow solver.

Dispersed Multiphase Formulation

This section describes the mathematical formulation of the Dispersed Multiphase model. The model includes the following:

- Transport Equations
- Phase Interactions

Transport Equations

For each dispersed phase, STAR-CCM+ solves transport equations for the conservation of mass and momentum.

STAR-CCM+ can also solve for energy transport if required.

The continuity equation (for a phase i) is given as:

$$\frac{\partial}{\partial t}(\alpha_i \rho_i) + \nabla(\alpha_i \rho_i \mathbf{v}_i) = s_i^{\alpha} \tag{2078}$$

and the momentum equation is given as:

$$\frac{\partial}{\partial t}(\alpha_i \rho_i \mathbf{v}_i) + \nabla(\alpha_i \rho_i \mathbf{v}_i \mathbf{v}_i) = \bar{s}_i^{\,v} \tag{2079}$$

Eqn. (2078) is solved for as a volume fraction correction equation, and Eqn. (2079) is solved as a velocity correction equation.

These equations are essentially the same equations as are solved in Eulerian Multiphase (EMP) cases. However, the Dispersed Multiphase (DMP) model solves only for the dispersed phases, while Eulerian Multiphase solves for all phases. In the DMP method, the background phase is handled by a single phase flow solver.

Therefore, if you set up a particular simulation in both EMP and DMP, the EMP case has one more Eulerian phase than the DMP case (this extra phase is the background phase in DMP). In the EMP case, the volume fractions sum to unity. In the DMP case, the sum of volume fractions of the dispersed phases does not sum to unity, as the volume fraction of the background phase is not accounted for.

Phase Interactions

A phase interaction is defined as any process involving the dispersed phase that requires knowledge of the background continuous phase or a fluid film. For example, the gravity force $\rho\bar{g}$ is not a phase interaction, but the effect of buoyancy, (included through the pressure gradient force), is a phase interaction.

STAR-CCM+ can model the following phase interactions:

- Drag Force
- Heat Transfer
- Pressure Gradient Force
- Film Impingement

Drag Force

The drag force (N) acting on one single particle is defined as:

$$\mathbf{f}_d = \frac{1}{2} C_d \rho A_p |\mathbf{v}_s| \mathbf{v}_s \tag{2080}$$

where A_p is the projected area and $\mathbf{v}_s$ is the particle slip velocity. The continuous force $\mathbf{f}_d'$ (in N/m^3), acting on a phase i, is:

$$\mathbf{f}_{d,\,i}' = \frac{1}{2} C_{d,\,i} \left(\frac{6\alpha_i}{4D_i} \right) \rho_i |\mathbf{v}_{s,\,i}| \mathbf{v}_{s,\,i} \tag{2081}$$

where α_i is the volume fraction of the phase and D_i is the diameter of the particles.

Three methods are available for the calculation of the drag coefficient C_d:

- Constant
- Field Function
- Schiller-Naumann correlation, which is defined as:

$$C_D = \begin{cases} \frac{24}{Re_p}\left(1 + \frac{1}{6} Re_p^{2/3}\right), & Re_p \le 1000 \\ 0.424, & Re_p > 1000 \end{cases} \tag{2082}$$

with the Reynolds particle number defined as:

$$Re_p = \frac{\rho |\mathbf{v}_s| D}{\mu} \tag{2083}$$

The drag is one-way coupled: the background fluid does not feel the reacting force that the particles exert. The drag acts as a contribution to $\bar{s}_i^{\mathrm{v}}$ in Eqn. (2079).

Heat Transfer

Heat transfer causes the particles to heat or cool according to the temperature of their environment. For a single particle, the heat transfer (in W/m^2s) is defined as:

$$Q_c = h_c A(T_p - T_c) \tag{2084}$$

where h_c is the heat transfer coefficient, A is the area of the particle, and T_p and T_c are the particle and background temperatures, respectively. In a continuous approach, the heat transfer can be expressed as:

$$\overline{Q}_{c,i} = \bar{h}_{c,i}\left(\frac{6\alpha_i}{D_i}\right)(T_i - T_c) \tag{2085}$$

which is defined as positive for heat flowing from the particles to the background.

There are three methods available for the heat transfer coefficient:

- Constant
- Field Function
- Ranz-Marshall correlation, making use of the usual formulation of the heat transfer in terms of the Nusselt number:

$$h = \frac{kNu_p}{D} \tag{2086}$$

  and correlating the Nusselt number as:

$$Nu_p = 2(1 + 0.3Re_p^{1/2}\ Pr^{1/3}) \tag{2087}$$

The heat transfer is one-way coupled: the background fluid does not change temperature, even when the particles do.

Pressure Gradient Force

The pressure gradient force in a continuous sense (N/m^3) is defined as:

$$\bar{F}_p = -\nabla p_{static} \tag{2088}$$

This term contains the buoyancy force $-\rho_c g$, through the definition of the hydrostatic pressure as $p_{hs} = -\rho_c g h$. The pressure gradient force acts as a contribution to $\bar{s}_i^v$ in Eqn. (2079).

Film Impingement

Particles that reach the free-surface between the background fluid and a liquid film can be set to impinge onto the liquid film. All of the particle mass, energy, and momentum is transferred to the liquid film.

The mass source $(kg/m^2 s)$ due to impingement of a phase i is:

$$\dot{m}_{imp} = \rho_i \mathbf{v}_i \cdot A_{fs} \tag{2089}$$

which is essentially the outgoing mass flux. Corresponding sources are also applied to the momentum and energy transport equations.

Using the VOF Multiphase Model

This section describes:

- How the VOF multiphase model is used to solve problems involving immiscible fluid mixtures, free surfaces, and phase contact time.
- How to select a VOF multiphase simulation.
- How to define the composition of an immiscible mixture.
- How to define phase interactions in VOF simulations.
- How to set up initial, boundary, and region conditions for each phase.
- How to set up the VOF wall porosity model for single-component phases.
- How to model evaporation and condensation.
- How to model boiling.
- How to model gas dissolution.
- How to model cavitation.
- How to model melting and solidification.
- The mathematical formulation of the VOF model and related numerical techniques.

What Is the VOF Multiphase Model?

Volume Of Fluid (VOF) is a simple multiphase model. It is suited to simulating flows of several immiscible fluids on numerical grids capable of resolving the interface between the phases of the mixture, as shown below. In such cases, there is no need for extra modeling of inter-phase interaction, and the model assumption that all phases share velocity, pressure, and temperature fields becomes a discretization error.



Illustration of unsuitable grids (left) and suitable grid (right) for two-phase flows using the VOF model

Due to its numerical efficiency, the model is suited for simulations of flows where each phase constitutes a large structure, with a relatively small total contact area between phases. A good example of this type of flow is sloshing flow in a water tank, where the free surface always remains smooth. If the tank movement becomes pronounced, this results in breaking waves, large numbers of air bubbles in the water, and water droplets in the air. The method would then require a fine mesh (at least three cells across each droplet/bubble) to produce small modeling errors.

The spatial distribution of each phase at a given time is defined in terms of a variable that is called the volume fraction. A method of calculating such distributions is to solve a transport equation for the phase volume fraction (see Basic VOF Model Equations). The method uses the STAR-CCM+ Segregated Flow model.

In order to obtain sharp interfaces between the phases, use the 2nd-order discretization scheme and set the appropriate value of the Sharpening Factor parameter (described below). For more information, see The HRIC Convection Discretization Scheme.

Thermal effects can also be included with the Segregated Multi-Phase Temperature model.

Modeling Change of State

Melting of ice and boiling of water are everyday examples of transitions between solid, liquid, and gaseous states of matter. Depending on the conditions under which the change of state takes place, a number of models are available in STAR-CCM+ that provide a mathematical description of these phenomena.

The following sections describe how to work with the change of state:

- Boiling
- Cavitation
- Melting and solidification

The Volume of Fluid (VOF) node is activated when a VOF multiphase simulation is selected.



Volume of Fluid (VOF) Properties

Convection	For guidance on selecting a convection, see Diffusion Term.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.

Volume of Fluid (VOF) Expert Properties

CFL_l	The lower Courant number limit in the HRIC convection discretization scheme. This value is Cu_l in Eqn. (2156).
CFL_u	The upper Courant number limit in the HRIC convection discretization scheme. This value is Cu_u in Eqn. (2156).
Sharpening Factor	This factor is used to reduce numerical diffusion in the simulation. The valid values are 0.0 through 1.0. When the sharpening factor is set to 0.0, there is no reduction in numerical diffusion. This value is the default because, in most cases, the recommended 2nd-order discretization scheme is sufficient to achieve a sharp interface between the two phases. When the sharpening factor is set to 1.0, there is no numerical diffusion and, therefore, a very sharp interface is obtained. For some flows, such as surface-tension dominated flows or tank sloshing flows, it is recommended that you increase the sharpening factor to improve the resolution of the interface between the phases. However, this additional sharpening can result in a non-physical alignment of the free surface with the grid lines. If a non-zero value is specified, an additional term (Eqn. (2147)) is added to the VOF transport equation (Eqn. (2146)). This value is C_α in Eqn. (2147)
Angle Factor	The angle factor in the HRIC convection discretization scheme for interface tracking. This value is C_θ in Eqn. (2157).

What Is the Segregated Multi-Phase Temperature Model?

The Segregated Multi-Phase Temperature node, which has properties, lets you control thermal effects in your simulation provided you include it when you

select the VOF multiphase model.



Segregated Multi-Phase Temperature Properties

Convection	For guidance on selecting a convection, see Diffusion Term.	
	1st-order	Selects the first-order convection.
	2nd-order	Selects the second-order convection.

Segregated Multi-Phase Temperature Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Setting Up a VOF Multiphase Simulation

To use the VOF model for solving immiscible fluid problems:

1. In the *Physics Model Selection* dialog, select the Implicit Unsteady solver option from the *Time* group box.
2. Specify the material model as a multiphase fluid by selecting Eulerian Multiphase from the *Material* group box.
3. Select Volume of Fluid (VOF) from the *Eulerian Multiphase Model* group box.
4. Make the appropriate selections from the *Viscous Regime* group box and from the *Space* group box.

   The *Physics Model Selection* dialog then looks as shown below.



5. Make an appropriate selection from the *Optional Models* group box:
   - If gravity forces significantly influence the solution, select Gravity.
   - To include thermal effects in your simulation, select Segregated Multiphase Temperature.

Defining the Mixture Composition

The Eulerian Multiphase model is used for both the Multiphase Segregated Flow model and the Volume of Fluid model. This model controls the definition of Eulerian phases within the physics continuum. It contains no Eulerian phases in its initial state.

The procedure for creating and defining Eulerian phases is described in the section, Defining Eulerian Phases, Models, and Materials.

In a VOF multiphase simulation, a typical tree structure appearing under the Models node of a given continuum is as shown below.



STAR-CCM+ automatically selects the multiphase fluid model and sets up a default immiscible fluid mixture. You define the bulk properties and phase composition of the mixture in two separate stages.

Note the following:

- The most common situation is a two-phase, liquid-gas mixture.
- Default names, material properties, and (Constant) property values are assigned to materials retrieved from the database.

  If none of these materials correspond to the phases present in your problem, do the following:

  - Rename the appropriate materials by right-clicking its node in the tree.
  - Open the corresponding Properties node and specify new settings as necessary.
- The order in which phases are defined is not important in principle.

  However, if you want to model surface tension in the simulation at some point, you are advised to define the liquid phase as the first component, as shown above.
- For laminar flows, it is possible to select the Non-Newtonian Generalized Power Law for the dynamic viscosity of liquid materials in a VOF simulation.

Controlling Phase Interactions in VOF Simulations

This section covers how to create and define VOF phase interactions. Selecting the Volume of Fluid (VOF) model in the *Physics Model Selection* dialog automatically activates the Multiphase Interaction model and adds a corresponding node to the physics continuum models node. The VOF Phase Interaction model, which controls phase interactions in VOF simulations, is defined within this node.

The procedure for creating a phase interaction is described in the Creating a Phase Interaction section.

Selecting the VOF Phase Interaction Model

To select VOF phase interaction models:

1. Select VOF Phase Interaction from the *Phase Interaction Topology* group box.

Further VOF models are available in the *Optional Models* group, as shown below.



2. Select the appropriate models from the *Optional Models* group box, and then click Close.

Defining VOF Phase Interactions

The VOF Phase Interaction child node is used to define the phase interactions. This node appears in the Multiphase Interaction > Phase Interactions > Phase Interaction 1 > Models node, as shown below.



From here, you can set the Primary and Secondary Phase properties in the Properties window.

What Is Surface Tension?

Immiscibility of two fluids is a result of strong cohesion forces between their molecules and depends on the nature of the fluids. The experimentally determined surface tension coefficient σ expresses the ease with which the fluids can be mixed. σ is defined as the amount of work necessary to create

a unit area of free surface. Surface tension effects are included in a phase interaction when the Surface Tension Force model is chosen in the *Phase Interaction Model Selection* dialog.



Marangoni convection develops along a free surface when a concentration or temperature gradient exists across that interface. These gradients cause the surface tension coefficient to vary across the free surface. To model this effect, activate the Marangoni Convection property of the Surface Tension Force node.

These effects can be neglected if the spatial variation of surface tension coefficients is small. If constant surface tension coefficients are set for the phase interaction, activating this property has no effect.

Surface tension always exists between a given pair of fluids. Its magnitude depends on the nature of the fluids in contact as well as on temperature. For immiscible fluids, the value is always positive. The corresponding value node (Surface Tension in the following screenshot) appears as a material property for the phase interaction. It is entered as a scalar profile.



In the current implementation, each phase interaction is assigned its own surface tension coefficient. This coefficient is used to calculate the surface tension force between each of the defined phases in the phase interaction.

This implementation implies that, in general:

- The model works correctly in multiphase flows (more than two phases can be present).
- The surface tension coefficients that are assigned to each phase interaction can be different.

Calculation of the free surface curvature is sensitive to mesh quality. Therefore, to get good results even on coarse grids, the use of Cartesian meshes is recommended.

Surface Tension Force Properties

Marangoni Convection	Controls whether the tangential surface tension force Eqn. (2165) is taken into account.	
	Activated	Includes the effect of spatial variations in the surface tension coefficient.
	Deactivate	Excludes the effect of spatial variations in the surface tension coefficient.

Using the Contact Time Model

The Contact Time phase interaction model estimates how long two VOF phases have been in contact with each other. A typical application is estimating the probability of surface reactions such as oxidation. The transported quantity is the contact area multiplied by the contact time.



To include contact time in a phase interaction, choose the Contact Time model in the *Phase Interaction Model Selection* dialog.

Properties

None.

Contact Time Phase Interaction Model Expert Properties

Min. Interface Fraction	Contact time is only computed when interfacial area exceeds the specified fraction of the cube root of the cell volume.
Convection	For guidance on selecting a convection, see Convection Term. 1st-order — Selects the first-order convection. 2nd-order — Selects the second-order convection.
Min. Volume Fraction	Time is only accumulated if the volume fraction of the host phase exceeds the specified threshold.

See also:

Setting the Contact Time Condition Method

Contact Time Field Functions Reference

Setting the Contact Time Condition Method

When you activate the Contact Time phase interaction model, you can specify the contact time condition as a scalar profile.

Contact time is accumulated where both phases are in contact and the selected Method returns a nonzero Value.

In the following screenshot, the example shows this object available in the Multiphase Interaction > Phase Interactions > Tin_Air > Models > Contact Time node.



For example, if you want to determine the contact time between an air and a melt phase when the melt phase is in the liquid state, you can use a field function defined as:

```
($SolidVolumeFractionMelt < 1) 1 : 0.
```

Once the solid volume fraction of the melt reaches unity in the simulation, the field function returns zero and the contact time area stops accumulating.

A transport equation is solved to accumulate Contact Time Area, which is the Contact Area multiplied with the Contact Time of the phase. The field function Contact Time gives the local value of the accumulated Contact Time Area divided by the current value of the Contact Area. Hence, it is only an approximation of the actual contact time.

Contact Time Field Functions Reference

The following primitive field functions are made available to the simulation when the Contact Time phase interaction model is used. They have a typical set of base properties.

Contact Area of *Phase*	The contact area of the phase.
Contact Time Area of *Phase*	The contact area multiplied by the contact time of the phase.
Contact Time of *Phase*	The contact time of the phase, which equals Contact Time Area divided by Contact Area. This is only an estimation; its accuracy depends on the simulation. For example, consider a fluid rising in a vertical tube that has a decreasing cross-sectional area. When the area is initially large, the contact area multiplied by contact time is also large. At the end, the free surface of the fluid reaches the height with a small cross-sectional area. If the accumulated value of contact area multiplied by contact time is divided by the recent smaller cross-sectional area, the calculated contact time can exceed the simulation time.

Using the Interface Momentum Dissipation Model

The Interface Momentum Dissipation model adds extra momentum dissipation in the proximity of the free surface to dissipate parasitic currents. The momentum dissipation is proportional to the velocity gradients and the interface artificial viscosity.

Parasitic currents are caused by discretization errors that arise due to solution discontinuities across the free-surface interface. An interface between phases represents a discontinuity in the solution domain. At this interface, all material properties, (for example, density, viscosity, and thermal conductivity) are discontinuous. All gradients of dependent variables, (for example, gradients of velocity, pressure, and temperature), are also discontinuous. In the case of flows that are dominated by surface tension, pressure, too, is discontinuous across the interface.

Discontinuities in solution fields can cause discretization errors which do not scale with the mesh size. In the case of free-surface flows, these errors can cause some spurious or parasitic currents. In many cases these currents are negligible relative to the other transport processes that are present and they do not affect the solution. However, as the parasitic currents can affect the solution in some cases, the Interface Momentum Dissipation model can be used to reduce them.

You access the Interface Momentum Dissipation Model from the *Phase Interaction Model Selection* dialog for the relevant pair of phases.



The Interface Artificial Viscosity node provides a choice between constant or field function for the method that defines the artificial viscosity.

Larger values of interface artificial viscosity are better at removing parasitic currents, but they can affect the physics of the problem on a particular mesh. Ensure that you specify values of interface artificial viscosity which are sufficient to remove or reduce parasitic currents, but are not large enough to affect the quality of the results.

What Is the Segregated VOF Solver?

The Segregated VOF Solver, which has properties and is subject to common menu items, controls the solution update for the phase volume fractions. More specifically, it solves the discretized volume-fraction conservation equation for each phase present in the flow.



Properties for Segregated VOF Solver

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7. For the theoretical background, see Implicit Iteration, p. 6362.

Solver Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and memory usage. It is recommended that you do not change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Temporary Storage Retained	Retains or deletes the data that is in temporary storage at the end of the iteration.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Update reconstruction gradients with each iteration.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.	

Setting Initial Conditions

As the VOF model uses multiple Eulerian phases, it is necessary to set the initial volume fraction of each phase. This procedure is described in Setting Initial Volume Fraction.

Setting Region Conditions

Following the selection of the VOF model, the node of every region that accesses this continuum contains an additional Phase Source Option sub-node.



The Phase Source is a region source. It specifies the rate at which a unit volume of the phase is added or removed per unit volume of the region (Volume/Time/Volume). In the Physics Conditions of a region, you can set up user-specified phase sources by activating the Phase Source Term property.

Setting Phase Sources

Once the Phase Source Option is activated, you can specify Phase Sources of each of the defined phases with the options of this scalar array profile:

- Option Constant defines sources uniformly for all phases.



- Option Composite defines the source on each phase separately.



Setting Boundary Conditions

As the VOF model uses multiple Eulerian phases, it is necessary to set the volume fraction of each phase at appropriate boundaries. This procedure is described in the section, Setting Volume Fraction At Boundaries.

If surface tension effects are included in the simulation, all **closed** (wall) boundaries also contain an additional Contact Angle sub-node. This node lets you specify the capillarity effect at each of these boundaries.

What Is Capillarity?

Capillarity is a phenomenon in which the free surface is observed to be elevated or depressed when it comes into contact with a solid wall. Capillarity can be explained by considering the effects of two opposing forces:

- Adhesion, the attractive (or repulsive) force between the fluid and solid wall molecules
- Cohesion, the attractive force between the fluid molecules.

What Is a Contact Angle?

The contact angle θ describes the influence of a solid wall on the free surface. This value node (Contact Angle in the following screenshot) is activated when the Surface Tension option is selected in the *Phase Interaction Model Selection* dialog and after the Phase Interaction has been defined. It is entered as a scalar profile.



Its magnitude depends on the pair of fluids and the solid in contact and on temperature. In the current implementation, it is assumed that the specified value is with respect to the **Primary Phase** of the respective VOF Phase Interaction. Values of contact angle smaller than 90° mean that the phase is wetting the wall, as illustrated part (a) of the following Figure.



Free surface in contact with a solid wall

In order to comply with the usual convention on contact angle values, that are shown above, define the liquid phase as the first of the components of the mixture.

Working with Wall Porosity

The wall porosity model lets you define phase-specific porous boundaries within a VOF Multiphase simulation. That is, it permits material that belongs to single-component phases to leave through selected boundaries.

For each phase, you can control whether it leaves though a particular boundary or not. The Wall Porosity node represents this model in the simulation tree.



Wall Porosity Expert Property

Max. Mass Extraction Factor	Specifies the maximum mass fraction that can be removed from a cell. The default is 0.5. This property prevents the model from extracting more mass from a cell than is available in that particular cell. It is recommended that you do not change the value of this property unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

The following sections describe:

- Selecting the wall porosity model
- Setting up a case with wall porosity

The mathematical formulation and field functions reference for the wall porosity model are also available.

Selecting the Wall Porosity Model

To use STAR-CCM+ for solving problems that involve VOF phase-specific wall porosity:

1. Select the physics models that are required for a VOF multiphase simulation.
2. Set up a Eulerian phase.
3. In the *Phase Model Selection* dialog, choose Wall Porosity in the *Optional Models* group box.



4. Click Close.

   The Wall Porosity node appears in the simulation tree.



5. If necessary, set the appropriate Max. Mass Extraction Factor.

Note: Ensure that at least one phase of the multiphase mixture is compressible when using the wall porosity model. Otherwise, mass extraction in a completely enclosed region (for example enclosed by wall boundaries or flow-stopped cells) leads to infinite negative absolute pressures causing a floating point exception.

Setting up the Wall Porosity Model

To define the phase-specific wall porosity, several parameters must be set.

- Ambient pressure
- Porous inertial resistance coefficient
- Viscous inertial resistance coefficient

Once the wall porosity model has been selected, Phase Conditions nodes appear within the regions and boundaries that are associated with the physics continuum. If a baffle interface exists in the model, a Phase Conditions node is also added underneath the Interfaces node.



Do the following:

- To specify the ambient pressure for the region, use the region Phase Conditions node.

  See Setting Ambient Pressure for Wall Porosity Model

- To specify the porous inertial and viscous resistance coefficients for the porous wall boundaries and contact interfaces, use the boundary Phase Conditions node.

  See Setting Wall Porosity Coefficients at Walls and Contact Interfaces

- To specify the porous inertial and viscous resistance coefficients for the porous baffle interface, use the Interfaces > Phase Conditions node.

  See Setting Wall Porosity Coefficients at Porous Baffle Interfaces

The porous inertial resistance coefficient α and the porous viscous coefficient β are added to all the wall boundaries in regions for which the wall porosity model is active. The wall porosity model has no effect in either of the following scenarios:

- The porous inertial and viscous resistance coefficients are both zero (the default values).
- The ambient pressure is set to the Pressure field function, hence forcing Δp to be always zero.

Setting Ambient Pressure for Wall Porosity Model

The ambient pressure is used to calculate the pressure jump Δp (Eqn. (2090)) from which the phase extraction mass is calculated. Ambient pressure is entered as a scalar profile. The ambient pressure is specified relative to the reference pressure. It is not an absolute value.

To set the ambient pressure for a region:

- Select Phase Conditions > [Phase] > Physics Values > Ambient Pressure and set its scalar profile method and value.



Setting Wall Porosity Coefficients at Walls and Contact Interfaces

The porous inertial resistance coefficient α and the porous viscous coefficient β determine the pressure drop across the porous boundary (see Eqn. (2090)).

To set the wall porosity coefficient values at a phase-specific porous wall boundary or contact interface:

1. Select the Regions > [Region] > Boundaries > [Boundary] > Phase Conditions > [Phase] > Physics Conditions > Wall Porosity Specification node.

2. In the *Wall Porosity Specification - Properties* window, set the Method to Specified.



A Physics Values node is added to the simulation tree.

3. Open the Physics Values node and select the Porous Inertial Resistance node. Set the scalar profile method and the value for the porous inertial resistance coefficient.



4. Similarly, select the Viscous Inertial Resistance node and set the scalar profile method and the value for the porous viscous resistance coefficient.

Setting Wall Porosity Coefficients at Porous Baffle Interfaces

The porous inertial resistance coefficient α and the porous viscous coefficient β determine the pressure drop across the porous boundary (see Eqn. (2090)).

To set the wall porosity coefficient values at a phase-specific porous wall boundary or contact interface:

1. Select the Interfaces > [Interface] > Phase Conditions > [Phase] > Physics Conditions > Wall Porosity Specification node.

2. In the *Wall Porosity Specification - Properties* window, set the Method to Specified.



A Phase Conditions node is added under Regions > [Region] > Boundaries > [Interface].



3. Open the Phase Conditions node and select the [Phase] > Physics Values > Porous Inertial Resistance sub node. Set the scalar profile method and the value for the porous inertial resistance coefficient.



4. Similarly, select the Viscous Inertial Resistance node and set the scalar profile method and the value for the porous viscous resistance coefficient.

Wall Porosity Formulation

The implementation of the VOF phase-specific wall porosity model is analogous to that of a porous baffle. The velocity across the porous boundary is calculated from the known pressure drop across it.

This pressure drop is given as:

$$\Delta p = -\rho(\alpha|v_n| + \beta)v_n \tag{2090}$$

where ρ is the density of the phase, α and β are user supplied constants and v_n is the normal velocity across the porous boundary. STAR-CCM+ calculates the pressure drop from a specified external ambient pressure on the region. Eqn. (2090) is then applied to calculate the normal velocity. This normal velocity is used to predict the amount of mass leaving through the boundary face within a given time-step.

Wall Porosity Field Function Reference

The following field functions are available for every Eulerian single-component gas phase with the wall porosity model selected. These objects all share base properties.

VOF Porous Wall Flux	The actual mass flow per unit area through the porous wall (kg/m^2s).
VOF Porous Wall Target Normal Velocity	The target velocity normal to the porous wall v_n calculated according to (Eqn. (2090)) and the pressure drop Δp. This function is not the actual normal velocity, which is equal to VOF Porous Wall Flux divided by the phase density.

Modeling Evaporation/Condensation

The particular phase-change type that is described in this section concerns evaporation and condensation and describes:

- What are evaporation and condensation?
- Setting up a simulation with evaporation and condensation
  - Defining evaporation-condensation mixture components
  - Defining phase interactions
- The mathematical formulation of the evaporation-condensation model
- VOF Evaporation/Condensation field functions reference

What Are Evaporation and Condensation?

Evaporation is a form of vaporization, the phase change from liquid to gas, which occurs when liquid molecules or atoms with high enough kinetic energy escape through the interface between the two phases.

Condensation is the reverse process of vaporization and is the phase change from gas to liquid.



The VOF: Evaporation and Condensation model deals with evaporation (without boiling) and condensation taking place at the free surface interface between a liquid and gas phase within the context of the Volume Of Fluid method. Both phases are modeled as multi-component mixtures that can include non-interacting (inert) species [376], [377]. Evaporation and condensation are treated as hydrodynamically limited, meaning that:

- The phases are considered to be in equilibrium at the interface, and
- The driving force is the diffusion of species.

The grid must be sufficiently fine to model the transport of species away from and toward the interface accurately.

Raoult's law is used to describe the phase equilibrium.

The Evaporation/Condensation node represents the VOF Evaporation-Condensation model and is subject to common menu items.



Evaporation/Condensation Properties

Connectivity	Identifies the evaporating/condensing pairs in the liquid and the gas phases.

Setting Up a Simulation with Evaporation and Condensation

To use STAR-CCM+ for solving problems that involve evaporation and condensation:

1. Select the physics models that are required for a VOF multiphase simulation.
2. Define the multi-component Eulerian phases.
3. Select the Evaporation/Condensation model in a phase interaction.

Defining Evaporation and Condensation Mixture Components

For a simulation with evaporation and condensation in STAR-CCM+, create two Eulerian Phases. The primary Eulerian Phase must be a liquid, and the other phase must be a gas. Both Eulerian Phases must be multi-component phases.

1. Create a Eulerian Phase, and select Multi-Component Gas from the *Material* group.
2. Select Segregated Species from the *Optional Models* group and then select an appropriate model from the *Equation of State* group.
3. Right-click the Multi-Component Gas > Gas Components node and choose Select Mixture Components...
4. In the *Select Mixture Components* dialog which appears, select the gas components required. Here, Air, Heptane, and Octane are selected.



5. Repeat the creation steps for a second phase where evaporation is likely to occur, but select Multi-Component Liquid as the material.

Select only the liquid components in the *Select Mixture Components* dialog. In this example, select Heptane and Octane only.

The simulation tree appears as shown below.



The saturation pressure is available from the Component Properties of each liquid component. The latent heat of vaporization is incorporated automatically through the heat of formation property, available for both gas and liquid components.

The final step in setting up an evaporation/condensation case is to define the evaporation/condensation phase interactions.

Selecting the Evaporation/Condensation Model from a Phase Interaction

To select the Evaporation and Condensation model:

1. Create a phase interaction involving both Eulerian phases.
2. Select the VOF Phase Interaction model in the *Phase Interaction Model Selection* dialog.
3. Set the properties of the VOF Phase Interaction node.

   The Primary Phase must be the multi-component liquid and the Secondary Phase must be the multi-component gas. The phases must be assigned before the Evaporation/Condensation model becomes available in STAR-CCM+.
4. Reopen the *Phase Interaction Model Selection* dialog.
5. Choose the Evaporation/Condensation model and click Close.



The Evaporation/Condensation node appears in the simulation tree.



Set Component pairs using the properties of the Evaporation/Condensation model node:

1. Select the Evaporation/Condensation model node.
2. Click the ellipsis to the right of the Connectivity property in the *Properties* window.

3. In the *Connectivity* dialog that appears, match the liquid and vapor phase components using the drop-down menus provided



Evaporation-Condensation Formulation

The evaporation rate (kg/s m^2) for a particular component i is:

$$\dot{m}_i = -\frac{\rho_G D_{G,i} \left.\frac{\partial Y_{G,i}}{\partial n}\right|_s}{1 - \sum_{j=1}^{N_v} Y_{G,j}^s} \tag{2091}$$

The denominator expresses the influence of the Stefan flow and Y^s indicates the component mass fraction at the surface. N_v is the total number of interacting components, that is, components undergoing evaporation or condensation. The G subscript denotes a gas phase variable.

The vapor pressure p_i of a component is:

$$\tilde{p}_i = a_i p_i^* \tag{2092}$$

where p_i^* is the vapor pressure of the pure component i and a_i the activity, which is dependent on the liquid mixture. This expression can be written as:

$$\tilde{p}_i = \gamma_i X_{L,i} p_i^* \tag{2093}$$

Raoult's law states that $\gamma_i \approx 1$, an approximation that is used here. Given the vapor pressures of the components, their molar fractions at the interface can be derived from:

$$X_{G,i}^s = \frac{\tilde{p}_i}{p} \tag{2094}$$

where p is the ambient pressure. The conversion to interface vapor mass fraction is done through:

$$Y^s_{G,i} = \frac{X^s_i W_i}{\sum_{j=1}^{N_v} X^s_j W_j + \sum_{j=1}^{N_{G,p}} X^s_j W_j} \tag{2095}$$

where N_G is the total number of components in the gas phase. While the molar fraction of the evaporation components is known, it is unknown for the inert species (which do not undergo evaporation or condensation). A good approximation is the introduction of a background molar weight calculated a short distance away from the interface:

$$W_{bg} = \frac{\sum_{j=1}^{N_{G,p}} X_j W_j}{\sum_{j=1}^{N_{G,p}} X_j} \tag{2096}$$

The interfacial background molar fraction is:

$$X^s_{bg} = 1 - \sum_{j=1}^{N_v} X^s_j \tag{2097}$$

Therefore the interfacial vapor mass fraction can be approximated as:

$$Y^s_{G,i} \approx \frac{X^s_i W_i}{\sum_{j=1}^{N_v} X^s_j W_j + X^s_{bg} W_{bg}} \tag{2098}$$

The numerical solution of Eqn. (2091) begins from the definition of the following scalar function ϕ_i, for every interacting component:

$$\phi_i = \alpha_L Y^s_{G,i} + \alpha_G Y_{G,i} \tag{2099}$$

where α_L is the liquid volume fraction, and α_G is the gas volume fraction. The approximated evaporation rate $\dot{M}_{i,c}$ for one cell can be obtained as:

$$\dot{M}'_{i,c} \approx -\frac{\rho_G D_{G,i} \nabla\phi_i \nabla\alpha_L V_c}{1 - \sum_{j=1}^{N_v} Y^s_{G,j}} \tag{2100}$$

using the fact that the integral of $\nabla\alpha_L$, the gradient of liquid fraction, over a certain volume expresses the area of the interface it contains [378], [379].

Comparing Eqn. (2091) and Eqn. (2100), it is seen that $\nabla\phi_i$ expresses the effect of $\nabla Y_{G,i}$. A lowest order analysis shows that the relative error of the evaporation rate, in quasi-steady state, behaves as:

$$\varepsilon_{rel} \approx \frac{1}{H|\nabla\alpha|} \tag{2101}$$

where H expresses the thickness of the boundary layer above the interface.

As such, the accuracy depends on how well the location of the interface is determined in comparison with the boundary layer. When the interface is not smeared, $|\nabla\alpha|^{-1}$, the reciprocal of the magnitude of the gradient of volume fraction, is a measure of the mesh size. This condition highlights the need for a sufficiently fine grid, to model the transport of species away from and towards the interface accurately. Similar considerations are valid for gradients of the temperature and pressure near the interface.



The sources for species, phases, and continuity are, based on Eqn. (2100), distributed over those cells that make up the interface, that is with non-vanishing gradient of the volume fraction, as shown in the above graph depicting the typical behavior of α_L (solid line) and $-|\nabla\alpha_L|$ (dashed line) at a free surface interface, along the surface normal coordinate n.

In order to avoid undershoot or overshoot of these sources, $\dot{M}'_{i,c}$ is weighted along the interface according to:

$$\dot{M}_{i,c} = \begin{cases} 2\alpha_L \dot{M}'_{i,c} & \text{if } M'_{i,c} > 0 \\ 2\alpha_G \dot{M}'_{i,c} & \text{if } M'_{i,c} < 0 \end{cases} \tag{2102}$$

Evaporation-Condensation Bibliography

[376] J. Schlottke and B. Weigand. Direct numerical simulation of evaporating droplets. Journal of Computational Physics, 227(10):5215–5237, 2008.

[377] S.S. Sazhin. Advanced models of fuel droplet heating and evaporation. Progress in Energy and Combustion Science, 32:162–214, 2006.

[378] S. Hardt and F. Wondra. Evaporation model for interfacial flows based on a continuum-field representation of the source terms. Journal of Computational Physics, 227(10):5871–5859, 2008.
[379] J.U. Brackbill, D.B. Kothe, and C. Zemach. A continuum method for modeling surface tension. J. Comput. Phys., 100(2):335–354, 1992.

VOF Evaporation/Condensation Field Functions Reference

The following field functions are available for every interacting gas components: the evaporation rate, the evaporation latent heat, and the evaporation mass fraction. These objects all share base properties.

Evaporation Rate	The rate of evaporation in kg/s for every cell.
Evaporation Latent Heat	The latent heat of vaporization in J/kg if evaporation/condensation takes place in the diagnosed cell. Available when Temporary Storage Retained is selected for the Segregated VOF solver.
Evaporation Mass Fraction	The saturated vapor mass fraction when evaporation/condensation takes place in the diagnosed cell. Available when Temporary Storage Retained is selected for the Segregated VOF solver.

Modeling Boiling

The particular phase-change type that is described in this section concerns boiling and describes:

- What is boiling?
- Selecting a simulation with boiling
- Defining boiling mixture components
- Defining Phase Interactions
- What is the Rohsenow Boiling model?
- What is the Transition Boiling model?
- What is the HTCxArea?
- What is boiling temperature?
- What is latent heat of vaporization?
- What is the film-boiling vapor volume fraction?

- Setting boundary conditions
- The mathematical formulation of the boiling model

What Is Boiling?

Boiling is a rapid vaporization of a liquid. It typically takes place when a liquid is heated to the boiling point (saturation temperature of the liquid T_{sat}). Its saturation vapor pressure then becomes equal to or larger than the pressure of the surrounding liquid.

Vapor can form from a liquid in the following locations:

- At a vapor-liquid interface away from surfaces
- At a solid surface with pre-existing vapor or gas pockets
- In the bulk of the liquid due to density fluctuations

The VOF Boiling model is the gateway, or entry point to the boiling modeling capabilities of STAR-CCM+. The VOF Boiling node represents this model and is subject to common menu items.



If you are only interested in modeling the improved heat transfer due to boiling, you can simulate boiling with a single-component liquid.

Selecting a Simulation with Boiling

To use STAR-CCM+ for solving problems that involve boiling, you select the following models. You can then select the boiling models in the Multiphase Interaction > Phase Interactions node:

1. In the *Physics Model Selection* dialog, specify the material model as a multiphase fluid by selecting Eulerian Multiphase in the *Material* group box.
2. Select Volume of Fluid (VOF) from the *Eulerian Multiphase Model* group box.
3. Select the Implicit Unsteady solver option from the *Time* group box.
4. Make an appropriate selection from the *Viscous Regime* group box.
5. In the *Optional Models* group box, select Gravity and Segregated Multiphase Temperature.

The *Physics Model Selection* dialog then looks as shown below.



6. Make any further appropriate selections from the other optional physics model groups.

The next step is to define the boiling mixture components.

Defining Boiling Mixture Components

In a simulation with boiling, STAR-CCM+ automatically sets up a default immiscible fluid mixture. You define the bulk properties and phase composition of the mixture.

Note: Each phase must be single-component (that is, pure liquid and pure vapor). The VOF Boiling model does not support multi-component phases.

To define the boiling mixture components:

1. Create a Eulerian phase and select the liquid material model.

2. If necessary, replace the material, such as H2O in the example below.

   STAR-CCM+ reads the material properties from its material database.
3. Create a second Eulerian phase and select the gas material.
4. Replace the default material so that the second Eulerian phase represents the gaseous form of the liquid phase.



The continuum must consist of at least two phases, that is the boiling liquid and its vapor. You can add any number of additional phases to complete the physics set-up.

The boiling model has to know the boiling temperature for the liquid phase and the latent heat of vaporization. The latent heat of vaporization is the difference between the enthalpies of the vapor and the liquid at the saturation temperature.

The final step in setting up a boiling case is to define the boiling phase interactions.

What Is Boiling Temperature?

The boiling or saturation temperature T_{sat} of a substance is the temperature at which it can change its state from a liquid to a gas throughout the bulk of the liquid at a given pressure. It is only available once you have set up a boiling phase interaction.



Selecting the Boiling Models and Defining Phase Interactions

The *Phase Interaction Model Selection* dialog lets you select the boiling models: the VOF Phase Interaction model and the VOF Boiling model. The Rohsenow Boiling and Transition Boiling models are available for selection in the *Boiling Models* group.



The boiling models appear in the Multiphase Interaction > Phase Interactions > Phase Interaction 1 > Models node. The Transition Boiling model is shown below.



The phase interaction is defined in the *Properties* window of the VOF Phase Interaction node. The liquid phase corresponds to the Primary Phase and the gas phase corresponds to the Secondary Phase.

What Is the Rohsenow Boiling Model?

When Rohsenow Boiling is selected in a simulation, STAR-CCM+ creates the Rohsenow Boiling node in the simulation tree, as shown below.



To specify the Rohsenow Boiling model properties, use the *Properties* window of this node.

Rohsenow Boiling Model Properties

C_qw	The empirical coefficient C_{qw} in the Rohsenow expression for the wall heat flux, Eqn. (390). Varies with the liquid-surface combination.
C_ew	The fraction C_{ew} of the wall heat flux due to boiling that is used in creation of vapor bubbles. Used in Eqn. (391).
Sc_t	The turbulent Schmidt number that is associated with the turbulent diffusion of vapor bubbles.

Rohsenow Boiling Model Expert Properties

Resolve liquid-vapor interface	The boiling model assumes that the vapor bubbles are smaller than the mesh size and so the interface between the liquid and the vapor phase is not resolved. If, however, large clouds of vapor are expected, and the mesh is sufficiently fine that the interface between the liquid and the vapor phase can be resolved, activate this flag.	
	Activated	The liquid-vapor interface is resolved.
	Deactivated	The liquid-vapor interface is not resolved.
n_p	The Prandtl number exponent n_p (1.73 by default) in Eqn. (390).	
URF_whf	Under-relaxation factor for heat flux due to boiling, at walls with prescribed heat flux.	

What Is the Transition Boiling Model?

When Transition Boiling is selected in a simulation, STAR-CCM+ creates a

Transition Boiling node in the simulation tree, as shown below.



To specify the Transition Boiling model properties, use the *Properties* window of this node.

Transition Boiling Model Properties

C_ew	The fraction C_{ew} of the wall heat flux due to boiling that is used in creation of vapor bubbles. Used in Eqn. (391).
Sc_t	The turbulent Schmidt number that is associated with the turbulent diffusion of vapor bubbles.
q_max	The positive empirical constant q_{max} in the Transition Boiling model, Eqns. (2105) to (2107).
k1	The positive empirical constant K_1 in the Transition Boiling model, Eqns. (2105) to (2107).
delT1	The positive empirical constant ΔT_1 in the Transition Boiling model, Eqns. (2105) to (2107).
k2	The positive empirical constant K_2 in the Transition Boiling model, Eqns. (2105) to (2107).
delT2	The positive empirical constant ΔT_2 in the Transition Boiling model, Eqns. (2105) to (2107).
phi	The constant ϕ in the Transition Boiling model, Eqns. (2105) to (2107).

Transition Boiling Model Expert Properties

Resolve liquid-vapor interface	The boiling model assumes that the vapor bubbles are smaller than the mesh size and so the interface between the liquid and the vapor phase is not resolved. However, if large clouds of vapor are expected, and the mesh is sufficiently fine that the interface between the liquid and the vapor phase can be resolved, activate this flag.	
	Activated	The liquid-vapor interface is resolved.
	Deactivated	The liquid-vapor interface is not resolved.
URF_whf	Under-relaxation factor for heat flux due to boiling, at walls with prescribed heat flux.	

What is the HTCxArea?

The HTCxArea node has properties and appears as a child of the Rohsenow Boiling or Transition Boiling model node. It represents the heat transfer coefficient (W/m^2K) between the vapor bubbles and the surrounding liquid, which is multiplied by the specific contact area (contact area per unit

volume, m^2/m^3) between the two.



HTCxArea Properties

Method	Selects the method that is used to enter the HTCxArea coefficient.	
	Constant	The coefficient is entered as a constant in W/m^3K. A Constant node is added as a child to this node.
	Field Function	The coefficient is entered as a field function that returns a value in W/m^3K. A field function method node is added as a child to this node.

Setting Boundary Conditions

Either of the following objects becomes available when you work with a boiling model:

- Alpha film boiling—activated by the Rohsenow Boiling model
- Boiling scale factor—activated by the Transition Boiling model

What Is Alpha Film Boiling?

Alpha Film Boiling controls the vapor volume fraction that is associated with the start of film boiling. This value node (Alpha_filmBoiling in the following screenshot) is activated as soon as the Rohsenow Boiling option is selected in the *Phase Interaction Model Selection* dialog. It is entered as a scalar profile.



What Is the Boiling Scale Factor?

The scale factor S in Transition Boiling allows you to scale q_{max} up or down for a given boundary. This value node (BoilingScaleFactor in the following screenshot) is activated as soon as the Transition Boiling option is selected in the *Phase Interaction Model Selection* dialog. It is entered as a scalar profile.



What Is the Film Boiling Vapor Volume Fraction?

The film-boiling vapor volume fraction $\alpha_{filmBoiling}$ is the value of volume fraction of the vapor phase which indicates full transition to film boiling. If the VOF boiling model is active, you can specify $\alpha_{filmBoiling}$ for each wall boundary. A value of $\alpha_{filmBoiling}$ larger than 1 implies that there is no film boiling along these wall boundaries.

For values of volume fraction less than $\alpha_{filmBoiling}$, a transition function is:

$$f(\alpha_{vapor}) = 1 - \min(1, \alpha_{vapor}/\alpha_{filmBoiling})^8 \tag{2103}$$

This function controls the fraction of boiling at the wall that is considered to be nucleate boiling. Here α_{vapor} is the vapor volume fraction at the wall.

Boiling Formulation

This section describes:

- Condensation and boiling at a vapor-liquid interface away from surfaces
- Boiling at a solid surface

- The Rohsenow Boiling model
- The Transition Boiling model

Condensation and Boiling at a Vapor-Liquid Interface Away from Surfaces

The heat transfer to the interface between the phases is used to compute the rate of evaporation or condensation $\dot{m}_{ec}$. The following are assumed:

- The vapor bubbles are on the saturation temperature T_{sat}
- The temperature of the liquid T_l is approximately the temperature of the mixture T
- All the heat flux from the liquid to the interface is used in mass transfer (that is, evaporation or condensation):

$$\dot{m}_{ec} = \frac{\mathrm{HTCxArea}(T - T_{sat})}{h_{lat}} \tag{2104}$$

  where h_{lat} is the latent heat of vaporization. $\mathrm{HTCxArea}$ is the heat transfer coefficient between the vapor bubbles and the surrounding liquid, which is multiplied by the specific contact area (contact area per unit volume) between the two.

Boiling at a Solid Surface

When a liquid is in contact with a solid surface maintained at temperatures T_{wall} above the saturation temperature of the liquid, boiling eventually occurs at that liquid-solid interface. In this case, boiling occurs in three characteristic stages:

- Nucleate boiling involves a creation and growth of vapor bubbles on a heated surface, which rise from discrete points on a surface. The temperature of the surface is only slightly above the saturation temperature of the liquid. In general, the number of nucleation sites increases with an increasing surface temperature. An increased surface roughness can create more nucleation sites, while an exceptionally smooth surface can result in superheating.
- Film boiling occurs when the critical heat flux is exceeded and a continuous vapor film covers the heated surface. The vapor layer has a lower thermal conductivity so the vapor layer typically “insulates” the surface.
- Transition boiling occurs at surface temperatures between the maximum attainable in nucleate and the minimum attainable in film boiling.

There are two distinct choices for modeling boiling in STAR-CCM+, namely the Rohsenow and Transition Boiling models. The former uses the Rohsenow nucleate boiling model applicable for boiling at relatively low solid temperatures. A film boiling model is also integrated into the Rohsenow implementation in STAR-CCM+ to facilitate its use with boiling at high solid temperatures. The Transition Boiling model has expressions for nucleate and transition boiling. The boiling model that used to be available in previous versions of STAR-CCM+ has now been formally named the Rohsenow Boiling Model. Once VOF Boiling is chosen in the continuum model selection, you can choose between the two models.

The Rohsenow Boiling Model

This model works with nucleate and film boiling.

The Nucleate Boiling Model

See Single-Phase Boiling Formulation.

The Film Boiling Model

At a sufficiently high wall temperature, a layer of vapor is created above the heated wall. As a result the liquid is no longer in contact with the heated wall. The evaporation occurs at the vapor liquid interface instead of at the heated wall. The vapor film that is created between the liquid and the heated wall acts as an insulator, which slows down the heat transfer process considerably. This process is known as film boiling.

The film boiling model, which is used in a combination with the VOF multiphase model, assumes that the film thickness is resolved in one or more cells. That is, the vapor occupies one or more cells close to the heated wall. In this case, the expressions for the wall heat flux are the same as in the case of single-phase flows. If it is not possible to have a mesh which resolves the film thickness, you specify the value of the vapor volume fraction $\alpha_{filmBoiling}$, which indicates when the film boil starts. For values of the vapor volume fraction smaller than $\alpha_{filmBoiling}$ the nucleate boiling is assumed, and for larger values it is assumed that the film boiling takes place. If the film is resolved and the cell next to the heated wall is filled with vapor, the results do not depend on values of $\alpha_{filmBoiling}$ from 0 through 1.

The Transition Boiling Model

You can specify the constants of this boiling model. The equations are as follows:

$$q_{boiling}(\Delta T) = q_{\max} S \phi \left(\frac{\Delta T}{\Delta T_1}\right)^{K_1} \qquad 0 \le \Delta T \le \Delta T_1 \tag{2105}$$

$$q_{boiling}(\Delta T) = q_{\max} S\left[1 - 4(1-\phi)\left(\frac{\Delta T - \Delta T_{\max}}{\Delta T_2 - \Delta T_1}\right)^2\right] \qquad \Delta T_1 \le \Delta T \le \Delta T_2 \quad (2106)$$

$$q_{boiling}(\Delta T) = q_{\max} S\phi\left(\frac{\Delta T}{\Delta T_2}\right)^{-K_2} \qquad \Delta T_2 \le \Delta T \quad (2107)$$

The first of the above equations simulates nucleate boiling, while the remaining two simulate transition I and transition II boiling. The three regimes are shown graphically in the following figure.



Different regimes of the new boiling transition boiling model

There are five empirical constants in the above model, namely, K_1, K_2, ΔT_1, ΔT_2, and $q_{\max}$, all of which are positive. ΔT_{max} is $\frac{1}{2}(\Delta T_1 + \Delta T_2)$. The constant ϕ is set equal to 0.75. Further, ΔT_2 must be greater than ΔT_1. (In the above expressions K_2 is positive even though it is a negative exponent because the sign is accounted for in the implementation.) The scale factor S in the above expressions is not treated as an independent sixth constant, as it allows you to scale $q_{\max}$ up or down for a given boundary. All the constants are likely to be application-specific. Hence, it is recommended that default values for the constants be tuned with credible measurements. The default values are from fits of boiling experiments of Ellion [381] using an apparatus with a particular geometry.

As mentioned in the previous section, $\alpha_{filmBoiling}$ (that you specify in the input) is the metric used to toggle between the nucleate and film boiling regimes in the Rohsenow Boiling Model. However, in the Transition Boiling Model, $\alpha_{filmBoiling}$ is not used; instead ΔT_1 and ΔT_2 are the metrics that switch between the three expressions (or regimes).

The total heat flux at the wall is the sum of the heat fluxes due to boiling, convection of the liquid-vapor mixture and thermal radiation. In the above equations, $q_{\max}$ is the maximum of the boiling heat flux component; it is not the critical heat flux which is a maximum of the total heat flux. The point on the curve corresponding to the minimum of the total heat flux is the Leidenfrost point and it occurs in the Transition II boiling regime. At the Leidenfrost point, vapor completely covers the surface and heat transfer is primarily driven by mixture convection and radiation with a negligible boiling contribution.

In the CFD implementation of the above boiling model, the boiling heat flux contribution to the total heat flux is considered to be zero when there is no liquid to boil in any cell that is next to the wall. Further, if the fluid temperature T_f is above the saturation temperature T_{sat} the boiling heat flux reduces by the factor $(T_w - T_f)/(T_w - T_{sat})$.

Boiling Bibliography

[380] Butterworth, D. and Hewitt, G.F. 1974. "Two-Phase Flow and Heat Transfer", Oxford University Press, Oxford.

[381] Ellion, M.E. 1953. "A study of the mechanism of boiling heat transfer", Ph.D. Dissertation, California Institute of Technology.

[382] Rohsenow, W.M. 1952. "A Method of Correlation Heat Transfer Data for Surface Boiling of Liquid", *Trans. ASME*, **74**, 969.

[383] Saiz Jabardo, J.M., Fockink da Silva, E., Ribatski, G. and de Barros, S.F. 2004. "Evaluation of the Rohsenow Correlation Through Experimental Pool Boiling of Halocarbon Refrigerants on Cylindrical Surfaces", *J. of the Braz. Soc. of Mech. Sci. & Eng*, **XXVI**, No. 2.

Modeling Gas Dissolution

The particular phase-change type concerns gas dissolution. This section describes:

- What is Gas Dissolution?
- How to set up a gas dissolution simulation.
- How to define Gas Dissolution model parameters.
- How to set the seed properties for the dissolving fluid.
- How to set up initial and boundary conditions for each fluid phase.

- The mathematical formulation of the Gas Dissolution model.

What is Gas Dissolution?

Liquid can dissolve a gas up to a certain maximal concentration (saturation). When there are changes in the pressure or temperature (or both), the liquid can become over-saturated or under-saturated. In the over-saturated case, gas bubbles will form and grow in the liquid; in the under-saturated case the gas bubbles will shrink and disappear. For a multi-component gas phase, the bubbles of free gas are a mixture of those components that come out of the solution.

The presence of dissolved and free gas in a liquid affects other properties, such as the rates of cavitation.

Setting up a Gas Dissolution Simulation

To use STAR-CCM+ for solving gas dissolution problems:

1. Set up a VOF Multiphase simulation.

   For more information, see Setting Up a VOF Multiphase Simulation.

2. Define the bulk properties and phase composition of the mixture.

   For more information, see Defining the Mixture Composition.

3. Create a Eulerian phase, and then do the following:

   a. Under the new phase node, right click the Models node and click Select models...

   b. In the *Phase 1 Model Selection* dialog, select Multi-Component Liquid from the *Material* group box.



c. Select Segregated Species from the *Optional Models* group box.

d. Select Constant Density from the *Equation of State* group box.

e. Click Close.

The selected models are added to the phase.



4. Repeat the previous step for a second phase but, in the *Phase 2 Model Selection* dialog, select Multi-Component Gas from the *Material* group box.

5. Right-click the Phase 1 > Models > Multi-Component Liquid > Liquid Components node, and then click Select Mixture Components.

6. In the *Select Mixture Components* dialog, select the appropriate components.

   The liquids that are available in the STAR-CCM+ material database do not include dissolved gasses. To specify a dissolved gas as a liquid component, you add the appropriate material to the material database.

7. Repeat the previous step for the gas phase.

8. Set the appropriate material properties relevant to the simulation.

9. Define the appropriate phase interactions.

   For more information, see Controlling Phase Interactions in VOF Simulations.

   Before you can select the Gas Dissolution model, you set the primary and secondary phases for the phase interaction.

10. Select the new VOF Phase Interaction model node.



11. In the *Properties* window, do the following:
    - Set the Primary Phase to the appropriate liquid phase.
    - Set the Secondary Phase to the appropriate gas phase.

    After you have set the primary and secondary phases for the phase interaction, you can select the Gas Dissolution model.

12. Right-click the Phase Interaction 1 > Models node and choose Select models...

13. In the *Phase Interaction 1 Model Selection* dialog, select Gas Dissolution from the *Optional Models* group box.

The *Phase Interaction 1 Model Selection* dialog appears as shown below.



14. Click Close.

    The Gas Dissolution model appears under the Multiphase Interaction > Phase Interactions > Phase Interaction 1 > Models node.

Setting Gas Dissolution Model Parameters

When a gas dissolution model is selected, STAR-CCM+ creates a node that is called Gas Dissolution in the simulation tree, as shown below.



Setting the Gas Dissolution Connectivity

The Gas Dissolution model has one property, Connectivity. This property identifies the gas component that corresponds to each dissolved gas component in the liquid mixture.

To specify the connections between the liquid and gas phases for each component:

1. Select the Multiphase Interaction > Phase Interactions > Phase Interaction 1 > Models > Gas Dissolution node.
2. In the *Properties* window, click to the right of the Connectivity property.
3. In the *Gas Dissolution - Connectivity* dialog, for each liquid component of dissolved gas, select the corresponding gas component.



For gas dissolution models, you ignore the gas components of the solvent liquids in the mixture. To ignore a particular component of the liquid mixture, set the corresponding gas component to None.

4. Click OK.

Setting the Gas Dissolution Equilibrium Coefficient Method

For each liquid component of dissolved gas, you can specify the method by which the equilibrium coefficient is found. The default for gas components that are dissolved in liquids is Henry's Law.

1. Expand the Models > Gas Dissolution > Equilibrium Coefficient > Composite node, and click the component that you want to modify.
2. In the *Properties* window, set the Method property to the appropriate option:
   - Henry's Law

     This option applies when the dissolved fraction is small, so it is the default for gas components that are dissolved in liquids.
   - Raoult's Law

     This option applies when the component is a large fraction of the liquid mixture. It does not usually apply to dissolved gas components.
3. If you use Henry's Law, you can specify the appropriate coefficient value in the Henry's Law Coefficient property.

   The default value is `4.26 E9 Pa`, which is the value for oxygen that is dissolved in water at 298.15 K.

Modeling Cavitation

The particular phase-change type concerns cavitation. This section describes:

- What is the Cavitation model?
- Physical parameters relevant to cavitation.
- Scaling cavitation bubble growth rates.
- How to set up a cavitation simulation.
- How to set the saturation pressure for the liquid phase components.
- How to set the seed properties for the cavitating fluid.
- How to set up initial and boundary conditions for each fluid phase.
- The mathematical formulation of the cavitation model.

The Cavitation Model

When the static pressure at a particular location within a liquid falls below the saturation vapor pressure of the liquid, the liquid undergoes a phase change called cavitation. This phase change creates cavitation bubbles that are filled with the vapor of the liquid.

When cavitation bubbles collapse, they release large amounts of energy into relatively small volumes. The collapsing bubbles create spots of high temperature and emit shock waves which are sources of noise. The collapsing bubbles can produce extensive erosion of the solid surfaces which are in contact with the cavitating regions.

The Cavitation model is based on the Rayleigh-Plesset equation and uses the inertia controlled growth model for bubble formation.

For more information on bubble growth rates, see Bubble Growth Rate Model for Cavitation.

The Cavitation model has no properties for single-component phases. The equilibrium coefficient for liquid components in a Cavitation simulation is found using Raoult's Law. There are no values to set.

For multi-component phases, you specify the connections between the liquid and gas phases for each component. For more information, see Cavitation Model Properties.

The Cavitation model node is shown below.



Cavitation Model Properties

Connectivity	Specifies the gas component that corresponds to each liquid component. This property applies to multi-component phases only.
ScalingFac+	Specifies the scaling factor for the positive bubble growth rate. You can set a constant value or use a field function.
ScalingFac-	Specifies the scaling factor for the negative bubble growth rate. You can set a constant value or use a field function. For more information, see Scaling Bubble Growth Rates.

Parameters Relevant to Cavitation

Cavitation is a complex process that is influenced by many factors such as fluid static pressure, saturation temperature and pressure, surface tension, viscosity, density, thermal conductivity, fluid velocity, and distribution of cavitation seeds. The difference between the static pressure p and the saturation pressure p_{sat} divided by the dynamic pressure of the incoming flow is called the *cavitation number*:

$$N_{cav} = \frac{p - p_{sat}}{\frac{1}{2}\rho_l U^2} \tag{2108}$$

The cavitation number is an indicator for likelihood of cavitation occurrence.

In order for cavitation to occur, the cavitation bubbles generally need a surface on which to nucleate. Impurities in the liquid can provide this surface. These pre-existing bubbles are called seeds and start to grow when the conditions for cavitation are met.

Cavitating flows are typically considered to be isothermal. However, if heat transfer becomes important, cavitation can be considered to behave in a manner similar to boiling. In this case, specify the latent heat that is released or consumed to support the cavitation correctly.

Scaling Bubble Growth Rates

The Rayleigh-Plesset equation is an approximation that is derived for idealized conditions. However, the specific conditions of your simulation can be less than ideal (for example, insufficient mesh and time resolution, unresolved bubble to bubble interactions and flow details).

If you want to move the simulation results in a direction that matches observed experimental results, you can scale the bubble growth rate by specifying the appropriate scaling factors for positive and negative bubble growth rates. Using the positive scaling factors lets you increase the cavitation rate if, for example, the mesh is not fine enough to resolve sharp corners that consequently cause STAR-CCM+ to underpredict the minimum pressure. Negative scaling factors also let you delay bubble collapse if, for example, the bubbles travel inside turbulent eddies where the local pressure is smaller than the average pressure that STAR-CCM+ calculates, and the bubble collapse is consequently slower than what the Rayleigh-Plesset equation predicts.

For more information, see The Cavitation Model and Setting up a Cavitation Simulation.

Setting up a Cavitation Simulation

The Unsteady (transient) approach is the recommended choice for a cavitation analysis. Cavitation is an intensive process, containing large sources in the momentum and continuity equations, which would require low under-relaxation parameters in a steady analysis.

To use STAR-CCM+ for solving cavitation problems:

1. Set up a VOF Multiphase simulation.

   For more information, see Setting Up a VOF Multiphase Simulation.

2. Define the bulk properties and phase composition of the mixture, including the vapor phase that is created when cavitation occurs.

   For more information, see Defining the Mixture Composition.

3. Create a Eulerian phase, and then do the following:

   a. Under the new phase node, right click the Models node and click Select models...

   b. In the *Phase 1 Model Selection* dialog, select the appropriate option from the *Material* group box:

      - Liquid, for a single-component model.
      - Multi-Component Liquid, for a multi-component model.



   c. If you are setting up a multi-component model, select Segregated Species from the *Optional Models* group box.

   d. Select Constant Density from the *Equation of State* group box.

   e. Click Close.

The selected models are added to the phase.



4. Repeat the creation step for a second phase but, in the *Phase 1 Model Selection* dialog, select the appropriate Gas option from the *Material* group box:
   - Gas, for a single-component model.
   - Multi-Component Gas, for a multi-component model.

   The liquid and gas phases must be either of the following: both single-component materials or both multi-component materials. You cannot have one phase as a multi-component material and the other phase as a single-component material.
5. Do one of the following:
   - For a single-component model:
     i. Right-click the Models > Liquid > H2O node, and then click Replace with....
     ii. In the *Replace Material* dialog, select the appropriate liquid.
   - For a multi-component model:
     i. Right-click the Models > Multi-Component Liquid > Liquid Components node, and then click Select Mixture Components.
     ii. In the *Select Mixture Components* dialog, select the appropriate components.
6. Repeat the previous step for the gas phase.

In some cases, the gaseous form of the solvent is not available in the STAR-CCM+ material database. To specify the gaseous component, you can add the appropriate material to the material database.

7. If you are working with multi-component materials, set the appropriate material properties.

   The Saturation Pressure is a key material property.

   Before you can select the Cavitation model, you set the primary and secondary phases for the phase interaction.

8. Define the appropriate phase interactions.

   For more information, see Controlling Phase Interactions in VOF Simulations.

9. Select the VOF Phase Interaction model node and do the following:
   - Set the Primary Phase to the appropriate liquid phase.
   - Set the Secondary Phase to the appropriate gas phase.

10. Right-click the Phase Interaction 1 > Models node and choose Select models...

11. In the *Phase Interaction 1 Model Selection* dialog, select Cavitation from the *Optional Models* group box.

    The *Phase Interaction 1 Model Selection* dialog appears as shown below.



12. Click Close.

    The Cavitation model appears under the Multiphase Interaction > Phase Interactions > Phase Interaction 1 > Models node.

13. If you want to scale the positive and negative bubble growth rates, select the appropriate node beneath the Cavitation Model node and specify the scaling value.

    The scaling factors are always positive and are independent: you can set either or both, as appropriate to your simulation.

For more information, see Cavitation Model Properties and Scaling Bubble Growth Rates.

Mapping Liquid Phase Components to Gas Phase Components

Multi-component simulations require that you map the components in the liquid phase to their corresponding components in the gas phase.

To map components:

1. Select the Models > Cavitation node.
2. In the *Properties* window, click to the right of the Connectivity property.
3. In the *Cavitation - Connectivity* dialog, for each liquid component, select the corresponding gas component.



For cavitation models, you can ignore the dissolved gas components of the liquid mixture. To ignore a particular component of the liquid mixture, set the corresponding gas component to None.

4. Click OK.

Setting the Saturation Pressure for the Liquid Phase Components

The Saturation Pressure p_{sat} is a key material property that the Rayleigh-Plesset model requires.

You set this pressure as an absolute value using the Saturation Pressure node contained within the Material Properties manager node for the liquid phase. For a multi-component liquid, you set the saturation pressure for each component of the solvent. It is not relevant to the dissolved gas components.

In the example that is shown below, a constant value has been set: other methods are also available for this property.



Gas Dissolution and Cavitation Modeling Principles

Many phase change processes can be described as bulk processes that start from certain seeds (imperfections that act as nuclei). A homogeneous approach, which is based on the presence of seeds and their subsequent growth and collapse as bubbles, uses this assumption to calculate an appropriate mass transfer rate. The mass transfer rate is based on the material and thermodynamic properties of the entire control volume [385].

Setting the Seed Density and Seed Diameter

The seed density and seed diameter can only be specified per phase interaction, and not per rate model. If multiple rate models use the same phase pair, for example simultaneous cavitation and gas dissolution between a liquid and a gas, the results are consistent with the general formulation. That is, the local radius of the bubbles cannot be different depending on which rate model is used.

If you use two different phase interactions, between the same two phases, for modeling cavitation and dissolution, you can specify the seed properties individually for each phase interaction. This option allows you to tune the models, but in a manner that is inherently inconsistent.

Setting the Seed Density

The Seed Density n_0 is a liquid-dependent constant and is defined as the number of bubble seeds per unit volume of liquid (see Gas Dissolution and Cavitation Modeling Principles). Its value is set in the Multiphase Material > Multiphase Material > Material Properties > Seed Density node.



Seed Density Properties

Method	The method that is used to define the value for the seed density.	
	Constant	The material property is input as constant. A Constant node is added as a child to this node.
	Field Function	The material property is entered using a field function. A Field Function node is added as a child to this node.
Dimensions	The dimensionality of the property, expressed in terms of the basic dimensions, such as length, time (read-only).	

Setting the Seed Diameter

The average Seed Diameter ($2\ R_0$) is used to approximate the spectral seed distribution in a liquid (see Gas Dissolution and Cavitation Modeling Principles). Its value is set in the Multiphase Material > Multiphase Material >

Material Properties > Seed Diameter node.



Seed Diameter Properties

Method	The method that is used to define the value for the seed density.	
	Constant	The material property is input as constant. A Constant node is added as a child to this node.
	Field Function	The material property is entered using a field function. A Field Function node is added as a child to this node.
Dimensions	The dimensionality of the property, expressed in terms of the basic dimensions, such as length, time (read-only).	

Setting Initial and Boundary Conditions

Following the phase specification of the mixture, the Initial Conditions node contains an additional Volume Fraction sub-node. This node lets you define the initial distribution of each phase, including the gas (if present). STAR-CCM+ obtains the volume fraction of the last-defined phase automatically, by subtracting the sum of the volume fractions of the other phases from 1.0.

For multi-component phases, you specify the boundary and initial conditions for the mass fraction or molar fraction of each component of the phase.

The node of every **open** boundary in the simulation also contains a Volume Fraction sub-node. This node lets you define the boundary condition on the volume fractions of each phase, including the gas. As with initial conditions, the boundary value of the last-defined phase is calculated automatically, by subtracting the sum of the boundary values of the other phases from 1.0.

If surface tension effects are included in the simulation, all **closed** (wall) boundaries also contains an additional Contact Angle sub-node (of which the **Primary Phase** is a reference).

Gas Dissolution and Cavitation Formulation

This section includes:

- An overview of mass transfer models
- The general formulation of seed-based models
- Determining the gas dissolution rate

- Determining the cavitation bubble growth rate
- Modeling the change in enthalpy
- A summary of the formulation notation
- Gas Dissolution and Cavitation bibliography

Overview of Mass Transfer Models

This section describes a VOF method to handle a particular class of physical models: models that describe phase change, and corresponding mass transfer, between two physical phases (for example, a liquid and gaseous phase). In general, there are two ways to handle phase change:

- Interface resolution

  This method requires the capturing of fine details, such as the micro-bubbles in a fluid, with multiple cells. Therefore it is numerically expensive.

  The Evaporation/Condensation model uses an interface resolving method.

- Homogeneous approach

  This method makes certain assumptions regarding the interface between the phases, on a subcell level. This method is numerically more efficient, but has limitations regarding the physics that can be resolved.

  The Boiling, Gas Dissolution, and Cavitation models use a homogeneous approach.

STAR-CCM+ uses the same homogeneous seed-based approach for both cavitation and gas dissolution models. Although these zero-dimensional approaches do not capture all of the physical phenomena present, they are proven methods for use in industrial design studies.

General Formulation of Seed-based Mass Transfer Models

This section outlines the general approach for handling seed-based models, which are based on [385]. The two interacting phases are denoted as l and g, with mass transfer positive from l to g.

The volume fraction of the gas can be expressed as:

$$\alpha_g = \frac{V_g}{V} = \frac{V_g}{V_g + V_l} \tag{2109}$$

where V is the control volume, and V_g and V_l denoting the respective parts of the volume that the gas and the liquid phases occupy.

Within the control volume, the gas phase is assumed to be present in the form of bubbles. Each bubble has the same radius R. The number density of seeds is defined as n_0, which corresponds to the number of bubbles per unit volume.

Therefore the following relationship holds between the gas and liquid phases:

$$V_g = n_0 V_l \frac{4}{3}\pi R^3 \tag{2110}$$

The volume fraction of gas, α_g, becomes:

$$\alpha_g = \frac{n_0 \frac{4}{3}\pi R^3}{1 + n_0 \frac{4}{3}\pi R^3} \tag{2111}$$

The gas bubble radius can be expressed as:

$$R^3 = \frac{\alpha_g}{n_0 \frac{4}{3}\pi (1 - \alpha_g)} \tag{2112}$$

The seed diameter $D_0 = 2R_0$ (which is user-specified) provides a lower limit R_{min} for the bubble size.

The mass transfer rate per unit volume $kg/m^3 s$ can be written as:

$$\dot{m} = n_0 \alpha_l 4\pi \rho_g R^2 \frac{DR}{Dt} \tag{2113}$$

such that liquid disappears with a volume rate $m^3/m^3 s$:

$$\dot{V}_l = -\frac{\dot{m}}{\rho_l} \tag{2114}$$

and gas is generated with a volume rate:

$$\dot{V}_g = \frac{\dot{m}}{\rho_g} \tag{2115}$$

Gas bubbles form in the liquid, and increase in radius as more liquid changes phase:



The remaining unknown factor is $\frac{DR}{Dt}$, giving the increase in bubble radius as the bubble flows.

That factor cannot be determined in a general fashion and requires input from a physical model. The model estimates the bubble radius change rate, which is based on properties of the flow and the materials, within the control volume.

If the phases are multi-component, the rate model breaks down the global mass transfer rate $\dot{m}$ in rates $\dot{m}_i$ for each individual component. These component mass transfer rates appear as source terms in the species transport equations.

Bubble Growth Rate Model for Gas Dissolution

For single-component bubbles, the approximated Epstein-Plesset [387], [388], [386] formulation is assumed to hold:

$$\frac{DR}{Dt} = \frac{D_{i,l}(c_{i,l,\infty} - c_{i,l,s})}{\rho_g R} \tag{2116}$$

Due to the small fraction of $c_{i,l}$ present, this formulation can be written as:

$$\frac{DR}{Dt} = \frac{D_{i,l}\rho_l(Y_{i,l,\infty} - Y_{i,l,s})}{\rho_g R} \tag{2117}$$

To determine the saturation mass fraction, Henry's law is used, in the following form:

$$X_{i,g,s} = \frac{H_i}{p^{abs}} X_{i,l,s} \tag{2118}$$

or

$$X_{i,l,s} = \frac{p^{abs}}{H_i} X_{i,g,s} \tag{2119}$$

where H_i is the Henry's law constant.

This expression is similar to Raoult's law, but is valid in the limit of $X_{i,l,s} \rightarrow 0$, whereas Raoult's law holds in the opposite limit.

In single-component gas dissolution, $X_{i,g,s} = 1$.

To convert molar fraction to mass fraction, use the assumption:

$$\sum_{j}^{N} X_{j,l} W_j \approx \sum_{j}^{N_{sv}} X_{j,l} W_j \tag{2120}$$

with:

$$\sum_{j}^{N_{sv}} X_j \approx 1 \tag{2121}$$

where N_{sv} is the number of solvent components. The dissolved species barely contribute to the molar mass of the liquid mixture, due to their low concentration. Therefore, a good approximation is:

$$Y_{i,l,s} \approx (X_{i,l,s} W_i) / \left(\sum_{j}^{N} X_{j,l,\infty} W_j \right) \tag{2122}$$

Eqn. (2117) determines the bubble radius change rate. In the single-component case, the partial pressure of this component is equal to the total gas pressure. The extension to multi-component is readily achieved by assuming ideal gas behavior. Each component contributes to the pressure according to its molar fraction in the free gas, and the species distribution within the bubble maintains a constant profile. In that case, $X_{i,g,s}$ becomes a transported quantity that is influenced by the gas dissolution rates.

The bubble radius rate change is computed as:

$$\frac{DR}{Dt} = \sum_{i}^{N_{sl}} \frac{D_{i,l}\rho_l(Y_{i,l,\infty} - Y_{i,l,s})}{\rho_g R} = \sum_{i}^{N_{sl}} \left(\frac{DR}{Dt}\right)_i \tag{2123}$$

Bubble Growth Rate Model for Cavitation

For both single-component materials and multi-component materials, you can scale the bubble growth rate and collapse rate. You specify the appropriate scaling factors for positive and negative bubble growth rates in the Cavitation model settings.

For more information, see Scaling Bubble Growth Rates.

Single-Component Material

The single-component cavitation bubble growth rate is estimated using the *inertia controlled growth* model [385]:

$$\left(\frac{DR}{Dt}\right)^2 = \frac{2}{3}\left(\frac{p_{sat} - p_\infty}{\rho_l}\right) \tag{2124}$$

where p_{sat} is the saturation pressure corresponding to the temperature at the bubble surface, p_∞ is the pressure of the surrounding liquid and ρ_l is the liquid density. Eqn. (2124) is a simplification of the more general Rayleigh-Plesset equation which takes into account the inertia, viscous, and surface tension effects:

$$R\frac{d^2R}{dt^2} + \frac{3}{2}\left(\frac{dR}{dt}\right)^2 = \frac{p_{sat} - p_\infty}{\rho_l} - \frac{2\sigma}{\rho_l R} - 4\frac{\mu_l}{\rho_l R}\frac{dR}{dt} \tag{2125}$$

μ_l is the liquid viscosity and σ is the surface tension coefficient. Research results suggest that for most practical applications, it is not necessary to account for the viscous and surface tension effects [384].

Multi-Component Material

Eqn. (2124) can still be used for a mixture of multiple components. In the context of mixture saturation pressure, the equation becomes:

$$\left(\frac{DR}{Dt}\right)^2 = \frac{2}{3}\left(\frac{p_{sat,m} - p_\infty}{\rho_l}\right) \tag{2126}$$

with

$$p_{sat,\,m} = \sum_{i}^{N_c} X_{i,\,l,\,s} p^*_{sat,\,i} \tag{2127}$$

This expression is Raoult's Law, with $p^*_{sat,\,i}$ indicating the saturation pressure of the pure component. The interfacial molar fraction of the liquid components is approximately their value in the bulk, so $X_{i,\,l,\,s} \approx X_{i,\,l,\,\infty}$.

This formulation neglects the effect of other gases that are present in the bubbles (for example, as a consequence of gas dissociation) and which have their own partial pressure in the bubble. Homogeneous models compute only the pressure of the entire phase mixture, rather than computing a pressure for each individual phase. This single pressure puts restrictions on including the effect of cavitation-inert gases on the cavitation bubble growth rate. The implicit assumption is that cavitation, once it occurs, is so fast that it overwhelms other effects. Important deviations can be expected near the end of the collapse process, as the presence of 'slow gases' dampens the collapse [384]. However, those deviations are probably no larger than the already omitted effects of thermal effects, surface tension, and viscosity. This limitation seems reasonable for most practical applications.

The overall cavitation rate is broken down into rates for the individual components. Unresolved interfacial properties also affect this overall rate. A reasonable assumption is to use the following:

$$\dot{m}_{cav,\,i} = \dot{m}_{cav} \left((X_{i,\,l,\,\infty} p^*_{sat,\,i}) \Big/ \left(\sum_{i}^{N_c} X_{i,\,l,\,\infty} p^*_{sat,\,i} \right) \right) \tag{2128}$$

This expression represents a weighting that is based on the individual fugacity of each component, in the context of ideal gases and ideal liquids.

Modeling Change of Enthalpy

Due to the lack of interface resolution, the effects of latent heat on the mass transfer rate are only partially modeled. The enthalpy is described using a Heat of Formation for all components, as follows:

$$\Delta H_i = \Delta H_{for,\,i} + \int_{T_{ref}}^{T} C_{p,\,i}\, dT \tag{2129}$$

From this equation, you can compute an enthalpy of the phase, and then compute the enthalpy of the entire mixture. The latent heat of phase change for a specific component is then incorporated into the difference in component enthalpies at that particular temperature.

Rapid evaporation (for example, as a consequence of cavitation or gas dissociation) cools the interface between the phases. Homogeneous models do not resolve the interface, so no energy balance at the interface is incorporated: the energy balance is considered only on the level of the control volume. Therefore, conditions under which the thermal effects become important are unlikely to be simulated accurately.

Note: The cooling (or heating) effect is accurately included for conditions that are not thermally limited: the entire phase mixture does cool down as a consequence of net evaporation. The inaccuracies lie only in the calculation of the rate as a function of temperature.

Formulation Notation

This section describes the notation that is used in the Gas Dissolution and Cavitation Formulation sections.

N_{sv}	number of solvent components
N_{sl}	number of solute components
N_d	number of dissolving components
N_c	number of cavitating components
D	diffusion coefficient (m^2/s)
ρ	density (kg/m^3)
c	concentration (kg/m^3)
Y	mass fraction
X	molar fraction
W	molar weight (kg/mol)
α	volume fraction
n_0	seed density (m^{-3})
D_0	seed diameter (m)
R	bubble radius (m)
g	gas
l	liquid
s	surface
i	component index
L	interaction length scale (m)
∞	bulk
p_{sat}	saturation pressure (Pa)
$*$	pure component
$\dot{m}$	mass transfer rate per unit volume (kg/m^3 -s)
C_p	heat capacity at constant pressure (J/K)

Gas Dissolution and Cavitation Bibliography

[384] Brennen, C.E., *Cavitation and Bubble Dynamics*, Oxford University Press, 1995.

[385] Sauer, J., I*nstationaer kavitierende Stroemungen - Ein neues Modell, basierend auf Front Capturing VOF und Blasendynamik*, Dissertation, Universitaet Kalrsruhe, 2000.

[386] Ni, B., Pieprzak, J., *Transportation and transformation of air bubbles in aerated oil through an engine lubrication system*, SAE-Paper 2004-01-2915.

[387] Epstein, P., Plesset, M., *On the stability of gas bubbles in liquid-gas solutions*, J. Chemical Physics, Vol. 18, No. 11 (1985).

[388] Muzaferija, S., Peric, M., *Simulation of evolution of dissolved gas in liquids*, Internal document.

[389] Singhal, A.K., Athavale, M.M., Li, H., Jiang, Y., *Mathematical Basis and Validation of the Full Cavitation Model*, J. Fluids Engineering, Vol. 124 (2002).

[390] Nemec, T., Marsik, F., *The Classical Multicomponent Nucleation Theory for Cavitation in Water with Dissolved Gases*, Proc. 7th Int. Symp. on Cavitation, Paper No. 120, (2009)

Modeling Melting and Solidification

This section includes the following:

- Melting and Solidification
- Selecting a Simulation with Melting and Solidification
- Setting Melting-Solidification Model Parameters
- Working with the Melting-Solidification Flow Stop Model
- Using the Flow Stop Mass Compensation Model
- Flow Stop Field Functions
- Setting the Fraction Solid Curve
- Setting up the Slurry Viscosity Model
- Setting up the Mushy Zone Permeability Model
- What Is the Relative Solid Volume Fraction of the Phase?
- Melting-Solidification Formulation
  - Basic Melting-Solidification
  - Flow Stop Mass Compensation
  - The Slurry Viscosity Model - Metzner Method
  - Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model

- The Mushy Zone Permeability Model - Carman-Kozeny method
- Melting-Solidification Bibliography

Melting and Solidification

Melting is the process which changes the state of matter from solid to liquid. It takes place when a solid substance is heated to the melting temperature. The opposite process is called solidification. In STAR-CCM+ you do not need to specify a solid material for this process. Choose a liquid material, and then specify temperature values that cause the material to melt and/or to solidify. The Relative Solid Volume Fraction indicates the state of the fluid.

For pure substances, for example water, the melting and solidification usually take place at one temperature, that is the liquidus and the solidus are at the same temperature. Melting and solidification take place instantaneously. Depending on the temperature, there is either liquid or solid in the computational domain.

For impure substances, for example alloys, the melting and solidification can take place over a temperature range. In this case, the liquidus and solidus are not at the same temperature. In the temperature interval between them, the liquid is partly solidified, that is the solid and liquid phases exist simultaneously (like a slurry). This temperature interval can have a finite width up to several 100 degrees Kelvin. The melting and solidification process does not take place instantaneously but over a period of time. During this time, the phase mixture can influence the flow field.

Flow Resistance due to Solidification

The Melting and Solidification model only simulates the effect of phase change on the energy equation; it does not modify the momentum equations. For this reason, a further model is required to account for the change in material properties from the liquid state to the solid state.

Depending on the value of the relative solid fraction (specifying to which extent the liquid has solidified), there are three optional models available:

- **Slurry Viscosity Model**

  For low values of solid fraction the mixture is considered as a slurry, where solid grains are suspended in the melt nearly without interacting with each other. Their effect on the flow is often modeled as a relative solid fraction-dependent viscosity.

The Metzner method is implemented to increase the viscosity (slurry viscosity) up to the critical solid fraction (approximately 0.3) (see Eqn. (2135)), which models the solidified grains flowing with the fluid without much interaction. However, using a range of viscosity larger than about two orders of magnitude can be unstable.

- **Mushy Zone Permeability Model**

  When the solid volume fraction increases, equiaxed grains grow and start to agglomerate and to form dendritic regions. Dendrites also grow from cooled solid surfaces or the solidification front into the melt. A partially solidified stationary region permeated with dendrites is called a "mushy zone". The flow resistance in a mushy zone can be modeled similar to an isotropic porous medium, using the Carman-Kozeny Mushy Zone Permeability method (see Eqn. (2141)). This model is used for relative solid volume fractions above the specified critical solid fraction (approximately 0.3).

- **Flow Stop Functionality**

  To stop the flow on solidification completely, you can use the flow stop functionality. This model is used above the specified flow stop solid fraction (approximately between 0.4 and 1). The resistance that the Slurry Viscosity and Mushy Zone Permeability models add to moving fluids depends on the flow velocity. When the velocity is zero, the effect of both models is also zero. In the presence of pressure gradients, or body forces such as gravity, a creeping flow develops independently of the magnitude of slurry viscosity or mushy zone permeability.

  However, the flow stop functionality has limitations. You can only use the flow stop functionality in simulations where the cell volume remains constant. This includes simulations with rigid body motion and moving reference frames, but excludes simulations with morphing or user-defined vertex motion. For cases where the cell volume changes, use a field function to increase the dynamic viscosity with respect to the relative solid volume fraction by several orders of magnitude. Avoid sudden changes of viscosity, as these can destabilise the simulation.

Applicability of the Flow Resistance models

The slurry viscosity model and the mushy zone permeability model affect the flow properties of a (partially) solidified liquid only. You only need these models when you are simulating impure substances such as alloys. You do not need them, for example, when you simulate ice melting in a stationary container.

If you simulate the melting and solidification process of a pure substance, use the melting and solidification model only. Partial solidification is not possible for pure substances (see Melting and Solidification). To prevent creeping flows, you can additionally activate the flow stop functionality.

When simulating melting and solidification of impure substances, use the following table to determine which flow resistance models to apply.

Flow Velocities		Solidus/Liquidus temperature interval		Simulation Time		Flow Resistance Models
Low (< 1 m/s)		Small		Short (< 1 s)		Slurry Viscosity
Low (< 1 m/s)		Small		Long (> 3 s)		Slurry Viscosity Flow Stop
Low (< 1 m/s)		Large		Short (< 1 s)		Slurry Viscosity Mushy Zone Permeability
Low (< 1 m/s)		Large		Long (> 3 s)		Slurry Viscosity Mushy Zone Permeability Flow Stop
High (> 5 m/s)		Small		Short (< 1 s)		Slurry Viscosity (Flow Stop)
High (> 5 m/s)		Small		Long (> 3 s)		Slurry Viscosity Flow Stop
High (> 5 m/s)		Large		Short (< 1 s)		Slurry Viscosity Mushy Zone Permeability (Flow Stop)
High (> 5 m/s)		Large		Long (> 3 s)		Slurry Viscosity Mushy Zone Permeability Flow Stop

Selecting a Simulation with Melting and Solidification

Before you begin, activate the following models:

- In the *Physics Model Selection* dialog for the continuum:
  - Select Eulerian Multiphase from the *Material* group box.

- Select Volume of Fluid (VOF) from the *Eulerian Multiphase Model* group box.
- Select Segregated Multiphase Temperature from the *Optional Models* group box.

Note that the melting solidification model is available for liquid Eulerian phases only.

To solve problems that involve melting and solidification:

1. Right-click Continua > Physics 1 > Models > Eulerian Multiphase > Eulerian Phases > Phase 1 > Models and select Select models...
2. In the *Phase Model Selection* dialog:
   a. Select Liquid from the *Materials* group box.
   b. Select any option from the *Equation of State* group box.
   c. Select Melting-Solidification from the *Optional Models* group box.



Additionally, you can activate the Mushy Zone Permeability or the Slurry Viscosity model from the *Optional Models* group.

Setting Melting-Solidification Model Parameters

When a simulation with melting and solidification is selected, STAR-CCM+ creates a Melting-Solidification node in the simulation tree, under the Eulerian Phases > (Phase name) > Models manager.



This node lets you set the relative solid fraction at which the flow stops in a cell and the under-relaxation factor for melting and solidification. The flow stop is active only when the respective Flow Stop Solid Fraction is also specified for the Melting-Solidification Flow Stop model.

Additional material properties are added when the melting-solidification model is selected. The Fraction Solid Curve property is selected. When the Linear method is selected, the Liquidus Temperature and Solidus Temperature properties are also selected. These properties are shown in the following screenshot.



One of the user-defined phases must be the phase that can melt and solidify, and therefore is defined as liquid (water in the preceding example).

The melting-solidification model requires the fraction solid curve, the solidus, and liquidus temperatures and the latent heat of fusion.

It is good practice to use a field function for the Thermal Conductivity material property of the phase that melts and solidifies. The values of this property differ for liquid and solid states.

Melting-Solidification Properties

Flow Stop Relative Solid Fraction	The flow stop relative solid fraction defines the cell solid volume fraction at which the fluid loses its ability to flow. However, the flow is stopped in the cell only when the total amount of solidified fluid exceeds the Flow Stop Solid Fraction that is specified for the Melting-Solidification Flow Stop model. This property is deactivated for values greater than 1. Default value `1.1`.

Melting-Solidification Expert Properties

Linearize Melting	Linearizes the influence of melting.	
	Activated	Activates the expert property.
	Deactivated	Default option.
Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution.	

What Is the Latent Heat of Fusion?

The specific latent heat of fusion h_{fusion} is the amount of energy in the form of heat that a substance releases or absorbs during a change of state. It is represented by a material property node under the Models > Liquid > (Liquid

name) > Material Properties manager.



Working with the Melting-Solidification Flow Stop Model

When the Melting-Solidification model is selected for a Eulerian phase, the Melting-Solidification Flow Stop model is automatically added to the respective Physics continuum.



The Melting-Solidification Flow Stop model lets you specify the solid volume fraction at which the fluid flow is stopped in a cell.

Flow stopped cells have zero velocity (relative to the cell center velocity) and a pressure equal to the face-weighted average pressure of all adjacent non-stopped cells. However, if a flow stopped cell is only next to other flow stopped cells, the static pressure is equal to the reference pressure, as this model does not calculate stresses in the solid.

You can visualize flow stopped cells by using the Flow Stop Flag field function.

When you use the flow stop functionality, the following limitations apply:

- The flow stop relative solid fraction does not apply to regions that are subject to motion that affects the mesh volume (for example, Morpher or User Defined Vertex Motion.).

  In morphed regions, the flow stop flag for all cells is set to zero (that is, there are no flow-stopped cells at any time).
- The face between a stopped and a non-stopped cell is not a wall.

  No wall boundary conditions for flow (slip/no-slip, wall function...) or turbulence are applied to faces between stopped cells and non-stopped cells.

Activating the Melting-Solidification Flow Stop Model

The flow stop functionality is required when you want to prevent all flow (not allowing slow creeping flow) or the solid needs to withstand very high pressure gradients.

You normally set the Flow Stop Solid Fraction to high values (0.99) because the simulation already includes the effect of slurry viscosity and mushy zone which decelerate the flow for lower solid volume fractions. However, you may set a lower value when there is a large interval between the solidus and liquidus temperatures.

To activate the Melting-Solidification Flow Stop model:

1. In the appropriate phase, select the Melting-Solidification node and set the Flow Stop Relative Solid Fraction property to a value between 0 and 1.

   You set this property for each phase. The value depends on the properties of the material, and it specifies when the particular phase loses its ability to flow.
2. In the physics continuum, select the Melting-Solidification Flow Stop node and set the Flow Stop Solid Fraction property to a value between 0 and 1.

   This property specifies when a cell is blocked for the flow. It is a property of the continuum, so it depends on all of the phases that are present.

To deactivate this functionality:

- Set the Flow Stop Solid Fraction property to a value greater than 1.

Using the Flow Stop Mass Compensation Model

When the flow stop functionality of the melting-solidification model is used in combination with a compressible equation of state (Polynomial Density or User Defined EOS), fluxes through cell faces cannot compensate for the change in density in stopped cells. This effect causes the fluid mass in the

respective cell to change proportional to the density change. Over the simulation time, this change can have a significant effect on the total mass of the solidifying phase.

The purpose of the Flow Stop Mass Compensation model is to apply appropriate mass and energy sources to compensate for the unphysical mass change in stopped cells and maintain the total mass of each phase.

To prevent a flow deflection in a non-stopped cell next to a stopped cell, the face area weighted average static pressure of all non-stopped cells next to a stopped cell is applied to the stopped cell. To maintain the correct heat fluxes through the stopped cells, the temperature cannot be adjusted. The density, which is a function of pressure and temperature, determines the mass in a stopped cell that has a constant volume.

There is no way to compensate the mass change in stopped cells when none of the pressure, temperature, density, or cell volume values can be modified. Therefore the mass change is adjusted using sources that are applied to fluid phases in the cells on the surface of the contiguous stopped region.

To activate the Flow Stop Mass Compensation model:

1. Ensure that the Melting-Solidification Flow Stop model is activated.
2. In the physics continuum, select the Melting-Solidification Flow Stop node and, in the *Properties* window, activate the Flow Stop Mass Compensation option.
3. Set the Max Mass Extraction Factor to the appropriate value.

   This property specifies the maximum fraction of the phase mass that is extracted from a liquid cell that is next to a contiguous stopped region. When you apply the phase sources, the fraction that eventually leaves the cell as convective flux is unknown. Therefore, it is recommended that you set the value significantly less than 1. The default value is 0.5.

Flow Stop Field Functions

When you activate the Melting-Solidification Flow Stop model, the following field functions become available.

Flow Stop Flag	Lets you visualize the flow stopped cells.

Flow Stop Mass Compensation Error	This field function is registered for each VOF phase. It shows the total mass of the phase that could not be extracted from the fluid cells that were next to contiguous stopped regions. There are two possible causes of this error: • More mass is required to be extracted from the cell than the value specified by the Max Mass Extraction Factor. • The cell did not contain enough of the respective phase for extraction.

When the Temporary Storage Retained option is selected for the Segregated Flow solver, the following field functions become available.

Contiguous Stopped Region Index	A domain-wide unique index of contiguous stopped regions.
Flow Stop Mass Source	This field function is registered for each VOF phase. It shows the value of the mass source in fluid cells that are next to the stopped cells of a contiguous stopped region, in kg/s.

Setting the Fraction Solid Curve

The fraction solid curve defines the relative solid volume fraction variation with temperature (and absolute pressure for compressible simulations). A material property, that is called solid fraction curve, defines this path in STAR-CCM+. The following plot shows an example of a solid fraction

curve.


Monitor Plot 1
Solid Fraction Curve for Al-Cu Alloy.
Solid Volume Fraction
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
850
860
870
880
890
900
910
920
Temperature (K)

The Fraction Solid Curve node represents the solidification path property in the simulation tree.



Specify absolute temperature values when using the Polynomial in T, Table(T), or Table(T,P) methods. Solidus and liquidus temperatures are then assumed to be the minimum and maximum specified values. Both temperatures must exceed the minimum temperature that is specified under Reference Values. The solid volume fraction is 1 for the solidus temperature, and 0 for the liquidus temperature.

Fraction Solid Curve Properties

Dimensions — The dimensionality of the property that is expressed in terms of the basic dimensions, such as length, time (read-only).

Method — The method that is used to define the value for this material property. The default is Linear.

- Linear — The solid volume fraction changes linearly with temperature between the solidus and liquidus temperatures. (Eqn. (2133))
- Polynomial in T — The material property is entered using a Polynomial in T. A node by that name is added as a child to this node.
- Table(T) — The material property is entered using a table. A node by that name is added as a child to this node. Only applicable to some material properties.
- Table(T,P) — The material property is entered using a table. A node by that name is added as a child to this node. Only applicable to compressible cases.

What Are Solidus and Liquidus Temperatures?

It is assumed that melting and solidification take place when the substance is at a temperature between the solidus and liquidus temperatures. If the substance is at a temperature below the solidus temperature, the substance is entirely in the solid state. If the substance is at a temperature above the liquidus temperature, the substance is entirely in the liquid state. The

solidus and liquidus temperatures can have the same value for pure substances. The material property nodes under the Models > Liquid > (Liquid name) > Material Properties manager represent these temperatures.



The Liquidus Temperature and Solidus Temperature nodes are available only when choosing the Linear Fraction Solid Curve method. When using other methods, the solidus and liquidus temperatures are automatically computed as the lower and upper limits of the input that is provided for the respective method.

The liquidus and solidus temperatures must be within the Minimum and Maximum Allowable Temperatures that you set within the Reference Values for the physics continuum.

When using the linear fraction solid curve method, the liquidus temperature must be greater or equal to the solidus temperature. A minimum difference of 0.002 K between both is automatically applied. For other methods, the liquidus temperature must be at least 0.002 K higher than the solidus temperature.

Liquidus Temperature Properties

Dimensions
: The dimensionality of the property that is expressed in terms of the basic dimensions, such as length, time (read-only).

Method
: The method that is used to define the value for this material property.

  Constant
  : The material property is input as a constant. A Constant node is added as a child to this node.

Solidus Temperature Properties

Dimensions
: The dimensionality of the property that is expressed in terms of the basic dimensions, such as length, time (read-only).

Method
: The method that is used to define the value for this material property.

  Constant
  : The material property is input as a constant. A Constant node is added as a child to this node.

What Is the Linear Fraction Solid Curve Method?

With this method, the solid volume fraction changes linearly with temperature between the solidus and liquidus temperatures, as specified in

Eqn. (2133). The Linear node represents this method in the simulation tree.



Linear Properties

None

Setting up the Slurry Viscosity Model

For values of solid fraction lower than a critical solid volume fraction, the mixture is considered as a slurry, where solid grains are suspended in the melt with almost no interaction between them. Their effect on the flow is modeled as an additional viscosity using the slurry viscosity model.

Before you can activate the slurry viscosity model, select the melting solidification model (see Selecting a Simulation with Melting and Solidification). The slurry viscosity model is optional to the melting solidification model.

To select the slurry viscosity model:

1. Right-click Continua > Physics 1 > Models > Eulerian Multiphase > Eulerian Phases > [Phase] > Models and choose Select models...
2. In the *Phase Model Selection* dialog, select Slurry Viscosity from the *Optional Models* group box.



A Slurry Viscosity node is added to the Eulerian phase.



3. Set the method for the slurry viscosity model by selecting the Liquid > H2O > Material Properties > Slurry Viscosity node.
4. In the *Slurry Viscosity - Properties* window, select the method to use.

Method	Select the specific method to use for the slurry viscosity model.	
	Metzner	Select the Metzner method. If you select this option, a Metzner sub node appears where you can define the properties of the method (see The Slurry Viscosity Model - Metzner Method).
	Field Function	User-define a slurry viscosity method by use of a field function.

Depending on the selection, a Metzner node or a Field Function node appear.

- To define the properties of the Metzner method, select the Metzner node.



In the *Metzner - Properties* window, set the following properties::

Use Switching Function	Toggle the use of the switching function for the combined use of the Metzner slurry viscosity model and the Carman-Kozeny mushy zone model on and off (see Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model).

Switching Function: Critical Solid Fraction g_cr	The critical solid volume fraction α^*_{cr} (see Eqn. (2137)). The default value is 0.27 (for Aluminum alloy A201).
Crystal Constant A	The crystal constant A which depends on the aspect ratio and surface roughness of the crystal (see Eqn. (2135)). The default value is 0.3 (for Aluminum alloy A201).
Switching function constant s	The non-dimensional constant s (see Eqn. (2137)). The default value is 100 because the shape is rather insensitive for values from 80 through 120.

- To define the field function, select the Field Function node. Choose the appropriate field function that you have defined.

5. Finally, to set the under-relaxation factor for the property change per iteration, select the Average Slurry Viscosity node.



6. In the *Average Slurry Viscosity - Properties* window, set the under-relaxation factor if necessary.

Slurry Viscosity Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7.

A volume-weighted average slurry viscosity is computed based on the viscosities of each phase where melting and solidification is present.

For simulations with different phases, the slurry viscosity for one cell is taken as a volume-weighted average for the whole mixture. This volume-weighted average is represented by the Multiphase Mixture > Mixture > Mixture Properties > Slurry Viscosity node.



Setting up the Mushy Zone Permeability Model

When the solid volume fraction increases above a critical solid volume fraction, a mushy zone permeated with dendrites is formed. The flow resistance in a mushy zone is modeled as an isotropic porous medium.

When the Mushy Zone Permeability model is selected on one or more Eulerian phases, an Average Mushy Zone Permeability model node appears in the parent physics continuum. The contributions to flow resistance from each phase where melting and solidification occur are volume-weighted for the computed cell. These contributions are then added to the flow resistance source term of the momentum equation.

Before you can activate the mushy zone permeability model, select the melting solidification model (see Selecting a Simulation with Melting and Solidification). The mushy zone permeability model is optional to the melting solidification model.

To select the mushy zone permeability model:

1. Right-click Continua > Physics 1 > Models > Eulerian Multiphase > Eulerian Phases > Phase 1 > Models and select Select models...
2. In the *Phase Model Selection* dialog, select Mushy Zone Permeability from the *Optional Models* group box and click Close.



A Mushy Zone Permeability node is added under the Eulerian phase.



3. Set the method for the mushy zone permeability model by selecting the Liquid > H2O > Material Properties > Mushy Zone Permeabiltiy node.

4. In the *Mushy Zone Permeability - Properties* window, select the method to use.

Method	Select the specific method to use for the mushy zone permeability model.	
	Carman-Kozeny	Select the Carman-Kozeny method. If you select this option, a Carman-Kozeny sub node appears where you can define the properties of the method (see The Mushy Zone Permeability Model - Carman-Kozeny method).
	Field Function	User-define a mushy zone permeability method by use of a field function.

5. Depending on the selection, a Carman-Kozeny node or a Field Function node appear.

- To define the properties of the Carman-Kozeny method, select the Carman-Kozeny node.



In the *Carman-Kozeny - Properties* window, set the following properties:

Use Switching Function	Toggle the use of the switching function for the combined use of the Metzner slurry viscosity model and the Carman-Kozeny mushy zone model on and off (see Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model).

Switching Function: Critical Solid Fraction	The critical solid volume fraction α^*_{cr} (see Eqn. (2137)). The default value is 0.27 (for Aluminum alloy A201).
Secondary Dendrite Arm Spacing	The secondary dendrite arm spacing is used for the calculation of the Carman-Kozeny Shape Factor.
Switching Function: Constant s	The non-dimensional constant s (see Eqn. (2137)). The default value is 100 because the shape is rather insensitive for values from 80 through 120.
Carman-Kozeny Shape Factor c	The shape factor c_s is a function of the dendrite dimension (typically secondary dendrite arm spacing d) and a shape constant c: $c_s = c/d^2$. The default value for c is 180.

- To define the field function, select the Field Function node. Choose the appropriate field function that you have defined.

6. To set numerical parameters of the mushy zone permeability model, select the Models > Average Mushy Zone Permeability node.



7. In the *Average Mushy Zone Permeability - Properties* window, set the following properties:

Use Forchheimer's Term	Toggle the contribution of the Forchheimer's term to the flow resistance of the Carman-Kozeny mushy zone model on and off (see Eqn. (2141) in The Mushy Zone Permeability Model - Carman-Kozeny method).
Ergun Coefficient	The Ergun coefficient from Eqn. (2141). The default value is 0.55.
Max. Porous Viscous Resistance	Sets a maximum level of resistance. This setting applies only to the cells with a solid volume fraction greater than zero, in all regions to which the respective physics continuum is assigned.
Mushy Zone Permeability Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed solution supplants the old solution. The default value is 0.7.

For simulations with different phases, the Mushy Zone Permeability for one cell is taken as a volume-weighted average for the whole mixture. The Mixture Properties > Mushy Zone Permeability node represents this operation.



Relative Velocity Field for Mushy Zone Permeability Model

Generally, the mushy zone permeability model applies a momentum source term to stop the flow of the solidified liquid. This is done by reducing the convective velocity to zero with respect to the laboratory coordinate system.

For certain applications, it is useful that you can instead specify a reference velocity field for the solidified liquid. This means that you can specify a velocity field other than zero for the solidified liquid.

Examples of applications where the specification of a relative velocity field is useful are:

- Centrifugal casting

  Liquid melt is injected through a nozzle splashing onto the inner wall of a rotating cylinder. When solidification takes place, the solidified liquid must rotate at the same angular velocity as the rotating cylinder. This means that the solidified liquid has zero velocity relative to the rotating wall.

- Continuous/Strand casting

  Liquid melt is injected at the top of an open-base mold. In the mold, a thin sheet of metal next to the mold wall solidifies. The solidified liquid has to move at the same velocity as the strand is pulled away.

The velocity with which the mold wall and the solidified liquid move (with zero relative velocity between them) is called the pull velocity.

Setting up the Relative Velocity field for Mushy Zone Permeability Model

Before you begin, you have to set up a stationary moving reference frame motion or a rigid body motion model. The motion that you specify can either be translational or rotational.

For more information, see Working with Motions.

When the mushy zone permeability model is selected for an Eulerian phase, the Mushy Zone Reference Frame Specification node is automatically added to the Physics Conditions node of each region.



To select the reference frame for the relative velocity field:

1. Select the Mushy Zone Reference Frame Specification node.
2. In the *Mushy Zone Reference Frame Specification - Properties* window, set the following reference frame options:

Option	Select the reference frame for the relative velocity field for the mushy zone permeability model.	
	Region Reference Frame	The relative velocity is specified with respect to the region reference frame. The region reference frame is the motion reference frame, which is selected in Physics Values > Motion Specification > Reference Frame. When you select this option, the Mushy Zone Relative Velocity node is added to the Physics Values node of the region.

Lab Frame	The relative velocity is specified with respect to the laboratory reference frame. When you select this option, the Mushy Zone Velocity node is added to the Physics Values node of the region.
Local Reference Frame	The relative velocity is specified with respect to a local reference frame with which the region is associated. When you select this option, the Mushy Zone Relative Velocity node is added to the Physics Values node of the region. Additionally, the Mushy Zone Local Reference Frame Specification node is added to the Physics Values node of the region.

Depending on your selection of the mushy zone reference frame specification, you have three options to specify the relative velocity field:

To specify the relative velocity field with respect to a region reference frame:

- Select the Physics Values > Mushy Zone Relative Velocity node and set its scalar profile method and value.



To specify the velocity field with respect to the laboratory reference frame:

- Select the Physics Values > Mushy Zone Velocity node and set its scalar profile method and value.



To specify the relative velocity field with respect to the local reference frame:

1. Select the Mushy Zone Local Reference Frame Specification node.
2. In the *Properties* window, select the local reference frame that resides in Tools > Reference Frames to describe the motion.

3. Select the Mushy Zone Relative Velocity node and set its scalar profile method and value.



You need to observe the following difference between a stationary moving reference motion and a rigid body motion when you specify the relative velocity value:

- For a rigid body motion, the mesh actually moves in the ReferenceFrame for Rotation. Therefore, you specify a relative velocity value of zero with respect to the Coordinate System in the *Mushy Zone Relative Velocity - Properties* window, so that the solidified liquid moves with the velocity specified for the rigid body motion.
- For a moving reference frame motion, the simulation is stationary and the mesh does not actually move. Therefore, you specify a non-zero relative velocity value with respect to the Coordinate System in the *Mushy Zone Relative Velocity - Properties* window, so that the solidified liquid moves with the velocity specified for the moving reference frame motion.

What Is the Relative Solid Volume Fraction of the Phase?

The relative solid volume fraction is a field function that is associated with the melting and solidification phase. This fraction indicates how much of the phase is in the solid state, as defined by Eqn. (2131).

Melting-Solidification Formulation

This section contains the following:

- Basic Melting-Solidification
- Flow Stop Mass Compensation
- The Slurry Viscosity Model - Metzner Method
- Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model
- The Mushy Zone Permeability Model - Carman-Kozeny method

Basic Melting-Solidification

The Melting-Solidification model that is implemented in STAR-CCM+ does not track the liquid-solid interface explicitly. Instead, the model uses an enthalpy formulation to determine the distribution of the solidified portion of the liquid-solid phase. Additional components of this process are the Slurry Viscosity and Mushy Zone Permeability models, and the flow stop functionality.

For basic melting-solidification, the enthalpy of the liquid-solid phase h_{ls}^* includes the latent heat of fusion h_{fusion}:

$$h_{\mathrm{ls}}^* = h_{\mathrm{ls}} + (1 - \alpha_s^*)h_{\mathrm{fusion}} \tag{2130}$$

where h_{ls} is the sensible enthalpy. The relative solid volume fraction α_s^* is defined as the portion of the volume of the liquid-solid phase which is in the solid state. In the enthalpy model, the relative solid volume fraction α_s^* is a function of temperature:

$$\alpha_s^* = \begin{cases} 1 & \text{if} \quad T^* < 0 \\ f(T^*) & \text{if} \quad 0 < T^* < 1 \\ 0 & \text{if} \quad 1 < T^* \end{cases} \tag{2131}$$

where T^* is the normalized temperature that is defined as:

$$T^* = \frac{T - T_{\mathrm{solidus}}}{T_{\mathrm{liquidus}} - T_{\mathrm{solidus}}} \tag{2132}$$

The function $f(T^*)$ is called the fraction solid curve. For a linear dependence between α_s^* and T^*, the solidification path is defined as:

$$f(T^*) = 1 - T^* \tag{2133}$$

The solidification path can also be specified as a table, in which case linear dependencies are assumed between the points in the table. The values of temperature in the table have to increase monotonically as the values of the relative solid fraction decrease monotonically.

If melting and solidification take place at one temperature (that is, $T_{\text{solidus}} = T_{\text{liquidus}}$), a linear solidification path is assumed and a small temperature interval of 0.002 K is automatically introduced.

The physical properties of the liquid-solid phase can be constant, or can be a function of temperature and the relative solid volume fraction. A choice of methods exists for preventing fluid movement in solidified regions.

Flow Stop Mass Compensation

The mass change $\Delta m_{p,r}$ of a phase (index $_p$) in a contiguous stopped region (index $_r$) since the last time step (denoted by an apostrophe, ') and the current iteration of the current time step, accumulated over all cells c_r of the contiguous stopped region is:

$$\Delta m_{p,r} = \sum_{c_r} (\rho_{p,c} - \rho'_{p,c}) V_c Y_p \tag{2134}$$

This expression sums, for all cells, the density change of the phase in the cell $\rho_{p,c} - \rho'_{p,c}$ in the current time step times the volume of the cell V_c and the volume fraction of that phase Y_p.

All fluid cells that are directly connected to the solid cells of the contiguous stopped region require a mass source equal to

$\frac{-\Delta m_{p,r}}{dt}$ to be applied per time step dt.

Once the phase mass rate is determined, it is applied to the liquid that is next to the contiguous solid region. The phase mass rate is weighted by the ratio of the face area between the liquid and the solid cell to the total surface area of the contiguous stopped region. The result is a small velocity component normal to the solid surface, which is similar to what is observed in cases of shrinkage or expansion of a solid.

The Slurry Viscosity Model - Metzner Method

In the early stages of solidification of an alloy, small crystals nucleate and grow in the melt. The mixture of the two phases is called a slurry. These crystals move with the liquid without interacting with each another. The presence of the solidified crystals in the liquid increases the flow resistance. This increased resistance is modeled by the slurry viscosity model where the increased flow resistance results from an increase in viscosity of the

phase mixture. The slurry viscosity model is applied for low values of the solid volume fraction ($\alpha_s^* \leq \alpha_{cr}$) whereby the upper limit of applicability is given by the critical solid fraction α_{cr}.

In the method of Metner [394], the variation in viscosity of the mixture is expressed as a function of the solid volume fraction α_s^* :

$$\mu^* = \mu_l \left[1 - \left(\frac{\alpha^* \cdot F_\mu(\alpha_s^*)}{A} \right) \right]^{-2} \tag{2135}$$

where μ^* is the effective viscosity of the mixture, μ_l is the dynamic viscosity of the liquid, F_μ is a switching function that is given by Eqn. (2136), and A is a crystal constant which depends on the aspect ratio and surface roughness of the crystal.

For details on the switching function F_μ, see Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model.

Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model

The Metzner slurry viscosity model and the Carman-Kozeny mushy zone permeability model are typically used in combination. Both models cover different ranges of solidification state. The applicability of each model depends on the extent to which the liquid has solidified. The state of solidification is expressed by the solid volume fraction α_s^* . The critical solid fraction α_{cr} separates the applicability ranges of the two models.

$\alpha_s^* \leq \alpha_{cr}$	Metzner Slurry Viscosity Model
$\alpha_{cr} < \alpha_s^*$	Carman-Kozeny Mushy Zone Permeability Model

For the Metzner slurry viscosity model and the Carman-Kozeny mushy zone model, a switching function ensures that each model is applied within its appropriate range of solid volume fraction. The switching function is applied as a multiplicative factor to the respective model equation (see Eqns. (2135) and (2138)).

The switching function for the slurry viscosity model is:

$$F_\mu(\alpha_s^*) = 0.5 - \frac{\arctan[s(\alpha_s^* - \alpha_{cr}^*)]}{\pi} \tag{2136}$$

where s is a switching function constant.

The switching function for the mushy zone permeability model is the inverse of Eqn. (2136):

$$F_K(\alpha_s^*) = 1 - F_\mu(\alpha_s^*) = 0.5 + \frac{\arctan[s(\alpha_s^* - \alpha_{cr}^*)]}{\pi} \tag{2137}$$

The switching function has a smooth transition near the non-dimensional critical solid volume fraction α_{cr}^* (see the figure below) where solid grains start to agglomerate to a significant extent:

The default value of the non-dimensional constant s is 100 because the shape is rather insensitive for values from 80 through 120. The default value for the critical solid volume fraction α_{cr}^* is 0.27 according to Baeckerud and others [391] for Aluminum alloy A201.

Oldenburger and Spera [395] initially suggested this switching function. The switching function is used here in a slightly modified formulation according to Chang and Stefanescu [393], on whose publication the implementation of this model is based.

The following figure depicts the switching function F_K for critical solid volume fraction $\alpha_{cr}^* = 0.27$ (default value).



The use of the switching function for both models is optional. It is recommended to activate the use of the switching function for both models when you use the Metzner method and the Karman-Cozeny method in combination. If the slurry viscosity is not reduced to zero above the critical

solid volume fraction through the usage of the switching function, it serves as a multiplicative factor to the effect of the mushy zone permeability. The Metzner slurry viscosity model can result in rather high slurry viscosities when used above the critical solid fraction.

If you combine the Metzner slurry viscosity model with a user-defined mushy zone permeability model, you can de-active the use of the switching function. Similarly, if you combine the Carman-Kozeny mushy zone permeability model with a user-defined slurry viscosity model, you can de-active the use of the switching function. Ensure that the effect of each model in its respective applicability range is accounted for within the user-defined field function.

The Mushy Zone Permeability Model - Carman-Kozeny method

For high solid volume fractions, equiaxed grains grow and start to agglomerate and to form dendritic regions. Dendrites also grow from cooled solid surfaces or the solidification front into the melt. A partially solidified stationary region permeated with dendrites is called a “mushy zone”.

The Carman-Kozeny permeability model provides a mushy zone for flow resistance, eliminating the need to set up porous media. In addition, the Carman-Kozeny Mushy Permeability method updates the velocity field in the melting and solidification process.

A common engineering application for this model is high-pressure die casting where melt is pressed through ducts with wall temperatures below solidus temperature at high velocities. The melt solidifies at the walls and dendrites grow from the solidification front into the melt. This partially solidified region causes an additional pressure loss in the duct that influences the melt mass flow. This pressure loss can be simulated using the Carman-Kozeny Mushy Permeability method.

The resistance (active only in cells with a solid volume fraction greater than zero) is applied to all regions to which the respective physics continuum is assigned.

The flow resistance in a mushy zone can be modeled similar to an isotropic porous medium. Therefore, to quantify the permeability, the Carman-Kozeny equation [392] for flow through a porous medium is used:

$$K = \frac{(1-\alpha_s^*)^3}{\alpha_s^{*2} F_K c_s} \tag{2138}$$

In this expression, $K\,[m^2]$ is the permeability and α_s^* is the non-dimensional solid volume fraction. $c_s = c/d^2\,[m^{-2}]$ where c is a shape constant, with a default value of 180, that is divided by the secondary dendrite arm spacing $d\,[m]$ and a non-dimensional switching function F_K given by Eqn. (2137).

For details on the switching function, see Switching Function for Metzner Slurry Viscosity and Carman-Kozeny Mushy Zone Permeability Model.

The momentum equation in continuous integral form is given by Eqn. (201).

The contribution of the mushy zone permeability to the momentum equation goes into $\mathbf{f}_p$:

$$\mathbf{f}_p = -\mathbf{P}\cdot\mathbf{v} \tag{2139}$$

where $\mathbf{P}$ is the porous resistance tensor, which is given by

$$\mathbf{P} = \mathbf{P}_\mathrm{v} + \mathbf{P}_\mathrm{i}\|\mathbf{v}\|\,. \tag{2140}$$

$\mathbf{P}_\mathrm{v}$ and $\mathbf{P}_\mathrm{i}$ are the viscous (linear) and the inertial (quadratic) resistance tensors, respectively.

The mushy zone permeability contributes to $\mathbf{P}_\mathrm{v}$ through

$$\mathbf{P}_{\mathrm{v,\ Mushy\ Zone}} = \left(\frac{\mu}{K} + \frac{C_F\rho_l\alpha_l^*}{\sqrt{K}}\|\mathbf{v}\|\right). \tag{2141}$$

where

$\frac{\mu}{K}$	Darcy's term
$\frac{C_F\rho_l\alpha_l^*}{\sqrt{K}}\|\mathbf{v}\|$	Forchheimer's term

$\mathbf{P}_\mathrm{v}$ can be considered a drag coefficient. The drag coefficient is given by the Darcy's and Forchheimer's terms, which are often referred to as the "viscous-drag" and "form-drag" terms.

Here, μ is the dynamic viscosity, C_F is the Ergun coefficient, ρ_l is the liquid phase density, and α_l^* is the liquid fraction given by

$$\alpha_l^* = (1 - \alpha_s^*)\,. \tag{2142}$$

The Forchheimer's term is optional and its contribution to the mushy zone permeability can be deactivated. The Forchheimer's term must be considered when relatively high interdendritic fluid-flow velocities are expected. Such high velocities can be important during the intensification phase in die or squeeze casting.

Melting-Solidification Bibliography

[391] Baeckerud, L., Chai, G. and Tamminen, J. 1992. *Solidification Characteristics of Aluminum Alloys 2*, American Foundry Society, Inc. Des Plaines, Illinois.

[392] Carman, P. C. 1937. "Fluid Flow Through Granular Beds", *Transactions of the Institution of Chemical Engineers*, **15**, pp. 150-166.

[393] Chang, S. and Stefanescu, D. M. 1996. "A Model for Macrosegregation and Its Application to Al-Cu Castings", *Metallurgical and Materials Transactions*, **27A**, pp. 2708-2721.

[394] Metzner, A. B. 1985. "Rheology of Suspensions in Polymeric Liquids", *Journal of Rheology*, **29**(6), pp. 739-775.

[395] Oldenburg, C. M., and Spera, F. J. 1992. "Hybrid Model for Solidification and Convection", *Numerical Heat Transfer, Part B*, **21**, pp. 217-229.

[396] Swaminathan, C.R. and Voller, V.R. 1992 "A general enthalpy method for modeling solidification processes", *Met. Trans. B*, **22**(B), pp. 651-664.

[397] Teskeredzic, A., Demirdzic, I. and Muzaferija, S. 2002. "Numerical method for heat transfer, fluid flow, and stress analysis in phase-change problems", *Numerical Heat Transfer*, Part B, **42**, pp. 437-459.

[398] Voller, V.R. 1997 "An overview of numerical methods for solving phase change problems", in: *Advances in Numerical Heat Transfer*, W.J. Minkowycz and E.M. Sparrow Eds., Taylor & Francis, pp. 341-378.

VOF Formulation

This section describes:

- The mathematical formulation of the VOF model.
- The normalized variable diagram.
- The HRIC convection discretization scheme.
- The mathematical formulation of the surface tension model.
- The latent heat of vaporization.

Basic VOF Model Equations

The VOF model description assumes that all immiscible fluid phases present in a control volume share velocity, pressure, and temperature fields. Therefore, the same set of basic governing equations describing momentum, mass, and energy transport in a single-phase flow (see

Modeling Flow Using a Segregated Approach) is solved. The equations are solved for an equivalent fluid whose physical properties are calculated as functions of the physical properties of its constituent phases and their volume fractions.

The equations are:

$$\rho = \sum_i \rho_i \alpha_i \tag{2143}$$

$$\mu = \sum_i \mu_i \alpha_i \tag{2144}$$

$$c_p = \sum_i \frac{(c_p)_i \rho_i}{\rho} \alpha_i \tag{2145}$$

where:

$$\alpha_i = \frac{V_i}{V}$$

is the volume fraction and ρ_i, μ_i and $(c_p)_i$ are the density, molecular viscosity and specific heat of the i th phase.

The conservation equation that describes the transport of volume fractions α_i is:

$$\frac{d}{dt}\int_V \alpha_i dV + \int_S \alpha_i(\boldsymbol{v} - \boldsymbol{v}_g) \cdot d\boldsymbol{a} = \int_V \left(s_{\alpha_i} - \frac{\alpha_i}{\rho_i}\frac{D\rho_i}{Dt} \right) dV \tag{2146}$$

where s_{α_i} is the source or sink of the i th phase, and $D\rho_i/Dt$ is the material or Lagrangian derivative of the phase densities ρ_i.

If a non-zero sharpening factor is specified, an additional term is added to the VOF transport equation:

$$\nabla \cdot (\boldsymbol{v}_{c_i} \alpha_i (1 - \alpha_i)) \tag{2147}$$

where:

- α_i is the volume fraction of phase i
- $\boldsymbol{v}_{c_i}$ is defined as follows:

$$\boldsymbol{v}_{c_i} = C_\alpha \times |\boldsymbol{v}| \frac{\nabla \alpha_i}{|\nabla \alpha_i|}$$

- C_α is the sharpening factor that is specified in the Volume of Fluid (VOF) node properties.
- $\boldsymbol{v}$ is the fluid velocity

If there is a large time variation of phase volume fractions α_i, there is a large time variation of the mixture density ρ which features in the continuity equation, Eqn. (42). Since this unsteady term cannot be linearized in terms of pressure and velocity, it acts as a large source term which can be "unpleasant" for a numerical treatment within a segregated solution algorithm. Therefore, Eqn. (42) is rearranged in the following, non-conservative form:

$$\int_A (\boldsymbol{v} - \boldsymbol{v}_g) \cdot d\mathbf{a} = \sum_i \int_V \left(s_{\alpha_i} - \frac{\alpha_i}{\rho_i} \frac{D\rho_i}{Dt} \right) dV \tag{2148}$$

In the case when phases have constant densities and have no sources, the continuity equation reduces to $\nabla \cdot \boldsymbol{v} = 0$.

The Normalized Variable Diagram

The normalized variable diagram (NVD) is useful for analyzing boundedness properties of convective discretization schemes. Part (a) of the figure below shows three cells in the vicinity of a cell face *f*, across which the velocity $\boldsymbol{v}_f$ is known. The nodal variable values are labeled α_D, α_C and α_U, representing the downwind, central, and upwind positions relative to each other.



Upwind, downwind, and central cells that are used in the analysis (left) and convection boundedness criterion in the NVD diagram (right)

The normalized variable $\xi(\boldsymbol{r}, t)$ in the vicinity of face *f* is defined as:

$$\xi(\boldsymbol{r}, t) = \frac{\alpha(\boldsymbol{r}, t) - \alpha_U}{\alpha_D - \alpha_U} \tag{2149}$$

The normalized face value:

$$\xi_f = \frac{\alpha_f - \alpha_U}{\alpha_D - \alpha_U} \tag{2150}$$

calculated by any differencing scheme that uses only nodal values of α at points U, C, and D, can be written in the form:

$$\xi_f = \xi_f(\xi_C) \tag{2151}$$

where:

$$\xi_C = \frac{\alpha_C - \alpha_U}{\alpha_D - \alpha_U} \tag{2152}$$

To avoid non-physical oscillations in the solution, α_C (and consequently α_f) has to be locally bounded between α_U and α_D, meaning:

$$\alpha_U \le \alpha_C \le \alpha_D \quad \text{or} \quad \alpha_D \le \alpha_C \le \alpha_U \tag{2153}$$

If this criterion is satisfied at every point in the domain, the entire solution is free of non-physical oscillations.

The boundedness criterion for convection differencing schemes can be presented in the NVD diagram, showing ξ_f as a function of ξ_C, as the hatched region in part (b) of the above Figure (including the line $\xi_f = \xi_C$). It can also be expressed through the following conditions:

- For $0 \le \xi_C \le 1$ the bounded region lies above the line $\xi_f = \xi_C$ and below $\xi_f = 1$.
- For $\xi_C < 0$ and $\xi_C > 1$, ξ_f is equal to ξ_C.

NVD is concerned with convective transport alone. If sources or sinks are present, the conditions that are given in Eqn. (2153) can be violated. The importance of the boundedness criterion is especially clear in the case of variables which have physical bounds; for example, the phase volume fraction cannot become negative, or larger than unity.

The HRIC Convection Discretization Scheme

An important quality of an immiscible phase mixture (for example, air and water) is that the fluid components always remain separated by a sharp interface. A simple higher-order scheme (for example, CD or LUD) would fail in approximating large spatial variations of phase volume fractions, which are best represented by the Heaviside unit step function. On the other hand, the High-Resolution Interface Capturing (HRIC) scheme is designed to mimic the convective transport of immiscible fluid components, resulting in a scheme that is suited for tracking sharp interfaces.

The scheme is based on the NVD diagram that was shown earlier. The normalized face value ξ_f is calculated as:

$$\xi_f = \begin{cases} \xi_C & \text{if} \quad \xi_C < 0 \\ 2\xi_C & \text{if} \quad 0 \le \xi_C < 0.5 \\ 1 & \text{if} \quad 0.5 \le \xi_C < 1 \\ \xi_C & \text{if} \quad 1 \le \xi_C \end{cases} \tag{2154}$$

The calculated value of ξ_f is further corrected according to the local Courant number:

$$Cu = \frac{\boldsymbol{v} \cdot \boldsymbol{a}_f}{V_{P_C}} \delta t \tag{2155}$$

This correction is to take into account the availability criterion. This criterion says that the amount of one fluid that is convected across a cell face during a time-step is always less than or equal to the amount available in the donor cell. This correction is made according to the expressions:

$$\xi_f^* = \begin{cases} \xi_f & \text{if} \quad Cu < Cu_l \\ \xi_C + (\xi_f - \xi_C)\dfrac{Cu_u - Cu}{Cu_u - Cu_l} & \text{if} \quad Cu_l \le Cu < Cu_u \\ \xi_C & \text{if} \quad Cu_u \le Cu \end{cases} \tag{2156}$$

and it plays a role only in transient simulations. Default values for Cu_l and Cu_u are 0.5 and 1. They are introduced to control blending of HRIC and UD schemes depending on the Courant number. For values of $Cu < Cu_l$ HRIC is used, for $Cu_l < Cu < Cu_u$ a blend of HRIC and UD is used, and for $Cu_u < Cu$ UD is used.

The blending is introduced in the case when a large time variation of the free surface shape is present, and the time-step is too large to resolve details of it. It brings stability and robustness to the scheme, and is especially important when you use the second-order discretization in time. In this case, smaller values of Cu_u and Cu_l help to promote convergence. Smaller values activate UD sooner, and the calculation is more stable. At the same time, the interface is more "smudged" in regions with large Cu numbers.

For problems which have a steady-state solution, the values of Cu_u and Cu_l must be large, such that HRIC is used irrespective of the time-step selected. Any number larger than the maximum CFL is appropriate.

The front sharpening aspect of the downwind scheme is just what is needed when the interface is perpendicular to the flow. However, when the interface is parallel to the flow direction, the downwind scheme tends to

wrinkle it. The final correction of ξ_f is based on the angle θ between the normal to the interface $\boldsymbol{n}_i$ and the cell-face surface vector $\boldsymbol{a}_f$, as shown below:

$$\xi_f^{**} = \xi_f^{*}(\cos\theta)^{C_\theta} + \xi_C(1-(\cos\theta)^{C_\theta}) \tag{2157}$$

The default value of the angle factor C_θ is 0.05. If the free surface is not smooth and not following the grid lines, increase its value.

The cell-face value ξ_f is now calculated as:

$$\alpha_f^{HRIC} = \xi_f^{**}(\alpha_D - \alpha_U) + \alpha_U \tag{2158}$$



Interface between two fluids and the notation used.

Surface Tension Formulation

The surface tension force is a tensile force tangential to the interface separating two fluids. It works to keep the fluid molecules at the free surface in contact with the rest of the fluid. Its magnitude depends mainly on the nature of the fluid pair and on temperature. For a curved interface, the surface tension force $\boldsymbol{f}_\sigma$ can be resolved into two components:

$$\boldsymbol{f}_\sigma = \boldsymbol{f}_{\sigma,n} + \boldsymbol{f}_{\sigma,t} \tag{2159}$$

where:

$$\boldsymbol{f}_{\sigma,n} = \sigma K \boldsymbol{n} \tag{2160}$$

and:

$$\boldsymbol{f}_{\sigma,t} = \frac{\partial\sigma}{\partial t}\boldsymbol{t} \tag{2161}$$

and σ is the surface tension coefficient. $\boldsymbol{n}$ is the unit vector normal to the free surface and directed from liquid to gas. $\boldsymbol{t}$ is the unit vector in the tangential direction to the free surface. K is the mean curvature of the free surface.

For a constant surface tension coefficient σ, the tangential force is zero and the surface tension results in a force normal to the interface ($\boldsymbol{f}_\sigma = \boldsymbol{f}_{\sigma,n}$).

When the surface tension coefficient varies along the surface, which can be caused by temperature differences, the tangential part does not vanish. In this case, Marangoni or Bénard convection can develop tangential to the free surface.

In the VOF multiphase model case, the surface tension force is calculated according to the continuum surface force (CSF) model. That is, it uses the smooth field of the phase volume fraction α_i to calculate a vector normal to the interface:

$$\boldsymbol{n} = \nabla\alpha_i \tag{2162}$$

The curvature of the interface can therefore be expressed in terms of the divergence of the unit normal vector $\boldsymbol{n}$, as follows:

$$K = -\nabla\bullet\frac{\nabla\alpha_i}{|\nabla\alpha_i|} \tag{2163}$$

Now the normal component of the surface tension force $\boldsymbol{f}_{\sigma,n}$ can be expressed as:

$$\boldsymbol{f}_{\sigma,n} = -\sigma\nabla\bullet\left(\frac{\nabla\alpha_i}{|\nabla\alpha_i|}\right)\nabla\alpha_i \tag{2164}$$

When the surface tension coefficient σ is constant, $\boldsymbol{f}_\sigma = \boldsymbol{f}_{\sigma,n}$.

For a variable surface tension coefficient, the tangential force is evaluated as:

$$f_{\sigma,t} = (\nabla\sigma)_t|\nabla\alpha_i| \tag{2165}$$

where $(\nabla\sigma)_t$ is the gradient of the surface tension coefficient in the tangential direction.

Latent Heat of Vaporization

The specific latent heat of vaporization h_{lat} is the amount of heat that is required to convert unit mass of a liquid into the vapor without a change in temperature. The temperature at which vaporization takes place is called saturation temperature T_{sat}.

The latent heat is also the difference between the enthalpies of the vapor $h_v(T_{sat})$ and the liquid $h_l(T_{sat})$ at the temperature of saturation:

$$h_{lat} = h_v(T_{sat}) - h_l(T_{sat}) \tag{2166}$$

$$h_v(T_{sat}) = h_{0_v} + C_{p_v}(T_{sat} - T_{0_v}) \tag{2167}$$

$$h_l(T_{sat}) = h_{0_l} + C_{p_l}(T_{sat} - T_{0_l}) \tag{2168}$$

In problems with boiling, the heats of formation h_{0_l} and h_{0_v}, and the standard state temperatures T_{0_l} and T_{0_v} of the liquid and the vapor phase must be specified such that the above expression is satisfied. The latent heat is specified in this way. For stability reasons, set either h_{0_v} or h_{0_l} to zero, and the other to an appropriate value.

VOF Bibliography

[399] J.U. Brackbill, D.B. Kothe and C. Zemach, A Continuum method for modeling surface tension, J. Comp. Physics, 100, 335-354 (1992).

[400] S. Muzaferija and M. Peric, Computation of free surface flows using interface-tracking and interface-capturing methods, Chap. 2 in O. Mahrenholtz and M. Markiewicz (eds.), Nonlinear Water Wave Interaction, Computational Mechanics Publications, WIT Press, Southampton, 1999.

Working with VOF Waves

VOF waves are used to simulate surface gravity waves on a light fluid-heavy fluid interface. They are for three-dimensional cases only, and are used with the Volume of Fluid (VOF) multiphase model in STAR-CCM+. They are typically used with the 6-DOF Motion model for marine applications. When created, VOF Waves provide field functions that can be used to initialize the VOF calculation and to provide suitable profiles at boundaries. This procedure is exemplified in the *Boat in Head Waves* tutorial.

The VOF Waves model, when selected, also activates a region condition and a boundary condition to control damping of the waves near boundaries in order to reduce wave reflections.

The VOF Waves model node has properties, and contains the Waves manager node, which has its own pop-up menu, as shown in the following screenshot.



VOF Waves Model Properties

Damping Constant 1	Specifies f_1 in Eqn. (2179), used in VOF wave damping.
Damping Constant 2	Specifies f_2 in Eqn. (2179), used in VOF wave damping.
Damping Exponent	Specifies n_d in Eqn. (2180), used in VOF wave damping.
Enable Theory Comparison	When this option is activated, the VOF Waves field functions are also updated in the inner fluid domain. (When this option is deactivated, the VOF Waves field functions are only updated on the boundaries of the fluid domain.) Hence, it is possible to plot the entire scalar field of the VOF Waves field functions. This plot shows the solution of Stokes' wave theory of the selected wave with respective order. You can compare this theoretical solution to a plot of the volume fraction obtained through the solution of the transport equation.

Selecting the VOF Waves Model

The VOF Waves model is available for three-dimensional simulations only. To activate this model, select Eulerian Multiphase for the material model, then select Volume of Fluid (VOF). The VOF Waves option appears in the *Optional*

Models group box.



Creating VOF Waves

Five types of VOF wave can be created using the right-click menu of the Waves manager node: Flat, FirstOrder, FifthOrder, Superposition, and Irregular.



When created, VOF waves appear as their own node within the Waves manager having a default name. In common with other objects, they can be renamed or deleted.

What Is a Flat Wave?

A flat wave represents a calm plane of water, that is, the least amount of current. It is created by right-clicking on the Waves manager node and choosing New > Flat. When created, its node appears within the Waves manager. Edit the properties of this node as required.



Flat Wave Properties

Point on Water Level	For a flat VOF wave, this property defines the position of the water surface, that is, the water level.
Vertical Direction	The vector normal to the water surface.
Current	Velocity of the water or other heavy fluid.
Wind	Velocity of the air or other light fluid.
Light Fluid Density	This property is required for the hydrostatic field function that is created automatically with the wave. The default value is for air. If you are simulating any other substance, change this value manually.
Heavy Fluid Density	This property is required for the hydrostatic field function that is created automatically with the wave. The default value is for water. If you are simulating any other substance, change this value manually.

What Is a First Order Wave?

A first order wave is modeled with a first order approximation to the Stokes theory of waves. This approximation generates waves that have a regular periodic sinusoidal profile. It is created by right-clicking on the Waves manager and selecting New > FirstOrder. When created, edit the properties of

its node to give the required waveform.



First Order Wave Properties

Point on Water Level	The mean water level on the leading front of the wave.
Vertical Direction	The vector normal to the mean water surface.
Current	Velocity of water or other heavy fluid.
Wind	Velocity of air or other light fluid.
Advancing Direction	Direction traveled by the wave.
Amplitude	Height of the wave crest.
Specification Type	This property sets the quantity to use in defining the length or duration of the wave. **Wave Length**: The wave is defined by its wavelength. A WaveLengthSpecification node is added as a child node. **Wave Period**: The wave is defined by its time period. A Wave Period Specification node is added as a child node.
Deep Water Approximation	When this option is activated (default), the dispersion relation for first order waves is independent of the water (or other heavy fluid) depth (see Eqn. (2175)). When this option is deactivated, Eqn. (2174) applies.
Light Fluid Density	This property is required for the hydrostatic field function that is created automatically with the wave. The default value is for air. If you are simulating any other substance, change this value manually.

Heavy Fluid Density — This property is required for the hydrostatic field function that is created automatically with the wave. The default value is for water. If you are simulating any other substance, change this value manually.

What Is a Fifth Order Wave?

A fifth order wave is modeled with a fifth order approximation to the Stokes theory of waves. This wave more closely resembles a real wave than one generated by the first order method. The wave profile and the wave phase velocity depend on water depth, wave height and current. This wave theory is only valid for Ursell numbers less than 30 [402]. The formulation can be found in Fenton [403]. This wave is created by right-clicking on the Waves manager and selecting New > FifthOrder. When created, edit the properties of its node to give the required waveform.



Fifth Order Wave Properties

Point on Water Level	The mean water level on the leading front of the wave.
Vertical Direction	The vector normal to the mean water surface.
Current	Velocity of water or other heavy fluid.
Wind	Velocity of air or other light fluid.
Advancing Direction	Direction traveled by the wave.
Additional Velocity	This velocity is added to the combined velocity value that the current, wind, and wave predict. When solving for the flow around a moving vessel in the reference frame of the vessel, set this value to the inverse of the velocity of the vessel.
Specification Type	This property sets the quantity to use in defining the length or duration of the wave. **Wave Length**: The wave is defined by its wavelength. A WaveLengthSpecification node is added as a child node. **Wave Period**: The wave is defined by its time period. A Wave Period Specification node is added as a child node.
Wave Height	Vertical distance from crest to trough.
Water Depth	Depth of the body of water being simulated.
Light Fluid Density	Needed for the hydrostatic field function. The default value is for air. If you are simulating any other substance, change this value manually.

Heavy Fluid Density	Needed for the hydrostatic field function. The default value is for water. If you are simulating any other substance, change this value manually.

What Is a Superposition Wave?

A superposition wave is a linear superposition of different "partial" waves. It can be used to simulate more complex wave phenomena, for example a cross sea or spectral waves. (A cross sea is a sea state with two wave systems traveling at oblique angles. This state can occur when water waves from one weather system continue despite a shift in wind.)

This wave is created by right-clicking on the Waves manager node and choosing New > Superposition. When created, its node appears within the Waves manager. Edit the properties of this node as required.



This node contains a SubWaves sub-node, which in turn contains superposing VOF waves.

Superposition Wave Properties

Point on Water Level	The mean water level on the leading front of the wave.
Vertical Direction	The vector normal to the mean water surface.
Current	Velocity of water or other heavy fluid.
Wind	Velocity of air or other light fluid.
Light Fluid Density	Needed for the hydrostatic field function. The default value is for air. If you are simulating any other substance, change this value manually.
Heavy Fluid Density	Needed for the hydrostatic field function. The default value is for water. If you are simulating any other substance, change this value manually.

What Is a First-Order Superposing VOF Wave?

The first-order superposing VOF wave is the wave type that is used for superposition of waves, that is, to form a superposition wave. It is used to describe one “partial” wave. These waves are stored in the SubWaves

manager node, which you right-click to create new superposing waves in the pop-up menu. Edit the properties of first-order superposing waves as required.



First-Order Superposing VOF Wave Properties

Advancing Direction	Direction traveled by the wave.
Amplitude	Height of the wave crest.
Specification Type	This property sets the quantity to use in defining the length or duration of the wave. **Wave Length**: The wave is defined by its wavelength. A WaveLengthSpecification node is added as a child node. **Wave Period**: The wave is defined by its time period. A Wave Period Specification node is added as a child node.
Phase	This value is the initial wave phase at t = 0s for the particular partial wave. It is used here to describe the phase difference between the different partial waves. The wave phase is the fraction of a wave cycle that is measured in radians.
Enabled	Uses a checkbox to include (activated) or exclude (deactivated) the effects of this wave in the simulation.

What Is an Irregular Wave?

An irregular wave can be used to describe a short-term sea state by a wave spectrum, that is, the power spectral density function of the vertical sea surface displacement. The Pierson-Moskowitz spectrum and JONSWAP spectrum are frequently applied for wind seas. Both spectra describe wind sea conditions that often occur for the most severe sea states [404]. This

wave is created by right-clicking on the Waves manager node and choosing New > Irregular. When created, its node appears within the Waves manager. Edit the properties of this node as required.



Irregular Wave Properties

Point on Water Level	This value is the mean water level on the leading front of the wave.
Vertical Direction	The vector normal to the mean water surface.
Current	Velocity of water or other heavy fluid.
Wind	Velocity of air or other light fluid.
Advancing Direction	Direction traveled by the wave.
Significant Wave Height	Defined as the average height (trough to crest) of the highest one-third waves.
Peak Wave Period	The wave period that is determined by the inverse of the frequency at which a wave energy spectrum has its maximum value.
Spectrum Type	This property sets the quantity to use in defining the length or duration of the wave. **Pierson-Moskowitz**: The Pierson-Moskowitz spectrum was proposed originally for a fully developed sea. **JONSWAP**: Adds a JONSWAPSpecification node as a child node. The JONSWAP spectrum extends the Pierson-Moskowitz spectrum to include fetch limited seas, describing developing sea states. The JONSWAP spectrum is expected to be a reasonable model for the following (in SI units): $3.6 < \frac{\text{Peak Wave Period}}{\sqrt{\text{Significant Wave Height}}} < 5$

Light Fluid Density	Needed for the hydrostatic field function. The default value is for air. If you are simulating any other substance, change this value manually.
Heavy Fluid Density	Needed for the hydrostatic field function. The default value is for water. If you are simulating any other substance, change this value manually.

Expert Properties

Number of Frequencies	Number of partial waves that are used to create the irregular wave. The wave amplitudes of the partial wave are computed according to the wave spectrum.

What Is the JONSWAP Specification?

The JONSWAP specification lets you set peak shapes and spectral widths for the random wave. The JONSWAPSpecification node represents this specification. This node appears when the Spectrum Type property of the irregular wave is set to `JONSWAP`. Edit the properties of this node as required.



JONSWAPSpecification Properties

Peak Shape Parameter	Sets the peak shape, γ in Eqn. (2178). Average value is γ = 3.3.
Spectral Width A	Sets σ_a (see Eqn. (2178)). Average value is 0.07.
Spectral Width B	Sets σ_b (see Eqn. (2178)). Average value is 0.09.

Specifying Wavelength and Period

First and fifth order waves require the specification of either the wavelength or the wave period. The Specification Type property of the first and fifth order waves is used to select which quantity to set.

Setting the Wave Length

The WaveLengthSpecification node is used to set the wavelength. The corresponding wave period is displayed when the simulation has been initialized.

WaveLengthSpecification Properties

WavePeriod	The wave period T in Eqn. (2172) (read only).
WaveLength	The wavelength λ in Eqn. (2173). Enter using appropriate units of length.

Setting the Wave Period

The Wave Period Specification node is used to set the wave period. The corresponding wavelength is displayed when the simulation has been initialized.

Wave Period Specification Properties

Wave Period	The wave period T in Eqn. (2172). Enter using appropriate units of time.
Wave Length	The wavelength λ in Eqn. (2173) (read only).

Region Conditions Reference

The VOF Waves model activates the following region condition:

- VOF wave damping option

VOF Wave Damping Option

The VOF Waves model includes a VOF Wave damping capability. This means that a VOF wave can be damped in the vicinity of selected boundaries to reduce wave oscillation near those boundaries. Activating the damping introduces vertical resistance to the vertical motion.

In the particular example that is illustrated here, damping is applied toward the outlet for a distance of 100 meters. This quantity is entered in the Value property of the Constant sub-node of the VofWave Damping Length node, which is a scalar profile for the region.



In addition, the VOF wave damping option of the outlet boundary is activated.

These settings are visible in the scalar scene in the following screenshot. The wave damping takes place along a distance of 100 meters up to the outlet (at right).



According to Eqn. (2179) there is a resistance term added to the equation for w-velocity [401]. In order to be consistent with all effects, the effect of the damping is included in the wave profiles that are used for initialization.

Properties

VOF Damping Option	Specifies whether VOF damping is used for the region.	
	No Damping	Damping is off.
	Damping	Damping is activated. A VOF Damping Length node is added to the Physics Values node.

Region Values Reference

The following region value can be activated when the VOF Waves model is used:

- VOF damping length

VOF Damping Length

This region value becomes available when the VOF Wave Damping Option region condition has its VOF Damping Option property set to `Damping`.

This setting is the distance from the VOF wave damping-enabled boundary at which the damping starts. (See Eqn. (2180), the starting point for wave damping.)

This value is entered as a scalar profile.

Boundary Conditions Reference

The VOF Waves model activated the following boundary condition:

- VOF wave damping boundary option

VOF Wave Damping Boundary Option

When the VOF Waves model activates this boundary value, it is only available in velocity inlet, pressure outlet, and symmetry boundaries.

Properties

Method	Specifies whether VOF damping is used for the boundary.	
	No	Damping is off.
	Yes	Damping is activated for the boundary.

Field Functions for VOF Waves

The following field functions are made available to the simulation when VOF waves are created. Some of these field functions are added when any VOF wave is created, while others appear as a result of the creation of a particular type. They define initial conditions and boundary conditions for objects floating in water or any other liquid.

Hydrostatic Pressure of VofWave x	Hydrostatic pressure of a particular VOF wave, where x equals the total number of waves of this type in the simulation. This setting can be used to initialize the pressure field and on pressure boundaries.

Velocity of Vof Wave x	The velocity vector field of a particular VOF wave. Appears when a VOF wave is added to the simulation, where x equals the total number of waves in the simulation. This setting can be used for initial conditions and for the velocity specification on boundaries.
Volume Fraction of Heavy Fluid of VofWave x	Volume fraction of heavy fluid (such as water) of a particular VOF wave, where x equals the total number of first- or fifth-order waves in the simulation. This setting can be used for the initial volume fraction and for the volume fraction specification at boundaries.
Volume Fraction of Light Fluid of VofWave x	Volume fraction of light fluid (such as air) of a particular VOF wave, where x equals the total number of first- or fifth-order waves in the simulation. This setting can be used for the initial volume fraction and for the volume fraction specification at boundaries.

VOF Waves Formulation

First Order VOF Waves

The equation for horizontal velocity is:

$$\upsilon_h = a\omega\cos(\mathbf{K}\cdot\mathbf{x}-\omega t)e^{Kz} \tag{2169}$$

The equation for vertical velocity is:

$$\upsilon_v = a\omega\sin(\mathbf{K}\cdot\mathbf{x}-\omega t)e^{Kz} \tag{2170}$$

The equation for surface elevation is:

$$\eta = a\cos(\mathbf{K}\cdot\mathbf{x}-\omega t) \tag{2171}$$

where a is the wave amplitude, ω is the wave frequency, $\mathbf{K}$ is the wave vector, K is the magnitude of the wave vector, and z is the vertical distance from the mean water level.

The wave period T is defined as:

$$T = \frac{2\pi}{\omega} \tag{2172}$$

The wavelength λ is defined as:

$$\lambda = \frac{2\pi}{K} \tag{2173}$$

The dispersion relation (between wave period T and wave length λ) for first order waves in finite water depth d is:

$$T = \left[\frac{g}{2\pi\lambda}\tanh\left(\frac{2\pi d}{\lambda}\right)\right]^{-1/2} \tag{2174}$$

Whereas for infinite water depth, the dispersion relation is:

$$\lambda = \frac{gT^2}{2\pi} \tag{2175}$$

This means that the wave shape is independent of depth.

Fifth Order VOF Waves

The fifth order VOF waves are based on work by Fenton [403].

The Ursell number U_R [402] is defined as:

$$U_R = \frac{H\lambda^2}{d^3} \tag{2176}$$

where H, is the wave height λ is the wavelength and d is the depth of the water.

Irregular VOF Waves

The irregular wave can use either the Pierson-Moskowitz spectrum or the JONSWAP spectrum [404].

The Pierson-Moskowitz spectrum $S_{PM}(\omega)$ is defined as:

$$S_{PM}(\omega) = \frac{5}{16}(H_S^2\omega_p^4)\omega^{-5}\exp\left(-\frac{5}{4}\left(\frac{\omega}{\omega_p}\right)^{-4}\right) \tag{2177}$$

where $\omega_p = (2\pi)/(T_p)$ represents the angular spectral peak frequency.

The JONSWAP spectrum $S_J(\omega)$ is expressed as:

$$S_J(\omega) = A_\gamma S_{PM}(\omega)\gamma^{\exp\left(-0.5\left(\frac{\omega-\omega_p}{\sigma\omega_p}\right)^2\right)} \tag{2178}$$

where $S_{PM}(\omega)$ is the Pierson-Moskowitz spectrum, γ is the non-dimensional peak shape parameter, $A_\gamma = 1 - 0.287\ \ln(\gamma)$ is a normalizing factor, and σ is the spectral width parameter: σ_a for $\omega \le \omega_p$ and σ_b for $\omega > \omega_p$.

Wave Damping

Damping of waves is possible by introducing resistance to vertical motion. STAR-CCM+ implements the method by Choi and Yoon [401], adding a resistance term to the equation for w-velocity:

$$S_z^{\mathrm{d}} = \rho(f_1 + f_2|w|)\frac{e^\kappa - 1}{e^1 - 1}w \tag{2179}$$

with:

$$\kappa = \left(\frac{x - x_{\mathrm{sd}}}{x_{\mathrm{ed}} - x_{\mathrm{sd}}}\right)^{n_d} \tag{2180}$$

where x_{sd} is the starting point for wave damping (propagation in the x-direction); x_{ed} is the end point for wave damping (boundary); f_1, f_2 and n_d are parameters of the damping model; and w is the vertical velocity component.

VOF Waves Bibliography

[401] Choi, Junwoo and Sung, Bum Yoon. 2009. "Numerical simulations using momentum source wave-maker applied to RANS equation model", *Coastal Engineering*, **56** (10), pp. 1043 - 1060.

[402] Det Norske Veritas. 2007. "Environmental Conditions and Environmental Loads", Recommended Practice DNV-RP-C205, April 2007, pp. 27.

[403] Fenton, John D. 1985. "A Fifth-Order Stokes Theory for Steady Waves", *J. Waterway, Port, Coastal and Ocean Eng.*, **111** (2), pp. 216-234.

[404] Det Norske Veritas. 2007. *Environmental Conditions and Environmental Loads*, Recommended Practice DNV-RP-C205, April 2007 p. 33.

Modeling Dynamic Fluid Body Interaction

Use the DFBI (Dynamic Fluid Body Interaction) module to simulate the motion of a rigid body in response to pressure and shear forces the fluid exerts, and to additional forces you define. STAR-CCM+ calculates the resultant force and moment acting on the body due to all influences, and solves the governing equations of rigid body motion to find the new position of the rigid body.



Multiple rigid bodies can be modeled, with optional body couplings between them. The available couplings are the catenary coupling and the linear spring coupling.

The DFBI manager node only appears when one of the following DFBI motions has been created and applied to a region:

- DFBI Rotation and Translation
- DFBI Embedded Rotation
- DFBI Morphing
- DFBI Superposed Rotation

What Is a Rigid Body?

A rigid body (defined using a 6-DOF Body in STAR-CCM+) can be defined as an object in which the relative distance between internal points does not change. Using the standard Cartesian coordinate system as a reference frame, it is possible for a rigid body to move (translate) along each of the three axes, and to rotate about the axes.



In the figure above:

- T_x, T_y, and T_z are the translations in the x, y, and z directions respectively.
- R_x, R_y, and R_z are the rotations about the x, y, and z axes respectively.

For rigid bodies, it is sufficient to model the motion of the center of mass of the body alone. The relative motion of any other part of the body can be extrapolated from this center of mass. It is necessary to know the moments of inertia of the body about a fixed reference point (which is normally the center of mass) before the rotational motion can be known.

What Is a DFBI Rotation and Translation Motion?

This motion solves the equations of rigid body motion for all 6-DOF Bodies. When it is active, the 6-DOF Solver and 6-DOF Motion nodes are added to the Solvers manager. The DFBI Rotation and Translation motion node, which has no properties, is subject to common menu items.



What Is a DFBI Embedded Rotation Motion?

The DFBI Embedded Rotation motion involves a floating object in a sphere or cylinder. The simulation includes a sliding internal interface between the sphere and the rest of the mesh, and the sphere rotates. This approach allows for larger rotating angles while creating a fine mesh only within the sphere. This approach is an alternative to a higher resolution for the entire mesh, which would be more computationally expensive.

The DFBI Embedded Rotation motion node is subject to common menu items.



For this model, the geometrical center of the sphere must coincide with the initial center of mass of the 6-DOF body.

What Is a DFBI Morphing Motion?

The DFBI Morphing motion is a variant of the Morphing motion that is used for DFBI (Dynamic Fluid Body Interaction). It activates customized boundary conditions so that you can use morphing in DFBI without making time-consuming adjustments.

The DFBI Morphing motion node is subject to common menu items.



What Is a DFBI Superposed Rotation Motion?

The DFBI Superposed Rotation option superimposes an additional fixed body rotation in addition to the DFBI motion. This option makes it possible to model propellers attached to marine vessels, for example. The DFBI Superposed Rotation motion node has properties and is subject to common menu items.



DFBI Superposed Rotation Motion Properties

Direction	The direction of the axis. Specified in the Body Local Coordinate system.
Origin	The origin of the axis. Specified in the Body Local Coordinate system.
Rotation Rate	The angular velocity of the rotation about the axis.

What Is the Six DOF Solver?

The 6-DOF Solver computes fluid forces, moments, and gravitational forces on a 6-DOF body. Pressure and shear forces are integrated over the surfaces of the 6-DOF bodies. More forces and moments can be added.



The forces and moments acting on the 6-DOF body are used to compute the translational motion of the center of mass of the body and the angular motion of the orientation of the body.

Six DOF Solver Properties

Maximum Number of Iterations — Maximum number of iterations during times in which the mesh moves.

Six DOF Solver Expert Properties

Solver Frozen — Prevents or allows the solver from updating the solution while iterating. Activating this feature should not be your first choice, as it is primarily a "debugging" tool and not supported for normal simulation work.

- Activated — The solver does not update the solution.
- Deactivated — The solver updates the solution.

Sub-Stepping Enabled — Enables or disables the application of a smaller time step for the solution of external forces on a 6-DOF body. This option can be activated to overcome convergence problems due to strong non-linear external forces such as catenary forces. This option should not be used generally, but only to enhance stability for these special cases. Note that with sub-stepping enabled, the fluid mesh and the fluid forces are updated in each inner iteration and not sub-step-wise.
Sub-stepping is presently only available for the Body Motion Option Free Motion (see 6-DOF Body Properties).

- Activated — Sub-stepping is enabled. A Sub-Stepping subnode is created where you specify the Number of Sub-Steps during a time step. If you specify a large number of sub-steps, you may need to increase the number of inner iterations to ensure that the sub-steps converge.
- Deactivated — Sub-stepping is not enabled.

What Is the 6-DOF Motion Solver?

The 6-DOF Motion Solver moves the vertices of the grid according to motion of the body that is calculated by the 6-DOF solver.



6-DOF Motion Solver Properties

None. See Expert Properties.

6-DOF Motion Solver Expert Properties

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. Activating this feature should not be your first choice, as it is primarily a "debugging" tool and not supported for normal simulation work.

Activated — The solver does not update the solution.

Deactivated — The solver updates the solution.

Working with 6-DOF Bodies

The 6-DOF Bodies manager node, which has its own pop-up menu, is the manager object for working with 6-DOF bodies in the simulation tree. It is added to the tree once you select one of the DFBI motions for a region.



This section explains the following:

- What is a 6-DOF body?
- Creating 6-DOF bodies
- What is a Free Motion?
- What is One-DOF Rotating Motion?
- What is One-DOF Translating Motion?
- What are External Forces and Moments?
- 6-DOF Coordinate Systems

What Is a 6-DOF Body?

The 6-DOF body object defines the surface of the floating body for the calculations of the 6-DOF solver. The 6-DOF body is created with a DFBI motion. It has properties and is subject to common menu items.



If a DFBI Rotation and Translation motion is used, add only one 6-DOF body to a simulation. If more that one is added, the solver generates an error message. For the DFBI Embedded Rotation motion and the DFBI Morphing motion, there is no limitation on the number of 6-DOF bodies.

6-DOF bodies can be renamed, copied, and deleted in the same way as other simulation objects. However, when a 6-DOF body is copied, any existing linear spring forces or catenary forces in its External Force and Moment manager are ignored. This is because these forces are managed by their respective coupling elements which exist outside of all 6-DOF body nodes. Other forces and moments are copied.

6-DOF Body Properties

Body Surface	A list of boundaries that represent the surface of the body. Opens the in-place part selector for editing. (Integration of shear force and pressure is over the body surface.)
Body Mass	Mass of the object.
Release Time	Time before calculation of body motion begins. It is best to allow some time for the fluid flow to initialize. The number of seconds depends on the time model you are using and the mesh size. In an unsteady model, the typical range would be 10 to 50 time-steps. If a time-step size is equivalent to 0.01 seconds, 50 time-steps would amount to 0.5 seconds.
Ramp Time	Facilitates a more robust solution by reducing oscillation. At release time, force and moments are applied to the body and can cause a shock effect. Applying a ramp time applies the forces and moments proportionally across the interval to reduce the shock effect. The value of this property is typically 10 times the release time.
Body Motion Option	Selects the type of body motion from a drop-down list. Free Motion is the default motion type.

Free Motion	Selects the Free Motion.
One-DOF Rotating Motion	Selects the One-DOF Rotating Motion.
One-DOF Translating Motion	Selects the One-DOF Translating Motion.
Four-DOF Maneuvering Motion	Selects the Four-DOF Maneuvering Motion.
Planar Motion Carriage	Selects the Planar Motion Carriage.

Pop-Up Menu of the 6-DOF Bodies Manager

New	Adds a 6-DOF body to the 6-DOF Bodies node.
Edit	Executes the customize object action to allow editing of the selected object and its children using a separate modeless stay-on-top dialog.
Paste	Executes the paste action to replace the properties of the current object with properties from the clipboard.
Refresh	Updates the selected object or display to reflect the latest changes to the simulation.

See also: common menu items.

Creating 6-DOF Bodies

To create a 6-DOF body, right-click the DFBI > 6-DOF Bodies manager node and select New.



This action adds a 6-DOF body object to the simulation tree, with its own Free Motion subnode.

What are Initial Values?

You must specify initial values for the 6-DOF body, which are used to define the body at the start of the simulation.

These initial values comprise:

- Angular Velocity
- Center of Mass
- Moment of Inertia
- Orientation
- Velocity
- Attachment Point

You set the initial values of the 6-DOF body in the subnodes of the Initial Values node. Not all of the listed initial values must be set for each motion type. Depending on the selection of the Body Motion Option in the Body 1 node, the subnodes adjust dynamically.

The following screenshot shows the subnodes when a Free Motion Body Motion Option is selected in the Body 1 node.



You specify all of the initial values of the 6-DOF body with respect to an arbitrary coordinate system.

Angular Velocity Properties

Value	Specify the initial angular velocity of the 6-DOF body.
Coordinate System	Selects the coordinate system with respect to which the initial angular velocity is specified.

Center of Mass Properties

Position	Specify the initial position of the center of mass of the 6-DOF body.
Coordinate System	Selects the coordinate system with respect to which the initial center of mass is specified.

Moment of Inertia Properties

Diagonal Components	The diagonal components of the initial Moments of Inertia tensor (Eqn. (2183)).	
Coordinate System	Selects the coordinate system with respect to which the components of the Moments of Inertia tensor are specified. The coordinate system must be Cartesian.	
Off Diagonal Components	The off-diagonal components of the initial Moments of Inertia tensor (Eqn. (2183)).	
Use Center of Mass	Activated	The Moments of Inertia tensor is given with respect to the initial center of mass of the body.
	Deactivated	The Moments of Inertia tensor is given with respect to the origin of the selected Coordinate System.

Orientation Properties

Coordinate System	Select the coordinate system which defines the initial orientation of the 6-DOF body. This coordinate system is used to define the initial body coordinate system. The coordinate system must be Cartesian.

Velocity Properties

Value	Specify the initial velocity of the 6-DOF body.
Coordinate System	Selects the coordinate system with respect to which the initial velocity is specified.

Attachment Point Properties

Value	Specify the point where the 6-DOF body is initially attached to the planar motion carriage.
Coordinate System	Selects the coordinate system with respect to which the attachment point is specified.

What Is a Free Motion?

A free motion object lets you control which direction of motion is calculated, and which ones to leave frozen. The Free Motion node has properties and is subject to common menu items.



These motions are specified with respect to the laboratory coordinate system. Rotations are defined as positive according to the right-hand convention around each coordinate axis. For the purposes of rotation, the laboratory coordinate system can be considered as though its origin were moved to the center of mass of the 6-DOF body.

Free Motion Properties

X Motion	Calculate motion along the X-axis of the laboratory coordinate system.
Y Motion	Calculate motion along the Y-axis of the laboratory coordinate system.
Z Motion	Calculate motion along the Z-axis of the laboratory coordinate system.
X Rotation	Calculate rotation about the X-axis of the laboratory coordinate system (as transposed to the center of mass of the 6-DOF body).
Y Rotation	Calculate rotation about the Y-axis of the laboratory coordinate system (as transposed to the center of mass of the 6-DOF body).
Z Rotation	Calculate rotation about the Z-axis of the laboratory coordinate system (as transposed to the center of mass of the 6-DOF body).

What Is One-DOF Rotating Motion?

A One-DOF Rotating Motion object lets you rotate a body around a specified axis. The One-DOF Rotating Motion node has properties and is

subject to common menu items.



Each one-DOF rotating body can have its own axis of rotation; it is therefore possible to simulate several bodies having different rotation axes (for example, wind turbines in a wind farm).

Optionally, you can use a motion limiter to constrain the angular motion of the body between a minimum and a maximum angle (for example, for valve applications). The body is stopped when a motion limit is reached: the kinetic energy of the body is then destroyed. This physical situation is extreme and can have a significant influence on flow field residuals and convergence. Therefore, if a motion limiter is used, preferably use the default option: a damped motion limiter with which the body stops more smoothly.

You can also specify the angular damping length over which the body is stopped. The angular damping length is effective if the time step is small enough for the body to make several steps over the damping length.

The axis of rotation must be constant in time. Otherwise, the body has more than one degree of freedom, which is beyond the scope of the One-DOF Rotating Motion. Only use steady coordinate systems in the Coordinate System property; unsteady coordinate systems lead to unsteady axes of rotation. Body coordinate systems cannot be used, since these coordinate systems also change with time.

To report the rotation angle of the body, use the Rigid Body 1-DOF Angle report. Use the rotation angle report only with bodies undergoing One-DOF Rotating Motion since it reports zero for all other motions. For these other cases, use the 6-DOF Body Orientation report.

One-DOF Rotating Motion Properties

Rotation Axis Origin	The origin of the rotation axis.
Rotation Axis Direction	The direction of the rotation axis.

One-DOF Rotating Motion Expert Properties

Coordinate System	The coordinate system in which the origin and the direction of the rotation axis are specified.
Motion Limits	Whether to activate the motion limiter (Eqn. (2186) or Eqn. (2187)). Activated — Activates the motion limiter. Deactivated — Deactivates the motion limiter.
Minimum Angle	The minimum angle α_{min} at which the body motion is stopped if the motion limiter is enabled. The angle is measured with respect to the initial orientation of the body. The initial angle is 0. See Eqn. (2186).
Maximum Angle	The maximum angle α_{max} at which the body motion is stopped if the motion limiter is enabled. The angle is measured with respect to the initial orientation of the body. The initial angle is 0. See Eqn. (2187).
Angular Damping Length	The angular length λ (before the minimum and maximum angle) over which the stopping process takes place if the motion limiter is enabled.

What Is One-DOF Translating Motion?

A One-DOF Translating Motion object lets you translate a body along a specified direction. The One-DOF Translating Motion node has properties and

is subject to common menu items.



Each one-DOF translating body can have its own direction of motion; it is therefore possible to simulate several bodies having different translation directions.

Optionally, you can use a motion limiter to constrain the translating motion of the body between a minimum and a maximum displacement (for example, for valve applications). The body is stopped when a motion limit is reached: the kinetic energy of the body is then destroyed. This physical situation is extreme and can have a significant influence on flow field residuals and convergence. Therefore, if a motion limiter is used, preferably use the default option: a damped motion limiter with which the body stops more smoothly.

You can also specify the damping length over which the body is stopped. The damping length is effective if the time step is small enough for the body to make several steps over the damping length.

The direction of motion must be constant in time. Otherwise, the body has more than one degree of freedom, which is beyond the scope of the One-DOF Translating Motion. Only use steady coordinate systems in the Coordinate System property; unsteady coordinate systems lead to unsteady directions of motion. Body coordinate systems cannot be used, since these coordinate systems also change with time.

One-DOF Translating Motion Properties

Translation Origin	The origin of the translational motion. This point is treated as the initial body position, which can be different from the initial center of mass.
Translation Direction	The direction of the translational motion.

One-DOF Translating Motion Expert Properties

Coordinate System	The coordinate system in which the origin and the direction of the translational motion are specified.
Motion Limits	Whether to activate the motion limiter. (Eqn. (2188) and Eqn. (2189))
	Activated — Activates the motion limiter.
	Deactivated — Deactivates the motion limiter.
Minimum Position	The minimum displacement d_{min} at which the body motion is stopped if the motion limiter is enabled. The displacement is measured with respect to the initial position of the body. The initial displacement is 0.
Maximum Position	The maximum displacement d_{max} at which the body motion is stopped if the motion limiter is enabled. The displacement is measured with respect to the initial position of the body. The initial displacement is 0.
Damping Length	The length λ (before the minimum and maximum position) over which the stopping process takes place if the motion limiter is enabled.

What is Four-DOF Maneuvering Motion?

A Four-DOF Maneuvering Motion object lets you specify the speed and direction of a body in the X-Y plane, while leaving the body free to maneuver in the other four degrees of freedom. The Four-DOF Maneuvering Motion node has properties and is subject to common menu items.



Each Four-DOF maneuvering body can have its own speed and direction of motion; it is therefore possible to simulate several bodies having different velocities. For each body, the speed and direction in the X-Y plane is constant in time. The option uses the laboratory coordinate system throughout.

When you run an Orientation report for Four-DOF Maneuvering Motion, use the Rotation X-Y-Z option for the Euler Angle Convention property.

Four-DOF Maneuvering Motion Properties

Course Velocity	The speed of the body in the X-Y plane. The default is zero.
Course Angle	The direction of motion of the body in the X-Y plane, given as the counter-clockwise angle from the X-axis. The default is zero.

What is Planar Motion Carriage?

The Planar Motion Carriage option simulates a captive motion in the X-Y plane of the laboratory coordinate system. This mechanism drives the body along a prescribed trajectory in the X-Y plane, while the body is optionally allowed to move freely in the directions of heave and pitch. There are three Planar Motion methods available:

- Planar Motion Mechanism (PMM)
- Rotating Arm
- General Planar Motion

The Planar Motion Carriage node has properties and is subject to common menu items.



The option uses the laboratory coordinate system throughout. The Initial Orientation Coordinate System of the body must be aligned with the laboratory system.

When you run an Orientation report for the Planar Motion Carriage, use the Rotation X-Y-Z option for the Euler Angle Convention property.

Planar Motion Carriage Properties

Initial Body Position	The initial coordinates of the body position, defined as the location where the body is attached to the Planar Motion Carriage. The carriage translates this position according to the selected Planar Motion Option. Rotation of the body is specified or calculated with respect to this location.	
Planar Motion Option	The type of planar motion that you want your rigid body to perform.	
	Planar Motion Mechanism	Select the Planar Motion Mechanism (PMM). An additional Planar Motion Mechanism sub-node is added to the simulation tree where you can define the properties for this planar motion option. For more information, see What is Planar Motion Mechanism?
	Rotating Arm	Select the Rotating Arm Motion. An additional Rotating Arm sub-node is added to the simulation tree where you can define the properties for this planar motion option. For more information, see What is Rotating Arm Motion?
	General Planar Motion	Select the General Planar Motion. An additional General Planar Motion sub-node is added to the simulation tree where you can define the properties for this planar motion option. For more information, see What is General Planar Motion?
Free Heave	When activated, this option allows the body to move along the Z-axis.	
	Activated	The body is free to move on the Z-axis. This is the default.
	Deactivated	No Z-axis motion allowed.
Free Pitch	When activated, this option allows the body to rotate around its own Y-axis.	
	Activated	The body is free to rotate around its Y-axis. This is the default.
	Deactivated	No Y-axis rotation allowed.

What is Planar Motion Mechanism?

As part of the Planar Motion Carriage, the Planar Motion Mechanism (or PMM) [405] simulates the motion of a rigid body in the X-Y plane of the laboratory coordinate system along a sinusoidal path. The body moves with a constant velocity in X direction and performs an oscillatory motion in Y direction.

The Planar Motion Mechanism node has properties and is subject to common menu items.



Planar Motion Mechanism Properties

Velocity X	The magnitude of the velocity along the X-axis. The default is 1.0 m/s.
Amplitude Y	The maximum displacement of the body along the Y-axis. The default is zero. This corresponds to the sway amplitude y_0 in Eqn. (2228).
Oscillation Frequency	The number of cycles per second. The default is zero. This corresponds to the sway frequency f in Eqn. (2228).
Additional Drift Angle	You can specify an additional drift angle of the body. If specified, this value is added to the yaw angle. The drift angle must be a constant value.
Pure Yaw	With Pure Yaw, the body rotates back and forth in the X-Y plane (yaw) so that the same face (the bow, for a boat) always leads along the sinusoidal path. The Additional Drift Angle is added if specified. The yaw angle is given by Eqn. (2229). Activated: Body moves with pure yaw. Deactivated: Body does not change orientation as it sways. This is the default.

What is Rotating Arm Motion?

As part of the Planar Motion Carriage, the Rotating Arm motion simulates the motion of a rigid body in the X-Y plane of the laboratory coordinate system in a rotating action. The result is a circular path of the body in the X-Y plane.

The Rotating Arm node has properties and is subject to common menu items.



Rotating Arm Properties

Center of Rotation	The coordinates of the point, $\mathbf{r}_c$, around which the rigid body rotates with respect to the Coordinate System (see Eqn. (2231)).
Drift Angle	The drift angle of the body. The drift angle must be a constant value. The drift angle is added to the yaw angle, which is due to the rotating arm motion.
Angular Velocity	The angular velocity, ω, of the rotating arm (see Eqn. (2231)).
Coordinate System	The coordinate system in respect to which the center of rotation is specified.

What is General Planar Motion?

As part of the Planar Motion Carriage, the General Planar Motion simulates the motion of a rigid body in the X-Y plane of the laboratory coordinate system in a user-defined action. You prescribe a trajectory for the motion in the X-Y plane and a yaw angle. The trajectory and the yaw angle are functions of time. The corresponding heave and pitch motions are computed by STAR-CCM+.

The General Planar Motion node has properties and is subject to common menu items.



General Planar Motion Properties

Translation X	The translational motion of the rigid body in the X direction with respect to the laboratory coordinate system. You specify the translation as a function of time. To do this, you enter the function as an Expression (see Setting Values Using an Expression).
Translation Y	The translational motion of the rigid body in the Y direction with respect to the laboratory coordinate system. You specify the translation as a function of time. To do this, you enter the function as an Expression (see Setting Values Using an Expression).
Rotation Z	The rotational motion of the rigid body around the Z axis with respect to the laboratory coordinate system. You specify the rotation angle as a function of time. To do this, you enter the function as an Expression (see Setting Values Using an Expression).
User Defined Derivatives	To calculate the motion of a body defined by the General Planar Motion mechanism, the first and second order time derivatives of the prescribed trajectory are required. By default, STAR-CCM+ computes these derivatives numerically. If the accuracy of this numerical computation is considered insufficient, you can specify the first and second order derivatives explicitly.

Activated	Specifies that the first and second-order time derivatives of the user-defined trajectory are also user-defined. A new Planar Motion Derivatives sub-node is added to the simulation tree.
De-activated	The first and second order time derivatives of the user-defined General Planar Motion trajectory are computed numerically.

Planar Motion Derivatives Properties

Velocity X	Velocity of the rigid body in X direction of the laboratory coordinate system.
Velocity Y	Velocity of the rigid body in the Y direction of the laboratory coordinate system.
Angular Velocity Z	Angular velocity of the rigid body around the Z axis with respect to the laboratory coordinate system.
Acceleration X	Acceleration of the rigid body in X direction of the laboratory coordinate system.
Acceleration Y	Acceleration of the rigid body in Y direction of the laboratory coordinate system.
Angular Acceleration Z	Angular acceleration of the rigid body around the Z axis of the laboratory coordinate system.

What Are External Forces and Moments?

The External Forces and Moments node, which has its own pop-up menu, manages all forces and moments that act on the 6-DOF body. These forces and moments include the force due to gravity, and the pressure and shear forces from the fluid.



The Fluid Force and Moment and Gravity Force nodes are present by default. Additional forces and moments can be added to these using the manager node.

Fluid Force and Moment

The Fluid Force and Moment node, which has properties and is subject to common menu items, is added by default to the External Forces and Moments node. This node represents the forces and moments that act on the 6-DOF body due to the fluid surrounding it.



An option is provided to smooth the fluid forces and moments before they are applied to the body. A moving weighted average is used as a low pass filter; you set the number of time steps that should be included in this moving average.

Properties

Enable Smoothing	This option causes fluid forces and moments to be averaged over a specified time interval before they are applied to the body.	
	Activated	Forces and moments are smoothed.
	Deactivated	Forces and moments are applied directly.
Number of smoothing time steps	Sets the number of time steps that are used in the moving average filter.	
Use Static Pressure	Controls the type of pressure to use.	
	Activated	Switches to the use of static pressure.
	Deactivated	Leaves the choice at the default working pressure.
Enabled	Enables the action of fluid forces and moments on the 6-DOF body.	
	Activated	Fluid forces and moments are applied.
	Deactivated	Fluid forces and moments are ignored.

Gravity Force

The Gravity Force node, which has properties and is subject to common menu items, is added by default to the External Forces and Moments node. This node represents the force due to gravity that acts on the 6-DOF body.



Properties

Enabled	Enables the action of gravity on the 6-DOF body. The gravity model also needs to be selected in the physics continuum—otherwise enabling this option has no effect.	
	Activated	Gravity force acts on the 6-DOF body.
	Deactivated	Gravity force is ignored.

Adding an External Force or Moment

There are several types of external influence that can be added to the 6-DOF body using this manager node:

- Force acting on the Center of Mass
- Moment
- Generalized Force
- Propulsion Force
- Torsional Spring Moment
- Damping Moment
- Damping Force

There are also two forces that are added in response to body couplings created using the Body Couplings manager node:

- Catenary Force

- Linear Spring Force

To add an external force or moment, right-click the External Forces and Moments node and select New followed by the particular type of action to create:



The appropriate sub-node for the external force or moment is added to the manager. These subnodes have their own properties that should be set in the *Properties* window.

Pop-Up Menu of the External Forces and Moments Manager

New	Adds a force or moment of the selected type.	
	Force CM	Adds a Force CM node.
	Moment	Adds a Moment node.
	Force	Adds a Force node.
	Propulsion Force	Adds a Propulsion Force node.
	Damping Moment	Adds a Damping Moment node.
	Damping Force	Adds a Damping Force node.

See also: common menu items.

External Force Acting on the Center of Mass

This node, which has properties and is subject to common menu items, is added when you choose New > Force CM from the External Forces and Moments manager pop-up menu. It represents a user-defined force acting on the 6-DOF body at its center of mass.



The application of this force is illustrated in the diagram below. It does not generate any moment about the center of mass.



Force CM Properties

Force	The external force entered as a vector quantity, along with the appropriate unit of force.
Enabled	Enables the action of the force on the 6-DOF body. Activated — The force is active. Deactivated — The force is not active.

Force CM Expert Properties

Coordinate System	The coordinate system in which the vector of the force is specified, chosen from the drop-down list.

External Moment

This node, which has properties and is subject to common menu items, represents an external moment acting on the 6-DOF body with respect to the body position. It is added when you select New > Moment from the External Forces and Moments manager pop-up menu.



The moments are specified around each axis of the selected coordinate system. However, they act with respect to the position of the body, as illustrated in the following diagram:



Moment Properties

Moment	Components of the external moment entered as a vector quantity, along with the appropriate unit of moment.
Enabled	Enables the action of the moment on the 6-DOF body. Activated — The moment is active. Deactivated — The moment is not active.

Moment Expert Properties

Coordinate System	The coordinate system in which the vector of the moment is specified, chosen from the drop-down list.

External Force

This node, which has properties and is subject to common menu items, represents a general external force acting on the 6-DOF body at a specified location. It is added when you select New > Force from the External Forces and Moments manager pop-up menu.



This force results in a moment acting around the body center of mass, as shown in the following diagram:



External Force Properties

Force	The external force entered as a vector quantity, along with the appropriate unit of force.
Position	The point at which the force acts on the 6-DOF body.
Enabled	Enables the action of the force on the 6-DOF body. Ticked — The force is active. Cleared — The force is not active.

External Force Expert Properties

Coordinate System1	The coordinate system in which the Position coordinates are specified, chosen from the drop-down list.
Coordinate System 2	The coordinate system in which the Force vector is specified, chosen from the drop-down list.

External Propulsion Force

This node, which has properties and is subject to common menu items, represents an external force acting on the body, in the same way as the External Force. The difference is that the external propulsion force is specified using a force magnitude and direction rather than specifying the force components directly. This node is added when you select New > Propulsion Force from the External Forces and Moments manager pop-up menu



Propulsion Force Properties

Position	The point at which the force acts on the 6-DOF body.
Thrust	The magnitude of the force as a constant scalar value with appropriate unit of force.
Direction	Direction in which the force acts. Given as a vector.
Enabled	Enables the action of the force on the 6-DOF body. Activated — The force is active. Deactivated — The force is not active.

Propulsion Force Expert Properties

Coordinate System1	The coordinate system in which the Position coordinates are specified, chosen from the drop-down list.
Coordinate System2	The coordinate system in which the Direction vector is specified, chosen from the drop-down list.

Torsional Spring Moment

The Torsional Spring Moment node, which has properties and is subject to common menu items, represents the effect of a torsional spring acting on a rotating body. The torsional moment always acts to rotate the body into the angular position that is given by the relaxation angle. The magnitude of this moment is proportional to the difference between the current angular position of the body and the relaxation angle (as shown in Eqn. (2238)).

The Torsional Spring Moment node is added when you select New > Torsional Spring Moment from the External Forces and Moments manager pop-up menu.



The torsional spring moment can currently be added to bodies with one rotational degree of freedom only. Such bodies include:

- One-DOF rotating bodies, and
- Bodies undergoing Free Motion with only one rotational degree of freedom selected (and all other degrees of freedom cleared).

The action of the damping moment is illustrated in the following diagram.



The axis of rotation depends on the selected motion type. For a Free Motion, the axis of rotation is the axis of the laboratory coordinate system around which the body is allowed to rotate. For the One-DOF Rotating Motion, you specify the axis of rotation in the node of this motion type.

Torsional Spring Moment Properties

Spring Constant	The spring constant, k in Eqn. (2238), of the torsional spring.
Relaxation Angle	The angle, α_r in Eqn. (2238), at which the spring moment becomes zero. The angle is measured with respect to the initial orientation of the body. The initial angle is 0.
Enabled	Enables the effect of the torsional spring moment on the one-DOF rotating body. Activated — The torsional spring moment is active. Deactivated — The torsional spring moment is not active.

Damping Moment

The Damping Moment node has properties and is subject to common menu items. This node represents the effect of a moment acting around the center of mass of the body in opposition to the angular velocity of the body. The magnitude of the moment, which always acts to decelerate the body, is proportional to the angular velocity (as shown in Eqn. (2239)). This node is added when you select New > Damping Moment from the External Forces and Moments manager pop-up menu.



The action of the damping moment is illustrated in the following diagram.



Damping Moment Properties

Damping Constant	Constant $a_{d,n}$ in Eqn. (2239). See note on use of damping constants.
Enabled	Enables the effect of the damping moment on the 6-DOF body. Activated — The moment is active. Deactivated — The moment is not active.

Damping Force

The Damping Force node has properties and is subject to common menu items. This node represents a force acting on the center of mass of the body in the opposite direction to the body's velocity vector. It therefore always acts to decelerate the body. A linear damping function, that is given by Eqn. (2240), is used to define this force. The force is added when you select New > Damping Force from the External Forces and Moments manager pop-up menu.



The following diagram illustrates the action of the damping force.



Damping Force Properties

Damping Constant	Constant $a_{d,f}$ in Eqn. (2240). See note on use of damping constants.
Enabled	Enables the effect of the damping force on the 6-DOF body. Activated — The force is active. Deactivated — The force is not active.

Use of Damping Constants

Choosing too large a value for the damping constant might lead to numerical instability of the 6-DOF Body motion. If numerical instability is encountered, reduce either the damping constant or the time step.

Catenary Force

This force is automatically created when a catenary coupling is created from the Body Couplings manager node. The name of the coupling is shown in square brackets next to the name of the catenary force. The catenary force

cannot be created, deleted, copied, or renamed within the External Forces and Moments manager node. It is managed entirely by its corresponding catenary coupling.



Catenary Force Properties

Coupling Element	The name of the corresponding catenary coupling element in the Body Couplings manager node (read-only).

Linear Spring Force

The linear spring force node is automatically created when a linear spring coupling is created from the Body Couplings manager node. The name of the associated coupling is shown in square brackets next to the name of the linear spring force node. Unlike earlier versions of STAR-CCM+, the linear

spring force cannot be created, deleted, copied, or renamed within the External Forces and Moments manager. It is managed entirely by its corresponding linear spring coupling.



Linear Spring Force Properties

Coupling Element	The name of the corresponding linear spring coupling element in the Body Couplings manager node (read-only).

6-DOF Coordinate Systems

STAR-CCM+ automatically creates a new Cartesian coordinate system with the default name `Body n-CSys` when a new 6-DOF body with the name `Body n` is created. This coordinate system (which can be referred to as the Body Local Coordinate System) shows the current position and orientation of the 6-DOF body throughout the simulation.

The Initial Orientation Coordinate System is used to define the correct orientation of the 6-DOF body in its initial position.

The origin of the Initial Orientation Coordinate System need not necessarily coincide with the initial position of the 6-DOF body. The Initial Orientation Coordinate System can be set to any appropriate coordinate system defining the initial orientation of the body, including the Laboratory system. Upon initialization, the origin of the Body Local Coordinate System is automatically placed at the initial position of the body.

These coordinate systems are illustrated in the diagram below.



The rigid body motion is decomposed into a translation of a point, and a rotation about this point. The body position is the position of this point. The position of the rigid body depends on the selected body motion type. For example, for a Free Motion, the body position is the center of mass of the body. For a One-DOF Rotating Motion, it is the origin of the rotation axis.

Working with 6-DOF Body Couplings

The Body Couplings manager node is used to create and contain pre-defined types of body-to-body and body-to-environment coupling elements.



Two types of coupling can currently be created:

- Catenary Coupling
- Linear Spring Coupling
- Contact Coupling

All of these couplings are created by right-clicking the Body Couplings manager node and choosing New followed by the desired type.



What Is a Catenary Coupling?

This coupling element, which a catenary coupling node and its properties represent, models a quasi-stationary, elastic catenary between a pair of bodies, or between one body and the environment. The formulation provides the underlying equations.



Examples of where this element could be used are in modeling a tugboat and its payload, or a moored vessel and its associated tether. Quasi-stationary means that the catenary has its steady state shape at each instant of time. This model is appropriate for body motions which are sufficiently slow compared to the wave velocity in the catenary. The catenary coupling between two bodies is depicted in the following diagram:



$\mathbf{f}_1$ and $\mathbf{f}_2$ are tangential to the catenary curve; the horizontal components are equal, but with opposite sign. For a coupling between a body and the environment, the system becomes as follows:



There is a known limitation of the catenary coupling method: the end points of the coupling must not become vertically aligned, as this results in a singularity in the underlying catenary equation. It is therefore necessary to ensure that the end points of the catenary do not coincide in the vertical direction during the simulation. Violation of this requirement results in an error message. In case of vertical alignment failure, the expert property Alignment Failure Offset can be adjusted to prevent such failures.

The catenary coupling facility has two methods for specifying the catenary relaxation length:

- Relaxation Length—The catenary relaxation length is user-defined.
- Pretension—The catenary relaxation length is calculated using a pretension force and estimated relaxation length.

Catenary Coupling Properties

Catenary Mode	Specifies the method to define the catenary mode.	
	Relaxation Length	Selects the Relaxation Length mode, which has its own properties.
	Pretension	Selects the Pretension mode, which has its own properties.
Mass Per Unit Length	The mass per unit length of the catenary without loading.	
Stiffness	The force per unit change of length of the elastic catenary.	
Enabled	Sets whether the coupling is active or not.	
	Activated	The coupling is accounted for in the solution procedure.
	Deactivated	The coupling is not considered.
Object 1	Sets the type of anchor to which the first end of the catenary is attached. This should be different from Object 2.	
	Environment	Attaches the first end point to a location in the environment as defined by Position 1.
	<Body>	Attaches the first end point to the 6-DOF Body selected from this list. Position 1 defines the location to which the catenary is attached.
Object 2	Sets the type of anchor to which the second end of the catenary is attached. This should be different from Object 1.	
	Environment	Attaches the second end point to a location in the environment as defined by Position 2.
	<Body>	Attaches the second end point to the 6-DOF Body selected from this list. Position 2 defines the location on the body to which the catenary is attached.

Position 1	Anchoring position of the first end point of the catenary on Object 1.
Position 2	Anchoring position of the second end point of the catenary on Object 2.

Catenary Coupling Expert Properties

Coordinate System 1	The coordinate system in which Position 1 is specified, chosen from the drop-down list. Only coordinate systems that are related to Object 1 can be chosen.
Coordinate System 2	The coordinate system in which Position 2 is specified, chosen from the drop-down list. Only coordinate systems that are related to Object 2 can be chosen.
Alignment Failure Offset	Parameter that is used to prevent vertical alignment failure of the catenary. Whenever the horizontal distance between the two catenary end points gets smaller than this value, one of the end points is shifted by a small horizontal offset so that the total horizontal distance is equal to the Alignment Failure Offset. It is a relative value that is measured against the relaxation length of the catenary.

What Is the Relaxation Length Catenary Mode?

This mode assumes that you specify the catenary relaxation length explicitly.

Relaxation Length Mode Properties

Relaxation Length	Specifies the catenary length without loading.

What Is the Pretension Catenary Mode?

When using the pretension catenary mode, the catenary relaxation length is calculated so that the pretension force at position 1 equals the user-defined pretension. However, there can be more than one solution to this calculation, or possibly no solution. Therefore, a user-defined, estimated

relaxation length is used to determine which of these multiple solutions is closest to the expected length. If the Estimated Relaxation Length property is set to 0, the shortest catenary length is used.

In some cases, the specified pretension may result in no possible solution. If this occurs, adjust the catenary properties until a result that meets all input parameters is found.

Pretension Catenary Mode Properties

Pretension Force	Specifies the pretension force that is applied at Position 1 of the catenary.
Estimated Relaxation Length	Specifies the estimated catenary relaxation length. This property is used to determine the calculated relaxation length, as there may be more than one solution to the catenary calculation. A value of 0 uses the shortest calculated catenary relaxation length.
Calculated Relaxation Length	The calculated catenary relaxation length (read-only)

What Is a Linear Spring Coupling?

This coupling element models a linear (that is, non-torsional) spring between two bodies or one body and the environment. It carries the main functionality of springs in DFBI. The Linear Spring node, which has properties and is subject to common menu items, represents this linear spring. It is added when you select New > Linear Spring from the Body Couplings manager pop-up menu.



The action of the linear spring coupling between a body and a point in the environment is illustrated in the following diagram; the formula is given in the formulation section.



A linear spring coupling between two bodies is also permitted; this system is depicted below.



Linear Spring Coupling Properties

Relaxation Length

Constant scalar value l_{eq} in Eqn. (2243). No spring force is exerted when the distance between Position 1 and Position 2 is equal to this length.

Spring Constant

Constant scalar value k_s in Eqn. (2241). This property relates the distance by which the spring is stretched or compressed to the resulting force.

Enabled

Sets whether the coupling is active or not.

- Activated — The coupling is accounted for in the solution procedure.
- Deactivated — The coupling is not considered.

Object 1

Sets the type of anchor to which the first end of the spring is attached. This should be different from Object 2.

- Environment — Attaches the first end point to a location in the environment that is defined by Position 1.
- <Body> — Attaches the first end point to the 6-DOF Body selected from this list. Position 1 defines the location to which the spring is attached.

Object 2

Sets the type of anchor to which the second end of the spring is attached. This should be different from Object 1.

- Environment — Attaches the second end point to a location in the environment that is defined by Position 2.
- <Body> — Attaches the second end point to the 6-DOF Body selected from this list. Position 2 defines the location on the body to which the spring is attached.

Position 1	The position vector of the point at which the spring is attached to Object 1, specified in Coordinate System 1.
Position 2	The position vector of the point at which the spring is attached to Object 2, specified in Coordinate System 2.

Linear Spring Coupling Expert Properties

Coordinate System 1	The coordinate system in which Position 1 is specified, chosen from the drop-down list. Only coordinate systems that are related to Object 1 can be chosen.
Coordinate System 2	The coordinate system in which Position 2 is specified, chosen from the drop-down list. Only coordinate systems that are related to Object 2 can be chosen.
No Repelling Force	Determines if repelling forces between Position 1 and Position 2 are considered. If the force would be repelling, and this option is ticked, the force is set to zero.

Activated	Only attracting forces between Position 1 and Position 2 are considered.
Deactivated	Both attracting and repelling forces are considered.

Setting Up Contact Couplings

A 6-DOF body contact coupling represents a contact prevention force between a 6-DOF body and a boundary or between a 6-DOF body and another 6-DOF body. This contact prevention force acts on the 6-DOF body in a direction normal to the boundary and prevents the 6-DOF body from contacting the boundary or another 6-DOF body. This force is only active when the distance between the body and the boundary or another body falls below a user-defined effective range.

To set up a contact coupling:

1. Right-click the DFBI > 6-DOF Bodies > Body Couplings node and select New > Contact.

   The Body Couplings > Contact 1 node and the 6-DOF Bodies > Body 1 > External Forces and Moments > Contact Force [Contact 1] node are added to the simulation tree.

2. Select the Body Couplings > Contact 1 node and set the appropriate properties.

   See Body Couplings Contact Properties



3. Select each of the child nodes in turn and set their properties.
   - Normal Force Properties.
   - Tangential Force Properties.
   - Contact Surface Properties.

To activate the Minimum Contact Distance field function:

- Select the Solver > 6-DOF Solver node and activate Temporary Storage Retained.

For more information, see Contact Coupling Field Functions.

Body Couplings Contact Properties

Object 1

Specifies the 6-DOF body to which the contact prevention force is applied.

Object 2

Specifies the object Object 1 couples to.

- Environment — Applies the contact prevention force to a set of boundaries. The boundaries are specified using the Contact 1 > Contact Surface node.
- <Body> — Applies the contact prevention force to another body.

Enabled

Determines if the contact coupling is applied between Object 1 and Object 2.

- Activated — Contact prevention force is activated.
- Deactivated — Contact prevention force is deactivated.

Effective Range

Specifies the distance at which the contact prevention force starts to act on the 6-DOF body. The contact prevention force is only active when the body falls below the effective range.

Normal Force Properties

Elastic Coefficient

k_1 in Eqn. (2248). To estimate this coefficient, see Eqn. (2249).

Damping Coefficient

k_2 in Eqn. (2248). To estimate this coefficient, see Eqn. (2253).

Tangential Force Properties

Friction Coefficient	μ in Eqn. (2255).
Tanh Coefficient	k_t in Eqn. (2255).

Contact Surface Properties

Inputs	Specifies the boundaries of the environment to which contact is to be prevented. These Inputs must not be set to the boundaries of a rigid body. If you want to prevent contact between two rigid bodies, you set Object 1 and Object 2 under the Body Couplings > [Contact 1] node each to a rigid body (see Body Couplings Contact Properties).

Contact Coupling Field Functions

The following field functions are made available when activating contact coupling:

Minimum Contact Distance of [Contact]	Obtains the minimum distance between each face of the body surface and the nearest face on the opposing boundary, for the last iteration. If sub-stepping is used, the minimum distance is the global minimum of all sub-steps of the last iteration. The minimum distance is projected onto the boundary normal. There is one Minimum Contact Distance available for each contact coupling in the simulation, so that results can be evaluated for each coupling separately. For body-body couplings, the Minimum Contact Distance is only available for the first of the two bodies. To access this field function, select the Solvers > 6-DOF Solver node and activate Temporary Storage Retained. Perform at least one iteration to populate the field. When you close the simulation and then restore it, the Minimum Distance field does not contain any values. You must perform at least one iteration to populate it again.

DFBI Formulation

This section provides the following formulation:

- Governing Equations of Free Motion
- Governing Equation of One-DOF Rotating Motion
- Governing Equation of One-DOF Translating Motion
- Motion Limiter for One-DOF Rotating Motion
- Motion Limiter for One-DOF Translating Motion
- Resultant Force and Moment
- Ramp Time
- Torsional Spring Moment
- Damping Moment
- Damping Force
- Linear Spring Coupling
- Catenary Coupling
- Contact Coupling

Governing Equations of Free Motion

The equation for the translation of the center of mass of the body is noted in the global inertial coordinate system:

$$m\frac{d\mathbf{v}}{dt} = \mathbf{f} \tag{2181}$$

where m represents the mass of the body, $\mathbf{f}$ is the resultant force acting on the body and $\mathbf{v}$ is the velocity of the center of mass. The equation of rotation of the body is formulated in the Body Local Coordinate System with the origin in the center of mass of the body:

$$\mathbf{M}\frac{d\vec{\omega}}{dt} + \vec{\omega} \times \mathbf{M}\vec{\omega} = \mathbf{n} \tag{2182}$$

where $\mathbf{M}$ is the tensor of the moments of inertia, $\vec{\omega}$ is the angular velocity of the rigid body and $\mathbf{n}$ is the resultant moment acting on the body. The tensor of the moments of inertia is expanded as:

$$\mathbf{M} = \begin{bmatrix} M_{xx} & M_{xy} & M_{xz} \\ M_{xy} & M_{yy} & M_{yz} \\ M_{xz} & M_{yz} & M_{zz} \end{bmatrix} \tag{2183}$$

As this tensor is a symmetric, it can be defined by two vectors: one specifying the principal components, (M_{xx}, M_{yy}, M_{zz}) along the diagonal, and another specifying the off-diagonal components, (M_{xy}, M_{xz}, M_{yz}).

Governing Equation of One-DOF Rotating Motion

The equation of motion governing the rotation of a one-DOF rotating body is

$$M\frac{d\omega}{dt} = n \tag{2184}$$

where n is the overall moment acting on the body with respect to the axis of rotation. M is the moment of inertia with respect to the rotation axis. ω is the rotation angular velocity with respect to the rotation axis.

Governing Equation of One-DOF Translating Motion

The equation of motion governing the displacement of a one-DOF translating body is

$$m\frac{dv}{dt} = f \tag{2185}$$

where m represents the mass of the body, f is the resultant force acting on the body and v is the velocity of the center of mass.

Motion Limiter for One-DOF Rotating Motion

The motion limiter destroys the rotational energy of a one-DOF rotating body when the specified maximum angle α_{max} or minimum angle α_{min} is reached. The stopping process starts when the rotation angle α is in the interval $[\alpha_{min}, \alpha_{min} + \lambda]$ or $[\alpha_{max} - \lambda, \alpha_{max}]$, where λ is the damping length.

A damping moment $n_{damping}$ is applied to the body and destroys the rotational energy of the body until the limiting angle is reached.

For the minimum angle,

$$n_{damping} = \frac{1}{2}\frac{M\omega^2}{\alpha - \alpha_{min}} \tag{2186}$$

where M is the moment of inertia with respect to the axis of rotation. ω is the angular velocity of the body.

For the maximum angle,

$$n_{damping} = -\frac{1}{2}\frac{M\omega^2}{\alpha_{max} - \alpha} \tag{2187}$$

To yield a more realistic stopping process, the full damping moment is not applied to the body. A smaller deceleration is applied at the beginning and a higher deceleration at the end.

Motion Limiter for One-DOF Translating Motion

The motion limiter destroys the kinetic energy of a one-DOF translating body when the specified maximum displacement d_{max} or minimum displacement d_{min} is reached. The stopping process starts when the displacement d is in the interval $[d_{min}, d_{min} + \lambda]$ or $[d_{max} - \lambda, d_{max}]$, where λ is the damping length.

A damping force $f_{damping}$ is applied to the body and destroys the kinetic energy of the body until the limiting position is reached.

For the minimum displacement,

$$f_{damping} = \frac{1}{2}\frac{mv^2}{d - d_{min}} \tag{2188}$$

where m is the mass of the body and v is the magnitude of the velocity of the body.

For the maximum displacement,

$$f_{damping} = \frac{1}{2}\frac{mv^2}{d - d_{max}} \tag{2189}$$

To yield a more realistic stopping process, the full damping force is not applied to the body. A smaller deceleration is applied at the beginning and a higher deceleration at the end.

Governing Equations for Constrained Motions

It is sometimes useful to consider kinematic constraints for the simulation of body motion. These constraints are physical constraints to allow certain degrees of freedom and to restrict others.

The analytical expression of the holonomic constraint equations for a rigid body can generally be written as:

$$\Phi_k(\mathbf{d}, t) = 0 \qquad k = 1, 2, \ldots, p, \quad p < 6 \tag{2190}$$

where t is the time, $\mathbf{d}$ is the dependent coordinate vector, and p is the total number of constraints:

$$\mathbf{d} = \begin{bmatrix} r_x \\ r_y \\ r_z \\ \varphi \\ \theta \\ \psi \end{bmatrix}$$

The number of degrees of freedom of the rigid body is therefore reduced to:

$$f = 6 - p \tag{2191}$$

f is the number of degrees of freedom. f' denotes the minimum number of variables (independent variables) which would have described the motion of the rigid body completely. Only $f = 6 - p$ equations are obtained using the independent coordinates. The method used here is based on the Jacobian matrix.

Eqn. (2190) can be re-arranged to define the body position (dependent coordinates) $\mathbf{r}^g$ and the orientation matrix $\mathbf{T}^g$ in the laboratory coordinate system as a function of independent coordinates and time:

$$\mathbf{r}^g = \mathbf{r}(\mathbf{q}, t), \quad \mathbf{T}^g = \mathbf{T}(\mathbf{q}, t) \tag{2192}$$

where $\mathbf{q}$ is the vector of the selected independent coordinates:

$$\mathbf{q} = \begin{bmatrix} q_1 \\ q_2 \\ \ldots \\ q_f \end{bmatrix}$$

Differentiating the above equations, gives:

$$\mathbf{v}^g = \frac{\partial}{\partial \mathbf{q}}\mathbf{r}\dot{\mathbf{q}} + \frac{\partial}{\partial t}\mathbf{r} = \mathbf{J}_T\dot{\mathbf{q}} + \mathbf{v}' \tag{2193}$$

and

$$\boldsymbol{\omega} = \frac{\partial}{\partial \mathbf{q}}\mathbf{t}\dot{\mathbf{q}} + \frac{\partial}{\partial t}\mathbf{t} = \mathbf{J}_R\dot{\mathbf{q}} + \boldsymbol{\omega}' \tag{2194}$$

The $3 \times f$ Jacobian matrix for translation motion is:

$$\mathbf{J}_{\mathrm{T}} = \frac{\partial}{\partial \mathbf{q}}\mathbf{r} = \begin{bmatrix} \frac{\partial r_x}{\partial q_1} & \frac{\partial r_x}{\partial q_2} & \cdots & \frac{\partial r_x}{\partial q_f} \\ \frac{\partial r_y}{\partial q_1} & \frac{\partial r_y}{\partial q_2} & \cdots & \frac{\partial r_y}{\partial q_f} \\ \frac{\partial r_z}{\partial q_1} & \frac{\partial r_z}{\partial q_2} & \cdots & \frac{\partial r_z}{\partial q_f} \end{bmatrix} \tag{2195}$$

and the $3 \times f$ Jacobian matrix for rotation motion is:

$$\mathbf{J}_{\mathrm{R}} = \frac{\partial}{\partial \mathbf{q}}\mathbf{t} = \begin{bmatrix} \frac{\partial t_x}{\partial q_1} & \frac{\partial t_x}{\partial q_2} & \cdots & \frac{\partial t_x}{\partial q_f} \\ \frac{\partial t_y}{\partial q_1} & \frac{\partial t_y}{\partial q_2} & \cdots & \frac{\partial t_y}{\partial q_f} \\ \frac{\partial t_z}{\partial q_1} & \frac{\partial t_z}{\partial q_2} & \cdots & \frac{\partial t_z}{\partial q_f} \end{bmatrix} \tag{2196}$$

Combining Eqn. (2195) with Eqn. (2196) yields:

$$\mathbf{J} = \begin{bmatrix} \mathbf{J}_{\mathrm{T}} \\ \mathbf{J}_{\mathrm{R}} \end{bmatrix} \tag{2197}$$

$\partial\mathbf{t} = \begin{bmatrix} \partial t_x & \partial t_y & \partial t_z \end{bmatrix}^{\mathrm{T}}$ is an instant vector about which the rigid body rotates. The following tensor gives the components of this vector:

$$\partial\mathbf{T} \cdot \mathbf{T}^{\mathrm{T}} = \begin{bmatrix} 0 & -\partial t_z & \partial t_y \\ \partial t_z & 0 & -\partial t_x \\ -\partial t_y & \partial t_x & 0 \end{bmatrix} \tag{2198}$$

The matrix elements of Eqn. (2196) can be calculated as follows:

$$\frac{\partial}{\partial q_1}\mathbf{T} \cdot \mathbf{T}^{\mathrm{T}} = \begin{bmatrix} 0 & -\frac{\partial t_z}{\partial q_1} & \frac{\partial t_y}{\partial q_1} \\ \frac{\partial t_z}{\partial q_1} & 0 & -\frac{\partial t_x}{\partial q_1} \\ -\frac{\partial t_y}{\partial q_1} & \frac{\partial t_x}{\partial q_1} & 0 \end{bmatrix} \tag{2199}$$

Differentiating Eqns. (2193) and (2194) yields the accelerations of the rigid body:

$$\dot{\mathbf{v}}^{g} = \mathbf{J}_{T}\ddot{\mathbf{q}} + \mathbf{L}_{T} + \dot{\mathbf{v}}' \tag{2200}$$

and

$$\dot{\boldsymbol{\omega}}^{g} = \mathbf{J}_{R}\ddot{\mathbf{q}} + \mathbf{L}_{R} + \dot{\boldsymbol{\omega}}' \tag{2201}$$

Here

$$\mathbf{L}_{T} = \frac{\partial}{\partial \mathbf{q}}(\mathbf{J}_{T}\dot{\mathbf{q}}) + 2\frac{\partial}{\partial \mathbf{q}}\mathbf{v}' \tag{2202}$$

and

$$\mathbf{L}_{R} = \frac{\partial}{\partial \mathbf{q}}(\mathbf{J}_{R}\dot{\mathbf{q}}) + 2\frac{\partial}{\partial \mathbf{q}}\mathbf{v}' + \frac{\partial \mathbf{J}_{R}}{\partial t} \tag{2203}$$

with

$$\dot{\mathbf{v}}' = \frac{\partial^{2}}{\partial t^{2}}\mathbf{r}, \qquad \dot{\boldsymbol{\omega}}' = \frac{\partial^{2}}{\partial t^{2}}\mathbf{t} \tag{2204}$$

The governing equation under constraint conditions is given by:

$$\overline{\mathbf{M}}\ddot{\mathbf{q}} + \overline{\mathbf{N}} = \mathbf{F}^{g} \tag{2205}$$

where

$$\overline{\mathbf{M}} = \begin{bmatrix} m & \mathbf{J}_{T} \\ \mathbf{M}^{g} & \mathbf{J}_{R} \end{bmatrix}, \tag{2206}$$

$$\overline{\mathbf{N}} = \begin{bmatrix} m \cdot \mathbf{L}_{T}\dot{\mathbf{q}} + m \cdot \dot{\mathbf{v}}' + m[(\mathbf{J}_{R}\ddot{\mathbf{q}} + \mathbf{L}_{R}\dot{\mathbf{q}} + \dot{\boldsymbol{\omega}}') \times \mathbf{r}_{PC} + (\mathbf{J}_{R}\dot{\mathbf{q}} + \boldsymbol{\omega}') \times ((\mathbf{J}_{R}\dot{\mathbf{q}} + \boldsymbol{\omega}') \times \mathbf{r}_{PC})] \\ \mathbf{M}^{g} \cdot \mathbf{L}_{R}\dot{\mathbf{q}} \cdot \mathbf{M}^{g} \cdot \dot{\boldsymbol{\omega}}' + (\mathbf{J}_{R}\dot{\mathbf{q}} + \boldsymbol{\omega}') \times \mathbf{M}^{g} \cdot (\mathbf{J}_{R}\dot{\mathbf{q}} + \boldsymbol{\omega}') \end{bmatrix},$$

where m is the mass of the body, $\mathbf{M}^{g}$ is the tensor of the moments of inertia in the laboratory coordinate system, and $\mathbf{r}_{PC}$ is a position vector to describe the distance between the center of rotation and the center of mass. Only for the planar motion carriage mechanism, $\mathbf{r}_{PC} \neq 0$. For all other motions, $\mathbf{r}_{PC} = 0$.

The vector of forces and moments $\mathbf{F}^{g}$ is defined as:

$$\mathbf{F}^{g} = \begin{bmatrix} \mathbf{f}^{g} \\ \mathbf{n}^{g} \end{bmatrix} \tag{2207}$$

$\mathbf{F}^{g}$ can be split into two parts:

$$\mathbf{F}^{g} = \mathbf{F}^{i,g} + \mathbf{F}^{q,g} \tag{2208}$$

$\mathbf{F}^{i,g}$ stands for the impressed forces (or moments) acting on the body. These forces can either be obtained from the fluid flow or any external sources. $\mathbf{F}^{q,g}$ stands for the forces (or moments) which maintain the given constraints.

The forces (or moments) $\mathbf{F}^{q,g}$ are still unknown. $\mathbf{F}^{q,g}$ can be eliminated by substituting Eqn. (2208) into Eqn. (2205), multiplying the matrix $\mathbf{J}^{T}$ to both sides of Eqn. (2205), and by using:

$$\mathbf{J}^{T}\mathbf{F}^{q} = 0 \tag{2209}$$

As a result, you get:

$$(\mathbf{J}^{T}\overline{\mathbf{M}})\ddot{\mathbf{q}} + \mathbf{J}^{T}\overline{\mathbf{N}} = \mathbf{J}^{T}\mathbf{F}^{i,g} \tag{2210}$$

This equation has $f = 6 - p$ variables.

Governing Equations for Four-DOF Maneuvering Motion

Consider a solid body, for example a ship, moving with constant velocity V and a constant course angle ψ. Four degrees of freedom are free, for the simulation of a ship this can be sway, heave, roll, and pitch motion. They can be selected as independent variables. The independent coordinate can be written as:

$$\mathbf{q} = \begin{bmatrix} y \\ z \\ \varphi \\ \theta \end{bmatrix}, \quad f = 4 \tag{2211}$$

The solid body position $\mathbf{r}$ and the orientation $\mathbf{T}$ can be expressed by:

$$\mathbf{r} = \begin{bmatrix} r_x \\ r_y \\ r_z \end{bmatrix} = \begin{bmatrix} -y \cdot \sin\psi + V \cdot t \cdot \cos\psi \\ y \cdot \cos\psi + V \cdot t \cdot \sin\psi \\ z \end{bmatrix}, \tag{2212}$$

$$\mathbf{T} = \begin{bmatrix} \cos\theta\cos\psi & \sin\varphi\sin\theta\cos\psi & \cos\varphi\sin\theta\cos\psi + \sin\varphi\sin\psi \\ \cos\theta\sin\psi & \sin\varphi\sin\theta\cos\psi + \cos\varphi\cos\psi & \cos\varphi\sin\theta\sin\psi - \sin\varphi\cos\psi \\ -\sin\theta & \sin\varphi\cos\theta & \cos\varphi\cos\theta \end{bmatrix} \tag{2213}$$

According to Eqns. (2195) and (2196), the Jacobian matrices for translation and rotation motions are for Four-DOF motions:

$$\mathbf{J}_T = \begin{bmatrix} -\sin\psi & 0 & 0 & 0 \\ \cos\psi & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \tag{2214}$$

$$\mathbf{J}_R = \begin{bmatrix} 0 & 0 & \cos\psi\cos\theta & -\sin\psi \\ 0 & 0 & \sin\psi\cos\theta & \cos\psi \\ 0 & 0 & -\sin\theta & 0 \end{bmatrix} \tag{2215}$$

According to Eqns. (2202), (2203), (2200), (2201), (2204), the other matrices and vectors are for Four-DOF motions:

$$\mathbf{L}_T = 0, \tag{2216}$$

$$\mathbf{L}_R = \begin{bmatrix} 0 & 0 & 0 & -\dot{\varphi}\cos\psi\sin\theta \\ 0 & 0 & 0 & -\dot{\varphi}\sin\psi\sin\theta \\ 0 & 0 & 0 & -\dot{\varphi}\cos\theta \end{bmatrix}, \tag{2217}$$

$$\mathbf{v}' = \begin{bmatrix} V\cdot\cos\psi \\ V\cdot\sin\psi \\ 0 \end{bmatrix}, \quad \boldsymbol{\omega}' = 0 \tag{2218}$$

$$\dot{\mathbf{v}}' = 0, \quad \boldsymbol{\omega}' = 0 \tag{2219}$$

Governing Equations for Planar Motion Carriage

The planar motion carriage has only two free degrees of freedom. The other four are restricted by the prescribed motion of the planar motion carriage. The independent coordinate is given by:

$$\mathbf{q} = \begin{bmatrix} z \\ \theta \end{bmatrix}, \quad f = 2 \tag{2220}$$

The solid body position $\mathbf{r}$ can be written as:

$$\mathbf{r} = \begin{bmatrix} r_x \\ r_y \\ r_z \end{bmatrix} = \begin{bmatrix} r_x(t) \\ r_y(t) \\ z(t) \end{bmatrix} \tag{2221}$$

Eqn. (2221) describes the position vector in a general form. The orientation is described by Eqn. (2213) with the roll angle φ fixed at zero and the yaw angle $\psi = \psi(t)$. Depending on the planar motion carriage method that has been selected, the position vector and the orientation matrix can vary. The first planar motion carriage method that has been implemented is the Planar Motion Mechanism (PMM).

For more details on the formulation of this method see Governing Equations for Planar Motion Mechanism (PMM).

For the planar motion carriage, the Jacobian matrices can be written as:

$$\mathbf{J}_\mathrm{T} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \end{bmatrix} \tag{2222}$$

$$\mathbf{J}_\mathrm{R} = \begin{bmatrix} 0 & -\sin\psi \\ 0 & \cos\psi \\ 0 & 0 \end{bmatrix} \tag{2223}$$

According to Eqns. (2202), (2203), (2200), (2201), (2204), the other matrices and vectors can be expressed as:

$$\mathbf{L}_\mathrm{T} = 0, \tag{2224}$$

$$\mathbf{L}_\mathrm{R} = \begin{bmatrix} 0 & -\cos\psi\dfrac{\partial\psi}{\partial t} \\ 0 & -\sin\psi\dfrac{\partial\psi}{\partial t} \\ 0 & 0 \end{bmatrix}, \tag{2225}$$

$$\mathbf{v}' = \begin{bmatrix} \dfrac{\partial r_x}{\partial t} \\ \dfrac{\partial r_y}{\partial t} \\ 0 \end{bmatrix}, \qquad \boldsymbol{\omega}' = \begin{bmatrix} 0 \\ 0 \\ \dfrac{\partial\psi}{\partial t} \end{bmatrix}, \tag{2226}$$

$$\dot{\mathbf{v}}' = \begin{bmatrix} \frac{\partial^2 r_x}{\partial t^2} \\ \frac{\partial^2 r_y}{\partial t^2} \\ 0 \end{bmatrix}, \qquad \dot{\boldsymbol{\omega}}' = \begin{bmatrix} 0 \\ 0 \\ \frac{\partial^2 \psi}{\partial t^2} \end{bmatrix} \tag{2227}$$

Governing Equations for Planar Motion Mechanism (PMM)

The Planar Motion Mechanism is a method of the Planar Motion Carriage that prescribes the motion in the X-Y plane in the form of a sinusoidal sway motion.

The solid body position $\mathbf{r}$ can be written as:

$$\mathbf{r} = \begin{bmatrix} r_x \\ r_y \\ r_z \end{bmatrix} = \begin{bmatrix} V_0 \cdot t \\ y_0 \cdot \sin(2\pi f \cdot t) \\ z \end{bmatrix} \tag{2228}$$

where V_0 is the forward velocity of the body, y_0 is the sway amplitude, and f is the sway frequency.

The orientation $\mathbf{T}$ corresponds to Eqn. (2213) with the roll angle $\varphi = 0$.

If pure yawing is activated, the yawing angle is given by:

$$\psi = \psi_0 \cdot \cos(2\pi f \cdot t) \tag{2229}$$

with:

$$\psi_0 = \operatorname{atan}\left(\frac{y_0 \cdot 2\pi f}{V_0}\right), \tag{2230}$$

else $\psi = 0$.

Governing Equations for Rotating Arm Motion

The Rotating Arm Motion mechanism is a method of the Planar Motion Carriage that prescribes the motion in the X-Y plane in the form of a rotation. The result is a circular path of the body in the X-Y plane.

The solid body position $\mathbf{r}$ can be written as:

$$\mathbf{r} = \begin{bmatrix} r_x \\ r_y \\ r_z \end{bmatrix} = \begin{bmatrix} r_{c,x} + (r_{0,x} - r_{c,x})\cos\omega t - (r_{0,y} - r_{c,y})\sin\omega t \\ r_{c,y} + (r_{0,x} - r_{c,x})\sin\omega t - (r_{0,y} + r_{c,y})\cos\omega t \\ z \end{bmatrix} \tag{2231}$$

where:

$$\mathbf{r}_c = \begin{bmatrix} r_{c,x} \\ r_{c,y} \\ r_{c,z} \end{bmatrix} \text{ : Center of rotation of the rotating arm}$$

$$\mathbf{r}_0 = \begin{bmatrix} r_{0,x} \\ r_{0,y} \\ r_{0,z} \end{bmatrix} \text{ : Initial body position}$$

The initial body position corresponds to the location where the body is initially attached to the rotating arm.

The yaw angle is given by:

$$\psi = \omega \cdot t \tag{2232}$$

Governing Equations for General Planar Motion

The General Planar Motion mechanism is a method of the Planar Motion Carriage that prescribes the motion in the X-Y plane in the form of a user-defined trajectory.

The solid body position $\mathbf{r}$ can be written as:

$$\mathbf{r} = \begin{bmatrix} r_x \\ r_y \\ r_z \end{bmatrix} = \begin{bmatrix} x(t) \\ y(t) \\ z \end{bmatrix} \tag{2233}$$

The yaw angle is given by:

$$\psi = \psi(t) \tag{2234}$$

where x, y, ψ are user-defined functions of time.

Resultant Force and Moment

The resultant force and moment acting on the rigid body are obtained from the fluid pressure and shear forces acting on each boundary face of the body.

Pressure forces: $$\mathbf{f}_p = \sum_f p_f \mathbf{a}_f$$

Pressure moments: $$\mathbf{n}_p = \sum_f [\mathbf{r}_f \times (p_f \mathbf{a}_f)]$$

Shear forces: $$\mathbf{f}_\tau = -\sum_f \boldsymbol{\tau}_f a_f$$

Shear moments: $$\mathbf{n}_\tau = -\sum_f [\mathbf{r}_f \times (\boldsymbol{\tau}_f a_f)]$$

Gravity force: $$\mathbf{f}_g = m\mathbf{g}$$

In the above equations, p_f is the pressure acting on face *f*, $\mathbf{a}_f$ is the area vector of face *f*, $\mathbf{r}_f$ is the distance vector from the body center of mass to the center of face *f*, $\boldsymbol{\tau}_f$ is the shear stress acting on face *f*, and $\mathbf{g}$ is the gravity vector.

The overall resultant force and moment acting on the body are therefore:

$$\mathbf{f} = f_r\left(\mathbf{f}_p + \mathbf{f}_\tau + \mathbf{f}_g + \sum \mathbf{f}_{ext}\right) \tag{2235}$$

$$\mathbf{n} = f_r\left(\mathbf{n}_p + \mathbf{n}_\tau + \sum \mathbf{n}_{ext}\right) \tag{2236}$$

where f_r is the ramping factor that is defined in Eqn. (2237) below. $\mathbf{f}_{ext}$ represents an external user-defined force acting on the body, and $\mathbf{n}_{ext}$ represents an external moment (which you can define directly or as the result of an external user-defined force).

Ramp Time

For 6-DOF bodies, with the values of release time t_s and ramp time t_r a ramping function is defined as follows:

$$f_r = \begin{cases} 0 & : \quad t < t_s \\ \dfrac{t - t_s}{t_r} & : \quad t_s \le t < t_s + t_r \\ 1 & : \quad t \ge t_s + t_r \end{cases} \tag{2237}$$

This ramping function $f_r(t)$ is applied to external forces and moments on the body, forces, and moments that are induced by the fluid flow on the body and to the gravity force.

Torsional Spring Moment

The torsional spring moment M acting on a rotating body is defined as:

$$M = -k(\alpha - \alpha_r) \tag{2238}$$

where k is the spring constant, α is the rotation angle of the body and α_r is the relaxation angle at which the moment vanishes.

Damping Moment

The damping moment acts with respect to the center of mass of the body in opposition to the angular velocity vector. It defined as:

$$\mathbf{n}_d = -a_{d,n}\,\vec{\omega} \tag{2239}$$

where $a_{d,n}$ is the angular damping constant.

Damping Force

The damping force acts with respect to the center of mass of the body in opposition to the velocity vector. It is defined as:

$$\mathbf{f}_d = -a_{d,f}\,\mathbf{v} \tag{2240}$$

where $a_{d,f}$ is the linear damping constant.

Linear Spring Coupling

The linear spring coupling provides a restoring force, which means that it acts in the direction that tends to restore the spring to a state of equilibrium. The forces acting at the two end points of the spring are given according to Hooke's law, and are defined as:

$$\begin{aligned} \mathbf{f}_1 &= -k_s(\mathbf{r}_s - \mathbf{r}_{eq}) \\ \mathbf{f}_2 &= -\mathbf{f}_1 \end{aligned} \tag{2241}$$

where

$$\mathbf{r}_s = \mathbf{x}_1 - \mathbf{x}_2 \tag{2242}$$

and

$$\mathbf{r}_{eq} = l_{eq}\,\mathbf{r}_s / |\mathbf{r}_s| \tag{2243}$$

In these equations, $\mathbf{x}_1$ is the location to which one end of the spring is attached, $\mathbf{x}_2$ is the opposite end of the spring, $\mathbf{f}_1$ is the force acting at $\mathbf{x}_1$, $\mathbf{f}_2$ is the reaction force acting at $\mathbf{x}_2$, and l_{eq} is the relaxation or equilibrium length of the spring. When $l_{eq} \equiv |\mathbf{r}_s|$, the spring force is zero. Some of these parameters are illustrated in the following diagram:



In this diagram, one end of the spring, $\mathbf{x}_1$, is attached to a 6-DOF body and the other end, $\mathbf{x}_2$, is attached to the environment. This is only one of several possible configurations. It is also possible to attach $\mathbf{x}_2$ to a 6-DOF body and $\mathbf{x}_1$ to the environment, or to attach $\mathbf{x}_1$ to one 6-DOF body and $\mathbf{x}_2$ to another.

Catenary Coupling

The catenary coupling models an elastic, quasi-stationary catenary (such as a chain or towing rope), which is hanging between two end points, being subject to its own weight in the gravity field. In a local cartesian coordinate system, the shape of the catenary is given by the following set of parametric equations:

$$
\begin{aligned}
x &= au + b\sinh(u) + \alpha \\
y &= a\cosh(u) + \frac{b}{2}\sinh^2(u) + \beta \\
&\text{for } u_1 \le u \le u_2
\end{aligned}
\tag{2244}
$$

where

$$a = \frac{c}{\lambda_0 g}$$

$$b = \frac{ca}{DL_{eq}}$$

$$c = \frac{\lambda_0 L_{eq} g}{\sinh(u_2) - \sinh(u_1)}$$

In these equations, g is the gravitational acceleration, λ_0 and L_{eq} are the mass per unit length and the relaxation length of the catenary, respectively, under force-free conditions. D is the stiffness of the catenary, and α and β are integration constants depending on the position of the two end points and the total mass of the catenary. The curve parameter u is related to the inclination angle ϕ of the catenary curve by the following equation:

$$\tan\phi = \sinh(u) \tag{2245}$$

The parameter values u_1 and u_2 represent the positions of the catenary's end points p_1 and p_2 in parameter space.

The forces $\mathbf{f}_1$ and $\mathbf{f}_2$ acting at the two end points of the catenary are directed along the local tangent vectors of the catenary curve at the parameter values u_1 and u_2, respectively. They are given by the following expressions:

$$
\begin{aligned}
\mathbf{f}_{1,x} &= c \\
\mathbf{f}_{1,y} &= c\sinh(u_1)
\end{aligned}
\tag{2246}
$$

and

$$
\begin{aligned}
\mathbf{f}_{2,x} &= -c \\
\mathbf{f}_{2,y} &= -c\sinh(u_2)
\end{aligned}
\tag{2247}
$$

These forces, and the resulting catenary curve, are illustrated in the diagram below.



Contact Coupling

The DFBI contact coupling prevents contact of a DFBI body with a boundary. This boundary can either be a boundary of the surrounding computational domain (Body-Environment coupling) or the surface of another DFBI body (Body-Body coupling). To achieve contact prevention, the model calculates an external force normal to the boundary that acts on the DFBI body.

The contact force depends on the distance between the DFBI body surface and the opposing boundary. The contact coupling is a near-field coupling, that is, the force decreases to zero when the distance between the body and the boundary is larger than a user-specified effective range. If the distance is smaller than the effective range, a repulsive contact force between the DFBI body and the boundary is created. The contact force is distributed per face over the body surface and has two force components:

- Normal component, which is normal to the opposing boundary. This force component prevents contact.
- Tangential component, which is tangential to the opposing boundary. This component models the friction of the DFBI body sliding along the boundary.

The normal component of the contact force on face f of the body surface is given by:

$$\mathbf{f}_n(d_f) = a_f[k_1(d_0 - d_f) - k_2\dot{d}_f]\mathbf{n}_{bf} \tag{2248}$$

where

k_1	Elastic coefficient
k_2	Damping coefficient
a_f	Face area

d_f	Distance between face centroid and boundary
$\dot{d}_f$	Rate of change of the distance between face centroid and boundary
d_0	Effective range
$\mathbf{n}_{bf}$	Normal of the boundary face next to face f

The contact force has an elastic term and a damping term that are specified by the elastic coefficient k_1 and the damping coefficient k_2, respectively.

The elastic coefficient k_1 is essential to stop the DFBI body well before contact is made. To estimate a good value for this coefficient, you can apply the following relations:

- Body-Body couplings:

$$k_1 = \frac{m_1 m_2}{m_1 + m_2} \cdot \frac{1}{A} \cdot \frac{v_{n,\,rel}^2}{(d_0 - d_{min})^2} \tag{2249}$$

- Body-Environment couplings:

$$k_1 = \frac{m_1}{A} \cdot \frac{v_{n,\,rel}^2}{(d_0 - d_{min})^2} \tag{2250}$$

where

m_1, m_2	Body mass of Body 1 and Body 2
A	Estimated area of contact
$v_{n,\,rel}$	Normal relative impact velocity of the two coupled objects
d_{min}	Minimum distance. Body should stop before this distance is reached.

You can estimate the time that is needed to stop the DFBI body by applying the contact force as follows:

- Body-Body couplings:

$$\tau = \frac{\pi}{2}\sqrt{\frac{m_1 \cdot m_2}{(m_1 + m_2)Ak_1}} \tag{2251}$$

- Body-Environment couplings:

$$\tau = \frac{\pi}{2}\sqrt{\frac{m_1}{Ak_1}} \tag{2252}$$

To properly resolve the process of stopping, you must set the time step of the simulation sufficiently smaller than τ. If the simulation time step is too large to resolve the process, you can enable 6-DOF Solver sub-stepping. Sub-stepping allows STAR-CCM+ to resolve the stopping process in DFBI. The above relations are only valid for weak damping.

To estimate a value for the damping coefficient, you can use the following relation:

- Body-Body coupling:

$$k_2 = 2\varsigma\sqrt{\frac{(m_1 \cdot m_2)k_1}{(m_1 + m_2)A}} \tag{2253}$$

- Body-Environment coupling:

$$k_2 = 2\varsigma\sqrt{\frac{m_1 k_1}{A}} \tag{2254}$$

ς is a constant factor describing the amount of damping. For $\varsigma = 1$, damping is critical. Note that very large values of ς in combination with a large time step can cause numerical instability of the damping force. To avoid numerical instability, use a sufficiently small value of ς. When in doubt, use $\varsigma \ll 1$.

The tangential component of the contact force on face f of the body surface is given by:

$$\mathbf{f}_t = -\mu\left|\mathbf{f}_n\right|\tanh(k_t\mathbf{v}_t) \tag{2255}$$

where

μ	Friction coefficient
k_t	tanh coefficient
$\mathbf{v}_t$	Relative velocity component of face f tangential to the boundary

This tangential force component models the kinetic friction between the body and the boundary. Static friction is not accounted for. The tanh function makes the friction force stable for small relative sliding velocities.

For sliding velocities $\mathbf{v}_t \gg 1/k_t$, the friction force depends only on the friction coefficient, the normal force, and the direction of the relative velocity. Depending on k_t, the friction force remains constant with increasing sliding velocity. For smaller sliding velocities, the friction force additionally depends on the velocity magnitude.



The total contact force on face f is given by:

$$\mathbf{f}_f = \mathbf{f}_n + \mathbf{f}_t \tag{2256}$$

The total force creates a moment $\mathbf{m}_f$ around the current body position $\mathbf{r}_b$:

$$\mathbf{m}_f = (\mathbf{r}_f - \mathbf{r}_b) \times \mathbf{f}_f \tag{2257}$$

where $\mathbf{r}_f$ is the location of the face centroid of face f. The total force and the total moment on the body are each the sum of all contributions of all faces of the body surface.

DFBI Bibliography

[405] Ohmori, T. 1998. "Finite-volume simulation of flows around a ship in maneuving motion", *J. Mar. Sci. Tchnol.*, 3, pp. 82-93.

Modeling with Harmonic Balance

Some unsteady flows have a regularly repeating flow pattern, that is, they are time-periodic. Consider the flow from a fan blade passing across the entrance to a duct. Measurements of the instantaneous flow just within the duct would show a regularly repeating pattern. If the flow disturbances are sufficiently large, and propagate to the end of the duct, measurements of the unsteady flow at any point within the duct show repeating patterns. Such time-periodic patterns can be expressed using Fourier series.

Fourier series are series of cosine and sine terms that when taken together, reproduce a periodic function.



This periodic function, which is also known as a trigonometric series, can be written in simple form as:

$$a_0 + \sum_{n=1}^{\infty} (a_n \cos((nx) + b_n \sin(nx)))$$

The STAR-CCM+ harmonic balance method is a mixed frequency/time domain technique for computational fluid dynamics. It enables the efficient calculation of time periodic flows. The unsteady solution is stored at a small number of fixed *time levels* (solution snapshots) over one time period of the unsteady flow in a single blade passage in each blade row; thus the solution

is periodic by construction. The individual time levels are coupled to one another through a spectral operator representing the time derivative term in the Navier-Stokes equation, and at the boundaries of the computational domain through the application of periodic and nonreflecting boundary conditions.

Applications

The harmonic balance method allows for the rapid computation of unsteady flow solutions that are periodic in time. As a result, this method is useful for a wide variety of applications, including:

- Turbomachinery
- Tonal aeroacoustics
- Flutter and forced response
- Rotorcraft

In addition to quickly computing the periodic unsteady solution, the harmonic balance method can also account for inter-blade phase angle effects. This additional benefit means the harmonic balance method is well-suited for turbomachinery simulation. With the harmonic balance method, it is possible to represent how the solution varies from one blade to the next within a given row of a turbomachine while only meshing a single blade in that row. The need to mesh only one blade in each row dramatically reduces computational cost.

Benefits

The main benefit of the Harmonic Balance method is that you can solve unsteady flow problems using steady flow solver techniques. This means that the method is often faster than conventional time-accurate time-marching approaches.

Limitations

The method is designed to work with temporally periodic flows, or even some “nearly” periodic flows. It is not suited for transient problems that are non-periodic in time.

Currently, K-Epsilon, K-Omega, and Spalart-Allmaras turbulence models are available for use with Harmonic Balance. However, the setup and use of these turbulence models are identical to how the models would be used with conventional unsteady calculations.

Working with the Harmonic Balance Method in STAR-CCM+

The following sections provide information on the Harmonic Balance method in STAR-CCM+ as follows:

- Setting up a harmonic balance simulation
- What is the harmonic balance model?
- What is the harmonic balance flow and energy model?
  - Working with blade rows
  - Working with mode tables
- What are the harmonic balance turbulence models?
- What is a wake?
- What is the harmonic balance flutter motion?
- What is the implicit harmonic balance solver?
- Post-processing harmonic balance field functions
- Filtered mode field functions
- Harmonic balance formulation
  - Harmonic balance bibliography
- Harmonic balance field functions

Setting Up a Harmonic Balance Simulation

The following steps outline the actions that are required to set up a simulation with the Harmonic Balance model:

- Activating the harmonic balance model
- Activating the harmonic balance flow and energy model
  - Working with blade rows
  - Working with mode tables
- Activating the harmonic balance Spalart-Allmaras models
- Defining harmonic balance boundary conditions
- Setting up a simulation with harmonic balance flutter motion
- Post-processing the harmonic balance solution data

What Is the Harmonic Balance Model?

The Harmonic Balance model is represented in the object tree by the Harmonic Balance node, and is activated in the physics models selection dialog.



Activating the Harmonic Balance Model

To activate the Harmonic Balance model:

- In the *Physics Model Selection* dialog, select the Harmonic Balance model in the *Time* group.



If the Auto-select recommended models checkbox is ticked, the Harmonic Balance Flow and Energy model is selected automatically.

When the Harmonic Balance model is selected, the standard Coupled Flow or Segregated Flow models cannot be selected.



The Implicit Harmonic Balance solver is a coupled solver that couples the conservation of mass, momentum, and energy at all the time levels. Time levels are points in time, chosen by the solver, when a solution is computed.

The Harmonic Balance model works only with the following models within their respective categories:

- the Harmonic Balance *Time* model,
- the Gas *Material* model,
- the Ideal Gas *Equation of State* model, and
- the RANS K-Epsilon, K-Omega, and Spalart-Allmaras *Reynolds-Averaged Turbulence* models.

From the other model groups, the Harmonic Balance method works with:

- two-dimensional or three-dimensional models,
- Stationary motion,
- Reference Frame objects,
- Inviscid, Laminar, or Turbulent viscous regimes, and

- the Passive Scalar model; the passive scalars are treated as steady and use the time-mean flow field from the unsteady Harmonic Balance simulation.

You can also simulate blade vibration by selecting the Harmonic Balance Flutter motion instead of Stationary.

What Is the Harmonic Balance Flow and Energy Model?

The Harmonic Balance Flow and Energy model operates on similar principles as the Coupled Flow and Coupled Energy models, and uses the Implicit Harmonic Balance solver. The Harmonic Balance Flow and Energy model is controlled by properties similar to the Coupled Flow model properties, including discretization schemes and secondary gradients.

The Harmonic Balance Flow and Energy node represents this model in the object tree, and is subject to common menu items. This node includes the objects for working with blade rows, and working with mode tables.



Once you have activated this model, you can create blade rows. Using multiple blade rows in a simulation requires the use of user-specified or default filtered mode tables. Once the blade row has been created, it can be renamed at any time.

A blade row represents a region or a set of regions that models one or more blade passages, such as that of a stator or a rotor of a turbomachine, with a fixed rotation rate and number of blades.

After creating the blade rows, and initializing the solution, the model computes a mode table. This table tabulates all the unsteady frequencies and nodal diameters to be considered in the simulation for every unsteady mode.

Harmonic Balance Flow and Energy Properties

Discretization	Selects the discretization scheme that is used.	
	1st-order	Selects the first-order upwind discretization scheme
	2nd-order	Selects the second-order upwind discretization scheme.

Expert Properties

Unless you are thoroughly familiar with the theoretical aspects of this model, and the discretization techniques that are used in STAR-CCM+, we recommend that you do not change any properties within the *Expert* category. The values in that category reflect both the model's design and discretization approaches that have been optimized for accuracy and performance.

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Positivity Rate Limit	Maximum rate at which temperature is allowed to decrease.	
Minimum Absolute Pressure	Minimum allowable absolute pressure.	
Explicit Relaxation	A scaling factor that is used to relax all corrections explicitly before they are applied to the variables.	
Preconditioning Enabled	Enable/disable preconditioning.	

	Activated	Enable preconditioning.
	Deactivated	Disable preconditioning.
Flow Boundary Diffusion	Enable/disable the flow-boundary diffusion fluxes (or viscous fluxes for flow models)—see Eqn. (2558). This property is enabled by default.	
	Activated	Enable flow-boundary diffusion.
	Deactivated	Disable flow-boundary diffusion.

Activating the Harmonic Balance Flow and Energy Model

The Harmonic Balance Flow and Energy model is selected automatically when activating the harmonic balance (time) model.

Working with Blade Rows

You can model one or more blade rows in a Harmonic Balance simulation. First, you define blade rows in the Blade Rows manager that is located within the Harmonic Balance Flow and Energy node. After this step, you

associate each blade row with its corresponding region in the simulation. The Blade Row Specification property on a region is used to make this connection.



This section describes the following:

- What is the Blade Rows Manager?
- What is a blade row?
- Setting blades and pitches
- How to create a blade row
- Matching blade rows to regions
- What are the available options for computing blade row frequencies?

What Is a Blade Row?

A blade row represents a region or a set of regions that models a fixed number of blades within a specified pitch.

The Blade Rows Manager

The Blade Rows manager node contains all the individual blade row nodes, the Filtered Modes Table node, and the User Specified Modes Table node. This node is controlled with properties, and supports node actions.

Blade Rows Manager Pop-Up Menu

Use this pop-up menu to add blade row objects to the simulation.

New	Adds a blade row to the simulation.

Blade Rows Manager Expert Properties

Frequency Compute Option	Specifies whether blade row frequencies are computed using disturbance from neighboring blade rows, or from all blade rows. See Blade Row Frequencies Computation.	
	Consider Only Neighbors	Consider disturbances from neighboring rows only.
	Consider All Blade Rows	Consider disturbances from all blade rows. The relative rotation rate between any two blade rows must be non-zero. This is insured by the Fictitious Rotation Rate Factor.
	See also User Specified Mode Tables.	
Fictitious Rotation Rate Factor	Specifies the proportionality constant for the offset applied to blade frequencies. The default value is 0.001. See Choosing a Fictitious Rotation Rate Factor and Fictitious Rotation Rate Computation.	

Blade Row Nodes

A Blade Row node represents a separate blade row in the simulation. This node is controlled with properties, and is subject to common menu items. Blade row objects can be created, renamed, and deleted.



In order to specify the number of blades, the Blades Per Pitch child node is provided under the Blade Row node.

Blade Row Properties

Index	The sequential number of the addition of this blade row to the simulation, beginning at 1 (read-only).
Modes	The maximum value that a modal index m_i can take in this blade row, in the expressions for frequencies (Eqn. (2278)) and nodal diameter (Eqn. (2279)). This value subsequently determines how many unsteady frequencies are considered in this blade row. When you change the number of modes in any of the blade rows in the simulation, STAR-CCM+ reinitializes the new set of time-level solutions in all the blade rows. This reinitialization retains harmonic solutions at all the existing nodes.
User Specified Modes	A table column number. If a user-specified table is used as the modes table, then this field allows the user to select the column in the table that corresponds to the modal index m_i of the current blade row. If there are multiple blade rows in the simulation, this property must be defined for all blade rows, otherwise the default filtered modes table is used.
Time-levels	The number of unique time instants that are used to sample the solution in the blade row (read-only).
Number of Frequencies	The number of unique frequencies in the blade row (read-only).
Regions	The regions in the current blade row (read-only).
Interfaces	The interfaces in the current blade row (read-only).

Unique Positive Frequences	A list of the unique computed frequencies for the blade row (read only).

Blade Row Expert Properties

Maximum Nodal Diameter	This property sets the maximum allowable nodal diameter, which limits the modes included in the default filtered modes table.
Maximum Frequency	The maximum allowable frequency ω_i for the current blade row, which limits the modes included in the filtered modes table.

Creating Blade Rows

To create a blade row:

- Right-click the Harmonic Balance Flow and Energy > Blade Rows node, and select New from the pop-up menu.



A new blade row appears under the Blade Rows node. The default name Blade Row is concatenated with a number to ensure that the name is unique.

A child node Blades Per Pitch is created under the Blade Row node.



The properties of the blade row object include the ability to set the number of modes.

Setting Blades and Pitches

The properties of the Blades Per Pitch node are used to specify the Pitch P_{user} and Number of Blades B_{user}. With these inputs, STAR-CCM+ computes the number of blades:

$$B = \frac{2\pi B_{\text{user}}}{P_{\text{user}}} \tag{2258}$$

This number is used in the computation of the nodal diameters, and frequencies that occur in the blade row.

The Number of Blades B_{user} refers to the number of blades that are contained within the Pitch P_{user}.

Three-Dimensions

In three-dimensional problems, the Pitch P_{user} is best left at its default value of 2π radians, and the Number of Blades B_{user} is specified as the real number of blades on the full wheel.

For example, if there are 12 blades on the full wheel, specify the Pitch as 2π radians, and the Number of Blades as 12 so that $B = 12$. If the Pitch is specified as 2π radians, the Number of Blades should always be equal to the real number of blades on the wheel, irrespective of the number of blade passages that are modeled in the geometry.



Two-Dimensions

In two-dimensions, the Number of Blades B_{user} is best left at its default value of 1. The Pitch P_{user} is taken as the distance between one blade and its neighboring blade, or the distance between the upper and lower periodic boundaries. This value is specified in units of length. This distance could be inferred from the Periodic Transformation node of the periodic interface.

If more than one blade is modeled within the geometry, the Number of Blades is equal to the number of blade passages modeled.



For example, if only one blade passage is modeled with a pitch of 0.2 m (which could be the value that is shown in the Periodic Transformation node), you would specify the Pitch as `0.2 m`, and the Number of Blades as 1.

However, if two blade passages are modeled in the geometry with the same pitch, then you would specify the Pitch as `0.4 m`, and the Number of Blades as 2.

Blades Per Pitch Properties

Pitch	Specifies the pitch in angular units in three-dimensions and in length units in two-dimensions (linear length in Y direction).
Number of Blades	Specifies the number of blades within the given pitch. The number of blades should be an integer.

Defining the Region for a Blade Row

When you are done with defining the blade rows in the Blade Rows manager of the physics continuum, you must associate these rows with the regions that represent them. The Blade Row Specification node, that is located within the Physics Values folder for the region, is used to make this connection.



Blade Row Specification Properties

Blade row	Assigns a blade row to a region selected from the drop-down menu.

Choosing a Fictitious Rotation Rate Factor

When you use Consider All Blade Rows as your Frequency Compute Option, every blade row must be assigned a different frequency for the purposes of computation. This rule remains true even if some blade rows have identical frequencies in the User Specified Mode Table. Every blade is assigned an offset f proportional to the Fictitious Rotation Rate Factor ε and to the maximum user-specified rotation rate Ω_{max}.

As the computation works down the blade rows, the first time it comes to a blade that has the same frequency as a previous blade row, it adds an offset of f. The next time that it comes to a blade that repeats a former frequency, it adds $2f$, and the next time $3f$, and so on. This addition guarantees that no two blades have identical frequencies.

For example, consider the blade rows described in the following table:

	Stator	**Rotor**	**Stator**	**Rotor**	**Stator**
User-Specified Rotation Rate	0	Ω	0	Ω	0
New Rotation Rate	0	Ω	f	$\Omega + 2f$	$3f$

When the computation reaches the middle stator, it adds an offset of f because the starting rotation is the same as that of the first stator, 0.

When the computation reaches the second rotor, it adds an offset of $2f$ because the starting rotation is the same as that of the first rotor.

When the computation reaches the last stator, it adds an offset of $3f$ because the starting rotation is the same as that of the first two stators.

To set the Fictitious Rotation Rate factor, select the Blade Rows node, and set the value for the Fictitious Rotation Rate property.



The default value of f is 0.001. It is best to make the value small, so the departure from the original value stays small. See Fictitious Rotation Rate Computation.

The offset f is used only for entries in the modes table, not in computing the flow solution. The use of f is required when using the Consider All Blade Rows option under the Frequency Compute Option. Using f ensures that the unsteady solution at higher harmonics is exchanged between any two blade rows with the same rotation rate. The Fictitious Rotation Rate factor has no influence on the solution when the Consider Only Neighbors option is used.

Working with Mode Tables

For a single blade row, modes consist of the fundamental frequency and the harmonics, that is, integral multiples of the fundamental frequency.

For simulations with multiple blade rows, modes consist of the fundamental frequencies of each blade row interacting with neighboring blade rows, and the harmonics that result.

Mode tables are used to filter the modes included in the simulation when obtaining a solution. There are two methods of defining a mode table:

- the default filtered modes table that is calculated by the code, and
- a user-specified mode table.

Filtered Modes Table

The filtered modes table is created by default when the solution is initialized, and defines the modes that are included in the simulation. The Maximum Nodal Diameter and Maximum Frequency expert properties define the inclusion of modes in the table. The table consists of the following types of columns:

- m1, m2... : Modal indices for blade rows
- w1 [Row *n*], w2 [Row *n*]... : Modal indices for the wake blade row that is attached to modeled blade row *n*

  A wake blade row indicates a wake boundary condition that is applied to a stagnation inlet or a pressure outlet boundary. This boundary condition represents an external spatial disturbance that is applied to a modeled blade row.
- N: Nodal diameter
- omega1, omega2... : Frequencies for blade rows
- omegaw1 [Row *n*], omegaw2 [Row *n*]... : Frequencies for the wake blade row that is attached to the modeled blade row *n*

Using the tabulating feature, the information in the table can be displayed. For example:



	m1	m2	w1 [Row 1]	N	omega1	omega2	omegaw1 [Row 1]
0	-1	-1	-1	-40	-280	120	120
1	-1	-1	0	-30	-180	120	120
2	-1	-1	1	-20	-80	120	120
3	-1	0	-1	-22	-100	120	120
4	-1	0	0	-12	0	120	120
5	-1	0	1	-2	100	120	120
6	-1	1	-1	-4	80	120	120
7	-1	1	0	6	180	120	120
8	-1	1	1	16	280	120	120

In the above example, there are two blade rows and a wake blade row with the following properties:

	Wake Blade Row	Blade Row 1	Blade Row 2
Number of blades	10	12	18
Rotation rate (radians/s)	0	10	0
Number of modes	1	1	1

The nodal diameter is:

$$N = \sum_{i=1}^{n} m_i B_i \tag{2259}$$

where n is the number of blade rows, m_i are the nodal indices, and B_i the number of blades. STAR-CCM+ calculates the frequencies from the rotation rate for each blade, the effects of rotation from neighboring blades, and the disturbances coming from upstream and downstream.

The following illustrates how the Maximum Nodal Diameter properties can affect the filtered nodes table. Using the example above, if the Maximum Nodal Diameter is set to 25, the following table is created.


Tabular Data

	m1	m2	w1 [Row 1]	N	omega1	omega2	omegaw1 [Row 1]
0	-1	-1	1	-20	-80	120	120
1	-1	0	-1	-22	-100	120	120
2	-1	0	0	-12	0	120	120
3	-1	0	1	-2	100	120	120
4	-1	1	-1	-4	80	120	120
5	-1	1	0	6	180	120	120
6	-1	1	1	16	280	120	120
7	0	-1	0	-18	-180	0	0
8	0	-1	1	-8	-80	0	0
9	0	0	-1	-10	-100	0	0

Export... Close

All modes where $N > 25$ have been removed from the table. Definition of the Maximum Frequency property has a similar effect, with respect to *omega1* and *omega2*.

User Specified Mode Tables

This feature allows you to define a custom mode table for use within the simulation. Implementing a user-specified mode table overrides the information in the default filtered mode table.

Use the user specified mode table option to retain only a specific set of modes needed to capture the significant unsteady effects. To do this, first exporting the default filtered modes table and then removing the entries (the rows) that are not desired.

The user-specified mode table should only contain information on blade row modal indices (*m1* and *m2*) as the nodal diameter (*N*), and blade row frequencies (*omega1* and *omega2*) are defined by existing simulation properties.

To define a user-specified mode table, first import the table into STAR-CCM+, and then apply it to the physics model as follows:

- Select the Physics > Models > Harmonic Balance Flow and Energy > Blade Rows > User Specified Mode Table node.

- In the *Properties* window, set the Filtered Modes Table property to the required table in the drop-down menu.



Then, match the table column containing the modal indices to each blade row:

- Select the Blade Row node under the Blade Rows manager.

- In the *Properties* window, set the User Specified Modes property to the corresponding column in the table.



What Are the Harmonic Balance Turbulence Models?

These variants differ from their standard counterparts only in that they use the technique of the coupled Implicit Harmonic Balance solver.

Harmonic Balance Turbulence Model	Standard Model
K-Epsilon Models	
HB AKN K-Epsilon Low-Re	AKN K-Epsilon Low-Re
HB Realizable K-Epsilon	Realizable K-Epsilon
HB Realizable Two-Layer K-Epsilon	Realizable Two-Layer K-Epsilon
HB Standard K-Epsilon	Standard K-Epsilon
HB Standard K-Epsilon Low-Re	Standard K-Epsilon Low-Re
HB Standard K-Epsilon Two-Layer	Standard K-Epsilon Two-Layer
HB V2F K-Epsilon	V2F K-Epsilon
K-Omega Models	
HB Standard K-Omega	Standard K-Omega
HB SST (Menter) K-Omega	SST (Menter) K-Omega
Spalart-Allmaras Models	
HB Standard Spalart-Allmaras	Standard Spalart-Allmaras
HB High-Reynolds Number Spalart-Allmaras	High-Reynolds Number Spalart-Allmaras

Activating the Harmonic Balance Turbulence Models

To use turbulence with the Harmonic Balance method:

1. Right-click the Continua > Physics 1 > Models node and choose Select models...
2. Select the appropriate models. See Activating the Harmonic Balance Model.
3. Select Turbulent in the *Viscous Regime* group.

   Three types of turbulence, K-Epsilon, K-Omega, and Spalart-Allmaras, become available in the model selection dialog (*Physics 1 Model Selection* in the following screenshot), if you have the Auto-select recommended models checkbox activated.



Choose between the K-Epsilon, K-Omega, and Spalart-Allmaras models. Variants of each type of turbulence exist for use with Harmonic Balance.

- To use Realizable Two-Layer K-Epsilon:
  - Select K-Epsilon Turbulence in the *Reynolds-Averaged Turbulence* group.

    The HB Realizable Two-Layer K-Epsilon model is selected automatically.
- To use any of the other K-Epsilon models:
  - Deactivate the Auto-select recommended models checkbox.
  - Select K-Epsilon Turbulence in the *Reynolds-Averaged Turbulence* group.

    The following K-Epsilon models become available:
    - HB AKN K-Epsilon Low-Re
    - HB Realizable K-Epsilon

    - HB Realizable K-Epsilon Two-Layer
    - HB Standard K-Epsilon
    - HB Standard K-Epsilon Low-Re
    - HB Standard K-Epsilon Two-Layer
    - HB V2F K-Epsilon
- To use either of the K-Omega models:
  - Deactivate the Auto-select recommended models checkbox.
  - Select K-Omega Turbulence in the *Reynolds-Averaged Turbulence* group.

    The following K-Omega models become available:
    - HB Standard (Wilcox) K-Omega
    - HB SST (Menter) K-Omega
- To use Standard Spalart-Allmaras:
  - Select Spalart-Allmaras in the *Reynolds-Averaged Turbulence* group.

    The Harmonic Balance Standard Spalart-Allmaras turbulence model is selected automatically.
- To use High-Reynolds Number Spalart-Allmaras:
  - Deactivate the Auto-select recommended models checkbox.
  - Select Spalart-Allmaras in the *Reynolds-Averaged Turbulence* group.

    Both Harmonic Balance Spalart-Allmaras models, High-Reynolds Number and Standard, become available for selection.

Defining a Wake

The wake is an unsteady boundary condition that is specified on a stagnation inlet or a pressure outlet boundary. The steady state values for pressure, temperature, and velocity are specified separately by specifying the total pressure, total temperature, and flow angle for the stagnation inlet boundary, and the static pressure and static temperature for the pressure outlet boundary.

The wake specification adds the unsteady components of these variables, based on pre-specified, or user-specified profile functions. STAR-CCM+ automatically computes and extracts the unsteady part of the user-specified wake data and discards the mean values.

The following boundary conditions are available for the Harmonic Balance Flow and Energy model:

- Non-reflecting option
- Wake specification

These boundary conditions are in addition to the ones for the steady part of the solution, which you control by specifying flow direction, total pressure and total temperature.

What Is the Non-Reflecting Option?

This boundary condition is the unsteady version of the Non-Reflecting Option for the Coupled Flow model. Ticking its checkbox makes the Wake Specification boundary condition available, which has seven different options for specifying a wake.



Non-Reflecting Option Properties

Non-Reflecting Option	Specifies whether to use the non-reflecting option.	
	Activated	Use the non-reflecting option.
	Deactivated	Do not use the non-reflecting option.

What Is the Non-Reflecting Mode Specification?

This boundary value is activated automatically when the Non-Reflecting Option boundary condition is enabled and is the multiple frequency version of the Non-Reflecting Boundary Condition for the Coupled Flow model. This boundary condition prevents spurious numerical reflection of the solution into the solution domain.

You can specify the number of modes to retain using the Number of modes property. The number of modes you specify should be less than the number of cells in the circumferential direction.



Properties

Number of modes	Specifies number of modes to retain for the non-reflecting boundary condition. The number of modes should be less than the number of cell faces in the circumferential direction.

What Is the Wake Specification?

This node allows you to determine how the wake for a boundary is specified. It is activated when the Harmonic Balance Flow and Energy model and the Non-Reflecting Option boundary condition are enabled.



The Sine, Gaussian, and Hodson options are standard function shapes that are chiefly for academic use, while the other options are used more commonly for applications.

The wake can only be specified at the inlet boundary or the outlet boundary. However, the STAR-CCM+ GUI does not prevent you from entering values for the wake at both the inlet and outlet boundaries.

Properties

Method Selects the method for specifying the wake, either one of the standard shapes, or one based on user table data.

None	The option is turned off for the boundary. No nodes are added to the Physics Values node.
Sine	This option uses a standard sine curve for unsteadiness. Selecting it adds a Sine Wake node to the Physics Values node.
Gaussian	This option uses a standard Gaussian shape for unsteadiness. Selecting it adds a Gaussian Wake node to the Physics Values node.
Hodson	This option uses a standard Hodson shape for unsteadiness. Selecting it adds a Hodson Wake node to the Physics Values node.
Fourier Data	This option lets you use Fourier data to add unsteady pressure, temperature, and velocity. Selecting it adds a Wake Fourier Coefficients node to the Physics Values node.
Spatial Annular Data	This option works in cylindrical space (r, θ), or (r, θ, z) for three-dimensional cases, to add unsteady pressure, temperature, and velocity. Selecting it adds a Spatial Annular Wake Data node to the Physics Values node.
Spatial Linear Data	This option works in Cartesian space (x, y), or (x, y, z) for three-dimensional cases, to add unsteady pressure, temperature, and velocity. Selecting it adds a Spatial Linear Wake Data node to the Physics Values node.

Understanding Wake Options

The following boundary values are available for the Harmonic Balance Flow and Energy model:

- Gaussian wake
- Hodson wake
- Sine wake
- Spatial annular wake data

- Spatial linear wake data
- Wake Fourier coefficients

You must set the Pitch and the Number of Blades in the Blades Per Pitch child node of the Sine, Gaussian, Hodson, Fourier Data, Spatial Annular Data, or Spatial Linear Data wakes. These values are required for computing blade row frequencies.

Blades Per Pitch Properties

Pitch	Specifies the pitch in angular units in three-dimensions and in length units in two-dimensions (linear length in Y direction).
Number of blades	This property defines number of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the number of blades on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the number of blades on the downstream blade row is required.

Gaussian Wake

The Gaussian wake function is defined in Eqn. (2274).



Properties

Width	The wake width, W. This property is a fractional wake width and is a non-dimensional number. For example, a width of $W = 0.2$ would mean that the width of the unsteady disturbance (as defined in the wake functions) would be 20% of the pitch.
Defect	The wake defect, D. This property is the amplitude of the unsteady wake function, which defines a velocity defect distribution along the circumferential direction. For instance, specifying the defect D as 0.2 would mean that the amplitude of the unsteady wake function would be equal to $0.2V$, where V refers to the magnitude of the absolute velocity averaged over all the cell faces in the circumferential direction.
Rotation Rate	This property defines angular velocity of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the angular velocity on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the angular velocity on the downstream blade row is required.

Hodson Wake

The Hodson wake function is defined in Eqn. (2275).



Properties

Width	The wake width, W. This property is a fractional wake width and is a non-dimensional number. For example, a width of $W = 0.2$ would mean that the width of the unsteady disturbance (as defined in the wake functions) would be 20% of the pitch.
Defect	The wake defect, D. This property is the amplitude of the unsteady wake function, which defines a velocity defect distribution along the circumferential direction. For instance, specifying the defect D as 0.2 would mean that the amplitude of the unsteady wake function would be equal to $0.2V$, where V refers to the magnitude of the absolute velocity averaged over all the cell faces in the circumferential direction.
Rotation Rate	This property defines angular velocity of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the angular velocity on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the angular velocity on the downstream blade row is required.

Sine Wake

The Sine wake function is defined in Eqn. (2276).



Properties

Defect	The wake defect, D. This property is the amplitude of the unsteady wake function, which defines a velocity defect distribution along the circumferential direction. For instance, specifying the defect D as 0.2 would mean that the amplitude of the unsteady wake function would be equal to $0.2V$, where V refers to the magnitude of the absolute velocity averaged over all the cell faces in the circumferential direction.
Rotation Rate	This property defines angular velocity of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the angular velocity on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the angular velocity on the downstream blade row is required.

Spatial Annular Wake Data

The table that you select must specify the cylindrical components of the velocity field for the wake, along with pressure and temperature. For three-dimensional cases, the table must provide:

- Tangential coordinate
- Radial coordinate
- Normal velocity component
- Tangential velocity component
- Radial velocity component
- Temperature
- Pressure

The column headings in the table you provide are mapped to the required quantities using the properties of the node. In two-dimensional cases, the radial coordinate and radial velocity component are not required.

Properties of Spatial Annular Wake Data

Annular Table: Theta	Selects the table column containing the tangential coordinate.
Annular Table: P	Selects the table column containing the pressure data.
Annular Table: Vn	Selects the table column containing the normal component of velocity.
Annular Table: Vt	Selects the table column containing the tangential component of velocity.
Annular Table: Vr	Selects the table column containing the radial component of velocity.
Annular Table: T	Selects the table column containing the temperature data.
Annular Table: R	Selects the table column containing the radial coordinate.
Spatial Wake Data Table	Selects the table to use.
Number of curves	Specifies the number of curves (radial or axial strips) in the table.
Number of data points per curve	Specifies the number of points in the circumferential direction for each curve in the table.
Number of modes to extract	Specifies the number of Fourier modes to compute from the table data.

Rotation Rate	This property defines angular velocity of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the angular velocity on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the angular velocity on the downstream blade row is required.

Expert Properties

Wake Data Position Option	Determines whether wake data is positioned on the boundary, or at an offset to it.	
	At the boundary	The wake data is positioned at the boundary.
	At a specified location	The wake data is offset by an amount that is given by the Wake Data Position, which is applied in the direction normal to the wake boundary
Wake Data Position	The distance, positive, or negative, by which the wake data is offset from the boundary in the direction normal to the boundary.	

Spatial Linear Wake Data

The table that you select must specify the cartesian components of the velocity field for the wake, along with pressure and temperature. For three-dimensional cases, the table should provide:

- X coordinate
- Y coordinate
- Z coordinate
- X component of velocity
- Y component of velocity
- Z component of velocity
- Temperature
- Pressure

The column headings in the table you provide are mapped to the required quantities using the properties of the node. In two-dimensional cases, the Z coordinate and Z component of velocity are not required.

Properties of Spatial Linear Wake Data

Linear Table: P	Selects the table column containing the pressure data.
Linear Table: T	Selects the table column containing the temperature data.
Linear Table: Vx	Selects the table column containing the X component of velocity.
Linear Table: Vy	Selects the table column containing the Y component of velocity.
Linear Table: Vz	Selects the table column containing the Z component of velocity.
Linear Table: X	Selects the table column containing the X coordinate.
Linear Table: Y	Selects the table column containing the Y coordinate.
Linear Table: Z	Selects the table column containing the Z coordinate.
Spatial Wake Data Table	Selects the table to use.
Number of curves	Specifies the number of curves (radial or axial strips) in the table.
Number of data points per curve	Specifies the number of points in the circumferential direction for each curve in the table.

Number of modes to extract	Specifies the number of Fourier modes to compute from the table data.
Rotation Rate	This property defines angular velocity of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the angular velocity on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the angular velocity on the downstream blade row is required.

Expert Properties

Wake Data Position Option	Determines whether wake data is positioned on the boundary, or at an offset to it.	
	At the boundary	The wake data is positioned at the boundary.
	At a specified location	The wake data is offset by an amount that is given by the Wake Data Position, which is applied in the direction normal to the wake boundary
Wake Data Position	The distance, positive, or negative, by which the wake data is offset from the boundary in the direction normal to the boundary.	

Wake Fourier Coefficients

The table that you select needs a special format. It specifies the non-zero Fourier modes to control the wake, and the column headers must be as follows:

```
'Pnr','VNnr','VTnr','VRnr','Tnr',
'Pni','VNni','VTni','VRni','Tni'
```

where `n` in each of the above column headers corresponds to the non-zero modes, that is, `n` can take values 1, 2, 3, and so on. In the above, `P`, `VN`, `VT`, `VR`, and `T` refer to pressure, velocity components (normal, tangential, radial), and temperature. The entries `r` and `i` refer to the real and imaginary components of these modes. For two-dimensional cases, `VR` would not be required. Also, the order in which the columns appear does not matter, as long as they are named as described above.

For example, if you wish to specify two modes, the table contains the following data:

```
'P1r','VN1r`,'VT1r`,'VR1r`,'T1r`,
'P1i`,'VN1i`,'VT1i`,'VR1i`,'T1i`,
'P2r`,'VN2r`,'VT2r`,'VR2r`,'T2r`,
'P2i`,'VN2i`,'VT2i`,'VR2i`,'T2i`
```

While specifying the number of modes to use from this table, STAR-CCM+ allows you to specify an integral number from 0 to the number of modes present in the file. For the two-mode table example shown above, you could choose the number of modes as 0 (implying "use no wake data from file table"), 1 ("use only the first mode from the table"), or 2 ("use both the modes from the table").

Properties

Fourier Coefficients Table	Selects the table to use.
Number of modes to use	Specifies the number of modes to use from the table data.
Rotation Rate	This property defines angular velocity of blades in the blade row corresponding to the wake. If the wake is defined at an inlet, the angular velocity on the upstream blade row is required. Similarly, if the wake is defined at an outlet, the angular velocity on the downstream blade row is required.

Expert Properties

Wake Data Position	Specifies the spatial location of the wake table data.
Wake Data Position Option	Determines where the wake data are specified.

At the boundary	The wake data are at the boundary.
At a specified location	The wake data are at a spatial location other than the boundary.

What Is the Harmonic Balance Flutter Motion?

Harmonic Balance Flutter motion is used in simulations that involve the Harmonic Balance method with blade vibration. The Harmonic Balance Flutter node represents this motion, and is subject to

common menu items.



Harmonic Balance Flutter Properties

Morphing Order

The morpher can operate in two different modes: one region at a time, or one motion at a time. This property selects the mode to use.

- **Regionwise** — Each region that is associated with this motion is morphed independently, one region at a time. This option is used, for example, in a fluid-solid case, where the solid region has a solid displacement motion and the fluid region has a morphing motion.
- **Motionwise** — All regions that are associated with this motion are morphed at the same time. This option is recommended for use in cases where internal interfaces are present between regions, if the floating morpher boundary method is applied to the boundaries on either side of the interface, and both regions are set to use the same morphing motion.

Linear Fitter

This option extracts a linear transformation from the prescribed motion, and uses this transformation to move vertices outside of the main morpher. Any remaining motion is treated using the multiquadric procedure as usual. If the linear part constitutes most the motion and the deformation part is small, the main morpher is expected to perform more robustly if the Morph From Zero property is enabled on the Morpher solver.

- **Activated** — Linear motion is treated separately from multiquadric interpolation.
- **Deactivated** — All motion is treated in multiquadric interpolation.

Harmonic Balance Flutter Expert Properties

Automatic Thin-out	Applies an additional stage of control vertex thinning after the user-defined thin factors have been applied. Tick the Morph from Zero property on the Morpher solver when using this option. Activated: Automated thinning is applied. Deactivated: No automated thinning is applied. Only user-defined thin factors are accounted for.
Automatic Thin-out Cl	This property controls the target spacing between control points. If the value is smaller, the spacing is smaller resulting in denser control point distribution. If the value is larger, it coarser control point distribution.

Setting Up a Simulation with Harmonic Balance Flutter Motion

To use blade vibration in a simulation, follow the steps below:

- Activate the Harmonic Balance Flow and Energy continuum model.
- Create a Harmonic Balance Flutter motion.
- Apply the Harmonic Balance Flutter motion to a region using the Motion Specification region value.
- Specify the flutter parameters and, if necessary, control points and their associated displacements.
- Specify the properties of the Blade Flutter Solver node added to the Solvers manager.
- Set up the properties for the Morphing Option boundary condition.
- Finally, specify the Harmonic balance flutter boundary values.

To monitor the vibratory force, or the real or imaginary part of that force, use the Generalized Vibratory Force report.

Specifying Flutter Parameters

You specify flutter parameters in the properties of the Physics Values > Flutter Parameters node of a region.



Before setting the flutter parameters:

- Set the Motion Specification of the region to flutter motion.

  You can select flutter motion for different blade rows. There is no limit to the number of blade rows that can be modeled with flutter. However, in most practical cases, you use the flutter model only in one row, even in a multi-blade row case.

- Associate the region to a blade row.

To set the flutter parameters:

- Select the Regions > [Region] > Physics Values > Flutter Parameters node, and set the Flutter Frequency, Flutter Phase Angle, and Modes properties.

Flutter Parameters Properties

Flutter Frequency	The frequency ω of unsteadiness. See Eqn. (2281).
Flutter Phase Angle	The inter-blade phase angle σ. See Eqn. (2281).
Modes	The number of modes to retain for the blade flutter.

What Is the Blade Flutter Solver?

The Blade Flutter solver displaces the computational mesh at each of the time levels, based on user-specified blade displacements, and stores the mesh metrics (cell volumes and centroids, face areas, and centroids) at each time level for use by the Implicit Harmonic Balance Solver.

These metrics are computed and stored only once, during the first iteration. They are recomputed when the flutter boundary or region conditions change. You can force the recomputation of the mesh metrics by right-clicking the Blade Flutter solver node, and selecting Reset Solver.

The Blade Flutter solver is activated when you select the Harmonic Balance Flutter motion. The Blade Flutter solver node represents this solver, and has properties and a pop-up menu.



Blade Flutter Solver Properties

Mode Shape Scale Factor

Scales the user-specified blade displacement magnitudes.
The default value of this property is 1, indicating that no scaling is applied to the input mode shape.
This scaling is required as the mode shapes that are extracted from third-party codes are normalized in a way that gives a non-physical amplitude, which can be orders of magnitude too large in most cases.

Expert Properties

Verbosity

Provides additional output for this value while the simulation is running. This property may be useful for debugging problems when they occur.

None	Leaves output at its default.
Low	Adds messages as to which region is morphed.
High	Generates precise equation information about the morphing process as the solver converges.

Maximum Iterations

The number of iterations the morpher solver is allowed to take to calculate the vertex displacements to the specified tolerance (set by the Morpher Tolerance property).

Skip Thinned Vertex Corrector

Whether to disable the Thinned Vertex Corrector. This option is provided for backward compatibility only. Keep this property deactivated at all times.

Activated	Disables the Thinned Vertex Corrector
Deactivated	Enables the Thinned Vertex Corrector

Morpher Tolerance

Non-dimensional parameter that is used to specify acceptable accuracy in morphing. This value is multiplied by the longest edge length of the bounding box enclosing the morphing region.
Morpher tolerance is the non-dimensional magnitude of the noise that enters in the displacement field. If your grid size is O(10)m, the noise level for the default setting of 1e-7 is 1e-6 m. This noise value must be well below your smallest grid size (typically layer thickness). Larger tolerance values allow the morpher to run faster.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. Do not activate this feature as your first choice: it is primarily a debugging tool and not supported for normal simulation work.

Activated The solver does not update the solution.

Deactivated The solver updates the solution.

Blade Flutter Solver Pop-Up Menu

Reset Solver

Initializes the Blade Flutter solver. This action has the same effect as initializing and clearing the solution combined, except the only solution data that are cleared are that of the Blade Flutter solver.

See also: common menu items.

Morphing Option

This node allows you to determine how the morphing for a boundary is specified. It is activated when the Harmonic Balance Flutter motion is enabled.



Properties

Method Sets the morphing condition.

Fixed	This option adds the Morpher: ThinFactor physics value node. The Harmonic Balance Flutter motion and the Morpher motion share this node. Points on this boundary are treated as control points with zero displacement.
In Plane	This option adds the Morpher Plane Options physics value node. The Harmonic Balance Flutter motion and the Morpher motion share this node. Points on this planar boundary are allowed to slide in the plane. This is implemented by damping the displacements, and is much less expensive than using control points.
Displacement	Points on this boundary have specified displacement. This option, which is available when the Harmonic Balance Flutter motion is enabled, adds the following nodes to the Physics Values node: • Imaginary Displacement • Morpher: ThinFactor (also used by the Morpher motion) • Real Displacement
Floating	The option is turned off for the boundary. No nodes are added to the Physics Values node. Points on this boundary have no boundary condition; they move wherever the other control points dictate.

Harmonic Balance Flutter Boundary Values

The following boundary values are available for the Harmonic Balance Flutter motion when the Morphing Option Method is set to Displacement:

- Morpher: Imaginary Displacement

- Morpher: Real Displacement



The physical displacements are calculated according to Eqn. (2277) at each time level. The real and imaginary displacements specify the amplitude of the periodic function in Eqn. (2277).

In addition, the Harmonic Balance Flutter motion and the Morpher motion share the following boundary values:

- Morpher: ThinFactor
- Morpher Plane Options

What Is the Imaginary Displacement?

This boundary value is activated when the Morphing Option boundary condition is set to `Displacement`. It is input as a vector profile and is used to specify the amplitude of physical displacements, cf Eqn. (2277).

What Is the Real Displacement?

This boundary value is activated when the Morphing Option boundary condition is set to `Displacement`. It is input as a vector profile and is used to specify the amplitude of physical displacements, cf Eqn. (2277).

What Is the Implicit Harmonic Balance Solver?

The Implicit Harmonic Balance solver controls the solution update for the Harmonic Balance Flow and Energy model. It is used for implicit integration, using a coupled algebraic multi-grid method. It has properties and is subject to common menu items.



Implicit Harmonic Balance Solver Properties

Courant Number	The value of CFL in Eqn. (140) that controls the local time-step at each iteration. The default Courant number is 5. In general, and depending on the initial conditions, it might be necessary to start at a lower Courant number, and then increase it.

Expert Properties

Reduced Frequency Limiter	Governs stability for convergence, especially for large far-field cells. The entry specifies ω_r in the ratio ω_r/ω_{max}, which limits the magnitude of the time-step.

For information about properties in the *Expert* category, see Expert Properties, p. 6385.

Post-Processing Harmonic Balance Solution Data

Post-processing for harmonic balance solution data offers the following features:

- The solution is known on every blade in the row, not just on the meshed blades.
- The solution is known for all times, so you can derive views of blade positions and field values for any time during the run.
- The solution can be represented and viewed in the frequency domain. This can be helpful for understanding the unsteady flow features.

To examine and analyze the results of a harmonic balance simulation, create a harmonic balance solution view and use it as the basis for creating scenes, reports, and plots. See Applying Recorded Solution Views to Scenes and subsequent sections. Field functions and an annotation specific to harmonic balance are available.

Example Post-Processing Procedure

The following screen shot shows a scene of three blade rows, with four passages for each.



This scene was produced using the following procedure:

1. Right-click Continua > Physics > Models > Harmonic Balance Flow and Energy and select New HB Solution View. This creates an HB Solution View node under the Solution Views node.
2. Open the scalar scene, right-click on it, and select Apply Representation > HB Solution View.
3. Right-click on the color bar and select Harmonic Balance > HB: Turbulent Viscosity > Time Function.
4. Select Solution Views > HB Solution View > Blade Row and, in each Blade Row sub-node, set the Blade Passage Selection property to `[-2,1]`.
5. Select Solution Views > HB Solution View > Physical Time and set the Physical Time property to `5.0E-4 s`.

By varying the settings of the HB Solution View node, you can derive scenes for other physical times, or in time level, Fourier, or filtered modes. You can also derive plots and reports, as with other solution views.

What Is the HB Solution View?

The HB Solution View node controls the properties of harmonic balance solution view for all selected blade passages of all blade rows in the simulation.



To create a harmonic balance solution view:

1. Right-click Continua > Physics > Models > Harmonic Balance Flow and Energy and select New HB Solution View. This creates an HB Solution View node under the Solution Views node.
2. Select the HB Solution Views node and select a view mode from the Mode property drop-down menu. Set a new value for the Flutter Amplification property at this time, if necessary.
3. Select the HB Solution Views > Blade Rows > Blade Row *n* nodes, one by one, and set their properties.

The blade passages selected for all blade rows are shown in a *Scalar* scene under Scenes > Scalar Scene *n*.

HB Solution View Properties

Representation	The name of the HB Solution View node. (Read only.)	
Mode	Specifies the mode in which to visualize the fields selected for the individual blade rows.	
	Physical Time	The physical time at which to visualize the blade row fields. This is the default mode.
	Time Level	The time level at which to visualize the blade row fields.
	Fourier Mode	The number of the mode (harmonic) at which to visualized the blade row fields.
	Filtered Mode	The number of the filtered mode (harmonic) at which to visualize the blade row fields.
Flutter Amplification	Magnification factor for blade vibration. The default is 1 (no amplification).	

See Harmonic Balance Solution View Field Functions for a listing of the available fields.

Physical Time Mode

In this mode, the solution view shows the blade row fields at a given physical time. Cycle Time is displayed to help you select the point in the rotation for visualization.

Physical Time Mode Properties

Physical Time	The physical time at which to visualize the blade row fields.
Cycle Time	The time taken to complete one rotation of the rotor. (Read only.)

Time Level Mode

In this mode, the solution view shows the blade row fields at a given time level—that is, at a given point in the repeating periodic flow.

Time Level Mode Properties

Time Level	The time level at which to visualize the blade row fields.

Fourier Mode

In this mode, the solution view shows the blade row fields for a given harmonic when the periodic flow is analyzed into a Fourier series.

Fourier Mode Properties

Mode	The number of the mode (harmonic) at which to visualized the blade row fields.	
Output Option	The feature of the Fourier mode being visualized.	
	Real	The real part of the Fourier mode.
	Imaginary	The imaginary part of the Fourier mode.
	Magnitude	The magnitude of the Fourier mode.
	Phase	The phase of the Fourier mode.

Filtered Mode

In this mode, the solution view shows the blade row fields for a given mode taken from the Filtered Modes Table under the Harmonic Balance Flow and Energy node.

Filtered Mode Properties

Mode	The number of the filtered mode (harmonic) at which to visualize the blade row fields.	
Output Option	The feature of the filtered mode being visualized.	
	Real	The real part of the filtered mode.
	Imaginary	The imaginary part of the filtered mode.
	Magnitude	The magnitude of the filtered mode.
	Phase	The phase of the filtered mode.

Blade Row

One Blade Row node appears under the HB Solution Views > Blade Rows node for each blade row. Enter the selection as an ordered pair in square brackets, with a negative integer for the number of lagging blades and a positive integer for the number of leading blades. Examples:

- `[0, 0]` is the original blade alone. This is the default.
- `[-9, 8]` is nine lagging blades through eight leading blades.

Blade Row Properties

Blade Passage Selection	The blades in the passage selected for visualization.
Blade Passage Range	The number of blades available in the blade row. See Eqn. (2258). (Read only.)

Filtered Mode Field Functions

Filtered mode field functions allow you to visualize the Fourier coefficients corresponding to each mode (or harmonic) for multi-row harmonic balance cases. As the total number of modes that are made available in a multi-stage simulation could be a large number, you have the option to selectively choose the modes of interest. To control the number of such field functions to be made available for visualization:

- Right-click the Continua > Physics 1 > Models > Harmonic Balance Flow and Energy > Blade Rows > Filtered Modes Table node and select Tabulate.



- In the table that appears, review the values and note the numbers of the mode rows that you would like to post-process.



Tabular Data

	m1	m2	m3	N	omega1	c
0	-1	-1	-1	-23.5619	2.4607E+3	-4.92
1	-1	-1	0	-13.7445	2.4607E+3	-1.84
2	-1	-1	1	-3.927	2.4607E+3	1.230
3	-1	0	-1	-15.708	0	-4.92
4	-1	0	0	-5.8905	0	-1.84
5	-1	0	1	3.927	0	1.230
6	-1	1	-1	-7.854	-2.4607E+3	-4.92
7	-1	1	0	1.9635	-2.4607E+3	-1.84
8	-1	1	1	11.781	-2.4607E+3	1.230
9	0	-1	-1	-17.6715	2.4607E+3	-3.07
10	0	-1	0	-7.854	2.4607E+3	0
11	0	-1	1	1.9635	2.4607E+3	3.075
12	0	0	-1	-9.8175	0	-3.07
13	0	0	0	0	0	0

Export... Close

- To close the table, click Close.
- With the Filtered Modes Table node still selected, enter the row numbers in the Modes to visualize property.



After you enter this list, a new set of field functions, with Filtered Mode in their names, gets generated for the variables of interest such as density, velocity, temperature, and pressure. These field functions can be used to visualize the modal content of a variable corresponding to the chosen mode.

In the following example, the chosen modes are 0, 4, and 25. The variable Density has a subset called Filtered Mode. Each filtered mode (being a complex number) has four components: real, imaginary, magnitude, and phase. These components correspond to the real and imaginary parts of the complex number, the complex magnitude, and the phase, respectively.



Harmonic Balance Formulation

This section describes the following formulation for the Harmonic Balance method:

- Trigonometric series
- Time derivative
- Harmonic balance equations
- Wake functions
- Real and imaginary displacements
- Blade row frequencies computation for multi-row simulation
- Blade row nodal diameter computation for multi-row simulation
- Fictitious rotation rate computation for multi-row simulation
- Fluttering motion of a rotor

Trigonometric Series

The starting point for the Harmonic Balance method is the representation of a real, time-periodic function $w(\mathbf{x}, t)$ with period T by a trigonometric series in m harmonics:

$$w(\mathbf{x}, t) = w_0(\mathbf{x}) + \sum_{n=1}^{m} w_{2n-1}(\mathbf{x})\cos\omega nt + w_{2n}(\mathbf{x})\sin\omega nt \tag{2260}$$

where:

m is the number of user-specified modes or harmonics,

$w(\mathbf{x}, t)$ is the time domain or physical solution at the spatial location $\mathbf{x} = (x, y, z)$ and time t,

$N = 2m + 1$ is the number of time levels,

w_0 is the mean solution at position $\mathbf{x}$, and,

w_n are the Fourier coefficients.

ω is the user-specified excitation frequency, and

$$T = \frac{2\pi}{\omega} \tag{2261}$$

This can also be represented in complex form as:

$$w(\mathbf{x}, t) = \sum_{n=-m}^{m} w_n(\mathbf{x})e^{j\omega nt} \tag{2262}$$

where the summation goes from $-m$ to m, and $w_n(\mathbf{x})$ and $w_{-n}(\mathbf{x})$ are complex conjugate pairs.

At each point in space (x, y, z) Eqn. (2260) contains N coefficients w_n that can be uniquely determined by N time-samples of the continuous function, $w_k^* = w(t = t_k)$, as follows:

$$w_0 = \frac{1}{N}\sum_{k=0}^{N-1} w_k^* \tag{2263}$$

$$w_{2n-1} = \frac{2}{N}\sum_{k=0}^{N-1} w_k^* \cos\omega nt_k;\ n = 1\ldots m \tag{2264}$$

$$w_{2n} = \frac{2}{N}\sum_{k=0}^{N-1} w_k^* \sin\omega nt_k;\ n = 1\ldots m \tag{2265}$$

where the k-th discrete time level is evaluated at time

$$t_k = \frac{kT}{N} \tag{2266}$$

The N vector of functional values, w_k^* at the discrete time levels can be recovered from the N vector of coefficients w_k by use of Equations (2260), (2261), and (2266) as follows:

$$w_k^* = w_0 + \sum_{n=1}^{m} \left[w_{2n-1} \cos 2\pi \frac{nk}{N} + w_{2n} \sin 2\pi \frac{nk}{N} \right]; \; k = 0 \ldots N \tag{2267}$$

Note that it is really w_k^* that is stored in STAR-CCM+. The values w_n are computed when needed internally, or for reporting and plotting purposes when you request them.

Time Derivative

From Equations (2261), (2266) and (2267), the time derivative of the discrete solution at each time level is given by:

$$\frac{\partial w_k^*}{\partial t} = \omega \sum_{n=1}^{m} n \left(w_{2n} \cos 2\pi \frac{nk}{N} - w_{2n-1} \sin 2\pi \frac{nk}{N} \right) \tag{2268}$$

Substituting Equations (2264) and (2265) into Eqn. (2268) yields

$$\frac{\partial w_k^*}{\partial t} = \frac{2\omega}{N} \sum_{n=1}^{m} n \left(\cos 2\pi \frac{nk}{N} \sum_{j=0}^{N-1} w_k^* \sin 2\pi \frac{nj}{N} - \sin 2\pi \frac{nk}{N} \sum_{j=0}^{N-1} w_k^* \cos 2\pi \frac{nj}{N} \right) \tag{2269}$$

We can represent the N vector of time-derivatives given by Eqn. (2269) as

$$\frac{\partial w^*}{\partial t} = [\mathbf{D}] w^* \tag{2270}$$

where the transform operator $[\mathbf{D}]$ has elements D_{ij} given by:

$$D_{ij} = \frac{2\omega}{N} \sum_{n=1}^{m} n \left(\cos 2\pi \frac{ni}{N} \sin 2\pi \frac{nj}{N} - \sin 2\pi \frac{ni}{N} \cos 2\pi \frac{nj}{N} \right) \tag{2271}$$

Diagonal elements D_{ii} are zero.

Harmonic Balance Equations

The transform operator is used to provide a time derivative source term for the conservation equations at each time level. Applying the initial governing equation for the coupled solver, Eqn. (106), to each time level and using the time derivative operator gives:

$$\oint [\mathbf{F}^* - \mathbf{G}^*] \cdot d\mathbf{a} = \int_V \mathbf{H}^* dV - \int_V \mathbf{S}^* dV \tag{2272}$$

where $\mathbf{S}^* = \mathbf{D}W^*$ is the time derivative source term. Eqn. (2272) is a steady flow equation which can be solved using a pseudo-time marching technique to converge to steady state.

Harmonic Time Functions

The value of any variable Q at a user-specified time t is given by:

$$Q(t) = \sum_{m=-NF}^{NF} Q_m e^{i(\omega_m t + sN_m d\Theta)} \tag{2273}$$

where:

- NF is the number of frequencies (or harmonics) in the blade row in question.
- Q_m is the Fourier coefficient associated with the mth harmonic.
- ω_m is the frequency associated with the mth harmonic in Eqn. (2278).
- N_m is the nodal diameter for the mth harmonic in Eqn. (2279).
- $d\Theta$ is the pitch of the modeled blade passage.
- s is 0 for the modeled passage, +1 for the phase-lagged up passage, and -1 for the phase-lagged down passage.

When the Mean Flow Reconstruction option is turned on, the only terms retained while computing $Q(t)$ in Eqn. (2273) are the terms corresponding to the mean mode (that is, $m = 0$) and to any mode that is artificially introduced because of the Consider All Blade Rows option.

Wake Functions

The Gaussian, Hodson, and Sine wake functions are defined as follows, respectively:

$$f(\eta) = D \exp\left(\frac{-\eta^2}{2W^2}\right) \tag{2274}$$

$$f(\eta) = D\left[\max\left(0, \left[1 - \left|\frac{\eta}{W}\right|^{\frac{3}{2}}\right]\right)\right]^2 \tag{2275}$$

$$f(\eta) = D\sin(2\pi\eta) \tag{2276}$$

where η is a normalized pitch that varies between -0.5 to 0.5, D is the defect and W is the width.

Real and Imaginary Displacements

If v_r and v_i are the real and imaginary displacements that you specify, then the physical displacements dx at every time level k are computed as:

$$\begin{aligned} dx(t_k) &= (v_r, v_i) \exp\left(\frac{i(2\pi k)}{N}\right) + (v_r, -v_i) \exp\left(\frac{-i(2\pi k)}{N}\right) \\ &= 2v_r\left(\cos\frac{(2\pi k)}{N}\right) - 2v_i\left(\sin\frac{(2\pi k)}{N}\right) \end{aligned} \tag{2277}$$

where t_k represents the physical time at time level k (where k varies from 0 to $N-1$), N is the number of time levels, and (v_r, v_i) represents a complex number with v_r and v_r being the real and imaginary parts of the complex number.

Blade Row Frequencies Computation

For a case with N blade rows, the complete set of frequencies for blade row j is given by:

$$\omega^j = \sum_{\substack{i = 1 \\ i \neq i}}^{N} m_i B_i (\Omega_j - \Omega_i) \tag{2278}$$

where, for blade row or wake i:

- m_i is the modal index and can take all values between $-M_i$ and M_i. M_i is the number of modes.
- B_i is the number of blades.
- Ω_i is the rotation rate.

When using the Consider All Blade Rows option, i represents all blade rows and wakes, except blade row j. When using the Consider Only Neighbors option, only the adjacent blade rows that share a common indirect interface with blade row j are considered.

Blade Row Nodal Diameter Computation

For a case with N blade rows, the complete set of nodal diameters that correspond to each frequency computed above (in Eqn. (2278)) for blade row j is given by:

$$N^j = \sum_{\substack{i = 1 \\ i \neq i}}^{N} m_i B_i \tag{2279}$$

where, for blade row or wake i:

- m_i is the modal index and can take all values between $-M_i$ and M_i. M_i is the number of modes.
- B_i is the number of blades.

When using the Consider All Blade Rows option, i represents all blade rows except blade row j. When using the Consider Only Neighbors option, only the adjacent blade rows that share a common indirect interface with blade row j are considered.

Fictitious Rotation Rate Computation

When you use Consider All Blade Rows as your Frequency Compute Option, every blade must be assigned a different frequency for the purposes of computation, even if some have identical frequencies in the User Specified Mode Table. The Fictitious Rotation Rate Factor ε determines the offset from the originally specified rotation rate. The offset is an integer multiple of f where:

$$f = \varepsilon \Omega_{\max} \tag{2280}$$

$\Omega_{\max}$ is the maximum user-specified rotation rate in the blade row. The default value of ε is 0.001. You can set the value of ε in the Fictitious Rotation Rate Factor property of the Blade Rows manager node, but it is best to make the value small, so that the departure from the original value stays small.

Fluttering Motion of a Rotor

Consider the fluttering motion of a rotor. In the rotor frame of reference, an inter-blade phase angle σ, and a frequency ω describe the motion. For example,

$$h_i = h_0 e^{j\omega t} e^{(j\sigma\theta_i)/G_\theta} \tag{2281}$$

where G_θ is the blade-to-blade gap (in radians). Written in terms of nodal diameters, Eqn. (2281) becomes

$$h_i = h_0 e^{j\omega t} e^{jN\theta_i} \tag{2282}$$

where $N = \sigma/G_\theta$. This motion represents a traveling wave, which is like the motion of a rotor. In this case, we have an N nodal diameter wave moving with a rotational rate of

$$\Omega_{rel} = -\omega/N = -(\omega G_\theta)/\sigma \tag{2283}$$

This rotor rate is relative to the frame of reference of the fluttering rotor. Therefore, the rotation rate in the absolute reference frame is

$$\Omega_{abs} = \Omega_{rotor} - \omega/N \tag{2284}$$

To compute the frequencies and nodal diameters, another row is added to the mode table. For example, consider a stator with 36 blades and a rotor with 42 blades, rotating at 366 rad/s. Assume that the rotor blade is fluttering with an inter-blade phase angle of 180 ° with a frequency of 2000 rad/s. To construct the mode table:

1. Compute the nodal diameter N for the flutter motion:

$$N = \sigma/G_\theta = (180°)/((360°)/42) = 21$$

2. Compute the rate of disturbance Ω_{abs}:

$$\Omega_{abs} = \Omega_{rotor} - \omega/N = 366 \text{ rad/s} - (2000 \text{ rad/s})/21 = 270.76 \text{ rad/s}$$

3. Construct the mode table with three rows:

$$(B_0, B_1, B_2) = (21, 36, 42)$$

$$(\Omega_0, \Omega_1, \Omega_2) = (270.76, 0.0, 366.0) \text{ rad/s}$$

Harmonic Balance Bibliography

[406] Hall, Kenneth C., Thomas, Jeffrey P. and Clark, W. S. 2002. "Computation of Unsteady Nonlinear Flows in Cascades Using a Harmonic Balance Technique." *AIAA Journal* **40**:5 pp. 879-886.

[407] Giles, M. B. 1988 "Calculation of Unsteady Wake/Rotor Interaction", *J Propulsion and Power*, **4**:4 pp. 356–362

Harmonic Balance Field Functions Reference

The following primitive field functions include Harmonic Balance variants of density, entropy, pressure, temperature, and velocity. They, along with frequency, are activated by the Harmonic Balance Flow and Energy model. Modified diffusivity and turbulent viscosity are activated by one of the Harmonic Balance Spalart-Allmaras models.

These objects all share base properties, though the harmonic time functions also use the Harmonic Balance Time Field Function Properties.

[Value] at Time-level X	Solutions that are computed at each discrete time level (total number of time levels N = 2M + 1, where M is the number of modes)
[Value] Fourier Mode 0	The zeroth harmonic or the "mean" flow.
[Value] Fourier Mode X Imag	The Xth imaginary component of the Fourier coefficients.
[Value] Fourier Mode X Real	The Xth real component of the Fourier coefficients.
[Value] Harmonic Time Function	The flow quantity at a particular physical time (in seconds) specified by an entry in this field function object's Time property. To calculate the flow using only the mean mode, use the Mean Flow Reconstruction property. See Eqn. (2273).

[Value] Harmonic Time Phase Lagged Down Function	The flow quantity at a particular physical time (in seconds) specified by an entry in this field function object's Time property, which is phase-lagged against the direction of rotation. To calculate the flow using only the mean mode, use the Mean Flow Reconstruction property. See Eqn. (2273).
[Value] Harmonic Time Phase Lagged Up Function	The flow quantity at a particular physical time (in seconds) specified by an entry in this field function object's Time property, which is phase-lagged in the direction of rotation. To calculate the flow using only the mean mode, use the Mean Flow Reconstruction property. See Eqn. (2273).
[Value] Level X Phase Lagged Down	The Xth time level solution which is phase-lagged with respect to the user-specified phase angle, sigma, using -sigma for the phase-lagging.
[Value] Level X Phase Lagged Up	The Xth time level solution which is phase-lagged with respect to the user-specified phase angle, sigma, using +sigma for the phase-lagging.
Displacement at Time-level k	Displacement at the kth time level, $dx(t_k)$ in Eqn. (2277).
Fundamental Frequency	The greatest common divisor of all the unsteady frequencies present in the blade row.

Harmonic Balance Time Field Function Properties

Time	The physical time, in seconds, at which flow is calculated.
Inverse Distance Weight	Determines whether to use inverse distance weighting to interpolate node values. Activated: Interpolate node values using inverse distance weighting. Deactivated: Use nearest neighbor as node values. This is the default.
Mean Flow Reconstruction	When activated, calculation of the variable in question uses only terms for the mean mode and any mode that is artificially introduced because of the Consider All Blade Rows option. See Eqn. (2273). Activated: Calculates values using mean and introduced mode terms. Deactivated: Calculates values using all terms. This is the default.

Harmonic Balance Solution View Field Functions

The following primitive field functions are for visualization in harmonic balance solution views.

- HB: Density
- HB: Entropy
- HB: Grid Relative Velocity
- HB: Mach Number
- HB: Mass Flow Rate
- HB: Mass Imbalance
- HB: Modified Diffusivity
- HB: Pressure
- HB: Speed of Sound
- HB: Temperature
- HB: Total Enthalpy
- HB: Total Pressure
- HB: Total Temperature
- HB: Turbulent Viscosity
- HB: Turbulent Viscosity Ratio
- HB: Velocity

Modeling Solid Stress

This section covers the following:

- What is the Finite Volume Solid Stress model?
- Modeling Dynamic Problems
- Modeling Rayleigh Damping
- Modeling Large Displacements
- Setting up the Finite Volume Solid Stress model
- What is the FVA Stress Solver?
- Formulation

Also included is the following reference information:

- Shared conditions and values
- Field functions activated by this model

What Is the Finite Volume Solid Stress Model?

The Finite Volume Solid Stress model enables finite volume stress analysis in STAR-CCM+. The object has no properties of its own, but is subject to common menu items.



This section covers the following topics:

- Understanding solid stress modeling
- Comparing finite element analysis and finite volume analysis

Understanding Solid Stress Modeling

This section explains the following:

- What is solid mechanics?
- What is stress?
- How is mass conserved using solid stress?
- Understanding the differences between Static/Dynamic and Steady/Transient
- What constitutes a “good” mesh in finite volume stress analysis?
- Finite volume stress analysis best practices

What Is Solid Mechanics?

Solid mechanics is the study of the deformation, strain, and stresses in a solid material. The prediction of stresses in a material under load can be used to assess the potential for failure and provide information of the durability and estimated lifetime of the part.

What Is Stress?

Stress is a measure of the force divided by the area over which the force is applied as the area shrinks to zero, as presented in the formulation for stress.

How Is Mass Conserved Using Solid Stress?

Unlike the fluid equations, the solid stress is solved in a Lagrangian framework. That is, the mesh follows the material points in the material. This means that the mass in an particular cell will remain unchanged as it deforms because it physically represents the same material. This cell mass is determined as the product of the original density of the material and the original volume. Therefore the density in the deformed state will be the original mass divided by the current volume of the cell. In the "small displacement" FV stress solver, STAR-CCM+ does not actually move the solid mesh because it is assumed that the deformations are so small that these changes in the mass would be negligible. However, with the "large displacement" capability, STAR-CCM+ continually adjusts the solid vertices of the mesh according to the displacements.

Understanding the Differences between Static/Dynamic and Steady/Transient

Nomenclature used in solid mechanics is different than in fluid mechanics and there is not an exact one-to-one relationship between the solid mechanics "Static" and fluid mechanics "Steady" or solid mechanics "Dynamic" and fluid mechanics "Transient". In solid mechanics "Static" implies that the inertial terms are neglected but this does not imply that the solution is independent of the initial conditions. In fluid mechanics, "Steady" assumes that the inertial terms are neglected and the solution is independent of the initial conditions.

However "Static" elastic models do not depend on the history of the loading, the initial conditions, or the rate of loading so these materials can be run in the "steady" mode to arrive at the "static" solution. Fully coupled Fluid/Solid interactions must be conducted in "Transient" mode using real time since velocity of the solid boundary impacts the flow. However the "Steady" mode can be used to couple elastic solids to fluids but the

mechanical coupling is only one way – that is, the fluid pressure is applied to the solid boundary but the displacement of the solid does not affect the fluid.

In fluid mechanics, "Transient" implies that the solution depends on initial conditions and the time – the inertial terms are included. In solid mechanics a load can be transient but the solution can still be considered "Static". That is, the typical time scales to load or unload a part are very long compared to the time scales for flexural or sound waves to travel through the part. Therefore the inertial terms are not required. In truly "Dynamic" problems, the transients due to the low frequency flexural waves or high frequency sound waves are important and the inertial terms are important. Only the implicit transient solver is applicable.

What Constitutes a "Good" Mesh in Finite Volume Stress Analysis?

This section provides an overview of how to assess the quality of a mesh in finite volume analysis:

- Mesh resolution
- Cell quality
- Thin members under bending loads

Mesh Resolution

In many respects the requirements on finite volume meshes to resolve displacement and stresses accurately are similar to those for finite element meshes. Regions with stress concentrations need a proper resolution of the mesh. Fillets, the outer diameter of holes, parts that neck down, re-entrant corners, cracks, or welded surfaces where part of the two surfaces are not fully "wetted" will require some local refinement of the grid. For example, a fillet between two parts which are joined at 90 degrees will require at least 4 cells to span the quarter circle of the radius. A hole should have at least 16 cells to span the circumference. Near the root of a crack, the cells must be highly refined in a region about 2 crack lengths near the surface, with most of the refinement near the root and coarsening as the distance from the crack root increases. One should remember that a elastic model with a crack is "singular"--that is, the stresses are infinite at the tip of the crack. No matter how much one refines the mesh, the numerical stresses in the cell center will continue to increase. Here the principles of "Fracture Mechanics" must be applied to assess the severity of the loads on the crack--that is, whether or not the loads on the crack will have a tendency to cause catastrophic failure or slowly grow the crack under cyclic loads.

There are some differences in the requirements on the finite volume and finite element meshes. The finite volume mesh should have at least two cells that span between surfaces. Unlike a single finite element, at least two cells are needed to form a gradient across the span. At least three cells are recommended.

Cell Quality

The cell quality is the most important factor influencing the convergence and, to a lesser extent, the accuracy of the finite volume method. The finite volume stress analysis is more sensitive to the quality of the mesh than typical CFD. STAR-CCM+ offers many measures of the quality of the grid including skewness, cell-to-cell volume change, and least square gradient. The quality of the grid over all the cells is registered in one field function that is an indicator of bad cells. If less than 0.1% of the cells in the solid stress model are considered poor, then this mesh can be considered of high quality. If the number of cells of poor quality are greater than 1% then this grid is generally poor and more attention to the meshing is required. However, meshes where the number of cells are greater than 1% may be sufficiently good if the normalized residuals converge below 1e-3.

Thin Members under Bending Loads

If a shell is under bending loads then typically more finite volume cells will be required than in a finite elements, in the plane of the shell as well as across the thickness. In most situations with bending stress it is a good rule of thumb to have cells with aspect ratio near one. As the ratio of the distance between supports to the thickness of the shell relative becomes large, greater than 10, keeping the cell aspect ratio high will be extremely important else the members will behave much too stiffly. For this reason, the finite volume stress analysis is not recommended when the ratio of the distance between supports to the thickness becomes greater than 20. This is not to say that, with sufficient cell density between the supports, the bending displacements and stresses cannot be resolved.

This rule of thin members under bending applies to “open” structures. “Closed” structures such as a long pipe or a long conduit can support high bending loads down the axis with the thickness of the pipe or conduit quite small. In this case, the axial stresses do not change appreciably over the thickness of the pipe or conduit wall but rather across the breadth of its diameter. The rule would still apply if large bending occurred in the cross-section of the closed structure.

The user should also be aware that thin members of open structures in bending will converge slowly unless the cell quality is good.

Finite Volume Stress Analysis Best Practices

- FVA (Finite Volume Stress Analysis) has no FEA counterparts for structural elements such as rods, beams, membranes, and shells. These elements may be modeled as three-dimensional components in FVA but some care must be taken to resolve them. **For example a shell in bending should have at least 5 cells across the thickness and the cell aspect ratio should be near 1.** Shells and beams in bending will also generally converge very slowly. Therefore applications involving sheet metal in bending are not a good application for FVA.
- When generating grids for stress analysis, you should pay attention to stress risers such as reentrant corners, cracks and holes. This is especially the case if you are concerned about the stresses in these regions and need an accurate stress prediction. The grids need to be sufficiently fine to resolve the stresses. The default reference value for defining the surface curvature is 36 points/circle and this should be sufficient for most fillets. But remember that reentrant comers with no fillets may be singular and, no matter how fine the grid, the stresses in the cells adjacent to the singularity will continue to rise as the grid resolution increases. Before constructing the grid, it is a good practice to consider the boundary regions needed to have properly set boundary conditions for the solid stress – not just the solid temperature.
- Traction boundaries (or traction-free boundaries) must be carefully treated in FVA and often require special attention when there are bad cells adjacent to these boundaries. Generally it is a good idea to have a extrusion layer on these boundaries. **But the extrusion layer thickness need only be about the nominal size of the cells in that general area.** The idea here is to make orthogonal cells without making them too thin. The cell aspect ratio should not be much larger than 4 and should be around 1 if this extrusion is in a thin part in bending. For some geometries the extrusion layers must be necessarily small because there is only a thin layer of material between opposing boundaries. The cells, in this case, should be relatively small in the lateral direction so that cells which span the thickness do not have too large an aspect ratio. The solid/fluid and solid/solid interfaces are treated numerically as traction boundary conditions so rules for extrusions apply in this case as well.
- There must be at least 3 cells that span between opposing solid boundaries. **It is a very common error of new users to have plates that are only 1 to 2 cells thick.** If thin solid material lies between two other solid materials, then this layer must also be at least 2 cells thick. **If you have only 1 cell across the thickness then the solution may not converge!**

- The boundary conditions for a static analysis must be sufficient to constrain the part from rigid body rotation and translation. The symptom of ill-constrained models is usually seen as very large displacements or rotations and divergence in the residuals. **This is a very common error of new users in both FVA and FEA. The solution will likely not converge.**
- It is prudent to check the quality of the cells as indicated by the Solid Stress Bad Cell Indicator. At the initiation of a run, a summary report is provided indicating the overall quality of the solid region mesh for stress analysis. This report gives the number of cells which are considered poor and cell remediation is applied. **If the mesh has more than about 1% bad cells then be very suspicious of the grid and the results. Meshes with less than 0.1% bad cells are good quality.** Be mindful of the number of poor cells which have no good neighbor cells, because the remediation will be less benign in these areas.
- When thermal stresses are included it is important **not** to restrain a part in shear. The fixed boundary condition restrains all components of the displacement. This constraint will unrealistically prevent the part from thermally expanding and will generate unrealistic thermal stresses at the fixed boundary. It is as if one is constraining the part to rigidly adhere to an adjacent part which is infinitely stiff and does not thermally expand. It is acceptable to apply this constraint, but be suspicious of the stress results in the near vicinity of the constraint--especially if the temperature is well above the zero thermal expansion reference temperature. It is better to restrain the part with a fixed normal condition which allows the part to slide freely along this surface. This will prevent translation or rotation but it will generally allow the part to thermally expand realistically. Using symmetry boundaries is also useful for constraining the part. It is also useful to use a displacement boundary defined by a field function of the local temperature. You can assume that the part is freely expanding with the temperature. It is also important not to over-constrain a part. A part prevented from expanding globally can generate unrealistically high compression stress. For example, you could heat up a pipe and fix both ends. This may produce unrealistically high compression stress since in reality the physical parts that would produce that constraint will also expand. Remember that attached parts can also expand, bolts and washers can slip, male and female threads can slip. Often it is best to simply restrain the part to prevent translation and rotation but still allow the part to expand freely.
- Avoid using tetrahedral solid meshes for stress (sparse use is acceptable).
- Solid Stress Divergence is caused by:

a. Static models which have not been properly constrained from rigid body motion

b. Bad quality cells not identified by the automatic cell quality check

c. Almost incompressible materials (elastic materials with Poisson ratio of near 0.5)

d. Crack roots near relatively poor cells

If the problems diverge immediately then the most likely cause is most likely [a]. Remember it requires that all rigid body motions be constrained in static models. If reason [b] is causing this problem then normally you will find that the residuals will go down and then start to diverge. **This is a very common problem** in FVA since the convergence is **highly dependent on the grid quality** and depends on **how well STAR-CCM+ identifies the bad cells** that are in the model. For most CFD users the first inclination is to under-relax the solid momentum, **but this is not recommended**. **It will also generally diverge with under-relaxation but it will only diverge slower.** The best strategy for removing the divergence problem caused by unidentified bad cells is to activate solver procedures to cap the solid stress effective stress. This is under the expert properties of the FVA Stress Solver. You must define a limiting effective stress that would be considered much larger than the expected peak stress due to the load. The default peak stress is 1 GPA. A good rule of thumb would be to use a value less than 10% of the Young's modulus. Divergence can also occur if the Poisson ratio is too close to 0.5, and you should not use a value greater than 0.45. Finally, if cracks are in the model and the cracks have poor cells, then divergence might be expected. The areas near the cracks need to be meshed properly.

- Slow convergence is caused by:

  a. Bad quality cells not identified by the automatic cell quality check

  b. SOR and simultaneous solution of flow, thermal and stress in static problems

  c. Thin members under relatively large bending

  Bad quality cells may slow or stop convergence. If you see the residual oscillating up and down then generally the accelerated convergence strategy will not work. Selecting the stress cap in the expect mode of the solver may be effective in improving the convergence rate.

  The Successive Over Relaxation SOR under the FVA Stress Solver will be less effective under conditions when solving for the temperature and the displacement simultaneous in steady problems. You may choose to turn off SOR in this case. Often it is best in these thermal stress

problems to solve the thermal problem with the solid stress frozen. Once the energy has converged, then turn on the stress solver and freeze the energy solver.

Problems with thin members in bending will be slow to converge--especially when the distance between the supports is greater than about 10 times the thickness of the member. It would be advisable to monitor the displacement near the point of maximum bending displacement. If it is slowly tending to change then you should continue with iterations. This rule applies to "open" thin members where the in-plane stress varies dramatically through the thickness. "Closed" members, like a thin pipe, are an exception to the rule since the bending stresses across the pipe thickness do not change dramatically, and the pipe is of course much stiffer than a pipe with a slot cut down the axis.

- STAR-CCM+ assumes that displacements, rotations and strains are small. The vertices of the solid do not change with time--only the displacements of the cells and face centers do. For large displacements, vertices in the solid mesh are moved based on interpolation of the solid stress displacement and its gradient solution at cell centers. In general, strains should be less than 10% and the rotations should not be more than 10 degrees, to satisfy the small deformation assumption in the FV formulation.
- When you create a solid/solid interface and the solids are different materials, the FV Stress treats this as a material discontinuity and preserves the discontinuity in the strains across this material interface. These materials are treated as if they were welded at the interface (continuous displacements). Contact is not supported. You can imitate contact in static problems by constraining the part, running the solution and determining where the part goes into tension. In the regions the boundary constraints can be removed by setting the area to "free" and then running again.
- When you create a solid/solid interface or a solid/fluid interface then it is prudent to make sure all faces from the original surface are matched. It is possible that some faces are unmatched to the opposing surface and they will be left in the original boundary surface. The default boundary conditions will be adiabatic free surface in the solid and an adiabatic wall in the solid. If these faces are isolated then they will appear much like adiabatic cracks and these cracks can produce unusually high stresses at the solid/solid interface.
- In a transient problem coupled to the fluid, the change in the displacement is introduced as a swept flux into the fluid but the grid does not change. **In steady-state problems the displacement of the solid does not change the fluid at all.** In a steady-state problem, the

pressure and the heat from the fluid affects the solid and the temperature of the solid affects the fluid temperature, but the displacement of the solid has no impact on the fluid.

For more information, see other parts of this section, as well as the tutorial on finite volume stress.

Comparing Finite Element Analysis and Finite Volume Analysis

This section is intended for FEA users who may be interested in using the Finite Volume Analysis (FVA) for stress analysis. For those CFD users who are unfamiliar with FEA or stress analysis in general, this section represents a brief "crash course" in both FEA and FVA.

Those users who are familiar with FEA methods will find some similarities and differences in the Finite Volume stress analysis (FVA) methods. For example, the necessity of properly constraining the model to remove any rigid body translations and rotations in static analysis is applicable to both FEA and FVA. The current FVA formulation is comparable to ***continuum*** formulations in FEA. But unlike FEA, there is no counterpart to structural constructs such as rods, beams, membranes and shells in the FVA formulation. This document will be limited to comparisons to the FEA continuum elements, as follows:

- Degrees of Freedom (DOF)
- Elements versus cells
- Mechanical loads
- Thin shells under bending loads
- Thermal loads
- Constraints
- Material models
- Geometric nonlinearities
- Contact
- Fluid structure interaction
- Solver technology
- Cell quality
- Details of the storage requirements

Degrees of Freedom (DOF)

What is quite different between FEA and FVA is the location where the displacements, strains and stress are represented in discrete form. In FEA the displacements are evaluated at the nodes while the stresses and strains are evaluated at the Gauss points. In FVA the displacements and strains are known at the cell centers. The strain is interpolated to the face centers and stress is computed on the face centers. For post-processing the stress is evaluated at the cell center. In FEA the total number of degrees of freedom (DOF) corresponds to the number of nodes multiplied by the DOF per node. In a standard displacement FEA three-dimensional continuum formulation, the number of DOF per node is 3. In the FVA three-dimensional formulation, the total number of DOF is 3 times the number of cells. So the total number of DOF in a three-dimensional continuum is about equivalent between FEA and FVA when roughly the number of FEA nodes is the same as the number of FVA cells.

Because FEA computes displacements at the nodes, an estimate of the deformation gradient can be established using only one element and the shape functions on that element. In FVA the deformation gradients are determined by least square formulation and use the neighbor cell displacements to arrive at the gradient at the cell center. The important distinction is that **FVA cannot represent a structure such as a plate with only one cell across the thickness**. There must be at least two cells, and to be second-order accurate there needs to be three cells. This is an important consideration when building the FVA mesh.

Elements versus Cells

Most analysts familiar with continuum FEA understand that there are a number of different choices of FEA elements whose accuracy depends on the nature of the problem and the material properties. First- and second-order “brick” elements with reduced integration elements are appropriate for problems with bending but may require “hour glass” control. First- and second-order brick elements with full integration elements will be overly stiff in bending but do not have hour glass problems. Hybrid formulations are introduced for incompressible or almost incompressible materials like rubber. The first-order tetrahedral element is generally not used in stress analysis because this constant strain element behaves overly stiff in bending. But all first-order elements (including the tetrahedral) will converge to the grid independent solution of a compressible material with no reentrant corners (without singularities such as cracks). So the difference in the elements is not merely defined by their shape or order but how ***well they represent the solution when the grid is still relatively coarse*** and whether or not the element has unrealistic distortion modes such as hour glassing.

The current FVA stress analysis offers not only tetrahedrals, wedges, and bricks but also general polyhedral. All are of comparable order (between first and second order for general automatic mesh). The difference in FV cells is strictly due to the geometry and is primarily related to the quality of the shape of the cell and how many neighbors are available to compute the deformation gradients (strains). As in CFD, the tetrahedral is the poorest cell because it has few neighbors to compute the gradient. Unlike the FEA element, the traction boundary condition requires the extrapolation of the displacement to the boundary face center. Therefore the gradients at cells with faces adjacent to traction boundaries are not as reliable unless they have other faces to estimate the gradient well. Since the tetrahedrals have only four neighbors, one or more traction boundaries on a tetrahedral cells severely degrades the accuracy. ***Therefore tetrahedral cells are not recommended – especially any tetrahedral cells with traction loads applied at boundary faces.*** The extrapolation of displacement at the traction boundary affects all cells and it is also important that cells adjacent to traction boundaries be as ***orthogonal*** as possible. Like FEA continuum elements, high aspect ratio cells are less accurate and have poor stability compared to the low aspect ratio cells. Extrusion of meshes at the surface to obtain "orthogonal" cell is beneficial. However, unlike CFD which may require a thin cell to model a boundary layer, no such small layer is needed in stress analysis. ***So an orthogonal extrusion is good but thin extrusions are not desirable***. If regions of the part are very thin, then it may be necessary to reduce the lateral size of the cell to maintain a good aspect ratio and still apply a few cells across the thickness. This will be especially true if this thin part is in bending because it may be excessively stiff if the grid is coarse. Just like some continuum elements in FEA, coarse FVA cells with large aspect ratios may be overly stiff under bending.

Mechanical Loads

In FEA, mechanical loads on continuum elements can be applied as discrete nodal force loads or as a distributed traction stress. In FEA, the given distributed traction is multiplied by the nodal shape function and integrated over the faces of the three-dimensional continuum model to arrive at a nodal load. In FVA, a distributed traction stress can be applied. The traction at the center of the cell face multiplied by the cell face area is the cell face force. In this sense, the FVA concepts for force are more straightforward. In STAR-CCM+, a force is applied to a boundary region via an applied traction. If the applied force is constant over the surface, then the required constant traction is simply the force divided by the area of the face. **Just like FEA, continuum model loads cannot be expressed via moments in the FVA model**.

In FEA models, the free surface are those faces which have no loads applied to them (zero traction). A traction free loading is the "natural" boundary condition and no special considerations are needed to maintain the zero

traction condition. In contrast, FVA has no counterpart of "natural" boundary condition so there is no particular distinction between zero traction and nonzero traction in terms of the formulation. Both are traction conditions with one simply having a traction of zero. For this reason the FVA model must identify the boundary regions of the traction boundary condition whether it is zero or not. In STAR-CCM+ the default condition is the free boundary condition.

The traction boundary condition is particularly difficult for FVA because an estimate of the displacement on the boundary face is still needed to compute the gradient (strain tensor) at the center of the adjacent cell. This gradient is needed to define properly the complete stress tensor on the face. Remember a zero traction boundary condition does not imply zero stress tensor on the face, much less a zero strain tensor. The boundary face displacement must be extrapolated using the gradient at the cell center, but the boundary face center is used to estimate the gradient. This iterative process to define the displacement may not converge – especially when the cell in question has few neighbors or it is not orthogonal or is poorly shaped. The quality of a cell depends not only on its geometry, but the number of neighbors and orientation of its neighbors, and the boundary condition that may be applied to a particular cell. The cell quality metric makes all these considerations in attempting to measure the quality of the cell.

Thin Shells under Bending Loads

As stated previously the FVA has no counterpart to shell elements in FEA. Often in FEA continua three-dimensional continuum elements are applied to solid parts with relatively thin members under bending loads. In this situation it is well known that linear tetrahedrals behave more stiff than in reality unless the grid is extremely fine. Similarly linear brick elements with full integration exhibit this same behavior. Usually under these conditions an under-integrated linear brick element is a reasonably good choice with at least two elements through the thickness of the thin member. Under-integration quadratic bricks also perform reasonably well. **Under the same conditions, the FVA brick will require a denser grid than the under-integrated linear brick element to give comparable stiffness behavior.** Generally for thin members where the distance between supports relative to the thickness of the plates is greater than ~10, the FVA cell stiff behavior will be more pronounced. To give acceptable stiffness behavior, the FVA cell aspect ratio needs to be near a value of one for the cell. Furthermore the iterative FVA solver will generally require more iterations to converge with thin members under bending. Under such loads, convergence is greatly accelerated by the SOR technique, but it is particularly important for the FVA mesh to be free of poor cells if the SOR technique is to be effective.

Thermal Loads

The application of thermal loads in FVA is exactly the same as in FEA. Proper selection of constraints for problems considering the thermal loads are needed for both types of analysis in order not to produce unrealistic thermal loads. So for instance, applying a fixed displacement boundary condition may produce excessive thermal stress in the part at the fixed boundary. This is because it does not necessarily consider the fact that the constraint may be unrealistic since it assumes that the adjacent part does not also expand to the thermal load. Therefore the considerations for constraining parts in thermal loads are identical in FVA and FEA.

Constraints

FEA offers the capability to constrain certain DOF on a node while allowing DOF on the same node to be unrestrained. This complete flexibility in specifying the constraint is not found in the FVA solver. The FVA constraints are in two forms – either all displacement components are constrained on a face or the displacement is constrained only in the normal direction while allowing the tangential components to be unrestrained. **There is also no FVA counterpart to FEA "multiple point constraints"**. Multiple point constraints allow a constraint of a DOF based on the DOF on other nodes which may be remote. Unless relatively stiff continuum cells are actually linking the nodes, the FVA solver cannot treat this restraint. In this respect the FVA is not unlike an **explicit** FEA solver.

Material Models

FVA and FEA can, in principle, accept most any material model linear or nonlinear. The FEA methodology is formulated to accept most any type of material model. This is accomplished by separating the functions of the solver from the material model. The function of the material model is twofold:

1. Compute the stresses given the current increment of the strain, temperature, and other state variables which describe the internal degrees of freedom such as plastic strain
2. Return a tangent modulus (the change in the stress relative to the change in the strain) relative to the current state of stress of the model

In this sense, the FEA model and the FVA model are no different. FEA evaluates the material model at Gauss points while the FVA model must evaluate the material model at the center of internal faces as well as centers of faces where the displacement is applied. In addition FVA evaluates the material on opposite face centers at bonded material interfaces.

FEA of course has a long history and its material model library is large. The FVA library in STAR-CCM+ includes only the linear elastic material but nothing fundamental is preventing the introduction of new material models into this library in the future.

Geometric Nonlinearities

Some FEA codes allow for geometric nonlinearities which arise from the definition of the strain when the strain or the local rotation of the material is large. In FEA the geometric nonlinearity is usually treated with either Total or Updated Lagrangian. Total Lagrangian methodology operates on the original configuration and thus the grid remains fixed. It is generally restricted to elastic materials and often applied to hyperelastic materials such as rubber. Updated Lagrangian is more generally applicable to all types of material (such as materials which undergo plastic deformation). It requires that the mesh be moved with the deformation in increments. Consequently, the Updated Lagrangian models may become excessively distorted under large deformations and may require remeshing and interpolation.

The large displacements model in STAR-CCM+ considers ***large displacements but is restricted to small strains and deformations***. The Updated Lagrangian formulation used in the FV Stress model allows for large solid mesh motion and is especially applicable to FSI.

Contact

Currently STAR-CCM+ has no contact model.

Fluid Structure Interaction

As currently implemented in STAR-CCM+, the modeling of fluid-structure interaction (FSI) can handle both small and large displacements. For small displacements, (displacements must be small compared to the typical size of the cell), an approximation may be used to describe the boundary motion into or away from the fluid, even though the mesh is fixed. In this case a "swept" flux is computed to estimate the apparent flux of fluid mass on the fixed boundary. For large displacements, (deformation and strain must be small but displacement can be large), "swept" flux is computed based on actual motion of the solid boundary. All other interactions, such as the heat flux, are treated as normal. As a result "thermal stress" problems, where the temperature of the solid is greatly affected by the efficiency of the fluid heat transfer coefficients, can be readily solved with this framework. Flutter problems can also be solved using large displacements, so long as the strain is not too large (less than 10%). Structural problems involving moderate deformation of the solid can result in geometric nonlinearities, which are modeled with the large displacement capability in STAR-CCM+.

It is important to remember that large displacements of the solid can occur even if the solid is perfectly rigid. These large motions imply that the fluid mesh will undergo large motions requiring either grid morphing, cell layer insertion/deletion, remeshing/interpolation or all the above. These mesh motion problems will be ***present in all problems with solid motion regardless if the solid is treated via some FEA code like ABAQUS or the solid motion is treated within STAR-CCM+.*** The mesh motion in the fluid is particularly difficult to describe when two solid bodies come in contact – such as a ball-socket stop valve.

That said, there are also many important benefits for treating solid stress in FVA for FSI problems. Part of this is due to the fact that the interpolation of data based to/from the solid and fluid is completely automated in STAR-CCM+. There is no need for a additional code such as MpCCI to perform that data passing and interpolation. But the most important reason for treating the solid and fluid in STAR-CCM+ is the commonality of data structures and the solver technology. The cost per iteration per cell for the solid is much the same as for the fluid, and therefore data passed between the codes can be passed more frequently than could be expected when passing data between a CFD and FEA code. More about this in the next section.

Solver Technology

The major discriminating factor between FEA and FVA is how the discrete equations are actually solved. Most implicit FEA solvers use a Newton Method which "linearizes" the problem and a "direct" solver is used to determine the solution of a linearized problem. The direct solver must perform an "elimination" of the "stiffness" matrix and this elimination process constitutes the major CPU and memory cost. Modern FEA solvers found now in most commercial codes use efficient sparse matrix, multifrontal solvers. These modern codes automatically order the node numbering to reduce the CPU time and memory required to eliminate the matrix. **The direct solvers are also robust and are less sensitive to the quality of the grid.** But none the less, typically the CPU and memory costs are almost quadratic functions of the number of nodes. This means that much of the memory needed for large DOF problems must be stored on the disk rather than RAM. **The eliminated matrix requires much more memory than the matrix itself because of "fill-in".**

Most FEA commercial codes offer iterative solvers – usually based on some form of a conjugate gradient solver employing a pre-conditioner using incomplete elimination of the matrix. Iterative solvers require considerably less storage requirements than the direct solver, but the iterative solvers are not as robust and are sensitive to grid quality. The storage requirements are

basically linear with the number of nodes, but the CPU time of the iterative solvers is not linear with the number of DOF even when the mesh quality is good.

Finally another choice offered by many commercial codes is the explicit FEA solver. **Explicit solvers do not require the stiffness matrix to be evaluated or stored**. Therefore the storage requirements are related primarily to the number of elements and the complexity of the material model. But the explicit model is only stable when the time-step is less than the time it takes a shear or acoustic wave to travel across one cell. Thus generally many steps are required to reach a static equilibrium.

Like the implicit FEA schemes, the FVA stress analysis is implicit and has no restriction on the time-step size to remain stable. The FVA employs an iterative solver which is much different than the iterative solvers used in FEA but more like the segregated solvers found in CFD. The FVA solver uses a very rough estimate of the stiffness matrix for the "preconditioner" or linearization. The preconditioner is based solely on the nearest neighbors of each cell even though the actual "stencil" of cells which directly influence a cell can be much larger. In addition, since the FVA uses a segregated solver, the matrix (generally referred to as coefficients in CFD) needs only store terms of the stiffness matrix which involve one component. For more details see the last section, but it can be shown that the FVA ***iterative*** solver will be require ***about 10 times less memory*** to store the required matrix than the typical FEA ***iterative*** solvers with the same number of DOF. And remember that iterative FEA solvers require ***much less memory*** than a direct FEA solver which is typically used in FEA. So the memory requirements for the FVA implicit iterative solver lie some where between an FEA explicit solver and an FEA implicit, iterative solver. The memory and CPU required by the FVA stress solver is essentially linear with the number of DOF provided the grid quality is good.

Unlike the iterative solvers found in FEA, the FVA iterative solver updates any material nonlinearity on the fly at each iteration. The FEA linearizes the solution and then iterates to obtain the solution of this linear solution. After obtaining the linear solution the FEA takes another Newton's step, computes a new linearization, and then uses the iterative solver to converge to the next step. But the FVA iterative solver updates the linearization matrix (or more precisely the precondition matrix) each iteration so that when it converges all the nonlinearities are taken into account. In this sense, the linearized iterative solver and the Newton's iteration are rolled into one outer iteration.

Unlike the FEA pre-conditioners, the FVA precondition matrix is not easily inverted. So like the CFD segregated solvers, the FVA employs an "inner" iterative solver (the AMG solver). This inner solver need not be converged to machine accuracy but only sufficiently converged to provide a

reasonably "smooth" solution. The outer iteration then updates the solution base on the inner iteration solution and updates the residual including any changes in the material nonlinearity. The inner iteration solver's default tolerance is typically smaller than the CFD momentum tolerance (0.03 as opposed to 0.1) Unlike the CFD solver, under-relaxation of the FVA stress solver is generally not recommended since it tends only to slow convergence and not add to the stability.

In addition, since the FVA iterative stress analysis solver has so much in common with the CFD segregated solvers, the stress analysis memory requirements will be comparable to CFD models of a similar number of cells. In addition since the solver technology is common to both the CFD solver and the stress analysis, the scalability on parallel computing is comparable.

Thus there is a trade-off. The FVA elements likely require more elements to describe properly bending stress in thin members but the costs per cell is much, much less per typical cost per element in FEA. ***FVA offers the potential for analyzing much larger models than in FEA so the overall computer costs relative to the accuracy can be much lower than FEA.***

Cell Quality

Like FEA iterative solvers, the FVA solver stability and convergence will be sensitive to the quality of the cells. But unlike current FEA solvers, there is some leniency in how the FVA solver treats bad cells to avoid polluting the solution and causing the iterative solver to diverge. There is a principle in solid mechanics which says that stresses at points are independent of the actual load distribution provided the load is applied at some distance. Only the total force and moment applied has an effect on the stress distribution at some distant point removed from the load. Therefore there is a strong ground to support that even though some poor cells may produce locally incorrect results, these poor cells will not affect the overall behavior in the part and will not affect the local stress distribution at points distant from the poor quality cells. Therefore the FVA solver includes a unique strategy to identify poor quality cells and then acts on these cells to minimize their impact on the surrounding cells and on the convergence of the iterative solver.

Details of the Storage Requirements

Consider a cube made up of elements/cells which are also cubes. Let us say that there are N linear elements or cells which make up the cube. Then the requirements to store the stiffness matrix (or the eliminated stiffness matrix for the direct solver) are shown in the following table. This Direct Solver estimate is based on a banded matrix solver but the even the more modern direct solvers require storage. The reordering schemes are likely better and

so we have actually used an estimate which is only slightly non-linear. The FEA iterative solver matrix storage is related to the number of DOF on any row of the matrix. But these schemes use an incomplete elimination of the entire stiffness matrix. In addition to the matrix there are other requirements, for example storing locations, displacements, stress and strain tensors which are basically linear with the number of elements.

Formulation	Matrix storage
FEA implicit with Direct Solver	$14N^{5/3}(14N^{1.5})$
FEA implicit with Iterative Solver (Incomplete elimination)	$121N$
FVA implicit with Segregated Iterative Solver	$6N$

Stiffness Matrix Storage Estimates

Modeling Dynamic Problems

Dynamics problems require the addition of the acceleration (inertial) terms into the equations of motion, Eqn. (2307). STAR-CCM+ offers two different implicit time integration methods, which are the 1st order Backward Euler method and the 2nd order Newmark method common to FEA. The Backward Euler method introduces more numerical damping than the Newmark method for a given step size, whereas the Newmark method offers no numerical damping. As a result, for time dependent quasi-static models the Backward Euler is recommended since it will quickly damp any initial transients, while the Newmark method is recommended for dynamic problems due to its accuracy and absence of damping.

In transient problems the effective stiffness matrix is modified by the inertial terms, for example for the Backward Euler method the effective stiffness matrix becomes:

$$\mathbf{K}' = \mathbf{K} + \frac{\mathbf{M}}{\Delta t^2} \tag{2285}$$

In this case, even if the structure is completely free, the effective matrix is not singular, since the mass matrix $\mathbf{M}$ is always symmetric positive definite. In the finite volume method the mass matrix is always a diagonal matrix. When at least one of the six rigid body degrees of freedom (DOF) are unconstrained, the stiffness matrix is singular. The addition of the mass matrix term can lead to better conditioning of the effective matrix ***provided the step size is sufficiently small***. In this case, "sufficiently small" means the

time step is equal to or smaller than the period of the fundamental mode of vibration of the structure. In fluid structure interaction (FSI) problems, the step size to provide a non-singular effective stiffness matrix *may* be smaller than that which would be needed to resolve transients in the fluid when the fundamental natural frequency of the structure is large compared to the fluid frequencies (such as vortex shedding frequencies or buoyancy induced frequencies in free surface flows).

Modeling Rayleigh Damping

This section describes the Rayleigh Damping model, which can be used to impose viscous damping on a solid structure. The topics covered are as follows:

- What is the Rayleigh Damping model?
- Selecting the Rayleigh Damping model
- Setting Rayleigh damping parameters

The methodology for rayleigh damping is provided in the formulation section.

What Is the Rayleigh Damping Model?

Rayleigh damping can be used to mimic the actual structural response to dynamic excitation by using approximate damping coefficients to model the viscous damping effects in the structure. These damping coefficients can be estimated based on the dominant modes in the structure, as established in a modal analysis of the structure. The Rayleigh Damping model can also be used to include some numerical damping to damp unwanted numerical oscillations.

The effects of mass and stiffness on the viscous damping are approximated in the Rayleigh Damping model via two damping coefficients, where one damps the low frequency and the other damps higher frequency. The Rayleigh Damping model assumes that the viscous damping matrix for a system is proportional to the mass and stiffness matrices as described in the Rayleigh damping methodology.

The object has properties, which are used to set the damping coefficients, and is subject to common menu items.



Selecting the Rayleigh Damping Model

The Rayleigh Damping model can be selected in the *Optional Physics Models* box in the *Physics Model Selection* dialog after the Finite Volume Solid Stress model has been activated.



Setting Rayleigh Damping Parameters

The Rayleigh damping parameters are used to damp low frequency and high frequency oscillations in the solid structure. Since the finite volume method does not compute the stiffness matrix, an equivalent law is used as described in the Rayleigh damping methodology.

The Rayleigh damping parameters are properties of the Continua > Solid Continuum > Models > Rayleigh Damping node:

- Low frequency damping is specified using the Mass Damping Coefficient, which is $1/\tau_M$ in Eqn. (2315). When this value is set to zero, that is $\tau_M \rightarrow \infty$, there will be no low frequency damping. This coefficient can be used to damp both shear and dilation waves.

- High frequency damping is specified using the Stiffness damping coefficient, which is τ_K in Eqn. (2315). When this value is set to zero there will be no high frequency damping. It can be used to remove high frequency oscillations induced by pressure waves in the fluid.

Setting both of these coefficients to zero will cause no damping to be included in the model, and is equivalent to running the simulation without the Rayleigh Damping model selected.

Selecting appropriate values for the Rayleigh damping parameters depends on the availability of modal data for the structure. Where modal data is available, the Rayleigh parameters can be calculated based on the most significant modes of oscillation, as described by Chowdhury and Dasgupta [408].

Alternatively, you can estimate the Rayleigh damping parameters using the following equation:

$$\zeta = \frac{1}{2\tau_m\omega} + \frac{\tau_k\omega}{2} \tag{2286}$$

where ζ is the damping coefficient, $1/\tau_M$ is the mass damping coefficient, τ_K is the stiffness damping coefficient and ω is the angular frequency.

Rayleigh Damping Bibliography

[408] Chowdhury, I. and Dasgupta, S. P. 2003. "Computation of Rayleigh Damping Coefficients for Large Systems", *The Electronic Journal of Geotechnical Engineering*, **8**, 8C.

Modeling Large Displacements

This section describes the large displacement model, which can be used in fluid-structure interaction (FSI) simulations. The following topics are covered:

- What is the Large Displacement model?
- Enabling the Large Displacement model

What Is the Large Displacement Model?

The Large Displacement model, which is enabled via the Solid Displacement motion, allows a solid mesh to be geometrically updated based on the solid stress displacement. This model can be used to allow for the large displacements associated with rigid body motions, such as in 6 DOF cases, and is particularly useful for FSI simulations, where it is important that the fluid mesh conforms to the large motions of the solid structure. In FEA this capability is often referred to as "Non-Linear Geometry".

It is well known that structures can stiffen in tension, for example the strings on a violin become stiffer when they are tightened. This is often known as "Geometric stiffening". Similarly, compression tends to decrease the stiffness, which can lead to buckling, for example. When using the Large Displacement model, the formulation requires strains that may give rise to geometric stiffening to be small (less than 10%), however, displacements in the solid can be large. In cases where the loading has a significant effect on the apparent stiffness of the structure, the Large Displacement FV solid stress may have difficulty reaching convergence and the numerical solution may show sudden jumps in the computed stresses. In these cases a smaller solid displacement under-relaxation factor should be used. If it is necessary to set the under-relaxation value to below 0.01 for a severe case, the results should be viewed very cautiously. A similar approach should be adopted for structures that are buckling due to a compressive load.

The FV large displacement formulation is completely analogous to the FEA formulation termed "Updated Lagrangian", in that the geometry (vertex coordinates) is updated after each iteration in a time step for a transient analysis, and after each iteration in a steady-state static analysis.

When the Solid Displacement motion is selected for the region, a Solid Displacement node will be added to the region values. The object has a single property, which is described below, and is subject to common menu actions.



Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the solid mesh is geometrically updated based on the computed solid stress displacement and gradients. The default value is `0.8`.

Enabling the Large Displacement Model

The Large Displacement model is enabled by specifying a Solid Displacement motion on the solid region. To do this, follow the steps described in What Is the Solid Displacement Motion?. The Solid Displacement motion can be applied to a region for which the Finite Volume Solid Stress continuum model has been activated.

Setting Up the Finite Volume Solid Stress Model

The Finite Volume Solid Stress model can be used in a simulation either with solid regions only or in conjunction with fluid regions. It can be run in either implicit steady (static) mode, or implicit unsteady (dynamic) mode. The static mode neglects the inertial terms (acceleration), while still allowing for the boundary conditions, including the solid/fluid interface conditions, to be functions or time.

This section explains the following:

- Selecting the Finite Volume Solid Stress model
- Applying a body force
- Working with interfaces in solid stress
- Using the summary report for bad cells

Other aspects of a simulation with solid stress are described in separate sections:

- Solvers
- Initial conditions and boundary conditions

Selecting the Finite Volume Solid Stress Model

The Finite Volume Solid Stress model is available for selection in the *Optional Physics Models* group after you enable the Solid model during material model selection.



Once you select the Finite Volume Solid Stress model, the Linear Isotropic Elastic material model is selected automatically, provided the automatic selection of recommended models is activated.



Selecting the cell quality remediation model is highly recommended for most all solid stress analysis. The solid stress convergence is particularly sensitive to the quality of the grid and the cell quality remediation is ideal for detecting and correcting for poor cells in the model. Generally, the effects of the remediation are only local to the poor cells – as can be expected from Saint-Venant's principle.

Applying a Body Force

Body forces are forces that act throughout the volume of a body, such as weight due to gravity. They can be applied to solid regions in order to model forces acting throughout the region, and are specified in STAR-CCM+ in the units of N/m^3.

Body forces can be enabled on a region by selecting the Regions > Solid Region > Physics Conditions > Solid Source Option node and setting the Method to Specified.



This adds two nodes to the Physics Values for the region:

- Solid Stress Source

  The Solid Stress Source option is used to specify body forces acting on the solid region. It is entered as a vector profile. To specify several body forces acting on a region, you can enter the value of the resultant body force.

- Solid Stress Source Force Derivative

  The Solid Stress Source Force Derivative can be used to stabilize the solution.

Working with Interfaces in Solid Stress

It is possible to create interfaces between a solid and a fluid region, or between two solid regions. The procedures to create these interfaces are identical to that of the Fluid/Fluid interface. It is assumed that two regions share a boundary surface.

One way in which interfaces can be used in solid stress is to model fluid structure interaction (FSI). This involves creating an interface between the fluid and solid domains across which STAR-CCM+ will pass field data, such as pressures, heat fluxes, temperatures and displacements. Coupling between the fluid and solid domains is carried out automatically by STAR-CCM+ when a fluid-solid interface has been set up, and the finite volume solid stress model has been enabled for the solid region. Depending on the physics models you have selected, one or both of the following coupling controls will be available under the FSI interface node:

- FSI Coupling Method

  This lets you control whether data will be exchanged across the FSI interface after each time-step, or after each inner iteration.

- FSI Under-Relaxation

  This lets you specify the extent to which the mesh will be morphed after each time-step for an FSI case using large displacements.

At an interface between the fluid and solid domains it is also possible to apply an additional load that will act on the solid domain.

Applying an Interface Load

An additional load can be applied at a fluid-solid interface that will act on the solid domain. It can be used to represent an additional load on the solid structure that may not have been modeled explicitly. For example, for a boat on water where all surfaces of the boat are interfaces with either air or water, an additional load may be applied to a surface at the rear of the boat to model the force from a motor, even though the motor itself is not present in the model. The load can be a traction, pressure or force and is specified via

the Interfaces manager node. You can select which type of load to apply at the interface by selecting from the options for the Interfaces > In-place 1 > Physics Values > Interface Load node.



The value of the interface load is specified by selecting the corresponding node under the Interfaces > In-place 1 > Physics Values node, that is, either Traction, Force or SolidStressPressure. Values for these nodes are specified in the same way as for prescribed tractions, forces or pressures on a solid boundary. Note that you can specify load on direct or indirect interfaces.



Using the Summary Report for Bad Cells

To activate the report for solid stress bad-cell statistics, select the Cell Remediation Model. This will add a field function named Solid Stress Bad Cell Indicator, which is an extension of the Bad Cell Indicator. However, different indices are used to indicate the severity of poor cells. The indices appear as follows:

```
GOOD                    = 0,
BAD                     = 1,
BAD_NO_GOOD_NEIGHBOR    = 2,
GOOD_NO_GOOD_NEIGHBOR   = 3,
BAD_STRESS              = 4,
```

When you initialize the solution or begin a run, a cell quality report will include the following information:

- Total number of cells in each region
- Total number of bad cells and the percentage
- Number of bad cells without good neighboring cells, and the percentage
- Number of good cells without good neighboring cells, and the percentage

Bad cells which are identified by high stress (an user input value) are not reported as they are based on the stress solution.

What Is the FVA Stress Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the Finite Volume Solid Stress model in all the continua for which the model is activated.



This section also explains the effect of other solver parameters on solid stress modeling.

FVA Stress Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value (1.0) means that no under-relaxation is used.
Least Square Quality Tolerance	Deformation gradients are computed using a "Least Square Method" and a metric is employed to determine if the cell quality is sufficient to compute the local gradient accurately. The Least Square Quality Tolerance defines the metric value below which the cell is pronounced poor and remedial actions are taken on this cell. Smaller values will allow more cells to pass – larger values will mark more cells as poor. The default value is 0.1. Some models will run satisfactorily with tolerances as low as 0.01. The Least Square Quality Tolerance is perhaps the single most important parameter affecting the convergence of the iterative FVA solver.
Successive Over-Relaxation	Controls whether or not to accelerate convergence. Ticked: The solver will accelerate convergence. Cleared: The solver will not accelerate convergence.

Expert Properties

The properties in the *Expert* category are designed to be used in a temporary debug situation, at the expense of simulation accuracy and for additional memory usage. We do not recommend that you change any of these

properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients and the cell gradients. These quantities will become available as field functions during subsequent iterations.

Ticked — Retain temporary storage.

Cleared — Do not retain temporary storage.

Identify High Stress Cells

Controls whether the Limiting Equivalent Stress property is used.

Ticked — Activate Limiting Equivalent Stress.

Cleared — Leave Limiting Equivalent Stress deactivated.

Limiting Equivalent Stress

If this property is activated, then the specified value (default `1000 MPa`) is used to flag cells which have equivalent stress above this value. These flagged cells will be treated as bad cells. This marking of high-stress cells will be implemented after 20 iterations to clear the effect of the initial condition.
This feature should be used as a last resort. You should run an analysis without activating it and monitor the stress in the solution. If a diverging cell is identified by high equivalent stress values, then activate this feature and use the solution value as an input for this property.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. Activating this feature should not be your first choice, as it is primarily a “debugging” tool and not supported for normal simulation work.

Ticked	The solver will not update the solution.
Cleared	The solver will update the solution.

Setting Parameters of Other Solvers

The other solvers that become available in STAR-CCM+ involve choices that affect the quality of a simulation with solid stress:

- AMG iteration tolerance
- Implicit dynamic time integration solvers

AMG Iteration Tolerance

The default value of the AMG iteration (convergence) tolerance is generally smaller than what is required for the fluid solver. The default is 0.01. Lower values may be necessary for structures with thin members under bending loads.

Implicit Dynamic Time Integration Solvers

In dynamic analysis, there are two choices to integrate the implicit system of equations:

- Euler, First-Order
- Newmark, Second-Order

This selection would be made from the Temporal Discretization property of the Implicit Unsteady Solver.

The Euler Method is not as accurate as the Newmark Method so a smaller time-step is required to obtain the same accuracy. For given time-step size, the Euler method incurs more numerical damping than the Newmark method for modes of vibration which are at a higher frequency than a time-step can resolve. However, some users may find the Euler method more stable and useful for damping out high-frequency, small-wavelength oscillations that are of little consequence to overall transient behavior.

Solid Stress Formulation

This section includes the following formulation:

- Stress
- Deformation

- Strain
- Thermal strain
- Finite volume stress analysis methodology
- Rayleigh damping methodology

Stress

Stress is a measure of the force divided by the area over which the force is applied, as the area shrinks to zero. In a free-body diagram of a part, a cross-section of the part is made to expose the internal forces acting on the part. At a point the planes of the cross-section can be cut at many angles – each exposing a different set of internal forces. To describe the stress completely at a point on any plane requires a full tensor description as shown in the free-body of a cube. This tensor is often written in a matrix form or in a vector form:

$$\sigma = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix} \tag{2287}$$

$$\sigma = \begin{bmatrix} \sigma_{xx} & \sigma_{yy} & \sigma_{zz} & \sigma_{xy} & \sigma_{yz} & \sigma_{zx} \end{bmatrix}^T \tag{2288}$$



Since the sum of moments must go to zero as the cube shrinks to zero, the stress tensor is symmetric. The off diagonal terms are call shear stresses since they act tangentially to the face. The diagonal terms are referred to as normal stresses since they act normal to a face. The traction stress on any plane defined by its unit normal is given as the matrix product of the vector:

$$\tau = \sigma \bullet \mathbf{n} \begin{Bmatrix} \tau_x \\ \tau_y \\ \tau_z \end{Bmatrix} = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix} \begin{Bmatrix} n_x \\ n_y \\ n_z \end{Bmatrix} \tag{2289}$$



On a certain plane, three orthogonal planes, the shear traction stress is zero and the normal traction stresses on these planes are called the principal stress $\{\sigma_1 \quad \sigma_2 \quad \sigma_3\}$. They are ordered by their magnitude $\sigma_1 \geq \sigma_2 \geq \sigma_3$ σ_1. is the maximum principal stress, σ_3 is the minimum principal stress and σ_2 is the intermediate principal stress. Solid mechanics uses the convention where stresses are positive when in tension and negative in compression so the maximum principal is the most tensile. For this reason it is a often used as measure of tensile failure in a material. It turns out that the principal stresses are the eigenvalues of the stress matrix. In elementary strength of materials, these principal stress are described by the Mohr Circle in two-dimensional stress fields.

On one plane, the shear stress traction is a maximum. The maximum shear stress is the maximum of the differences of the principal stresses:

$$\tau_{\max} = \frac{(\sigma_1 - \sigma_3)}{2} \tag{2290}$$

This is an important measure of ductile failure associated with plastic behavior.

The mean stress is:

$$\sigma_m = \frac{(\sigma_{xx} + \sigma_{yy} + \sigma_{zz})}{3} = \frac{(\sigma_1 + \sigma_2 + \sigma_3)}{3} \tag{2291}$$

If the material is in a hydrostatic state of stress then all principal stresses are equal and the mean stress is interpreted as the negative of the pressure.

If a material is in a state of plane stress, then $\sigma_{zz} = \sigma_{xz} = \sigma_{yz} = 0$. Very thin materials with no loading on the planar surfaces are in a state of plane stress.

Deformation

The deformation of a part is defined by a vector quantity which describes how much a material point on the part is deformed or displaced from its original undeformed configuration. A perfectly rigid body has points that all move according to a single displacement and rotation vector, but for a deformable body, every point in the body is prescribed a different deformation. If a material point in the original undeformed coordinates is $\mathbf{x}_0$ then the displacement of this point in the deformed configuration is $\mathbf{u}(\mathbf{x}_0)$ and the new coordinates of the point are $\mathbf{x} = \mathbf{x}_0 + \mathbf{u}(\mathbf{x}_0)$. Often in fixed, stiff structures the deformations are generally small and some useful approximations can be made to simplify the mathematical model to describe the deformation. The "small displacement" formulation in STAR-CCM+ relies on the assumption that the displacement is small compared to the cell size. The "large displacement" formulation removes this restriction.

The initial release of the solid mechanics features in STAR-CCM+ will rely on the assumption of small displacement.

The deformation gradient $\nabla\mathbf{u}$ measures how the deformation changes from point to point:

$$\nabla\mathbf{u} = \begin{bmatrix} \frac{\partial u}{\partial x_0} & \frac{\partial u}{\partial y_0} & \frac{\partial u}{\partial z_0} \\ \frac{\partial v}{\partial x_0} & \frac{\partial v}{\partial y_0} & \frac{\partial v}{\partial z_0} \\ \frac{\partial w}{\partial x_0} & \frac{\partial w}{\partial y_0} & \frac{\partial w}{\partial z_0} \end{bmatrix} \tag{2292}$$

where the Cartesian components of material point $\mathbf{x}_0$ are $\{x_0 \quad y_0 \quad z_0\}$ and the Cartesian components of the displacement field $\mathbf{u}$ are $\{u \quad v \quad w\}$.

Strain

Strain is the measure of how much a body deforms. If two material points on an undeformed body are originally a distance of l_0 and in some deformed configuration the two points are distance l, the strain ε is then measured as:

$$\varepsilon_G = \frac{(l - l_0)}{l_0} \tag{2293}$$

or:

$$\varepsilon_A = \frac{(l - l_0)}{l} \tag{2294}$$

The former is what is referred to as a Green strain and the latter is called the Almansi strain. When the strain is small, these measures of strain are almost identical. Like the stress, to describe the strain completely at a point requires a tensor:

$$\varepsilon = \begin{bmatrix} \varepsilon_{xx} & \varepsilon_{xy} & \varepsilon_{xz} \\ \varepsilon_{yx} & \varepsilon_{yy} & \varepsilon_{yz} \\ \varepsilon_{zx} & \varepsilon_{zy} & \varepsilon_{zz} \end{bmatrix} \tag{2295}$$

$$\varepsilon = \begin{bmatrix} \varepsilon_{xx} & \varepsilon_{yy} & \varepsilon_{zz} & \varepsilon_{xy} & \varepsilon_{yz} & \varepsilon_{zx} \end{bmatrix}^T \tag{2296}$$

The strain tensor is also symmetric, the diagonal terms are called normal or extensional strains and the off-diagonal terms are called shear strains. For small strains the relationship between deformation and strain is:

$$\varepsilon = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^T) \tag{2297}$$

$$= \left\{ \frac{\partial u}{\partial x_0} \quad \frac{\partial v}{\partial y_0} \quad \frac{\partial w}{\partial z_0} \quad \frac{1}{2}\left(\frac{\partial u}{\partial y_0} + \frac{\partial v}{\partial x_0}\right) \quad \frac{1}{2}\left(\frac{\partial w}{\partial y_0} + \frac{\partial v}{\partial z_0}\right) \quad \frac{1}{2}\left(\frac{\partial u}{\partial z_0} + \frac{\partial w}{\partial x_0}\right) \right\}^T$$

If one inscribed a 90-degree angle on the undeformed material, then the shear strains can be geometrically interpreted as half the change in this angle in the deformed state. For this reason, the engineering strain has a slightly different definition:

$$\varepsilon_{\text{eng}} = \begin{bmatrix} \varepsilon_{xx} & \varepsilon_{yy} & \varepsilon_{zz} & \gamma_{xy} = 2\varepsilon_{xy} & \gamma_{yz} = 2\varepsilon_{yz} & \gamma_{zx} = 2\varepsilon_{zx} \end{bmatrix}^T \tag{2298}$$

Adopting the practice in most finite element codes, STAR-CCM+ output strains follow the engineering convention where the shear strains are twice the conventional shear strains. The volume strain is the change of volume relative to the original volume, and is the sum of the normal strains:

$$\frac{\Delta V}{V} = \varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz} \tag{2299}$$

A material is in a state of plane strain when $\varepsilon_{zz} = \varepsilon_{xz} = \varepsilon_{yz} = 0$.

Thermal Strain

Thermal strain is the measure of the elongation of the material due to changes in the temperature. In an unconstrained state, a homogeneous linear elastic isotropic material with no surface loads is at a uniform temperature, and the material will freely expand and produce no stresses. If the part is constrained, if the part is made of different materials, or if the temperature is changing in the part then thermal stresses will ensue.

In a linear *isotropic* elastic material the thermal strain is a linear function of temperature and impacts only the volume change in the material. It uniquely described by the coefficient of thermal expansion α and the expansion reference temperature T_{ref} :

$$\varepsilon_T(T) = \alpha(T - T_{ref}) \tag{2300}$$

The expansion reference temperature is the temperature at which the thermal strain is zero. In the linear elastic material the mean stress is related to the volumetric strain:

$$\sigma_m = K(\varepsilon_{xx} + \varepsilon_{yy} + \varepsilon_{zz} - 3\varepsilon_T(T)) \tag{2301}$$

This is the linear elastic analog of the equation of state in a fluid relating the pressure, the density and the temperature.

In the full tensor notation the stress is:

$$\sigma = 2G(\varepsilon - \varepsilon_T\delta) + \lambda(tr(\varepsilon) - 3\varepsilon_T)\delta \tag{2302}$$

The thermal strain properties for a solid material are prescribed under Continua.

Finite Volume Stress Analysis Methodology

The differential form of the solid momentum equation is:

$$\rho\ddot{\mathbf{u}} = \nabla \bullet \sigma + \mathbf{b} \tag{2303}$$

where $\mathbf{u}$ is the displacement field, σ is the stress tensor field and $\mathbf{b}$ is the body force.

A general constitutive relation between the stress tensor and the strain tensor is assumed:

$$\sigma = \sigma(\varepsilon, T, \kappa) \tag{2304}$$

where T is the temperature and κ represents other internal variables. An assumption of small strain allows the strain to be expressed as:

$$\varepsilon = \frac{1}{2}\left(\frac{\partial \mathbf{u}}{\partial \mathbf{x}} + \frac{\partial \mathbf{u}^T}{\partial \mathbf{x}}\right) = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^T) \tag{2305}$$

The finite volume form is determined by integration of parts to obtain:

$$\int_V \rho \ddot{\mathbf{u}} dV = \int_A \sigma d\mathbf{s} + \int_V \mathbf{b} dV \tag{2306}$$

For all cells i in the domain V, the discrete form of the momentum equation becomes:

$$\rho_i \ddot{\mathbf{u}}_i V_i = \sum_{j \in f(i)} \sigma_j \bullet \mathbf{s}_j + \mathbf{b}_i V_i \quad \forall i \tag{2307}$$

where the sum is over all the faces j that form cell i with normal surface vector $\mathbf{s}_j$. If one defines the surface force on face j as:

$$\mathbf{f}_j = \sigma_j \bullet \mathbf{s}_j \tag{2308}$$

then Eqn. (2307) becomes:

$$\rho_i \ddot{\mathbf{u}}_i V_i = \sum_{j \in f(i)} \mathbf{f}_j + \mathbf{b}_i V_i \quad \forall i \tag{2309}$$

For the purposes of illustration we will consider only one material on a single domain. In addition, to simplify the initial analysis we will consider only a static or quasi-static problem and drop the inertial term. In this case the sum of the forces must be zero.

$$\mathbf{r}_i = \sum_{j \in f(i)} \sigma_j \bullet \mathbf{s}_j + \mathbf{b}_i V_i \quad \forall i \tag{2310}$$

In order to compute the stress tensor on the centroids of the faces we need to establish the discrete form of the strain tensor ε and the deformation gradient $\nabla \mathbf{u}$. The deformation gradient is computed at cell centers based on the difference in displacement between neighboring points. In an identical manner to viscous fluxes in the fluid momentum, the deformation gradients are interpolated to the face between two cell centers as:

$$\nabla \mathbf{u}|_f = \overline{\nabla \mathbf{u}} + \underbrace{(\delta \mathbf{u} - \overline{\nabla \mathbf{u}} \bullet \mathbf{d})}_{\text{at least first order}} \otimes \frac{\mathbf{s}}{\mathbf{d} \bullet \mathbf{s}} \tag{2311}$$

Where $\overline{\nabla \mathbf{u}}$ is the mean of the gradient computed at the adjoining face centers $\delta \mathbf{u}$ is the difference in the displacement between two adjoining cell centers, and $\mathbf{d}$ is the vector between the two adjoining cell centers. This deformation gradient is at least first-order accurate. The strain (Eqn. (2305)) on the face center is determined and Eqn. (2304) provides the stress on the face.

An approximation of the stiffness matrix can be established including only the terms which depend on $\delta \mathbf{u}$ in the deformation gradient along with the tangent modulus $\partial \sigma / \partial \varepsilon$ supplied by the material model.

The approximation of the stiffness matrix can be accumulated over all faces:

$$K_{ik} = -\frac{\partial \mathbf{r}_i}{\partial \mathbf{u}_k} = -\sum_{j \in f(i)} \frac{\partial \mathbf{f}_j}{\partial \mathbf{u}_k} \quad \forall i, k \in N(i) \tag{2312}$$

It can then be employed in a Newton's method to correct for the displacement:

$$\mathbf{K}_{ik} \Delta \mathbf{u}_k = \mathbf{r}_i^n \quad , \quad \Delta \mathbf{u}_k = \mathbf{u}_k^{n+1} - \mathbf{u}_k^n \tag{2313}$$

This linear system of equations is solved using the AMG iterative solver.

In transient problems the inertial terms are included so that the residual becomes:

$$\mathbf{r}_i = \sum_{j \in f(i)} \sigma_j \bullet \mathbf{s}_j + (\mathbf{b}_i - \rho_i \ddot{\mathbf{u}}_i) V_i \tag{2314}$$

The acceleration terms are modeled either with a 1st order (Backward Euler) or a 2nd order (Newmark) method. Both schemes provide an implicit integration of the system of equations, subject to the initial conditions for the displacement and the velocity. Both methods are absolutely stable regardless of step size.

Rayleigh Damping Methodology

Rayleigh damping is often added in dynamic problems to mimic either real physical damping or to include some numerical damping to damp unwanted numerical oscillations. The damping can be expressed as:

$$\mathbf{C} = \frac{\mathbf{M}}{\tau_M} + \tau_K \mathbf{K} \tag{2315}$$

where $\mathbf{C}$ is the viscous damping matrix, $\mathbf{M}$ is the mass matrix, $\mathbf{K}$ is the stiffness matrix, $1/\tau_M$ is the mass damping coefficient (alpha), and τ_K is the stiffness damping coefficient (beta).

In the FV stress analysis we do not compute the stiffness matrix; an equivalent law for viscous stress is:

$$\boldsymbol{\sigma}_d = 2G\dot{\boldsymbol{\varepsilon}} + \lambda tr(\dot{\boldsymbol{\varepsilon}})\boldsymbol{\delta} \tag{2316}$$

Where the rate of change of the strain, $\dot{\boldsymbol{\varepsilon}}$, is simply from the deformation rate gradient:

$$\dot{\boldsymbol{\varepsilon}} = \frac{1}{2}(\nabla\dot{\mathbf{u}} + \nabla\dot{\mathbf{u}}^T) \tag{2317}$$

The total stress is then the sum of the elastic and viscous damping parts:

$$\boldsymbol{\sigma} = \boldsymbol{\sigma}_e + \boldsymbol{\sigma}_d \tag{2318}$$

By supplying the elastic model with pseudo strain of the form,

$$\tilde{\boldsymbol{\varepsilon}} = \boldsymbol{\varepsilon} + \tau_K\dot{\boldsymbol{\varepsilon}} \tag{2319}$$

the elastic model will return the total stress due to *both the elastic and viscous damping*.

Shared Conditions and Values

The following sections provide descriptions and reference information for objects that may need to be set for the Finite Volume Solid Stress model:

- Region conditions
- Region values
- Boundary conditions
- Boundary and initial condition values

Region Conditions Reference

This section details region conditions that are available for the Finite Volume Solid Stress model:

- Solid Source Option

Solid Source Option

This node lets you specify a body force on the solid region.

Properties

Method The method that will be used to apply a body force on the region.

None	This is the default option and will cause no body forces to be applied to the region.
Specified	This will add a Solid Stress Source node and a Solid Stress Source Force Derivative node to the Values node, which must be entered as vector profiles.

Region Values Reference

This section details region values that are available for the Finite Volume Solid Stress model:

- Solid Stress Source
- Solid Stress Source Force Derivative

Solid Stress Source

This node lets you apply a body force to the solid region. It is added when the Solid Source Option region condition is set to `Specified` and is entered as a vector profile.

Solid Stress Source Force Derivative

This node represents linearization of the Solid Stress Source with respect to x-, y-, [z-] components of the body force, $\mathbf{b}$ in Eqn. (2306). The derivative of the user source term is used to linearize Eqn. (2306). Its value is set to zero by default. Providing a value for the derivative helps stabilize the solution.

The Solid Stress Source Force Derivative is added when the Solid Source Option region condition is set to `Specified` and is entered as a vector profile.

Boundary Conditions Reference

This section details boundary conditions that are available for the Finite Volume Solid Stress model:

- Stress/Displacement

Stress/Displacement

This node allows you to determine how the stress and displacement across the boundary will be specified.

Properties

Method — The method that will be used to define the stress and displacement at the boundary.

Displacement	This will add a Displacement node to the Values node, which must be entered as a vector profile.
Elastic Foundation	This will add a SolidStressElasticFoundation node to the Values node, which must be entered as a scalar profile.
Fixed	The displacement vector is constrained to be zero (all three components). No nodes will be added to the Physics Values node.
Force	This option will add a Force node to the Values node, which must be entered as a vector profile.
Free	Free indicates the applied surface traction is zero. Generally much of the surface of a part can be considered free of loads. No nodes will be added to the Physics Values node.
Mixed	The Mixed option will add three nodes to the Values node: a Displacement node, a Mixed Vector node and a Traction node. Each of these must be entered as a vector profile.
Normal Displacement	This option will add a Normal Displacement node to the Values node, which must be entered as a scalar profile.

Pressure	This option will add a Solid Stress Pressure node to the Values node, which must be entered as a scalar profile.
Traction	This option will add a Traction node to the Values node, which must be entered as a vector profile.

Boundary and Initial Condition Values

Boundary condition methods (values) for the Finite Volume Solid Stress model come in two types: constraints and surface loads. This section provides an introduction into the types of constraints and loads offered by STAR-CCM+ and how to apply a combination of these to a boundary. Insights into working with initial conditions are also provided.

This section also contains reference descriptions of boundary, initial condition and interface values that are available for the Finite Volume Solid Stress model:

- Displacement
- Force
- Mixed
- Normal Displacement
- Solid Stress Pressure
- Solid Stress Elastic Foundation
- Traction

The boundary values are elected by the boundary type, and/or by methods selected in corresponding nodes in the Conditions node.

What Are Constraints?

STAR-CCM+ offers three types of constraint boundary conditions that are applied to a boundary surface.

- Fixed
- Prescribed Displacement
- Prescribed Normal Displacement

These conditions may be applied to a boundary surface. The boundary surface may be a collection of cell faces or a little as one cell face.

Constraints are boundary conditions on a part which impose some constraint on its deformation. The deformation is prescribed as some value. For example one can impose a fixed constraint, which requires the displacement vector to be zero on a given boundary. In FEA a constraint is applied to set of nodes on a boundary but in FVA the constraint is applied to a set of boundary cell face centers. The smallest discrete constraint set is a single node in FEA while in FVA the smallest discrete constraint set is a single boundary cell face center.

Constraining a part properly is extremely important in FEA and FVA for identical reasons. If conducting a static analysis, the body must be constrained so that it is not allowed to freely rotate or translate as a whole. The constraints must remove the so-called "rigid body modes" or a linear solution is not possible. The stiffness matrix will be singular and the solvers will not converge. Not constraining a part to remove these rigid body modes is a common "rookie" mistake. If the solution appears to diverge beginning with the first iteration, carefully examine the constraint boundary conditions.

The predicted stress field near the constrained surface should be viewed with some suspicion. For example a fixed constraint is like welding a part to another part that is infinitely rigid and cannot thermally expand. High stresses can develop at such an interface as the parts are loaded mechanically or thermally. Other than welding or braising, most parts are constrained to move by friction. Examples include the contact friction in a bolted parts between washers and bolt heads, the contact between male and female threads, the contact surfaces local to the teeth of a clamp and the contact produced from one part lying on top of another part under the weight. The constraint boundary conditions are only approximations of such loadings. Luckily, Saint-Venant's principle applies. The stresses computed at a distance large relative to the size of the applied load will be insensitive to the details of the how the load is actually applied. For thermal stresses, the constraints must be carefully chosen so as not to over-constrain the part. Opposing constraints can confine a part producing unrealistically high thermal stresses.

What Are Loads?

Loads are surface or body forces on the solid material. The thermal stresses are commonly called loads. The surface loads are prescribed in STAR-CCM+ as boundary and/or interface conditions. There are five types of boundary surface loads.

- Free
- Prescribed Tractions
- Prescribed Pressure

- Prescribed Forces
- Elastic Foundation

Prescribed tractions, pressures and forces can also be applied to interfaces between solid and fluid domains.

Applying a Combination of Constraints and Loads

The Mixed boundary condition makes it possible to apply a combination of constraints and loads to a boundary. The constraints and loads are applied as prescribed displacements and prescribed tractions, respectively. These are mixed by applying either a constraint or a load in each of the global directions, and the configuration is set by a Mixed Vector. This is used to differentiate which components will be read from the displacement vector, and which from the traction vector.

Setting Initial Conditions

For dynamic simulations, appropriate initial conditions must be specified for displacement as well as velocity. **The default initial values for both displacement and velocity are set to null vectors ([0,0,0]).** These values can be changed through the initial condition node. Field functions may also be used to prescribe initial displacement and velocity profiles.

Displacement

Prescribed Displacement is a more general form where all three components of the displacement are prescribed a constant vector value or a tabular or field function vector value. A prescribed vector of zero displacement would be identical to a fixed displacement condition. The key point to remember here is that all three components of the vector are prescribed a value as opposed to only one or two components. This value node will be visible if `Displacement` is selected as the Method property for the Stress/Displacement node of the Conditions node. It is entered as a vector profile.

Force

The prescribed force is much the same as the prescribed traction. The force vector is divided by the area of the boundary on which it is prescribed. This "average" surface traction is then applied uniformly to the surface regardless of the local orientation of the normal to the surface. You should be wary of computed stresses near the applied surface loads--especially when a large force is applied only over a small surface area.

Gravity and centrifugal loads are body force loads. The gravity or loads due to an accelerating frame of reference are turned on via the gravity option. Centrifugal loads are activated by choosing a rotating reference frame. It is also possible to apply a body force directly on a solid region via the Solid Source Option physics condition.

This value node will be visible for a boundary if `Force` is selected as the Method property for the Stress/Displacement node (boundary) or Interface Load node (interface) of the Conditions node. It is entered as a vector profile.

Mixed

The mixed condition applies a combination of constraints and loads to a boundary. It is not possible to apply both a constraint and a load in any one direction, so the conditions are combined by prescribing a displacement constraint along one or two of the principal axes, and a traction load along the remaining axes. To achieve this, both a displacement vector profile and a traction vector profile are entered, and a Mixed Vector is used to specify which condition will act along each axis.

When using this boundary condition it is recommended that all vector profiles are entered with reference to the default laboratory coordinate system. In cases where the vector profile is referenced by a user-specified coordinate system, the vector will be resolved to the laboratory coordinate system before the mixed vector is applied.

When `Mixed` is selected as the Method property for the Stress/Displacement node of the Conditions node, three nodes will be visible:

- Displacement
- Mixed Vector
- Traction

Each of these will be entered as vector profiles.

What is the Mixed Vector?

The mixed vector is a dimensionless vector that acts as a reference for differentiating which vector components will be applied as displacements, and which as loads. It is entered as a vector profile, with each value equal to either 1 or 0 (when the value entered is neither 1 nor 0, the value is taken as 1.) A value of 0 corresponds to a displacement constraint, and this is read from the displacement vector. A value of 1 corresponds to a traction load, and this is read from the traction vector. By specifying the mixed vector in this way, a mixed vector of `[1, 0, 1]`, for example, is equivalent to saying

"apply a traction load equivalent to the x-component of the traction vector in the x-direction", "apply a displacement constraint equivalent to the y-component of the displacement vector in the y-direction" and so on.

Normal Displacement

In prescribed normal displacement, the component normal to the displacement is prescribed. The tangential components are left free and it is assumed that boundary has zero shear traction. The prescribed normal displacement can be a constant or a field function and a positive value indicates the direction away from the boundary (tensile) while negative is towards the boundary (compressive). This value node will be visible if `Normal Displacement` is selected as the Method property for the Stress/Displacement node of the Conditions node. It is entered as a scalar profile.

Solid Stress Elastic Foundation

The solid stress elastic foundation models a compliant foundation, the surface of which is considered frictionless. Where the solid region rests on the foundation the two surfaces will always remain in contact, so no gap can exist between the elastic foundation and the solid region. As a result, the foundation will place a tensile load on the region if the region attempts to pull away from the foundation.

The mechanical properties of the elastic foundation are defined by a spring constant K. If A is the area of the boundary, then the spring constant per unit area is K/A which is considered to be uniformly applied across the surface. The normal stress exerted by the foundation is:

$$\sigma_n = -\frac{K}{A}\delta_n \tag{2320}$$

where δ_n is the normal displacement as measured outward from the region boundary. A positive normal displacement gives rise to a compressive normal stress.

This value node will be visible if `Elastic Foundation` is selected as the Method property for the Stress/Displacement node of the Conditions node. It is entered as a scalar profile.

Solid Stress Pressure

For prescribed pressure, only the normal component of the surface traction is prescribed. The shear traction components are taken as zero. In this case a positive value of pressure indicates a compressive load while a negative value indicates tension. Therefore a Free condition is the same as the

prescribed traction vector load of [0,0,0] or a prescribed pressure load of zero. This value node will be visible if `Pressure` is selected as the Method property for the Stress/Displacement node (boundary) or Interface Load node (interface) of the Conditions node. It is entered as a scalar profile.

Traction

For prescribed traction, the surface traction is a prescribed vector in units of force/area (pressure). All three components are prescribed. This value node will be visible if `Traction` is selected as the Method property for the Stress/Displacement node (boundary) or Interface Load node (interface) of the Conditions node. It is entered as a vector profile.

Solid Stress Field Functions Reference

The following field functions are made available to the simulation when the Finite Volume Solid Stress model is used.

Solid Strain XX	The measure of strain ε_{xx} in Eqns. (2295) and (2296).
Solid Strain XY	The measure of strain ε_{xy} in Eqns. (2295) and (2296).
Solid Strain XZ	The measure of strain ε_{xz} in Eqns. (2295) and (2296). Activated only in three-dimensional simulations.
Solid Strain YY	The measure of strain ε_{yy} in Eqns. (2295) and (2296).
Solid Strain YZ	The measure of strain ε_{yz} in Eqns. (2295) and (2296). Activated only in three-dimensional simulations.
Solid Strain ZZ	The measure of strain ε_{zz} in Eqns. (2295) and (2296). Activated only in three-dimensional simulations.

Solid Stress Acceleration	The second time derivative of the displacement field: $\ddot{u} = \frac{\partial^2 u}{\partial t^2} = \frac{v^n - v^{n-1}}{\Delta t}$ Activated only when an unsteady time model is used.
Solid Stress Displacement	The displacement field $\mathbf{u}$ in Eqn. (2303).
Solid Stress Equivalent	The equivalent or von Mises stress. It is calculated as: $\sigma_{eq} = \sqrt{\frac{(\sigma_1 - \sigma_2)^2 + (\sigma_2 - \sigma_3)^2 + (\sigma_3 - \sigma_1)^2}{2}}$
Solid Stress Int. Principal	The intermediate principal σ_2 in the principal stress $\{\sigma_1 \quad \sigma_2 \quad \sigma_3\}$.
Solid Stress Max. Principal	The maximum principal stress σ_1 in Eqn. (2290).
Solid Stress Max. Shear	The maximum shear stress traction τ_{max} in Eqn. (2290).
Solid Stress Mean	The mean stress σ_m in Eqn. (2291).
Solid Stress Min. Principal	The minimum principal stress σ_3 in Eqn. (2290).
Solid Stress Normal Displacement	The normal component of the displacement field $\dot{u}.n$.
Solid Stress Pressure	The normal component of the traction on a given surface $(\sigma.n).n$.

Solid Stress Velocity	The first time derivative of the displacement field: $\dot{\vec{u}} = \frac{\partial u}{\partial t} = \frac{u^n - u^{n-1}}{\Delta t}$ Activated only when an unsteady time model is used.
Solid Stress XX	The measure of stress σ_{xx} in Eqns. (2287) and (2288).
Solid Stress XY	The measure of stress σ_{xy} in Eqns. (2287) and (2288).
Solid Stress XZ	The measure of stress σ_{xz} in Eqns. (2287) and (2288). Activated only in three-dimensional simulations.
Solid Stress YY	The measure of stress σ_{yy} in Eqns. (2287) and (2288).
Solid Stress YZ	The measure of stress σ_{yz} in Eqns. (2287) and (2288). Activated only in three-dimensional simulations.
Solid Stress ZZ	The measure of stress σ_{zz} in Eqns. (2287) and (2288). Activated only in three-dimensional simulations.
Thermal Strain	The measure of the thermal strain ε_T in Eqn. (2300).

Modeling Electromagnetism

In the current STAR-CCM+ release, two electric potential models and one magnetic vector potential model are available, which allow solving for the electric and magnetic fields. The Electromagnetism option is the gateway to these models. The Electromagnetism model node represents this model in the simulation tree and is subject to common menu items.



The Electrodynamic Potential model can be used to model electric currents and the Electrostatic Potential model can model electrostatic effects. As these models solve for different physics, they cannot be selected together. However, each of them uses the electric potential solver. The Magnetic Vector Potential model is used to model magnetic fields and can be combined with the Electrodynamic Potential model.

The Java API defines several compile-time constant expressions, which are common for electrodynamics. These constants can be accessed in a Java macro.

The formulation provides a brief theoretical background for these models. Information on field functions is also included.

Modeling Electric Currents

This section explains the following:

- What is the Electrodynamic Potential model?
- Region conditions
- Boundary conditions
- Interface condition
- Field functions
- Modeling Magnetohydrodynamics
- Modeling Ohmic Heating
- Modeling Electro-Deposition Coating

What Is the Electrodynamic Potential Model?

The Electrodynamic Potential model is available for Gas, Solid, and Liquid phases, and can also be combined with the VOF Multiphase model. You can also use the Multi-Part Solid model to set different electrical conductivities for different parts within the same physics continuum.

The Electrodynamic Potential model is compatible with the Multi-Component Gas and Multi-Component Liquid models, although the electrical conductivity that you specify is the effective conductivity of the mixture. STAR-CCM+ does not let you specify the electrical conductivity as a function of the conductivities of the components.

This model solves for the electric potential according to Eqn. (2389). The electrical conductivity can be specified for linear, isotropic materials as a scalar quantity. From the electric potential the electric field Eqn. (2361) and the electric current density Eqn. (2360) are calculated. This model can be coupled together with the Magnetic Vector Potential model, which includes the magnetic vector potential term in the calculation (Eqn. (2389)).

You can also use the electrodynamic potential model to simulate corrosion and determine the effectiveness of various cathodic protection scenarios for a number of different applications. STAR-CCM+ solves the Butler-Volmer equation Eqn. (2441) at the interface between the geometry that is corroding and the electrolyte.

Selecting the Electrodynamic Potential Model

To make the Electrodynamic Potential model available:

1. Select a space model.
2. Select a material model.

   The Electromagnetism model becomes available in the *Optional Models* group box.
3. Select Electromagnetism from the *Optional Models* group box.

   The Electrodynamic Potential model becomes available in the *Electromagnetism* group box.
4. Select Electrodynamic Potential from the *Electromagnetism Models* group box.
5. Close the dialog.

The list of selected models can be seen by selecting the Models node from the respective physics continuum.



This model also activates field functions.

Setting the Electrical Conductivity

The Electrodynamic Potential model requires you to specify the appropriate electrical conductivity values as material properties. The electrical conductivity quantifies the ability of the material to conduct electric current, and is measured in S/m. In most cases, the electrical conductivity is entered as a scalar profile. The exception is anisotropic electrical conductivity, which is entered as a tensor profile.

The Electrical Conductivity node represents this material property in the simulation tree.



Electrical Conductivity Properties

Method — Select the material property method to use to set the electrical conductivity.

Anisotropic	The electrical conductivity is anisotropic. This option is available only for solid materials.
Constant	The electrical conductivity is specified as a constant.
Electrochemical Species	Automatically computes the electrical conductivity using Eqn. (2448).
Field Function	The electrical conductivity is specified using a field function.
Polynomial in T	The electrical conductivity is specified using a Polynomial in T. This option is available only when an energy model is activated in the physics continuum.
Table (T)	The electrical conductivity is specified using a Table (T). The available temperature range in the table must be consistent with the Minimum/Maximum Allowable Temperature settings that are specified under the Reference Values node for the physics continuum. This requirement is specific to electrical conductivity - it is not a requirement for all Table (T) temperature tables. This option is available only when an energy model is activated in the physics continuum.
Tabular	The electrical conductivity is specified using a table interpolation method. Select the independent variables (usually temperature and concentration) that are used to interpolate the electrical conductivity. This option is available only when an energy model is activated in the physics continuum.

Setting the Solid Electrical Conductivity in Porous Media

Porous regions, such as fuel cell catalyst layers and battery electrodes, are modeled as volumetric regions where there is a separate electronic potential and ionic potential.

Note: The solid electrical conductivity is specified at the porous region level but the fluid electrical conductivity is specified at the fluid continuum level.

To define the solid electrical conductivity within a porous region:

1. Select the Solid Electrical Conductivity node within the Physics Conditions folder for the region.
2. In the *Properties* window, specify the appropriate tensor profile.



Setting a Transfer Current Density Region Source

The Electrodynamic Potential model activates the Transfer Current Density region source. In the Physics Conditions of a region, you can choose to specify the source in the *Properties* window.

To define the transfer current density source within a region:

1. Select the Transfer Current Density Source Option node within the Physics Conditions folder for the region.
2. In the *Properties* window, activate Transfer Current Density User Source Option.



An additional node, Transfer Current Density User Source, is added to the *Physics Values* folder for the region, where you can set its value. The default option uses a constant value. The alternative options are described in the scalar profile properties section.

Setting Region Boundary Conditions

Selecting the Electrodynamic Potential model activates the Electrodynamic Potential region boundary conditions. In the Physics Conditions of each boundary of a region, you can specify the boundary condition type in the *Properties* window. The default boundary condition is Insulator: no electric current penetrates the boundary.

To define the boundary conditions within a region:

1. Select the Electric Potential Specification node within the Physics Conditions folder for the region boundaries.
2. In the *Properties* window, set the Method property:
   - For the Neumann boundary condition, select one of the following:
     - Insulator
     - Specific Electric Current
     - Electric Current
     - Electric Current Density

- Electric Current Potential Characteristic



The appropriate corresponding nodes, Specific Electric Current, Electrical Resistance, Electric Current, Electric Current Density, or Current Voltage Characteristic and Electric Potential, is added to the *Physics Values* folder for the region boundary, where you can set the relevant value.

You can use the Electric Current Potential Characteristic to define a characteristic relation of electric current to electric potential drop at a boundary or an interface.

- For the Dirichlet boundary condition, select Electric Potential.



Three additional nodes, Electric Potential, Electrical Resistance, and Electrical Resistance Potential Derivative, are added to the *Physics Values* folder for the region boundary, where you can set the relevant values.

You can use the Electrical Resistance Potential Derivative to improve the convergence of the electric potential solver when you are modeling strongly non-linear electrical resistance.

The default option for these physics values uses a Constant method. The alternative method options are described in the scalar profile and vector profile properties section.

The Specific Electric Current and Electric Current Density methods both represent current density. The Specific Electric Current is entered as a scalar profile, while the Electric Current Density is entered as a vector profile.

As with boundaries, you can also set a specific electric current or total electric current value on an interface using the Electric Current Source Option node.



Setting the Electric Current Potential Characteristic

For some engineering applications, it is of interest to prescribe a characteristic relation of electric current to electric potential drop at a boundary or an interface. This characteristic relation can also be known as the potential-voltage characteristic or polarization curve.

An example is cathodic protection, in which the complexity of the electro-chemical corrosion process is reduced to the polarization curve of a specific setup. Much can be learned from the distribution of electric potential and electric current.

You can specify an Electric Current Potential Characteristic condition at boundaries and contact interfaces when the Electrodynamic Potential Model is selected.



You specify the Current Voltage Characteristic method under the Physics Values node.



Two methods are provided:

- Butler-Volmer
- Tabular Polarization Curve

Using the Butler-Volmer Method

You specify the appropriate property values in the *Butler-Volmer Properties* window.



See Butler-Volmer Properties

Using the Tabular Polarization Curve Method

The relationship between the potential drop across the boundary or interface, $\phi_1 - \phi_0$ (see Eqn. (2323)), and the specific electric current J_n can be specified by tabular data.

You specify the appropriate property values in the *Tabular Polarization Curve Properties* window.



See Tabular Polarization Curve Properties

Butler-Volmer Properties

Anodic Transfer Coefficient	The anodic apparent transfer coefficient, α_a. Used to determine the relation between specific electric current and overpotential in Eqn. (2321). Also, used to determine the current density in the Butler-Volmer equation, Eqn. (2441).
Cathodic Transfer Coefficient	The cathodic apparent transfer coefficient, α_c. Used to determine the relation between specific electric current and overpotential in Eqn. (2321). Also, used to determine the current density in the Butler-Volmer equation, Eqn. (2441).
Specific Exchange Current	This value represents specific electric exchange current, J_E, used in Eqn. (2321) and, i_0, used in the Butler-Volmer equation, Eqn. (2441).
Zero Electric Current Potential	The zero electric current potential, ϕ_E in Eqn. (2321). This value specifies the electric potential drop when there is zero current.

Tabular Polarization Curve Properties

Table	Allows you to select the appropriate imported polarization curve file table. This table provides pairs of potential drop ($\phi_1 - \phi_0$) and current values.
Electric Potential	The column of the table that specifies the electrical potential drop.
Specific Electric Current	The column of the table that specifies the specific current density.
Linearization Factor	A linearization factor that can enhance convergence. The specified factor is applied to the electric resistance potential derivative that is computed internally at boundaries. You can enhance convergence by increasing the linearization parameter from its default value of zero up to 1. For best results, it is suggested that you use a value similar to the under-relaxation factor used for the electric potential.
Zero Electric Current Potential	The zero current potential drop ϕ_E in Eqn. (2326)specifies the electric potential drop when there is zero current. The value is extracted from the imported file table and is defined by interpolating the condition $J_n(\phi_E) = 0$.

Electric Current Potential Characteristic Formulation

Butler-Volmer

The relation between specific electric current, j_n, and the over-potential η is:

$$j_n(\eta) = j_0(e^{\alpha_a F\eta/(R_u T)} - e^{-\alpha_c F\eta/(R_u T)}) \tag{2321}$$

where:

- α_a is the anodic apparent transfer coefficient
- α_c is the cathodic apparent transfer coefficient
- j_0 is the specific electric exchange current
- ϕ_E is the zero electric current potential

- $F = 9.648525879 \times 10^7 \quad \mathrm{C\ kmol^{-1}}$ is the Faraday constant
- $R_u = 8.3144621 \times 10^3 \quad \mathrm{J\ kmol^{-1} K^{-1}}$ is the universal gas constant

If a temperature T is not provided in the boundary / interface cell by an energy model, a default temperature of $T = 293K$ is used. A step change in the temperature across the interface should be avoided.

The over-potential is defined as:

$$\eta = \phi_1 - \phi_0 - \phi_E = j_n R_\eta \tag{2322}$$

For an interface, ϕ_0 and ϕ_1 denote the potential values on boundary 0 and boundary 1 of the interface respectively. For a boundary, ϕ_0 and ϕ_1 denote the potential of the solution at the boundary and the specified potential value respectively.

The zero-current potential drop is defined as:

$$\phi_E = \phi_1 - \phi_0 \tag{2323}$$

for which $\eta = 0$ and $j_n(\phi_E) = 0$ holds.

An efficient way of imposing a given relation between the electric current and the potential drop across a boundary or an interface is realized by utilizing the zero-current potential drop jointly with an internally-computed electrical resistance.

The electric resistance is defined as:

$$R_\eta = \frac{\eta}{j_n(\eta)} \qquad \forall j_n \neq 0 \tag{2324}$$

$$R_\eta = \frac{R_u T}{F} \frac{1}{(\alpha_a + \alpha_c) j_0} \qquad \forall j_n = 0 \tag{2325}$$

For boundaries, the electrical resistance potential derivative is also computed internally to enhance convergence.

Tabular Polarization Curve

The polarization curve has the following requirements:

- More than two pairs need to be provided.
- The values of the potential drop are strictly monotonically increasing.
- The values of current exhibit a single sign change from negative to positive values.

The over-potential is defined as:

$$\eta = \phi_1 - \phi_0 - \phi_E = j_n R_\eta \tag{2326}$$

The resistance is computed as:

$$R_\eta = \frac{\eta}{j_n(\eta)} \qquad \forall j_n \neq 0 \tag{2327}$$

$$R_\eta = \frac{1}{j'_n(\phi_E)} \qquad \forall j_n = 0 \tag{2328}$$

where $j'_n(\phi_E)$ denotes the derivation of the current at the zero current potential. For well-defined polarization curves, the resistance thus defined is always positive.

The polarization curve is an essential part of the simulation and should be specified with care. In particular, the polarization curve should cover the range of over-potential required for the simulation; otherwise it can be better to revert to a Butler-Volmer relation. To enhance convergence of sufficiently complex polarization curves it can help to decrease the under-relaxation parameter of the electric potential solver.

Setting Interface Conditions

The Electrodynamic Potential model activates the Electrical Resistance option on a contact interface. Equations (2396) and (2397) calculate the potential difference over the interface. The default option is to specify a constant value for the electrical resistance. The alternative options are described in the scalar profile properties section.



Energy that is dissipated by the electric current going through contact interfaces with the user-defined electrical resistance does not make any contribution to the source term (Eqn. (2382)) of the energy equation.

Baffle Interfaces and Porous Baffle Interfaces

The Electrodynamic Potential model supports Baffle and Porous Baffle interface types. These interface types are available when each of the two regions is either a fluid region or a porous region.

To specify the interface condition for a baffle or a porous baffle:

1. Select the Baffle Electrical Option node within the Physics Conditions folder for the interface.

2. In the *Properties* window, specify the appropriate option.



Baffle Electrical Option Properties

Method	Specifies whether the baffle conducts electricity.	
	Conducting	Current flows through the baffle. The appropriate electrical resistance can be set under the Physics Values node for the interface.
	Non-Conducting	No current flows through the baffle. This is the default setting.

Electrodynamic Potential Field Functions Reference

The following primitive field functions are made available to the simulation when the Electrodynamic Potential model is used. These objects all have the same base properties.

Boundary Specific Electric Current	See Eqn. (2335) and Eqn. (2394). See Specifying the Specific Electric Current and Electric Flux.
Electric Current Density	See Eqn. (2360).
Electrical Conductivity	See Eqn. (2360). This field function is not available for solid regions that have a specified electrical conductivity method of Anisotropic. For such regions, the electrical conductivity tensor components are defined separately at the region level for each region in the solid continuum. This field function is not available for porous regions, because the conductivity of porous regions is defined to be tensorial. For both of the above cases, only the Electrical Conductivity Tensor Component field function is available.
Electrical Conductivity Tensor Component	See Eqn. (2383). The effective electrical conductivity tensor Σ_{eff} for a porous region. This field function is also available for solid regions that have a specified electrical conductivity method of Anisotropic.

Modeling Magnetohydrodynamics (MHD)

This section contains the following:

- Using the One-Way Coupled MHD Model
- Setting the Prescribed Magnetic Flux Density
- Magnetohydrodynamic Field Functions Reference
- Low Magnetic Reynolds Number Formulation

Using the One-Way Coupled MHD Model

The One-Way Coupled MHD model models the interaction between the 3D Flow Model and the Electrodynamic Potential Model. This model is appropriate if the induced magnetic flux density is small compared to the prescribed magnetic flux density (see Eqn. (2330)). The model calculates the additional source terms to the flow and the electrodynamic potential equations.

To select the One-Way Coupled MHD model:

1. In the *Physics 1 Model Selection* dialog, do the following:
   - From the *Space* group box, select a 3-D space model.
   - From the *Material* group box, select the appropriate fluid.
   - From the *Flow* group box, select the appropriate flow model.
2. Select Electromagnetism from the *Optional Models* group box.
3. Select Electrodynamic Potential from the *Electromagnetism* group box.
4. Select One-Way Coupled MHD from the *Optional Models* group box.
5. Select the appropriate models from the following:
   - *Equation of State* group box
   - *Time* group box
   - *Viscous Regime* group box.
6. Close the *Physics 1 Model Selection* dialog.

The list of selected models can be seen by selecting the Models node from the respective physics continuum.



The One-Way Coupled MHD model has no properties, but it activates related field functions.

Setting the Prescribed Magnetic Flux Density

When you use the One-Way Coupled MHD model, you set the prescribed magnetic flux density in the appropriate regions.

To set the prescribed magnetic flux density in a region:

1. Expand the Regions node and, in the appropriate region, expand the Physics Values node.
2. Select the Prescribed Magnetic Flux Density node.

3. In the *Prescribed Magnetic Flux Density - Properties* window, set the Method property to the appropriate option.



4. Select the node under the Prescribed Magnetic Flux Density node and, in the *Properties* window, make the appropriate settings to specify the magnetic flux density.

Magnetohydrodynamic Field Functions Reference

The following primitive field functions are made available to the simulation when the Magnetohydrodynamic model is used. These objects all have the same base properties.

Electric Current Density	See Eqn. (2332).
Lorentz Force Density	See Eqn. (2329).

Magnetic Flux Density See Eqn. (2331).

Low Magnetic Reynolds Number Formulation

Magnetohydrodynamics describes the interaction between a conducting fluid in motion and the electric-magnetic fields.

The interaction is modeled using a body force which is added to the Navier-Stokes equation. This force is called the Lorentz force (or space force) and has the following form:

$$\boldsymbol{f}_L = \boldsymbol{J} \times \boldsymbol{B} \tag{2329}$$

In Eqn. (2329), $\boldsymbol{J}$ is the electric current density and $\boldsymbol{B}$ is the magnetic flux density, both of which are vector quantities. The unit for the Lorentz force, $\boldsymbol{f}_L$, is $[N/m^3]$.

The electrically conducting fluid in relative motion to the magnetic field induces electric current density which is calculated as $\sigma \boldsymbol{v} \times \boldsymbol{B}$.

Here, $\boldsymbol{v}$ is the flow velocity and σ is the electrical conductivity (see Eqn. (2360)).

The applicability of the one-way coupled model is restricted to cases where flow induced magnetic flux density $\boldsymbol{B}_i$ is significantly weaker than the prescribed one $\boldsymbol{B}_p$. A non-dimensional characteristic number that is defined as the ratio between the induced and prescribed magnetic flux densities is called the magnetic Reynolds number and is defined as:

$$Re_{mag} = \mu_0 \sigma U L \tag{2330}$$

where μ_0 is the permeability of vacuum (see Eqn. (2374)), U is the characteristic velocity and L is the characteristic length.

For small Re_{mag} the induced magnetic flux density can be neglected and the magnetic flux density becomes equal to the prescribed one, that is:

$$\boldsymbol{B} = \boldsymbol{B}_p \tag{2331}$$

Further, the electric current density is:

$$\boldsymbol{J} = \sigma(-\nabla\phi + v \times \boldsymbol{B}_p) \tag{2332}$$

Taking the incompressibility of the electric current density (see Eqn. (2379)), the transport equation for the electric potential becomes:

$$-\oint_A \sigma \nabla \phi d\mathbf{a} + \oint_A \sigma v \times \boldsymbol{B}_p d\mathbf{a} = \int_V S dV \tag{2333}$$

For a full list of terms and their definitions, see Coupled Flow and Energy Nomenclature and Electromagnetism Nomenclature.

Modeling Ohmic Heating

This section explains the following:

- What is the Ohmic Heating model?
- Setting up the Ohmic Heating model
- Field functions

What Is the Ohmic Heating Model?

The Ohmic Heating model is used to solve for the heat that is generated in the conductive material due to the flow of current, often referred to as the Joule effect. The Ohmic Heating model is an optional model to the Electrodynamic Potential model and energy model. This model solves the total energy equation with an additional source term given by Eqn. (2381) and Eqn. (2398). This term adds the total power dissipation of an existing electric current density in an electric field to the energy equation. It can be used with all models with which the energy model together with the Electrodynamic Potential model are compatible. The temperature dependency of the electrical conductivity that the Electrodynamic Potential Model registers can be modeled utilizing STAR-CCM+'s capability of using field functions for material properties.

To make the Ohmic Heating model available:

- Select the Electrodynamic Potential model.

  The Ohmic Heating model appears as an option in the *Optional Models* group box.

You can see the list of selected models by selecting the Models node from the respective physics continuum.



This model also activates related field functions.

Setting Up the Ohmic Heating Model

This section contains recommendations for setting up a simulation using the Ohmic Heating model.

Selecting Regions for Ohmic Heating

If a region where you solve for the electric potential is expected to have significant energy dissipation due to electric currents, consider the Ohmic Heating source term. The Ohmic Heating source term is proportional to the electric current density and the electric field. If this term is significant, compared to all other source term contributions in the energy equation, select the Ohmic Heating model.

Defining Electrical Conductivity

An important physical property to define when using the Ohmic Heating model is the Electrical Conductivity σ, that is defined in S/m (Siemens per meter). Electrical conductivity is dependent on temperature, so it has a strong effect on the accuracy of the heat calculation. The dependence of electrical conductivity on temperature must be taken into account. You cannot define electrical conductivity as a constant, but three other options are available:

- Polynomial in T (if you have a polynomial dependency)

- A field function
- Anisotropic

One method for defining the electrical conductivity for a given material is to obtain values for σ at different temperatures and states from literature. The electrical conductivity law can then be obtained using a regression analysis technique.

Setting Boundary Conditions

There are two options available to define the electric potential specification at the boundaries:

- **Electric Potential** ϕ, defined in Volts
- **Specific Electric Current** J, defined in A/m^2, also referred to as the **Current Density**

To allow the current to enter and leave the domain, set the electric potential or specific electric current at the "inlets" and "outlets" to the electrified domain, for example at the terminals. The current density is considered to be incompressible, meaning that it is the same everywhere in the domain. Therefore, the amount of current going into the domain is the same as the current leaving it.

When defining the boundary conditions, different techniques are available. For example, consider a wire for which the resistance R and the current I are known. In this case, there would be three possible ways to set the boundary conditions:

- Based on I and R it is possible to deduce the potential difference, and then apply electrical potentials at both ends of the wire to reflect the potential difference.
- Knowing I and the inlet area, determine J and apply it with a positive sign at the "inlet" and a negative sign at the "outlet".
- Knowing I and the inlet area, determine J and apply it with a positive sign at the inlet, then define a null electrical potential at the outlet to allow a null net flux.

Use the Current Density Source Option with care. It is only applicable when there is a production of current from electrochemical reactions or electromagnetic effects.

Ohmic Heating Field Functions Reference

The following primitive field function is made available when the Ohmic Heating model is used.

Specific Ohmic Heat Source	See Eqn. (2381).
Specific Ohmic Heat Surface Source	See Eqn. (2398).

Modeling Electro-Deposition Coating

It is possible to simulate e-coating in STAR-CCM+ with the Electro-Deposition Coating model. This model uses its own solver to generate results.

The process begins with activating the model. Material properties and other parameters are specified for setting up the coating. You would then select one or more boundaries to coat.

Post-processing of your results is feasible with the field functions activated by this model.

This section also includes the formulation of this model.

What Is the Electro-Deposition Coating Model?

The Electro-Deposition Coating model is an optional model to the Electrodynamic Potential model, flow model, and Implicit Unsteady model.

To make the Electro-Deposition Coating model available:

- Select the Electrodynamic Potential model.

  The Electro-Deposition Coating model appears as an option in the *Optional Models* group box.

The list of selected models can be seen by selecting the Models node from the respective physics continuum



This model also activates field functions.

The Electro-Deposition Coating model is a framework to obtain the paint layer thickness on a specified surface that is wetted by a paint liquid. Calibration is required for a specific paint liquid along with given process parameters. The model can be combined with the VOF Multiphase model to

cover gas bubbles resulting from dipping a body into a paint bath. The following figure illustrates the notation that is used in the model formulation.



The formulation makes the following assumptions:

- The paint layer phase is modeled by a function available at a coated surface (boundary), rather than represented by a volume mesh.
- The formulation assumes a quasi-steady state for the electric potential. The steady equation is solved together with possible time-dependent boundary conditions.
- The model is calibrated against actual process parameters.
- The electric potential field and the layer growth do not affect the flow field.

What Is the Electro-Deposition Coating Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the Electro-Deposition Coating model in all the continua for which this model is activated.



Electro-Deposition Coating Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the newly computed solution supplants the old solution.

Expert Properties

Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work.

Activated	The solver does not update the solution.
Deactivated	The solver updates the solution.

Setting up the Coating

The sub-nodes of the Electro-Deposition Coating model specify material properties and other characteristics for the coating.



These objects provide the following inputs to the calculations:

- Paint Layer — This object has a sub-node that works like the material properties manager: it contains material properties objects that are specific to this type of coating:
  - Coulomb Efficiency — c_{paint} in Eqns. (2335) and (2336)
  - Density — ρ_{paint} in Eqns. (2335) and (2336)
  - Electrical Resistivity — r_{paint} in Eqn. (2334)
- Minimum Accumulated Specific Surface Charge — scalar profile.
- Minimum Specific Electric Current — j_{min} in Eqn. (2335). Also set as an individual boundary value.
- Electrical Resistance Offset — g in Eqn. (2334). Also set as an individual boundary value.

Selecting a Boundary to Be Coated

To select a particular boundary for coating, you would select the Electro Deposition Coating Specification node in the Physics Conditions node of that boundary. The properties of this object provide two approaches for coating: using continuum values or specifying custom values for the boundary.



Electro-Deposition Coating Specification Properties

Method — Specifies the electro-deposition coating method that is used.

None	Leaves the boundary uncoated (default).
Current Density	Applies the coating. The coating uses the continuum values or the custom boundary values, depending on the setting of the Option property.

Option — Selects the source of values for electro-deposition coating.

Use Continuum Values	Uses values that are specified for the continuum.
Specify Custom Values	Uses values that are set for this boundary. To allow these settings, this option adds four profile nodes to the Physics Values node for the boundary: • Electrical Resistance Offset • Initial Paint Electrical Resistance • Initial Paint Layer Thickness • Minimum Specific Electric Current

Specifying Electrical Resistance Offset

The Electrical Resistance Offset boundary value is activated when the Option property of the Electro Deposition Coating Specification boundary condition is set to `Specify Custom Values`. It is input as a scalar profile.

Specifying Initial Paint Electrical Resistance

The Initial Paint Electrical Resistance boundary value is activated when the Option property of the Electro Deposition Coating Specification boundary condition is set to `Specify Custom Values`. It is input as a scalar profile.

Specifying Initial Paint Layer Thickness

The Initial Paint Layer Thickness boundary value is activated when the Option property of the Electro Deposition Coating Specification boundary condition is set to `Specify Custom Values`. It is input as a scalar profile.

Specifying Minimum Specific Electric Current

The Minimum Specific Electric Current boundary value is activated when the Option property of the Electro Deposition Coating Specification boundary condition is set to `Specify Custom Values`. It is input as a scalar profile.

Electro-Deposition Coating Field Functions Reference

The following primitive field functions are made available to the simulation when the Electro-Deposition Coating model is used. These objects all have the same base properties.

Paint Density	See Eqn. (2335).
Paint Coulomb Efficiency	See Eqn. (2335).
Paint Electrical Resistivity	See Eqn. (2334).
Paint Layer Thickness	See Eqn. (2335).
Paint Layer Thickness Growth Rate	See Eqn. (2335).
Paint Electrical Resistance Area	See Eqn. (2334).
Boundary Specific Electric Current	See Eqn. (2335).

Electro-Deposition Coating Formulation

The basic transport equation is the equation for the electric potential (Eqn. (2389)) with a zero source term, plus the boundary condition (Eqn. (2394)). Here σ is the electrical conductivity of the liquid bath. At anodes possibly time-varying values for ϕ^{∞} are applied. At coated cathodes, you set R^{User}, which is a function of the thickness of the metal sheet and its electrical conductivity. The potential in the center of the metal sheet is assumed to be zero. For R^{Model} the following starting equation is found to be appropriate:

$$\frac{dR^{\text{Model}}}{dt} = \frac{dh}{dt} r_{\text{paint}} + g \tag{2334}$$

Here h is the paint layer thickness in $[m]$. By default the offset scalar profile g is zero and the resistivity of the paint layer r_{paint} is assumed to be constant. In reality, it is most likely a more complicated function.

The paint layer growth rate (paint deposition rate) is a function of the boundary specific electric current j_n at the interface between the paint bath and paint surface. The ODE (Ordinary Differential Equation) for paint layer thickness along a coated body boundary is:

$$\frac{dh}{dt} = \frac{c_{\text{paint}}}{\rho_{\text{paint}}}(j_n - j_{min}) \tag{2335}$$

where j_n is calculated from the boundary conditions. The proportionality factor is the inverse of the electrochemical equivalent E which is:

$$E = \frac{\rho_{\text{paint}}}{c_{\text{paint}}} \tag{2336}$$

Here the material property ρ_{paint} is the density of the paint solid in $[kg/m^3]$ and c_{paint} is the solid paint coulomb efficiency in $[kg/C]$. The model parameter j_{min} is the minimum specific electric current ($_{[A/m^2]}$) which is required to observe any deposition at the boundary. On the other hand, the formulation assumes that the deposition rate is always positive or zero.

Measurements of polarization effects for different mixtures of paint indicate that there is a further mechanism to control the paint layer growth. A nonlinear function is integrated and the Minimum Accumulated Specific Surface Charge scalar profile is used to control the contribution to thickness building. Values greater than 0 based on individual paint calibration affect the simulated paint layer thickness distribution.

Modeling Magnetic Fields

This section describes the following:

- Using the Magnetic Vector Potential model
  - Region conditions
  - Boundary conditions
  - Magnetic Vector Potential solver
  - Field functions
- Using the Transverse Magnetic Potential model
  - Region conditions
  - Boundary conditions
  - Magnetic Vector Potential solver
  - Field functions
- Setting the magnetic permeability
- Modeling permanent magnets
- Modeling magnetic flux linkage

Using the Magnetic Vector Potential Model

The Magnetic Vector Potential model is available for three-dimensional Gas, Solid, and Liquid phases, and can also be combined with the VOF Multiphase model. You can also use the Multi-Part Solid model to set different permeabilities for different parts within the same physics continuum.

This model solves for the magnetic vector potential according to Eqn. (2390). From the magnetic vector potential, the magnetic flux $\mathbf{B}$ (Eqn. (2362)) and the magnetic field $\mathbf{H}$ (Eqn. (2359)) are calculated. For the Solid material model, it also activates the Permanent Magnet model. The Magnetic Vector Potential model can be used in combination with the Electrodynamic Potential model, which includes the electric potential term in the calculation (Eqn. (2390)).

Selecting the Magnetic Vector Potential Model

To make the Magnetic Vector Potential model available:

1. Select a material model.
2. In the *Space* group box, select Three Dimensional.
3. Select Electromagnetism from the *Optional Models* group box.

4. Select Magnetic Vector Potential from the *Electromagnetism Models* group box.

The list of selected models can be seen by selecting the Models node from the respective physics continuum.



This model also activates field functions.

Setting an Electric Current Density Region Source

The Magnetic Vector Potential model activates the Electric Current Density region source. In the Physics Conditions of a region, you can choose to specify the source in the *Properties* window. The Electric Current Density region source is a vector profile.

To define the electric current density source within a region:

1. Select the Electric Current Density Source Option node within the Physics Conditions folder for the region.

2. In the *Properties* window, set the option to Specified.



If you combine the Magnet Vector Potential model with the Electrodynamic Potential model, the Transfer Current Density Source Option is made available instead.

An additional node, Electric Current Density, is added to the *Physics Values* folder for the region, where you can set its value. The default option uses a constant value. The alternative options are described in the scalar profile properties section.

Setting Region Boundary Conditions

The Magnetic Vector Potential model activates the Magnetic Vector Potential region boundary conditions. In the Physics Conditions of the boundary of a region, you can specify the boundary condition type in the *Properties* window. The boundary condition is a vector profile.

To define the boundary conditions within a region:

1. Select the Magnetic Vector Potential Specification node within the Physics Conditions folder for the region boundaries.
2. In the *Properties* window, set the Method property:

Electric Current Sheet	Used for the Neumann boundary condition. This boundary condition specifies the normal derivative of the potentials along the boundary.
Magnetic Vector Potential	Used for the Imposed Magnetic Vector Potential boundary condition. This boundary condition is defined when the values of the potentials $\mathbf{A}$ and/or ϕ are explicitly specified on the boundary.
Tangential Magnetic Field	Used for the tangential magnetic field boundary condition.
Symmetry - Perfect Magnetic Conductor	Used for the Symmetry boundary condition. The normal component of the magnetic vector potential value $\mathbf{A}_{\infty}\big\|_{n}$ is set to zero.
Anti-Symmetry - Perfect Electric Conductor	Used for the Anti-Symmetry boundary condition. The tangential components of the magnetic vector potential value $\mathbf{A}_{\infty}\big\|_{t_1, t_2}$ are set to zero.

The corresponding node is added to the Physics Values node for the region boundary. This node lets you set the relevant value. The default option uses a constant method. The alternative methods are described in the scalar profile properties section.

The example below shows the Magnetic Vector Potential method set.



What is the Magnetic Vector Potential Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the magnetic vector potential model in all the continua for which the model is activated.

The Magnetic Vector Potential solver node is shown below.



Magnetic Vector Potential Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the new solution replaces the old one. The default value is 1.0. ω in Eqn. (2337).

Under-Relaxation Factor

The Magnetic Vector Potential solver uses standard under-relaxation:

$$\mathbf{A}^{n+1} = \mathbf{A}^{n} + \omega \mathbf{A}' \tag{2337}$$

where ω is the under-relaxation factor.

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and more memory usage. It is not recommended that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in

STAR-CCM+.

Reconstruction Freeze	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Update reconstruction gradients with each iteration.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero.	
	Activated	Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed.
	Deactivated	Do not reset reconstruction gradients. The gradients are recomputed at the next iteration if they have been zeroed.
Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration. These quantities become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Magnetic Vector Potential Field Functions Reference

The following primitive field functions are made available to the simulation when the Magnetic Vector Potential model is used. These objects all have the same base properties.

Magnetic Vector Potential	See Eqn. (2390).
Magnetic Flux Density	See Eqn. (2362).

Magnetic Field	See Eqn. (2359).
Electric Current Density	See Eqn. (2357).
Permeability	See Eqn. (2359).

When the Temporary Storage Retained property is activated for the Magnetic Vector Potential solver, the following solver-related field functions are made available:

- Magnetic Vector Potential Correction Ap Coefficient
- Magnetic Vector Potential Correction Residual
- Magnetic Vector Potential Gradient
- Magnetic Vector Potential Recon

Using the Transverse Magnetic Potential Model

The Transverse Magnetic Potential model is available for two-dimensional Gas, Solid, and Liquid phases, and can also be combined with the VOF Multiphase model.

This model solves for the transverse magnetic potential according to Eqn. (2390). The permeability can be specified for linear, isotropic materials as a scalar quantity. The magnetic flux density $\mathbf{B}$ (Eqn. (2362)) and the magnetic field $\mathbf{H}$ (Eqn. (2359)) are calculated from the magnetic potential-Z.

Selecting the Transverse Magnetic Potential Model

To make the Transverse Magnetic Potential model available:

1. Select a material model.
2. In the *Space* group box, select Two Dimensional.
3. In the *Optional Models* group box, select Electromagnetism.
4. In the *Electromagnetism* group box, select Transverse Magnetic Potential.

To see the list of selected models, select the Models node from the respective physics continuum.



This model also activates field functions.

Setting an Electric Current Density Region Source

The Transverse Magnetic Potential model activates the Electric Current Density region source. In the Physics Conditions of a region, you can choose to specify the source in the *Properties* window. The Electric Current Density region source is a scalar profile.

To define the electric current density source within a region:

1. Select the Electric Current Density Source Option node within the Physics Conditions folder for the region.
2. In the *Properties* window, set the Electric Current Density Source Option property to Specified.



An additional node, Electric Current Density, is added to the *Physics Values* folder for the region, where you can set its value. The default option uses a constant value. The alternative options are described in the scalar profile properties section.

Setting Region Boundary Conditions

The Transverse Magnetic Potential model activates the Transverse Magnetic Potential region boundary conditions. In the Physics Conditions of the boundary of a region, you can specify the boundary condition type in the *Properties* window. The boundary condition is a scalar profile.

To define the boundary conditions within a region:

1. Select the Magnetic Vector Potential Specification node within the Physics Conditions folder for the region boundaries.
2. In the *Properties* window, set the Method property:

Electric Current Sheet Z	Used for the Neumann boundary condition. This boundary condition specifies the normal derivative of the potentials along the boundary.
Magnetic Vector Potential Z	Used for the Dirichlet boundary condition. This boundary condition is defined when the values of the potentials $\mathbf{A}$ and/or ϕ are explicitly specified on the boundary.

The corresponding node is added to the Physics Values node for the region boundary. This node lets you set the relevant value. The default option uses a constant method. The alternative methods are described in the scalar profile properties section.

The example below shows the Magnetic Vector Potential Z method set for the Transverse Magnetic Potential model



Setting Interface Conditions

On periodic and repeating interfaces you have the option to turn the interface into an anti-periodic interface. The magnetic potential has the opposite sign at each side of an anti-periodic interface. For example, in electrical machine problems, the cross-sectional field analysis can be reduced to either an odd number or an even number of poles by employing anti-periodic or periodic interface conditions respectively.

For periodic interfaces, you can also specify the appropriate periodic contact properties. For more information, see the section *Specifying the Transformation and Axis Properties for Periodic Interfaces* in the *User Guide*.

To set the interface conditions for a repeating or periodic interface:

1. Select the Magnetic Vector Potential Z Periodicity node within the Physics Conditions folder for the interface.
2. In the *Properties* window, set the Method property:

Periodic	The magnetic potential has the same sign at each side of the interface. See Eqn. (2399). This is the default setting.
Anti Periodic	The magnetic potential has the opposite sign at each side of the interface. See Eqn. (2400).

The example below shows the Anti Periodic method set for a periodic interface.



What is the Transverse Magnetic Potential Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the transverse magnetic potential model in all the continua for which the model is activated.

The Transverse Magnetic Potential solver node is shown below.



Transverse Magnetic Potential Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the new solution replaces the old one. The default value is 1.0. ω in Eqn. (2338).

Under-Relaxation Factor

The Transverse Magnetic Potential solver uses standard under-relaxation:

$$A_z^{n+1} = A_z^n + \omega A_z' \tag{2338}$$

where ω is the under-relaxation factor.

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and more memory usage. It is not recommended that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in

STAR-CCM+.

Reconstruction Freeze	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Update reconstruction gradients with each iteration.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero.	
	Activated	Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed.
	Deactivated	Do not reset reconstruction gradients. The gradients are recomputed at the next iteration if they have been zeroed.
Temporary Storage Retained	Retains or deletes temporary storage at the end of the iteration. These quantities become available as field functions during subsequent iterations.	
	Activated	Retain temporary storage.
	Deactivated	Do not retain temporary storage.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.

Transverse Magnetic Potential Field Functions Reference

The following primitive field functions are made available to the simulation when the Transverse Magnetic Potential model is used. These objects all have the same base properties.

Magnetic Vector Potential-Z	See Eqn. (2390). Available for the Transverse Magnetic Potential model only.

Magnetic Flux Density	See Eqn. (2362).
Magnetic Field	See Eqn. (2359).
Electric Current Density-Z	See Eqn. (2357). Available for the Transverse Magnetic Potential model only.
Permeability	See Eqn. (2359).
Stored Magnetic Energy Density	For linear or soft materials which operate in the first quadrant of the B-H characteristic. See Eqn. (2401).
Stored Magnetic Coenergy Density	For linear or soft materials which operate in the first quadrant of the B-H characteristic. See Eqn. (2402).

When the Temporary Storage Retained property is activated for the Transverse Magnetic Potential solver, the following solver-related field functions are made available:

- Magnetic Vector Potential Correction-Z Ap Coefficient
- Magnetic Vector Potential Correction-Z Residual
- Magnetic Vector Potential-Z Gradient
- Magnetic Vector Potential-Z Recon

Setting the Magnetic Permeability

Magnetic permeability μ is a material property that is determined using the magnetic flux density $\mathbf{B}$ within the material and the magnetic field $\mathbf{H}$ ($\mu = \mathbf{B}/\mathbf{H}$). The permeability property appears in the magnetic vector potential equation (Eqn. (2390)). It is defined in the Material Properties node of a physics continuum.

For diamagnetic and paramagnetic materials, the magnetic permeability is assumed to be a constant. For ferromagnetic materials, the magnetization is a non-linear process.

When you set the magnetic permeability of a material in a physics continuum, you have the following options:

- Constant
- Field Function
- Table(B,H)

  You can use the B-H-mu table from the material database, or provide the appropriate B-H curve using a table.

  This method is the recommended table method for setting magnetic permeability.
- tabular

  You provide the appropriate B-H curve using a table.

The first entry in the table must be (0,0). When $\boldsymbol{B}$ becomes greater than the maximum flux density that is provided in the table, the B-H curve is extended at a slope of μ_0 from the last point.

To set the magnetic permeability:

1. Under Continua > [Physics Model] > Models > Solid > [Material] > Material Properties, select Magnetic Permeability.
2. Select the appropriate Method.



The remaining steps are for the Table(B,H) method. For more information on the alternative methods, see Using Constant Values, Using Field Functions, and Using Tabular Data.

3. Select Magnetic Permeability > Table(B,H) > Table(B,H).



4. In the Input table list, select the table to use.



If you want to use the B-H-mu table from the material database, select Material Database Table.

5. Set the Table: Magnetic Field and Table: Magnetic Flux Density properties to specify the appropriate columns of the selected table.



Specifying the Magnetic Permeability of a Material in the Material Database

The constant magnetic permeability property is available for every material in the material database. The nonlinear magnetic permeability property is available for solid materials only.

You specify the nonlinear permeability as a table of $\mathbf{B}$, $\mathbf{H}$ and μ values. To import a permeability table into the material database, the table must contain the following text in the column headings:

- MagneticField
- MagneticFluxDensity
- Permeability

The heading text is not case-sensitive. This requirement applies only to importing a table into the material database. If you import a permeability table as tabular data, there are no restrictions on the column headings.

1. Under Tools > Material Databases, right-click the database material node and select New Database Properties...

2. In the *Select Database Properties* dialog, activate Permeability - Table (B, H, mu).



3. Click OK.

The new Permeability node appears under the material node.



4. Right-click Permeability > Table(B,H,mu) > Material Database Table and select Import Material Table from a file.
5. In the *Open* dialog, select the appropriate table file and click Open.

   The selected table is loaded into the material database. This table is available when you use the Table(B,H) method to set the magnetic permeability property for a solid in a physics continuum.

Modeling Permanent Magnets

This section explains the following:

- What is the Permanent Magnet model?
- Selecting the Permanent Magnet model
- Setting the Remanent Magnetic Flux Density

What Is the Permanent Magnet Model?

The Permanent Magnet model is an optional model under the Magnetic Vector Potential model and the Transverse Magnetic Potential model for a solid material. When the Permanent Magnet model is used, it adds an extra source term to the Magnetic Vector Potential model (see formulation for magnetic vector potential).



Selecting the Permanent Magnet Model

To make the Permanent Magnet model available, ensure that Solid is selected in the *Material* group box before you select the Magnetic Vector Potential model. The Permanent Magnet model appears as an option in the *Optional Models* group box.



The following dialog shows the Permanent Magnet model selected.



Setting the Remanent Magnetic Flux Density

The remanent magnetic flux density is the magnetic field remaining in a permanent magnet after any external magnetic field is removed. You have the option of setting its values in the physics continuum or in a specific region. By default, the remanent magnetic flux density is specified using the physics continuum values.

To specify the remanent magnetic flux density in the physics continuum:

1. Navigate to the Permanent Magnet > Remanent Magnetic Flux Density > Constant node of the physics models.

2. Specify the value in the *Properties* window.



To specify the remanent magnetic flux density for a particular region:

1. Select the Remanent Magnetic Flux Density node within the Physics Conditions folder for the region.

2. In the *Properties* window, set the option to Specify Region Values.



An additional node, Remanent Magnetic Flux Density, is added to the *Physics Values* folder for the region, where you can set the appropriate values. The default option uses a constant value. The alternative options are described in the scalar profile properties section.

Modeling Magnetic Flux Linkage

This section explains the following:

- What is the Excitation Coil Model?
- Selecting the Excitation Coil Model
- Setting up an electric current source and coil region properties
- Formulation
- Field function reference

What Is the Excitation Coil Model?

The Excitation Coil model is an optional model under the Magnetic Vector Potential model and the Transverse Magnetic Potential model. This model is available for a solid material. This model supplies an electric current

source $\mathbf{J}$ to the magnetic vector potential equation Eqn. (2385) and calculates the specific magnetic flux linkage according to Eqn. (2343).



The Excitation Coil model also activates field functions.

Selecting the Excitation Coil Model

The Excitation Coil model is available once the following physics models are activated:

- Solid
- Electromagnetism
- One of the following:
  - Magnetic Vector Potential
  - Transverse Magnetic Potential

The Excitation Coil model appears as an option in the *Optional Models* group box.



Setting Up the Electric Current Source and Coil Region Properties

When you work with the Excitation Coil model, you supply the following properties in the coil region:

- Local Direction Field
- Electric Current Density Magnitude
- Excitation Coil Reference Cross Section or Excitation Coil Reference Height

Setting Up the Local Direction Field

You set the direction of the electric current density $\mathbf{J}$ that is supplied to the magnetic vector potential equation Eqn. (2385). The Local Direction Field node represents this value in the simulation tree.



When you are using the three-dimensional Magnetic Vector Potential model, you set the Local Direction Field using the appropriate vector profile method: the Feature Curve profile.

When you are using the two-dimensional Transverse Magnetic Potential model, the Local Direction Field always follows the Z-axis. You specify the sign by setting the Direction property to the appropriate value: Forward or Backward.

Setting Up the Electric Current Density Magnitude

The magnitude of the electric current $\mathbf{J}$ that is supplied to the magnetic vector potential equation Eqn. (2385) is specified as a scalar profile. The Electrical Current Density Magnitude node represents this value in the simulation tree.



The electric current density magnitude is set using the Ampere Turn method. The electric current I that goes though an individual wire is set using a scalar profile in the Electric Current manager. The number of turns n in the region and the number of sections n_a that cut though the coil region are entered as properties of the Ampere Turn manager.



When you are using the three-dimensional Magnetic Vector Potential model, the magnitude of the electric current density is calculated with respect to the cross section $\mathbf{a}_0$ of the part that is specified in the Excitation Coil Reference Cross Section node. The calculation is:

$$|\mathbf{J}| = \mathrm{J} = \frac{I \cdot n \cdot n_a}{\mathbf{a}_0} \tag{2339}$$

When you are using the two-dimensional Transverse Magnetic Potential model, the magnitude of the electric current density is calculated with respect to the cross section area of the selected region. This cross section area is computed automatically internally.

In order to provide the direct current resistance of the coil, you specify the area that is covered by conducting material under the Conductor Area node.



You can specify the conductor area as an absolute value or as a relative value (a percentage of the total coil area).

Selecting the Excitation Coil Reference Cross Section

When you are using the three-dimensional Magnetic Vector Potential model, you specify the cross section $\mathbf{a}_0$ in Eqn. (2342). You select the appropriate internal interface boundary, or a plane section that cuts through the coil region.

To select this part:

1. Select the Regions > [Coil Region] > Physics Values > Excitation Coil Reference Cross Section node.

2. In the *Properties* window, set Reference Part to the appropriate part.



The area $\mathbf{a}_0$ of the excitation coil reference cross section is annotated in the following scene. Here, the number of sections n_a cutting through the coil region is 1



Selecting the Excitation Coil Reference Height

When you are using the two-dimensional Transverse Magnetic Potential model, the cross section $\mathbf{a}_0$ in Eqn. (2342) is the area of the selected region.

The magnetic flux linkage is calculated from the volume integral of the specific magnetic flux linkage over the region. You specify the height of the region by setting the appropriate Excitation Coil Reference Height.

To set this value:

1. Select the Regions > [Coil Region] > Physics Values > Excitation Coil Reference Height node.

2. In the *Properties* window, set Reference Height to the appropriate value.



Setting Up the Excitation Coil Conductor Resistance Report

The Excitation Coil Conductor Resistance report evaluates the direct current resistance of the excitation coil region. This report is available for unsteady simulations, when you have specified the conductor area of the coil region using the Ampere Turn method.

To set up an Excitation Coil Conductor Resistance report:

1. Right-click the Reports node and then click New Report > Excitation Coil Conductor Resistance.



The new report is added under the Reports node.

2. Specify the regions that you want to include in the report.



Magnetic Flux Linkage Formulation

The magnetic flux linkage ψ that is obtained from the line integral of magnetic vector potential $\mathbf{A}$ around the contour of a coil is:

$$\psi = \int \mathbf{A} \cdot d\mathbf{l} \equiv \psi = \iint \mathbf{B} \cdot d\mathbf{a} \tag{2340}$$

$d\mathbf{l}$ is a vector element of line length and $d\mathbf{a}$ is a vector elemental area of surface.

A coil is regarded as a volumetric domain: for example, a thick wire or a closely wound bundle of a fixed cross section. In this case, $\mathbf{f}\frac{dV}{\mathbf{a}_0}$ replaces the line element $d\mathbf{l}$.

In STAR-CCM+, the magnetic flux linkage is obtained using:

$$\psi = \int \rho_\psi dV = \frac{n_{coil}}{\mathbf{a}_0}\int \mathbf{A}\cdot\mathbf{f}dV \tag{2341}$$

where n_{coil} is the number of turns of the coil and where $\mathbf{f}$ is a local direction vector field with $\|\mathbf{f}\| = 1$.

The area $\mathbf{a}_0$ of the coil cross section is:

$$\mathbf{a}_0 = \int \mathbf{f}\cdot d\mathbf{a} = \int (\mathbf{f}\cdot\mathbf{n}_a)d\mathbf{a} \tag{2342}$$

Therefore the specific magnetic flux linkage is:

$$\rho_\psi = \frac{n_{coil}}{\mathbf{a}_0}\mathbf{A}\cdot\mathbf{f} \tag{2343}$$

Electric Current Density

The electric current density in the coil is:

$$\mathbf{J}_{coil} = \frac{n_a}{\mathbf{a}_0}\mathbf{I}_{coil}\mathbf{f} \tag{2344}$$

where $\mathbf{I}_{coil}$ is the total current in the coil and n_a is the number of sections.

Electromotive Force

The electromotive force is defined as the time derivative of the magnetic flux linkage. Therefore, the specific electromotive force is:

$$\rho_\varepsilon = \frac{d\rho_\psi}{dt} \tag{2345}$$

Direct Current Resistance

The direct current resistance of the excitation coil is obtained by:

$$R = \frac{l}{\sigma a_{coil}} \tag{2346}$$

Here length $l = l_0 n_{coil}$, the area $a_{coil} = \frac{a_0 p_{coil}}{n_{coil} n_c}$, and p_{coil} gives the percentage of the area a_0 that is covered by conducting material.

Magnetic Flux Linkage Field Function Reference

Specific Electromotive Force	See Eqn. (2345).
Specific Magnetic Flux Linkage	See Eqn. (2343).
Local Direction Field	See Eqn. (2343).

Modeling Losses

This section explains the following:

- What is the Modified Steinmetz Model?
- Selecting the Modified Steinmetz Model
- Setting Modified Steinmetz Material Properties
- Setting Up the Modified Steinmetz Monitor
- Modified Steinmetz Field Function Reference
- Modified Steinmetz Formulation

What is the Modified Steinmetz Model?

The Modified Steinmetz model provides an additional monitor called the Modified Steinmetz monitor that calculates hysteresis and eddy-current losses using a modified Steinmetz formula based on magnetic flux density calculations (for more details, see Eqn. (2347)). These loss terms can then be used for a thermal analysis calculation.

Selecting the Modified Steinmetz Model

The Modified Steinmetz model is available once the following physics models are activated:

- Two Dimensional
- Solid
- Implicit Unsteady
- Electromagnetism
- Transverse Magnetic Potential or Magnetic Vector Potential
- Quasi Unsteady Losses

The Modified Steinmetz model appears as an option in the *Optional Models* group box.



Setting Modified Steinmetz Material Properties

When you activate the Modified Steinmetz model, you need to specify the appropriate material properties.



Modified Steinmetz Material Properties

Eddy-Current Loss Coefficient	The eddy-current loss coefficient in $\mathrm{W\,kg^{-1}}$. This is C_e in Eqn. (2347)
Hysteresis Loss Coefficient	The hysteresis loss coefficient in $\mathrm{W\,kg^{-1}}$. This is C_h in Eqn. (2347)
Steinmetz Coefficient A	a in Eqn. (2347)
Steinmetz Coefficient B	b in Eqn. (2347)

Setting Up the Modified Steinmetz Monitor

The Modified Steinmetz monitor becomes available after the Modified Steinmetz model is selected.



Setting the Modified Steinmetz Monitor Properties

When you have added a Modified Steinmetz monitor, you specify the appropriate property values in the *Modified Steinmetz Properties* window.



Editable Modified Steinmetz Monitor Properties

Specify the following properties:

Parts	The parts for which the Modified Steinmetz Monitor is active.
Enabled	Activate the Modified Steinmetz Monitor.
Field Function	The field function for which the peak value and the derivative are calculated. The read-only properties are set accordingly.
Excitation Frequency	The excitation frequency as a scalar physical quantity.

Read-only Modified Steinmetz Monitor Properties

The following properties are read-only and are set according to the field function that you select:

Sample Count	The number of samples collected.
Trigger	The trigger type. This value is set to Time Step by default. See Setting the Time Step Frequency Properties.
Max	The Magnitude Maximum value of the selected field function.
Min	The Magnitude Minimum value of the selected field function.
Peak	The Magnitude Peak value of the selected field function.
Time Derivative	The Magnitude Mean Square Time Derivative value of the selected field function.
Eddy-Current Loss	The Modified Steinmetz Eddy-Current field function.
Hysteresis Loss	The Modified Steinmetz Hysteresis Loss field function.
Total Loss	The Modified Steinmetz Total Loss field function.

Setting the Time Step Frequency Properties

You specify the appropriate property values in the *Time-Step Frequency Properties* window.



Time Step Frequency Properties

Time-Step Frequency	The number of time steps between updates.
Start Time-Step	The time step at which the monitor starts.
Enable Stop	Activates the Stop Time-Step option.
Stop Time-Step	The time step at which the monitor stops.

Modified Steinmetz Field Function Reference

When the Modified Steinmetz model is selected, the following field functions are activated:

Eddy-Current Loss Coefficient	See Eqn. (2347)
Hysteresis Loss Coefficient	See Eqn. (2347)
Magnetic Flux Density Magnitude Maximum	See Eqn. (2350)
Magnetic Flux Density Magnitude Minimum	See Eqn. (2350)
Magnetic Flux Density Magnitude Peak	See Eqn. (2350)
Magnetic Flux Density Magnitude Mean Square Time Derivative	See Eqn. (2351)
Modified Steinmetz Eddy-Current	See Eqn. (2347)
Modified Steinmetz Hysteresis Loss	See Eqn. (2347)
Modified Steinmetz Total Loss	See Eqn. (2347)
Steinmetz Coefficient A	See Eqn. (2347)
Steinmetz Coefficient B	See Eqn. (2347)

Modified Steinmetz Formulation

The core losses W in a cell i consisting of hysteresis and eddy-current losses are calculated using the modified Steinmetz formula:

$$W = C_h f B_{pk}^n + \frac{C_e}{2\pi^2}\overline{\left[\frac{dB}{dt}\right]^2} \qquad \left[\frac{W}{kg}\right] \tag{2347}$$

where the losses by hysteresis is represented by:

$$C_h f B_{pk}^n \tag{2348}$$

where:

- $n = a + bB_{pk}$, where a and b are the Steinmetz Coefficients
- C_h is the hysteresis loss coefficient
- f is the frequency of excitation

and the losses by eddy-currents are represented by:

$$\frac{C_e}{2\pi^2}\overline{\left[\frac{dB}{dt}\right]^2} \tag{2349}$$

where:

- C_e is the eddy-current loss coefficient

The peak magnetic flux density is:

$$B_{pk}\Big|_i = \frac{1}{2}\left(\max_{k = 1, N}|B|_i^k - \min_{k = 1, N}|B|_i^k\right) \tag{2350}$$

The average time derivative of magnetic flux density is:

$$\overline{\left[\frac{dB}{dt}\right]^2}\Bigg|_i = \frac{1}{N}\sum_{k=1}^{N}\left(\frac{|B|_i^k - |B|_i^{k-1}}{t\delta}\right) \tag{2351}$$

where:

- N is the rotor position
- $t\delta$ is the time step.
- The subscript i denotes a certain mesh cell.
- k denotes B at a certain time.

Modeling Electrostatics

This section explains the following:

- What is the Electrostatic Potential model?
- Region conditions
- Boundary conditions
- Field functions

What Is the Electrostatic Potential Model?

The Electrostatic Potential model, represented by the Electrostatic Potential node, is available for Gas, Solid and Liquid phases.

The Electrostatic Potential model is compatible with the Multi-Component Gas and Multi-Component Liquid models, although the permittivity that you specify is the effective permittivity of the mixture. STAR-CCM+ does not let you specify the permittivity as a function of the permittivities of the components.

It solves for the electric potential according to Eqn. (2370). From the electric potential the electric field $\mathbf{E}$ (Eqn. (2361)) and the electric flux density $\mathbf{D}$ (Eqn. (2358)) are calculated.

Selecting the Electrostatic Potential Model

To make the Electrostatic Potential model available:

- Select a space model.
- Select a material model. The Electromagnetism model will become available in the *Optional Models* group box.
- Select Electromagnetism in the *Optional Models* group box. The Electrostatic Potential model will become available in the *Electromagnetism* group box.
- Select Electrostatic Potential in the *Electromagnetism Models* group box.
- Close the dialog.

The list of selected models can be seen by selecting the Models node from the respective physics continuum



Setting an Electric Charge Density Region Source

The Electrostatic Potential model actives the Electric Charge Density region source. In the Physics Conditions of a region, you can choose to specify the source in the *Properties* window.

To define the electric current density source within a region:

- Select the Electric Current Density Source Option node within the Physics Conditions folder for the region.

- In the *Properties* window, tick the checkbox for the Electric Charge Density User Source.



An additional node, Electric Charge Density User Source, will be added to the *Physics Values* folder for the region, where you can set its value. The default option uses a constant value. The alternative options are described in the scalar profile properties section.

Setting Region Boundary Conditions

The Electrostatic Potential model actives the Electrostatic Potential region boundary conditions. In the Physics Conditions of a region's boundary, you can choose to specify the boundary condition type in the *Properties* window.

To define the boundary conditions within a region:

- Select the Electrostatics Potential Specification node within the Physics Conditions folder for the region boundaries.

- In the *Properties* window, set the Method property to either Specific Electric Flux or total Electric Flux for the Neumann boundary condition or Electric Potential for the Dirichlet boundary condition.



An additional node, Specific Electric Flux, Electric Flux or Electric Potential, will be added to the *Physics Values* folder for the region boundary, where you can set this value. The default option uses a constant value. The alternative options are described in the scalar profile properties section.

Note that as with boundaries, you can also set a specific electric flux or total electric flux value on an interface using the Electric Flux Source Option node.



Electrostatic Potential Field Functions Reference

The following primitive field functions are made available to the simulation when the Electrostatic Potential model is used. These objects all have the same base properties.

Electric Flux Density	See Eqn. (2358).
Permittivity	See Eqn. (2358) or Eqn. (2370).

What Is the Electric Potential Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the Electromagnetism models in all the continua for which this model type is activated.



Electric Potential Solver Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the old solution is supplanted by the newly computed solution. The default value (1.0) means that no under-relaxation is used. ω in Eqn. (2352).

Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and more memory usage. It is not recommended that you change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Temporary Storage Retained	Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.
	Activated — Retain temporary storage.
	Deactivated — Do not retain temporary storage.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. Activating this feature should not be your first choice, as it is primarily a "debugging" tool and not supported for normal simulation work.

Activated — The solver does not update the solution.

Deactivated — The solver updates the solution.

Electric Potential Under Relaxation Factor

The electric potential solver uses standard under-relaxation

$$\phi^{n+1} = \phi^{n} + \omega\phi' \tag{2352}$$

where ω is the under-relaxation factor.

Accessing Electromagnetism Constants

The following table lists the electromagnetism constants accessible through the Java API:

star.electromagnetism.commom.ElectromagnetismConstants

public static final double	AVOGADRO_CONSTANT	6.022136736E23 1/kmol
public static final double	ELECTRIC_CONSTANT (permittivity)	8.854187817620389E-12 F/m
public static final double	ELEMENTARY_CHARGE	1.602176487E-19 C
public static final double	FARADAY_CONSTANT	9.648525879918127E7 C/kmol
public static final double	MAGNETIC_CONSTANT (permeability)	1.2566370614359173E-6 H/m
public static final double	SPEED_LIGHT	2.99792458E8 m/s

These constants can be invoked inside Java macros. For example, the following command can be used to access the value of the speed of light in a vacuum:

```
double speedOfLight =
star.electromagnetism.common.ElectromagnetismConstants.SPEED
_LIGHT;
```

Electromagnetism Formulation

This section contains the following:

- Fundamental equations
- Constitutive relations
- Scalar and vector potentials
- Electrostatic and magnetostatic formulations
- Low frequency electromagnetic formulation
- Boundary and interface conditions
- Stored magnetic energy formulation
- Troubleshooting
- Nomenclature

- Bibliography

Fundamental Equations

The fundamental equations governing all macroscopic electromagnetic phenomena consist of four Maxwell's equations,

- Faraday's law of induction:

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla \times \mathbf{E} = 0 \tag{2353}$$

- Maxwell-Ampere law:

$$\frac{\partial \mathbf{D}}{\partial t} - \nabla \times \mathbf{H} = -\mathbf{J} \tag{2354}$$

- Gauss' law for electricity:

$$\nabla \cdot \mathbf{D} = \rho \tag{2355}$$

- Gauss' law for magnetism:

$$\nabla \cdot \mathbf{B} = 0 \tag{2356}$$

and the conservation of electric charge, known as the *continuity equation*, which is obtained by taking the divergence of Eqn. (2354) and using Eqn. (2355)

$$\nabla \cdot \mathbf{J} + \frac{\partial \rho}{\partial t} = 0 \tag{2357}$$

where:

t is time;

$\mathbf{E}$ is the electric field intensity;

$\mathbf{D}$ is the electric flux density;

$\mathbf{H}$ is the magnetic field intensity;

$\mathbf{B}$ is the magnetic flux density;

$\mathbf{J}$ is the electric current density; and

ρ is the electric charge density.

Only three of the above five equations are independent -- either the first three Eqn. (2353)-(2355), or the first two Eqn. (2353) and (2354) with (2357) are a set of independent equations. The remaining two, Eqn. (2356) and (2357) or (2355) and (2356), can be derived from the independent equations.

Constitutive Relations

There are three independent equations out of five fundamental equations that are valid for any continuum. However, they do not form a closed set of equations as the number of equations (three) is less than the number of unknowns (five). Using the following constitutive relations, which describe properties of the medium being considered, you can solve for the unknowns

$$\mathbf{D} = \varepsilon \mathbf{E} \tag{2358}$$

$$\mathbf{B} = \mu \mathbf{H} \tag{2359}$$

$$\mathbf{J} = \sigma \mathbf{E} \tag{2360}$$

where ε is the permittivity, μ is the permeability and σ is the conductivity. These parameters are scalars for homogeneous isotropic media, functions of position for inhomogeneous media and tensors for anisotropic media.

Eqn. (2360) is known as Ohm's law. In Eqn. (2360), the direction of the electric current density is from the areas of high to the areas of low values of the electric potential, which is known as the technical direction of electric current density. However, if electric current density is considered as the flow of electrons through conductive materials, the direction of this flow is opposite to the technical direction as defined above. This direction is known as the physical direction of the electric current. In STAR-CCM+, the technical direction of electric current density is always used.

Scalar and Vector Potentials

Electromagnetic problems are often expressed in terms of the electric scalar potential ϕ and the magnetic vector potential $\mathbf{A}$:

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t} \tag{2361}$$

$$\mathbf{B} = \nabla \times \mathbf{A} \tag{2362}$$

To uniquely determine the magnetic vector potential, gauge conditions are introduced. For the low frequency formulation, the Coulomb gauge is used:

$$\nabla \cdot \mathbf{A} = 0 \tag{2363}$$

The transport equation for the electric potential ϕ is

$$-\oint_A \sigma \nabla\phi \, d\mathbf{a} - \oint_A \sigma \frac{\partial \mathbf{A}}{\partial t} d\mathbf{a} = \int_S S \, dV \tag{2364}$$

where σ is the electrical conductivity. $S = J_{tc}$ is the volumetric source term [A/m^3] that is called transfer current density.

Caution: When the transfer current density is applied externally to account for unresolved physics, the applied source terms must be in balance within the solution domain. Alternatively, there must be at least one boundary patch on the solution domain with a Dirichlet boundary condition (prescribed voltage) for electric potential. This condition is to avoid violating the incompressible nature of the electric current.

Electrostatic and Magnetostatic Formulations

For fields that do not vary with time, there is no interaction between magnetic and electric fields. This independence can be seen by rewriting Eqn. (2353), (2354) and (2357) as:

$$\nabla \times \mathbf{E} = 0 \tag{2365}$$

$$\nabla \times \mathbf{H} = \mathbf{J} \tag{2366}$$

$$\nabla \cdot \mathbf{J} = 0 \tag{2367}$$

Eqn. (2366) is known as Ampere's law.

Electrostatic Fields

Electrostatics introduces a new fundamental quantity called electric charge.

Electrostatic fields are described in Eqn. (2355) and (2365). By using the constitutive relation Eqn. (2358), Eqn. (2355) can be rewritten as:

$$\nabla \cdot (\varepsilon \mathbf{E}) = \rho \tag{2368}$$

where $\varepsilon = \varepsilon_r \varepsilon_0$. Here, ε_r is called relative permittivity and ε_0 is the vacuum permittivity with the value of:

$$\varepsilon_0 = 8.8543 \times 10^{-12} \quad \frac{F}{m} \tag{2369}$$

Eqn. (2368) can also be written in terms of the electric scalar potential Eqn. (2361), with the time derivative equal to zero

$$-\nabla \varepsilon \nabla \phi = \rho \tag{2370}$$

Eqn. (2365) is always satisfied due to the vector identity $\nabla \times (\nabla \phi) = 0$. Eqn. (2370) defines the electrostatic mathematical model expressed in terms of the electric potential equation.

The electric field in an electrostatic formulation is given by the force acting on an electric charge q:

$$\mathbf{E} = \frac{\mathbf{F(r)}}{q} \tag{2371}$$

The electric charge has a quantized, atomic structure and appears only in multiples of the elementary electric charge, which is essentially the electric charge of an electron (with the negative sign).

$$e = 1.602176487 \times 10^{-19} \ C \tag{2372}$$

Magnetostatic Fields

Magnetostatic fields are described in Eqn. (2356), (2366). By using the constitutive relation Eqn. (2359) and the magnetic vector potential Eqn. (2362), Eqn. (2366) can be rewritten as:

$$\nabla \times \left(\frac{1}{\mu} (\nabla \times \mathbf{A}) \right) = \mathbf{J} \tag{2373}$$

where $\mu = \mu_r \mu_0$. Here, the dimensionless quantity μ_r is called relative permeability and μ_0 is the vacuum permeability or magnetic constant with the value of:

$$\mu_0 = 4\pi \cdot 10^{-7} \ \frac{H}{m} \tag{2374}$$

Using vector analysis, you can rewrite Eqn. (2373) as:

$$\nabla \cdot \left(\frac{1}{\mu} \nabla \mathbf{A} \right) = -\mathbf{J} + \mathbf{C}_\mu \tag{2375}$$

where the last term is the correction due to the variable magnetic permeability μ given by:

$$\mathbf{C}_\mu = \nabla \mathbf{A} \cdot \nabla \frac{1}{\mu} \tag{2376}$$

$\mathbf{C}_\mu$ is zero when μ is constant. Also, Eqn. (2356) is always satisfied due to $\nabla \cdot (\nabla \times \mathbf{A}) = 0$. Eqn. (2375) defines the magnetostatic mathematical model expressed in terms of the vector magnetic potential.

Low Frequency Electromagnetic Formulation

The low frequency formulation of Maxwell equations assumes that the conduction current $\mathbf{J}$ in the Maxwell-Ampere law Eqn. (2354) is many orders of magnitude larger than the displacement current $\frac{\partial \mathbf{D}}{\partial t}$. Therefore, the latter time derivative term can be omitted from the equation. This assumption reduces the Maxwell-Ampere law to Ampere's law Eqn. (2366).

The low frequency formulation is valid for many industrial applications, such as electric machines, transformers, or electric switches.

Electric Scalar Potential Equation

By expressing Ohm's law Eqn. (2360) in terms of the electric scalar and magnetic vector potentials, Eqn. (2361) becomes:

$$\mathbf{J} = \sigma\mathbf{E} = -\sigma\left(\nabla\phi + \frac{\partial \mathbf{A}}{\partial t}\right) \tag{2377}$$

Furthermore, for electric currents within conducting media the electric current density is incompressible in nature, simplifying Eqn. (2357) to:

$$\nabla \cdot \mathbf{J} = 0 \tag{2378}$$

By combining Eqn. (2377) and (2378), and using Coulomb gauge Eqn. (2363), the electrodynamic formulation of the electric potential equation is derived as:

$$\nabla\sigma\nabla\phi + \nabla\left(\sigma\frac{\partial \mathbf{A}}{\partial t}\right) = 0 \tag{2379}$$

For constant conductivity σ, the second term in Eqn. (2379) vanishes.

Additionally, the electric current I is defined as flux of electric current density through surface A:

$$I = \int_A \mathbf{J} d\mathbf{a} \tag{2380}$$

Joule's Law defines the heat that is generated by electric currents flowing in conducting media. The density of the power that is generated in this way is:

$$\frac{P}{d^3 r} = \mathbf{J} \cdot \mathbf{E} \tag{2381}$$

Therefore, the total power that is dissipated from the volume of conducting media where an electric current exists is:

$$P = \int_v \mathbf{J} \cdot \mathbf{E} \, \mathrm{d}V \tag{2382}$$

Electrical Conductivity in Porous Regions

Porous regions, such as fuel cell catalyst layers and battery electrodes, are modeled as volumetric regions where there is a separate electronic potential and ionic potential.

The effective electrical conductivity tensor σ_{eff} for a porous region is computed as:

$$\sigma_{eff} = \chi\sigma_{fluid} + (1-\chi)\sigma_{solid} \tag{2383}$$

where:

- χ is the porosity of the porous region
- σ_{fluid} is the scalar conductivity specified for the fluid continuum
- σ_{solid} is the electrical conductivity tensor for the solid part of the porous region. This value is specified in the Physics Values node for the region.

Note that the formulation for effective electrical conductivity does not include contributions from tortuosity and constrictivity, which vary depending on the geometric properties of the porous media. In such cases where you have accounted for these factors and have already computed an effective electrical conductivity that you want to use, the effective electrical conductivity can be obtained by modifying the Solid Electrical Conductivity. In this case, you can compute the solid electrical conductivity profile by setting:

$$\sigma_{solid} = \frac{\sigma_{eff} - \chi\sigma_{fluid}}{1-\chi}$$

When the Electrodynamics model is active in a porous region, the Ohmic Heating model is also active in the region. The ohmic heating source term is unchanged: in all regions the ohmic heating source term is computed as:

$$S_{T,\,ohm} = \mathbf{J} \cdot \mathbf{E} \tag{2384}$$

where $S_{T,\,ohm}$ is the ohmic heating source term, $\mathbf{J}$ is the electric current vector, and $\mathbf{E}$ is the electric field vector.

Magnetic Vector Potential Equation

By using Eqn. (2377), you can rewrite Eqn. (2373) as:

$$\sigma\frac{\partial \mathbf{A}}{\partial t} = \nabla \cdot \left(\frac{1}{\mu}\nabla \mathbf{A}\right) - \sigma\nabla\phi - \mathbf{C}_\mu + \mathbf{S} \tag{2385}$$

where term $\mathbf{S}$ can contain an external current $\mathbf{J}_{ex}$, which is the contribution from moving parts:

$$\mathbf{J}_{mp} = \sigma \mathbf{v}_p \times \mathbf{B} = \sigma \mathbf{v}_p \times \nabla \mathbf{A} \tag{2386}$$

where $\mathbf{v}_p$ is the part velocity from the permanent magnet. The permanent magnet current density is:

$$\mathbf{J}_{mp} = \nabla \times \frac{\mathbf{B}_r}{\mu} \tag{2387}$$

where $\mathbf{B}_r$ is the remanent magnetic flux density.

The magnetic flux density $\mathbf{B}$ is given by Eqn. (2362).

The magnetic field is:

$$\mathbf{H} = \frac{1}{\mu}(\mathbf{B} - \mathbf{B}_r) \tag{2388}$$

Eqn. (2385), together with Eqn. (2379), determine the low frequency electromagnetic mathematical model expressed in terms of the magnetic vector and electric scalar potentials.

If you select the Quasi Unsteady Losses model for an implicit unsteady simulation, the unsteady term in Eqn. (2385) is ignored ($\frac{\partial \mathbf{A}}{\partial t} = 0$) and a steady solution is computed for the Magnetic Vector Potential.

Integral Form

Eqn. (2379) and (2385) can be rewritten in integral form that is suitable for Finite Volume discretization:

$$-\oint_A \sigma \nabla \phi \, d\mathbf{a} - \oint_A \sigma \frac{\partial \mathbf{A}}{\partial t} d\mathbf{a} = \int_V \mathbf{S} \, dV \tag{2389}$$

$$\int_V \sigma \frac{\partial \mathbf{A}}{\partial t} dV - \oint_A \frac{1}{\mu} \nabla \mathbf{A} \cdot d\mathbf{a} = -\int_V \sigma \nabla \phi \, dV - \int_V \mathbf{C}_\mu \, dV + \int_V \mathbf{S} \, dV \tag{2390}$$

where V is an arbitrary part of the continuum bounded by the surface A with the surface vector $d\mathbf{a}$.

In Eqn. (2389), a volumetric source term is called *transfer current density* and can be applied externally to account for unresolved physics. By doing this, extra care must be taken not to violate the incompressible nature of the electric current. This condition implies that the so-applied source terms

must be in balance within the solution domain. Alternatively, there must be at least one boundary patch on the solution domain with the Dirichlet boundary condition (prescribed voltage) for electric potential.

Eqn. (2389) and (2390) are a set of mutually coupled equations with two unknown fields, $\mathbf{A}$ and ϕ. After obtaining the potential fields, the original five electromagnetic fields can be calculated from the constitutive relations and Eqn. (2361) and (2362).

Transverse Magnetic Potential Equation

In transverse magnetic mode, the magnetic field lies in a plane and the magnetic vector potential has only one component perpendicular to that plane. Let that plane be the x-y plane, and Eqn. (2390) simplifies into a scalar transport equation for A_z:

$$\int_V \sigma \frac{\partial A_z}{\partial t} - \int_A \frac{1}{\mu} \nabla A_z \cdot d\boldsymbol{a} = \int_V S_z dV \tag{2391}$$

S_z is as defined in Eqn. (2385): it has the same physical contributions but only for the scalar z component.

The magnetic flux density $\mathbf{B}$ is given by Eqn. (2362) with A_x and A_y always zero.

The magnetic field is:

$$\mathbf{H} = \frac{1}{\mu}(\mathbf{B} - \mathbf{B}_r) \tag{2392}$$

If eddy-current carrying regions are modeled, the resulting current from Eqn. (2391) is scaled such that the zero net current condition:

$$\int \mathbf{J} d\boldsymbol{a} = 0 \tag{2393}$$

is fulfilled for each region separately.

Boundary and Interface Conditions

For a problem to be well-posed, boundary conditions must be specified on all solution domain boundaries. A boundary condition can be Dirichlet, Neumann, tangential magnetic field, or symmetry plane. You can also specify an interface condition for a contact interface. The following describes the available boundary and interface conditions.

Dirichlet Boundary Conditions

A Dirichlet boundary condition is defined when the values of the potentials $\mathbf{A}$ and/or ϕ are explicitly specified on the boundary.

Specifying the Electric Potential and Electric Resistance

Define an electric potential value ϕ_∞ and set $\phi_f = \phi_\infty$ on the boundary face.

Additionally, the Dirichlet conditions can be seen as a series connection of resistances. The following relationship is implemented:

$$\phi_f = \phi_0 - \frac{\phi_0 - \phi_\infty}{R^0 + R^{\text{User}} + \sum_i R_i^{\text{Models}}} R^0 \tag{2394}$$

Here ϕ_f is the boundary face centroid value and ϕ_0 is the potential value of the face neighboring cell centroid. R^0 is approximated by (δ/σ) where δ is:

$$\mathbf{s} \cdot \frac{\boldsymbol{a}}{|\boldsymbol{a}|}$$

where $\boldsymbol{a}$ is the face area vector, $\boldsymbol{s}$ is the vector between the cell face and the cell centroids, and σ is the electrical conductivity of the phase.

If the *Electric Resistance* $R^{\text{User}} = 0$ then Eqn. (2394) reduces to a basic Dirichlet boundary condition that specifies the potential value on the boundary, $\phi_f = \phi_\infty$. The dimensions of the user-specified electric resistance are $\Omega m^2 = Vm^2/A$. Other models can also contribute to the Dirichlet boundary by an additional Electric Resistance R^{Model}. These terms are documented together with the corresponding models.

It is possible to specify the derivative of the electrical resistance with respect to the electric potential ϕ_0 in equation Eqn. (2394), that is, $\partial R/\partial\phi_0$. The derivative of the user-defined electrical resistance with respect to the electric potential is used for linearization. Its value is set to zero by default. Providing a value for the derivative helps to stabilize the solution and can speed up convergence of the solution.

The Electrical Resistance Potential Derivative property is available within the Electrodynamic Potential model as a Dirichlet boundary condition. It is entered as a scalar profile.

Specifying the Magnetic Vector Potential

The following boundary conditions are supported:

- Symmetry - Perfect Magnetic Conductor

The normal component of the magnetic vector potential value $\mathbf{A}_{\infty}\big|_{n}$ is set to zero. The tangential components $(\mathbf{A}\big|_{t_1,\, t_2})$ are left free.

- Anti-Symmetry - Perfect Electric Conductor

  The tangential components of the magnetic vector potential value $\mathbf{A}_{\infty}\big|_{t_1,\, t_2}$ are specified and set on the boundary face:

$$\mathbf{A}_{f}\big|_{t_1,\, t_2} = \mathbf{A}_{\infty} = 0$$

  Normally you would set $\mathbf{A} = 0$ along the boundary to keep the magnetic flux from crossing the boundary. The normal component $\mathbf{A}\big|_{n}$ is left free.

- Imposed Magnetic Vector Potential

  When a strong external magnetic field is applied, such that the magnetic vector potential can written as:

$$\mathbf{A} = \mathbf{A}_{\text{imp}} + \mathbf{A}_{\text{ind}}$$

  where

  - $\mathbf{A}_{\text{imp}}$ denotes the externally imposed field
  - $\mathbf{A}_{\text{ind}}$ denotes the additional magnetic field induced by electric currents

  all the components of the magnetic vector potential at the boundary are set explicitly. That is:

$$\mathbf{A}\big|_{x,\, y,\, z} = (\mathbf{A}_x, \mathbf{A}_y, \mathbf{A}_z)$$

  STAR-CCM+ assumes that the induced magnetic vector potential $(\mathbf{A}_{\text{ind}})$ in the domain decays sufficiently quickly to be neglected at the boundary. No electric current should flow through this boundary.

For information on setting these boundary conditions, see Setting Region Boundary Conditions.

Specifying the Transverse Magnetic Potential

The magnetic vector potential-Z value $\mathbf{A}_{z\infty}$ is specified and set on the boundary face $\mathbf{A}_{zf} = \mathbf{A}_{z\infty}$. Normally you would set $\mathbf{A}_z = 0$ along the boundary to keep the magnetic flux from crossing the boundary.

Tangential Magnetic Field Boundary Conditions

Since the diffusion term is calculated at faces as:

$$\nu[(\nabla\mathbf{A}) - (\nabla\mathbf{A})^T] \cdot \mathbf{a} = \mathbf{H} \times \mathbf{a} \tag{2395}$$

the tangential magnetic field boundary condition assigns the cross-product of the specified magnetic field $\mathbf{H}$ with the face area vector $\mathbf{a}$ to the diffusion term associated with this boundary face.

Neumann Boundary Conditions

A Neumann boundary condition specifies the normal derivative of the potentials along the boundary.

Specifying the Specific Electric Current and Electric Flux

Specify the flux (specific electric current) at a cell boundary face. The electric current density can be specified as a scalar profile or a vector profile. By default, the specific electric current J_n at the boundary is set to zero (the Insulator boundary condition). No electric current penetrates the boundaries where this condition is selected.

STAR-CCM+ uses the convention for specifying the boundary specific electric current on boundaries, as illustrated in the following figure. A negative flux value specified at a wall boundary face means that the electric current density goes in the direction of the face normal vector $(J_n = -\mathbf{J} \cdot \mathbf{n} < 0)$. To get a unique solution, a Dirichlet boundary condition must be used somewhere.



When modeling electrostatic fields, specify the specific electric flux D_n instead of the specific electric current J_n. Note that the flux values that are provided by the corresponding field function BoundarySpecificElectricCurrent have the opposite sign.

In STAR-CCM+ versions previous to 6.04, an opposite convention was used. If a simulation written by versions previous to 6.04 is loaded into version 6.04 or newer, you must change the specific electric current scalar profile accordingly.

Specifying the Electric Current

Define the total current on a boundary from Eqn. (2380). The sign convention is the same as for the specific electric current.

Specifying the Electric Current Sheet

Specify the tangential components of the electric current sheet on the boundary. For example, at an interface with a highly permeable metal, to force the magnetic flux to cross the boundary at an angle of $90°$, you specify $J_s = 0$.

The normal component of the current sheet is always treated as zero.

Specifying the Transverse Electric Current Sheet

The electric current sheet-Z condition specifies the normal derivative of potential along the boundary. For example, at an interface with a highly permeable metal, to force the magnetic flux to cross the boundary at a $90°$ angle, you would set:

$$\frac{\partial \mathbf{A}_z}{\partial n} = 0$$

Interface Conditions

As with boundaries, the total and specific electric current and flux can be set on contact interfaces. When setting electric current surface sources, the sign convention is the same as for Neumann boundaries, except that the sign of the sources is always chosen with respect to boundary-0.

Electric Resistance

Additionally, Electric Resistance can be specified on the contact interface. The potential values are:

$$\phi_{f0} = \phi_0 - \frac{\phi_0 - \phi_1}{R^0 + R^{\text{User}} + R^1} R^0 = \phi_0 - J_{n_0} R^0 \tag{2396}$$

$$\phi_{f1} = \phi_1 - \frac{\phi_1 - \phi_0}{R^0 + R^{\text{User}} + R^1} R^1 = \phi_1 - J_{n_1} R^1 \tag{2397}$$

On wall boundaries with an Electric Resistance, the power loss:

$$P = J_n^2 R^{\text{User}} \tag{2398}$$

makes a contribution to the face neighboring cell center source term of the energy equation. On contact interfaces, the contribution to the cell center source term of the energy equation is weighted by the effective thermal conductivities of the neighboring physics continua.

Transverse Magnetic Vector Potential

On periodic and repeating interfaces there is an option to turn the interface into an anti-periodic interface.

The magnetic potential has the opposite sign at each side of an anti-periodic interface. For example, in electrical machine problems, the cross-sectional field analysis can be reduced to either an odd number or an even number of poles by employing anti-periodic or periodic interface conditions respectively.

In such configurations, the periodic interface condition can be written as:

$$A_z(r,\theta) = A_z(r,\theta + \frac{2k\pi}{p}),\ k = 1, 2, 3... \tag{2399}$$

and the anti-periodic interface condition can be written as:

$$A_z(r,\theta) = -A_z(r,\theta + \frac{2(k-1)\pi}{p}),\ k = 1, 2, 3... \tag{2400}$$

where p is the even or odd number of pole pairs and r, θ are polar coordinates.

Stored Magnetic Energy

For linear or soft materials which operate in the first quadrant of the B-H characteristic, the stored magnetic energy density $\rho_{w_m}(B)$ and the stored magnetic coenergy density $\rho_{w'_m}(B)$ are given by:

$$\rho_{w_m}(B) = \int_0^B H'(B')\mathrm{d}B' \tag{2401}$$

and

$$\rho_{w'_m}(B) = \int_0^H B'(H')\mathrm{d}H' \tag{2402}$$

Troubleshooting

Meshing recommendation

The use of an extruder mesh at the boundaries can improve the results. Additional guidelines on mesh quality are available in the chapter on checking the mesh.

Electric potential solution converges slowly

The under-relaxation factor can be adjusted to control the convergence rate. The default Under-Relaxation factor is set to 1.0 and should be kept at this value whenever possible. However, if the magnetic vector potential solution does not converge, a reduction of the under-relaxation factor can stabilize the convergence procedure.

Magnetic vector potential solution does not converge with the given B-H curve

If the magnetic permeability is set using the tabular method and the provided B-H curve is monotonous and convex $\left(\frac{dB}{dH} > 0 \text{ and } \frac{d^2B}{dH^2} < 0\right)$, then the magnetic field $\mathbf{H}$ must be set as the interpolating variable in the method. On the other hand, if the curve is concave, the magnetic flux density $\mathbf{B}$ must be the interpolating variable.

Energy solution is diverging during initial iterations

The Specific Ohmic Heat Source (Eqn. (2381)) depends quadratically on the electric potential gradients. Assuming the problem is well-posed then nevertheless these gradients can be significantly wrong at the beginning of the iteration process and produce a large source term dominating the energy equation. What to look for:

- Examine the critical areas where the gradients of the electric potential are large (field functions electric field or ohmic heat source term). Maybe mesh quality can be improved in these areas.
- Reduce the maximum temperature.
- Decouple the simulation by freezing the energy solver and getting a converged solution for the electric field first and continue afterwards with an active energy solver.

See also the Troubleshooting sections for the energy equations (coupled and segregated).

A message appears about maximum temperature

`Max. Temperature limited to ...` — This message does not necessarily mean that the solution is diverging. See the previous note about what to look for.

Electric potential solution is diverging for Electro-Deposition Coating

What to look at in electro-deposition coating:

- Reduce the time-step size.
- A reduction of the under-relaxation factor for the Electro-Deposition Coating Solver can stabilize the convergence procedure even without changing the under-relaxation factor of the Electric Potential Solver. On the other hand, a higher under-relaxation factor can accelerate the convergence rate.

Electromagnetism Nomenclature

The following is a guide to the terms used in this section.

Variable	Definition
ϕ	electric potential $[V]$
ε	permittivity $[A^2 s^2/Nm^2]$
ε_0	vacuum permittivity, permittivity of free space or electric constant $[A^2 s^2/Nm^2]$
ε_r	relative permittivity
μ	magnetic permeability $[H/m]$
μ_0	vacuum permeability, permeability of free space or magnetic constant $[H/m]$
μ_r	relative permeability
ν	magnetic reluctivity
ρ	charge density $[C/m^3]$
σ	electrical conductivity $[S/m]$
ψ	magnetic flux linkage $[Wb]$
ρ_ψ	specific magnetic flux linkage $[Wb/m^3]$
$\mathbf{A}$	magnetic vector potential $[Wb/m]$
$\mathbf{B}$	magnetic flux density, magnetic induction $[T]$
$\mathbf{D}$	electric flux density, electric displacement field, electric induction $[C/m^2]$
$\mathbf{E}$	electric field, electric field intensity $[V/m]$
$\mathbf{H}$	magnetic field, magnetic field intensity $[A/m]$
$\mathbf{J}$	electric current density $[A/m^2]$
D_n	specific electric flux $[C/m^2]$
E	electrochemical equivalent $[kg/C]$
F	Faraday constant $[c/kmol]$

J_n	specific electric current $[A/m^2]$
J_{tc}	transfer current density $[A/m^3]$
N_A	Avogadro constant
P	power $[W]$
c_0	speed of light in vacuum $[m/s]$
e	elementary electric charge, charge of an electron $[C]$
t	time $[s]$
r	electrical resistivity $[\Omega m]$

Electromagnetism Bibliography

[409] Ida, Nathan. 2003. *Engineering Electromagnetics*, 2nd ed., Springer Publishing Company, Inc.

[410] Jackson, John David. 1999. *Classical Electrodynamics*, 3rd ed., Wiley, New York, NY.

[411] Streitberger, Hans Joachim and Doessel, Karl-Friedrich, Eds., 2008. *Automotive Paints and Coatings*, 2nd ed., Wiley-VCH.

Electromagnetism Field Functions Reference

The following primitive field functions are made available to the simulation when electromagnetism models are used. These objects all have the same base properties.

Electric Potential	See Eqn. (1448).
Electric Field	See Eqn. (1448).

Solver-related field functions, common for STAR-CCM+ solvers when Temporary Storage retained is active, are:

- Electric Potential Correction Ap Coefficient
- Electric Potential Correction Residual
- Electric Potential Gradient
- Electric Potential Recon
- Magnetic Vector Potential Correction Ap Coefficient
- Magnetic Vector Potential Correction Residual
- Magnetic Vector Potential Gradient
- Magnetic Vector Potential Recon

Modeling Electrochemistry

Subsequent sections provide the following information:

- Modeling Li-Ion Battery Cells
- Modeling Corrosion
- Modeling Ionic Species Flux
- Electrochemistry nomenclature

Modeling Li-Ion Battery Cells

STAR-CCM+ offers simulation tools that help you to understand the properties and behavior of Li-ion electrode materials. Since the electrode active materials—especially the cathode active materials—are usually the most expensive components in a Li-ion battery cell, understanding their properties and behavior is crucial for cell manufacturers.

The ability to model the micro-structure of an electrode and predict its electrochemical reactions help to determine the quality and performance of a cell design. For example, it is possible to determine how tortuous a porous electrode can be, and analyze or parameterize particle size and distribution, both key factors for an electrode.

Subsequent sections provide the following information:

- What is the Li-Ion Battery Cell model?
- What is the Li-Ion Electric Potential model?
- What is the Li-Ion Concentration model?
- What is the Solid Electrolyte Interface (SEI)?
- Setting up a Li-ion battery cell with a lumped electrode
- Reporting Li-ion cell parameters
- Troubleshooting Li-ion battery cell simulations

What Is the Li-Ion Battery Cell Model?

The Li-Ion Battery Cell model models lithium-ion battery cells at the microscale level. The following annotated diagram depicts a typical Li-Ion cell. During the discharge process, lithium de-intercalates from the anode active material where it dissociates at the Solid Electrolyte Interface (SEI)

into a lithium cation and an electron. The Li^+ ion migrates through the separator and is inserted into the cathode active material. Charging is the reverse process.



Note: The Li-Ion Battery Cell model is available with the double precision version of STAR-CCM+ only.

The Li-Ion Battery Cell model serves as a gateway between two Li-Ion physics models:

- The Li-Ion Electric Potential model

  The Li-Ion Electric Potential model solves for the electric potential.

  Select this model for all the battery model physics continua that are applied to a lithium-ion battery cell.

- The Li-Ion Concentration model

  The Li-Ion Concentration model accounts for the transport of lithium in the (solid) electrodes, and transport of lithium ions in the (liquid) electrolyte.

  Select this model only for the physics continua that describe the corresponding electrodes and the electrolyte.

At the SEI and in the electrolyte, both physics models are coupled in the following manner:

- The interface conditions for both electric fluxes and concentration depend on the same SEI normal component of electric current density. The SEI normal current is computed from the Butler-Volmer relation, which is in turn a function of both solid/liquid electric potentials and solid/liquid concentrations.

- The electric potential is solved using Eqn. (2410) which is dependent on the solution for the Lithium/Salt Concentration c.
- Concentration transport depends on electric potential, see Eqn. (2407) and Eqn. (2412).

The model transport equations for the solid phase (Eqn. (2409) and Eqn. (2411)) do not contain terms which are directly coupling the electric potential with the Li-ion concentration.

The Li-Ion Battery Cell node represents the Li-Ion Battery Cell model in the simulation tree.



Subsequent sections provide the following information:

- Selecting the Li-Ion Battery Cell model.

Selecting the Li-Ion Battery Cell Model

You select the Li-Ion Battery Cell model using the *Physics Model Selection* dialog.

Before selecting the Li-Ion Battery Cell model:

1. Select Implicit Unsteady from the *Time* group box.

   The Li-Ion Battery Cell Model is available only in implicit transient simulations.

2. Select a material model from the *Material* group box.

The Li-Ion Battery Cell model is available as an optional model.



To activate the Li-Ion Battery Cell model:

- Select Li-Ion Battery Cell from the *Optional Models* group box.



The Li-Ion Concentration and Li-Ion Electric Potential models are available in the Optional Models group box.

What Is the Li-Ion Electric Potential Model?

The Li-Ion Electric Potential model solves for electric potential, and is similar to the Electrodynamic Potential Model. Select the Li-Ion Electric Potential model for all the physics continua that are applied to a lithium-ion battery cell.

Note: The Li-Ion Electric Potential model is available with the double precision version of STAR-CCM+ only.

The Li-Ion Electric Potential node represents the Li-Ion Electric Potential model in the simulation tree.



The Li-Ion Electric Potential model uses the Electric Potential Solver.

Subsequent sections provide the following information:

- Selecting the Li-Ion Electric Potential model.
- Setting the electrical conductivity.
- Setting the region boundary conditions.
- The Li-Ion Electric Potential model formulation.
- The Li-Ion Electric Potential model field function reference.

Li-Ion Electric Potential Properties

None.

Li-Ion Electric Potential Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Selecting the Li-Ion Electric Potential Model

You select the Li-Ion Electric Potential model using the *Physics Model Selection* dialog.

Before selecting the Li-Ion Electric Potential model:

- Select the Li-Ion Battery Cell model.

The Li-Ion Electric Potential model is available as an optional model.



To activate the Electric Potential model:

- Select Li-Ion Electric Potential from the *Optional Models* group box.



Setting the Electrical Conductivity

The Li-Ion Electric Potential model adds a material property: electrical conductivity (σ in Eqn. (2403), or κ_0 in Eqn. (2405)). The electrical conductivity quantifies the ability of the material to conduct electric current, and is measured in S/m. Electrical conductivity is entered as a scalar profile.

The Electrical Conductivity node represents this material property in the simulation tree.



If you use the Tabular method, select the independent variables (usually temperature and concentration) that are used to interpolate the electrical conductivity.

Setting the Region Boundary Conditions

The Li-Ion Battery Cell model activates the boundary conditions of a Li-Ion Battery Cell region. Specify the boundary condition type and set the boundary condition values for the region as part of the problem setup.

If you set the Specific Electric Current or Electric Current Density boundary condition, you are explicitly specifying the electric current that is flowing across that boundary. For example, if you set a value of 1 A/m^2, the normal component of the electric current density is set to a constant value of 1 A/m^2 along the entire boundary.

For the Electric Current boundary condition, you are specifying the total electric current that is flowing across the boundary. However, the profile of the specific electric current along the boundary is not necessarily constant. It is the integral value of the current across the boundary that is fixed at the specified value; the specific electric current can change as the solution develops.

Note: You cannot set independent boundary conditions for the liquid electrolyte phase. This phase has an interface with solid active electrodes. If the Li-Ion Electric Potential and Li-Ion Concentration models are activated in both regions, the SEI interfaces are created automatically. However, one region must be liquid and the other solid; SEI interfaces are not created for solid/solid and liquid/liquid interfaces.

To define the boundary conditions:

1. Select the Regions > [Region] > Boundaries > [Boundary] > Physics Conditions > Electric Potential Specification node.
2. In the *Properties* window, set the Method property:
   - For a Neumann boundary condition:
     - If you know the total electric current, set Method to Electric Current.

       The electric current is defined in A. This boundary condition type specifies the total electric current through the boundary. Current entering the domain is a positive electric current. Current leaving the domain is a negative electric current.

- If you know the electric current density, set Method to Specific Electric Current or Electric Current Density.



The electric current density is defined in A/m^2. The total current through the boundary is the product of the electric current density with the area of the boundary. The Specific Electric Current is entered as a scalar profile, while the Electric Current Density is entered as a vector profile.

- For a Dirichlet boundary condition:
  - If you know the electric potential, set Method to Electric Potential.

    The Electric Potential ϕ is defined in V.
  - If you know the electric current density, set Method to Electric Current Density.

    The electric current density is defined in A/m^2.

An additional node—Specific Electric Current, Electric Current, Electric Current Density, or Electric Potential—is added to the Physics Values manager for the region boundary.

3. Select the Regions > [Region] > Boundaries > [Boundary] > Physics Conditions > [Specific Electric Current, Electric Current, Electric Current Density, or Electric Potential] node.

4. Enter the scalar profile value for the node.



Li-Ion Electric Potential Model Formulation

In the solid electrodes and current collectors, the electric current density is defined using Ohm's law as

$$\mathbf{J} = \sigma \mathbf{E} \tag{2403}$$

where $\mathbf{E} = -\nabla\phi$.

Within the electrolyte phase, the electric current density $\mathbf{J}$ is due to diffusion and migration of ions (salt). Using concentrated solution theory [412], $\mathbf{J}$ can be written as

$$\mathbf{J} = \kappa\mathbf{E} - \frac{\kappa}{F}\left(\frac{s_+}{n\nu_+} + \frac{t_+^0}{z_+\nu_+} - \frac{s_0 c}{n c_0}\right)\nabla\mu_e \tag{2404}$$

where the first term of the right-hand side corresponds to migration, and the second term to diffusion. The effective electrical conductivity κ in the liquid phase accounts for an eventual non-unity porosity χ and tortuosity τ within the separator.

$$\kappa = \frac{\chi}{\tau}\kappa_0 \tag{2405}$$

The Li-Ion Battery Cell model assumes a $LiPF_6$ (EC:DMC) electrolyte. Using $LiPF_6$ chemistry to parametrize Eqn. (2404) [413][414]:

$$\mathbf{J} = \kappa\mathbf{E} + \frac{2R_u T\kappa}{F}(1 - t_+^0)\left[1 + \frac{d(\ln f_\pm)}{d(\ln c)}\right]\nabla(\ln c) \tag{2406}$$

or

$$\mathbf{J} = \kappa\mathbf{E} + \kappa_\mathrm{d}\nabla(\ln c) \tag{2407}$$

where $\kappa_\mathrm{d} = \frac{2R_u T\kappa}{F}(1 - t_+^0)\left[1 + \frac{d(\ln f_\pm)}{d(\ln c)}\right]$.

STAR-CCM+ uses this electrical current density definition for the electrolyte that is used in the Li-Ion Battery Cell Model.

Due to charge conservation, the electric current density field must be divergence free:

$$\nabla \cdot \mathbf{J} = 0 \tag{2408}$$

Applying Eqn. (2403) and Eqn. (2407) (the electric current density definition) to Eqn. (2408), integration over a control volume, and subsequent application of Gauss's theorem yields the transport equations for electric potential

$$\oint_S \sigma\nabla\phi \cdot d\mathbf{s} = 0 \tag{2409}$$

within the (solid) electrode and current collector phases, and

$$\oint_S \kappa\nabla\phi \cdot d\mathbf{s} = \oint_S \kappa_\mathrm{d}\nabla(\ln c) \cdot d\mathbf{s} \tag{2410}$$

within the (liquid) electrolyte phase.

Li-Ion Electric Potential Model Field Function Reference

The following primitive field functions are made available to the simulation when the Li-Ion Electric Potential Model is used. These objects share base properties.

Electric Potential	ϕ in Eqn. (2409).
Electric Field	$\mathbf{E}$ in Eqn. (2403).
Electric Current Density	$\mathbf{J}$ in Eqn. (2403).

Electrical Conductivity	σ in Eqn. (2403), and κ_0 in Eqn. (2407).
Effective Electrical Conductivity	σ in Eqn. (2403), and κ in Eqn. (2407).
Boundary Specific Electric Current	SEI normal component of electric current density J_n
SEI Surface Overpotential	η in Eqn. (2428).

What Is the Li-Ion Concentration Model?

The Li-Ion Concentration model accounts for the transport of lithium in the (solid) electrodes, and transport of lithium ions in the (liquid) electrolyte. You select this model only for the physics continua that describe the corresponding solid electrodes and the liquid electrolyte. This selection is different to the Li-Ion Electric Potential model, which is activated in all the battery model physics continua.

Note: The Li-Ion Concentration model is available with the double precision version of STAR-CCM+ only.

The Li-Ion Concentration node represents the Li-Ion Concentration model in the simulation tree.



Subsequent sections provide the following information:

- Selecting the Li-Ion Concentration model.
- What is the Lithium/Salt Concentration solver?
- Setting the lithium/salt diffusivity.
- Setting the lithium cation transference number.
- Setting the concentrated solution correction to salt diffusivity.
- Setting the non-ideal solution correction to salt migration.
- Setting the region boundary conditions.
- The Li-Ion Concentration model formulation.
- The Li-Ion Concentration model field function reference.

Li-Ion Concentration Properties

None.

Li-Ion Concentration Expert Properties

Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.

Selecting the Li-Ion Concentration Model

You select the Li-Ion Concentration model using the *Physics Model Selection* dialog.

Before selecting the Li-Ion Concentration model:

- Select the Li-Ion Battery Cell model.

The Li-Ion Concentration model is available as an optional model.



To activate the Li-Ion Concentration model:

- Select Li-Ion Concentration from the *Optional Models* group box.



What Is the Lithium/Salt Concentration Solver?

This solver, which has properties and is subject to common menu items, controls the solution of the Li-Ion Concentration model in all the continua for which the model is activated.



Lithium/Salt Concentration Properties

Under-Relaxation Factor	At each iteration, governs the extent to which the newly computed solution supplants the old solution. The default value (1.0) means that no under-relaxation is used.

Lithium/Salt Concentration Expert Properties

The properties in the *Expert* category are used in a temporary debug situation, at the expense of simulation accuracy and higher memory usage. It is recommended that you do not change any of these properties unless you are thoroughly familiar with the discretization techniques used in STAR-CCM+.

Temporary Storage Retained	Retains or deletes the data that is in temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations. Activated — Retains temporary storage. Deactivated — Does not retain temporary storage. (Default)
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This feature is primarily a "debugging" tool and is not supported for normal simulation work. Activated — The solver does not update the solution. Deactivated — The solver updates the solution. (Default)

Setting the Lithium/Salt Diffusivity

The Lithium/Salt Diffusivity (D in Eqn. (2411)) is the diffusion coefficient of the Lithium/salt, and is measured in m^2/s. This material property quantifies the strength of diffusion of the Lithium/salt. Lithium/salt diffusivity is entered as a scalar profile.

The Lithium/Salt Diffusivity node represents this material property in the simulation tree.



If you use the Tabular method, you select the independent variables (usually temperature and concentration) that are used to interpolate the Lithium/salt diffusivity.

Setting the Lithium Cation Transference Number

The Lithium Cation Transference Number (t_+^0 in Eqn. (2412)) is a dimensionless quantity. The transference number of an ion is defined as the fraction of the current that an ion carries in a solution of uniform composition. The Lithium Cation Transference Number is entered as a scalar profile.

The Lithium Cation Transference Number node represents this material property in the simulation tree.



If you use the Tabular method, you select the independent variables (usually temperature and concentration) that are used to interpolate the Lithium Cation Transference Number.

Setting the Concentrated Solution Correction to Salt Diffusivity

The Concentrated Solution Correction to Salt Diffusivity $(d(\ln c_0)/d(\ln c)$ in Eqn. (2413)) is a dimensionless material property. This property accounts for the change in the density of the solution with salt concentration. This material property is available only for the liquid phase, and is entered as a scalar profile.

The Concentrated Solution Correction to Salt Diffusivity node represents this material property in the simulation tree.



If you use the Tabular method, you select the independent variables (usually temperature and concentration) that are used to interpolate the Concentrated Solution Correction to Salt Diffusivity.

Setting the Non-Ideal Solution Correction to Salt Migration

The Non-Ideal Solution Correction to Salt Migration, or thermodynamic factor, $(d(\ln f_{\pm})/d(\ln c))$ is a dimensionless material property. This property accounts for the effect of moderately dilute or concentrated solutions on salt transport. This material property is available only for the liquid phase, and is entered as a scalar profile.

According to Newman [412], electrolytes can be described as either dilute (or ideal) solutions, moderately dilute solutions, and concentrated solutions. For moderately dilute solutions and concentrated solutions, the corrections to dilute solution theory include:

- activity coefficient variations
- variation of ionic diffusion coefficients with concentration

Moderate dilute solution theory accounts for electrochemical potentials and activity coefficients, whereas concentrated solution theory accounts additionally for concentration-dependent diffusion coefficients.

The Non-Ideal Solution Correction to Salt Migration node represents this material property in the simulation tree.



Setting the Region Boundary Conditions

The Li-Ion Battery Cell model activates the boundary conditions of a Li-Ion Battery Cell region. You specify the boundary conditions type and set the boundary condition values for the region as part of the problem setup.

Note: You cannot set independent boundary conditions for the liquid electrolyte phase. This phase requires an interface with solid active electrodes. If the Li-Ion Electric Potential and Li-Ion Concentration models are activated in both regions, the SEI interfaces are created automatically. However, one region must be liquid and the other solid; SEI interfaces are not created for solid/solid and liquid/liquid interfaces.

To define the boundary conditions for the concentration:

1. Select the Regions > [Region] > Boundaries > [Boundary] > Physics Conditions > Lithium Wall Concentration Specification node.
2. In the *Properties* window, set the Method property:

- To specify a Neumann boundary condition, set Method to Flux.



A Lithium Concentration Flux node is added to the Physics Values manager for the region boundary.

- To specify a Dirichlet boundary condition, set Method to Concentration.



A Lithium/Salt Concentration node is added to the Physics Values manager for the region boundary.

3. Select the Regions > [Region] > Boundaries > [Boundary] > Physics Conditions > [Lithium Concentration Flux or Lithium/Salt Concentration] node.

4. Enter the scalar profile value for the node.



Li-Ion Concentration Model Formulation

A simple species diffusion transport equation is solved for lithium (Lithium/Salt Concentration) transport within the electrode phase:

$$\frac{\partial}{\partial t}\int_V c\,dV = \oint_S D\nabla c \cdot d\boldsymbol{s} \tag{2411}$$

where D is a constant defined in Eqn. (2413).

An additional transport term for lithium cation migration is considered in the (liquid) electrolyte phase, so that

$$\frac{\partial}{\partial t}\int_V \chi c\,dV = \oint_S D\nabla c \cdot d\boldsymbol{s} - \int_V \frac{\boldsymbol{J} \cdot \nabla t_+^0}{F}\,dV \tag{2412}$$

The Effective Lithium/Salt Diffusivity Coefficient D accounts for non-unity values of porosity χ and tortuosity τ within the separator, and considers the "Concentrated Solution Correction to Salt Diffusion Coefficient", written as

$$D = D_0 \frac{\chi}{\tau}\left[1 - \frac{d(\ln c_0)}{d(\ln c)}\right] \tag{2413}$$

Convective transport in electrolyte is not modeled.

Li-Ion Concentration Model Field Function Reference

The following primitive field functions are made available to the simulation when the Li-Ion Concentration Model is used. These objects share base properties.

Lithium/Salt Concentration	c in Eqn. (2411).
Lithium/Salt Diffusivity	D in Eqn. (2411) and D_0 in Eqn. (2413)
Effective Lithium/Salt Diffusivity	D in Eqn. (2411) and Eqn. (2413).
Lithium Cation Transference Number	t_+^0 in Eqn. (2412).
Concentrated Solution Correction to Salt Diffusivity	$\frac{d(\ln c_0)}{d(\ln c)}$ in Eqn. (2413).
Non-Ideal solution Correction to Salt Migration	$\frac{d(\ln f_\pm)}{d(\ln c)}$ in Eqn. (2406).

What Is the Solid Electrolyte Interface (SEI)

The Solid Electrolyte Interface (SEI) is the interface between the solid electrode material and the liquid electrolyte.

Note: *Interface* refers to a location. *Interphase* refers to the processes that take part between the two phases of solid active material, and the liquid electrolyte phase (that is, interphase reaction).

SEI are created between interfacing liquid-solid regions, where both the Li-Ion Concentration Model, and the Li-Ion Electric Potential Model are selected on both sides.

Subsequent sections provide the following information:

- Setting up the Butler-Volmer relation parameters.
- Setting up the electrode equilibrium potential curves.
- Guidelines for simulating lithium depletion and saturation.
- Solid Electrolyte Interface formulation.

Setting Up the Butler-Volmer Relation Parameters

Once the corresponding physics continua have been assigned to regions, you set the parameters for the SEI through the Physics Values > Butler-Volmer relation parameters node.



Butler-Volmer Relation Parameters Properties

Resistance	R_0 in Eqn. (2430).
SEI Activation Energy	E_{SEI} in Eqn. (2430).
Capacitance	C in Eqn. (2426).
Rate Constant	k_0 in Eqn. (2429).
Kinetic Activation Energy	E_a in Eqn. (2429).
Maximum Concentration in Electrode	$c_{s,\,\mathrm{max}}$ in Eqn. (2427).
Reference Concentration in Electrolyte	$c_{l,\,\mathrm{ref}}$ in Eqn. (2427).
Anodic Transfer Coefficient	α^a in Eqn. (2426).
Cathodic Transfer Coefficient	α^c in Eqn. (2426).
Electrolyte Reaction Order	α^3 in Eqn. (2427).
Electrode Reaction Order	α^1 in Eqn. (2427).
Vacancies Reaction Order	α^2 in Eqn. (2427).
Under-Relaxation Parameter	At each iteration, governs the extent to which the newly computed solution supplants the old solution. The default value (1.0) means that no under-relaxation is used.

Electrode Equilibrium Potential

U_{eq} in Eqn. (2428).

See Setting Up the Electrode Equilibrium Potential Curves.

Linear — Values are linearly interpolated for the interval $y = [0, 1]$.

Tabular — Uses a table as input.

Butler-Volmer Relation Parameters Expert Properties

Projection Order

Sets the projection order that is used to compute the liquid concentration c_l and solid concentration c_s at the interface.

0th-order — Uses corresponding cell center values. Use this method when lithium depletes or saturates.

1st-order — Uses a first order projection. (Default)

Prevent Electrode Depletion

Whether to correct the user-defined equilibrium voltage curves to approach ∞ asymptotically when $y = 0$, and $-\infty$ when $y = 1$.
Using this option ensures that a local chemical equilibrium exists under all external load conditions, where $\eta = 0$ in Eqn. (2428).

Activated — Activates the correction. (Default)

Deactivated — Deactivates the correction.

Prevent Electrolyte Depletion

Whether to correct the SEI Surface Overpotential to become $\eta = 0$ when the electrolyte depletes of salt. See Eqn. (2428).

Activated — Activates the correction. (Default)

Deactivated — Deactivates the correction.

Setting Up the Electrode Equilibrium Potential Curves

Electrode equilibrium potentials $U_{eq} = f(y)$ are parametrized in terms of the electrode stoichiometry y for both positive and negative electrodes, where

$$y = \frac{c}{c_{s,\,\max}} \qquad y \in [0, 1] \tag{2414}$$

The electrode equilibrium potential is set as a property of the Butler-Volmer relation parameters node, using one of two methods:

- Linear
- Tabular

Linear

This method inputs a simplified electrode equilibrium potential U_{eq}. You supply the values of U_{eq} for the points $y = 0$ and $y = 1$. The values of for the interval $y = [0, 1]$ are linearly interpolated.



For Li-Ion cells, you can approximate the electrode equilibrium potential by a constant, for example 3.7 V for the cathode. When choosing such a setup, the driving force for discharge does not increase during the charging and the discharging process. If Prevent Electrode Depletion is deactivated, this circumstance can lead to numerical problems, where an equilibrium state cannot be attained.

Linear Electrode Equilibrium Potential Properties

Electrode Equilibrium Potential for y = 0	The value of U_{eq} at $y = 0$.
Electrode Equilibrium Potential for y = 1	The value of U_{eq} at $y = 1$.

Tabular

Use a table as input for the electrode equilibrium potential U_{eq}. The electrode equilibrium potential must decrease monotonically with increasing electrode stoichiometry y.



Tabular Electrode Equilibrium Potential Properties

Table	The table that contains the values of U_{eq} as a function of y.
Relative Concentration in Electrode (y)	Column entry for the electrode stoichiometry y. See Eqn. (2414).
Electrode Equilibrium Potential	Column entry for the electrode equilibrium potential U_{eq}.

Guidelines for Simulating Lithium Depletion and Saturation

High charge or discharge currents can lead to strong gradients of Lithium/Salt Concentration within a lithium-ion battery cell. The enhanced transport, and the concentration differences that are required to sustain the gradients then rapidly deplete or saturate (locally) electrodes or electrolyte.

The Li-Ion Battery Cell model can handle cases of lithium depletion and saturation within solid electrodes, as well as the depletion of salt in electrolyte.

When lithium depletes or saturates, the numerical scheme becomes less stable when Eqn. (2426) is evaluated using a 1st-order projection of the Lithium/Salt Concentration towards the SEI. To stabilize the algorithm, use the 0th-order projection.

To account for the resulting low order of the scheme, use a mesh with sufficiently refined prism layers at the SEI.

Simulation of lithium depletion or saturation requires that an equilibrium state can be found. At that equilibrium state, the driving force for reactions at the SEI vanishes, that is $\eta = 0$ in Eqn. (2428). Make sure that such an equilibrium exists by providing appropriate electrode equilibrium potential profiles. If this requirement is not met, activate Prevent Electrode Depletion.

When the electrolyte depletes, activating Prevent Electrode Depletion corrects the SEI surface overpotential to become $\eta = 0$ when the electrolyte depletes of salt.

Solid Electrolyte Interface Formulation

Subsequent sections provide the following information:

- Electric potential interfacing condition.

- Concentration interfacing condition.
- Butler-Volmer relation.

Electric Potential Interfacing Condition

Conservation of the electric charge that is exchanged across the SEI gives the interface condition for the Boundary Specific Electric Current:

$$j_{n,s} = -j_{n,l} \tag{2415}$$

where $j_{n,s}$ is the Boundary Specific Electric Current on the electrode (solid)

$$j_{n,s} = \mathbf{J}_s \cdot \boldsymbol{n}_s \tag{2416}$$

and $j_{n,l}$ is the Boundary Specific Electric Current on the electrolyte (liquid)

$$j_{n,l} = \mathbf{J}_s \cdot \boldsymbol{n}_l \tag{2417}$$

The Boundary Specific Electric Current is the SEI normal component of the Electric Current Density.

The normal vector always points outward of the considered electrode or electrolyte domain:

$$\boldsymbol{n}_s = -\boldsymbol{n}_l \tag{2418}$$

Using the definition for the specific electric current (Eqn. (2415)), the potential boundary condition on the electrode side is

$$\left.\frac{\partial \phi}{\partial n}\right|_{\mathrm{SEI}} = -\frac{j_{n,s}}{\sigma} \tag{2419}$$

Using Eqn. (2419), the potential boundary condition on the electrolyte side is

$$\left.\frac{\partial \phi}{\partial n}\right|_{\mathrm{SEI}} = -\frac{j_{n,l}}{\kappa} + \frac{\kappa_d}{\kappa}\frac{\partial(\ln c_1)}{\partial n} \tag{2420}$$

Concentration Interfacing Condition

By conservation of the flux of Lithium $\mathbf{N}_+$ across the SEI,

$$\mathbf{N}_+ \cdot \boldsymbol{n}_s = -\mathbf{N}_+ \cdot \boldsymbol{n}_l \tag{2421}$$

On the electrolyte side, the flux of lithium ions is expressed as

$$\mathbf{N}_+ = \frac{\mathbf{J}}{F} = -D\nabla c\frac{1}{1-t_+^0} \tag{2422}$$

Taking the inner product of Eqn. (2422) with the SEI normal vector $\boldsymbol{n}_l$, the interfacing condition reads

$$\left.\frac{\partial c}{\partial n}\right|_{\mathrm{SEI}} = -\frac{j_{n,\,l}}{FD}(1 - t_+^0) \tag{2423}$$

On the solid electrode side, the lithium non-ionic flux is expressed as

$$\mathbf{N}_+ = \frac{\mathbf{J}}{F} = -D\nabla c \tag{2424}$$

Taking the inner product of Eqn. (2424) with $\boldsymbol{n}_s$, the concentration interfacing condition becomes

$$\left.\frac{\partial c}{\partial n}\right|_{\mathrm{SEI}} = -\frac{j_{n,\,s}}{FD} \tag{2425}$$

Butler-Volmer Relation

The Butler-Volmer relation is used to compute the SEI normal component of the Electric Current Density $j_{n,\,s} = -j_{n,\,l}$, written as

$$j_{n,\,s} = j_0\left(e^{\frac{\alpha_a F\eta}{R_u T}} - e^{\frac{-\alpha_c F\eta}{R_u T}}\right) + C\frac{\partial}{\partial t}(\phi_s - \phi_l) \tag{2426}$$

$$j_0 = FK\left(\frac{c_s}{c_{s,\,\max}}\right)^{\alpha_1}\left(1 - \frac{c_s}{c_{s,\,\max}}\right)^{\alpha_2}\left(\frac{c_l}{c_{l,\,\mathrm{ref}}}\right)^{\alpha_3} \tag{2427}$$

$c_{l,\,\mathrm{ref}}$ is a molar reference concentration of lithium in the electrolyte, for example.

The SEI Surface Overpotential η is computed as

$$\eta = \phi_s - \phi_l - U_{\mathrm{eq}} - j_{n,\,s}R_{\mathrm{SEI}} \tag{2428}$$

The rate coefficient K depends on temperature according to:

$$K = K_0 e^{\frac{-E_a}{R_u T}} \tag{2429}$$

Similarly, the electrical resistance of the SEI is dependent on temperature:

$$R_{\mathrm{SEI}} = R_0 e^{\frac{-E_{\mathrm{SEI}}}{R_u T}} \tag{2430}$$

Setting Up a Li-Ion Battery Cell with a Lumped Electrode

In many cases, researchers are not interested in the simulation of a full Li-ion battery cell, but only in a simulation of a single electrode. For these cases, simulations are performed for a cell with a two-dimensional lithium foil working electrode, instead of a three-dimensional one. For simulations with a lithium foil working electrode, do the following:

- Approximate the lithium foil by a hexahedral block. The mesh for this block can be coarse.
- Use a linear Electrode Equilibrium Potential for the SEI, and set the reference foil Electrode Equilibrium Potential $U_{eq} = 0$:
  - For Electrode Equilibrium Potential for y = 0, set the value to `0`
  - For Electrode Equilibrium Potential for y = 1, set the value to `0`.
- Deactivate Prevent Electrode Depletion.
- Set the rate constant to a high value, for example `1.0E-3 kmol/m2 s`.
- To make sure that the resolved, non-foil electrode is capacity limiting, set Maximum Concentration in Electrode to a high value, for example $c_{s,\,max} = 1E30$.

Reporting Li-Ion Cell Parameters

The Li-Ion cell report can be run once an initialized setup for a Li-Ion Battery Cell Model is provided. The Li-Ion cell report gives the following values:

- Capacity CAP
- State of charge SOC
- Open circuit voltage OCV

Subsequent sections provide the following information:

- Setting up a Li-ion cell report
- Using the Li-ion cell report state table
- Li-ion cell report formulation

Setting Up a Li-Ion Cell Report

To set up a Li-Ion Cell report:

1. Create a Li-Ion cell report.

2. Select the Reports > [Li Ion Cell Report] node.



3. In the *Properties* window, set the following values:
   - Voltage at 100% SOC—the voltage at which the battery is fully charged.
   - Voltage at 0% SOC—the voltage at which the battery is considered to be fully discharged. (The lower cut-off voltage.)
   - Positive Electrode Parts—the parts that constitute the positive electrode.
   - Negative Electrode Parts—the parts that constitute the negative electrode.
4. Optionally, set the Monitor Value to the quantity that is monitored when creating a monitor.

Using the Li-Ion Cell Report State Table

The Li-Ion Cell Report State Table provides the necessary information for:

- the initial conditions for the Lithium/Salt Concentration in both electrodes, or
- the conditions to reinitialize the Lithium/Salt Concentration at different values of state of charge SOC.

The State Table tabulates the open circuit voltage OCV, the stoichiometries, and average concentration values for both electrodes for different values of SOC.

To determine the appropriate initial (or reinitialization) conditions for Lithium/Salt Concentration in both electrodes:

1. Create an initial setup with a constant Lithium/Salt Concentration of $c = c_{s,\,\max}/2$ as initial value in both electrode regions.
2. Optionally, select the Reports > [Li-Ion Cell Report] node, and set the SOC Step Increment expert property.

   This expert property sets the resolution of the state table.
3. Run the Li-Ion cell Report.

   The state table opens in the output window.



Output

3DEChemTutorial × | Li-Ion Cell 1 Report ×

SOC	OCV (V)	Stoichiometry+	Stoichiometry-	Average Concentration+ (kmol/m^3)	Average Concentration- (kmol/m^3)
0.000000e+00	3.700000e+00	8.382050e-01	2.796539e-01	2.212861e+01	7.382863e+00
2.000000e-01	3.843621e+00	6.944295e-01	3.687494e-01	1.833294e+01	9.734984e+00
4.000000e-01	3.915766e+00	5.506539e-01	4.578449e-01	1.453726e+01	1.208711e+01
6.000000e-01	4.001221e+00	4.068784e-01	5.469404e-01	1.074159e+01	1.443923e+01
8.000000e-01	4.068669e+00	2.631028e-01	6.360359e-01	6.945915e+00	1.679135e+01
1.000000e+00	4.200000e+00	1.193273e-01	7.251314e-01	3.150240e+00	1.914347e+01

   The state table gives the conditions for a cell at a given SOC.
4. Use the value of `Average Concentration+` at the required SOC for the concentration at the cathode.
5. Use the value of `Average Concentration-` at the required SOC for the concentration at the anode.

Li-Ion Cell Report Formulation

Subsequent sections provide the following information:

- Open Circuit Voltage
- State of Charge
- Capacity

Open Circuit Voltage

The Open Circuit (Cell) Voltage OCV is defined as the cell voltage that is approached asymptotically as the cell rests for an infinite period. Empirically, this quantity is measured as the voltage difference between the positive and negative current collectors after the cell is rested for a few hours.

For a cell in perfect equilibrium (that is, where the stoichiometry is constant), the OCV is computed as

$$OCV = U_{\text{eq},+}(\bar{y}) - U_{\text{eq},-}(\bar{x}) \tag{2431}$$

where $U_{\text{eq},+}(\bar{y})$ and $U_{\text{eq},-}(\bar{x})$ are the relative galvanic equilibrium potentials of the positive and negative electrode active materials. $\bar{y}$ and $\bar{x}$ denote their respective (average) stoichiometry (or utilization).

$$\bar{y} = \frac{\int_{+} c dV}{c_{s,\,\max,+} V_{+}} = \frac{\overline{c_{+}}}{c_{s,\,\max,+}} \tag{2432}$$

$$\bar{x} = \frac{\int_{-} c dV}{c_{s,\,\max,-} V_{-}} = \frac{\overline{c_{-}}}{c_{s,\,\max,-}} \tag{2433}$$

State of Charge

The State of Charge SOC is defined as 1 when the battery is fully charged, and 0 when the battery is fully discharged. You define the voltage values corresponding to $SOC = 0$ and $SOC = 1$ in the properties of a Li-ion cell report. The SOC corresponding to other voltage values is linearly interpolated from these two specified voltage values.

Note: Overcharge, as well as deep-discharge, leads to $SOC \notin [0, 1]$.

Capacity

The capacity of the battery cell depends on the maximum amount of lithium (ions) that can be transferred from one electrode to the other, in the range $OCV = [OCV_{SOC=0}, OCV_{SOC=1}]$. A lumped model is used within the Li-Ion Cell Report to compute the cell capacity:

$$CAP = n_{-}\Big|_{SOC=0}^{1} F = (n_{-}\Big|_{SOC=1} - n_{-}\Big|_{SOC=0})F \tag{2434}$$

where n_{-} is the number of lithium moles that are stored in the negative electrode.

Troubleshooting Li-Ion Battery Cell Simulations

Lithium/Salt Concentration does not change in solid active materials

Make sure that all solid active materials are electrically connected and have electrical connections to the current collectors.

Convergence problems when initializing with an arbitrary potential distribution

Just like velocity and pressure in an incompressible CFD computation, Lithium/Salt Concentration and Li-Ion Electric Potential are closely related. When an arbitrary potential distribution (for example, the default zero field) is chosen as initial condition, the solution must be converged within the first time-step.

In some cases, this convergence is disrupted during the first inner iterations. The value of the SEI surface overpotential η (see Eqn. (2428)) is far off the solution value, leading to a spurious computation of the SEI-specific electric current (Eqn. (2426)). This spurious computation affects the Li-Ion Concentration solver.

To overcome this problem, deactivate the Li-Ion Concentration Solver during the first few inner iterations, by setting the under-relaxation factor to zero.

Convergence problems when initial Lithium/Salt Concentration set to zero or saturation value (maximum)

A uniform zero Lithium/Salt Concentration in electrolyte corresponds to a pure (that is, non-conducting) solvent. Such a setup is physically possible, but breaks the operability of the Li-ion battery cell, and is thus considered an input error.

When the electrode Lithium/Salt Concentration values are exactly $c_s = 0$, or $c_s = c_{s,\,\max}$, the derivative of the SEI-specific electric current with respect to Lithium/Salt Concentration becomes infinite (see Eqn. (2427)). The solver becomes unstable.

Such extreme initial conditions are set when investigating the formation process of a Li-ion battery cell. The implementation in STAR-CCM+ cannot handle such cases.

However, simulations in close vicinity of these settings are possible: set initial conditions close to these limits.

Convergence problems when Lithium/Salt Concentration depletes locally in electrolyte, or depletes/saturates in solid electrodes

When lithium depletes or saturates, the numerical scheme becomes less stable when Eqn. (2426) is evaluated using a 1st-order projection of the Lithium/Salt Concentration towards the SEI. To stabilize the algorithm, use the 0th-order projection.

Convergence problems when setting zero resistance in Butler-Volmer relation parameters

To overcome these issues:

- Start with a non-zero resistance.
- Compute some inner iterations until the residuals for Li-Ion Electric Potential and Lithium/Salt Concentration have decreased substantially.
- After reaching a converged solution for the first time, set the resistance to zero again.

Li-Ion Cell Report: Cannot find battery cell states where OCV equals voltages at 0% and 100% SOC

This error can occur when a monitor is used for the Li-Ion Cell Report.

Lithium/Salt Concentration is a conserved quantity. However, due to non-converged simulations, Lithium/Salt can be erroneously introduced into the simulation. Due to this accumulation or loss of lithium, the Li-ion battery cell is incapable of assuming states where the OCV corresponds to the user-defined values at 0% and 100% SOC.

This error points at issues where the solution has not properly converged.

Li-Ion Battery Cells Bibliography

[412] Newman, J., and Thomas-Alyea, K. 2004. *Electrochemical Systems*. J. Wiley.

[413] Spotnitz, R. 2006. Porous Electrode Theory, ECS Short Course: Battery Design and Simulation.

[414] Doyle, C.M. 1995. “Design and Simulation of Lithium Rechargeable Batteries.” PhD Thesis, Department of Chemical Engineering, University of California, Berkeley.

Modeling Corrosion

You can use the Electrodynamic Potential model in STAR-CCM+ to help you to understand and predict the effect of corrosion.

For metal structures that are prone to corrosion, the challenge is to maximize the design life of the structure by preventing corrosion. The ability to understand and predict corrosion behavior using STAR-CCM+ lets you compare corrosion prevention methods and select the optimal solution for each specific scenario—even with complex geometries.

One way to prevent corrosion is using cathodic protection. The concept of cathodic protection is straightforward. However, finding the optimal level of protection is a more complex consideration. Insufficient cathodic protection results in corrosion. Over-protection can result in negative effects such as hydrogen embrittlement, which weakens high-strength steels and increases the likelihood of the structure fracturing.

Using STAR-CCM+, you can optimize the location and size of the sacrificial anodes that prevent corrosion. This process makes sure that the level of cathodic protection that a structure receives is most effective in prolonging the life of the structure. Simulating the cathodic protection of a structure is also cost effective as it prevents wasting expensive anodes unnecessarily.

You can use corrosion modeling in STAR-CCM+ for modeling cathodic protection on:

- ship hulls
- oil rigs and gas pipes
- bridges
- reinforced concrete
- underground structures

This section provides the following information about modeling corrosion:

- What Is Corrosion?
- What is Cathodic Protection?
- What is the Butler Volmer Method?
- What is the Tabular Polarization Curve Method?
- Modeling Cathodic Protection
- Corrosion Formulation
- Troubleshooting Corrosion
- Corrosion Bibliography

What Is Corrosion?

Corrosion is an electrochemical process which occurs naturally when metals react to form compounds that are more stable than their pure elemental form. During corrosion, a number of electrochemical surface reactions occur at the interface between the metal and the electrolyte. Metal atoms lose electrons and become ions which are soluble in the electrolyte. When metals lose electrons, the reaction is known as oxidation—an example of an anodic reaction. The following reaction shows the oxidation of iron:

$$Fe \rightleftharpoons Fe^{2+} + 2e^{-}$$

The electrolyte accepts electrons from the metal and forms ions. When substances accept electrons, the reaction is known as reduction—an example of a cathodic reaction. The following reaction shows the reduction of water:

$$2H_2O + 2e^{-} \rightleftharpoons H_2 + 2OH^{-}$$

The ions that are produced during corrosion form compounds that are, in most instances, soluble in the electrolyte. The corrosion rate depends on the potential drop at the interface between the metal and the electrolyte. Overall, corrosion causes metals to break down and so weaken metal structures and reduce their lifespan.

In the diagram below, iron, Fe, is corroding in the presence of water, H_2O, and oxygen, O_2, and losing electrons, e^-, to form iron ions, Fe^{2+}. The water and oxygen accept electrons from the iron to form hydroxide ions, OH^-:



One method of preventing corrosion is by cathodic protection.

See also:

Corrosion Formulation

What is Cathodic Protection?

Cathodic protection is a process that is used to prevent a metal from corroding by creating an electrochemical cell in which the protected metal forms the cathode. There are several ways to create an electrochemical cell for the purpose of cathodic protection, including:

- Actively imposing a direct electrical current on the metal from a power supply
- Using galvanic anodes

Using Galvanic Anodes

During cathodic protection using galvanic anodes, metal anodes are attached to the metal structure that requires protection. The anodes are made of a metal with an electrostatic potential greater than that of the metal structure. Since the anodes are less noble than the metal, the electrolyte potential is greater near the anodes than near the metal structure. This potential difference results in ion current flow from the anodes to the metal structure. Ions and electrons are produced at the anode, the ions travel through the electrolyte, and electrons travel through the metal circuit. Ions and electrons then recombine at the cathode. The potential drop through the metal circuit is usually much smaller than the potential drop through the electrolyte, since the metal has a much higher conductivity.

The current flow that is supplied to the metal structure prevents the metal structure from undergoing oxidation.

What is the Butler Volmer Method?

The Butler-Volmer method lets you solve a Butler-Volmer equation Eqn. (2441) at the interface boundary between a corroding metal and an electrolyte. The Butler-Volmer equation combines Tafel equations Eqn. (2440) for both anodic and cathodic reactions which occur when reactions are reversible.

When this method is selected, the boundary conditions at the interface between the corroding metal and the electrolyte are set automatically. However, you can adjust the boundary conditions manually if necessary.

See also:

Using the Butler-Volmer Method

Butler-Volmer Properties

Modeling Corrosion

What is the Tabular Polarization Curve Method?

The tabular polarization curve method lets you import polarization curves when modeling corrosion for species that do not follow the same pattern as generalized by the Butler-Volmer equation Eqn. (2441). You can import specific data that defines the precise current for each potential value at

monatomically increasing intervals. STAR-CCM+ uses the polarization curves to automatically calculate the corrosion potential for each metal when the current is zero.

When this method is selected, set the boundary conditions at the interface between the corroding metal and the electrolyte manually.

See also:

Using the Tabular Polarization Curve Method

Tabular Polarization Curve Properties

Modeling Corrosion

Modeling Cathodic Protection

The following procedure describes how to model the effects of cathodic protection of a metal structure within a region that represents an electrolyte.

Before modeling cathodic protection:

Make sure that you have:

- Polarization curves for every metal (when using the tabular polarization curve method).

  STAR-CCM+ uses the polarization curves to automatically calculate the corrosion potential for each metal when the current is zero.
- The conductivity of the electrolyte.

To Model Cathodic Protection:

1. Prepare a volume mesh that represents the electrolyte and the wetted surfaces of the corroding metal and the anode.

   You do not need to include meshed regions for solid parts—only the surfaces of the metal regions that are in contact with the electrolyte are relevant.
2. Set up a physics continuum for the electrolyte.

   In addition to the fluid models, activate the following electromagnetic models:

   - Electromagnetism

- Electrodynamic potential



3. Set the physics model properties.

   Make sure that you specify Electrical Conductivity and Electric Potential values that correspond with the material that you are modeling.



4. Set Boundary Conditions

   For each boundary that represents a metallic surface:

   a. Select the Physics Condition > Electric Potential Specification to Electric Current Potential Characteristics.

   b. Set Physics Values > Current Voltage Characteristic to Butler-Volmer or Tabular Polarization Curve.

      A corresponding sub-node is added to the tree.

   c. Select the corresponding sub-node and set the properties. See Using the Butler-Volmer Method or Using the Tabular Polarization Curve Method.

5. Define the interface boundary between the solid and the electrolyte.
   - Select the boundary on the electrolyte region first
   - Select the boundary on the solid region next

Note: If the residuals do not converge, see Troubleshooting Corrosion.

See also:

What Is the Electrodynamic Potential Model?

Setting Region Boundary Conditions

Using the Butler-Volmer Method

The Butler-Volmer method simulates corrosion using a Butler-Volmer equation for the electrochemical reactions on the interface between the metal and the electrolyte.

When Current Voltage Characteristic is set to Butler-Volmer, the boundary conditions are defined automatically—which you can alter as required (see Butler-Volmer Properties):



See also:

What is the Butler Volmer Method?

Setting Region Boundary Conditions

Using the Tabular Polarization Curve Method

The Tabular Polarization Curve method allows you to simulate corrosion using tabular data that defines the current and electrochemical potential at the interface between the metal and the electrolyte.

1. To define the polarization curve for the metal structure that requires protection, import a .csv table that contains monatomically increasing potential values and corresponding values for current.



The values for current must include positive and negative current values so that STAR-CCM+ can calculate the potential at zero current.

2. Select the Current Voltage Characteristic > Tabular Polarization Curve node and in the *Properties* window:

   a. Set Table to the imported file table

   b. Set the remaining properties as you require (See Tabular Polarization Curve Properties).

See also:

What is the Tabular Polarization Curve Method?

Setting Region Boundary Conditions

Butler-Volmer Properties

The default Butler-Volmer properties are selected automatically for a boundary when the Physics Value, Current Voltage Characteristic, for that boundary is set to Butler-Volmer.

See Butler-Volmer Properties

Corrosion Formulation

The Electrochemical surface reactions that occur during corrosion either produce electrons (anodic), or consume electrons (cathodic). Any number of these reactions can occur at any given time at the interface between the metal that requires protection and the electrolyte.

A general form of electrochemical surface reactions is represented as:

$$j_{n,s} = f(\eta)\ ;\ \begin{Bmatrix} j_{n,s} > 0 \\ \eta > 0 \end{Bmatrix} \equiv anodic\ and \begin{Bmatrix} j_{n,s} < 0 \\ \eta < 0 \end{Bmatrix} \equiv cathodic \tag{2435}$$

Where $j_{n,s}$ is the current density and η is the overpotential.

The driving force behind these reactions is the difference in Gibbs free energy between the reactants and the products, ΔG°, which equates to the equilibrium potential, U_{corr}. For the electrochemical reaction:

$$(s_a A \leftrightarrow s_b B) + ne^- \tag{2436}$$

$$U_{corr}\left[volts\right] = \frac{\Delta G^{\circ}}{nF} \tag{2437}$$

Where the Faraday constant, F =96,485,258.799 C/kmol and n is the magnitude of electric charge per mole of electrons.

The potential difference, U, at the interface is due to the difference in potential of the metal and the electrolyte. In STAR-CCM+, the potential difference is defined as:

$$U = \phi_s - \phi_l \tag{2438}$$

Where ϕ_s is the potential of the solid and ϕ_l is the potential of the liquid electrolyte.

If the potential difference at the interface is larger than the corrosion equilibrium potential, $U > U_{corr}$, the reaction is anodic. If the potential difference is less than the corrosion equilibrium potential, $U < U_{corr}$, the reaction is cathodic. When the equilibrium potential is zero, the anodic and cathodic reactions balance and the net flux is zero.

The rate at which the anodic and cathodic reactions occur is related to an electric potential driving force, called the overpotential. The overpotential depends upon the electrochemical potential difference at the interface and the equilibrium potential. In corrosion reactions, the overpotential, η, is defined as:

$$\eta = \phi_s - \phi_l - \phi_E \tag{2439}$$

The corrosion overpotential definition is equivalent to that of the overpotential which is defined for the electrodynamic potential model in Eqn. (2322). Where:

- ϕ_s is equivalent to ϕ_1

  In corrosion modeling, this value represents the potential of the boundary on the solid region side of the interface which corresponds to the second boundary that you select when defining an interface.

- ϕ_l is equivalent to ϕ_0

  In corrosion modeling, this value represents the potential of the boundary on the electrolyte region side of the interface which corresponds to the first boundary that you select when defining an interface.

In Eqn. (2439), ϕ_E represents the corrosion equilibrium potential U_{corr}.

Anodic reactions (that produce electrons) are related to positive overpotential. Cathodic reactions (that consume electrons) are related to negative overpotential. The current density, i, in these reactions is positive for anodic reactions and negative for cathodic reactions.

The flow of charged particles determines the current, that is, electrons in the metal and ions in the electrolyte.

The current density can be related to the electrochemical potential of a corroding metal using either Tafel equations or a Butler-Volmer equation:

- Tafel equations:

$$j_{n,s} = \pm j_o \exp\left(\pm\alpha \frac{nF}{R_u T}\eta\right) \tag{2440}$$

Where j_o is the exchange current density and α is the charge transfer coefficient.

- Butler-Volmer equation:

When a reaction is reversible, both anodic and cathodic Tafel equations are combined to form the Butler-Volmer equation:

$$j_{n,s} = j_o\left\{\exp\left(\alpha_a \frac{nF}{R_u T}\eta\right) - \exp\left(-\alpha_c \frac{nF}{R_u T}\eta\right)\right\} \tag{2441}$$

Where α_a and α_c are the anodic and cathodic charge transfer coefficients respectively.

The exchange current density and equilibrium potential are often dependent on other variables such as the concentrations of reactants and products, temperature, and surface contaminants.

See also:

Electrochemistry Nomenclature

Troubleshooting Corrosion

If residuals do not converge while running a corrosion simulation, there are several amendments that you can make to improve the solution and achieve convergence.

Since the boundary conditions in corrosion simulations are non-linear, you can adjust the solver settings to account for the non-linear conditions:

To Adjust the Solver Parameters:

- Select the Solvers> Electric Potential > Under-Relaxation Factor Ramp > AMG Linear Solver node.
  - Activate Enable direct Solver

- Set Acceleration method to Bi Conjugate Gradient Stabilized



- Select the Solvers > Electric Potential node and reduce the Under-Relaxation Factor.



- For the boundary that represents the corroding surface, select the Physics Values > Current Voltage Characteristic > Tabular Polarization Curve node, and increase the Linearization Factor.

  Ideally, set the linearization factor to the same value as the under-relaxation factor.



Corrosion Bibliography

[415] Krupa, Matthew, et al. "Cathodic Protection of Highly Complex and Shielded Components." *Corrosion*. 2006.

[416] Zou, Y., J. Wang, and Y. Y. Zheng. "Electrochemical techniques for determining corrosion rate of rusted steel in seawater." *Corrosion Science* 53.1 (2011): 208-216.

Modeling Ionic Species Flux

You can use the Electrochemical Species model within a liquid or gas to model the transport of ionic species, including, diffusive, convective, and migrative/drift fluxes.

There are many applications that rely on effective migration of charged particles. In general, charged particles are referred to as ionic species which can represent positively or negatively charged ions. Modeling the flow and concentration of ionic species using STAR-CCM+ provides a convenient and cost effective approach to comparing the efficiency of various engineering setups. This ability is even more useful when raw materials for electrochemical applications are rare or expensive.

By modeling all relevant transport mechanisms for ionic species, it is possible to predict the heterogeneity of ion concentrations within an electrolyte. Migration of ionic species is proportional to an electric field. You can only define the migration of ionic species when the Electrochemical Species model is selected along with one of the following models:

- Electrodynamic Potential model

  Using the Electrochemical Species model with the Electrodynamic Potential model, allows you to model ionic species flux in electrochemical applications such as:

  - paint delamination
  - pitting corrosion
  - crevice corrosion
  - etching
  - flow batteries
- Electrostatic Potential model

  Using the Electrochemical Species model with the Electrostatic Potential model, allows you to model ionic species flux in electrochemical applications such as:

  - ionic wind simulations
  - electrostatic precipitators
  - electrophoretic spray-painting
  - plasma modeling

If the Electrochemical Species model is used without either the Electrodynamic Potential or Electrostatic Potential model, STAR-CCM+ can only calculate diffusive and convective flux, and not migrative flux. You can define convection, diffusion, and migration terms using the electrochemical species solver.

To model electrohydrodynamic forces that arise from local ionic charge densities, use the Coulomb Force model in addition to the Electrochemical Species model and Electrostatic Potential model. See Coulomb Force Formulation.

This section explains the following:

- Selecting the Electrochemical Species Model
- Setting up the Electrochemical Species Model
- Ionic Species Flux Region Physics Conditions
- Ionic Species Flux Boundary and Interface Conditions
- Electrochemical Species Model Expert Properties
- Electrochemical Species Model Formulation
- Electrochemical Species Model Bibliography

See also:

Electrochemistry Nomenclature

Selecting the Electrochemical Species Model

The Electrochemical Species model uses transport equations to solve the flux of ionic species in electrolytic solutions.

Note: For improved convergence, use a double precision version of STAR-CCM+.

Before selecting the Electrochemical Species model, set up and create a volume mesh.

To model convection, diffusion, and migration of ionic species using the Electrochemical Species model, create, or select an existing physics continuum and select the physics models that you require.

1. Select one of the following material models which are compatible with the Electrochemical Species model.
   - Liquid

Use the Liquid model for modeling corrosion applications.

- Multi-component Liquid
- Gas
- Multi-component Gas

Note: The Electrochemical Species model does not support Multiphase models.

2. If Multi-component Liquid or Multi-component Gas is selected, select Non-reacting.
3. Select either Segregated Flow or Coupled Flow.
4. Select Constant Density.
5. Select suitable *Space, Time*, and *Viscous Regime* models.
6. Select Electrochemical Species.
7. Select Electromagnetism.

   then select either of the following *Electromagnetism* models:

   - Electrodynamic Potential model

     Select this model when simulating applications such as corrosion and etching.
   - Electrostatic Potential model

     Select this model when simulating non-corrosion applications, especially if you want to model ionic charge densities or other setups which do not require electroneutrality.

   If you do not select an Electromagnetism model, the migrative flux cannot be solved—only diffusive and convective flux are solved.
8. select any extra *Optional Models*:
   - If the Electrostatic Potential model is already selected and you want to model the coulomb force that is exerted due to non-zero ionic charge densities, select the Coulomb Force model.
   - If you want to set a constant temperature without storing or solving an additional transport equation, select the Segregated Fluid Isothermal model.
9. Define the parameters for each of the selected physics models.

   See Setting up the Electrochemical Species Model.

Setting up the Electrochemical Species Model

1. Select the [Physics Continuum] > Models > Electrochemical Species node and if necessary set any *Expert Properties* that you require.

See Electrochemical Species Model Expert Properties.

2. Right-click the Electrochemical Species > Electrochemical Species Components node and select one of the following options:
   - Select Mixture Components

     In the *Select Mixture Components* dialog, select the mixture components from the Material Databases, then click Apply and Close.
   - Import Species

     In the *Select file and material database* dialog, select the file which contains the species that you are importing. Specify the Material Database to which these species are to belong, then click OK.

     See Using Material Databases.
3. Expand the Electrochemical Species Components node and for each of the Electrochemical Species Components, specify the Material Properties Methods.

   Certain material properties are specific to the electrochemical species model and allow you to define the properties of the ionic species directly. You can also define ionic species further under the Tools > Material Databases node.

   The following material properties are available to specify when the Electrochemical Species model is selected:
   - Using the Nernst-Einstein method for Charged Species Mobility
   - Using the Electrochemical Species method for Electrical Conductivity
4. If you use the Electrochemical Species model with the Electrodynamics Potential model:
   a. Expand the Models > [Liquid] > Material Properties node.
   b. Select the Electrical Conductivity node and set Method to Electrochemical Species.

      See Using the Electrochemical Species method for Electrical Conductivity.
5. Specify region conditions.

   See Ionic Species Flux Region Physics Conditions
6. Specify boundary conditions.

   See Ionic Species Flux Boundary and Interface Conditions

Ionic Species Flux Region Physics Conditions

To model ionic species flux effectively in a region, you can define specific source terms to solve in the transport equations that are suited to the application that is being modeled.

When modeling ionic species flux, you can set the following region source terms:

- Electric Charge Density User Source

  Part of the electric charge density which is used in the evolution equation for electric potential, Eqn. (2455), is a result of molar concentrations. This electric charge density is computed using Eqn. (2450). STAR-CCM+ allows you to define a region source, S_{Estat}, that provides additional electric charge density.

  See Setting an Electric Charge Density Region Source.

- Electrochemical Species Source Term

  See Defining the Electrochemical Species Source Term.

- Transfer Current Density Region Source

  See Setting a Transfer Current Density Region Source.

Defining the Electrochemical Species Source Term

Defining the electrochemical species source term allows you to provide additional physics in Eqn. (2444) to represent a specific molar concentration.

- Electrochemical Species Sources

  Use this node to specify the species source, S_{c_i}, of each of the defined species as a scalar array profile.

- Electrochemical Species Source Jacobian

  Use this node to specify the Jacobian, $d(S_{c_i})/d(c_i)$, of each of the defined species as a scalar array profile.

Ionic Species Flux Boundary and Interface Conditions

Defining specific parameters for ionic species at the boundary face can help to stabilize the solution and speed up convergence.

The following boundary conditions are relevant to the Electrochemical Species model:

- Electric Potential Specification

Different methods are provided depending on the chosen electromagnetic model. See either Electrodynamic Potential model or Electrostatic Potential model.

- Wall Electrochemical Species Option

  See Wall Electrochemical Species Option

Wall Electrochemical Species Option

The wall electrochemical species option on a boundary allows you to specify the molar concentration flux, and the derivative of the molar concentration flux of the same species. You can also define a specific molar concentration at the boundary, or define the boundary as impermeable to electrochemical species. This boundary condition is available to specify when the Electrochemical Species model is selected.

Wall Electrochemical Species Option Properties

Method	Specifies the species concentration or flux at the boundary for the Electrochemical Species model.	
	Impermeable	Specifies zero flux at the boundary.
	Specified Flux	Creates physics values sub-nodes for each electrochemical species: • Molar Concentration Flux • Molar Concentration Flux Derivative The Molar Concentration Flux Derivative sub-node allows you to specify the derivative of the molar concentration flux that is used for linearization. You can specify the Method for these options as a Constant or Composite array. For a stable set-up, make sure that the flux derivative has a negative value.
	Specified Value	Creates a Molar Concentration node which allows you to specify the molar concentration as a Constant or Composite array.

See also:

Molar Concentration Flux Derivative Formulation

Electrochemical Species Model Expert Properties

The following expert model *Properties* are set to default options which are appropriate for corrosion-related applications. Maintain the default settings unless you are certain that you can achieve a consistent setup.

Convection	Selects the convection scheme that is used when computing the convective flux at a boundary face. See Boundary Face Discretization.	
	1st-order	Selects the first-order upwind convection scheme. See Eqn. (2551).
	2nd-order	Selects the second-order upwind convection scheme. See Eqn. (2552).
Secondary Gradients	Neglect or include the boundary secondary gradients for diffusion and/or the interior secondary gradients at mesh faces.	
	On	Include both secondary gradients.
	Off	Exclude both secondary gradients.
	Interior Only	Include the interior secondary gradients only.
	Boundaries Only	Include the boundary secondary gradients only.
Electroneutrality option	Specifies if the overall ionic charge density is strictly enforced as zero or not. Maintain the default option (Enforced) when simulating corrosion applications.	

	Not Enforced	A transport equation is solved for each registered electrochemical species and a uniform zero ionic charge density is not enforced. However, depending on the setup of other parameters, the overall ionic charge density may already equal zero. Select this option for non-corrosion applications that use the Electrostatic Potential model.
	Enforced	Ensures that the overall ionic charge density is equal to zero. The transport equation for the species that is ordered last in the species is not solved. See Reordering Mixture Components. Instead, the electroneutrality condition, $\rho_{ion} = 0$, is used (See Eqn. (2450)). This option is appropriate for applications that feature electroneutrality, such as corrosion and etching. See Eqn. (2446).
Transport	Defines how the Electrochemical Species model is coupled to other models and so defines the transport equations that STAR-CCM+ solves.	
	Species <> Electric Potential	Two-way migration of flux contributions between the Electrochemical Species and Electric Potential models.
	Electric Potential > Species	One-way migration of drift flux contributions from the Electric Potential model to the Electrochemical Species model.

Species > Electric Potential	One-way transfer of contributions from the Electrochemical Species model to the Electric Potential model. • When using the Electrostatic Potential model, the contributions are for the space charge density. • When using the Electrodynamic Potential model, the contributions are for the diffusive current.
No Electric Potential Coupling	The Electrochemical Species model and the Electric Potential model act independently.

Electrochemical Species Model Formulation

The formulation of transport equations that are solved depends on which model you couple with the Electrochemical Species model. If necessary, you can define the coupling option using the Electrochemical Species model expert property, Transport. See Electrochemical Species Model Expert Properties.

The options are:

- Coupling Electrochemical Species and Electrodynamic Potential Models
- Coupling Electrochemical Species and Electrostatic Potential Models

The following sections provide further extensions of Electrochemical Species model formulation:

- Coulomb Force Formulation
- Molar Concentration Flux Derivative Formulation

See also:

Electrochemistry Nomenclature

Coupling Electrochemical Species and Electrodynamic Potential Models

The formulation in this section is solved when modeling applications such as corrosion and etching.

The Electrochemical Species model solves a transport equation for molar concentrations.

Continuous form:

$$\frac{\partial c_i}{\partial t} = -\nabla \cdot N_i \tag{2442}$$

$$N_i = -D_i \nabla c_i - z_i F u_i c_i \nabla \phi + c_i \mathrm{v} \tag{2443}$$

Finite Volume or integral form:

$$\frac{\partial}{\partial t}\int_V c_i \mathrm{d}V + \oint_A c_i \mathrm{v} \cdot \mathrm{da} = \oint_A D_i \nabla c_i \cdot \mathrm{da} + \oint_A z_i F u_i c_i \nabla \phi \cdot \mathrm{da} + \int_V S_{c_i} \mathrm{d}V \tag{2444}$$

When the Electrochemical Species model is active, the Electrodynamic Potential model solves a transport equation for the ionic space charge density.

Continuous form:

$$\frac{\partial \rho_{ion}}{\partial t} = -\nabla \cdot \mathrm{J} = 0 \tag{2445}$$

$$\mathrm{J} = -\kappa \nabla \phi - F \sum_i z_i D_i \nabla c_i + \rho_{ion} \mathrm{v} \tag{2446}$$

Assuming electroneutrality, $\rho_{ion} = 0$, the finite volume, or integral form reads:

$$0 = \oint_A \kappa \nabla \phi \cdot \mathrm{da} + F \sum_i z_i \oint_A D_i \nabla c_i \cdot \mathrm{da} + \int_V S_{\rho_{ion}} \mathrm{d}V \tag{2447}$$

Substituting expressions for the ionic conductivity, κ,

$$\kappa = F^2 \sum_i z_i^2 u_i c_i \tag{2448}$$

the Nernst-Einstein relation,

$$u_i = \frac{D_i}{R_u T} \tag{2449}$$

and the relation between the ionic space charge density and molar concentrations,

$$\rho_{ion} = F\sum_i z_i c_i \tag{2450}$$

the following flux expressions are derived:

$$N_i = -D_i \nabla c_i - z_i F + \frac{D_i}{R_u T} c_i \nabla \phi + c_i \mathrm{v} \tag{2451}$$

$$\mathbf{J} = -F^2 \sum_i z_i^2 \frac{D_i}{R_u T} c_i \nabla \phi - F\sum_i z_i D_i \nabla c_i + F\sum_i z_i c_i \mathrm{v} \tag{2452}$$

Comparison of Eqn. (2451) and Eqn. (2452) produces a relation between the molar species fluxes and the electric current densities:

$$\mathbf{J} = F\sum_i z_i N_i \tag{2453}$$

Due to the electroneutrality, $\rho_{ion} = 0$, that is assumed in corrosive systems, the convective term in Eqn. (2446) is ignored.

Coupling Electrochemical Species and Electrostatic Potential Models

When coupling Electrochemical Species and Electrostatic Potential models, the following equation,

$$\nabla \cdot \mathrm{E} = \rho_{ion} / \varepsilon_0 \tag{2454}$$

replaces Eqn. (2445) to determine the electric potential.

The continuous form is written as:

$$\oint_A \varepsilon_0 \nabla \phi \cdot \mathrm{da} = \int_V \rho_{ion} \mathrm{d}V + \int_V S_{Estat} \mathrm{d}V \tag{2455}$$

Here, ρ_{ion} is non-zero and is computed from Eqn. (2450).

The transport equations for the electrochemical species are expressed in terms of a species number density:

$$\frac{\partial N_P}{\partial t} = -\nabla \cdot \mathrm{N} \tag{2456}$$

$$\mathrm{N} = -N_P \mathrm{K} \nabla \phi - D \nabla N_P + N_P \mathrm{v} \tag{2457}$$

See Cagnoni *et al.* (2013) [417].

The species number density, N_P, is equal to the molar concentration of the ionic species, c, multiplied by Avogadro's number. Therefore, number density represents the number of ionic species particles per volume.

Often, only a single ionic species with a unity charge number ($z = 1$) is considered.

For a single ionic species, i, the relation:

$$qN_P = \rho_{ion} = qN_A c = Fc \tag{2458}$$

(by definition of Faraday's constant) is used to rewrite Eqn. (2456) and Eqn. (2457) as a transport equation for charge density:

$$\frac{\partial \rho_{ion}}{\partial t} = -\nabla \cdot \mathrm{J} \tag{2459}$$

$$\mathrm{J} = -\rho_{ion}\mathrm{K}\nabla\phi - D\nabla\rho_{ion} + \rho_{ion}\mathrm{v} \tag{2460}$$

as in Neimarlija *et al.* (2009) [418].

You can also obtain Eqn. (2459) and Eqn. (2460) by multiplying Eqn. (2442) and Eqn. (2443) with Faraday's constant, F, and charge number, z_i, applying Eqn. (2450), and substituting ionic mobility as $\mathrm{K}_i = Fz_i u_i$.

Although K and u_i are commonly termed ion mobility, their dimension and units differ by the dimension of Faraday's constant ($[\mathrm{m}^2/\mathrm{V\,s}]\,\mathrm{vs}\,[\mathrm{m}^2\mathrm{kmol}/\mathrm{J\,s}]$).

Coulomb Force Formulation

The Coulomb force that is exerted on ionic species in an electric field leads to a volumetric force term which is written as:

$$\mathrm{f}_{Coulomb} = \rho_{ion}\mathrm{E} \tag{2461}$$

See Cagnoni *et al.*, (2013) [417].

The ionic space charge density here is computed from Eqn. (2450) rather than considering a single ionic species as in Eqn. (2458).

Molar Concentration Flux Derivative Formulation

You can define the molar concentration flux derivative by specifying the wall electrochemical species option Method on a boundary.

If you specify a flux derivative:

$$\frac{\partial N_i}{\partial c_i}, \tag{2462}$$

the derivative of the boundary flux with respect to the concentration at the cell centre can then be computed. This term can then be incorporated into the construction of the Jacobian matrix.

where:

N_i = molar flux of species i [kmol/m^2 s]

c_i = molar concentration of species i at the wall boundary [kmol/m^3]

See also:

Ionic Species Flux Boundary and Interface Conditions

Electrochemical Species Model Bibliography

[417] Cagnoni, D., Agostini, F., Christen, T., Parolini, N., Stevanovi c, I. & de Falco, C. 2013 Multiphysics simulation of corona discharge induced ionic wind. *Journal of Applied Physics* **114** (23).

[418] Neimarlija, N., Demirdzic, I. & Muzaferija, S. 2009 Finite volume method for calculation of electrostatic fields in electrostatic precipitators. *Journal of Electrostatics* **67** (1), 37 — 47.

[419] Newman, J. and Thomas-Alyea, K. E. 2004. *Electrochemical Systems.*, 3rd edn. J. Wiley. ISBN: 978-0-471-47756-3.

[420] Georgiadou, Maria and Richard Alkire. "Anisotropic Chemical Pattern Etching of Copper Foil." *J. Electrochem Soc* (1994): 679-689.

Electrochemistry Nomenclature

The following is a guide to the symbols used in this section.

Scalars

α_1	Electrode reaction order [-]
α_2	Vacancies reaction order [-]
α_3	Electrolyte reaction order [-]
α_a	Anodic transfer coefficient [-]
α_c	Cathodic transfer coefficient [-]
$\Delta G°$	Gibbs free energy of reactants and products [J/mol]
η	SEI surface overpotential [V]
κ	Ionic conductivity [S/m] or [A/V m]
μ_e	Chemical potential of electrolyte [J/mol]
ν_+	Number of cations into which a molecule of electrolyte dissociates [-]
$\overline{c_+}$	Uniform/Volume average Lithium/Salt concentration in positive electrode [kmol/m^3]
$\overline{c_-}$	Uniform/Volume average Lithium/Salt concentration in negative electrode [kmol/m^3]
∇	Del operator [1/m]
$\overline{x_{init}}$	Volume average initial stoichiometry of negative electrode active materials [-]
$\bar{x}$	Volume average stoichiometry of negative electrode active materials [-]
$\overline{y_{init}}$	Volume average initial stoichiometry of positive electrode active materials [-]
$\bar{y}$	Volume average stoichiometry of positive electrode active materials [-]
ϕ	Electric potential [V]
ϕ_s	Electric potential of the solid [V]
ϕ_l	Electric potential of the liquid electrolyte [V]
ϕ_{corr}	Corrosion potential [V]
ϕ_s^{corr}	Corrosion potential of the solid [V]
ϕ_l^{corr}	Corrosion potential of the liquid electrolyte [V]
σ	Effective electrical conductivity [S/m] or [A/V m]
σ_0	Electrical conductivity [S/m] or [A/V m]
τ	Tortuosity, non-unity only in separator [-]
C	Double-layer capacitance [F/m^2] [C/V m^2]

c	Molar concentration of a single electrolyte [kmol/m^3]
c_0	Molar concentration of solvent [kmol/m^3]
c_i	Molar concentration of species i [kmol/m^3]
c_l	Molar concentration of lithium cations in electrolyte (liquid) [kmol/m^3]
c_s	Molar concentration of lithium in electrode material (solid) [kmol/m^3]
$c_{l,\,\mathrm{ref}}$	Molar reference concentration of lithium in electrolyte [kmol/m^3]
$c_{s,\,\max}$	Molar concentration of lithium in saturated electrode material [kmol/m^3]
CAP	Call capacity [C] [A s]
D	Effective lithium diffusivity [m^2/s]
D_i	Diffusivity of species i [m^2/s]
D_0	Reference lithium diffusivity [m^2/s]
ε	permittivity [A^2 s^2/Nm^2]
E	Electric field $-\kappa\nabla\phi$ [V/m]
E_a	Kinetic activation energy [J/kmol]
E_{SEI}	SEI activation energy [J/kmol]
F	Faraday's constant 96,485,258.799 [C/kmol]
f	Force [N/m^3]
$f_{\pm}$	Mean molar activity coefficient of an electrolyte [-]
j_0	Exchange current density [A/m^2]
j_n	SEI normal component of electric current density [A/m^2]
$j_{n,\,s}$	Current density [A/m^2]
J_i	Non-linear flux
κ	Ionic conductivity [S/m]
K	Ion mobility [m^2/V s]
K	Rate coefficient [equiv/s m^2]
K_0	Reference rate constant [equiv/s m^2]
N_i	Molar flux density of species i [kmol/m^2 s]
n	Number of electrons that are involved in electrode reaction [-]
n_+	Li+ cations in positive electrode active material [kmol]
n_-	Li+ cations in negative electrode active material [kmol]
N	Number flux [1/m^2 s]
N_A	Avagadro constant 6.022136736 e^{26} [1/kmol]
N_P	Number density [1/m^3]
OCV	Open circuit voltage [V]
$OCV_{SOC\,=\,0}$	Open circuit voltage at SOC = 0
$OCV_{SOC\,=\,1}$	Open circuit voltage at SOC = 1
q	Elementary charge 1.602176487 e^{-19} [A s]
$\dot{Q}$	Volumetric heat release [J/m^3]
ρ_{ion}	Ionic space charge density [A s/m^3]
R_u	Universal gas constant [8314.4621 J/kmol K]

R_0	Reference electrical resistance of SEI [Ω m^2]
R_{SEI}	Electrical resistance of SEI [Ω m^2]
s_+	Stoichiometric coefficient of cation in electrode reaction [-]
s_0	Stoichiometric coefficient of solvent in electrode reaction [-]
S_{c_i}	Electrochemical Species source term.
S_{Estat}	Electric charge density
SOC	State of charge. $SOC \in [0, 1]$ under standard operation conditions. [-]
T	Temperature [K]
t	Time [s]
t_+^0	Transference number of cations with respect to the velocity of the solvent species [-]
u_i	Mobility of species i [m^2 kmol / J s]
U	Potential difference
$U_+(y)$	Relative galvanic equilibrium potential of positive active electrode material [V]
$U_-(y)$	Relative galvanic equilibrium potential of negative active electrode material [V]
U_{eq}	Relative galvanic equilibrium potential of active electrode material [V]
v	Velocity [m^2 /s]
V_+	Volume of positive electrode active material [m^3]
V_-	Volume of negative electrode active material [m^3]
χ	Porosity, non-unity only in separator [-]
z_+	Charge number of cation [-]
z_i	Charge number of species i [-]
SEI	Solid electrolyte interface

Tensors or Column Vectors

$\mathbf{E}$	Electric field $-\kappa\nabla\phi$ [N/C] or [V/m]
$\mathbf{J}$	Electric current density [A/m^2]
$\mathbf{N}_+$	Lithium (cation) molar flux [kmol/m^{2}s]

Cartesian Vectors

$\boldsymbol{n}_l$	Normal vector pointing outwards from liquid region [-]
$\boldsymbol{n}_s$	Normal vector pointing outwards from solid region [-]

Modeling Batteries

The battery modeling capability in STAR-CCM+ is used to simulate batteries. This process involves designing individual battery cells and battery cycling procedures in the external software package Battery Design Studio® (BDS). Currently, only spirally wound cylindrical cells and stacked cells that are designed in BDS can be used in STAR-CCM+. Once the battery model has been activated in STAR-CCM+, the output files from BDS can be imported into the simulation and used to set up battery cells, battery modules, and electrical loads. The battery module generation process is then used to first create geometry and connector parts and then regions for the battery objects. A battery mesher is provided to mesh the battery objects. It is also possible to create a volume that automatically encases the battery, which can be used to define a fluid continuum around the battery. A circuit solver is used to connect battery modules with electrical loads.

This section also includes formulation and a reference section for battery model field functions.

Licensing Requirements

For battery modeling, STAR-CCM+ requires a `batterysim` license in addition to the other license features.

Battery Design Studio®, which is supplied for use with the STAR-CCM+ battery modeling capability, requires a `batterydesignstudio` license.

The `batterysim` license cannot be used for Battery Design Studio®.

What Is Battery Modeling?

Battery modeling is the process of simulating the electrical and thermal response of a battery over the course of a load cycle. In STAR-CCM+, this is done using a Battery Simulation Module (BSM), which uses a combination of an electrical solver and a fluid/thermal solver. The electrical solver calculates the electrical behavior of the battery cell while the fluid/thermal solver calculates the heat transfer within the cell and its surroundings.

Battery Structure Terminology

You can model two types of battery cells in STAR-CCM+: stacked and cylindrical spiral cells. Each cell type is associated with its own basic terminology. The following images depict the general terminology that is used to describe each cell structure.





- Electrodes—the cell components that conduct the electrochemical reaction (every electrode has a tab). Electrodes are comprised of coating on a metal foil. The positive electrode is typically a metal oxide or metal phosphate on an aluminum foil while the negative electrode is typically a graphite paste on copper.
- Separator—a layer that is made of non-conductive material that keeps two electrodes separated, thus preventing a short-circuit. A separator allows ionic transport.
- Tabs—flat connectors that protrude from each electrode
- Stack—the electrodes and separators within the battery cell, excluding tabs
- Package/Can—an outer layer that is used to house cell components
- Package sealant layer—the layer that contains the active chemicals
- Posts—the parts that connect electrode tabs to other conducting objects, such as another battery cell (also known as busbars)
- Terminal leads—current collectors from either the positive or negative electrode to the can
- Center pin—a support that is left in the core of a spiral cell to support its shape, typically made of steel or plastic
- Insulators—typically polyolefin disks that are used to prevent the ends of the jellyroll from touching the bottom or top of the package

- Gasket—an insulating sealant material that is used to electronically isolate the top cover from the can
- Top cover—a top cover for the package, typically made of nickel plated steel
- Safety vent—a rupture disk that is used to release pressure
- PTC—Positive Temperature Coefficient device that is put in series between the positive terminal lead and the top cover. The electrical resistance of the PTC greatly increases with temperature. In this way, it functions as a safety device that prevents current flow when the battery is hot, and limits high currents that can heat up the PTC.

Understanding the Battery Modeling Process

Battery modeling in STAR-CCM+ closely follows the physical structure of the battery, defined by either a stacked or spirally wound cylindrical assembly of the battery cell. The battery modeling process starts with defining a battery cell design in Battery Design Studio® (BDS), which is a separate software for battery design and simulation. In addition to cell design, BSM also incorporates the battery models and numerical algorithms of BDS into STAR-CCM+. That is, most BDS models have been made a part of STAR-CCM+ and can be used with a more detailed thermal model that is unique to STAR-CCM+.

Once the battery cell is assembled in BDS, it is imported into STAR-CCM+. For a stacked cell, the model only includes the stack and the tabs. In STAR-CCM+, the collection of tabs from all the positive electrodes in the cell is referred to as the *positive tab* for the battery cell. Similarly, the collection of

negative tabs from all the negative electrodes is referred to as the *negative tab* for the battery cell. The following screenshot shows an example of a single battery cell that has been imported into a simulation.



For a cylindrical spiral cell, the model includes a jellyroll that is made up of a positive and negative electrodes. The jellyroll is wound on a cylindrical mandrel. Also imported are the negative and positive tabs, a washer that joins all the tabs together, caps (internal posts) that connect the washer to the actual posts, and a can and end plates that enclose the cell. The following screenshot shows an example of a single cell in STAR-CCM+. The can is

hidden to show the interior parts.



In STAR-CCM+, sets of battery cells can be connected together to form a battery module.

What Is an Ecell?

Ecells define the structured electrical grid on the battery cell and are used in STAR-CCM+ to run the electrical solver portion of the battery modeling process. For each ecell, the electrical solver implements a one-dimensional battery model to compute voltage as a function of time, current, and temperature.

Ecells model actual cells, which are made up of one or more pairs of positive and negative electrodes. Each ecell consists of a positive plate, a unit cell, and a negative plate to represent a positive electrode, separator, and negative electrode of an actual cell. The ecell representation is done internally and cannot be viewed in STAR-CCM+.

Initially, the number of ecells in the x, y, and z directions for a stacked cell, and in the x and y directions for a spiral cell is specified in the imported BDS file. For a spiral cell, the ecells are not of equal size and the specified number in BDS is automatically adjusted based on the layout of the battery cell and tabs. The adjusted number is imported into STAR-CCM+ and cannot be

modified. For a stacked cell, the ecells are of equal volume and their number can be changed in STAR-CCM+.

For stacked cells, part of the modeling strategy of STAR-CCM+ is to allow you to decide the extent to which the discretization model gets simplified in order to reduce computational expense. For instance, in the following image an actual cell with one pair of positive and negative electrodes is represented in STAR-CCM+ in terms of ecells on the x-y plane. In this case, there is not any spatial variation of heat generation in the z direction.



For an actual cell with more than one pair of positive and negative electrodes, the STAR-CCM+ representation can again be modeled as one ecell in the z direction, meaning that the spatial variation of heat generation is only on the x-y plane.



Alternatively, this actual cell can be modeled in STAR-CCM+as two, three or four cells in the z-direction. Using this modeling strategy results in spatial variation of heat generation in the x, y and z directions.



Further information on using ecells can be found in Setting the resolution of the electrical mesh.

The thermal solution within the battery stack is independent of the ecell discretization. Instead, it uses its own user-defined, unstructured grid resolution. The communication between the electrical and thermal solutions is done using internal mapping between the electrical mesh and the thermal mesh.

What Is an Electrical Solver?

The BSM electrical solver calculates the electrical voltage and heat generation as well as, if applicable, chemical solution for each battery cell. It computes the voltage and heat generation on a grid of ecells incorporating the battery unit cell model type and parameters that are specified in the same imported BDS file that defines the battery cell design. The available unit cell models are NTG, RCR, and DISTNP. Regardless of whether a 1D, 2D, or 3D representation of the model is chosen in BDS, STAR-CCM+ uses the equivalent 3D representation.

What Is a Thermal Solver?

The BSM thermal solver calculates the temperature field within the battery cell stack and throughout the rest of the thermal model and cooling system.

The thermal solver takes the internal heat generation values calculated by the electrical solver and outputs the local temperature for each thermal grid cell. These values are then mapped to the ecells and used as inputs for the electrical solver.

The thermal solver uses a fixed time step that is controlled by user-defined iteration parameters.

What Is Battery Modeling Solver Coupling?

In STAR-CCM+, the battery modeling process involves running the electrical and thermal solvers sequentially for each thermal time step, starting with the electrical solver. The thermal time step is determined through a user-defined number of inner iterations. To increment the time variable, the thermal solver needs to complete this number of inner iterations. From this point forward, we refer to the thermal time step as simply time step.

At each time step, the coupled solver process is:

- Map the stack temperature distribution from the previous time step to the electrical mesh. If the current time step is the first, then start from the default initial state.
- Solve the electrical solution using the electrical solution from the previous time step as the initial condition. For the first time step, the solver starts from a default initial state. The first time step may take substantially longer than subsequent time steps as the electrical solver iterates as many times as necessary to guarantee convergence.
- Map the heat generation field from the electrical mesh to the thermal mesh.

- Solve for the thermal solution by using the number of inner iterations that you define.
- Check the solution stopping criteria and either stop or proceed to the next time step.

You control the convergence behavior of the thermal solver by setting the time step parameters. For the electrical solver, because the equations are solved at smaller, adaptive time increments, the BSM handles the solution convergence automatically.

To set an appropriate time step, you should consider the rate of change of the electrical load. The time step must be small enough so that the loading function gets resolved correctly, meaning that if the electrical load changes too abruptly between time steps the electrical solver may fail to converge. Furthermore, if the electrical load is specified using a program file and its time step is much finer than that defined for the thermal solver, the heat generation output from the electrical solver to the thermal solver might not be correctly captured. This occurs if the electrical load changes significantly several times within one time step.

The Battery Modeling User Procedure

The general workflow for battery modeling involves the steps listed below. A battery modeling tutorial is available in the tutorial guide, which introduces and walks you through the main concepts.

- Design and optimize a battery cell in Battery Design Studio® (BDS).

  This step includes saving battery cell data in a text battery model (`.tbm`) file or an encrypted battery model (`.ebm`) file, and program data in a program file (`.prg`) file, if this is the way you want to specify the load for the electrical solver.
- Activate the battery model in the physics continuum.
- Import the battery cell created in BDS into the STAR-CCM+ simulation.
- In STAR-CCM+,create and set up a battery module:
  - Specify the number of battery cells in the battery module.
  - Arrange the battery cells within the module.
  - Define the electrical cell stack resolution.
  - Prepare terminal connections and tab connections.
  - Set resistance values.
- Generate the battery module parts and connector parts.

This step may include specifying an external casing part around the battery module.

- Generate the battery module regions.
- Set up the battery mesh.

  Ensure that the external casing is assigned to a different mesh continuum than the battery.
- Generate the volume mesh.
- If necessary, set up ohmic heating for the connector parts.
- Apply an electrical load to the battery by importing the program file (`.prg`) defined in BDS or by using a table.
- Connect the battery module and the electrical load in a circuit.
- Run the simulation.
- Post-process results using the battery field functions.

Selecting the Battery Model

Selecting the Battery model within a Physics continuum reveals the Batteries manager node in the simulation tree. The Battery model is only available for solid, three-dimensional continua, and can only be used with the Implicit Unsteady solver.

To select the Battery model in the *Physics Model Selection* dialog, do the following:

- Select Three Dimensional from the *Space* group box.
- Select Implicit Unsteady from the *Time* group box.
- Select Solid from the *Material* group box.

  The Battery model is available for selection from the *Optional Models* group box.



After selecting the Battery model, you are required to select an *Energy* model and an *Equation of State* model.

What Is the Battery Solver?

The Battery solver controls the solution of the Battery model in the continua for which it is activated. The Battery solver has properties and is subject to common menu actions.



Expert Properties

Solver Frozen — Prevents or allows the solver from updating the solution while iterating. Activating this feature should not be your first choice, as it is primarily a "debugging" tool and not supported for normal simulation work.

- Activated — The solver does not update the solution.
- Deactivated — The solver updates the solution.

Working with Battery Design Studio

Battery Design Studio® (BDS) is a software package for designing, simulating, and analyzing batteries. It provides tools for costing and sizing batteries, visualizing cell designs, generating test procedures and correlating model results to performance, safety, and life predictions. BDS serves as an interface for materials suppliers, battery developers, and engineers to collaborate on the battery design.

BDS can be used to design stacked cells and spirally wound cylindrical and prismatic cells. In BDS, you can specify the physical construction of the cell, then use a simulation model to estimate how that cell behaves under various use conditions.

In order to run a STAR-CCM+ simulation on a battery, you need to create a battery cell file (.tbm) in BDS. Additionally, you can select to specify the electrical load using a program file (.prg), which is also created in BDS.

This section includes some basic information on defining the battery and program file in BDS. For further information, refer to the BDS documentation. Battery Design Studio® is only available for Windows.

Creating Battery (.tbm or .ebm) Files in BDS

In Battery Design Studio® (BDS), all the information about the battery is stored in the Text Battery Model file (`.tbm`). This information includes the physical description of the battery as well as the model parameters. STAR-CCM+ supports stacked cells and spirally wound cylindrical and prismatic cells.

For improved security, you can create an encrypted battery file. An encrypted battery file (an `.ebm` file) can be used in the same way as a `.tbm` file, but the file contents cannot be viewed or modified.

Creating a new battery cell

To create a new battery cell in BDS:

- In the BDS workspace, right-click the Batteries node on the left side of the window.
- In the context menu, click *Add new Battery* and then select the appropriate battery type. The default options for each battery type are:
  - *LiIon Pouch* > *LiIonPouch*

    Creates a stacked cell.
  - *LiIon Spiral* > *Spiral-36650-NTG3D*

Creates an idealized cell that you can use for modeling a cylindrical spiral cell. It is set up using the NTG3D model.

- *LiIon Spiral* > *Spiral-36650-NTG3D-Prismatic*

  Creates an idealized cell that you can use for modeling a prismatic spiral cell. It is set up using the NTG3D model.

The new battery cell is added under the Batteries node and the corresponding battery definition window is displayed in the workspace. You can rename the battery and specify the appropriate configuration settings.

For more information on configuration settings, see Configuring a Stacked Cell and Configuring a Spiral Cell.

Building the battery file

Once you have configured the battery in BDS, you need to build the battery file. The build process recalculates the values displayed in the Report tab and writes the battery configuration information to a `.tbm` file.

To build the battery file:

- Under the Batteries node, right-click the battery that you want to build, and then select Build.



When the build process is complete the message **Build Successful** is displayed. If any errors in the cell setup are detected, an error message is displayed.

Encrypting the battery file

If you want to distribute your battery file without the contents being exposed, you can encrypt the battery file. You must build the battery file as a `.tbm` file before you can create the encrypted battery file.

To encrypt the battery file:

- Under the Batteries node, right-click the battery that you want to encrypt, and then select Encrypt Cell.

The *Encrypt Options* dialog appears:



- In the Output File box, specify the location and name of the encrypted file that you want to create.
- Click Encrypt.

The encrypted battery file is added under the Batteries node and is saved to the specified location. You cannot modify the encrypted file; when you open the file in BDS you can only view the data sheet. If you want to modify the file contents, you need to make the appropriate changes in the `.tbm` file and then encrypt the file again.

Configuring a Stacked Cell

When you create a new stacked cell in Battery Design Studio® (BDS), the appropriate battery definition window is displayed in the workspace.

The following shows the battery definition of a LiIon pouch battery.



You can use the nodes on the left of the battery window to define characteristics and properties of the battery, including the dimensions of the cell.

Defining the Dimensions of a Stacked Cell

The cell dimensions are specified in BDS within the battery cell definition. For a stacked configuration, the battery cell imported into STAR-CCM+ includes the positive and negative electrodes and tabs, and also the separators. More information on the parameters that STAR-CCM+ uses is given below. Other cell components, such as packaging and busbars, are not imported into STAR-CCM+.

The cell dimensions depend on the following attributes:



To see a summary of the values that are used to define the cell size, select the Report node in the battery window. The cell dimensions are in the *Stack Properties* section.



Electrode Height and Width

You can define the electrode size in the battery window. In the left panel, select the appropriate node, + Electrode or - Electrode. Then, in the right-side of the window enter the appropriate height and width values for the electrode.

When you set the electrode sizes, ensure that the width of the negative tab is larger than the width of the positive tab.



Tab Height, Width, and Offset Values

You can also define the tab dimensions in the + Electrode or - Electrode node. You can specify the tab height and width, and also define the offset of the tab from the left edge of the electrode.

The tab offset values are specified from the edge of each electrode individually. Since the negative electrode is typically wider than the positive electrode, you need to set the relative position of the electrodes in order to align the tabs correctly.

Separator Offset Values

You can define the size of the separators in the battery window. In the left panel, select the Builders > Simple node, and then select the Offsets tab.

The *X* and *Y* offset values of the Negative wrt separator property specify how far the separator layer extends from edges of the *negative* electrode. For example, if the negative electrode was 88 mm wide, specifying an Xoffset value of 1 mm would result in a separator layer 90 mm wide.



Relative Position of Electrodes

You can also define the relative position of electrodes in the Offsets tab of the Builders > Simple node. The Positive wrt negative setting defines the position of the positive electrode relative to the negative electrode. The relative position is defined in the *X* and *Y* offset values. These values specify the distances from the left edge of the negative electrode to the left edge of the positive electrode, and from the top edge of the negative electrode to the top edge of the positive electrode.

By default, all these values are set to 1 mm. You need to change these values if you have changed the width or height of the electrodes, or if you want to define wider or narrower separators.

Electrode Thickness

To review the total thickness of each electrode, select the + Electrode > Formulation node, or the - Electrode > Formulation node. The total thickness is displayed under the *Computed Electrode Properties*, next to Thickness (w/collector). This value includes the thickness of the collector and coating for a single electrode.



Separator Thickness

You can define the thickness of the separators in the battery window. To set the separator thickness, select the Separator node. In the table, modify the appropriate values in the Thickness column.



Number of Electrodes in Stack

You can define the number of electrode layers in the stack. In the battery window, select the Builders > Simple node, then select the Stack tab. The Total # of electrodes including ends box specifies the total number of electrodes in the cell.

The selected option for *End Electrodes* affects the values that you may specify. If you select Both Negative or Both Positive, the total number of electrodes must be odd. If you select Pos(right)/Neg, the total number of electrodes must be even.

The number of separators in the cell is equal to the (*total # electrodes + 1)*, that is, a separator layer between adjacent electrodes and an additional separator at either end of the stack.



Configuring a Spiral Cell

You can use Battery Design Studio® (BDS) to design spirally wound cylindrical and prismatic cells: the significant difference is the shape of the mandrel around which the cell is wound. In cylindrical cells, the mandrel is a cylinder; in prismatic cells, the mandrel is a geometric prism.

When you create a new cylindrical spiral or prismatic spiral cell in BDS, the appropriate battery definition window is displayed in the workspace.

The following shows the battery definition of the default LiIon cylindrical spiral battery.



The following shows the battery definition of the default LiIon prismatic spiral battery.



You can use the nodes on the left of the battery window to define characteristics and properties of the battery.

For both the cylindrical spiral cell and the prismatic spiral cell, you are recommended to start with the default dimensions. However, you can specify the tab configuration and the numerical settings that are used to map data to the thermal mesh in STAR-CCM+.

For more information, see Tab Vertical Alignment, Tab Type, Tab Angular Alignment, and Number of Spokes.

Tab Vertical Alignment

STAR-CCM+ supports same-face tabs for spiral cells, that is, battery cell configurations where both the positive and negative tabs are on the same side of the cell.

You can check the vertical alignment of the tabs in the Builders > Detailed node, in the Vertical Alignment tab. The example below shows the default cylindrical spiral cell, in which the tabs have opposite orientation.

If you want to modify the orientation of the tabs, make the appropriate changes in the Vertical tab orientation box.



Tab Type

Both the default cylindrical spiral cell and the default prismatic spiral cell have the following tab types defined:

- Positive electrode: Coated foil with center tab
- Negative electrode: Coated foil with 1 tab on side

However, you can specify a different type for either or both electrodes. To modify the tab type assigned to a particular electrode, select the appropriate node, +Electrode or -Electrode. From the *Type* drop-down menu, select the new tab type.



If you select a tab type that uses multiple tabs, you need to specify the number of tabs to use.



- To specify the number of tabs, in the Fix # tabs box, enter the appropriate number.

Tab Angular Alignment

You can specify the angular alignment of the tabs in the winding. For spiral cells, the tabs are used to orient the local coordinate systems (used for mapping) within the cell. If you do not specify the tab alignment, the local coordinate system origin and the Y-axis (axially along the cell) are defined, but there is no way to define the X-axis and Z-axis of the cell.

The alignment angle is with respect to the mandrel and the zero angle is the positive X-axis. The angle is measured in the anti-clockwise direction for cylindrical cells and the clockwise direction for prismatic cells, as viewed from the positive Y-direction.

In previous versions of BDS, the tab angle was measured with respect to the start of the electrode. The tab angle is now measured with respect to the mandrel, regardless of the starting point of the electrode.



- To align the tabs in the winding, activate the Align tabs in winding box and then, in the Alignment angle, degrees box, specify the angle at which the tabs begin.

You can use the tab alignment setting to specify a particular angle (A_{sep}) between the tabs. To do this, specify the alignment angles of the tabs in the positive electrode (A_{pos}) and the negative electrode (A_{neg}) such that the difference $A_{pos} - A_{neg} = A_{sep}$.

Number of Spokes

Spokes are used in the discretization of the spiral cell and affect how the cell is wrapped. The number of spokes determine the location on the cell from where the data is mapped to the thermal mesh. You can reduce the number of spokes to speed up calculations or increase it to improve the accuracy of the model.

For a cylindrical cell, you can specify the number of spokes. The default value in the provided cylindrical cell is 50. The range of valid values is 50 - 1000.

To set the number of spokes for a cylindrical cell, select the Builders > Detailed node, and then select the Numerical Settings tab.



- In the Number of Spokes box, specify the appropriate number.

If necessary, you can modify the number of spokes value in STAR-CCM+, but the same range of valid values applies.

For a prismatic cell, you specify the number of spokes in relation to the size and shape of the mandrel.

To set the number of spokes in a prismatic cell, select the Builders > Detailed node, and then select the Numerical Settings tab.



- In the Side Increment box, specify the appropriate value.

  The number of spokes on each side of the mandrel is the side length divided by the increment. You can view and modify the mandrel width and thickness in the Winding Stack tab.

- In the Corner Angle box, specify the appropriate angle.

  The number of spokes on each corner is 90 divided by the specified angle. The minimum value is 10 degrees (9 divisions) and the maximum is 22.5 degrees (4 divisions). If you specify an angle that does not divide evenly into 90, BDS automatically uses the nearest integral number of spokes.

Preparing Program (.prg) Files

In Battery Design Studio® (BDS), procedures for cycling batteries are stored in program files (`.prg`). In a program file, you can specify periods when the battery is charging, discharging, and resting. The program file is one way to apply an electrical load to the battery in STAR-CCM+.

To create a procedure in BDS, right-click the Procedures node and select Add new procedure > IET Cycler.



A new cycling procedure is added under the Procedures node, and a program definition window is displayed in the workspace.



When the procedure is created, a `.prg` file is automatically created in the workspace directory. You can use the table in the program definition window to set up the cycle.

Selecting a Step Type

In the Step column, click the <add> button and select from the following types of step:

- Advance

  Increments the cycle counter.

- Rest

  Sets the applied current to zero.

- Charge

  Sets the sign of the applied current to negative (the convention for BDS charging).

- Discharge

  Sets the sign of the applied current to positive (the convention for BDS discharging).

- End

  Ends the program.

Setting Control Conditions

After selecting a step type, you can specify the control, end, and limit conditions for the step. The Control Condition identifies the variable to control during the step, such as current, voltage, power, or resistance. The End Condition identifies the value that, when reached, terminates the step.

Setting End Conditions

STAR-CCM+ fully supports Time end conditions. Voltage end conditions are partially supported.

The Time end condition has units of seconds. To specify an end condition of 1 hour, for example, select Time in the *Type* column and enter `3600` in the *Value* column.

The Voltage end condition has units of volts. To specify an end condition of 10 volts, for example, select Voltage in the *Type* column and enter `10` in the *Value* column.

If the step type is a discharge load, the voltage end value is interpreted as a minimum voltage. The simulation ends with a message indicating that the minimum cell voltage limit has been reached.

If the step type is a charging load, the voltage end value is interpreted as a maximum voltage. The simulation ends with a message indicating that the maximum cell voltage limit has been reached.

When a voltage end condition is specified, the battery simulation continues with the specified load type until the end of the STAR-CCM+ time steps, or until the voltage end condition is met, whichever comes first. Any remaining lines of input in the program file are not used.

Setting Limit Conditions

STAR-CCM+ uses Voltage limit conditions to specify a warning message.

The Limit Condition is an optional parameter that allows you to specify maximum or minimum values. The most common use of a limit condition in BDS is to limit the voltage of the battery during the cycle. Within STAR-CCM+ a battery can consist of multiple modules, with each module containing more than one battery cell. In STAR-CCM+, the current limit condition option uses the maximum and minimum values to print a

warning message when any battery cell exceeds the specified limits. The circuit solver does not switch to a new load type as is the case in single cell battery simulation in BDS.

If the step type is a discharge load, the limit value is interpreted as a minimum voltage. The circuit solver prints a warning message for every time step in which one or more battery cells reach the lower limit voltage.

If the step type is a charging load, the limit value is interpreted as a maximum voltage. The circuit solver prints a warning message for every time step in which one or more battery cells reach the limit voltage.

The simulation does not stop when a limit condition is reached. You can specify a limit condition on each line of input in the program file that is used to define a load.

Setting Up Reports

In the Report column of the program file, you can specify the frequency at which data is recorded. You can specify the frequency by time and/or voltage change.

Setting Temperatures

In the Temp column, you can specify the temperature at which to carry out the step. If this column is left blank, BDS uses the Initial Battery Temperature. This value is defined at the bottom of the program file window.

Working with a Battery Cell

A battery cell is a single unit that is used to store energy. Battery cells are defined in Battery Design Studio® (BDS), and stored in a .tbm file or in an encrypted .ebm file. Battery cells are used in STAR-CCM+ to create battery modules, which contain a number of battery cells that are connected in series and/or parallel. STAR-CCM+ supports stacked cells and spirally wound cylindrical and prismatic cells.

Creating a Battery Cell

To use a battery cell in STAR-CCM+, you import the .tbm or .ebm file into the simulation. By default, all of the battery parts that are specified in the file are imported. However, if you want to use your own geometry components for some parts, you can choose to leave out the corresponding parts from the imported file.

Creating a Battery Cell

To create a battery cell:

- Right-click the Batteries > Battery Cells node and then select Create from Tbm...



- In the *Import Battery Data From File* dialog, select the appropriate .tbm or .ebm file and then click Open.



- In the *Import Battery Options* dialog, select the battery parts that you want to import.

  By default, all parts of the battery cell are selected. However, you can exclude some parts and replace them with your own geometry.

  For more information, see Selecting Optional Parts for a Stacked Cell, Selecting Optional Parts for a Spiral Cell, and Adding External Posts to a Spiral Cell.

  The active material (the stack for stacked cells and the jellyroll for spiral cells) is always imported.

- Click OK.

STAR-CCM+ automatically recognizes the cell models and updates the appropriate properties of the battery cell, and reads in the electrical mesh settings that are specified in the battery model.

Selecting Optional Parts for a Stacked Cell

For stacked cells, the optional parts are as shown below:



- If you want to exclude some parts from the imported file, deselect the appropriate parts.

The active material (the stack) does not appear on the list of optional parts; it is always imported.

Selecting Optional Parts for a Spiral Cell

For spiral cells, the optional parts are as shown below:



- If you want to exclude some parts from the imported file, deselect the appropriate parts.

The active material (the jellyroll) does not appear on the list of optional parts; it is always imported.

Adding External Posts to a Spiral Cell

If you are importing a cell that has same face tabs, you have the option to add external posts to the battery model as part of the import process. The external posts may be complex CAD entities that you have created outside BDS. You need to import the appropriate parts into STAR-CCM+ before you can add them to the battery model.

- Under Select Post Parts to Import, click the appropriate option: -ve Post or +ve Post.

  The *Set Connection Geometry Parts* dialog appears:



- In the *Set Connection Geometry Parts* dialog, select the parts that you want to add, and then click OK.

The selected parts are displayed under the Spiral External Posts node, under the new Cylindrical Cell node.

Configuring a Battery Cell

When the battery cell is imported into the simulation, STAR-CCM+ reads the dimensions of the cell from the .tbm or .ebm file and generates an idealized battery cell 3D-CAD model. This model is displayed in the *Graphics* window in a *3D-CAD View* scene.

The corresponding Battery Cell node, Cylindrical Cell node, or Prismatic Cell node is added to the tree under the Batteries > Battery Cells node.



The new Battery Cell node, Cylindrical Cell node, or Prismatic Cell node has properties and a pop-up menu, and is also subject to common menu items.

Battery Cell Pop-Up Menu

The following options are available when you right-click on a battery cell node.

Reload...	Reloads the battery cell data from a .tbm or .ebm file. This option allows you to update the battery cell if you change the .tbm or .ebm file.
Export...	Exports the battery cell data to a .tbm or .ebm file. This option allows you to recreate a file from the battery cell information that is stored within the simulation. You cannot choose the export format: it is the same as the format of the original file that you imported.

Stacked Cell Properties

The following properties are available for stacked battery cells.

Electrical Mesh Dimensions	Defines the user-specified electrical mesh resolution.
TBM File	Displays the path to the imported .tbm file or .ebm file.

Unit Cell Model	Identifies the battery model that is defined in the .tbm file for the imported battery cell. For information on the available battery models, see Battery Modeling Formulation. For empirical models, you can override the cell capacity or state of charge settings. You cannot override any settings for a battery cell that uses a physical model. For more information, see Setting Cell Data Overrides.

Spiral Cell Properties

The following properties are available for spiral battery cells.

TBM File	Displays the path to the imported .tbm file or .ebm file.
Unit Cell Model	Identifies the battery model that is defined in the .tbm file for the imported battery cell. For information on the available battery models, see Battery Modeling Formulation. For empirical models, you can override the cell capacity or state of charge settings. You cannot override any settings for a battery cell that uses a physical model. For more information, see Setting Cell Data Overrides.
Specified Electrical Mesh Dimensions	Defines the user-specified electrical mesh resolution. The default is the mesh that is specified in the .tbm file. You can modify the mesh dimensions if necessary. For information on modifying the mesh dimensions, see Setting the Resolution of the Electrical Mesh.

Actual Electrical Mesh Dimensions	Defines the electrical mesh resolution that is computed from the specified layout of a spiral cell and the specified electrical mesh dimensions. STAR-CCM+ uses this resolution. This property is read only.
Number of Spokes	Defines the number of spokes. Spokes are a component of the electrical mesh. You can reduce the number of spokes to speed up calculations or increase it to improve the accuracy of the model. For a cylindrical cell, you can specify the number of spokes. The range of valid values is 50 - 1000. You can also change this value in the Specified Electrical Mesh Dimensions property. For more information on configuring spokes for a spiral cell, see Number of Spokes.
Positive Post Parts	If you have added external posts to the battery model, this property specifies the positive post parts. You can modify the selection as necessary. If you have not added any external posts, this property is empty. If you want to add external post parts, click the ... button and then, in the selection dialog, select the appropriate parts. For more information, see Adding External Posts to a Spiral Cell.

Negative Post Parts

If you have added external posts to the battery model, this property specifies the negative post parts. You can modify the selection as necessary.

If you have not added any external posts, this property is empty.

If you want to add external post parts, click the ... button and then, in the selection dialog, select the appropriate parts.

For more information, see Adding External Posts to a Spiral Cell.

Setting the Resolution of the Electrical Mesh

The electrical mesh defines cells on the battery stack or jellyroll strip. Each cell in the electrical mesh represents an ecell in the electrical calculation. Although you must define the electrical mesh dimensions, it is created internally and cannot be visualized.

The initial resolution for the electrical mesh depends on the battery model that is selected in BDS, and is read into STAR-CCM+ during the import process. The electrical mesh is used during the simulation only for the electrical calculations; a thermal mesh is used for the energy calculations.

You can change the resolution of the electrical mesh to suit the requirements of your simulation.

Under the Battery Cells node, select the appropriate battery cell. The electrical mesh dimensions are displayed in the *Properties* window.

For a stacked plate cell, you can specify the number of mesh cells in the X, Y, and Z directions.



For a cylindrical cell, you can specify the number of mesh nodes in the X and Y directions. You can also specify the number of spokes.



Within BDS, the specified electrical mesh dimensions are used to compute the resolution that STAR-CCM+ uses for its electrical solver. The actual dimensions are derived from the specified resolution and the minimum number that is feasible based on the layout of the jellyroll and tabs.

Typically, the electrical mesh is less dense than the thermal mesh or the surrounding fluid mesh. For many applications, it is recommended that you start by using a coarse electrical mesh and increase the mesh density only if necessary. For a stacked cell, a coarse electrical mesh is 10x10x1 or less. The solution time varies linearly with the addition of battery cells to a battery module.

For a stacked cell, you set the value of the electrical mesh resolution in the Z-direction, according to the thickness of the battery cell. For thin cells, there is no advantage in using more than one ecell as there is not normally a significant temperature gradient in the normal direction.

For a spiral cell, the *minimum* mesh resolution is determined automatically in BDS based on the layout of the cell. If the specified resolution is lower than this minimum, BDS increases it accordingly. For a stacked cell, the minimum mesh resolution is 1x1x1.

The thermal mesh for the battery cell is defined separately in STAR-CCM+ in a mesh continuum. Mesh settings for other parts of the battery, such as tabs and connectors, are also defined in a mesh continuum. These settings are covered in generating battery mesh.

Viewing the Battery Cell Properties

The Cell Properties node under each battery cell node contains the properties that are imported from the .tbm file. The properties are organized in a tree structure that is similar to the cell design tree in BDS.

For more information on the BDS cell design tree and the properties that you can specify when you design a battery cell, see the *Battery Design Studio User Guide*.

Most of these properties are read-only and you cannot change them in STAR-CCM+. The only exceptions are the cell capacity and state of charge settings, for which you can set override values. For more information, see Setting Cell Data Overrides.

If you want an overview of the properties of a particular battery cell, you can view the cell summary report. For more information, see Viewing a Cell Summary Report.

Stacked Cell Properties

The properties tree for a stacked cell is shown below:



The stacked cell properties tree contains the following major nodes:

Builders	Simple Builder	Specifies how the components are placed together to make the completed cell. The simple builder specifies the orientation of the plates within the cell, the size and configuration of the stack, and any protective overwrap that is given to the electrode stack. The simple builder also specifies the offsets that the separator introduces between the electrodes and the pouch.
Models	Electrolyte	The General Electrolyte model is always specified. This model describes the behavior of the electrolyte within the cell.

	IET	One IET model is specified per cell. This model describes the electronic and thermal behavior of the cell. STAR-CCM+ supports the NTG model, the RCR model, and the DISTNP model. Some other models that are available in BDS are not supported. For more information on the IET models that are available in Battery Design Studio, see the *Battery Design Studio User Guide*.
Physical Cell Description	Cell Voltage	Specifies the current state of the cell, how the lithium is distributed between the two electrodes, and the allowable voltage range of the cell.
	Electrolyte	Specifies the formulation of the chemical mix that makes up the electrolyte. An electrolyte is a combination of multiple salts and solvents, and each ingredient is specified separately. For more information, see Electrolyte Properties.
	Internals	Specifies the fillers that are inserted in cell packaging alongside the active components. These items are used to provide shape and stability to the components within the packaging and are effectively inert.
	Negative Electrode Positive Electrode	Specifies the physical and thermal properties of the electrode, the formulation (the chemical substance that makes up the electrode), the tabs, and the collector that connects the electrode to the tab. The busbars are also specified. The two busbars span the cell, joining tabs together. For more information, see Electrode Properties.
	Package	Specifies the packaging of the cell. The packaging affects the thermal performance of the cell.
	Separators	Specifies the separator that is used to separate the various chemically reactive components within the cell to limit reaction rates.

Electrolyte Properties

STAR-CCM+ supports the CAEBAT electrolyte property sets that are available in Battery Design Studio. These property sets are based on a model that was developed by K. Gering and were generated as part of the Computer-Aided Engineering of Batteries (CAEBAT) program.

For more information, see *Abstract #593, 224th Meeting of the Electrochemical Society*:

https://ecs.confex.com/ecs/224/webprogram/Abstract/Paper20427/B4-0593.pdf

For more information on the electrolyte property sets that are available in Battery Design Studio, see the section *Using Predefined Electrolyte Property Sets* in the *Battery Design Studio User Guide*.

Electrode Properties

The example below shows the expanded Negative Electrode node:



For more information on the electrode properties, see the *Battery Design Studio User Guide*.

Cylindrical Cell Properties

The properties tree for a cylindrical cell is shown below:



The cylindrical cell properties are similar to the stacked cell properties. The cylindrical cell properties tree contains the following major nodes:

Builders	Detailed Builder	Specifies how the components are placed together to make the completed cell. The detailed builder specifies the orientation of the cell tabs, the winding and vertical alignment of the cell.
Models	Electrolyte	See the electrolyte description in the stacked cell properties tree.
	IET	See the IET description in the stacked cell properties tree.
Physical Cell Description	Cell Voltage	See the cell voltage description in the stacked cell properties tree.

Electrolyte	See the electrolyte description in the stacked cell properties tree.
Internals	See the internals description in the stacked cell properties tree.
Negative Electrode Positive Electrode	Specifies the physical and thermal properties of the electrode, the formulation (the chemical substance that makes up the electrode), the tabs, and the collector that connects the electrode to the tab. The tape that connects the tab to the electrode is also specified. For more information, see Electrode Properties.
Package	See the package description in the stacked cell properties tree.
Separators	See the separators description in the stacked cell properties tree.

Viewing a Cell Summary Report

A cell summary report provides an overview of the battery cell properties. This information is taken directly from the <REPORT> section of the .tbm file.

To view a cell summary report for a particular battery cell:

1. Right-click Batteries > Battery Cells > [Battery Cell] > Cell Properties.



2. In the pop-up menu, click Summary Report and then click one of the following:
   - Browser: Displays the summary report in your default browser.
   - To File: Saves the summary report as an HTML file in a specified location.

An example of a cell summary report is shown below:

Cell Report: BatteryCellProperties

1	**Cell Properties**	State of Charge, %	100
		Voltage, V	3.65879
		Capacity, Ahr	10.0331
		Energy, Whr	36.7092
		Energy Density, Whr/kg	248.748
		Energy Density, Whr/liter	572.008
		Weight, g	147.576
		Volume, cm^3	64.176
		Material Cost, $	4.95685
		Active Area, m^2	0.153768
		Unit Capacity, mAh/cm2	6.56695
		C/A Ratio, mAh/mAh	1.50999
		Electrolyte Mass, g	18.7778
		Electrolyte Volume, cm^3	15.9594
		Separator Area, m^2	0.16524
		Heat Capacity@25C, J/g-K	0.871747
		Cell Dimensions Length, mm	190
		Cell Dimensions Breadth, mm	100
		Cell Dimensions Thickness, mm	4.6156
2	**Active Material Properties**	Energy Density, Whr/kg	288.99
		Energy Density, Whr/liter	609.261
		Weight, g	127.026
		Volume, cm^3	60.252
		Height, mm	153
		Width, mm	90
		Thickness, mm	4.3756
3	**Electrode Properties**		
+-1	**Negative Electrode**	Average Voltage, V	0.141296
		Stoichiometry at formation	0.9
		Unit Capacity, mAh/cm^2	6.56695
		Thickness (w/collector), um	300

Setting Cell Data Overrides

You can override the cell capacity or state of charge settings for a battery cell that uses an empirical model (such as the NTG model or the RCR model). However, you cannot override any settings for a battery cell that uses a physical model (such as the DISTNP model).

If you change the cell capacity or state of charge settings during a simulation, the new values do not take effect immediately. To apply the new values to the simulation, you must clear the solution.

When you load a simulation into STAR-CCM+, the restored simulation uses the cell capacity or state of charge settings that were in effect when the simulation was saved. However, if you clear the solution, any override settings are lost: the simulation reads the cell capacity and state of charge settings from the original .tbm file. If you want to use the same values again, you must re-enter the appropriate cell capacity or state of charge override settings.

To override the cell capacity or state of charge settings:

1. Under the appropriate battery cell node, expand the Cell Properties > Models > IET > Model Name > Override Cell Data node.



2. Select the Capacity node or the Initial SOC node.
3. Set the Cell Data Override Option property to Override.

The corresponding node is added. The example below shows setting the override for the Initial SOC node.



4. Select the Capacity or SOC node and set the corresponding override property:
   - Value (for Capacity): Set the cell capacity in Ah or coulombs.
   - Percentage (for SOC): Set the cell state of charge as a percentage of its capacity.

Note: When you run a simulation, if the battery cell state of charge reaches 0% or 100%, the state of charge of the cell becomes fixed at that value. The simulation continues to run, but uses the constant value that the model computes at the fixed state of charge. Obviously, a real battery cell does not behave in this way.

For more information, see What Is the Circuit Solver?.

Setting Up a Battery Module

The battery module is the core component in battery module simulations. It contains sets of battery cells, and is connected directly to the electrical load.

The process for setting up a battery module involves the following steps:

- Creating a battery module
- Defining the number of cells in the module
- Setting up the module cells
- Setting up terminal connections
- Setting up tab connections
- Specifying battery module resistance values

Once these steps have been carried out, you can generate the parts, connector parts, and regions of the battery module.

What Is a Battery Module?

A simulation can contain several battery modules. A battery module consists of one or more battery cells. The cells within a module are electrically connected in a series/parallel configuration. A set of cells that are connected in series is referred to as a *string* and the strings of cells are connected in parallel. The number of cells that are connected in each string is N Series, and the number of strings that are connected in parallel within each module is N Parallel. The total number of cells in a module is the product of N Series and N Parallel. The circuit solver for each module automatically makes electrical connections according to the user-designated values for N Series and N Parallel.

The electrical connections between cells are not explicitly displayed in STAR-CCM+. However it is possible to observe the connections after an electrical time step solution using the Tab Voltage field function. Each battery cell in a module has a positive and negative tab. The Tab Voltage function displays the circuit voltage that is computed at each tab that is displayed in a module.

You can add physical connectors between the cells of the battery modules. The main purpose of adding physical tab and terminal connections is to model the thermal solution more accurately. The physical connectors that you add must be consistent with the electrical connections that the circuit solver makes automatically. Within a string of cells, tab connectors join a positive or negative tab of one cell to the complementary tab of an adjoining cell. The open tabs at either end of the string of cells are referred to as

terminals and these connections can be physically modeled using terminal connections. Both tab and terminal connections can be set up as illustrated below.



For spiral cells, a battery module looks as follows, showing two parallel strings of cells that are connected in series.



In the above examples, alternate cells have been rotated in order to align the positive tab of one cell with the negative tab of its neighbor.

Battery modules are stored under the Battery Modules manager node. Each Battery Module node has properties and a pop-up menu.



Battery Module Pop-Up Menu

The following options are available when you right-click on a battery module node.

Create Circuit Element	Creates a circuit element from the battery module. A circuit element is required to add the battery module to a circuit.
Generate Battery Parts	Generates the battery parts.
Generate Connector Parts	Generates the battery connector parts.
Create External Casing Part	Creates external battery casing part.
Generate Regions...	Generates the battery regions.
Setup Connector Ohmic Heating	Sets up connector ohmic heating.

Battery Module Properties

Mesh Continuum	Assigns a mesh continuum to the battery module.
Index	Provides an index value for the battery module.
Module Local Coordinate System	Sets the local coordinate system that is associated with the battery module.
N Series	Sets the number of battery cells that are connected in series in each string.
N Parallel	Sets the number of strings in the module. Each string is connected in parallel.

See also: common menu items.

Creating a New Battery Module

To create a battery module:

- Right-click the Battery Modules node and select New.



Defining the Number of Cells in the Module

The number of battery cells in the battery module is defined by specifying the values N Series and N Parallel. N Series in the number of cells that are connected in each series string, and N Parallel is the number of series strings that are connected in parallel. Strings are connected in parallel, and each string contains the same number of battery cells.

You can set the number of cells in the battery module in the properties of the Battery Module node: N Series and N Parallel. These properties are shown in the following diagram using stacked cells.



To set the number of cells in the battery module:

1. Select the Battery Modules > Battery Module node.
2. In the *Properties* window, enter the values of N Series and N Parallel.

The number of cells that are shown in the battery scene is not updated until you generate the battery parts.



You can also set N Series and N Parallel using the Battery Module > Model Configuration > Module Layout > Module Grid Size menu or by using the Setup Battery Module Cells option.

Defining the Module Configuration

For most cases, the Module Configuration node is the recommended module setup option. You specify the Battery Cell property before you can configure any module setup parameters. The other configuration options are not visible until you have selected a battery cell.

- Select the Module Configuration node and, in the Battery Cells property, specify the battery cell that the battery module contains.



STAR-CCM+ adds a Battery Module Configuration sub-node, along with Module Cell Orientation Conditions, Module Connection Values, and Module Layout subfolders to the Module Configuration node.

The Module Configuration node allows you to change the position of each individual battery module by specifying its global offset vector. The offset is defined in terms of the local coordinate system that is associated with the battery module. The offset is used to set the battery module layout using uniform spacing between cells.

- In the Global Offset property, specify the appropriate global offset vector.

If you define the position and orientation of the cells using the module configuration tool, STAR-CCM+ uses these settings instead of the settings that are set in the Battery Module Cells node.

Setting the Battery Module Configuration

The Battery Module Configuration node defines how the battery cells are configured in the module:

- Grid—cells are placed on a regular grid that defines the number of cells in series (rows) and parallel (columns).

- Staggered—spiral cells are placed in staggered rows. The number of cells in each series string determines the number of cells. The number of parallel strings defines the number of staggered rows in the configuration.
- End to End—spiral cells are stacked end to end based on the number of cells in series. The number of parallel strings sets the number of stacks.

The staggered and end to end options are only available for spiral cells.



The following screenshot shows the different configurations using three cells in series and two cells in parallel. From left to right, it shows grid, staggered, and end to end configurations.



Specifying the Module Cell Orientation

In a battery module, the cells are connected in an N Series X N Parallel grid. For a module with N Series = 6 and N Parallel = 3, there are 18 cells. Cell 0,0 is connected in series to Cell 0,1, Cell 0,2, Cell 0,3, Cell 0,4, and Cell 0,5. Similarly, Cell 1,0 to Cell 1,5 (and Cell 2,0 to Cell 2,5) are connected in series.



The tab locations in an individual battery cell are determined from the `.tbm` file. You determine the cell configuration within a battery module as described in Setting the Battery Module Configuration. Typically, the cells are oriented as shown in the previous figure. The negative tabs of cells 0,1 -

0,4 are connected to the positive tabs of cells 0,2 - 0,5 respectively. For the parallel connections, the positive terminals of Cell 0,0, Cell 1,0, and Cell 2,0 are connected to the negative terminal of the load. Similarly, the negative terminals of Cell 0,5, Cell 1,5, and Cell 2,5 are connected to the positive terminal of the load.

STAR-CCM+ offers an option to flip the cells within each string so that the negative tabs of each cell align with the positive tabs of the adjacent cells. Flipping the cells does not change the electrical connection: flipping the cells make tab connectors that follow the electrical circuit connections easier to model.

You set this property in the Module Cell Orientation Conditions > Cell Orientation Setting node. To flip the series cells, set Option to Flip Every Other.



You can also leave the tabs unaligned (for example, in setting up a string of cylindrical cells that uses the staggered cell arrangement).

- Set Option to Do Not Flip.

Note: Changing the cell orientation does not change the electrical connections. Take care to add tabs consistent with the underlying electrical connections.

To change which side of the displayed cell corresponds to the open positive tab, you can flip either the odd or the even cells. You set this property in the Module Connection Values > Series Connection Flip Option node.



The first cell in a row of a module (for example, `Cell 0,0`) is an odd cell. The second cell in a row of a module (for example, `Cell 0,1`) is an even cell.

- To flip the odd cells, set Option to Flip Starting With First Cell.
- To flip the even cells, set Option to Flip Starting With Second Cell.

Defining the Module Layout

In the sub-nodes of the Module Layout manager node, you can set the position of the cells in the configuration and the size of the battery module.



Setting the Module Positioning

The Module Positioning node lets you specify the properties that are appropriate to the module configuration. The settings for each of the possible battery module configurations are summarized in the following tables:

Module Grid Positioning Properties

Offset Type	Specifies if the parallel and series offsets are absolute values or are based on the size of the battery cell.
Parallel Offset	Specifies the parallel offset in the units that are appropriate to the specified Offset Type.
Parallel Direction	Specifies the vector direction that is applied to the parallel offset.
Series Offset	Specifies the series offset in the units that are appropriate to the specified Offset Type.
Series Direction	Specifies the vector direction that is applied to the series offset.
Local Coordinate System	Sets the local coordinate system that is associated with the battery module.

The default offset and spacing between cells is set to 0, which produces a valid configuration of cells in contact with each other. For a staggered configuration, if you use the default positioning, the cells are offset based on the exterior diameter of the can.

If you want to perform a global rotation on the battery module, create a coordinate system in the appropriate orientation and then apply it to the battery module. The coordinate system handles the rotation and the global offset.

Module Staggered Positioning Properties

Staggered Row Offset	Specifies the geometric offset between staggered rows.
Staggered Row Direction	Specifies the vector direction that is applied to each row of staggered cells. The default is [1.0, 0.0, 0.0].
Gap Between Cells	Specifies the distance between each cell in a staggered row.
Local Coordinate System	Sets the local coordinate system that is associated with the battery module.

Module End to End Positioning Properties

Stack Offset	Specifies the geometric offset between stacks.
Local Coordinate System	Sets the local coordinate system that is associated with the battery module.

Setting the Module Size

The Module Size node lets you specify the module size parameters. The default size of the battery module is the same as the number of cells in series and parallel that you set (the N Series and N Parallel properties) using the Battery Module node.

The settings for each of the possible battery module configurations are:

Module Grid Size	• N Series • N Parallel
Module Staggered Size	• Number of Cells in Staggered Row (the default value is the N Series setting) • Number of Staggered Rows (the default value is the N Parallel setting)

Module End to End Size	• Cells in Stack (the default value is the N Series setting) • Number of Cell Stacks (the default value is the N Parallel setting)

Setting Up Battery Module Cells

As an alternative to using the Module Configuration feature, you can define the position and orientation of each cell within the module using the battery module cell setup tool. This option lets you set non-uniform spacing in your cell formation.

This option also lets you align the series cells along different axes and define the rotation of cells that are flipped along any axis or angle. However it is slower to use as you set up each parallel string of cells manually.

If you configure your module using the Module Configuration node, these settings supersede the ones that are defined using the Battery Module Cells node.

To set up battery module cells:

1. Right-click the Battery Module > Battery Module Cells node and select Setup Battery Module Cells.



The *Setup Battery Module Cells* panel appears.



For a battery module containing multiple strings, you can specify the spacing between the parallel strings using the *Parallel String Spacing* box.

2. Enter a value for the Offset Vector.
3. Click Set.

Each successive parallel string is offset from the previous string by the specified amount.



After setting the spacing of each parallel string, you can define the spacing of the cells within each string. You work on each string independently. To select the string that you wish to work on, use the Select Parallel String box. Choose the local Cartesian coordinate system for your battery module from the Battery Module Coordinate System drop-down menu. You can create one interactively by either clicking the Define New Coordinate System button or using the coordinate systems tool.

In the *Cell Spacing* box, enter an Offset Vector for the cells in the selected string and click Set. The position of the cells in the string is updated. For modules containing multiple parallel strings, select the next string using the Select Parallel String box and enter the desired cell spacing as described above.



The table at the bottom of the dialog shows a summary of the position of each cell in the selected string. You can use this table to change the position of individual cells by typing a position vector directly into the Offset column. If a battery cell has been imported into the simulation, STAR-CCM+ automatically assigns this cell to each position in the battery module. If more than one cell has been imported into the simulation, you

can choose to assign a different battery cell to a particular position in the module. You assign a battery cell to a position by using the drop-down menu in the Battery Cell column.



In order to align the negative and positive electrodes of adjacent cells, flip some of the battery cells in the module:

1. Select the axis about which the cells are flipped from the Select Cell Rotation Axis menu in the *Cell Rotation Axis* box.
2. Activate the Flip? checkbox for any cells that you wish to flip.

   The rotation value in the Rotation column is updated automatically.



You can now also define a new Cartesian coordinate system for your battery cells:

1. Click the Define New Coordinate System button.

   In the *Edit* panel, change the axis definition of the coordinate system.
2. When you have completed the battery module setup, click Apply, and then click Close.

You can return to this panel to change the position of the cells, if necessary.

Setting Properties of Battery Module Cells

The cells that are defined in the battery module are stored under the Battery Module > Battery Module Cells node. You can define the position, type, and orientation of each cell in the module by specifying the properties of the

battery module cells node. These settings can also be applied using the *Setup Battery Module Cells* panel, and are updated at the end of the battery module cell setup.



Battery Module Cell Properties

Battery Cell	Selects the imported battery cell.
Orientation	Specifies whether the cell is flipped (rotated through 180 degrees). Activated — Cell is flipped. Deactivated — Cell is not flipped.
Rotation	Specifies the rotated position of the cell.
Translation	Specifies a position vector for the cell.
I Series	Indicates the *i* component for the cell index, in a module of size `[i, j]`.
I Parallel	Indicates the *j* component for the cell index, in a module of size `[i, j]`.

Creating an Unrolled View of a Battery Cell

You can create an unrolled view of the jelly roll for prismatic and cylindrical battery cells in a module. The unrolled view provides a picture of the internal connectivity of the electrical mesh and the battery cell structure.

The inner and outer e-cell centroids in the jelly roll mesh are averaged to develop a spiral representation. This representation is unrolled along the spiral path that is defined by the winding of the electrodes from the cell definition, and displayed as a 2D view.

STAR-CCM+ calculates the cell height as illustrated below:



The Vertical tab orientation settings in the battery cell can have a significant effect on the calculated jellyroll height. For example, the overhang that is caused by the difference in height of the coated and uncoated parts of the electrodes can increase the cell height significantly.

Note: In some batteries, the negative and positive electrodes can be different sizes. This size difference results in regions outside of the jelly roll that contain only one electrode and therefore contain no real e-cells. When the jelly roll is unrolled, the composite view has a non-rectangular outline.

To create an unrolled view of a prismatic or cylindrical battery cell:

1. Expand the Battery Modules > Battery Module > Battery Module Cells node.

2. Right-click the appropriate battery module cell and click Create Jelly Roll Unroll Part.



The new part appears under the Derived Parts node.

The new part is named after the battery cell and is appended with Unroll.



Jelly Roll Unroll Part Properties

Battery Module Cell	Selects the battery module cell.	
No Data Value	Specifies the value that is used when no e-cell centroid data is available.	
Plate	Specifies the plates that are unrolled and displayed.	
	Positive and Negative	Both the positive and negative plates are displayed.
	Positive	Only the positive plates are displayed.
	Negative	Only the negative plates are displayed.
E-cells	Specifies the e-cells that are unrolled and displayed. For any of the e-cell centroids that do not have values, the No Data Value is used.	
	Inner and Outer	Data is displayed if both the inner and the outer e-cell centroids have values. The average value is used.
	Inner	The values of the inner e-cell centroids are displayed.
	Outer	The values of the outer e-cell centroids are displayed.

Displaying an Unrolled View of a Battery Cell

You can display the unrolled view of a prismatic or cylindrical battery cell to see the values within the jelly roll.

To display an unrolled view of a prismatic or cylindrical battery cell:

1. Create a scalar scene or a vector scene.
2. Under the Displayers > Scalar 1 or Displayers > Vector 1 node, add the unrolled jelly roll as a derived part.
3. Select the field function that you want to display.

The example below shows the field functions that are available from the battery module.



The appropriate values are displayed:

4. If necessary, transform the unrolled view using the standard displayer properties.

   For example, you can scale the Y axis to make the displayed image easier to view.

Note the following:

- Field functions that are computed within STAR-CCM+ (for example, temperature) are mapped to the average spiral representation.
- Field functions that are computed from the battery model are averaged from the original e-cell centroid values.

Setting Up Terminal Connections

In a battery module, parallel strings can be connected using terminal connections, which lets you use imported geometry to define physical connections between the strings. These connections are used to compute the thermal solution only. Whether or not you set up physical connections between the parallel strings, an electrical connection always exists between the parallel strings in the circuit solver.

To set up terminal connections:

1. Import the terminal connection geometry into the simulation.

2. Right-click the Battery Module > Terminal Connections node and select Setup Terminal Connections.



3. In the *Setup Terminal Connections* panel, use the *Positive Terminal Connections* and the *Negative Terminal Connections* tables to select the parts for each connection.

   The Index value indicates which parallel string the terminal connection corresponds to, that is, 0 is the first parallel string, 1 is the second string, and so on.



4. When the terminal connection setup is complete, click Apply, and then click Close.

It is also possible to set up terminal connections using the terminal connections nodes.

Setting Properties of Terminal Connections

Terminal connections are stored under the Battery Module > Terminal Connections node. In the properties for the terminal node, you can review which terminals the connection is applied to, and also assign parts to the terminal connection.



Terminal Connection Properties

Post Parts	Selects the parts that define the post.
Strap Parts	Selects the parts that define the strap.
Copy Connection Parts	Specifies whether the post and strap parts are copied to other terminal connections. Activated: Parts are copied. Deactivated: Parts are not copied.

Reviewing Connected Terminals

You can review the terminals that a particular connection connects using properties of the Connector Strap and Connector Post nodes. These properties tell you which terminal of a particular battery cell the connection is applied to, for example, the positive terminal of battery module cell 0,0. To review the terminals that are assigned to the connection, select Connector Strap or

Connector Post under the Battery Module > Terminal Connections > Terminal Connection node. The terminal that the connection connects is displayed next to the Connected to property.



Setting Up Tab Connections

In a battery module, tab connections are used to connect battery cells within a string. Setting up tab connections allows you to use imported geometry to define physical connections between the cells. These connections are used to compute the thermal solution only. Whether or not you set up physical tab connections, the circuit solver always makes electrical connections between the cells in a series string.

Typically, the tab connector consists of three parts:

- First Post
- Second Post
- Strap

For stacked cells, the post parts must touch the tabs, as shown in the following diagram.



For cylindrical spiral cells, the posts must be in contact with the end plates at the location of the caps (internal posts). This configuration is demonstrated in the following screenshot.



You can prepare the tab connector geometry in an external CAD package, or use the 3D-CAD modeler available in STAR-CCM+. If you use the 3D-CAD modeler, you set up tab connections after generating battery parts.

To set up tab connections:

1. Import the connector geometry into STAR-CCM+.

2. Right-click the Battery Module > Tab Connections node and select Setup Tab Connections.



In the *Setup Tab Connections* panel, you can choose to use the same connector parts for each tab connection, or use different parts for each connection.

The positive tab on the first battery cell is always the open tab; the tab on the Nth cell has the negative tab open. This setup is consistent with the electrical connections that are set in the circuit solver.

You can check that the tab connectors are placed consistent with the electrical circuit by using the Tab Voltage field function.

Verify that the maximum voltage is at the open positive tab and the minimum voltage within each string is at the open negative tab. Additionally, the voltage that is displayed at tabs which are connected together is nearly identical (that is, the same color), assuming the resistance that is modeled for tab connections is small.

Using a Single Tab Connector

For battery modules whose cells are distributed in a regular pattern, you can import a single tab connector and use it multiple times to connect all the battery cells. In this method, STAR-CCM+ automatically copies and repositions the connector according to the arrangement of cells in the battery module.

To use this method, make sure that the *Connections Setup Method* is set to Copy Single Tab Connector.



In the First Post Parts, Strap Parts, and Second Post Parts boxes, select the corresponding tab connector parts. When you have selected the parts, click Apply and Close.

The tab geometry is multiplied and positioned between each tab pair when the tab connector parts are generated.

Selecting Parts for Each Connector

You can use different geometry parts to create each tab connection. This method is useful if the battery cells are not laid out in a regular pattern. To use this method, set the *Connections Setup Method* to Select Parts For Each Connector.



Connector parts are specified for each parallel string separately. To select the first set of connector parts, click the [] button for each option in the table, and select the corresponding post and strap parts. Once you have selected tab connectors parts for each tab connection in the table, increase the value in the Select Parallel String box and define the next set of tab connections. Continue until all tab connections have been defined.

When the settings are complete, click Apply and Close.

Setting Properties of Tab Connections

Tab connections are stored under the Battery Module > Tab Connections node. In the properties for each tab connection, you can select the parts that are assigned to the connection. Expanding each tab connection node also allows you to review which components each post and strap is connected to.



Tab Connection Properties

First Post Parts To Copy	Selects the parts that define the first post.
Second Post Parts To Copy	Selects the parts that define the second post.
Strap Parts To Copy	Selects the parts that define the strap.
Copy Connection Parts	Specifies whether the post and strap parts are copied to other tab connections. Activated — Parts are copied. Deactivated — Parts are not copied.

Specifying the Battery Module Resistance Values

The circuit solver allows you to set resistance values for the various connectors that are used within the battery module. By default, the resistance of each connector type is set to 0 Ω. You can set the resistances on the Circuit Model node as shown below.



In a battery, tab stems and tab connectors are welded together. These welds are not directly modeled as part of the geometry. They do not make perfect electrical contact, resulting in some electrical resistance between these parts.

You can model this electrical contact resistance using the Physics Values > Electrical Resistance node of a stem/connector interface. The electrical resistance value is entered as a scalar profile.



The Electrical Resistance option of a contact interface is made available only when the Electrodynamic Potential model is activated.

Circuit Model Properties

Straps Resistance	Sets the resistance of battery module straps.
Positive Terminal Resistance	Sets the resistance of positive terminals.
Negative Terminal Resistance	Sets the resistance of negative terminals.
Positive Posts Resistance	Sets the resistance of positive posts.
Negative Posts Resistance	Sets the resistance of negative posts.

Generating a Battery Module

Once the battery module has been fully defined, the next step is to generate the battery module. This process requires creating geometry parts and connector parts for the battery that is based on the layout of the battery cells and module definition, and generating regions from the necessary parts/contacts.

Generating Battery Parts

To generate the battery module parts:

- Right-click the Battery Modules > Battery Module node and select Generate Battery Parts.



The Geometry > Parts node is automatically populated with the battery parts that are designed in BDS. This action also activates the Generate Connector Parts and Generate Regions... options.



You can use individual parts to visualize results in a scalar scene when you run the simulation. This feature is useful if you wish to view results on a cell in the middle of the module.

- As with battery module parts, right-click the Battery Modules > Battery Module node and select Generate Connector Parts to generate connector parts, if tab connectors are used.

Before you generate regions, you imprint the parts that are in contact, as well as tab connectors and tabs. The imprinting is required to create all the necessary interfaces. Use the surface repair tool to find and imprint parts in contact, using an appropriate tolerance for your geometry. Furthermore, you can either edit or create new CAD data and merge/imprint your preferred module geometry to the idealized battery cell created by the battery simulation module.

Generating Battery Regions

After completing all the desired modifications to the battery module CAD at the Parts level, battery module regions and interfaces can be set up.

Before you generate the battery module regions, it is recommended that you use the surface repair tool to eliminate any possible surface errors. Before you continue, check that all the regions and interfaces have been properly created.

To generate the battery module regions:

- Right-click the Battery Modules > Battery Module node and select Generate Regions. .



The Generate Regions... action allows you to select the non-battery parts that you want to use to generate regions. You make the selection using the *Additional (non-Battery) Parts to Generate Regions* dialog.

If you did not create or import any non-battery parts, this dialog is empty.

One region is created for each type of battery component, for example all the battery stack parts are assigned to a single battery stack region.



When the battery module regions are generated, STAR-CCM+ recognizes each part of the battery. Where appropriate, interfaces between battery parts are created automatically.

Working with Battery Regions

Battery regions are regions that reference battery objects. They are created when the Generate Regions... action is selected in the battery module generation process.



By default, each battery region is defined as a battery energy source. These settings are defined automatically when the regions are generated, and do not need to be changed. Regions that are created for the tab connectors are not defined as battery energy sources.

Defining a Region as a Battery Energy Source

The thermal output of the battery during the load cycle is activated in a battery region using the battery energy source option. This setting is set up automatically on the region when the regions are generated and you do not need to change it. For a region specified as a battery energy source, the energy source type can be specified in the Physics Values for the region.

The Battery Energy Source Option node is available in the Physics Conditions for the region. It has properties and is subject to common menu items.



Battery Energy Source Option Properties

Battery Energy Source Term	Specifies whether the region is activated as a battery energy source.	
	Activated	Region is a battery energy source.
	Deactivated	Region is not a battery energy source.

Selecting a Battery Energy Source Type

The battery energy source type can be specified for a region that is activated as a battery energy source. STAR-CCM+ sets this option automatically when the battery module regions are generated and you do not need to change it.

The Battery Energy Source node is available in the Physics Values for the region. In the properties for the Battery Energy Source node, you can specify whether the region is a Core or Tab energy source. The Core option is used for the battery stack, while the Tab option is used for the tab roots and stems.



Battery Energy Source Properties

Battery Source Type	Specifies whether the region is a core or tab battery energy source.

Generating a Battery Mesh

The thin parts of a battery require specific meshing models and associated settings.

The STAR-CCM+ battery simulation module provides an option for automatically creating a mesh continuum that contains the models required to mesh a battery.

The mesh continuum is assigned to the battery module region by default.

You can change the pre-defined battery settings as necessary.

The mesh continuum contains:

- Thin Mesher and Surface Remesher models.

  The thin mesher recognizes thin parts of geometries and produces prismatic cells within the thin parts. For non-thin parts of the geometry, the thin mesher automatically changes to produce polyhedral or tetrahedral cells. The mesh remains conformal.

  You can control the relative or absolute thickness of the thin parts by activating the Customize Thickness Threshold property. See Thin Mesh Expert Properties.

- Reference values suitable for battery meshing.

  You can keep the default reference values in the battery mesh continuum or change them if required.

- A Volumetric Control.

  The volumetric control is created for the battery tabs and tab connectors with a reference value of 7.5% of the surface base size.

The typical steps for generating a battery mesh are:

1. Set up battery meshing
2. Create mesh continua for non-battery parts
3. Generate the volume mesh

Setting Up the Battery Mesh Continuum

To set up a battery mesh:

- Right-click the Batteries > Battery Modules > Battery Module node and select Setup Battery Meshing



A new Mesh 1 node appears under the Continua manager node. When a battery mesh continuum is created, it is assigned to a battery module region by default.

The battery mesh continuum contains the Surface Remesher and Thin Mesher models as well as a volumetric control for the tab parts.

To view the meshing models and volumetric control:

- Expand the Continua > Mesh 1 > Models node.
- Expand the Mesh 1 > Volumetric Controls node.

The nodes are listed in the tree as shown below:



Setting Up Connector Ohmic Heating

The Ohmic Heating model in tab connector regions can be set automatically for the battery module. The model accounts for the ohmic heating effect in the tab connectors when a current passes through them.

The ohmic heating for the connectors is defined using the Battery Module > Setup Connector Ohmic Heating option. This option becomes available once a battery mesh has been generated.



When ohmic heating in connectors is selected:

1. STAR-CCM+ creates a physics continuum, Battery Module Connectors: Ohmic Heating, for the Ohmic Heating model and assigns it to the tab connector regions.



2. STAR-CCM+ creates a Volumetric Control for the tabs and connectors with a default value of 4% of the surface base size. We recommend that you keep the size of the mesh on the tabs and connectors at roughly 4% of the tab stem height. The refinement of the mesh in the connectors helps with the convergence of the energy solver when the Ohmic Heating model is turned on.



3. STAR-CCM+ creates two reports for the current and interface surface area: BatteryCellCurrent and TabStem_Post_Intersect_Area. The battery cell current report is a volume average report that collects the current going through the cell. The battery module stack or the jellyroll region is automatically assigned to it. The tab stem/post intersect area report returns the area of the intersection between the stem and the post for the stacked cell, and end plate and post for the cylindrical cell.



4. STAR-CCM+ creates current density field functions from the reports: CurrentDensityToApplyToConnector_IN and CurrentDensityToApplyToConnector_OUT. The current flows into the post for the negative tab and flows out of the post for the positive tab. The field functions are automatically applied to the appropriate interfaces.



The two field functions compute the current density at the interfaces, with one for current flowing in and the other for current flowing out. These values are of equal magnitude but opposite sign. The definition for the positive current density is shown below. The convention is to have the positive current density represent the current flowing in.



Improving Convergence

As a result of non-converged initial electric potential, spurious heat is introduced into the Segregated Energy solver in the first iterations of an analysis. This behavior can lead to high temperatures in intermediate non-converged states.

To improve convergence of the Segregated Energy solver:

- Set a Solid Under-Relaxation Factor of `0.9999` for the Segregated Energy solver.



- Set a solid under-relaxation factor Linear Ramp



- Set the *Linear Ramp - Properties* as follows:

Start iteration	`6`
End iteration	`10`
Initial value	`1.0E-7`



Defining External Battery Casing

The external battery casing feature lets you automatically create a volume to fill the void between the battery parts and the surface of the part that you use for the casing. The casing part can be an external bounding box, or it can be a part that you specify, depending on the selected casing option.

External battery casing is specified using the Battery Module > External Casing Geometry node. To reveal this node in the simulation tree, you generate battery parts.



The external casing requires a separate mesh continuum to the one assigned to the battery module, as the Battery Mesher can only be used for battery objects. If you are using tab connectors with thin geometry, add extra volumetric controls for the connectors inside the external casing mesh continuum.

Selecting a Casing Option

The external casing option is selected using the Battery Module > External Casing Geometry > External Casing Specification node.

There are two methods for specifying the casing:

- Create a basic box around the module

- Use a geometry part that defines the casing



Selecting one of the options for the Battery External Casing Specification property allows you to define the size and shape of the external casing. Once you have fully defined the casing, you can create the external casing part.

Defining a Bounding Box

To create a bounding box around the battery module:

1. Select the Specify External Bounding Box casing option.

2. Select the Battery Module > Casing Values > External Casing Bounding Box node and set the position of the first and second corners of the box properties.



After defining the bounding box, the next step is to create the external casing part.

An example of a bounding box around a battery consisting of three series cells in a single parallel string is shown below.



The bounding box that is created around the battery module has a single external surface. If you wish to define one surface of the box as an inlet and another as an outlet, for example, you can split the surface using the standard techniques. If necessary, you can define an inlet and an outlet by

importing separate geometry parts. These parts can be interfaced with the bounding box before meshing, allowing you to create geometry such as that shown below.



Selecting a Casing Part

You can create external battery casing that is a specific shape by using a geometry part:

1. Create the geometry of the external casing; one method for creating the geometry is to use the embedded CAD modeler in STAR-CCM+, 3D-CAD.

   The casing part is a single volume. You do not need to subtract the battery module from the casing volume; STAR-CCM+ subtracts the module automatically when you create the casing part.

2. If the external casing part is created in a different simulation, import the casing geometry into the battery simulation.

   Make sure that the *Import Mode* is set to Create new Part.

3. Select the Specify External Casing Using Part casing option.

4. Select the Battery Module > Casing Values > External Casing Part node and set the External Casing Part property to the casing part.

After selecting the part that you want to use as the battery module casing, the next step is to create the external casing part.

An example of a custom external casing around a battery consisting of three series cells in a single parallel string is shown below.



Creating External Battery Casing

To create an external casing:

- Right-click the Battery Module node and select Create External Casing Part.



During the process of creating external casing, STAR-CCM+ creates a bounding box, or takes the selected casing part, and subtracts the battery module parts from it. The resulting part is named Battery Module: External Casing.



STAR-CCM+also generates a region for the external casing during the generate regions step. An example of the regions and interfaces that are created for a single cell battery module with an external casing is shown in the following screenshot.



The final step in setting up the external casing is to define the mesh settings for the external casing region in a separate mesh continuum. It is only possible to use the battery mesh continuum for battery regions, such as the regions for the stack, tab stems, and tab roots. When setting up the mesh continuum for the external casing region, a good technique is to use volumetric controls to refine the mesh around the tab parts, including root, stems, posts, and straps. This technique promotes a good quality mesh around the thin tab and connector parts.

Defining an Electrical Load

You can define the electrical load on the battery in the following ways:

- program file (.prg)

  Specifies a load for a particular duration. For example, 10 amps for 10 seconds, then 8 amps for the next 20 seconds, and then 6 amps for the next 30 seconds.

  For more information, see Using a Program File.

- table

  Each table row specifies a load at a particular time. The interval between specified time values can be flexible: you do not need to specify values at regular time intervals. However, the load values are linearly interpolated between the specified time values to provide data at the time step of the thermal solver.

  Because of this interpolation, using a table to define a load can provide more flexibility than using a program file.

  For more information, see Using a Table.

- scalar function

  Specifies a load as a scalar function. The scalar function can be a constant value or an expression which is evaluated over time. The expression can use field functions that you define.

  For more information, see Using a Scalar Function.

Once you have created the program file in BDS or a table, you can import the load into STAR-CCM+ and connect it to the battery.

Using a Program File

If you are using a program file:

1. Right-click the Batteries > Program Files node and select Load...



2. In the *Open* dialog, select the program file and click Open.



A program file is added under the Program Files manager node. The Program Files node has properties and a pop-up menu. You can edit the program file using the Program Data property.



Currently not all the functionality for end and limit conditions is available in STAR-CCM+.

- STAR-CCM+ fully supports Time end conditions. Voltage end conditions are partially supported.

  See Setting End Conditions.

- STAR-CCM+ uses Voltage limit conditions to specify a warning message.

  See Setting Limit Conditions.

Program File Pop-Up Menu

The following options are available when you right-click on a program file node.

Create Circuit Element	Creates a circuit element, which is used to connect the electrical load to the battery module.
Reload...	Reloads the program file data into the simulation.
Export...	Exports program file data to a .prg file.

Program File Properties

File name	Displays the path to the .prg file.
Program Data	Displays the data that is contained in the .prg file.

Reloading Program Data

If you have changed the program data outside of the simulation, for example, by modifying the .prg file in BDS, you must reload the data into STAR-CCM+. To reload the data, the name and path of the .prg file must be the same.

To reload the program data:

- Right-click the Program Files > <filename>.prg node and select Reload...



STAR-CCM+ rereads the contents of the .prg file, and loads the data into the simulation. The progress of the reload operation is displayed in the *Output* window.

Exporting a Program File

You can export data from a simulation to a .prg file. Exporting the file allows you to recreate the .prg file from the program data that is stored within the simulation.

To export program file data to a .prg file:

- Right-click the program file node and select Export...



The *Save* dialog allows you to save the data in a .prg file.

Using a Table

As an alternative to using a program file, you can use a table to define the electrical load. The table must have a time column, which is interpreted as the independent variable, and a load column, specifying current, power, or voltage. You can define the units in the column heading using parentheses, for example Time (min) or Power (kW). STAR-CCM+ converts these units to the default SI units: seconds for time, amps for current, watts for power, and volts for voltage. If no units are specified, the values are imported in their default units. The load values are linearly interpolated to provide data at the time steps that are required by the thermal solver.

Below is an example of a table for the electrical load. Although it shows three different load columns, only one such column is needed.



Tabular Data

	Time	Current	Voltage	Power
0	1	5.88	4.15	24.4
1	2	69.74	4	278.95
2	3	93.33	3.94	368.14
3	4	116.21	3.89	451.94
4	5	138.77	3.83	531.95
5	6	161.23	3.78	609.1
6	7	183.42	3.72	682.71
7	8	205.43	3.67	753.26
8	9	227.21	3.61	820.47
9	10	248.83	3.56	884.72
10	11	270.13	3.5	945.46

Export... Close

To specify the electric load using a table:

1. Load a table using the Tools > Tables node.
2. Right-click the Circuit > Circuit Solver > Circuit Elements node and select New Circuit Element > Table Circuit Element.



STAR-CCM+ adds a Table Circuit Element node to the Circuit Elements manager node.

3. In the *Properties* window of the Table Circuit Element node, select a load table and use the column headings of the table to set the load for the circuit solver.

Table Circuit Element Properties

Table	Sets the table that is used for the load.	
Time Data	Specifies the time column of the table.	
Load data	Specifies the load column of the table.	
Load Type	Specifies the load type that is read in.	
	Current	The load is current load.
	Power	The load is power load.
	Voltage	The load is voltage load.
	Resistance	The load is a resistor load.
	Inductance	The load is an inductor load.
	Capacitance	The load is a capacitor load.

Using a Scalar Function

As an alternative to using a program file, you can use a scalar function define the electrical load. The scalar function can be a constant value or an expression which is evaluated over time. The expression can use field functions that you define.

To specify the electrical load using a scalar function:

1. If necessary, set up the appropriate field functions using the Tools > Field Functions node.

2. Right-click the Circuit > Circuit Solver > Circuit Elements node and select New Circuit Element > Scalar Circuit Element.



STAR-CCM+ adds a Scalar Circuit Element node to the Circuit Elements manager node.

3. In the *Properties* window of the Scalar Circuit Element node, specify the scalar function and the load type.

Scalar Circuit Element Properties

Load	Specifies the scalar function that is used for the load.
Load Type	Specifies the load type of the scalar function.
	Current — The load is current load.
	Power — The load is power load.
	Voltage — The load is voltage load.
	Resistance — The load is a resistor load.
	Inductance — The load is an inductor load.
	Capacitance — The load is a capacitor load.

An example of the Circuit > Circuit Solver node with circuit elements of each load type is shown below:



Creating a Circuit

In order to simulate the response of the battery to an applied electrical load, it is necessary to connect the battery and the load in a circuit. The circuit is created by first creating circuit elements, then creating connections between these elements.

A circuit is formed from elements and connections, as shown in the diagram below.



The internal structure of the battery module is illustrated below.



When the simulation is run, the Circuit Solver computes the electrical solution for the circuit.

Creating Parallel Connections

To create a circuit that consists of several circuit elements connected in parallel, add all the terminals for a parallel connection to a single circuit connection.



In the above example, two connections are required. The first contains the positive terminal from each battery module and the negative terminal from the load. The second contains the negative terminal from each battery module, and the positive terminal from the load.

To create the first connection from the example above:

1. Multi-select the positive terminal from each battery module and the negative terminal from the load.

2. Right-click any selected node and then click Create Connection.



The new connection is added under the Circuit Connections node.



3. Repeat the previous two steps to create the second connection with the appropriate terminals, that is, the negative terminal from each battery module, and the positive terminal from the load.

What Is the Circuit Solver?

The circuit solver computes the electrical solution for the circuit. The circuit solver operates in conjunction with the electrical solver, where its role is to calculate the global battery voltage, and current in the case of a power load, based on the voltage of each battery cell.

The main computational cost of the circuit solver is in repeated calls to the electrical solver to obtain the battery cell voltage as the circuit solver iterates to a converged solution at each time step. The number of ecells within a battery cell directly impacts the time that is required to reach each battery cell solution. Furthermore, the iterative solution of the electrical circuit for each battery cell can be fast when the current load is changing slowly. On the other hand, if the specified load is a power load for which both the voltage and current must be determined, or a rapidly varying current load, the circuit solver takes more iterative steps to reach a converged solution. In the case of a power load, the specified load must not exceed the capacity of the battery circuit. In this case, no electrical solution exists and the electrical solver can fail to converge.

At each electrical time step, the circuit solver does the following to determine the electrical solution:

- updates the load
- iteratively finds the final current and voltage for each battery cell that satisfies the specified circuit load
- computes heat generation for circuit connectors

To calculate voltage and current values, the circuit solver calls the battery cell model for each battery cell in the circuit. You cannot specify the battery cell model time step. The initial time step that is used in the battery cell model is the same as the time step that is used in the circuit solver. However, when necessary, STAR-CCM+ automatically adjusts the battery cell model time step to stabilize the circuit solver and obtain a solution.

To determine a solution, the battery cell model calls the unit cell battery model. You can specify the maximum time step that the unit cell model can use.

The circuit load can be defined using either a separate program file that is created in BDS, or using STAR-CCM+ table or scalar load options. The program file contains one or more lines of input that specify a load type (voltage, current, resistance, or power), the numerical value of the load, and the time duration for that load. The table load option allows you to specify a load type (voltage, current, resistance, inductance, capacitance, or power), and then set the values using a table from which interpolated values can be

easily obtained. The scalar load option allows you to specify a load type and then set constant value loads, or use user-defined field functions to set time-dependent loads.

The local temperatures that are input values to each ecell are mapped from the thermal solution. To start the simulation, the local temperatures are taken from the default initial state. The electrical solver uses BDS solvers that have been integrated into STAR-CCM+ to incorporate the user-determined unit cell model. The electrical solver receives a current value from the circuit solver and computes the voltage and the internal heat generation of each ecell, and computes the output voltage for each battery cell for the circuit solver.

After each circuit solver time-step, the heat generation from each ecell is mapped to the thermal grid and used as input for the thermal solver.

The circuit solver uses an adaptive time stepping scheme such that there is an electrical solution for each thermal solution time step. The circuit solver and the thermal solver normally use the same time-step size, but the time-step sizes can be different.

If the electrical load is specified using a program file that is created in BDS, the electrical solver can take multiple substeps for each thermal time step and is able to evaluate the load at intermediate thermal time values. If the electrical load is specified using a table, the load values are linearly interpolated so that only one load value is evaluated for each thermal time step. It is possible to take substeps within each thermal time-step using a table load but the load value changes at the thermal time step only. Similarly, using a scalar load option, the electrical solver evaluates the load at each thermal time step.

Besides current and voltage, the ecells within each battery cell maintain a state of charge parameter. The normal operating range for this parameter is from 0 through 1: a state of 0 percent charge (depletion of the cell) and a state of 100 percent charge (fully charged cell). When one or more ecells reach either of these limit values, the simulation can continue but the results can become unphysical. For more information, see Working with the Battery Cell State of Charge.

The Circuit Solver manager node stores information relating to circuit elements and connections. It has properties and is subject to common menu items.



Properties

Verbosity	Sets the level of output information that is displayed while the solver is running. By default this option is set to None. Setting the level to Low displays circuit voltage, current, and power after each time step as well as CPU time and memory use for the circuit solver.
Electrical Substeps	Sets the number of Electrical Solver substeps per Thermal Solver time step. The default setting is 1 electrical substep per thermal solver time step. This option can be useful when there is detailed transient load input for the electrical load but the thermal solution does not require small time-steps. However, it can be expensive for DIST model runs to take multiple time-steps for each thermal time step.
Relative Tolerance	Specifies a convergence parameter that the circuit solver uses internally. You are advised not to change this parameter.

Specifying the Maximum Time Step for the Unit Cell Model

When necessary, STAR-CCM+ automatically adjusts the battery cell model time step to stabilize the circuit solver and obtain a solution.

For example, if a time step of 100 seconds is specified in the Implicit Unsteady solver, the battery cell model normally delivers results in 100 seconds. However, if STAR-CCM+ reduces the battery cell model time step (for example, to 25 seconds), the battery model calls the unit cell model to deliver the results at 25-second increments.

You can specify what the unit cell model does in this scenario: attempt the full 25-second time step, or use a specified maximum time step. For example, if you specify a maximum time step of 1 second (the default value), the unit cell model makes 25 substeps, each of 1 second, to reach the battery cell model time step (25 seconds). This substepping usually enhances stabilization but requires a longer time to compute. However, if you specify a maximum time step greater than 25 seconds, the unit cell model makes a single time step of 25 seconds.

The larger maximum time step setting allows the battery cell model to achieve the Implicit Unsteady time step quickly. You can use this setting in cases where CFD calculations are not needed, such as conduction heat transfer problems or a single cell analysis with no cooling system. In these cases, change is slow and large time steps can be taken without affecting the stability of the simulation.

To specify the maximum unit cell time step:

1. Under the Cell Properties node for the cell, expand the Models > IET > Distributed NP > Numerical Settings > Overall Battery Settings node.

2. Select the Max Unit Cell Time Step node and, in the *Properties* window, set the appropriate maximum time step value in seconds.



Working with the Battery Cell State of Charge

If a battery cell state of charge reaches 0% or 100%, the simulation results can become unphysical, but the simulation can continue to run.

If the battery cell uses a physical model (such as the DISTNP model), the simulation stops automatically when the state of charge reaches 0% or 100%.

However, if the battery cell uses an empirical model (such as the NTG model or the RCR model), the state of charge of the cell becomes fixed at 0% or 100%. The simulation continues to run, but uses the constant value that the model computes at the fixed state of charge. Obviously, a real battery cell does not do this.

If you want your simulation to stop when the state of charge of any battery cell reaches 0% or 100%, you can set the appropriate stopping criteria.

To set the stopping criteria:

1. Create a new minimum report and, in the *Properties* window, do the following:
   - Set Scalar Field Function to Battery SOC.
   - Set Parts to the jellyroll region.

     This setting is important because the state of charge is defined only on the jellyroll region. If you choose any other regions, the monitor cannot find the minimum state of charge value.



2. Create a monitor from the minimum state of charge report.



If you want to visually monitor the state of charge values, create a plot from the minimum state of charge report as well.

3. Create a stopping criteria from the minimum state of charge monitor and, in the *Properties* window, do the following:

   - Set Criterion Option to Minimum.

- To stop the simulation when the minimum state of charge is reached, activate Stop Outer Iterations.



4. Select the Minimum Limit node and, in the *Properties* window, set the Minimum Value as a percentage of the cell capacity.



What Is a Circuit Element?

A circuit element is a component of a circuit. Each circuit element has a positive and a negative terminal. When two or more circuit elements are defined, you can create a circuit by specifying connections between the terminals.

You can derive a circuit element from a battery module or a program file, or create a circuit element directly from the Circuit Elements node.

Circuit elements are stored under the Circuit > Circuit Solver > Circuit Elements manager node.



Each circuit element node has properties and is subject to common menu items.

Properties

Battery Assembly	Available for Battery circuit elements. These elements can be created directly or derived from battery modules. Allows you to specify the battery module that the circuit element is associated with.
Forward Voltage	Available for Diode circuit elements. These elements can only be created directly. Allows you to specify the forward voltage of the diode. If the voltage across the diode exceeds the specified value, the diode behaves like a linear resistor plus a series voltage source. The value of the On Resistance property specifies the resistance of the linear resistor. If the voltage across the diode is less than the specified value, the diode behaves like a linear resistor with low conductance specified by the value of the Off Conductance property.
On Resistance	Available for Diode circuit elements. Allows you to specify the resistance of the diode.
Off Conductance	Available for Diode circuit elements. Allows you to specify the conductance of the diode.
Program Data	Available for Program circuit elements. These elements can be created directly or derived from program files. Allows you to specify the program file that the circuit element is associated with.
Load	Available for Scalar circuit elements. These elements can only be created directly. Allows you to specify the load with a scalar function. The scalar function can be a constant value or an expression which is evaluated over time. The expression can use field functions that you define.
Load Type	Available for Scalar circuit elements and Table circuit elements. These elements can only be created directly. Specifies the load type that is read in. Current — The load is current load.

Power	The load is power load.
Voltage	The load is voltage load.
Resistance	The load is a resistor load.
Inductance	The load is an inductor load.
Capacitance	The load is a capacitor load.

Table	Available for Table circuit elements. These elements can only be created directly. Allows you to specify the load with a table. The complete list of table circuit element properties is listed under using tables.

What Is a Circuit Element Terminal?

Each circuit element has a positive and a negative terminal. The terminal nodes are stored under the circuit element node, and are used to create connections between different circuit elements. Terminal nodes are denoted with (+) for positive and (-) for negative, and have properties.



By default, STAR-CCM+ assumes that the positive terminal of the program file or table circuit element is grounded.

Properties

Connection	Displays the name of the connection to which the terminal is assigned.
Grounded	Specifies whether the terminal is grounded, or "earthed".
	Activated — Terminal is grounded.
	Deactivated — Terminal is not grounded.

Adding Circuit Elements

Circuit elements can be derived from battery modules and program files, or created directly from the Circuit Elements node.

To derive a circuit element from a battery module or program file:

- Right-click the Battery Modules > Battery Module node or the Program File > **`<program_file>`**.prg node and select Create Circuit Element.



To create a circuit element directly:

- Right-click the Circuit > Circuit Solver > Circuit Elements node, select New Circuit Element > and then select the appropriate option:
  - Battery Circuit Element
  - Diode Circuit Element
  - Program Circuit Element

- Scalar Circuit Element
- Table Circuit Element



The new circuit element is added under the Circuit > Circuit Solver > Circuit Elements node.



If you create a circuit element directly from the battery module or program file, the corresponding Battery Assembly or Program Data property is set automatically.

If you create a circuit element directly from the Circuit Element node, you must specify the necessary properties in the *Properties* window.

Connecting Circuit Elements

Circuit elements are connected by creating *circuit connections* within the circuit solver. The process involves connecting the terminals of different circuit elements. It is possible to connect circuit elements in parallel and in series. In a series connection, the positive terminal of one circuit element is connected with the negative terminal of another. In a parallel connection, the circuit connection contains three or more terminals, depending on the layout of the circuit.

There are two ways to create connections between circuit elements:

- Using the terminal nodes (Recommended)
- Using the circuit connections manager

Each circuit connection is stored under the Circuit > Circuit Solver > Circuit Connections manager node. Each Connection node has properties and is subject to common menu items.



Properties

Terminals	Displays the names of the terminals that are connected in this connection.

Using the Terminal Nodes to Create Connections

Before you begin, ensure that you have created circuit elements that you wish to connect in the circuit.

1. Expand the circuit element nodes under the Circuit Solver > Circuit Elements manager node.
2. Select the terminal nodes that you wish to connect by using `<Ctrl><Click>`, for example, (+) Battery Module and (-) Discharge Program in the following screenshot.
3. Right-click one of the selected terminals and select Create Connection.



The connection is added under the Circuit Solver > Circuit Connections node. You can change the terminals that are connected by this connection by editing the properties of the Connection node.



Using the Circuit Connections Manager to Create Connections

In this method, an empty circuit connection is created first and the circuit element terminals to include in the connection are defined later.

To create a connection:

1. Right-click the Circuit Solver > Circuit Connections node and select New.



The connection is added under the Circuit Solver > Circuit Connections node. Terminals must be added to the connection in order to activate it.

2. Select the Circuit Connections > Connection node.



3. In the *Properties* window, click the [...] button to the right of the Terminals property.

4. In the *Terminals* dialog, select the terminals to be connected.



The *Terminals* dialog only shows terminals that have not been used in any other connections.

5. Click OK to confirm the selection.

Reporting on Circuit Elements

You can report on specific circuit elements to monitor the voltage, current, or current derivative across the element.

The following reports are available:

- Circuit Element Current
- Circuit Element Current Derivative
- Circuit Element Voltage

To set up a circuit element report:

1. Right-click the Reports node, click New Report, and then click the appropriate report type.



The corresponding node is added under the Reports node.

2. Select the new report node and set the appropriate properties:
   - Circuit Element Current Report Properties
   - Circuit Element Current Derivative Report Properties
   - Circuit Element Voltage Report Properties

   The example below is for the Circuit Element Voltage report.



Circuit Element Current Properties

Circuit Element	Specifies the circuit element to report on. The drop-down list contains all of the circuit elements that are defined under the Circuit Solver > Circuit Elements node.
Units	Specifies the units to use for the current values. The default is A (Amps).

Circuit Element Current Derivative Properties

Circuit Element	Specifies the circuit element to report on. The drop-down list contains all of the circuit elements that are defined under the Circuit Solver > Circuit Elements node.
Units	Specifies the units to use for the current derivative values. The default is A/s (Amps per second).

Circuit Element Voltage Properties

Circuit Element	Specifies the circuit element to report on. The drop-down list contains all of the circuit elements that are defined under the Circuit Solver > Circuit Elements node.
Units	Specifies the units to use for the voltage values. The default is V (Volts).
Terminals	Specifies the terminal for which the voltage is measured.
Voltage Drop	Specifies whether the voltage is measured with respect to the other terminal or with respect to ground. This option is activated by default. Activated — The voltage at the selected terminal is measured with respect to the other terminal. Deactivated — The voltage at the selected terminal is measured with respect to ground.

Exporting Electrochemistry Data

You can export detailed electrochemistry data for a battery module at each specified time-step. Electrochemistry data is exported to `.pr2` files, which are text files that can be read into the BDS Profile Viewer. One `.pr2` file is generated for each ecell in each battery cell. For example, for a single cell battery module where the battery cell contains 3x3x1 ecells, 9 `.pr2` files are generated.

This feature is only available for battery modules containing battery cells that use a *Distributed* battery model. The Distributed model is selected in Battery Design Studio® (BDS) when you set up the .tbm file, for example Distributed 1D, Distributed 2D. You can check which battery model a battery cell uses by checking the Unit Cell Model property of the Battery Cell node. It is set to `Distributed` if a Distributed battery model was selected in BDS.

Electrochemistry data export is set up using the properties of the ElectroChemistry Export node of the battery module of interest. This node is added to the tree after setting up and generating the battery module, and once the battery mesh continuum has been assigned.



ElectroChemistry Export Properties

Export Directory	Sets the directory for the exported data files.
Export frequency	Sets the time interval at which to export data. The default value of 1 results in data being exported at each time-step. Setting the frequency to 4, for example, would export data at every fourth time-step.
Export start time	Sets a specific time at which data export begins (s).
Exported Battery Module Cells	Selects cells in the battery module for which to export data.

Evaluating Electrode Structure Tortuosity

This section outlines a technique that you can use to evaluate an electrode structure tortuosity with STAR-CCM+. This technique can be helpful when you are creating your own geometry from idealized particles.

You can use STAR-CCM+ to calculate the tortuosity of porous media from the data that is obtained with laboratory X-ray instruments. The tortuosity is evaluated from a mass transport perspective.

The analog between Fick's law for mass transport driven by a concentration gradient and Fourier's law for heat transfer across a temperature gradient is well documented (for example, [422]). The mathematical and physical similarities of the heat and mass transport calculations allow you to effectively utilize a heat-transfer simulation in STAR-CCM+ to predict the tortuosity of a pore network. This technique provides the same tortuosity value but has boundary conditions that are easier to implement.

STAR-CCM+ can use tomography data sets that are formatted as stereolithography (*.stl) files, which describe the interface between the solid and pore phases. STAR-CCM+ can then generate adaptive polyhedral volume meshes from the .stl files.

Using this microstructurally faithful modeling framework, you can run simulations by applying the boundary conditions that are suitable for the heat transfer analog, using nominal values for the thermal conductivity and temperature difference.

To minimize the effect of small scale heterogeneities, run the simulations on the largest volume that is fully bounded by the sample data. To measure the tortuosity in all three directions, you should run three separate simulations. For each simulation, set the two parallel faces that are normal to the direction being investigated to fixed temperatures, while specifying all of the other faces as adiabatic.

The heat transfer measured can then be put into a modified form of the heat transfer equation, and the tortuosity determined as follows:

$$\tau = \phi \frac{Q}{Q_p} \tag{2463}$$

where:

- Q = heat transferred through a solid block (W)
- Q_p = heat transferred through a porous media (W)
- τ = Tortuosity
- ϕ = Porosity

Battery Modeling Formulation

This section provides the following types of battery modeling formulation:

- NTG (Newman, Tiedeman, Gu) model
- RCR model
- Physics-based models
  - DistNP model
  - DistSegregated model

A bibliography is also included.

NTG (Newman, Tiedeman, Gu) Model

This model was proposed in a number of papers ([425], [427], [429]) and uses a polynomial representation to capture the voltage performance of a battery cell. You supply the model with the coefficients a_0 to a_6, which are then used during run time to predict how the battery cells respond to the imposed load. In its original form, this model had no temperature dependence but this formulation has been extended to include an activation energy, E_a, which can account for temperature effects. The model is most suited for constant current operation of battery cells, similar to those found in electric vehicles whereby the discharge does not contain high peaks of either charge or discharge. To create this model, you need relevant discharge curves from bench tests of the battery cell to which the coefficients are fitted. These coefficients should then be normalized by the active area of the battery cell:

$$V_{cell}(DoD, J) = \frac{J}{Y} + U + b(T - T_{ref}) \tag{2464}$$

$$DoD = \frac{\int Idt}{Capacity}, J = \frac{I}{A} \tag{2465}$$

$$U = a_0 + a_1 DoD + a_2 DoD^2 + a_3 DoD^3 \tag{2466}$$

$$Y = (a_4 + a_5 DoD + a_6 DoD^2)_e^{\frac{E_a}{R}\left(\frac{1}{T_{ref}} - \frac{1}{T}\right)} \tag{2467}$$

$$\frac{dU}{dT} = a_7 + a_8 DoD + a_9 DoD^2 + a_{10} DoD^3 \tag{2468}$$

$$Q = A \cdot \left(U - V_{cell} + T\frac{dU}{dT}\right) \tag{2469}$$

where the terms are defined as follows:

V_{cell}	working cell voltage $[V]$
DoD	depth of discharge (a fraction)
Y	conductance $[S/m^2]$
U	equilibrium voltage $[V]$
J	current density $[A/m^2]$
A	electrode area $[m^2]$
I	current $[A]$
T	temperature $[k]$
T_{ref}	reference temperature $[k]$
Q	electrode polarization heat generation $[W]$

RCR Model

This model uses a simple equivalent circuit consisting of a resistance R_o in series with a resistor R_p and capacitor C in parallel, along with a state-of-charge dependent voltage source V_o to represent the behavior of the battery cell. Due to the presence of the capacitor within the model, it is more suited to battery cells which contain high discharge or charge peaks, allowing for the temporal effects within the battery cell. This model can be created using the regression process within Battery Design Studio, depending on pulse power curves presented from testing. The five input parameters to the model can be represented by fifth-order polynomials. The coefficients should then be normalized by the active area of the battery cell [432].



$$V_0 = \sum a_i \cdot SOC^i \tag{2470}$$

$$R_o = e^{-\frac{E_o}{RT}} \sum R_{o,i} \cdot SOC^i, E_o = \sum E_{o,i} \cdot SOC^i \tag{2471}$$

$$R_p = e^{-\frac{E_p}{RT}} \sum R_{p,i} \cdot SOC^i, E_p = \sum E_{p,i} \cdot SOC^i \tag{2472}$$

$$\tau = e^{-\frac{E_\tau}{RT}} \sum \tau_i \cdot SOC^i, E_\tau = \sum E_{\tau,i} \cdot SOC^i \tag{2473}$$

$$V_L = V_o - I \cdot R_o - \left(\frac{I}{C}\right)\Delta t + (V_L - V_o + I \cdot R_o)_{t-\Delta t} \exp\left[\frac{-\Delta t}{\tau}\right] \tag{2474}$$

$$\frac{dV_o}{dT} = \sum \frac{dV_o}{dT_i} \cdot SOC^i \tag{2475}$$

$$Q = I \cdot \left(V_o - V_L + T\frac{dV_o}{dT}\right) \tag{2476}$$

Physics-based Models

The physics-based battery models in STAR-CCM+ are derived from the DualFoil model [423], [424]. This model is a detailed electrochemistry model that is used to predict battery cell performance. The physical formulation is common to both of the physical models that are implemented in STAR-CCM+: the DistNP model and the DistSegregated model.

The model is essentially one-dimensional in nature in that it resolves the liquid-phase as well as solid-phase electric charge transport (electric potential), and the liquid-phase salt transport from the negative current collector to the opposing positive current collector. An additional, one-dimensional transport equation is solved for the solid-phase lithium concentration, which models the lithiation of a respective active material particle. The equation is also solved for the flux of lithium cations across the Solid-Electrolyte Interface (SEI). The temperature is computed using an energy balance around the whole cell [421].

The physics-based model formulation is presented in the following sections:

- Potential Model
- Concentration Model
- Solid Electrolyte Interface Current Density

Potential Model

Solid Active Material

Integral equation

$$\oint_A -\sigma \nabla \phi_s \cdot d\boldsymbol{A} = -\int_V a J_n dV - \int_V a C \frac{\partial(\phi_s - \phi_l)}{\partial t} dV \tag{2477}$$

where the terms are defined as follows:

σ	electrical conductivity [S/m]
ϕ	electric potential [V]
a	specific interfacial area [m^2/m^3]
C	double-layer capacitance [Farad/m^2]
J_n	SEI normal component of electric current density [A/m^2]

Effects introduced by surface chemistry are modeled by the terms on the right-hand side of Eqn. (2477), where the second term corresponds to the double-layer capacitance effect.

(Electric) Current definition

In the solid electrodes and current collectors, the electric current density is defined using Ohm's law as:

$$\mathbf{J}_s = -\sigma \nabla \phi_s \tag{2478}$$

Electrolyte

Integral equation

$$\oint_A -\kappa \nabla \phi_l \cdot d\boldsymbol{A} = \oint_A -\kappa_d \nabla(\ln c_l) \cdot d\boldsymbol{A} + \int_V a J_n dV + \int_V a C \frac{\partial(\phi_s - \phi_l)}{\partial t} dV \tag{2479}$$

The effective electrical (ionic) conductivity κ in the liquid phase accounts for an eventual non-unity porosity ε and tortuosity τ within the separator:

$$\kappa = \kappa_0 \frac{\varepsilon}{\tau} \tag{2480}$$

$$\kappa_d = \frac{2RT\kappa}{F}(1 - t_+^0)\left[1 + \frac{d(\ln f_\pm)}{d(\ln c)}\right] \tag{2481}$$

(Ionic) Current definition

$$\mathbf{J}_l = -\kappa\nabla\phi_l + \kappa_d\nabla(\ln c_l) \tag{2482}$$

A schematic illustration describing the electric potential formulation, including the prescribed boundary conditions, is shown below:



Concentration Model

Solid Active Material

Integral equation

A simple species diffusion transport equation is solved for lithium (lithium/salt concentration) transport within the electrode phase:

$$\frac{\partial}{\partial t}\int_V c_s dV = \oint_A D_s \nabla c_s \cdot d\mathbf{A} \tag{2483}$$

where c_s is the molar concentration [kmol/m^3] and D_s is the effective lithium diffusivity [m^2/s] (if the Concentration Dependent Diffusion Coefficient setting is selected in Battery Design Studio):

$$D_s = -D_{s,0}\frac{Fy}{RT}\frac{dU_{eq}}{dy} \tag{2484}$$

The terms are defined in the Electrochemistry Nomenclature.

Electrolyte

Integral equation

An additional transport term for lithium cation migration is considered in the (liquid) electrolyte phase, so that:

$$\frac{\partial}{\partial t}\int_V \varepsilon c_l dV = \oint_A D_l \nabla c_l \cdot d\boldsymbol{A} - \int_V \frac{\boldsymbol{J}_l \cdot \nabla t_+^0}{F} dV + \int_V \frac{a\boldsymbol{J}_n}{F}(1 - t_+^0) dV \tag{2485}$$

where t_+^0 is the transference number of cations and where the effective lithium/salt diffusivity coefficient:

$$D_l = D_{l,0}\frac{\varepsilon}{\tau}\left[1 - \frac{d(\ln c_0)}{d(\ln c)}\right] \tag{2486}$$

accounts for non-unity values of porosity ε within the separator.

The model that is assumed for lithium (salt) transport is illustrated below:



In this figure, a* illustrates an internal relaxation current and b* illustrates salt transport during discharge.

Solid Electrolyte Interface Current Density

The Butler-Volmer relation is used to compute the solid electrolyte interface (SEI) electric current density, which is written as:

$$J_n = J_0\left(e^{\frac{\alpha_a F\eta}{RT}} - e^{\frac{-\alpha_c F\eta}{RT}}\right) \tag{2487}$$

where:

$$J_0 = Fk\left(\frac{c_s}{c_{s,\max}}\right)^{\alpha_1}\left(1 - \frac{c_s}{c_{s,\max}}\right)^{\alpha_2}\left(\frac{c_l}{c_{l,\mathrm{ref}}}\right)^{\alpha_3} \tag{2488}$$

The SEI Surface Overpotential η is computed as:

$$\eta = \phi_s - \phi_l - U_{\mathrm{eq}} - J_{n,s}R_{\mathrm{SEI}} \tag{2489}$$

The rate coefficient k depends on temperature according to:

$$k = k_0 e^{\frac{-E_a}{RT}} \tag{2490}$$

Similarly, the electrical resistance of the SEI is dependent on temperature:

$$R_{\mathrm{SEI}} = R_0 e^{\frac{-E_{\mathrm{SEI}}}{RT}} \tag{2491}$$

The Butler-Volmer relation is illustrated below:



J_a represents the anodic current contribution, while J_c represents the cathodic current contribution to J_n.

DistNP Model

This model is an extension of the DualFoil model. The STAR-CCM+ implementation allows for several active materials within the same model and also accounts for complex expressions to model diffusion coefficients.

This model uses a pentadiagonal version of Newman's band algorithm [431]. The DistNP model solves numerically for the solid-phase diffusion equation, instead of using Duhamel's superposition integral [423]. This allows the usage of concentration-dependent diffusion coefficients. Active material particles are modeled either as spherical, cylindrical, or planar.

In the present implementation, output quantities within the battery cell, such as lithium concentration, are not available from the STAR-CCM+ user interface. You can apply such a model but only the post-processing quantities that are generic to all three models are available.

To check the 1-D model data, you can export it to `.pr2` files, which are text files that can be read into the BDS Profile Viewer. For more information, see Exporting Electrochemistry Data.

The DistNP model requires many input parameters to describe the battery cell but these parameters are used in a model based on first principles and hence produce a more robust prediction. The model created should be checked in Battery Design Studio before proceeding in STAR-CCM+.

DistSegregated Model

This model is an extension of the DualFoil model. The STAR-CCM+ implementation employs an iterative, segregated solution algorithm, but allows only a single active material per electrode. Active material particles are modeled as spheres.

The solver for electric potential is numerically coupled in active electrodes and the electrolyte. The DistSegregated model has a fully segregated solution procedure for calculating the concentration in active electrodes and the electrolyte. The model uses "delta form".

Activating the DistSegregated Model

The DistSegregated model cannot be created in Battery Design Studio.

Note: The DistSegregated model can only be activated for Distributed models that have only one positive particle and one negative particle per electrode.

To activate the DistSegregated model in STAR-CCM+:

1. Open a suitable DistNP model .tbm file for editing.
2. Make the following changes to the file:

- In the `<MODELMAP>` section:

```
S-2009-A1-OS-thin-fast.tbm | Segregated-2009-A1-OS-thin-fast.tbm
5836 <MODELMAP>
5837     Electrolyte =   General Electrolyte
5838     IET =   Distributed 1D
5839 </MODELMAP>
```

  replace the **IET=** keyword `Distributed` with the keyword `DistributedSegregated`:

```
S-2009-A1-OS-thin-fast.tbm | Segregated-2009-A1-OS-thin-fast.tbm
5836 <MODELMAP>
5837     Electrolyte =   General Electrolyte
5838     IET =   DistributedSegregated 1D
5839 </MODELMAP>
```

  Leave the remaining text (`1D` in this example, but can be `3D` or `NP`) unchanged.

- Make the same change in the corresponding `<SIMMOD>` section.

  There can be several of these sections in the .tbm file, and they usually appear before the `<MODELMAP>` section. You edit only the variant (`1D`, `3D`, or `NP`) that is listed in the `<MODELMAP>` section (in this example, `1D`):

```
S-2009-A1-OS-thin-fast.tbm | Segregated-2009-A1-OS-thin-fast.tbm
5082 <SIMMOD>
5083 Distributed 1D
5084 LiIon\Pouch
5085 IET
5086     Negative Particle 1  Activation Ene
```

  Replace the keyword `Distributed` with the keyword `DistributedSegregated`, but leave the remaining text (`1D` in this example) unchanged:

```
S-2009-A1-OS-thin-fast.tbm | Segregated-2009-A1-OS-thin-fast.tbm
5082 <SIMMOD>
5083 DistributedSegregated 1D
5084 LiIon\Pouch
5085 IET
5086     Negative Particle 1  Activation Ene
```

3. Save the modified .tbm file under a suitable name.

Battery Modeling Bibliography

[421] Bernardi, D., Pawlikowski, E. and Newman, J. 1985. "A general energy balance for battery systems." *Journal of The Electrochemical Society.*, **132** (1), pp. 5–12.

[422] Bird, R. B., Stewart, E., and Lightfoot, E. N., Transport Phenomenon, ISBN 0-471-07392-X

[423] Doyle, M., Fuller, T. F. and Newman, J. 1993. "Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell." *Journal of The Electrochemical Society.*, **140** (6), pp. 1526–1533.

[424] Fuller, T.F., Doyle, M. and Newman, J. 1994. *J. Electrochem. Soc.*, **141**(1).

[425] Gu, H. 1983. *J. Electrochem. Soc.* **130** (7) pp. 1459-1464.

[426] Jang, Y.-I., Neudecker, B. J. and Dudney, N. J. 2001. "Lithium Diffusion in LixCoO2 (0.45 < x < 0.7) Intercalation Cathodes." *Electrochemical and Solid-State Letters.*, **4** (6), pp. A74–A77.

[427] Kim, U.S., Shin, Ch. B., Kim, C.S. 2009. *J. Power Src.* **189** pp. 841-846.

[428] Newman, J. and Thomas-Alyea, K. E. 2004. *Electrochemical Systems.*, 3rd edn. Wiley. ISBN: 978-0-471-47756-3.

[429] Newman, J. and Tiedeman, W. 1993. *J. Electrochem. Soc.* **140** (7) pp. 1961-1968.

[430] Ong, I. J. and Newman, J. 1999. "Double-layer capacitance in a dual lithium ion insertion cell." *Journal of The Electrochemical Society.*, **146** (12), pp. 4360–4365.

[431] J. Van Zee, G. Kleine, R. White, J. Newman, "Extension of Newman's Numerical Technique to Pentadiagonal Systems of Equations", in *Electrochemical Cell Design*, Edited by R. White, Plenum Press (1984).

[432] Verbrugge, M. and Tate, E. 2004. *J. Power Src.* **126** pp. 236-249.

Battery Field Functions Reference

The following primitive field functions are made available to the simulation when a battery module is prepared. Therefore the names in this table are all prefixed with the label Battery in the STAR-CCM+ user interface.

These objects all have the same base properties.

Cell Current	The current, I, passing between the posts of the battery cell (A).
Cell Voltage	The voltage drop, V, between the posts of the battery cell (V).

Electrical Solver Total Heat	The total generated heat from the electrical solver before it is mapped onto the thermal mesh (W).
Electrode Polarization Heat	Volume-based heat generated due to polarization of the two electrodes. This is initiated by the electrochemistry reaction (W/m^3).
Negative Current Density	A vector field denoting the flux of current through the negative electrode for each ecell (A/m^2).
Negative Electrochemical Current Density	A vector field denoting the flux of current through the negative electrode for each ecell in the axial direction of the cell (A/m^2). This field function is available for cylindrical and prismatic cells only.
Negative Plate Voltage	The voltage across the negative electrode for each ecell (V).
Ohmic Heat	Volume-based heat generated by Joule Effect for each ecell (W/m^3).
Positive Current Density	A vector field denoting the flux of current through the positive electrode for each ecell (A/m^2).
Positive Electrochemical Current Density	A vector field denoting the flux of current through the positive electrode for each ecell in the axial direction of the cell (A/m^2). This field function is available for cylindrical and prismatic cells only.

Positive Plate Voltage	The voltage across the positive electrode for each ecell (V).
Single True Cell Current	The current through a single true cell for each ecell (A).
SOC	The battery cell state of charge represents the current battery cell capacity as a percentage of its total capacity (%).
Tab Voltage	The voltage in each tab (V).
Thermal Mass	A material property which is an analogy for the heat capacity, *Thermal Mass* = *Cp x Mass* (J/K).
True Cells	Number of true cells in the battery pack. A true cell = 1 positive electrode + 1 negative electrode + 1 separator.
Volts	Equal to *(Positive Plate Voltage - Negative Plate Voltage)* for each ecell (V).
Volumetric Heat	Volume-based sum of all the generated heat: *Electrode Polarization Heat* (includes entropic heat) + *Ohmic Heat* (W/m^3).

Modeling Casting

This section contains the following:

- Casting Simulation with STAR-Cast
- Selecting the Casting Model
- Working With the Casting Material Database
- Using Imported Casting Materials and Properties in STAR-CCM+
- Using the Casting Shell Generator
- Using the Macro Porosity Model to Detect Shrinkage-Related Defects
- Defect Prediction Using Criteria Functions
- Artificial Compressibility
- Automated Time-Step Control
- Combined Thermal Contact Resistance
- Corrective Energy Source for Pure Thermal Simulations
- The Casting User Interface

Casting Simulation with STAR-Cast

The Casting Module, also referred to as STAR-Cast, is a casting computational software tool. STAR-Cast was jointly developed by CD-adapco and ACCESS. ACCESS is an independent research center that is associated with the Technical University of Aachen (RWTH) in Germany.

Casting simulations with STAR-Cast are performed as multiphase mould filling flows with solidification. STAR-Cast can be used as a design tool to assess the filling, solidification and cool-down process. It allows for defect prediction such as topological macro porosity prediction, micro porosity prediction, shrinkage due to phase change, trapped gas during filling as well as cold shots or misruns. Casting simulation enables the design of gating and runner systems before start of production.

About the Casting Module

In order to use the Casting Module in STAR-CCM+, you require a STAR-Cast license.

To check whether the Casting Module is available in STAR-CCM+:

1. In the STAR-CCM+ menu bar, select Tools > Add-ons.
2. In the *Add-on Licensing* dialog, verify that STAR-Cast is activated.



3. Click OK to close the *Add-on Licensing* dialog.

Casting Workflow

The process for using the Casting Module in STAR-CCM+ is outlined below:

1. Create a volume mesh.
2. Activate the Casting model in STAR-CCM+.

   See Selecting the Casting Model.
3. Use the Casting Material Database Editor to select the casting materials that you want to use in your simulation, and export the selected collection of materials to STAR-CCM+.

   See Working With the Casting Material Database.

4. Use the imported casting materials in physics continua and assign the appropriate material properties as boundary values.

   See Using Imported Casting Materials and Properties in STAR-CCM+.

5. Use the Casting Shell Generator to create shell molds that you can use in your casting simulations.

   See Using the Casting Shell Generator.

6. Use the Macro Porosity model to detect shrinkage-related defects in a cast part.

   See Using the Macro Porosity Model to Detect Shrinkage-Related Defects.

7. If necessary, use the available Criteria Functions to post-process the simulation results.

   See Defect Prediction Using Criteria Functions.

Selecting the Casting Model

The following steps describe how to select and use the Casting model:

1. Start STAR-CCM+ and create the appropriate physics continuum.
2. In the *Physics Model Selection* dialog, select Eulerian Multiphase from the *Material* group box.
3. Select Volume of Fluid (VOF) from the *Eulerian Multiphase Model* group box.

4. Select Casting from the *Optional Models* group box.



The Casting model appears under the Models node in the simulation tree.



The Casting node also appears under the Tools node in the simulation tree. This node contains the Material Database Editor node, which lets you access the Casting Material Database Editor.

See Using the Casting Material Database Editor.

If you run a pure thermal simulation, the Models > Casting node

5. If you run a pure thermal simulation, select the Models > Casting node and activate Apply Corrective Energy Source.

   See Activating the Corrective Energy Source for Pure Thermal Simulations

Working With the Casting Material Database

The accuracy of a casting simulation depends on the quality and completeness of the material data that you use. Therefore, a database that contains certified and qualified property data for casting is essential for high quality casting simulations. To meet these requirements, STAR-CCM+ provides the Casting Material Database. This database is a component of STAR-Cast.

The Casting Material Database contains three classes of material data:

- Licensed data: data accessible to the licensee. This data includes datasets for the most frequently posed problems.
- Proprietary data: private and confidential data belonging to individual owners. This data is either provided by ACCESS as a service for customers, or is made available by the owner. Access to this data is restricted to the owner and to ACCESS.
- User-defined: You can generate your own datasets, for which you take full responsibility.

  See Working with User-Defined Material Data.

The Casting Material Tree Structure

In the Casting Material Database, casting materials are categorized according to the tree structure shown below:



Materials are classified into three main groups, or material classes, as follows:

- Casting metals
- Mold materials
- Others

Each class is divided into a material group, and a material subgroup. The material designations are at a fourth level.

Using the Casting Material Database

You access the Casting Material Database with the Casting Material Database Editor.

See Using the Casting Material Database Editor.

The Casting Material Database Editor lets you do the following:

- Select a casting material from the Casting Material Database.

  See Selecting a Casting Material.
- View the physical properties of a selected casting material.

  See Viewing Available Material Properties for a Casting Material.
- Viewing the datasets that belong to a selected casting material.

You can extract graphs and material data information from datasets, and integrate this information into casting simulation reports.

See Viewing and Exporting Datasets.

- Compare the physical properties of different casting materials.

  See Comparing the Physical Properties of Different Materials.

- Export selected casting material data to STAR-CCM+.

  Material data that is exported to STAR-CCM+ can be included in your casting simulation.

  See Exporting Casting Material Data to STAR-CCM+.

- Create and modify user-defined casting material data.

  See Working with User-Defined Material Data.

Using the Casting Material Database Editor

The Casting Material Database Editor, also referred to as STAR-Cast mat, allows you to search for material data, view it, and export it to STAR-CCM+. The material data that you export to STAR-CCM+ can be included in your casting simulation.

Accessing the Casting Material Database Editor

To open the Casting Material Database Editor:

1. Activate the casting model.
2. Right-click the Tools > Casting > Material Database Editor node and select Casting Material Database Editor.



When you open the Casting Material Database Editor for the first time, the following disclaimer notice appears.



3. Read the provisions governing database rights and the disclaimer of liability carefully.
4. If you agree with the provisions, select the agreement radio button and click OK.

The Casting Material Database Editor appears.



You cannot access the STAR-CCM+ workspace while the Casting Material Database Editor is open.

The Casting Material Database Editor closes automatically when you export a collection of materials directly to STAR-CCM+.

To close the Casting Material Database Editor manually:

- Select File > Exit.

Changing the Appearance of the Casting Material Database Editor

You can change the appearance of the Casting Material Database Editor to suit your requirements. You can resize the window and set the text style, size, and font. You can also restore the default settings when necessary.

To specify the font to use in the Casting Material Database Editor:

1. Select Settings > Font.

2. In the *Select Font* dialog, make the appropriate changes to the font settings.



3. To confirm the new settings, click OK.

To restore the Casting Material Database Editor to the default font settings and window size:

- Select Settings > Reset GUI.

Getting Information on the Casting Material Database

To get help with using the Casting Material Database, do either of the following:

- Select Help > Help.
- Press `F1`.

To get product information specific to the Casting Material Database, do either of the following:

- Select Help > About.
- Click the General Information tab.

The General Information panel displays the units that the Casting Material Database uses, and lists the new material designations that have been added to the database recently.



Physical Property	STAR-Cast mat Units	Export Units (SI)
Density	[g/cm^3]	[kg/m^3]
Dynamic Viscosity	[mPa s]	[Pa s]
Elastic Modulus	[GPa]	[Pa]
Enthalpy	[J/g]	[J/kg]
Fraction Solid	[-]	[-]
Heat Transfer Coefficient	[W/m^2 K]	[W/m^2 K]
Permeability	[-]	[-]
Poisson Coefficient	[-]	[-]
Specific Heat Capacity	[J/g K]	[J/kg K]
Stress Strain Curves	[Pa]	[Pa]
Surface Tension	[N/m]	[N/m]
Thermal Conductivity	[W/m K]	[W/m K]
Thermal Diffusivity	[mm^2/s]	[m^2/s]
Thermal Expansion Coefficient	[10^(-6)/K]	[1/K]
Total Emissivity	[-]	[-]
Wetting Angle	[deg]	[deg]

Selecting a Casting Material

You select the casting materials that you want to use from the Casting Material Database Editor. By default, licensed and proprietary data is displayed, but user-defined data is hidden.

To display user-defined data:

- In the Casting Material Database Editor, on the *Settings* menu, activate Show User Defined Data.

  For more information, see Working with User-Defined Material Data.

All material designations, including user-defined ones, are displayed.

To select a casting material:

1. In the Casting Material Database Editor, click the Material Selection tab.

2. In the Material Selection panel, select the appropriate material designation by doing any of the following:
   - Choose the appropriate items from the four drop-down menus that reflect the tree structure.

     See Selecting a Material Designation Using the Drop-down Menus.
   - Select the appropriate material by using the icon lists.

     See Selecting a Material Designation Using the Icon Lists.
   - Run a string search for the material designation.

     See Selecting a Material Designation Using a String Search.
   - Enter a unique dataset number.

     See Selecting a Material Designation Using Its Dataset Number.

Details of the selected material are displayed in the Material Selection panel.

See Viewing Details of the Selected Material Designation.

You can also select materials from the Physical Properties panel. However, to do so, you must have previously selected another material designation.

See Selecting a Material Designation from a Material Property.

Selecting a Material Designation Using the Drop-down Menus

You can select a material designation from the four drop-down menus that reflect the tree structure.

To select a material designation:

1. Choose the material class from the Material Class drop down menu.
2. Select the appropriate material group from the Material Group menu.

3. Select the appropriate material subgroup from the Material Subgroup menu.
4. Select a material designation from the Material Designation menu.

Selecting a Material Designation Using the Icon Lists

You can select a material designation from the icon lists that are available to the right of the Material Class and Material Group drop-down menus.

To select a material designation from an icon list:

1. Click the appropriate icon list button.
2. In the *Available Materials* dialog, click the material designation that you want to use.



	Material Designation	Material Subgroup
1	AZ 91	Magnesium-aluminium alloys
2	CMSX-4®	Nickel-cobalt alloys
3	Cu-DHP	Cu-DHP
4	EN AB-AlSi7Mg0.3	Aluminium-silicon-alloys
5	EN AC-AlSi10Mg (a)	Aluminium-silicon-alloys
6	EN AC-AlSi12(Cu)	Aluminium-silicon-alloys
7	EN AC-AlSi7Mg0.3	Aluminium-silicon-alloys
8	EN AC-AlSi8Cu3	Aluminium-silicon-alloys
9	FeAl19	Iron-aluminium alloys
10	GE-48-2-2	Titanium-aluminium alloys
11	GJS 600	Cast iron with graphite(grey cast iron)

3. Click Select.

Selecting a Material Designation Using a String Search

You can select a material designation using a string search:

1. In the Material Designation box, enter the name of the material that you want (or a related string).



   If you want to see the complete set of material designations that are available in the Casting Material Database, leave the Material Designation box empty.

2. Click Search.

   The *Material Search* dialog appears.



3. In the *Material Search* dialog, click the material designation that you want.
4. Click Select.

Selecting a Material Designation Using Its Dataset Number

You can select a material designation by specifying the dataset number of one of its physical properties:

1. In the Dataset box, enter the appropriate dataset number.



2. Click Select.

The Datasets tab appears, displaying details of the selected dataset. The corresponding material is automatically selected in the Material Selection tab.

See Viewing the Details of a Dataset.

Viewing Details of the Selected Material Designation

After you have selected a material designation, the material number and the standard (for example, European, American, British) or a company-defined specification are displayed below the material drop-down menus in the Material Selection panel.



The Material Selection panel contains information text boxes that display general information about the selected material designation.



- Material Type: Nature of the material (as applicable) to which the Standard/Specification refers, or the specimens to which the measurements correspond.

  For example: castings, alloyed ingots for remelting, and alumina-based shell.

- Material Treatment: Describes how the material has been processed.

  For example: heat-treated and as cast condition.

- Equivalent Materials/Aliases/Umbrella Terms: A list of equivalent materials.

For example, alloys with similar composition or specifications, and aliases for the material designation in the Casting Material Database dialog.

- Description: The composition or specification of the material (as applicable) and other information about the composition and manufacturing or special properties.

  For example, porosity.

Viewing Available Material Properties for a Casting Material

The Physical Properties panel of the Casting Material Database Editor lets you view the available material properties for the selected material designation.

To view the material properties for a casting material:

1. In the Casting Material Database Editor, click the Physical Properties tab.



If a material designation is selected, the *Material Summary* group box displays the Material Designation, the Material Number (as applicable), and the Standard/Specification (as applicable).



The *Properties* group box lists physical properties under three headings: Fluid, Mechanical, and Thermal.



2. In the *Properties* group box, click the property that you want to view.

The Datasets panel opens, and displays dataset information for the selected property.

See Viewing the Details of a Dataset.

If data for a property is not available, the corresponding property button is grayed out.

Selecting a Material Designation from a Material Property

When a material designation is loaded in the Physical Properties panel, you can select a different material designation that has a particular property.

To select another material designation from a material property:

1. Click the icon list button that is next to the appropriate property button.

   A dialog listing all of the available materials for the corresponding property appears.



	Material Designation	Material Subgroup
1	Air	Air
2	CMSX-4®	Nickel-cobalt alloys
3	EN AB-AlSi7Mg0.3	Aluminium-silicon-alloys
4	EN AC-AlSi10Mg (a)	Aluminium-silicon-alloys
5	EN AC-AlSi7Mg0.3	Aluminium-silicon-alloys
6	EN AC-AlSi8Cu3	Aluminium-silicon-alloys
7	GE-48-2-2	Titanium-aluminium alloys
8	GJS 600	Cast iron with graphite(grey cast iron)
9	Inconel 718	Nickel-iron alloys
10	MAR-M-247	Nickel-cobalt-tungsten alloys
11	Pure silicon	Pure silicon

2. Click the material designation that you want to select.
3. Click Select.

   Details of the selected material designation are displayed in the *Material Summary* group box.

Viewing and Exporting Datasets

A dataset is a set of data values that represents the variation of a physical property with temperature. In the Casting Material Database, each dataset is assigned a unique number and is associated with a particular casting material designation. A particular physical property of a casting material designation can have two or more datasets that are associated with it.

You can view the details of a particular dataset, or view the details of all of the datasets that are associated with a particular property of a material.

See Viewing the Details of a Dataset.

To see the physical property data in graphical form, you can view and compare the dataset curves.

See Viewing Dataset Curves.

If two or more datasets are available for a particular property, you can select the dataset that is most appropriate for your simulation.

See Selecting a Dataset to Include in an Exported Material.

The units of measurement for a physical property in the Casting Material Database are expressed in the system of units that is the most commonly used in the literature. During the data export, the units are automatically converted according to the requirements of STAR-CCM+.

Viewing the Details of a Dataset

To view the details of a particular dataset:

1. In the Casting Material Database Editor, do one of the following:
   - In the Datasets tab, in the Dataset box, enter the appropriate dataset number.
   - In the Material Selection tab, in the Dataset box, enter the appropriate dataset number.
2. Click Select.

   The selected dataset is displayed. The material with which the dataset is associated is also selected and is displayed in the Material Selection tab.

To view the details of all of the datasets that are associated with a particular property of a casting material:

1. Select the appropriate casting material.

   Selecting a Casting Material

2. In the Physical Properties tab, in the *Properties* group box, click the appropriate physical property.

   The Datasets panel appears, displaying details of all of the datasets that are associated with the specified casting material and physical property.



Viewing Detailed Information about a Dataset

You can view a description of each dataset. You can also view the raw data for a dataset, including the error tolerances on each value.

The text box at the left of the Datasets panel contains information on each selected dataset. This information includes a description of the dataset and details of its origin.



To view the raw data for a dataset:

1. In the vertical bar to the left of the graph, select the dataset that you want.
2. Ensure that all other datasets are deselected.

   If two or more datasets are selected, the Show Table of Values button remains grayed out.
3. Click Show Table of Values.

The raw values of the dataset, including the error tolerance on each value, are displayed in tabular form.



Temp [°C]	Dynamic Viscosity [mPa s]	Tolerance [%]
615	1.38	10
700	1.2	10
800	1.1	10
900	1	10
1000	0.9	10

Viewing Dataset Curves

To see the temperature dependence of the physical property in graphical form, you can view the dataset curves. The symbol and ID number of each available dataset is shown in the vertical bar to the left of the physical property axis. You can select the datasets to display in the graph.

To display particular datasets:

- In the vertical bar, activate the appropriate dataset numbers.

To remove particular datasets from the display:

- Deactivate the appropriate dataset numbers.

To display all of the available datasets:

- In the lower menu bar, click Show Curves.

If you want to deactivate all of the datasets:

- Click Hide Curves.

The selected datasets are displayed.



Ordinates represent the physical property to which the plotted data relates. The abscissa denotes temperature. The graph displays liquidus and solidus temperatures (depending on the range of temperatures plotted), and latent heat information. A symbol on the graph represents each value in the dataset.

To determine the value of a particular dataset point:

- Hover the mouse over the appropriate dataset point.

  Its coordinates are displayed next to the pointer.

Changing the Appearance of a Graph

You can use the following options to change the appearance of the graph.

Grid	Toggles the background grid display.

ErrorBars

Toggles the display of error bars on the data points. The error bars reflect the tolerance that is specified on the data values.

In the drop-down list, select the type of error bars to use:

- Bars
- Tube

Zoom

Lets you zoom in on a particular area of the graph.

To select the area of interest, click one corner of the area and then drag the rectangle over the appropriate range.

To reset the range to the default value, deactivate Zoom.

Saving and Printing a Graph

You can save a graph as a PDF file that you can use later, for example, to include in a report. You can also print the graph directly to your printer. The generated output includes the dataset symbols and ID numbers to the left of the property axis.

Export

Saves the graph as a PDF file, using the default name `myPlot.pdf`.
The dataset information is saved automatically in a separate PDF file, using the same file name with `_Refs` appended.

Print

Prints the graph and the dataset information.

Selecting a Dataset to Include in an Exported Material

If two or more datasets are available for a material property, you can choose which dataset to use in your simulation. When you export the material as part of a material collection, the selected dataset is used to define the material property.

If a particular dataset is recommended, it is selected for export in the Export Data panel by default. If you want to use a different dataset, or if no dataset is selected, you can select the appropriate dataset.

To select a dataset for export:

1. In the vertical bar to the left of the graph, select the dataset that you want.
2. Ensure that all other datasets are deselected.

   If two or more datasets are selected, the Select Dataset for Export button remains grayed out.
3. Click Select Dataset for Export.

   The selected dataset number is shown in the *Selected Datasets* group box on the Export Data tab.

Comparing the Physical Properties of Different Materials

The Data Comparison panel lets you compare a single physical property for up to five different materials. This feature lets you compare the datasets that are available for particular materials, and select the materials and datasets that are the most appropriate to your simulation.

To compare a particular physical property for different materials:

1. In the Casting Material Database Editor, click the Data Comparison tab.
2. In the Physical Property box, click the down arrow and then select the physical property that you want to compare.



3. Select the materials that you want to compare:
   a. In the Material Designation box, enter all or part of the material name.

      If you want to list all available materials, leave the Material Designation box empty.
   b. Click Search to display the list of matched materials.
   c. In the *Material Search* dialog, select the materials that you want to compare.

      You can select up to five materials. If you attempt to select any more, the Select button is grayed out.

d. Click Select.

The selected materials are added to the comparison slots on the Data Comparison panel.



If a material designation has two or more datasets that are associated with the selected physical property, the recommended curve and its documentation are shown by default.

4. If you want to select another dataset, click the drop-down list under the material designation and select the appropriate dataset number.



You can remove or replace a material designation:

1. Click Remove below the material designation.

The material designation field is cleared and the associated dataset is removed.

2. If you want to display a different material, click Search and then select the appropriate material designation.

You can use the graphical display features to view and compare the dataset curves. You can also save a graph as a PDF file to use later, or print a graph directly to your printer.

For more information, see Viewing Dataset Curves.

Exporting Casting Material Data to STAR-CCM+

The Export Data panel lets you prepare the collection of materials for export to STAR-CCM+.

See Preparing a Collection of Materials for Export.

You can export the collection of materials directly to STAR-CCM+, or save it as a file that you can import later.

See Exporting the Collection of Materials.

Preparing a Collection of Materials for Export

In order to export material data from the Casting Material Database to STAR-CCM+, you prepare the appropriate collection of materials.

Adding a material to the collection

For each material that you want to export:

1. In the Casting Material Database Editor, in the Material Selection tab, select the material designation.

   See Selecting a Casting Material.

2. Review the datasets that are associated with the material and, if necessary, select the datasets that you want to use.

   If two or more datasets are available for a particular property, you can select the appropriate dataset in the Datasets panel.

   See Selecting a Dataset to Include in an Exported Material.

3. Click the Export Data tab.

The *Selected Datasets* group box shows the datasets that are currently selected for the material.



The letter R to the right of each dataset denotes that the recommended dataset is selected for that particular material property.

4. Verify that the selected datasets are appropriate for your simulation.

   If alternative datasets are available for a particular property, you can change the selection in the Datasets panel.

   See Selecting a Dataset to Include in an Exported Material.

5. If you want to see summary information for a selected dataset, click the blue button to the right of the dataset.

The dataset documentation summary is displayed in the *Dataset* dialog.



6. To add the current material designation (labeled at the top of the Export Data panel) to the collection of materials for export, click Add Material to Collection.

   The selected material is added to the Collection of Materials list.

Removing a Material from the Collection

If you want to remove a material designation from the collection of materials:

1. In the Collection of Materials list, select the material designation that you want to remove.
2. Click Remove Material from Collection.

Reviewing the Selected Materials

You can review the materials that are currently selected:

1. In the Collection of Materials list, click the material that you want to review.

A summary of the selected datasets for each material property is displayed in the lower text box.



Exporting the Collection of Materials

When you have prepared the appropriate collection of materials, you can export the materials and the selected physical properties from the Casting Material Database to STAR-CCM+.



To export the collection of materials directly to STAR-CCM+:

- Click Export Collection to STAR-CCM+.

  The Casting Material Database Editor closes. The material collection is saved as an XML file with a predefined name (`ExportToStarCCM_YYMMhhmmss.xml (Yr,Month,hr,min,sec)`) under the current working directory.

STAR-CCM+ automatically reads the exported material database and loads it under the Tools > Material Databases > Selected_Casting_Materials node.



To export the collection of materials to a file that you can import later:

1. Click the appropriate option:
   - Export Collection to File

     Saves the collection of materials as a .scm file.
   - Export Collection to XML

     Saves the collection of materials as an .xml file.
2. In the *Save File* dialog, specify the appropriate file name and location and then click Save.

The Casting Material Database Editor remains open. You can import the exported file into the STAR-CCM+ Material Database later.

See Using Imported Casting Materials and Properties in STAR-CCM+.

Working with User-Defined Material Data

By default, licensed and proprietary data is displayed, but user-defined data is hidden.

To display user-defined material data:

1. In the Casting Material Database Editor, on the *Settings* menu, activate Show User Defined Data.
2. Launch the *Material Search* dialog by running an empty string search.

All material designations, including user-defined ones, are displayed.

If you have write permission on the database and the directory in which the database resides, you can create and edit your own datasets in the Casting Material Database Editor.

You can create your own datasets in the Casting Material Database Editor. The user-defined material data is saved to a local database file called `amdb.sqlite3` in your current STAR-CCM+ working directory.

For more information, see Creating User-Defined Material Data.

Creating User-Defined Material Data

You create user-defined material data by copying an existing material designation under a new name, and then modifying its data as appropriate.

To create user-defined material data:

1. In the Casting Material Database Editor, select an existing material designation whose properties values are close to your user-defined properties.
2. In the *Material Selection* panel, click Copy to User Defined Data.
3. In the *Copy material to user defined data* dialog, in the User defined material designation box, enter the appropriate name.



Alternatively, you can leave the default name composed of the original material designation from which the data have been copied and some postfix notation.

4. Optional: Enter a comment in the Comment box.
5. Click OK to confirm the new details.

The user-defined material designation is added to the material subgroup that the original material designation belongs to.

6. Make the appropriate changes to the new user-defined material designation.

   You can import and edit the datasets that define the physical properties of the material. You can also modify the comments that are associated with the material and the datasets.

   See Editing User-Defined Material Data.

Editing User-Defined Material Data

To edit a user-defined material designation:

1. In the Casting Material Database Editor, select the user-defined material designation that you want to modify.
2. In the *Material Selection* panel, click Edit User Defined Data.

   The *Edit user defined data* dialog appears:



3. In the *Edit user defined data* dialog, use the following options to make the appropriate changes:

Edit comment	Lets you modify the user-defined data comment. • Enter the comment in the *Edit comment* dialog, and then click Save.
Delete user defined material	Lets you delete the user-defined material from the Casting Material Database.
Add new dataset to property	Lets you create a user-defined dataset for the selected physical property. See Creating a User-Defined Dataset.
Edit dataset comment	Lets you specify or modify the comment for the selected dataset. See Adding a Comment to a Dataset.
Edit dataset data	Lets you specify or modify the data values for the selected dataset. See Editing Dataset Data Values.
Delete dataset	Lets you delete the selected dataset from the Casting Material Database.

Creating a User-Defined Dataset

To create a user-defined dataset:

1. In the *Select physical property* drop-down list, select the appropriate physical property.



2. Click Add new dataset to property.

   The *Add new dataset to property* dialog appears.



3. In the *Add new dataset to property* dialog, specify the following details:

From File	Specify the database input file that you want to use. • Click Browse and then open the appropriate database input file.
Delimiter	Specify the delimiter that is used to separate columns in the database input file.

X Data Column	Specify the temperature column number.
Y Data Column	Specify the physical property column number.
Ignore Lines Beginning With	Lets you specify the character that indicates the beginning of a comment in the database input file. • Activate this option and then specify the appropriate comment character in the text box.

4. Click OK to confirm the settings.

   The specified database input file is imported, and the *Info* dialog indicates the dataset number for the new user-defined dataset.



Adding a Comment to a Dataset

To add or modify a comment for a dataset:

1. In the *Select physical property* drop-down list, select the appropriate physical property.
2. In the adjacent drop-down list, select the dataset that you want to modify.



3. Click Edit dataset comment.
4. In the *Edit dataset comment* dialog, edit the dataset comment, and then click Save.

Editing Dataset Data Values

To edit dataset data values:

1. In the *Select physical property* drop-down list, select the appropriate physical property.
2. In the adjacent drop-down list, select the dataset that you want to modify.
3. Click Edit dataset data.

The *Data editor* dialog appears.



	X	Y
1	0	2.678
2	24	2.675
3	48	2.672
4	73	2.667
5	99	2.663
6	125	2.658
7	151	2.654
8	177	2.649
9	201	2.644
10	227	2.639
11	252	2.633
12	277	2.626
13	302	2.62
14	325	2.615
15	352	2.61

To add a row:

- Click in the row above where you want to insert the new row, and then click Add Row.

  You cannot add a row above the first row.

To delete a row:

- Click in the row that you want to delete, and then click Del Row.

To save changes to the dataset:

- Click Save.

To exit without saving changes:

- Click Cancel.

Using Imported Casting Materials and Properties in STAR-CCM+

You can import casting materials and properties into the STAR-CCM+ Material Database.

See Importing Casting Materials into the STAR-CCM+ Material Database.

You can use the imported casting materials as single components, phases, or within multi-component mixtures in a physics continuum.

See Using Imported Physical Properties in Physics Continua.

You can use some of the imported material properties as boundary values in STAR-CCM+.

See Using Boundary Values Exported from the Material Database Editor.

Importing Casting Materials into the STAR-CCM+ Material Database

You import the casting material data from an XML file that is created in the Casting Material Database Editor.

For more information, see Exporting Casting Material Data to STAR-CCM+.

1. If necessary, create a new Material Database in STAR-CCM+.

2. Right-click the material database to which you want to add casting materials, and select Import Casting Materials and Properties.



3. In the *Open* dialog, select the appropriate XML file.

   The casting materials are imported into the material database and appear under a tree structure representing the material hierarchy of the Casting Material Database Editor.

For example, a material designation for an aluminum silicon alloy appears under Casting metals > Aluminium alloys > Aluminium-silicon-alloys.



Using Imported Physical Properties in Physics Continua

Once the selected casting materials are imported into STAR-CCM+, they can be selected as single components, phases or within multi-component mixtures in the Physics Continua.



When the Table(T) method is chosen for an imported physical property, the Input Table property defaults to `Material Database Table` provided the imported material has tabulated data in the STAR-CCM+ material database. Otherwise it is set to `None`.

Using Boundary Values Exported from the Material Database Editor

Some material properties from the Casting Material Database Editor are defined as boundary values by STAR-CCM+. These properties include:

- Emissivity
- Heat Transfer Coefficient

- Wetting Angle (Contact Angle).

To use these properties in STAR-CCM+:

1. Export the material data to a `.csv` file.



2. Create a table using the exported `.csv` file.
3. Create a field function and use the table interpolation syntax to generate a scalar expression for the material property.

   For casting properties, the `<scalar_expression>` is `$Temperature`.



4. Use this field function to specify the boundary value profile.



You can also use this workaround in the physics continuum for any material property which does not possess a Table(T) method.

Using the Casting Shell Generator

This section contains the following sections:

- About the Casting Shell Generator
- Accessing the Casting Shell Generator
- Preparing to Use the Casting Shell Generator
- Using the Casting Shell Generator

About the Casting Shell Generator

The Casting Shell Generator (also referred to as the Automatic Shell Generator, or ASG) is part of the STAR-Cast installation and is provided with the Casting Module in STAR-CCM+. The Casting Shell Generator is a geometry extrusion tool that allows you to create shell molds for investment casting simulations. To generate the shell mold, the Casting Shell Generator

extrudes the surface describing the geometry normally outwards thus creating a new region. The modeled shell region is thicker in concave areas and thinner in convex areas -- just like a real shell mold.

The Casting Shell Generator can be applied as part of the mesh generation process, when the shell mold is to be included in the simulation as a solid region for a thermal conjugate heat transfer analysis.

Accessing the Casting Shell Generator

To make the Casting Shell Generator accessible in STAR-CCM+, add the Shell Generator icon to the toolbar:

1. In the STAR-CCM+ menu, select Window > Toolbars > Customize.



2. In the *Customize Toolbars* dialog, under the ACCESS Shell Tool node, select Shell Generator.



3. Drag the Shell Generator icon onto one of the toolbars at the top of the STAR-CCM+ workspace and drop it at the appropriate position.

   The icon is added to the toolbar at the specified location.

4. In the *Customize Toolbars* dialog, click Close.

Preparing to Use the Casting Shell Generator

Before you use the Casting Shell Generator to create a shell, consider the following:

- A shell mold is extruded from a surface mesh. To create an extrusion of good quality, make sure that the surface to be extruded from has a smooth triangle growth rate. The faces of imported surface meshes (for example imported from STL format) can consist of long, sharp triangles,

which are not suitable for use with the Casting Shell Generator. The Remeshed Surface Representation provides the required triangulation quality and is typically applied.

See Selecting the Representation.

- You can improve the extruded shell mold quality (for example in case of pierced faces) by using the surface wrapper. You can activate this feature in the Casting Shell Generator, and the generated shell is automatically wrapped and remeshed as part of the shell generation process.

  See Setting the Wrapper/Remesher Options.

- Do not select the surface wrapper when generating multiple shell mold layers.

  See Setting the Shell Generation Options

Using the Casting Shell Generator

The Casting Shell Generator loads default values when it starts for the first time. Each time that the Casting Shell Generator is restarted, it loads the values that were last used.

If you want to restore the default values, delete the `.asgConfig` directory in the home directory.

To use the Casting Shell Generator:

1. In an open simulation, activate the Casting model.

   See Selecting the Casting Model.

2. Ensure that you have the appropriate surface mesh available.
3. Click the (Shell Generator) button in the STAR-CCM+ toolbar.

The *Generate Shells* dialog appears.



4. In the *Generate Shells* dialog, set the appropriate options in the following group boxes:
   - *Boundaries*

     See Selecting the Boundaries.
   - *Representations*

     See Selecting the Representation.
   - *ASG Options*

     See Setting the Shell Generation Options.
   - *Advanced ASG Options*

     See Setting the Advanced Shell Generation Options.
   - *Wrapper/Remesher Options*

     See Setting the Wrapper/Remesher Options.
5. Click OK to start the shell generation.

   The *Output* window indicates the progress of the shell generation process. The information that is displayed includes the thickness of the shell that has been calculated and the thickness that remains to be calculated.
6. If you want to cancel the shell generation process, click Abort Shell Generation.

When the shell generation process is complete, the *Output* window closes.

The generated shell is added to the simulation tree under the Regions node. The default name for the shell region is Region 1.



Selecting the Boundaries

The Casting Shell Generator lets you specify the boundaries on which the shell is generated.

The settings in the *Boundaries* group box are described below:

Regions	Lists the available regions and lets you choose the region from which to select boundaries.
Boundaries	Lists the boundaries that are in the selected region and lets you select the boundaries that you want to use. To add the selected boundaries to the list in the Selected panel, use the appropriate button: • Click >> to add all of the boundaries to the selection. • Click > to add only the selected boundaries.

Selected	Lists the boundaries that are currently selected for use. If you want to remove any boundaries from the list in the Selected panel, use the appropriate button: • Click < to remove the selected boundary. • Click << to remove all of the boundaries.

Selecting the Representation

The Casting Shell Generator extrudes directly from the surface mesh. Before you start the shell generation process, ensure that the surface mesh captures the geometry complexity in all desired detail. Select the surface mesh representation for generating the shell.

To select the surface mesh representation to use:

- In the *Representations* group box, in the Choose Representation drop-down list, select the appropriate surface mesh representation.

Setting the Shell Generation Options

The Casting Shell Generator lets you specify how the shell is generated. The most important settings are the average thickness of the generated shell and the number of layers.

The settings in the *ASG Options* group box are described below:

Import in	Lets you specify how the generated shell mold is imported: • Regions • Parts
Total thickness (mm)	Specifies the average thickness of the generated shell mold. The average thickness is the thickness on a plane. If the number of layers is greater than 1, the average thickness is divided linearly by the number of shells.
Number of Layers	Specifies the number of layers that the shell generator generates.

Change growing orientation	When activated, this option inverts the shell growing orientation. By default, the shell growing direction is the surface mesh orientation that is defined by STAR-CCM+.

Setting the Advanced Shell Generation Options

The Casting Shell Generator lets you control the thickness of the generated shell in the curved regions. The shell thickness is a function of curvature and a specified multiplication factor: a larger multiplication factor increases the shell thickness.

You can set the thickness of convex curves and concave curves separately. Even (not curved) regions are not affected.

The settings in the *Advanced ASG Options* group box are described below:

Multiply step length by	Specifies the factor by which the growing step length is multiplied. This value has a significant effect on shell generation speed and quality. The algorithm is faster for higher values, but the shell quality decreases. The best quality values are in the range 0.19 – 0.55.
Convex Curve	Specifies the factor by which the growing vectors are multiplied in the convex regions. The practical values for this setting are in the range from 0 through 25.
Concave Curve	Specifies the factor by which the growing vectors are multiplied in the concave regions. The practical values for this setting are in the range from 0 through 25.
Use mesh checks	When activated, this option turns off the internal checking algorithms.

Use non-uniform thickness	When activated, different shells can have different thicknesses.

Setting the Wrapper/Remesher Options

To improve the quality of the generated shell (for example in case of pierced faces), you can use the Surface Wrapper and Surface Remesher. When you activate this option, the Surface Wrapper and Surface Remesher are applied automatically as part of the shell generation process.

The settings in the *Wrapper/remesher options* group box let you activate the STAR-CCM+ surface wrapper and surface remesher, and set the appropriate base size and the wrapper scale factor.

The settings in the *Wrapper/remesher options* group box are described below:

Use wrapper/remesher after shell generation	Activates the surface wrapper and surface remesher in STAR-CCM+.
Base size (m)	Lets you specify the appropriate value.
Wrapper scale factor (%)	Lets you specify the appropriate value.
Specification method	Lets you specify the surfaces to wrap: • External • Largest Internal

Using the Macro Porosity Model to Detect Shrinkage-Related Defects

You can use the Macro Porosity model to detect shrinkage-related defects in a cast part. Most casting materials shrink during solidification and cooling. If isolated liquid melt regions occur during the solidification process of a cast part, shrinkage cannot be compensated by supplying extra melt, and pores form. The Macro Porosity model uses an auxiliary void phase to model these shrinkage-related cavities. The cavities can be vacuum bubbles, or they can be filled with gas that is dissolved in the liquid phase (hydrogen, for example). In either situation, the Macro Porosity model lets you specify the appropriate void pressure.

To select the Macro Porosity model:

1. Select the Casting model.
2. Select Multiphase Interaction from the *Optional Models* group box.
3. In the simulation tree, under the Models > Eulerian Multiphase node, create the Eulerian phase.
4. Under the Eulerian phase, open the *Phase Model Selection* dialog and select the following models:
   - Select Liquid from the *Material* group box.

     The Macro Porosity model is available for single-component liquids only.
   - Select User Defined EOS from the *Equation of State* group box.

     The Macro Porosity model is incompatible with the Constant Density and Polynomial Density models.
5. Click Close.
6. In the simulation tree, under the Eulerian phase, select the User Defined EOS node and, in the *Properties* window, activate Compressible.

   The Macro Porosity model is compatible with the User Defined EOS when the Compressible property is activated and when $d\rho/dp > 0$. The liquid phase must be compressible because otherwise, a positive VoidPhase source does not lead to a pressure increase in that cell.
7. Under the Eulerian phase, open the *Phase Model Selection* dialog and select Melting-Solidification from the *Optional Models* group box.
8. Select Macro Porosity from the *Optional Models* group box.



9. Click Close.

The Macro Porosity node is added to the Eulerian phase models in the simulation tree.



Macro Porosity Properties

Void Pressure	The pressure p_{void} in Eqn. (2493) below which a gaseous phase void is assumed to outgas from the liquid phase. The default value is 10000 Pa.

Macro Porosity Expert Properties

Under-Relaxation Factor	At each iteration, this property governs the extent to which the newly computed void phase fraction supplants the old solution. The default value is 0.8.
Scaling Factor	s in Eqn. (2499). Limited to $s \geq 0$.

What Is the Auxiliary Void Phase?

When the Macro Porosity model is selected, an additional Eulerian phase is created automatically. This phase is named Void with the name of the solidifying phase as a suffix.



By default, the auxiliary void phase is an ideal gas with the material properties of air. This gas is the fluid that fills the pores (cavities) of the respective solidifying phase. Even if the pores are evacuated because the liquid does not emit any gas, the void pressure must still be specified significantly above the Minimum Absolute Pressure. To check the value of the Minimum Absolute Pressure, go to the *Segregated Flow Properties* window under the Physics Continuum > Models > Segregated Flow > Minimum Absolute Pressure. If the difference between the void and minimum pressures is too small, the system mass is not conserved because the pressure correction is limited.

If the material properties of the auxiliary void phase are modified, make sure that the following prerequisites are met:

- $d\rho/dp$ must be larger than zero.
- This void phase must be a compressible low-density gas that has a much lower density than the liquid phase.

  As the ratio of liquid to void phase density gets closer to 1, more fluid is required to be converted from liquid to void phase to achieve a certain pressure increase in a cell.

You cannot delete the solidifying phase or the auxiliary void phase directly because an auxiliary phase interaction uses both phases. To delete the auxiliary void phase, deactivate the Macro Porosity model in the solidifying phase.

What Is the Auxiliary Phase Interaction?

When the Macro Porosity model is selected, a multiphase interaction is created automatically. This phase interaction is named MacroPorosity with the name of the solidifying phase as a suffix.



Two phase interaction models are selected: Macro Porosity Phase Interaction and VOF Phase Interaction (fixed phases). The primary and secondary phases that are selected in the VOF Phase Interaction (fixed phases) model are read-only properties. These two models cannot be deselected and you cannot select any other models in this phase interaction. However, you can create other phase interactions that act on the solidifying and auxiliary phases (for example, to model surface tension).

If you delete this phase interaction, the Macro Porosity model is automatically deactivated and the auxiliary void phase is deleted. When the Macro Porosity model is deactivated in the solidifying phase, both this phase interaction and the void phase are deleted.

Variation of the Scaling Factor

Consider a liquid region that is enclosed by flow-stopped cells and cooled by a volumetric heat source. The temperature decrease causes shrinkage that the void gas compensates for such that the mixture density remains constant. The pressure asymptotically approaches a constant value slightly below the target void pressure depending on the cooling rate and the

scaling factor. After reaching a constant pressure level, the volumetric heat source is deactivated, and the pressure asymptotically approaches the specified void pressure.

The scaling factor s is a multiplicative factor to the source rate of the void phase. The scaling factor influences the minimum pressure and the time at which it is reached. The lower the source rate, the longer it takes for the void phase to compensate the pressure decrease. Hence, the largest scaling factor—2.0 in the following plot—yields the largest minimum pressure value.



Macro Porosity Field Functions Reference

The following primitive field functions are available to the simulation when the Macro Porosity model is used.

Density of Void<Phase>	The density of the auxiliary void phase.
Dynamic Viscosity of Void<Phase>	The dynamic viscosity of the auxiliary void phase.

Mass Flow Rate of Void<Phase>	The mass flow rate of the auxiliary void phase.
Mass Flux of Void<Phase>	The face mass flux of the auxiliary void phase.
Mass Imbalance of Void<Phase>	The imbalance of auxiliary void phase mass in a particular cell.
Volume Fraction of Void<Phase>	The fraction of volume of the current cell that the auxiliary void phase occupies.

The following primitive field function is available when the Temporary Storage Retained expert property of the Segregated VOF Solver is activated.

Source Rate of Void<Phase>	The VOF phase source for the auxiliary void phase.

Macro Porosity Formulation

This section describes:

- The Macro Porosity model criterion
- The void phase source
- The void phase source linearization

This section also includes a nomenclature that defines the terms that are used in the Macro Porosity formulation.

Macro Porosity Model Criterion

When the static pressure in a cell falls below the user-specified void pressure, a fraction of the liquid phase is converted to the gaseous void phase. The total fluid mass is kept constant. At the same pressure, a certain amount of a light gas fills a larger volume than the same mass of a heavy liquid. Since the cell volume is constant, the pressure is increased. This mechanism is used to keep the pressure from decreasing significantly below the target void pressure.

The equation of state of an ideal gas, the void phase void, is:

$$p = \rho \frac{R_u}{M_{\text{void}}} T \tag{2492}$$

with the universal gas constant $R_u = 8314.4621$ J/kmol K, and M_{void} is the molecular weight in kg/kmol.

Mass and cell volume are constant, and pressure increases during the phase conversion. Therefore, the following relation is fulfilled even for the minimum temperature in any cell during the simulation that is run at $p = p_{\text{void}}$:

$$p_{\text{void}} \leq \rho_l \frac{R_u}{M_{\text{void}}} T_{\min} \tag{2493}$$

Hence, the Macro Porosity model can operate correctly when the molecular weight of the ideal gas void phase is:

$$M_{\text{void}} \leq \frac{\rho_l R_u T_{\min}}{p_{\text{void}}} \tag{2494}$$

However, this criterion is easily met for common casting material properties:

$$M_{\text{air}} = 28.9664 \leq 44898.1 = \frac{2700 \cdot 8314.4621 \cdot 100}{50000}$$

Void Phase Source

The rate at which the liquid phase is converted to the void phase without pressure decreasing below p_{void} must be determined.

The relative volume change of the cell that leads to a pressure increase up to p_{void} is computed. The result is the fraction of the cell volume that must be filled with void gas at $p = p_{\text{void}}$ and $T = T_{\text{cell}}$.

Δp is the difference between the user-specified void pressure p_{void} and the static pressure p in the cell:

$$\Delta p = p_{\text{void}} - p \tag{2495}$$

This pressure difference causes a change in mixture density in the cell, approximated with the derivative of mixture density with respect to pressure $\partial\rho/\partial p$, assuming that $\partial\rho/\partial p$ is constant for small pressure deviations:

$$\Delta\rho \approx \Delta p \frac{\partial\rho}{\partial p} \tag{2496}$$

The relation between a density change and the related volume change for constant fluid mass in a cell is:

$$m = \rho V$$

$$\rho V = (\rho + \Delta\rho)(V + \Delta V)$$

$$\frac{\rho V}{\rho + \Delta\rho} - V = \Delta V$$

$$\frac{\rho}{\rho + \Delta\rho} - 1 = \frac{\Delta V}{V}$$

$$-\frac{\Delta\rho}{\rho + \Delta\rho} = \frac{\Delta V}{V} \tag{2497}$$

Eqn. (2497) gives the fraction of the cell volume that must be filled with the void phase to compensate (inverted sign) the volume change of the mixture when the pressure is increased again to void pressure:

$$\frac{\Delta V}{V} = \frac{\Delta\rho}{\rho + \Delta\rho} \tag{2498}$$

A VOF-phase source is specified in STAR-CCM+ in the units $\frac{\text{Volume}}{\text{Volume-Time}} = \frac{1}{\text{s}}$. The relative volume change applied per unit time (second) and an under-relaxation factor gives the rate at which the void fluid is converted from the liquid phase to the gaseous phase.

$$\left.\frac{\Delta V}{V}\right|_{\text{void}} = \max\left[0, \frac{\Delta\rho}{\rho + \Delta\rho} s\right] \tag{2499}$$

$$= \max\left[0, \frac{(p_{\text{void}} - p)\frac{\partial\rho}{\partial p}}{\rho + (p_{\text{void}} - p)\frac{\partial\rho}{\partial p}} s\right]$$

Values without index are cell values of the mixture. The scaling factor s is limited to values $s \geq 0$. It allows an adaptation to the specific case where larger rates are required to increase the minimum pressure.

The rate is limited to positive values to prevent emitted gas being dissolved again.

Void Phase Source Linearization

The linearization of the void phase source term is the derivative of Eqn. (2499) with respect to pressure (assuming $\partial\rho/\partial p$ is constant):

$$\frac{\partial}{\partial p}\left(\frac{\Delta V}{V}\bigg|_{\text{void}}\right) = \frac{\partial}{\partial p}\left[\frac{(p_{\text{void}} - p)\frac{\partial \rho}{\partial p}}{\rho + (p_{\text{void}} - p)\frac{\partial \rho}{\partial p}} s\right] \tag{2500}$$

$$= -\frac{p\frac{\partial \rho}{\partial p}}{\left[\rho + (p_{\text{void}} - p)\frac{\partial \rho}{\partial p}\right]^2} s$$

Macro Porosity Nomenclature

The terms that are used in the Macro Porosity formulation are defined below.

ρ	Density
M_{void}	Molecular weight of the void phase
R_u	Universal gas constant
T	Temperature
V	Volume
p	Pressure
s	Scaling factor

Defect Prediction Using Criteria Functions

Criteria functions are post-processing items to help predict casting defects like macro- or micro-porosity and to estimate relevant microstructural quantities such as primary and secondary dendrite arm spacings. They are scalar and vector quantities available as field functions when the following models are selected:

- Criteria Functions model

  Select this model for any criteria function to become available.
- Primary Dendrite Arm Spacing model

  Select the Criteria Functions model before this model becomes available.
- Secondary Dendrite Arm Spacing models

  Select the Criteria Functions model before this model becomes available.
- Dimensionless Niyama model

You must select the Criteria Functions model and one of the Secondary Dendrite Arm Spacing models before this model becomes available. You are advised to use the Secondary Dendrite Arm Spacing (CR) (see Secondary Dendrite Arm Spacing (CR) model in combination with the Dimensionless Niyama model. The original formulation of the Dimensionless Niyama criterion [433] uses the secondary dendrite arm spacing computed from the cooling rate.

Criteria functions can be plotted in a scene for visualization where the distribution of these quantities allows conclusions about casting defects.

Some criteria functions are collected over the complete duration of the simulation and contain the complete solution history. Therefore, whenever any property of the Criteria Functions model or any property relevant to the melting/solidification process is modified, the simulation must be restarted from scratch. No solution history file or transient post file is required to evaluate Criteria Functions: all the required data is stored in the `.sim` file.

The main point of interest is to evaluate the criteria functions at the critical temperature when the macroscopic flow stops due to solidification. Liquid melt cannot be further supplied to compensate for shrinkage effects. However, this critical temperature is not known exactly because the macroscopic stop of flow depends on the casting process and the alloy. Therefore, a set of temperatures close to the estimated critical temperature is specified in the range between the solidus and liquidus temperature. The criteria functions are evaluated at the several specified critical temperatures.

Setting Up a Criteria Function

To set up a Criteria Function:

1. Select the Casting model.
2. Select Multiphase Interaction from the *Optional Models* group box.
3. In the simulation tree, under the Models > Eulerian Multiphase node, create the Eulerian phase.
4. Under the Eulerian phase, open the *Phase Model Selection* dialog and select the following models:
   a. Select Liquid from the *Material* group box.
   b. Select Melting-Solidification from the *Optional Models* group box.
   c. Select Solidification Time from the *Optional Models* group box.

      You must select the Solidification Time model first for the Criteria Functions model to become available.

d. Optionally, select Liquid Residence Time from the *Optional Models* group box.

e. Select Criteria Functions from the *Optional Models* group box.



5. In the *Phase Model Selection* dialog, within *Optional Models*, select the particular criteria function that you require:
   - Select Primary Dendrite Arm Spacing.
   - Select Secondary Dendrite Arm Spacing and then one of the following sub-models:
     - Secondary Dendrite Arm Spacing (CR)
     - Secondary Dendrite Arm Spacing (LST)
   - Select Dimensionless Niyama Criterion.

Note: You are advised to select the Secondary Dendrite Arm Spacing (CR) model in combination with the Dimensionless Niyama criterion, because only this model complies with the original formulation of the Dimensionless Niyama model [433].

6. Click Close.

The Criteria Functions node, the Liquid Residence Time node, and the Solidification node are added to the Eulerian phase models in the simulation tree. The Initial Conditions node is added to the Eulerian phase.



7. Select each child node of the Initial Conditions node in turn and set the initial conditions for:
   - Liquid Residence Time
   - Solidification Time

Criteria Functions Definitions

The following Criteria Functions are available:

- Liquid Residence Time
- Solidification Time
- Local Solidification Time
- Cooling Rate

- Isotherm Speed
- Temperature Gradient
- Mean Cooling Rate
- Solidification Velocity
- G/v
- Niyama Criterion
- Dimensionless Niyama Criterion
- Primary Dendrite Arm Spacing
- Secondary Dendrite Arm Spacing

Criteria Functions Properties

Critical Temperatures	The array of critical temperatures between the solidus and liquidus temperatures at which Criteria Functions are evaluated. When you specify one or more critical temperatures, the Solidification Time function is activated. You can enter any number of critical temperatures.
Input Units	The unit for the Critical Temperatures array.

Liquid Residence Time

The fluid enters the computational domain in a liquid state with a temperature above the solidus temperature. The Liquid Residence Time is the time between the fluid entering the computational domain and the moment that the temperature decreases below the solidus temperature.

This time is computed for every VOF phase for which you activate the Liquid Residence Time model (see Setting Up a Criteria Function).



Liquid Residence Time Properties

Max. Solid Volume Fraction	The residence time is accumulated only for solid volume fractions below the specified value.
Convection	The convection scheme to use. For guidance on selecting a convection scheme, see Convection Term. 1st-order: First-order convection scheme. 2nd-order: Second-order convection scheme.
Min. Volume Fraction	The residence time is accumulated only for liquid volume fractions above the specified value.

Liquid Residence Time is transported only by convection. The initial value is zero everywhere. A source term adding the time step to the residence time at every time step is applied to every cell in which the temperature of the VOF phase is above solidus temperature.

The Liquid Residence Time is zero for every cell with a phase volume fraction below the minimum volume fraction.

Note: The computation of the Liquid Residence Time is optional, because no other Criteria Function depends on this value.

Solidification Time

The Solidification Time is the time that is taken to cool down the VOF phase from liquidus temperature to solidus temperature.

This time is computed for every VOF phase for which you activate the Solidification Time model (see Setting Up a Criteria Function). The appropriate critical temperatures must be specified in the Criteria Functions model.



Solidification Time Properties

Convection	The convection scheme to use. For guidance on selecting a convection scheme, see Convection Term.	
	1st-order	First-order convection scheme.
	2nd-order	Second-order convection scheme.
Min. Volume Fraction	The residence time is accumulated only for liquid volume fractions above the specified value.	

Solidification time is transported only by convection. The initial value is zero everywhere. In every cell in which the VOF phase is between the liquidus and solidus temperatures, a source term that adds the time step to the solidification time is applied at every time step. The scalar value equals the accumulated time that the VOF phase spent within the computational domain in a partly solidified state.

The value of the passive scalar is reset to zero when the phase remelts. That is, when the temperature increases above liquidus temperature after being in the interval between liquidus and solidus temperature for more than one iteration. Hence, the Solidification Time represents only the time that the phase spent between liquidus and solidus temperature for the last time. It is zero for every cell where the temperature is above liquidus temperature or the volume fraction of the phase is below the minimum volume fraction.

Large solidification times indicate hot spots in casting; these spots are commonly due to porosity defects. On the other hand, a very short solidification time for cast iron indicates the presence of white solidified areas in casting.

This quantity is used to compute the Local Solidification Time.

Local Solidification Time

The scalar quantity Local Solidification Time is collected for all specified Criteria Temperatures below the liquidus temperature. The Local Solidification Time is the value of the Solidification Time passive scalar in the cell at the moment the temperature decreases below the respective Criteria Temperature. The field is initialized with zero and values are modified in the iteration when the temperature falls below the respective Criteria Temperature.

The Local Solidification Time $t|_{T_{crit}}$ measures how long the fluid was in a partially solidified state before finally solidifying in the cell. Short local solidification times indicate a fine microstructure, and thus enhanced mechanical properties for aluminum and steel alloys.

This quantity is used for the computation of:

- (spatial) Local Solidification Time Gradient $\nabla t|_{T_{crit}}$ — not available as a field function but required for the computation of Isotherm Speed.
- Isotherm Speed
- Mean Cooling Rate—requires only Local Solidification Time for $T_{crit} = T_{\text{solidus}}$.
- G/v—requires Local Solidification Time Gradient.
- Secondary Dendrite Arm Spacing

Cooling Rate

The scalar quantity Cooling Rate is the time derivative of temperature. The cooling rate is computed for all specified Criteria Temperatures. Values are initially zero and are only overwritten each iteration when the temperature falls below the respective Criteria Temperature.

$$\dot{T}\Big|_{T_{crit}} = \frac{\partial T}{\partial t}\Big|_{T_{crit}} \tag{2501}$$

Analyzing the cooling rate in the temperature interval of interest can provide information on local microstructure properties, and thus on local mechanical properties.

For example, a high cooling rate in the solidification interval of an aluminum alloy indicates a fine microstructure: this fine microstructure corresponds to better tensile strength in these casting areas. On the other hand, cooling rates in the solidification interval close to the eutectic temperature for cast iron indicate that cementite can appear.

This quantity is used for the computation of:

- Solidification Velocity
- Niyama Criterion

Isotherm Speed

The scalar quantity Isotherm Speed is an approximation of the propagation velocity of the isotherm at the respective Criteria Temperature.

$$v_{iso}\Big|_{T_{crit}} = \frac{\nabla t\big|_{T_{crit}}}{\left|\nabla t\big|_{T_{crit}}\right|^2} \tag{2502}$$

It is computed for all specified Criteria Temperatures.

Temperature Gradient

The vector quantity Temperature Gradient $\nabla T\big|_{T_{crit}}$ is the spatial derivative of temperature. It is computed for all specified Criteria Temperatures.

Analyzing the temperature gradient in the temperature interval of interest can provide information on local microstructure properties, and thus on local mechanical properties.

For example, a high temperature gradient in the solidification interval of an aluminum alloy indicates a fine microstructure: this fine microstructure corresponds to good tensile strength in these casting areas. On the other hand, high temperature gradients in the solidification interval close to the eutectic temperature for cast iron indicate that cementite can appear.

This quantity is used for the computation of:

- Solidification Velocity
- G/v
- Niyama Criterion
- Primary Dendrite Arm Spacing

Mean Cooling Rate

The scalar quantity Mean Cooling Rate is the average cooling rate over the entire solidification process.

$$\dot{\bar{T}} = \frac{T_{\text{solidus}} - T_{\text{liquidus}}}{t\big|_{T_{\text{solidus}}}} \tag{2503}$$

Solidification Velocity

The vector quantity Solidification Velocity is an approximation for the velocity of the solidification front. This quantity is computed for all specified Criteria Temperatures.

$$v_s\big|_{T_{crit}} = \frac{\dot{T}\big|_{T_{crit}} \nabla T\big|_{T_{crit}}}{\left|\nabla T\big|_{T_{crit}}\right|^2} \tag{2504}$$

The solidification velocity is important for the analysis of directional solidification of single crystal castings. Inhomogeneity of the solidification velocity pattern can indicate areas in casting that contain crystals of different orientation.

This quantity is used for the computation of Primary Dendrite Arm Spacing.

G/v

The scalar quantity G/v is the ratio between the Temperature Gradient G and the Solidification Velocity v and is an important measure for the resulting micro structure. This quantity is approximately the scalar product

of the Temperature Gradient and the gradient of the Local Solidification Time, and is computed for all Criteria Temperatures below liquidus temperature:

$$G/v|_{T_{crit}} = \nabla T|_{T_{crit}} \cdot \nabla t|_{T_{crit}} \tag{2505}$$

This criteria function is used in the analysis of interface stability and is thus used to predict the solidification type: planar, cellular, or dendritic. High G/v ratios increase the probability for planar interface solidification, and lower G/v ratios leads to cellular or dendritic solidification.

Niyama Criterion

The scalar quantity Niyama Criterion estimates the porosity probability from the ratio between the shrinkage rate and the feeding distance. The shrinkage rate is assumed to be proportional to the Solidification Velocity and the feeding distance is assumed to be proportional to the Temperature Gradient. The criterion is computed for all Critical Temperatures.

$$K_{\text{Niyama}} = \frac{\nabla T|_{T_{crit}}}{\sqrt{-\dot{T}|_{T_{crit}}}} \tag{2506}$$

The Niyama Criterion predicts central porosity areas that appear as a result of melt contraction in the last stages of solidification. This criterion applies to low-carbon steel alloys. The critical value of this criterion, below which shrinkage defects form, is determined experimentally.

Dimensionless Niyama Criterion

You can use the dimensionless Niyama criterion to directly predict the amount of shrinkage porosity during solidification of metal alloy castings. This criterion does not require any threshold value below which shrinkage porosity forms. You can use this criterion for pure thermal simulations without computing pressure or velocity fields. The criterion is computed for all Critical Temperatures. The implementation is based on [433].

Solid Fraction Curve Integral

First, the fraction solid curve is integrated:

$$I(\alpha^*_{l,cr}) = \int_{\alpha^*_{l,cr}}^{1} 180 \cdot \frac{(1-\alpha^*_l)^2}{\alpha^{*2}_l} \cdot \frac{d}{d\alpha^*_l}(T^*) \cdot d\alpha^*_l \tag{2507}$$

where α^*_l is the relative liquid volume fraction. α^*_l is computed from the relative solid volume fraction α^*_s as follows:

$$\alpha_l^* = 1 - \alpha_s^*$$

$\alpha_{l,cr}^*$ is the critical liquid volume fraction at which the melt pressure decreases to the critical pressure value and porosity begins to form. T^* is the normalized temperature and is defined as:

$$T^* = \frac{T - T_{solidus}}{T_{liquidus} - T_{solidus}} \tag{2508}$$

The integral in Eqn. (2507) is evaluated numerically using at least the user-specified Min. Number of Integration Points (see Dimensionless Niyama Criterion Properties). The integral is calculated only once when the Dimensionless Niyama model is activated or when the Fraction Solid Curve material property changes.

Note: If the fraction solid curve is specified as a function of pressure, a pressure of 1 atm is applied for its evaluation. STAR-CCM+ assumes that the fraction solid curve is measured at this pressure.

Computation of the Dimensionless Niyama Criterion

At the end of every time step, the value of the Dimensionless Niyama criterion is computed as:

$$Ny^* = \lambda_2 Ny \sqrt{\frac{\Delta P_{cr}}{\mu_{liq} \beta \Delta T_f}} \tag{2509}$$

where

ΔP_{cr}	Critical pressure drop
μ_{liq}	Dynamic viscosity at liquidus temperature
β	Solidification contraction
$\Delta T_f = T_{liq} - T_{sol}$	Freezing temperature range
λ_2	Secondary Dendrite Arm Spacing from Eqn. (2516)

The critical pressure drop is defined as:

$$\Delta P_{cr} = P_{liq} - P_{cr} \tag{2510}$$

where P_{liq} is the pressure at liquidus temperature. P_{liq} is the sum of the ambient pressure of the system and the local head pressure. P_{cr} is the critical pressure and is calculated by considering the mechanical equilibrium that is necessary for a stable pore to exist. Here, P_{cr} is an adjustable parameter.

The solidification contraction is defined as:

$$\beta = \frac{\rho_{sol} - \rho_{liq}}{\rho_{liq}} \tag{2511}$$

The liquid and solid densities, ρ_{liq} and ρ_{sol}, are considered constant during solidification.

The original formulation of the Dimensionless Niyama criterion [433] uses the secondary dendrite arm spacing computed from the cooling rate (see Secondary Dendrite Arm Spacing (CR)).

Computation of the Solid Fraction Curve Integral Target Value

The target value of the solid fraction integral is given by:

$$I(\alpha^*_{l,cr}) = Ny^{*2} \tag{2512}$$

Computation of the Pore Volume Fraction

Knowing the target value of the solid fraction integral, the critical liquid fraction $\alpha^*_{l,cr}$ can be calculated. $\alpha^*_{l,cr}$ is the integration limit of Eqn. (2507).

Computation of the Abolute Pore Volume Fraction

The pore volume fraction is:

$$\alpha_p = \beta'\alpha^*_{l,cr} = \frac{\beta}{\beta + 1}\alpha^*_{l,cr} = \frac{\rho_{sol} - \rho_{liq}}{\rho_{sol}}\alpha^*_{l,cr} \tag{2513}$$

Dimensionless Niyama Criterion Properties

Dynamic Viscosity @ T_liquidus	The dynamic viscosity at liquidus temperature for the melt material.
Critical Pressure Drop	Defined as the difference between the pressure at liquidus temperature and the critical pressure (see Eqn. (2510)).
Solidification Contraction	β in Eqn. (2511).
Min. Number of Integration Points	The minimum number of integration points used to numerically integrate Eqn. (2507).

Primary Dendrite Arm Spacing

The scalar quantity primary dendrite Arm spacing (also referred to as PDAS or DAS) is the distance between the axes of the dendrites. This quantity is computed only for liquidus temperature. It is possible to calculate the primary dendrite arm spacing from the Temperature Gradient and the Solidification Velocity. This calculation uses the material property Primary Dendrite Arm Spacing Coefficient K_1 and the field function properties Solidification Velocity Exponent a and Temperature Gradient Exponent b.

$$\lambda_1 = K_1\left(\left|v_s\right|_{T_{\text{liquidus}}}\right)^a\left(\left|\nabla T\right|_{T_{\text{liquidus}}}\right)^b \tag{2514}$$

Since both exponents are derived from theory, they are field function properties and not material properties. By contrast, the exponent that is used for the computation of the Secondary Dendrite Arm Spacing is a material property.



The primary dendrite arm spacing is used to estimate the local mechanical properties in columnar solidification. The finer the primary dendrite arm spacing is, the higher the mechanical properties.

Primary Dendrite Arm Spacing Properties

Coefficient	K_1 in Eqn. (2514).
Solidification Velocity Exponent	a in Eqn. (2514).
Temperature Gradient Exponent	b in Eqn. (2514).

Secondary Dendrite Arm Spacing (LST)

The scalar quantity Secondary Dendrite Arm Spacing (LST) (also referred to as SDAS-LST) is the distance between the arms of dendrites. This quantity is computed for all specified Criteria Temperatures. The secondary dendrite

arm spacing (LST) can be estimated from the Local Solidification Time using the material properties Secondary Dendrite Arm Spacing Coefficient K_2 and Secondary Dendrite Arm Spacing Exponent c.

$$\lambda_{2,\,LST}\Big|_{T_{crit}} = K_{2,\,LST}(t_{T_{crit}})^c \tag{2515}$$

The secondary dendrite arm spacing is used to estimate local mechanical properties in both columnar and equiaxed solidification. The finer the secondary dendrite arm spacing is, the higher the mechanical properties.

Secondary Dendrite Arm Spacing (LST) Properties

Coefficient	$K_{2,\,LST}$ in Eqn. (2515).
Exponent	c in Eqn. (2515).

Secondary Dendrite Arm Spacing (CR)

The scalar quantity Secondary Dendrite Arm Spacing (CR) (also referred to as SDAS-CR) is the distance between the arms of dendrites. This quantity is computed for all specified Criteria Temperatures. The secondary dendrite arm spacing (CR) can be estimated from the Cooling Rate using the material properties Secondary Dendrite Arm Spacing Coefficient K_2 and Secondary Dendrite Arm Spacing Exponent c.

$$\lambda_{2,\,CR}\Big|_{T_{crit}} = K_{2,\,CR}(\dot{T}\big|_{T_{crit}})^c \tag{2516}$$

Secondary Dendrite Arm Spacing (CR) Properties

Coefficient	$K_{2,\,CR}$ in Eqn. (2516).
Exponent	c in Eqn. (2516).

Artificial Compressibility

Typically, for casting simulations, the gas phase is considered to be compressible and the ideal gas law is applied. The melt is a liquid and has a constant density. Therefore, the multiphase mixture is not compressible

throughout the entire computational domain. This inhomogeneity results in regions with finite and infinite speeds of sound, which can lead to numerical instabilities and convergence problems.

To overcome such robustness issues caused by combining a constant density VOF melt phase with an ideal gas VOF gas phase, the artificial compressibility method can be applied. The artificial compressibility for the VOF melt phase ensures a compressible mixture in the entire fluid domain.

There are two options for defining the artificial compressibility of the melt:

- Pressure Dependent Density
- Pressure and Temperature Dependent Density

Pressure Dependent Density

If the melt density does not depend on temperature, such as for isothermal simulations, the density is defined as:

$$\rho_{\mathrm{Melt}} = \max\left(\rho_{\mathrm{Melt}_{min}}, \rho_{\mathrm{Melt}_{ref}} + \frac{d}{dp}(\rho_{\mathrm{Melt}}) \cdot p\right) \tag{2517}$$

where:

Symbol	Description	Value
ρ_{Melt}	Density of the melt	
$\rho_{\mathrm{Melt}_{min}}$	Minimum melt density	$1 kg/m^3$
$\rho_{\mathrm{Melt}_{ref}}$	Reference melt density that is linearly interpolated from the temperature-dependent density table of the Casting Material Database at a temperature that closely matches the Casting Temperature.	
$\frac{d}{dp}(\rho_{\mathrm{Melt}})$	Pressure derivative of the melt density	$10^{-5} Pa$
p	Relative pressure	

and

$$\frac{\partial}{\partial T}(\rho_{\mathrm{Melt}}) = 0$$

Pressure and Temperature Dependent Density

If the melt density depends on both pressure and temperature, the density is given as:

$$\rho_{\mathrm{Melt}} = \max\left(\rho_{\mathrm{Melt}_{min}}, \rho_{\mathrm{Melt}}(T) + \frac{d}{dp}(\rho_{\mathrm{Melt}}) \cdot p\right) \tag{2518}$$

where:

Symbol	Description	Value
ρ_{Melt}	Density of the melt	
$\rho_{\mathrm{Melt}_{min}}$	Minimum melt density	$1 kg/m^3$
$\rho_{\mathrm{Melt}}(T)$	Temperature-dependent density of the melt that is linearly interpolated from the Casting Material database.	
$\frac{d}{dp}(\rho_{\mathrm{Melt}})$	Pressure derivative of the melt density	$10^{-5} Pa$
p	Relative pressure	

and

$$\frac{\partial}{\partial T}(\rho_{\mathrm{Melt}}) \neq 0$$

is computed from the tabular melt density values of the Casting Material Database.

Automated Time-Step Control

To adjust the time-step automatically during the run, you can use the Segregated Time Step model. The Segregated Time Step model improves the quality of your results by automatically setting an appropriate time-step which reduces the temporal discretization error. This adjustment is particularly useful for cases with large variations of flow topology or large varying time scales of the physics.

To control the automatic time-step process, you create one or more time-step controls that act independently. Within each time-step control, you define the cells to which the control applies. You also define a CFL condition that must be met. There are two CFL conditions available:

- flow-based convective Courant number
- energy-based thermal diffusion

The automatic time-step control collects time-step proposals from all active time-step controls for each active CFL condition. The segregated time step model chooses the global minimum time-step size from the set.

The time-step size is evaluated as follows:

- If the specified CFL condition is met, the time-step size proposal is set to a value that is larger than the previous time-step size.
- If the CFL condition is exceeded, the time-step size proposal is set to a value that is smaller than the previous time-step size.

To avoid large jumps in time-step size, specify a maximum time-step change factor. The maximum time-step change factor is defined as the ratio of the current targeted time-step to the previous time-step. If the maximum time-step change factor is exceeded, the targeted time-step size is cut off at the maximum value permitted by the change factor.

You specify a physical start time for each time-step control. The time-step control remains inactive until the start time is reached.

You also specify a physical end time for each time-step control. The time-step control becomes inactive from the specified time onwards. The time-step control deviates from its original time-step size proposal in favor of reaching the specified end time in a certain number of time-steps. This procedure ensures that the specified end time is reached.

Setting up the Segregated Time Step Model

The Segregated Time Step model is only available when the Casting model is selected.

See Selecting the Casting Model

To select the Segregated Time Step model:

1. Right-click the Continua > Physics 1 > Models node and choose Select models...

   The *Model Selection* dialog appears.
2. In the *Optional Models* group box, activate Segregated Time Step.
3. Click Close.

The Segregated Time Step node is added to the simulation tree.



4. Set the Segregated Time Step Properties.
5. Right-click the Segregated Time Step > Time Step Controls node and select New.

A Time Step Control 1 node is added to the simulation tree.



6. Select the Time Step Control 1 node and set the Time Step Control Properties.
7. Depending on the activated CFL conditions, select each CFL condition child node in turn and set their CFL Condition Properties.
8. Select the Solvers > Implicit Unsteady node and set Time-Step to a value that corresponds to the maximum time-step.

   This value is the upper time-step size limit for the automatic time-step computation.

Segregated Time Step Properties

Verbose	Additional information appears in the output window during the simulation.

Time Step Control Properties

Start Time	Specifies the time at which the time-step control is included in the automatic time-step control.
End Time	Specifies the time after which the time-step control is excluded from the automatic time-step control. STAR-CCM+ adjusts time-steps automatically to reach this goal. End Time is only active if the specified value is larger than the specified Start Time.
Start Time Step	Size of the first proposed time-step size at Start Time.
Min. Time Step	Lower limit on time-step size. Exceptionally, the time-step size can be reduced below this value in order to meet the End Time target.
Max. Time Step Change Factor	Limits the proposed time-step size if the ratio of the current proposal to the previous time-step exceeds this value.
Flow-Based CFL Condition	Activates or deactivates the CFL condition based on the convective Courant number. This condition is not available for time-step controls within a solid continuum. When activated, a Flow-Based CFL Condition node appears in the simulation tree.
Energy-Based CFL Condition	Activates or deactivates the CFL condition based on thermal diffusion. This condition is only available if the Segregated Energy model is enabled for the continuum. When activated, a Energy-Based CFL Condition node appears in the simulation tree.

CFL Condition Properties

Target Mean CFL Number	Specifies the upper limit of the calculated weighted average CFL number that is used to compute a proposed time-step size.
Max. Allowed CFL Number	Specifies the upper limit of the calculated maximum CFL number that is used to compute a proposed time-step size.

Combined Thermal Contact Resistance

When meld solidifies during a casting process, the density of the melt increases with decreasing temperature. Due to mass conservation, the melt volume must shrink. This thermal shrinkage occurs both in the melt and in the surrounding mold. However, as shrinkage does not occur to an equal extent in the melt and in the solid, gaps form at the melt/mold interface. This gap significantly reduces heat transfer rates. To account for this effect, you can apply the combined thermal contact resistance at the interface.

The combined thermal resistance is given by:

$$R_{tot} = \frac{\left|T_{Melt} - T_{Mold}\right|}{\left|\dot{q}_{tot}\right|} \tag{2519}$$

where T_{Melt} is the melt temperature at the melt/mold interface, T_{Mold} is the mold temperature at the melt/mold interface, and $\dot{q}_{tot}$ is the total heat flux.

The total heat flux consists of a conduction and a radiation component:

$$\dot{q}_{tot} = \dot{q}_{cond} + \dot{q}_{rad} \tag{2520}$$

where $\dot{q}_{cond}$ is the conduction heat flux and $\dot{q}_{rad}$ is the radiation heat flux.

The conductive heat flux across the interface is computed as:

$$\dot{q}_{cond} = \alpha_{cond} \cdot (T_{Melt} - T_{Mold}) = \frac{T_{Melt} - T_{Mold}}{R_{con}} \tag{2521}$$

where α_{cond} is the conductive heat transfer coefficient at the melt/mold interface and R_{con} is the conductive thermal resistance at the melt/mold interface.

The radiative heat flux is given by the Stefan-Boltzmann law:

$$\dot{q}_{rad} = \sigma(\varepsilon_{Melt} T_{Melt}^4 - \varepsilon_{Mold} T_{Mold}^4) \tag{2522}$$

where $\sigma = 5.6704 \cdot 10^{-8}$ is the Stefan-Boltzmann constant, ε_{Melt} and ε_{Mold} are the melt and the mold emissivities, respectively. These emissivities are user-specified values (see Interfaces Properties).

Corrective Energy Source for Pure Thermal Simulations

If you run a pure thermal casting simulation with a temperature-dependent density, there is an unsteady contribution to the energy equation due to the change in density. The purpose of the corrective energy source is to undo this unsteady contribution to the energy equation. Activating the corrective energy source results in a temperature field that is equivalent to the solution you would obtain had the density remained constant over time.

In the first iteration in which the corrective energy source is activated, the density field of the fluid or the solid region is stored for reference. At the end of each time-step, the unsteady term due to the density change during that time step is evaluated and the contribution to the source and the linearization is applied with a negative sign. For liquid regions, also the contribution to mass imbalance is reverted.

- First order in time

  The source term due to density change using the solution at the current time level $n+1$ and from the previous time level n is given by:

$$S = \frac{(\rho VE)^n - (\rho VE)^{n+1}}{\Delta t} \tag{2523}$$

where:

S	Source term to the energy equation
ρ	Density
V	Volume
E	$(dh)/(dT)$ Derivative of enthalpy with respect to temperature
Δt	Time step size

The desired value for the source term is:

$$S = \frac{\rho_{ref}(VE)^n - \rho_{ref}(VE)^{n+1}}{\Delta t} \tag{2524}$$

Therefore, the corrective energy source term is:

$$S_{corr} = \frac{(\rho_{ref} - \rho^n) \cdot (VE)^n - (\rho_{ref} - \rho^{n+1}) \cdot (VE)^{n+1}}{\Delta t} \tag{2525}$$

- Second-order in time

  The source term due to density change is given by:

$$S = \frac{(\alpha^2 - 1)[(\rho VE)^n - (\rho VE)^{n+1}] + [(\rho VE^n) - (\rho VE)^{n-1}]}{(\alpha - 1)\alpha \cdot \Delta t^{n+1}} \tag{2526}$$

where:

$$\alpha = 1 + \frac{\Delta t^n}{\Delta t^{n+1}}$$

$$\Delta t^{n+1} = t^{n+1} - t^n$$
$$\Delta t^n = t^n - t^{n-1}$$

The desired value is:

$$S = \frac{(\alpha^2 - 1)[\rho_{ref}(VE)^n - \rho_{ref}(V_n E_n)^{n+1}] + [\rho_{ref}(VE)^n - \rho_{ref}(VE)^{n-1}]}{(\alpha - 1)\alpha \cdot \Delta t^{n+1}} \tag{2527}$$

Therefore, the corrective energy source term is:

$$S_{corr} = \frac{(VE)^n[\rho_{ref} \cdot \alpha^2 - \rho^n(2 - \alpha^2)] + (\alpha^2 - 1)(VE)^{n+1}(\rho^{n+1} - \rho_{ref}) + (VE)^{n-1}(\rho^{n-1} - \rho_{ref})}{(\alpha - 1)\alpha \cdot \Delta t^{n+1}} \tag{2528}$$

The unsteady term of the energy equation for fluid regions is given in Segregated Fluid Energy Formulation.

The source terms for fluid regions are equal to those for solid regions except for the contribution to mass imbalance:

$$S = S_{solid} - \Delta m \cdot H_t \tag{2529}$$

where H_t is the total enthalpy and Δm is the mass imbalance.

Activating the Corrective Energy Source for Pure Thermal Simulations

You activate the corrective energy source only for pure thermal simulations. The corrective energy source revokes the unsteady contribution to the energy equation due to the change in density.

To activate the corrective energy source:

- Select the Continua > [Continuum] > Models > Casting node and activate Apply Corrective Energy Source.



Even if activated, the corrective energy source term is not applied in case any of the following solvers is active:

- Pressure solver
- Velocity solver
- VOF solver

Cases using these solvers do not require correction as density variations are already compensated by fluxes.

The Casting User Interface

The Casting User Interface (CUI) is a dedicated user interface specifically designed to guide you through the setup of casting simulations. It is intuitive to use and easy to learn. The Casting User Interface lets you set up:

- Physical models
- Material models

- Initial conditions
- Boundary conditions
- Interfaces
- Visualization and reports

The currently supported casting processes are:

- Investment casting
- Sand casting
- Permanent die casting

When you have prepared the geometry and completed the meshing, you can start using the CUI. When you use the CUI to set up the casting simulation, the STAR-CCM+ simulation tree is populated and updated in the background. Making a selection or setting in the CUI can result in one or more actions in the simulation tree.

It is strongly recommended that you perform the entire casting setup in the CUI. Do not toggle between the simulation tree and the CUI. The CUI updates settings in the simulation tree. However, selections and entries from the simulation tree are not transferred back to the CUI.

Current Limitations

- The table editor only supports a maximum of 20 rows.
- It is not possible to change the units of the input values for tables.
- Variable values in the form of field functions are not available.

See also:

Accessing the Casting User Interface

Saving and Restoring the Casting User Interface Settings

Casting User Interface Reference

Accessing the Casting User Interface

To make the Casting User Interface (CUI) accessible in STAR-CCM+:

1. In the STAR-CCM+ menu, select Window > Toolbars > Customize.



2. Create a new toolbar to contain the CUI icon.
   a. In the *Customize Toolbars* dialog, click New Toolbar...
   b. In the *New Toolbar* dialog, enter an appropriate name for the new toolbar, for example, `Casting`.
   c. Click OK.

      The new toolbar appears on the right next to the existing toolbars at the top of the STAR-CCM+ workspace. This can take a few seconds.

3. In the *Customize Toolbars* dialog, under the Wizards node, select STAR-Cast User Interface.



4. Drag the STAR-Cast User Interface icon onto the newly created casting toolbar at the top of the STAR-CCM+ workspace and drop it at the appropriate position.

   The icon is added to the toolbar at the specified location.

5. In the *Customize Toolbars* dialog, click Close.

To open the Casting User Interface:

- Click the STAR-Cast User Interface icon in the *Casting* toolbar.

The CUI opens with the *Casting Processes and Simulation Type* panel as an embedded panel in the *simulation* window covering the simulation tree.



To detach the CUI from the simulation tree, click the Detach button. This docks the CUI to the right-hand side of the STAR-CCM+ workspace. Once the CUI is detached, an Attach button appears that allows you to re-attach the CUI to the simulation tree. Use the Close button to close the CUI.

Saving and Restoring the Casting User Interface Settings

After you have made model selections and settings in the CUI and have saved the simulation file (see Saving Simulations to File), an XML file is automatically saved to the local working directory where the simulation file resides. The naming convention of the XML file is `<simfilename>.cast.xml` where `simfilename` is replaced with the name of the sim file. The XML file contains all of the settings that you have specified in the CUI.

The next time that you open the simulation file `<simfilename>.sim` and start the CUI, the existing XML file `<simfilename>.cast.xml` is read. The CUI opens with the previously stored settings.

Casting User Interface Reference

The Casting User Interface consists of eleven different panels. Each panel covers a step in the setup of a casting simulation, as follows:

1. The Condition Settings Import/Export Panel
2. The Casting Processes and Simulation Type Panel
3. The Material Data Panel
4. The Regions Panel
5. The Boundaries Panel
6. The Interfaces Panel
7. The Process Parameters Panel
8. The Physical Models Panel
9. The Solvers Panel
10. The Criteria Functions Panel
11. The Post Processing Panel

You walk through the panels from top to bottom. The order is important because the CUI updates and adjusts the possible settings depending on previous selections. This behavior is reflected by the look of the icons and tabs: if a required selection was not made in a previous panel, the icons and/or the tabs remain grayed out.

In the panels, you enter the settings required for the casting simulation. For most panels, the procedure is:

1. Specify the conditions for a particular setup step, for example, the boundaries, in the Settings tab.

A condition is a set of properties that can be saved under a specific name.

2. Assign the saved condition to the item to be defined in the Selections tab, for example, to the boundary surface.

The advantage of this two-step approach is that different conditions can be saved for later use and can easily be exchanged in the Selections tab.

The Condition Settings Import/Export Panel

In the *Condition Settings Import/Export panel*, you can store or retrieve condition settings that you make in the CUI. These settings are stored in the cast.xml file. Use the *Condition Filter* section to select either all or certain conditions for export or import.

- To import condition settings from a cast.xml file, click Import.
- To export condition settings to a cast.xml file, click Export.

Condition Filter	The following condition settings can be selected for export or import: • Regions • Interfaces • Physical Models • Stopping Criteria • Boundaries • Process Parameters • Solver Settings
Conditions	Displays all the conditions that were either defined manually during the current setup and/or conditions that were imported from the cast.xml file.

The Casting Processes and Simulation Type Panel

In order to use the Casting User Interface, you must make the first three selections:

- Casting Process
- Process Type
- Simulation Type

The selections that you make here affect the options that are available in the following panels. Not all of the Process Type options listed in the following table are available for all Casting Process options.

Casting Process	Selects the casting process to be used.	
	Investment	Sets up an investment casting process.
	Sand	Sets up a sand casting process.
	Permanent Die	Sets up a permanent die casting process.
Process Type	Selects the process type to be used for the previously selected casting process. Note that not all process type options are available for each casting process. Which process type option becomes available depends on the selection that you have made under Casting Process. This option invokes the creation of a physics continuum for the melt in the simulation tree. Various physics models are activated, including: • Casting • VOF Multiphase The Cell Quality Remediation model is also activated. Selects the implicit unsteady segregated solver for a transient simulation. The flow is laminar for all process types.	
	Bridgman	Sets up a Bridgman type casting process. Activates a translational motion model.
	Centrifugal	Activates a rotational motion model.
	Gravity	Sets up a gravity type casting process.
	Tilt	Activates a rotational motion model.
Simulation Type	Selects the simulation type to be used for the casting process.	
	Filling and Solidification	Sets up a simulation including the pouring and the solidification process of the melt. Activates the segregated energy solver.

Isothermal Filling	The solidification process of the melt is not considered. Energy is not solved for. Note that the Isothermal Filling option is not available for the Bridgman process type option.
Pure Thermal	Sets up a thermal solidification analysis of the melt. Activates the segregated energy solver.

Gravity

Specification of the gravity vector.

Reference Pressure

Specification of the reference pressure value.

Casting Temperature

- For the Simulation Type Isothermal Filling, the Casting Temperature sets the continuum temperature of the Melt continuum. The simulation is performed at this temperature level. This means that the thermodynamic properties of the melt and the gas at this temperature level are applied. To obtain the constant material properties for the melt at this temperature, the thermodynamic properties provided in the Casting Material Database are evaluated.
- For the Simulation Type Filling and Solidification, the Casting Temperature is the default inflow boundary temperature of the melt for the boundary types Velocity Inlet Boundary, Mass flow Inlet Boundary, and Pressure Outlet Boundary. This value can be overwritten by respective entries in the Boundaries tab.
- For the Simulation Type Pure Thermal, the Casting Temperature is the default initial temperature of the melt. This value can be overwritten by respective entries in the Region tab.

Ambient Temperature

Effective radiation temperature of the environment. This temperature is needed for the gray thermal radiation model. If Simulation Type is Isothermal Filling, this option is deactivated.

The Material Data Panel

In the *Material Data* panel you select the materials used for the casting simulation. The selection encompasses materials for the melt as well as materials for the solid molds.

Use the ACCESS casting material database to select the materials.

To select casting materials:

- Click Launch Database...

  This opens the casting material database. For more information, see Working With the Casting Material Database.

The selected materials are displayed in the scroll list of the CUI in the Material Data Panel. The selected materials also automatically go to the Tools > Material Databases node of the simulation tree. A new material database called Selected_Casting_Material.mdb is created under that node.

You can import previously selected materials data from an existing casting material XML file, which has been created by the ACCESS casting material database.

To import an existing casting material database XML file:

- Click Import XML... and select the file that you would like to import.

The Regions Panel

In the *Regions* panel, you set the initial conditions of the simulation and assign the previously selected materials to the regions. The regions exist already as they are created during the geometry preparation and the mesh generation process. The *Regions* panel consists of two tabs: the Selections tab and the Settings tab. You assign the initial conditions that you define in the Settings tab to the existing regions in the Selections tab.

Regions Selections Tab

Regions

Contains the regions that have been created during the geometry preparation and the meshing process.

Materials

Selection of the previously selected materials from the casting material database.

Selecting a liquid casting material invokes:

- Creation of the corresponding liquid material within the Eulerian phase of the melt continuum previously created.
- Assigning of melt continuum to the selected region in the table.
- Activation of Density Limiting for the ideal gas model of the Eulerian gas phase. The Cut-off pressure is set to a value of 20% of the specified reference pressure. For more information, see Ideal Gas Model Expert Properties.
- Activation of Artificial Compressibility for the melt phase. For more information, see Artificial Compressibility. The type of artificial compressibility depends on the Simulation Type that is set in the *Casting Processes and Simulation Type* panel. For Isothermal Filling, Pressure Dependent Density is applied. For the Pure Thermal and Filling and Solidification, Pressure and Temperature Dependent Density is applied.
- Activation of the Surface Tension model (see Surface Tension Formulation).
- Activation of the Interface Momentum Dissipation model (see Using the Interface Momentum Dissipation Model).

Selecting a solid mold material invokes:

- Creation of a solid physics continuum containing the solid mold material.
- Grey thermal radiation model is activated. Assigning of solid mold continuum to the selected region in the table.

Conditions — Assigns the initial conditions to the selected region. The available conditions must be set before in the Settings tab. Depending on the Type selection that you have made in the Settings tab, only a suitable selection of either a fluid or a solid initial condition is available.

Regions Settings Properties

In the Settings tab, you define the initial conditions for either the fluid or the solid regions and save them under a user-defined name. These initial conditions can then be assigned to a region in the Regions table in the Selections tab.

Initial Conditions — Lists the initial conditions that are already defined.

Name — User-defined name under which the initial condition is saved.

Type — Specifies whether the initial condition is related to a fluid or a solid region.

Properties — Specifies the variable for which the initial condition is to be set.

- Temperature — Sets the initial temperature. The units are C for degrees Celsius, K for kelvins, and F for Fahrenheit.
- Velocity — Sets the initial velocity components with respect to the laboratory coordinate system in m/s.
- Volume Fraction — Sets the initial volume fraction. This is a dimensionless value between 0 and 1.

Options — The properties can be specified in different formats.

- Constant — Sets a constant value of the property. You specify this value in the corresponding Values cell of the Properties table.
- Table (Time) — Sets the property value in the form of a table as a function of time.

Table (XYZ)	Sets the property value in the form of a table as a function of coordinates with respect to the laboratory coordinate system.

Table Options

If either Table (Time) or Table (XYZ) has been selected, the Values cell is greyed out and a table editor appears. You can use the units selector to set the units of the variable. For Temperature, the table values adjust when the unit changes. The table can be input in two different ways.

Tabulate	Enter the tabular values manually in the table editor.
Import .csv	Import the tabular values from a file, which contains comma separated values.
Export .csv	The tabular values can be exported to a file containing comma separated values.

Note: If you selected Isothermal Filling as Simulation Type in the *Casting Processes and Simulation Type* panel, only the Temperature option set to Constant is valid. The other two options Table (Time) and Table (XYZ) can be selected, but they have no effect on the simulation.

The Boundaries Panel

In the *Boundaries* panel, you specify the boundary conditions. The panel consists of two tabs: the Selections tab and the Settings tab. First, you define the boundary conditions in the Settings tab. Then, you can assign these conditions to the actual boundary surfaces in the model.

Boundaries Selections Tab

Boundaries	Contains the boundary surfaces that have been created during the meshing process.
Conditions	Assigns the boundary conditions to the selected boundary surface. The available conditions must have been set before in the Settings tab. Depending on the Type selection that you have made in the Settings tab, only a suitable selection of either a fluid or a solid boundary condition is available. After a selection has been made, the corresponding properties are shown at the bottom of the panel.
Type	Specifies the type of boundary condition to use. There are different boundary condition types available depending on whether it is a fluid or a solid boundary.

Boundaries Settings Tab

Boundary Conditions	Lists the boundary conditions that are already defined.
Name	User-defined name under which the boundary condition is saved.
Type	Specifies the type of boundary condition to be used. There are different boundary condition types available depending on whether it is a fluid or a solid boundary. Apply boundary conditions termed Fluid to boundary surfaces encompassing a fluid. Apply boundary conditions termed Solid to boundary surfaces encompassing a solid. The boundary conditions have different properties

Adiabatic Fluid Wall Boundary	Boundary, which is impenetrable for fluid flow and energy.
Adiabatic Solid Wall Boundary	Boundary, which is impenetrable for energy.
Fixed Temperature Fluid Wall Boundary	Boundary, on which a fixed temperature is specified.
Fixed Temperature Solid Wall Boundary	Boundary, on which a fixed temperature is specified.
Fluid Wall Boundary	Convective boundary, for which an ambient temperature outside of the fluid domain is specified. The heat flux through the wall and the temperature at the wall are computed by use of the specified heat transfer coefficient (W/m^2 s). The wall porosity model can be applied at this boundary type.
Heat Flux to Fluid Wall Boundary	Boundary, on which a heat flux in W/m^2 is specified. The wall porosity model can be applied at this boundary type.
Heat Flux to Solid Wall Boundary	Boundary, on which a heat flux in W/m^2 is specified. The wall porosity model can be applied at this boundary type.

Mass Flow Inlet Boundary	Boundary, for which the inflow mass flow rate (kg/s) is specified. The inflow temperature and the volume fraction (dimensionless values between 0 and 1) are also specified.
Pressure Outlet Boundary	Outflow boundary for which the pressure is specified. Specified temperature and volume fraction are applied under inflow conditions, that is when recirculation occurs.
Solid Wall Boundary	Convective boundary, for which an ambient temperature outside of the solid domain is specified. The heat flux through the wall and the temperature at the wall are computed by use of the specified heat transfer coefficient (W/m^2 s). The emitted radiation can be specified through the emissivity parameter (dimensionless value between 0 and 1).
Symmetry Plane Boundary	This boundary condition can be used at symmetry planes of geometries. It has no properties.
Velocity Inlet Boundary	Specify the normal inlet velocity components (m/s).

Boundary Conditions Properties

Properties

Contact Angle	The contact angle in degrees between the wall and the melt. For details on the definition of the contact angle, see What Is a Contact Angle?.
Porous Inertia Coefficient	Coefficient of the wall porosity model. It corresponds to the α coefficient in Eqn. (2090). See Wall Porosity Formulation for details on the definition of the porous inertia coefficient.
Porous Viscous Coefficient	Coefficient of the wall porosity model (m/s). It corresponds to the β coefficient in Eqn. (2090). See Wall Porosity Formulation for details on the definition of the porous inertia coefficient.
Slip Condition	This property has two options: • No. A no-slip wall is applied. The fluid velocity at the wall equals the wall velocity. The wall can be at rest or you can specify a tangential wall velocity in the form of a velocity vector with respect to the laboratory coordinate system. • Yes. A slip wall is applied. The fluid has a different velocity than the wall velocity.
Temperature	Temperature at the boundary.
Emissivity	Ratio of the power that a body emits to the power it would emit as a black body at the same temperature (dimensionless). Emissivity is a parameter of the grey thermal radiation model.
Ambient Temperature	Assumed temperature outside of the computational domain.
Heat Transfer Coefficient	Heat transfer coefficient at the wall (W/m^2s) used to compute the temperature at the wall.
Heat Flux	Heat flux at the boundary (W/m^2)
Mass Flow Rate	Fluid mass flow rate (kg/s).

Volume Fraction	Volume fraction of a Lagrangian phase is the fraction of the local cell volume which that phase occupies. It takes values between 0 and 1. A value of 0 means that nothing of that phase is contained in the cell whereas a value of 1 means that the cell is completely filled with the phase. Specify the volume fraction in the form of a vector where the first entry stands for the melt, the second entry stands for the gas phase.

Options

The properties can be specified in different formats.

Constant	Sets a constant value of the property. Specify this value in the corresponding Values cell of the Properties table.
Table (Time)	Sets the property value in the form of a table as a function of time.
Table (XYZ)	Sets the property value in the form of a table as a function of coordinates with respect to the laboratory coordinate system.
Material DBS	Takes the property value from the Casting Material Database.

Table Options

If either Table (Time) or Table (XYZ) is selected, the Values cell is greyed out and a table editor appears. You can use the units selector to set the units of the variable. For Temperature, the table values adjust when the unit changes. The table can be input in two different ways.

Tabulate	Enter the tabular values manually in the table editor.
Import .csv	Import the tabular values from a file, which contains comma separated values.
Export .csv	The tabular values can be exported to a file containing comma separated values.

The Interfaces Panel

In the Interfaces panel, you specify the interfaces between the fluid and the solid regions. At the interfaces, variable values between fluid and solid regions are exchanged in the course of the iterations of the solvers. First, you define the interface conditions in the Settings tab. Then, you assign these conditions to the actual interface surfaces in the model in the Selections tab. The properties of the assigned conditions are then shown in the Properties table of the Selections tab.

Note that the settings made here in the *Interfaces* panel of the CUI, invoke simulation tree settings both on the interface nodes on the regions level as well as on the Interface node itself.

Interfaces Selections Tab

Interfaces	Contains the interfaces that have been created during the meshing process.
Conditions	Assigns the interface conditions to the selected interface. The available conditions must have been set before in the Settings tab. Depending on the Type selection that you have made in the Settings tab, only a suitable selection of the interface types is available. For example, at a fluid-solid interface you are only able to select an interface condition of the type Fluid to Solid. After a selection has been made, the corresponding properties are shown at the bottom of the panel.
Type	Specifies the type of interface to be used. There are different interface condition types available depending on whether it is a solid-fluid interface, a solid-solid interface or a flow-through (internal) interface.

Interfaces Settings Tab

Conditions — Lists the interface conditions that are already defined.

Name — User-defined name under which the interface condition is saved.

Type — Specifies the type of interface condition to use. There are different interface condition types available. The interface conditions have different properties.

- Fluid to Solid — Defines a solid-fluid interface.
- Solid to Solid — Defines a solid-solid interface.
- Flow Through — Defines an internal interface that permits fluid and energy flow through it.

Interfaces Properties

Properties

Contact Angle	The contact angle in degrees between the interface and the melt. For details on the definition of the contact angle, see What Is a Contact Angle?.
Porous Inertia Coefficient	The coefficient of the wall porosity model. It corresponds to the α coefficient in Eqn. (2090). See Wall Porosity Formulation for details on the definition of the porous inertia coefficient.
Porous Viscous Coefficient	The coefficient of the wall porosity model (m/s). It corresponds to the β coefficient in Eqn. (2090). See Wall Porosity Formulation for details on the definition of the porous inertia coefficient.
Slip Condition	This property has two options: • No. A no-slip condition is applied at the interface. The fluid velocity at the interface equals the interface velocity. The interface can be at rest or you can specify a tangential interface velocity in the form of a velocity vector with respect to the laboratory coordinate system. • Yes. A slip condition is applied at the interface. The fluid has a different velocity than the interface velocity.
Ambient Temperature	Ambient temperature for use with convection condition thermal specification.

Heat Transfer Coefficient	Heat transfer coefficient at the interface (W/m^2K) used to compute the temperature at the interface. In the simulation tree, this value is transformed into a contact resistance, which is the inverse of the heat transfer coefficient in m^2K/W. When the option Table (Temp.) is selected, a Combined Thermal Contact Resistance consisting of conduction and radiation is applied. Radiation is only accounted for if values greater than zero are specified for Fluid Emissivity and Solid Emissivity. For more information, see Combined Thermal Contact Resistance.
Fluid Emissivity	Radiation emissivity of the solidified melt (dimensionless).
Solid Emissivity	Radiation emissivity of the solid mold.

Options

The properties can be specified in different formats.

Constant	Sets a constant value of the property. Specify this value in the corresponding Values cell of the Properties table.
Table (Time)	Sets the property value in the form of a table as a function of time.
Table (XYZ)	Sets the property value in the form of a table as a function of coordinates with respect to the laboratory coordinate system.
Table (Temp.)	Set the property value in the form of a table as a function of temperature.
Material DBS	Takes the property value from the Casting Material Database.

Table Options

If either Table (Time) or Table (XYZ) has been selected, the Values cell is greyed out and a table editor appears. You can use the units selector to set the units of the variable. For Temperature, the table values adjust when the unit changes. The table can be input in two different ways.

Tabulate	Enter the tabular values manually in the table editor.
Import .csv	Import the tabular values from a file, which contains comma separated values.
Export .csv	The tabular values can be exported to a file containing comma separated values.

The Process Parameters Panel

In the *Process Parameters* panel, you specify process parameters, which are specific to the process type previously selected in the *Casting Process and Simulation Type* panel. The panel consists of two tabs: the Selections tab and the Settings tab. First, you define the process parameters in the Settings tab. Then, you assign the process parameter conditions to the selected regions in the Selections tab.

Process Parameters Selections Tab

Depending on the Process Type selection made in the *Casting Process and Simulation Type* panel, the tab appears differently. If the Process Type is Gravity, the Selections tab is empty and no assignments are required. If the Process Type is either Bridgman, Centrifugal, or Tilt, a region selection dialog appears that allows selection of the regions to which the process parameters are applied.

The Conditions selector allows selection of the process parameter conditions that have been previously defined in the Settings tab. The assigned properties are shown at the bottom of the panel.

Process Parameters Settings Tab

Conditions
: Lists the process parameter conditions that are already defined.

Name
: User-defined name under which the process parameter condition is saved.

Type
: Specifies the type of process parameter condition to be used. There are different process parameter condition types available. The process parameter conditions have different properties

  Bridgman Casting
  : Directed solidification and single crystal casting process where the cast part moves out of an oven. Activates a translational motion model.

  Tilt Casting
  : To fill the mold, the entire apparatus is tilted according to a tilting curve. Activates a rotational motion model.

  Centrifugal Casting
  : To fill the mold, the entire apparatus or the mold only is rotated according to a rotational curve. Activates a rotational motion model.

Process Parameters Properties

Properties	Process parameter properties.	
	Translational Velocity	The translational velocity for the Bridgman casting process. This is specified in the form of a vector with respect to the laboratory coordinate system.
	Axis Direction	The direction of the z-axis of a new local coordinate system around which the rotation is performed. This is defined with respect to the laboratory coordinate system.
	Axis Origin	The origin of the new local coordinate system for the rotation with respect to the laboratory coordinate system.
	Rotational Rate	Specify the rotation rate in rad/s.
Options	The properties can be specified in different formats.	
	Constant	Sets a constant value of the property. Specify this value in the corresponding Values cell of the Properties table.
	Table	Sets the property value in the form of a table as a function of time.
Table Options	If Table is selected, the Values cell is greyed out and a table editor appears. The table can be input in two different ways.	
	Tabulate	Enter the tabular values manually in the table editor.
	Import .csv	Import the tabular values from a file, which contains comma separated values.
	Export .csv	The tabular values can be exported to a file containing comma separated values.

The Physical Models Panel

In the *Physical Models* panel, you define and select the parameters and options for optional physical models as part of the casting simulation. These optional models are:

- Melting and Solidification Flow Stop Model

  See Working with the Melting-Solidification Flow Stop Model
- Mushy Zone Permeability Model

  See Setting up the Mushy Zone Permeability Model
- Slurry Viscosity

  See Setting up the Slurry Viscosity Model
- Surface Tension

  See Surface Tension Formulation
- Macro Porosity Model

  See Using the Macro Porosity Model to Detect Shrinkage-Related Defects
- Contact Time

  See Using the Contact Time Model

The *Physical Models* panel consists of two tabs: the Selections tab and the Settings tab. First, you define the physical model in the Settings tab. Then, you assign the physical models conditions to the selected physical models in the Selections tab.

Physical Models Selections Tab

Physical Models	Contains the physical models available.
Conditions	Assigns the physical model conditions to the selected physical model. The available conditions must be set before in the Settings tab. Only the relevant physical model types are available, corresponding to the Type selection that you made in the Settings tab. For example, for the Melting-Solidification model you are only able to select a physical model condition of the type Melting-Solidification. After a selection has been made, it is shown at the bottom of the panel.
Model	Lists the selected physical model. The properties, options and values are shown underneath.

Physical Models Settings Tab

Conditions	Lists the physical models conditions that are defined.
Name	User-defined name under which the physical model condition is saved.
Type	Specifies the type of physical model condition to be used. There are different process physical model types available. The physical model conditions have different properties

Melting and Solidification	Selects the Melting and Solidification Flow Stop model (see Working with the Melting-Solidification Flow Stop Model).
Mushy Zone	Selects the Mushy Zone Permeability model (see Setting up the Mushy Zone Permeability Model).
Slurry Viscosity	Selects the Slurry Viscosity Model (see Setting up the Slurry Viscosity Model).
Liquid Residence Time	Selects the Liquid Residence Time criteria function (see Liquid Residence Time).

Macro Porosity	Selects the Macro Porosity model (see Using the Macro Porosity Model to Detect Shrinkage-Related Defects).
Surface Tension	Selects surface tension (see Surface Tension Formulation).
Contact Time	Selects the Contact Time model (see Using the Contact Time Model). Determines the contact time between a gas and a melt phase only in the liquid state. Once the solid volume fraction of the melt reaches the Solid Volume Fraction Limit in the simulation, the contact time stops accumulating.

Physical Models Properties

Properties

Relative Solid Fraction	Defines the cell solid volume fraction (values between 0 and 1) at which the fluid loses its ability to flow. However, the flow is stopped in the cell only when the total amount of solidified fluid exceeds the Flow Stop Solid Fraction that is specified for the Melting and Solidification Flow Stop model. This property is deactivated for values greater than 1. Default value 1.1 (see Melting-Solidification Properties).
Solid Fraction	Specifies when a cell is blocked for the flow. Set the solid fraction to a value between 0 and 1. A typical value is 0.9.
Max. Viscous Resistance	Sets a maximum level of resistance ($kg/m^3 s$). This setting applies only to the cells with a solid volume fraction greater than zero, in all regions to which the respective physics continuum is assigned.
Secondary Dendrite Arm Spacing	Is used for the calculation of the Carman-Kozeny Shape Factor (m). The shape factor c_s is a function of the secondary dendrite arm spacing d) and a shape constant c: $c_s = c/d^2$. The default value for c is 180.

Critical Solid Fraction	The critical solid volume fraction α_{cr}^{*} (see Eqn. (1789)). The default value is 0.27 (for Aluminum alloy A201).
Max. Solid Volume Fraction	The residence time is accumulated only for solid volume fractions below the specified value.
Surface Tension	Specifies the value of the surface tension of the melt in N/m.
Void Pressure	The pressure p_{void} in Eqn. (2493) below which a vapour bubble forms with the liquid phase. The default value is 10000 Pa.
Solid Volume Fraction Limit	Sets the value of the solid volume fraction at which contact time stops accumulating. Below this value, contact time is accumulated.

Options

The properties can be specified in different formats.

Constant	Sets a constant value of the property. You specify this value in the corresponding Values cell of the Properties table.
Use Switching	Toggles the switching option for the Mushy Zone Permeability model on and off. If switching is activated, specify the non-dimensional constant s (see Eqn. (1789)). The default value is 100 because the shape is rather insensitive for values from 80 through 120.

The Solvers Panel

In the *Solvers* panel, you specify the numerical parameters and stopping criteria for the casting simulation.

The panel consists of three tabs: the Process Phases tab, the Solver Settings tab, and the Stopping Criteria tab. First, you define the solver settings conditions in the Solver Settings tab as well as the Stopping Criteria conditions in the Stopping Criteria tab. Then, you assign both of these conditions to a process phase in the Process Phases tab.

Solvers Process Phases Tab

Process Phase	Contains the process phases. You need to define a process phase first by entering a user-defined name in the Process Phase Name text field and clicking Insert. A process phase is used to divide the casting simulation into several distinct physical time-dependent phases such as, for example, the pouring phase and the solidification phase. Each process phase can have different requirements regarding the time step size and the numerical solver parameters since the physical processes being simulated in each phase work in different time-scales. For example, the time step size for the solidification process can be significantly larger than for the pouring process phase.
Solver Settings	Assigns the solver settings conditions to the selected process phase. The available conditions must have been set before in the Solver Settings tab. After a selection has been made, it is shown at the bottom of the panel.
Stopping Criteria	Assigns the stopping criteria conditions to the selected process phase. The available conditions must have been set before in the Stopping Criteria tab. After a selection has been made, it is shown at the bottom of the panel.
Process Phase Name	Specify a user-defined name for the process phase that you want to define.

You can use a Java macro to start the simulation using the specified process phases. Save the process phase settings to a Java macro by clicking Export... Start the simulation by clicking Run. This plays the previously saved process phase Java macro.

Solver Settings Tab

Conditions — Scroll list where the defined solver settings conditions are listed by name.

Name — User-defined name under which the solver settings condition is saved.

Solver Settings Properties

Properties — Numerical solver parameters.

- Time Steps — The time step size in s.
- Pressure URF — The under-relaxation factor for the pressure-correction equation (see Pressure Solver).
- Volume of Fluid URF — The under-relaxation factor for the segregated VOF solver (see What Is the Segregated VOF Solver?).
- Fluid Energy URF — The fluid under-relaxation factor for the segregated energy solver (see Energy Solver).
- Solid Energy URF — The solid under-relaxation factor for the segregated energy solver (see Energy Solver).
- Passive Scalar URF — The under-relaxation factor for the passive scalar solver.
- Melting/Solidification URF — The under-relaxation factor for the melting and solidification model (see Melting-Solidification Expert Properties).
- Mushy Zone URF — The under-relaxation factor for the mushy zone permeability model.
- Slurry Viscosity URF — The under-relaxation factor for the slurry viscosity model.

Options — The properties can be specified in different formats.

- Constant — Sets a constant value of the property. You specify this value in the corresponding Values cell of the Properties table.

CFL-based	Activates the Segregated Time Step model. For more information, see Automated Time-Step Control. The Segregated Time Step model requires the setting of additional properties. For more information, see Gravity Casting Properties.
Table	Sets the property value in the form of a table as a function of time.
Solver Off	Toggles the solution of the respective equation on and off.

Table Options

If Table has been selected, the Values cell is greyed out and a table editor appears. The table can be input in two different ways.

Tabulate	Enter the tabular values manually in the table editor.
Import .csv	Import the tabular values from a file, which contains comma separated values.
Export .csv	The tabular values can be exported to a file containing comma separated values.

Stopping Criteria Tab

Conditions	Lists the stopping criteria conditions that are already defined.
Name	User-defined name under which the stopping criteria condition is saved.

Stopping Criteria Properties

Maximum Physical Time	Specify the maximum physical time for the process phase in s.
Maximum Inner Iteration	Specify the maximum number of inner iterations for each unsteady time step.

The Criteria Functions Panel

In the *Criteria Functions* panel, you specify the critical temperatures at which the criteria functions are evaluated. For a complete list of the available criteria functions and their definitions, see Criteria Functions Definitions.

Criteria Temperatures	Lists the critical temperatures that you have specified. To define critical temperatures, enter a temperature value in the Temperature text field and click Add.
Units	The critical temperature units. Currently, degrees Celsius is the only option available.
Temperature	A critical temperature between the solidus and liquidus temperatures at which Criteria Functions are evaluated.

The Post Processing Panel

In the *Post Processing* panel, you specify which results of the casting simulation you want to visualize and report during the run.

The *Post Processing* panel consists of three tabs: the Global tab, the Probes tab, and the Transient tab.

Post Processing Global Tab

Cavity Plausibility Checks	Activating an option invokes the creation of one or several reports, monitors and plots in the simulation tree.	
	Velocity	Creates a maximum and minimum velocity report, monitor and plot for all melt regions.
	Pressure	Creates a maximum and minimum pressure report, monitor and plot for all fluid regions.
	Temperature	Creates a maximum and minimum temperature report, monitor and plot for all fluid regions.
Filling/Solidification Progress		
	Melt (Mass)	Creates a field function, report, monitor, and plot to analyze the volume integral of the melt mass (kg/m^3) as a function of time (s).
	Solidified (Mass)	Creates a threshold, report, monitor, and plot to analyze the volume integral of the solidified melt mass (kg) as a function of time (s).
	Volume Ratio (Melt/Cavity)	Creates a field function, report, monitor, and plot to analyze the ratio of melt volume to total volume of the fluid domain.
	Volume Ratio (Solid/Liquid)	Creates a field function, threshold, report, monitor, and plot to analyze the ratio of solidified melt volume to melt volume.
Scalar Scenes	Activating an option invokes the creation of a scalar scene.	
	Air	Creates a scalar scene visualizing the temperature field of air as a contour plot.

Melt	Creates a scalar scene visualizing the temperature field of melt as a contour plot.
Solidified	Creates a scalar scene visualizing the temperature field of solidified melt as a contour plot.

Post Processing Probes Tab

Probes	Lists the probes that are already defined.
Name	User-defined name under which the probe is saved.
Type	Specifies the type of probe.
	Point: A point probe lets you extract variable values at a point within the computational domain.
	Line: A line probe lets you extract variable values at several points along a line within the computational domain.
Properties	Specifies the input parameters of the different probe types.
	Origin: Specify the coordinates of the point within the computational domain where the point probe is located.
	Point 1/2: Specify the coordinates of the two points connecting the line on which the line probe point lie.
	Resolution: Specify the number of probe points along the probe line.
X-Y Plots	Creates a graph of the selected variable value at the probe over time.
	Name: Specify a name for the graph.

Scalar	Specify the scalar variable that you want to monitor in the graph.
Probes	Select a previously defined probe at which you want to monitor the selected Scalar.

Post Processing Transient Tab

Automatic Backup

Activate the Enable button to define automatic backup of sim files during the simulation.

Trigger Type	Specify whether you want the files saved based on iteration, time step, or time interval.
Number of Backup Files	Specify the maximum number of backup sim files to be saved.
Start	Specify the starting iteration, time step, or time interval.
Number of Iterations	Specify the iteration frequency by which the backup sim files are saved.
Update Frequency	Specify the time step frequency by which the backup sim files are saved.
Delta Time	Specify the time interval by which the backup sim files are saved.
Enable Stop	When activated, you can specify from which point the update event stops updating. Depending on the selected Trigger Type, you specify iteration, time step, or time interval.

Post Process

Activate the Enable button to define creation of a transient post-processing simh file.

Trigger Type	Specify whether you want the visualization data saved based on iteration, time step or time interval.
Start	Specify the starting iteration, time step, or time interval.
Number of Iterations	Specify the iteration frequency by which the visualization data are saved to the simh file.

Update Frequency	Specify the time step frequency by which the visualization data are saved to the simh file
Delta time	Specify the time interval by which the visualization data are saved to the simh file
Enable Stop	When activated, you can specify from which point the update event stops updating. Depending on the selected Trigger Type, you specify iteration, time step, or time interval.

Variables of Interest

Select the variables that you want to store to the simh file for visualization.

Scalar Functions	Selection of scalar field functions.
Vector Functions	Selection of vector field functions.

Casting Field Functions

Liquid VOF Phase Relative Mass Error Rate

The relative mass conservation error rate for all liquid phases in the computational domain. The relative mass error conservation rate is defined as: $\varepsilon = \left|(m_{target} - m_{actual})/(m_{target} \cdot \Delta t) \cdot 100\right|$, where m_{target} is the target liquid mass that should be in the domain, m_{actual} is the liquid mass that actually is in the domain, and Δt is the current time step. The target liquid mass is given by: $m_{target} = m_0 + \sum \dot{m}_{Bound} \cdot \Delta t$ where m_0 is the mass computed from the previous time step and $\dot{m}_{Bound}$ are the mass fluxes through the boundaries.

You can use this field function to define a monitor-based stopping criterion.

Macro Porosity Field Functions Reference

The following primitive field functions are available to the simulation when the Macro Porosity model is used.

Density of Void<Phase>	The density of the auxiliary void phase.
Dynamic Viscosity of Void<Phase>	The dynamic viscosity of the auxiliary void phase.
Mass Flow Rate of Void<Phase>	The mass flow rate of the auxiliary void phase.
Mass Flux of Void<Phase>	The face mass flux of the auxiliary void phase.
Mass Imbalance of Void<Phase>	The imbalance of auxiliary void phase mass in a particular cell.
Volume Fraction of Void<Phase>	The fraction of volume of the current cell that the auxiliary void phase occupies.

Criteria Functions Field Functions Reference

The following primitive field function is available when the Temporary Storage Retained expert property of the Segregated VOF Solver is activated.

Source Rate of Void<Phase>	The VOF phase source for the auxiliary void phase.

The following primitive field functions are available to the simulation when the Criteria Functions model is activated.

Cooling Rate of <Phase> @ <Critical Temperature>	The Cooling Rate of the melting-solidifying phase at the critical temperature.

Secondary Dendrite Arm Spacing of <Phase> @ <Critical Temperature>	The Secondary Dendrite Arm Spacing of the melting-solidifying phase at the critical temperature.
Primary Dendrite Arm Spacing of <Phase>	The Primary Dendrite Arm Spacing of the melting-solidifying phase.
G/v of <Phase> @ <Critical Temperature>	The G/v ratio of the melting-solidifying phase at the critical temperature.
Isotherm Speed of <Phase> @ <Critical Temperature>	The Isotherm Speed of the melting-solidifying phase at the critical temperature.
LiquidResidenceTime<Phase>	The Liquid Residence Time of the melting-solidifying phase.
Local Solidification Time of <Phase> @ <Critical Temperature>	The Local Solidification Time of the melting-solidifying phase at the critical temperature.
Mean Cooling Rate of <Phase>	The Mean Cooling Rate of the melting-solidifying phase.
Niyama Criterion of <Phase>@<Critical Temperature>	The Niyama Criterion of the melting-solidifying phase. (The field function Niyama Criterion of <Phase> in earlier releases now corresponds to Niyama Criterion of <Phase>@T_solidus.
Solidification Velocity of <Phase> @ <Critical Temperature>	The Solidification Velocity of the melting-solidifying phase at the critical temperature.
SolidificationTime<Phase>	The Solidification Time of the melting-solidifying phase.

Temperature Gradient of <Phase> @ <Critical Temperature>	The Temperature Gradient of the melting-solidifying phase at the critical temperature.
Dimensionless Niyama of <Phase> @ <Critical Temperature>	The Dimensionless Niyama Criterion of the melting-solidifying phase at the critical temperature and the solidus temperature.
Pore Volume Fraction of <Phase> @ <Critical Temperature>	The absolute pore volume fraction as given by Eqn. (2513) for the critical temperature and the solidus temperature.

Casting Bibliography

[433] K.D. Carlson, C. Beckerman, "Prediction of Shrinkage Pore Volume Fraction Using a Dimensionless Niyama Criterion", Metallurgical and Materials Transactions A, **40A**, 2009, pp. 163-175

Modeling Region Sources

This section provides some guidelines on how to use region sources for some common problems. Currently, guidance is provided on:

- Marking cells for which region sources become active.
- Setting fluid sources or sinks.
- Fixing the temperature of fluid cells.

Marking Cells for Region Sources

Typically it is not necessary to apply a region source to an entire region, but to a small zone of cells within the region. It is possible to use cell sets, and their associated field functions, to mark such a region of cells.

In the example that is shown below, it is desired to set an energy source on a vertical column of cells. The cell set defining this can be created using a geometric range where the limits are set just outside the column of interest.



The cell set was created on the Volume Mesh representation.



This cell set, which has been named Cell Set 1 by default, has a corresponding field function whose Function Name is a contracted form of the cell set name. In this instance, the Function Name is `CellSetVar_1`.



To use the cell set in another field function expression, simply call `$CellSetVar_1`. This call returns the value of 1 for all cells that are contained within the cell set, and 0 otherwise.

Setting Fluid Sources or Sinks

It is sometimes desirable to inject or extract fluid mass directly in simulations without having to define a geometrical representation of the real aperture that is used to do so. For example, the length scale of the actual geometry of the aperture may be too small to resolve adequately. Therefore, injecting the fluid from a much smaller region is an appropriate method. In STAR-CCM+, various region sources are provided to make this kind of approach possible. Cell sets can be used to define the cells to which the region source applies.

The general rule to follow in applying this method is that every transport equation active in the region should have a corresponding source or sink term of the form

$$S_\phi = \dot{m}'''\phi$$

where $\dot{m}'''$ is the mass flow rate of the injected/extracted fluid per unit volume, and ϕ is the corresponding transported quantity. For the continuity equation, $\phi = 1$.

For each transport equation, the appropriate region condition, and the number of corresponding value nodes that may be added, are given as follows:

- Continuity: Mass Source Option (providing up to four additional value nodes)
- Momentum: Momentum Source Option (providing up to four additional value nodes)
- Energy: Energy Source Option (providing up to four additional value nodes)
- Turbulence: Turbulence Source Option (providing up to two additional value nodes)
- Species: Species Source Option (providing two additional value nodes)

Caution: If you define a mass source or sink, you must also define a corresponding momentum, energy, turbulence, or any other transport variable source or sink. Otherwise, the physics of your simulation is not sensible. For example, if you define a mass sink without a corresponding energy sink, the temperature of the system increases: the same amount of energy stays in the system, with less mass to hold it.

Prescribing Constant Temperature in Fluid Cells

In modeling some applications, such as air conditioning units, it may be desirable to impose the temperature distribution within a range of fluid cells. This section describes a procedure that uses the discretized form of the energy equation to fix a target temperature that you specify. The steps are as follows:

1. Create a cell set containing the cells for which the temperature is to be fixed. This cell set has a corresponding field function with a default name. Here, the field function is `CellSetVar_1`.
2. Create a field function, `GREAT`, that contains a number sufficiently large to mask the contribution from the usual energy fluxes in the designated cells, but not too high to cause a floating point exception. For example, you could use 1000 times the maximum local value of $(mCp)/V$ or

kL/V, where m is the mass flow rate into the cell, Cp is the specific heat, V is the cell volume, k is the thermal conductivity, and L is some characteristic cell length. In this example, GREAT is set to `1e20`.



3. Create a field function for the user-defined temperature field. Here, for example, `T_target` is set to:

```
293 + 40 *$$Centroid[1] * (1 - $$Centroid[1])
```



4. Using these three field functions, it is possible to define a further field function for the Volumetric Heat Source, *Q*, that applies to cells in the cell set:

```
($CellSetVar_1) ? $GREAT*($T_target - $Temperature) : 0
```

5. The differential of this field function with respect to temperature is also required as a field function. Here, it is defined as dQdT, with the formula:

```
($CellSetVar_1) ? -$GREAT : 0
```



Having created the necessary field functions, the final step is to apply them to the relevant region values:

6. Set the Volumetric Heat Source of the region to the field function `Q`:



7. Set the Energy Source Temperature Derivative to the field function `dQdT`:



Remedying Cell Quality

The Cell Quality Remediation model helps you get solutions on a poor-quality mesh. This model identifies poor-quality cells, using a set of predefined criteria, such as Skewness Angle exceeding a certain threshold. Once these cells and their neighbors have been marked, the computed gradients in these cells are modified in such a way as to improve the robustness of the solution. In general, the effect of Cell Quality Remediation is confined to the immediate vicinity of poor-quality and/or degenerate cells, so that the influence on overall solution accuracy is minimal. However, a side effect is that the mesh quality metrics used to mark the poor-quality cells are recomputed each time the mesh changes. Therefore, in certain situations such as moving meshes, using the model can have a significant computational cost.

Subsequent sections provide the following information:

- Troubleshooting
- Field functions that this model activates

Troubleshooting

One or more cells have an invalid least squares tensor

You can encounter an error message similar to the following:

```
Invalid least squares tensor in cell with Local Index
193495 in region inlet on partition 19
```

This error comes from the least squares (reconstruction) gradient calculation. When the least squares tensor is inverted using the Jacobi method, this error indicates that the tensor is invalid, and the algorithm

cannot proceed. Since the least squares tensor is derived from geometrical considerations, the error signifies an invalid cell. Diagnose and repair or regenerate the mesh.

You can use the information in the error message to identify the problem:

- Create a user field function with the following definition:

  ```
  ($LocalCellIndex == 193495) && (PartitionID == 19)
  ```

- Create a threshold derived part where this quantity equals 1.

These steps help you visualize the cell. The cell can be flat, and/or have almost no neighboring cells. In addition, proximity to a poorly intersected interface can contribute to the bad cell.

Cell Quality Remediation Field Functions Reference

The following field functions are made available to the simulation when the Cell Quality Remediation model is used.

Bad Cell Indicator

This field function is activated with the Cell Quality Remediation Model, and it is defined as an integer value of 0, 1 or 2.
0: The cell is not considered bad.
1: The cell is considered bad.
2: The cell is a neighbor of a bad cell.

Boundary Sliver Cell Indicator

This field function is used to indicate cells that can be considered "boundary slivers" and hence could be removed from the solution domain without affecting the topology of the mesh. To be considered a boundary sliver, two conditions must be satisfied:
1) All vertices of the cell must lie on a boundary.
2) The Cell Quality must be below the specified tolerance (`1e-6` by default), indicating that the cell is too flat.
The field function has an integer value of either 0 (not a boundary sliver cell) or 1 (a boundary sliver cell).

Least Squares Quality

This field function is an indicator of the quality of a cell, using the physical location of a cell centroid relative to the cell centroid locations of its face-neighbors.



The field function value is based on the ratio of the maximum and minimum eigenvalue ratios of the symmetrical matrix:

$$\overline{\overline{M}} = \sum_{\substack{cell\\ faces\\ i}}^{''} \frac{\mathbf{ds}_i \otimes \mathbf{ds}_i}{\left|\mathbf{ds}_i\right|^2}$$

where $\mathbf{ds}_i$ represents the vector from the cell to the neighbor centroid. Two methods are optionally provided to evaluate the eigenvalue ratio: approximate and exact. The exact method uses a Jacobi transformation. The approximate method uses the Frobenius norm of the matrix to estimate the eigenvalue ratio. It is slightly more economical to evaluate than the exact method, and is guaranteed to be less than or equal to the exact value. A perfect cell has a Least Squares Quality computed by the exact algorithm equal to unity. The same perfect cell would have an approximate value of 1/N, where N is the dimension of the mesh (2 or 3). As the Least Squares Quality becomes small (< 0.1), the approximate and exact values yield almost identical results. An option also exists to include the boundary faces of the cell in the computation for boundary cells.

Zero Area Face Indicator

This field function has the value of unity if the area of any of the faces approaches zero; otherwise it has a value of zero. It is used to diagnose potential mesh problems.

Controlling Domain Decomposition

To make use of parallel processing, STAR-CCM+ breaks the computational domain into separate sections (called *sub-domains*) and assigns each to a separate process. This approach is called *domain decomposition*. Once the domain is decomposed, the processes work in parallel on their individual sub-domains, periodically communicating with each other.

The goals of domain decomposition are to balance the work as evenly as possible among the processes and to minimize the amount of communication that is needed between them. You can control domain decomposition through the partitioning solver.

The Partitioning Solver

The partitioning solver is responsible for controlling domain decomposition (or partitioning) for parallel simulations. Before any physics solvers are invoked on a given global iteration or time-step, the partitioning solver is first invoked to ensure that the domain decomposition is up-to-date.

The solver generally repartitions the domain when one of the following conditions is met:

- The cell connectivity has changed materially since the last partitioning.
- The computational loads have become imbalanced.
- The partitioning method is changed.

Partitioning Solver Operation

Changes in inputs to the partitioning solver – including but not limited to the specified partitioning method, cell connectivity, and load balancing settings and weights – are only applied when either:

- The partitioning solver is next run in the normal course of stepping or iterating,

  or

- The partitioning is manually updated with the Repartition Domain action.

The Repartition Domain action is included as a convenience; it is not needed to keep the partitioning up-to-date with the inputs. If the partitioning is out-of-date, STAR-CCM+ automatically repartitions the domain on the next global iteration or time-step. The Repartition Domain action lets you immediately inspect the impact of repartitioning before advancing the solution.

With respect to save-restore, the partitioning solver generally behaves like other STAR-CCM+ solvers – if the input settings are changed just before saving and exiting the simulation, they will not be applied to the saved simulation until the next time the solver runs after the simulation is restored. Restoring into a different process count only changes the partition count; it does not change the applied partitioning method, which remains the same as it was after the last time the partitioning solver ran before the simulation was saved.

Cell Connectivity Changes

The solver partitions the computational domain according to the connectivity graph between the cells. The solver partitions the graph to balance the work, and to minimize communication between the sub-domains. The solver usually balances work by balancing the number of computational cells and faces, but it can also consider work-per-cell from specific physics solvers. Communication between sub-domains is minimized by minimizing the number of cut edges in the connectivity graph.

Because performance depends on partitioning and partitioning depends on the cell connectivity, the solver recomputes the partitioning when a tangible change in cell connectivity takes place. For instance, if you add an interface between regions, the interface introduces new cell connectivity between the regions, which requires more information transfer. If the partitioning does not consider this connectivity, the resulting performance can be lower than it need be.

Currently, the domain is repartitioned when a cell connectivity change results from:

- Adding, removing, or modifying either:
  - An intra-region direct boundary interface (an interface within a region) when per-region partitioning is activated,

    or:
  - An intra-continuum direct boundary interface (an interface within a continuum) when per-continuum partitioning is activated.

  See What Is a Direct Interface?
- Splitting a region.
- Combining multiple regions.
- Modifying the mesh within a region such that it modifies the connectivity.

Load Balancing

The second role of the partitioning solver is to repartition the domain when the per-process loads become imbalanced. The goal is to keep loading as evenly balanced as possible.

The partitioning solver:

1. Periodically checks if specific physics computations have become imbalanced.
2. If computations have become imbalanced, the solver repartitions the domain according to the load information that those physics solvers provide.

Currently, load balancing is only available for Lagrangian transport. (See the Lagrangian Load Balancing section.)

Because load balancing itself has a cost, it is not always desirable on every global iteration. Usually, the computational solution changes gradually from iteration to iteration, so a partitioning based on the loads at a given global iteration is often suitable for a number of future iterations.

You can control the load balancing frequency using the Rebalance frequency property on the Load balancing options node within the partitioning solver. Load balancing and the associated Load balancing options node only become available when a physics solver that supports load balancing is active. Currently only the Lagrangian multiphase solver supports load balancing.

Partitioning Method

Every simulation has an overall computational domain that can be defined as the union of all the regions in the simulation upon which physics is being simulated.

One approach to partitioning the overall computational domain is a monolithic one. The simulation creates a graph of the entire domain and decomposes it in a single operation, producing a single set of self-contiguous subdomains. This approach, however, can lead to poor performance and scalability in some situations. For example, if one continuum has more complex physics than another, this approach can result in severe load imbalances.

For this reason, partitioning over domains of finer granularity is usually desirable. In STAR-CCM+, you can partition over the following extents:

- **Per-region partitioning:** Each region is partitioned individually. For each region, a connectivity graph for all the cells in the region is created, and that graph is partitioned to balance the work and minimize partition boundaries for that region. The result is that the cells in each region are divided evenly between processes, and the sub-domains within each region are contiguous.

Note: For a multi-region continuum, the sub-domains are not contiguous with respect to the containing continuum when per-region partitioning is employed.

- **Per-continuum partitioning:** Each continuum is partitioned individually. For each continuum, a connectivity graph is created for all the cells in all the regions in that continuum, and that graph is partitioned to balance the work and minimize partition boundaries for that continuum. The result is that the cells in each continuum are balanced, and the sub-domains within each continuum are contiguous.

In STAR-CCM+, per-region partitioning has been the historical approach. While per-region partitioning often balances loads successfully, it sometimes has the following disadvantages:

- The overall decomposition is more discontiguous from a process perspective, ultimately resulting in more inter-process communication in terms of both of data volume and message count.
- Connectivity across inter-region interfaces is not considered in the partitioning process, often resulting in both boundary face imbalances and increased inter-process communication.

These issues may adversely affect both performance and scalability.

Per-continuum partitioning, in contrast, mitigates or eliminates these issues. Because partitioning takes place over a larger extent, the per-process sub-domains are more contiguous. This reduces communication and other parallel overhead relative to per-region partitioning. As a result, per-continuum partitioning can significantly improve both performance and scalability relative to per-region partitioning in many situations.

In general, the benefits of per-continuum partitioning increase with both increasing number of regions per continuum and increasing process count.

Despite the general benefits of per-continuum partitioning and most applications usually performing better with per-continuum partitioning, some very specific STAR-CCM+ application types can still sometimes benefit from per-region partitioning. For instance, some harmonic balance applications still perform better with per-region partitioning at this time.

A change in partitioning method goes into effect the next time the partitioning solver runs. To apply the specified method immediately, use the Repartition Domain item from the partitioning solver popup menu.

Partitioning Solver Parameters

The partitioning solver can be found at the usual solver location in the simulation object tree.



Expert Properties

Partitioning method	Controls how the solver partitions the domain. A change in partitioning method becomes effective the next time the solver runs. See Partitioning Method.	
	Per-Region	Partitions each region individually.
	Per-Continuum	Partitions across all regions in each continuum. This is the default for new simulations.
Solver Frozen	Controls whether the Partitioning solver is active or inactive (frozen).	
	Activated	Partitioning solver is inactive or frozen.
	Deactivated	Partitioning solver is active (the default).

Partitioning Solver Pop-Up Menu Items

Repartition Domain	Immediately repartitions the domain according to the current mesh connectivity. This action is seldom needed because, if the partitioning is out-of-date, the domain is automatically repartitioned on the next global iteration or time-step. The action is included as an option to update the partitioning manually outside the context of advancing the solution.

See also: common menu items.

Load Balancing Options Node Properties

The Load balancing options node, which is located under the partitioning solver in the tree, becomes available only when a physics solver supporting it is active.

Rebalance frequency	Controls the frequency of load balancing in terms of global iterations or time steps. The default frequency is: • Once per iteration for steady state problems. • Once per time step for unsteady problems.

Solving Transport Equation

The mechanics of continua (both solid and fluid) are described by transport equations derived from the assumption that mass, momentum, energy and entropy are conserved in the continuum. The purpose of this chapter is to describe the basic methods used to discretize and solve these transport equations.

The sections below describe:

- finite volume discretization with reference to the transport of a simple scalar;
- gradient computation methods for both the reconstruction and the cell gradients;
- implicit iteration methods used to linearize and assemble the algebraic equation systems; and
- algebraic multigrid methods used to iteratively solve the discrete linear systems.

In addition, nomenclature and a bibliography are included. For further background material, readers are referred to [435] and [437].

Finally, reference information is provided on transport equation solvers.

Finite Volume Discretization

In the finite volume method, the solution domain is subdivided into a finite number of small control volumes, corresponding to the cells of a computational grid. Discrete versions of the integral form of the continuum transport equations are applied to each control volume. The objective is to obtain a set of linear algebraic equations, with the total number of unknowns in each equation system corresponding to the number of cells in the grid. (If the equations are non-linear, iterative techniques that rely on suitable linearization strategies must be employed.) The resulting linear equations are then solved with an algebraic multigrid solver.

This section describes the basics of the finite volume discretization methods that are used in STAR-CCM+. The transport of a simple scalar (as opposed to a vector or tensor) quantity is used for illustrative purposes. The sections below describe the governing equation and the discretization approach for its individual terms, specifically:

- the transient term,
- the source term,
- the convective term; and
- the diffusion term.

Governing Equation

Continuous Integral Form

The integral equation

$$\frac{d}{dt}\int_V \rho\chi\phi\, dV + \oint_A \rho\phi(\mathbf{v}-\mathbf{v}_g)\cdot d\mathbf{a} = \oint_A \Gamma\nabla\phi\cdot d\mathbf{a} + \int_V S_\phi\, dV \tag{2530}$$

represents the transport of a scalar quantity ϕ in a continuum

The terms in this equation are, from left to right, the transient term, the convective flux, the diffusive flux, and the volumetric source term.

Discrete Form

Applying the above equation to a cell-centered control volume for cell-0, the following is obtained:

$$\frac{d}{dt}(\rho\chi\phi V)_0 + \sum_f [\rho\phi(\mathbf{v}\cdot\mathbf{a} - G)]_f = \sum_f (\Gamma\nabla\phi\cdot\mathbf{a})_f + (S_\phi V)_0 \tag{2531}$$

where G is the grid flux that is computed from the mesh motion as in Eqn. (16). The following sections describe the approximations that are employed when writing each term of this discrete equation as functions of the cell variables.

Transient Term

The transient term is only included in transient calculations; it is not used as a device to obtain a steady-state solution. The Implicit Unsteady solver in STAR-CCM+ offers two temporal discretization options: first-order and second-order.

First-Order

The First-Order temporal scheme, also referred to as Euler Implicit, discretizes the unsteady term using the solution at the current time level, $n+1$, as well as the one from the previous time level, n:

$$\frac{d}{dt}(\rho\chi\phi V)_0 = \frac{(\rho\chi\phi V)_0^{n+1} - (\rho\chi\phi V)_0^{n}}{\Delta t} \tag{2532}$$

Second Order

A second-order temporal scheme discretization of the unsteady term uses the solution at the current time level, $n+1$, as well as the solutions from the previous two time levels, n and $n-1$, as follows:

$$\frac{d}{dt}(\rho\chi\phi V)_0 = \frac{\left\{(\alpha^2-1)[(\rho\chi\phi V)_0^{n+1} - (\rho\chi\phi V)_0^{n}] + [(\rho\chi\phi V)_0^{n-1} - (\rho\chi\phi V)_0^{n}]\right\}}{\alpha(\alpha-1)\Delta t^{n+1}} \tag{2533}$$

and:

$$\alpha = 1 + \frac{\Delta t^{n+1}}{\Delta t^{n}} \tag{2534}$$

$$\Delta t^{n+1} = t^{n+1} - t^{n} \tag{2535}$$

$$\Delta t^{n} = t^{n} - t^{n-1} \tag{2536}$$

On the first time-step of a second-order temporal simulation, a first-order discretization is used since only two time levels are available.

Source Term

The source term in Eqn. (2530) has been approximated in Eqn. (2531) by the product of the value of the integrand, S_ϕ, evaluated at the cell centroid, and the cell volume, V:

$$\int_V S_\phi \, dV = (S_\phi V)_0 \tag{2537}$$

This equation is the simplest formulation consistent with a second-order discretization.

Convection Term

The convective term at a face is discretized as follows:

$$[\phi\rho(\mathbf{v} \cdot \mathbf{a} - G)]_f = (\dot{m}\phi)_f = \dot{m}_f \phi_f \tag{2538}$$

where ϕ_f and $\dot{m}_f$ are the scalar values and mass flow rates at the face, respectively and G is the grid flux that is computed from the mesh motion as in Eqn. (16).

The manner in which the face value ϕ_f is computed from the cell values has a profound effect on the stability and accuracy of the numerical scheme. Several schemes are commonly used:

- First-Order Upwind
- Second-Order Upwind
- Central-Differencing
- Bounded Central-Differencing
- Hybrid MUSCL 3rd-Order/CD
- Hybrid Second-Order Upwind/Central
- Hybrid Second-Order Upwind/Bounded-Central

A description of the computation of the convective flux at a boundary face is also provided in this section.

First-Order Upwind

For a first-order upwind scheme, the convective flux is computed as:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}_f \phi_0 & \text{for } \dot{m}_f \geq 0 \\ \dot{m}_f \phi_1 & \text{for } \dot{m}_f < 0 \end{cases} \tag{2539}$$

This scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, the numerical dissipation that is thus introduced leads to "smearing" discontinuities, particularly when those discontinuities are not aligned with the grid lines. Generally, this option would only be selected if a solution using the second-order convection was unobtainable. It is also fairly common practice to obtain an initial solution using the first-order convection and then switch to the second-order one to obtain a more accurate, final converged solution.

Second-Order Upwind

For a second-order upwind scheme, the convective flux is computed as:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}_f \phi_{f,0} & \text{for } \dot{m}_f \geq 0 \\ \dot{m}_f \phi_{f,1} & \text{for } \dot{m}_f < 0 \end{cases} \tag{2540}$$

where the face values $\phi_{f,0}$ and $\phi_{f,1}$, are linearly interpolated from the cell values on either side of the face as follows:

$$\phi_{f,0} = \phi_0 + \mathbf{s}_0 \cdot (\nabla\phi)_{r,0} \tag{2541}$$

$$\phi_{f,1} = \phi_1 + \mathbf{s}_1 \cdot (\nabla\phi)_{r,1} \tag{2542}$$

where:

$$\mathbf{s}_0 = \mathbf{x}_f - \mathbf{x}_0 \tag{2543}$$

$$\mathbf{s}_1 = \mathbf{x}_f - \mathbf{x}_1 \tag{2544}$$

and $(\nabla\phi)_{r,0}$ and $(\nabla\phi)_{r,1}$ are the limited reconstruction gradients in cells 0 and 1, respectively.

The advantage of this scheme over the first-order upwind scheme is that it is nominally second-order accurate. However, the fact that the reconstruction gradients are limited helps to reduce local extrema and thus introduces more dissipation than a central-differencing scheme. Clearly, the accuracy of this scheme is always as good or better than the first-order upwind scheme. The downside is that, in some situations, the reduced numerical dissipation can result in poorer convergence properties than a first-order convection. Generally, the poorer convergence is an acceptable trade-off.

Central-Differencing

For a central-differencing scheme, the convective flux is computed as:

$$(\dot{m}\phi)_f = \dot{m}_f[f\phi_0 + (1-f)\phi_1] \tag{2545}$$

where the geometric weighting factor f is related to the mesh stretching. It would have a value of 0.5 for a uniform mesh.

Central-differencing is formally second-order accurate. However it is prone to dispersive error and is beset with stability problems for most steady-state situations. The dispersive errors make it problematical for discretizing positive-definite quantities (such as temperature or turbulent kinetic energy) where overshoots cannot be tolerated.

A significant advantage of central-differencing over second-order upwinding is that, when used to discretize velocity (not a positive definite quantity) it preserves turbulent kinetic energy. Therefore, it is a useful scheme in large eddy simulation (LES), where upwind schemes cause turbulent kinetic energy to decay unnaturally fast.

Bounded Central-Differencing

This scheme is the default for large eddy simulation (LES).

For a bounded central-difference (BCD) scheme as described in [436] and [434], the convective flux is computed as:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}\phi_{fou} & \text{for } \varsigma < 0 \text{ or } 1 < \varsigma \\ \dot{m}(\sigma\phi_{cd} + (1-\sigma)\phi_{sou}) & \text{for } 0 \le \varsigma \le 1 \end{cases} \tag{2546}$$

where, at face f:

- ϕ and ϕ_{fou} are for the first-order upwind scheme.
- ϕ_{sou} is for the second-order upwind scheme.
- ϕ_{cd} is for the central-differencing scheme.
- ς is the Normalized-Variable Diagram (NVD) value, computed based on local conditions.

A smooth and monotone function of ς is:

$$\sigma = \sigma(\varsigma) \tag{2547}$$

which satisfies $\sigma(0) = 0$ and $\sigma(\varsigma) = 1$ for $\varsigma_{ubf} \le \varsigma$. ς_{ubf} is called the "upwind blending factor" and its value ensures a proper balance between the accuracy and robustness of the scheme. Thus, smaller values of ς_{ubf} provide more accuracy, while larger values increase the robustness of the scheme.

To maintain boundedness, the bounded central-differencing scheme turns into a first-order upwind scheme when the convection boundedness criterion is not satisfied, for example when $\varsigma < 0$ or $1 < \varsigma$. The central-differencing scheme on the contrary, which is formally a

second-order accurate scheme, does not do that. Thus, the bounded central-differencing scheme can be more dissipative than the central-differencing scheme --especially on coarser meshes.

Overall, the bounded central-differencing scheme provides a good compromise between accuracy (much improved, when compared with the second-order upwind scheme) and robustness (due to its warranted boundedness). It is recommended as the scheme for LES of complex turbulent flows.

Hybrid MUSCL 3rd-Order/CD

This scheme is recommended for high-fidelity turbulence (LES/DES), aeroacoustic, and aerodynamic simulations. The scheme is valid for both steady and unsteady simulations, and has one model parameter σ, which is used to control the numerical dissipation in the scheme (see Recommendations).

As with the Bounded Central-Differencing scheme, this scheme uses a Normalized-Variable Diagram (NVD) value ζ to ensure the boundedness of the scheme by switching to the first-order scheme in regions of non-smooth flows.

When smooth local flow conditions are detected, the scheme is constructed as a blend between a MUSCL 3rd-order upwind and the 3rd-order central-differencing reconstruction schemes.

The convective flux is computed as:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}\phi_{fou} & \text{for } \zeta < 0 \text{ or } 1 < \zeta \\ \dot{m}(\sigma\phi_{muscl3} + (1-\sigma)\phi_{cd3}) & \text{for } 0 \le \zeta \le 1 \end{cases} \tag{2548}$$

The blending factor σ is user-controlled and must be decided based on the physical problem or model (and mesh quality considerations; see Recommendations).

The MUSCL 3rd-order upwind reconstructed value ϕ_{muscl3} is carefully limited at high speeds so as to not affect the formal order of accuracy while preventing spurious oscillations. The high speed limiting is based on WENO (Weighted Essentially Non-Oscillatory) principles and is only applied to the quadratic part of the face-value reconstruction. WENO-based averaging is computed based on three different stencils.

Note: The ϕ_{muscl3} accuracy is reduced to second order in the regions next to the strong shocks.

The switch between the unlimited and the WENO-limited ϕ_{muscl3} is done based on a Mach-number-based flux-limiter.

For incompressible simulations, ϕ_{muscl3} is *unlimited* (the flux limiter is 1), thus only relying on the use of the NVD value to ensure the boundedness of the scheme.

Overall, the Hybrid MUSCL 3rd-order/CD scheme provides improved (reduced) dissipation when compared with both second-order and Bounded Central-Differencing (BCD) schemes. At the same time, it is robust (due to its boundedness) and capable of simulating steady and unsteady flows from incompressible to high-speed compressible regimes.

Recommendations

- The default value for the model parameter σ is 1.0. Use lower values to reduce the numerical dissipation of the scheme.

Note: Values as low as 0.02 have been used successfully for simulation of turbulent flows (LES or DES) at low speeds.

- Set σ to 1 or close to 1 for cases with poor meshes or for high-speed flows. If you have good quality meshes, you can progressively reduce σ to prevent the solver from running into numerical instability.
- Cases which previously ran with BCD using a bounded differencing factor σ_{bcd} can also be run with Hybrid MUSCL 3rd-order/CD. You can use the same bounded differencing factor, or even lower factors in some cases.
- The reduced numerical dissipation of the MUSCL 3rd-order/CD scheme sometimes prevents steady cases from converging. If those cases converged with numerical dissipation in the 2nd-order scheme, and if the physical problem is inherently unsteady, run with the Unsteady model activated.
- The following parameters sometimes require adjustment when switching the Convective scheme from 2nd-order or BCD (used for LES and DES) to the Hybrid MUSCL 3rd-order/CD scheme:
  - The CFL number in the Coupled solver.
  - The Pressure or Velocity URFs in the Segregated solver.
  - Any other convergence-related parameters for flow solvers.

  STAR-CCM+ requires this adjustment to account for, or prevent, possible convergence problems due to less numerical dissipation. Adjustments are principally required for steady-state problems.

- The Hybrid MUSCL 3rd-order/CD scheme is not warranted to work in all situations or physical conditions. Sometimes, convergence of numerical models requires you to use some (reasonable) amount of numerical dissipation.
- Lack of convergence, or change in convergence, can indicate model issues, such as a simulation that should be run as unsteady, imperfect boundary conditions, or inconsistent problem formulations. These issues do not represent a bug.

Hybrid Second-Order Upwind/Central

The most accurate discretization schemes in STAR-CCM+ are the second-order upwind and central-differencing schemes. The second-order upwind is the most accurate scheme for steady-state calculations, while the kinetic energy-preserving properties of the central-differencing scheme make it appropriate for large eddy simulation (LES). Therefore, for detached eddy simulation (DES) it would be advantageous to blend both approaches as follows:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}_f\{\sigma\phi_{f,0} + (1-\sigma)[f\phi_{f,0} + (1-f)\phi_{f,1}]\} & \text{for } \dot{m}_f \geq 0 \\ \dot{m}_f\{\sigma\phi_{f,1} + (1-\sigma)[f\phi_{f,0} + (1-f)\phi_{f,1}]\} & \text{for } \dot{m}_f < 0 \end{cases} \tag{2549}$$

where σ is a blending factor that is chosen as appropriate for the flow regime. The face values $\phi_{f,0}$ and $\phi_{f,1}$, are linearly interpolated from the cell values on either side of the face using the reconstruction gradients, as with the second-order upwind schemes

Hybrid Second-Order Upwind/Bounded-Central

The following hybrid second-order upwind/bounded central scheme is the default for detached eddy simulation (DES) of complex turbulent flows:

$$(\dot{m}\phi)_f = \dot{m}(\sigma\phi_{sou} + (1-\sigma)\phi_{bcd}) \tag{2550}$$

where at face f, the values of ϕ are ϕ_{sou} for the second-order upwind scheme and ϕ_{bcd} for the bounded central-differencing scheme. σ is a blending factor that is chosen as appropriate for the flow regime (see the hybrid second-order upwind/central scheme).

Boundary Face Discretization

First-Order

When the first-order upwind scheme is used, the flux at a boundary face is evaluated as:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}_f\phi_0 & \text{for } \dot{m}_f > 0 \\ \dot{m}_f\phi_f & \text{for } \dot{m}_f < 0 \end{cases} \tag{2551}$$

where ϕ_0 is the cell value and ϕ_f is the face value as dictated by the boundary conditions.

Second-Order

When the second-order upwind, central-differencing, or hybrid second-order upwind/central schemes are used, the flux at a boundary face is evaluated as:

$$(\dot{m}\phi)_f = \begin{cases} \dot{m}_f\phi_{f,0} & \text{for } \dot{m}_f > 0 \\ \dot{m}_f\phi_f & \text{for } \dot{m}_f < 0 \end{cases} \tag{2552}$$

where $\phi_{f,0}$ is interpolated from the cell value using the limited reconstruction gradients in cell 0 (see Eqn. (2541)) and ϕ_f is the face value as dictated by the boundary conditions.

Diffusion Term

Let D_f be the discrete form of the diffusion term:

$$D_f = \sum_f (\Gamma\nabla\phi \cdot \mathbf{a})_f \tag{2553}$$

where Γ, $\nabla\phi$ and $\mathbf{a}$ represent the face diffusivity, gradient, and area vector, respectively.

Interior Face

To obtain an accurate second-order expression for an interior face gradient that implicitly involves the cell values ϕ_0 and ϕ_1, the following decomposition is used:



$$\nabla\phi_f = (\phi_1 - \phi_0)\vec{\alpha} + \overline{\nabla\phi} - (\overline{\nabla\phi} \cdot \mathbf{ds})\vec{\alpha} \tag{2554}$$

where:

$$\vec{\alpha} = \frac{\mathbf{a}}{\mathbf{a} \cdot \mathbf{ds}} \tag{2555}$$

$$\mathbf{ds} = \mathbf{x}_1 - \mathbf{x}_0 \tag{2556}$$

$$\overline{\nabla\phi} = \frac{(\nabla\phi_0 + \nabla\phi_1)}{2} \tag{2557}$$

The diffusion flux at an interior face can then be written:

$$D_f = \Gamma_f \nabla\phi_f \cdot \mathbf{a} = \Gamma_f[(\phi_1 - \phi_0)\vec{\alpha} \cdot \mathbf{a} + \overline{\nabla\phi} \cdot \mathbf{a} - (\overline{\nabla\phi} \cdot \mathbf{ds})\vec{\alpha} \cdot \mathbf{a}] \tag{2558}$$

where Γ_f is a suitable average value (normally a harmonic average) of the cell values.

Interior Secondary Gradients and Skewness Angle

The second and third terms in Eqn. (2558) represent the secondary gradient (or cross-diffusion) contribution. They are essential for maintaining accuracy on non-orthogonal meshes.

The formulation that is presented above assumes that the centroids of cells 0 and 1 lie on opposing sides of the face. It is further assumed that their location is consistent with the convention that the face area vector points out of cell 0. To prevent non-physical solutions, take care to use meshes for which the angle between $\mathbf{a}$ and $\mathbf{ds}$ is not greater than 90 degrees.

To ensure that a valid mesh is being used, a diagnostic tool that is provided in STAR-CCM+ allows this angle (termed skewness angle) to be computed in degrees and stored in adjacent cells. The value that is stored in each cell represents the largest skewness angle for each face of the cell.

As a further diagnostic tool that can be employed to check whether the secondary gradient contribution is destabilizing, these terms can be optionally omitted, in which case Eqn. (2558) reduces to:

$$D_f \approx \Gamma_f[(\phi_1 - \phi_0)\vec{\alpha} \cdot \mathbf{a}] \tag{2559}$$

Use this option with caution since the resulting loss of solution accuracy can be unacceptable.

Boundary Face

A similar decomposition is used at a boundary face:



$$D_f = \Gamma_f \nabla\phi_f \cdot \mathbf{a} = \Gamma_f[(\phi_f - \phi_0)\vec{\alpha} \cdot \mathbf{a} + \nabla\phi_0 \cdot \mathbf{a} - (\nabla\phi_0 \cdot \mathbf{ds})\vec{\alpha} \cdot \mathbf{a}] \tag{2560}$$

where:

$$\vec{\alpha} = \frac{\mathbf{a}}{\mathbf{a} \cdot \mathbf{ds}} \tag{2561}$$

$$\mathbf{ds} = \mathbf{x}_f - \mathbf{x}_0 \tag{2562}$$

The entire diffusion flux can be omitted at flow boundaries using the Flow Boundary Diffusion property of segregated flow and energy models.

Boundary Secondary Gradients and Skewness Angle

Similar to interior faces, the second and third terms in Eqn. (2560) represent the boundary secondary gradient (or cross-diffusion) contribution. Again, the angle between $\mathbf{a}$ and $\mathbf{ds}$ for boundary faces has to be greater than 0 and not greater than 90 degrees, to prevent non-physical solutions. The angle (termed boundary skewness angle) can be optionally computed and stored on the adjacent boundaries.

As with the interior, these boundary secondary gradients can be optionally neglected as a further diagnostic tool, in which case Eqn. (2560) reduces to:

$$D_f \approx \Gamma[(\phi_f - \phi_0)\vec{\alpha} \cdot \mathbf{a}] \tag{2563}$$

Use this option with caution since the resulting loss of solution accuracy can be unacceptable.

Gradient Computation

Gradients are used in several places within the transport equation solution methodology, including:

- Reconstructed field values at the cell faces
- Secondary gradients for diffusion terms
- Pressure gradients for pressure-velocity coupling in the segregated flow model
- Strain-rate and rotation-rate calculations for turbulence models

In STAR-CCM+, the steps that are involved in computing gradients are:

1. Computing the (unlimited) reconstruction gradients. This computation is done using either the Hybrid Gauss-LSQ or the Green-Gauss gradient method. Here, unlimited means that the gradients do not prohibit the reconstructed field variables on the cell faces from exceeding the minimum and maximum values of the neighboring cells.
2. Limiting the reconstruction gradients. The limited reconstruction gradients are used to determine scalar values at the cell faces. These scalar values are used in computing flux integrals.
3. Computing the cell gradients from the limited reconstruction gradients. This step is only required for the Green-Gauss method. For the Hybrid Gauss-LSQ method, the LSQ-based reconstruction gradients are linear-exact and are copied into cell gradients. This approach is more accurate for cell gradient calculations than the Green-Gauss method.

The Gradients model is automatically selected once the space model is defined.



Gradients Properties

Verbose	Provides more output while the simulation is running. This property can be useful for debugging problems when they occur.	
	Activated	Activates verbose output.
	Deactivated	Leaves output at its default.
Limiter Method	Specifies which limiter method to use.	
	Venkatakrishnan	Selects the Venkatakrishnan method that is given by Eqn. (2583).
	Modified Venkatakrishnan	Selects the Modified Venkatakrishnan method that is given by Eqn. (2585).
Gradient Method	Specifies which gradient method to use.	
	Hybrid LSQ-Gauss	Selects the Hybrid LSQ-Gauss method (default option).
	Green-Gauss	Selects the Green-Gauss method.
Custom Accuracy Level Selector	Specifies the accuracy of the numerical scheme at any level between 1.0 (1st-order) and 2.0 (2nd-order). This property is only active when a 2nd-order scheme is selected. Use any real value between 1.0 and 2.0. Default is 2.0, meaning no accuracy changes. See Accuracy Level.	

Expert Properties

Acceptable Field Variation (Factor)	Sets the allowed variation of the field due to TVB Gradient Limiting, as a fraction of the largest difference between local maximum and local minimum anywhere in the simulation. Default 0.05.
Maximum safe (positive) skewness angle (deg.)	Below the positive only skewness angle, no related specific corrections are applied to least-squares computed gradients. Default 75.0.

Normalized Flat Cells Curvature Factor	Acceptable ratio between the tangent of cell-skewness angle and the aspect ratio of the cell. Default 1.0.
Least-Squares Quality Criterion	Makes gradient corrections for cells with poor least-squares tensor quality. Default on and is read-only.
Cell Skewness Criterion	Makes gradient corrections for cells with large skewness angle. Default on.
Chevron-Cell Criterion	Makes gradient corrections for all chevron cells. Default on.
Least-Squares Tensor Minimum Eigenvalues Ratio	Minimum admissible value of ratio between the minimum and maximum eigenvalue of the least-squares tensor. Default 0.1.
Use TVB Gradient Limiting	TVB (Total Variation Bounded) limiting prevents excessive gradient limiting, which can cause a slow or stalled convergence. See Total Variation Bounded Gradient Limiting. Activated: Activates TVB gradient limiting. Deactivated: The default.
Minimum unsafe (positive) skewness angle (deg.)	Above the positive skewness angle and for all negative skewness angle cells, gradients are corrected to ensure robustness. Default 88.0.
Flat Cells Curvature Criterion	Makes gradient corrections for flat cells with sensible curvature. Default on.

Reconstruction Gradients

Three algorithms are used to compute reconstruction gradients that are based on the selected Gradients model:

- For the Hybrid Gauss-LSQ gradient method, the hybrid Gauss/weighted LSQ method is used for all variables
- For the Green-Gauss gradient method
  - the weighted least squares method is used for pressure
  - the "simple" Gauss method is used for all variables other than pressure

Hybrid Gauss-Least Squares Method

The unlimited (denoted by superscript u) reconstruction (denoted by subscript r) gradients in cell-0 are computed using the following hybrid formula:

$$(\nabla\phi)_r^u = \sum_f (\phi_n - \phi_0) w_f^0 \tag{2564}$$

with

$$w_f^0 = \beta w_f^{lsq} + (1-\beta) w_f^G \tag{2565}$$

$$w_f^{lsq} = \left[\sum_f \frac{\mathbf{ds} \otimes \mathbf{ds}}{\mathbf{ds}^2}\right]^{-1} \frac{[\mathbf{ds}]}{\mathbf{ds}^2} \tag{2566}$$

$$\mathbf{ds} = \mathbf{x}_n - \mathbf{x}_0 \tag{2567}$$

$$w_f^G = \frac{A_f}{V_0 + V_n} \tag{2568}$$

where $\mathbf{x}_0$ and $\mathbf{x}_n$ are the centroids of cell-0 and that of the neighbor cell n, addressed through face f. ϕ_0 and ϕ_n are data values in cell-0 and its neighbor, A_f is the face f area-vector, V_0 and V_n are the respective cell volumes, and β is the geometric Gauss-LSQ gradient blending factor field function.

Blending Factor for Hybrid Gauss-LSQ Gradient

The Gauss/LSQ gradient blending factor is used to determine how the two gradient method options are used in calculating the gradient. The following conditions apply to the bending factor in gradient calculation:

- Blending factor of 1 equates to using LSQ method
- Blending factor of 0 equates to using Green-Gauss method
- Blending factor from 0 through 1 equates to using a combination of LSQ and Green-Gauss, with:

$$\nabla\phi = (\beta)\mathbf{LSQ_grad}(\phi) + (1-\beta)\mathbf{GG_grad}(\phi) \tag{2569}$$

where β is the cell value of the gradient factor, $\mathbf{LSQ_grad}(\phi)$ is the LSQ computed gradient of ϕ and $\mathbf{GG_grad}(\phi)$ is the Green-Gauss computed gradient of ϕ.

If the Green-Gauss method is selected for gradient calculation, Gauss-LSQ gradient blending factor is not used.

The cell value of the blending factor is computed as the minimum of the following:

- Blending value is 1 if the least-squares tensor eigenvalues ratio of the cell is larger than the specified minimum value (set using Least-Squares Quality Criterion property)
- Blending value is 1 if the tangent of the skewness angle is smaller than the Normalized Flat Cells Curvature Factor times the aspect ration of the cell. It is less than 1, with a rapid decrease to zero otherwise (set using the Flat Cell Curvature Criterion property)
- Blending value is 1 if the skewness angle is smaller than the Maximum safe (Positive) skewness angle. Blending value is 0 if the skewness angle is larger than the Minimum unsafe (Positive) skewness angle. Blending value is from 0 through 1 with a linear variation, otherwise (Cell Skewness Criterion)
- Blending value is 0 if it is a Chevron-cell and 1 otherwise (Chevron-Cell Criterion)

The Gauss/LSQ gradient blending factor is only applicable with the Hybrid LSQ-Gauss method.

Weighted Least Squares Method

The unlimited reconstruction gradients in cell-0 are computed using the following weighted least squares formula:

$$(\nabla\phi)_r^u = \left[\sum_f \frac{\mathbf{ds}\otimes\mathbf{ds}}{\mathbf{ds}\cdot\mathbf{ds}}\right]^{-1}\left[\sum_f \frac{(\phi_0-\phi_n)\mathbf{ds}}{\mathbf{ds}\cdot\mathbf{ds}}\right] \tag{2570}$$

where

$$\mathbf{ds} = \mathbf{x}_n - \mathbf{x}_0 \tag{2571}$$

with $\mathbf{x}_0$ and $\mathbf{x}_n$ representing the centroids of cell-0 and that of the neighbor cell that is addressed through face f, and ϕ_0 and ϕ_n representing the data values in cell-0 and its neighbor.

Gauss Method

From the divergence theorem of Gauss:

$$\int_V \nabla\phi \, dV = \oint_A \phi \, d\mathbf{a} \tag{2572}$$

Written in discrete form, this divergence theorem can be used to compute the unlimited reconstruction gradients:

$$(\nabla\phi)_r^u = \frac{1}{V_0} \sum_f \phi_f \, \mathbf{a}_f \tag{2573}$$

where the face value is approximated by the arithmetic average of the adjacent cell values:

$$\phi_f = \frac{\phi_0 + \phi_1}{2} \tag{2574}$$

Reconstruction Gradient Limiting

The problem with simply reconstructing face values from the unlimited reconstruction gradients is that the reconstructed face values can fall outside the range of cell values found in neighboring cells (connected through faces). For this reason, STAR-CCM+ finds the minimum and maximum bounds of the neighboring cell values and uses these to limit the reconstruction gradients.

- Venkatakrishnan limiter, the default (Eqn. (2583))
- Min-Mod limiter (Eqn. (2584))
- Modified Venkatakrishnan limiter (Eqn. (2585))

The face value ($\phi_{f,0}$) reconstructed from the cell-0 value at any face centroid is given by:

$$\phi_{f,0} = \phi_0 + \mathbf{s}_0 \cdot (\nabla\phi)_{r,0} \tag{2575}$$

where $\mathbf{s}_0 = \mathbf{x}_f - \mathbf{x}_0$, $\mathbf{x}_f$ and $\mathbf{x}_0$ are the face and cell centroids, respectively, and $\nabla\phi_{r,0}$ is the reconstructed face value.

For each cell-0, a limited reconstruction gradient is needed, such that the reconstructed face value does not exceed the maximum and minimum of the neighboring cell centroid values, including the value in cell-0. A scale factor α is defined that expresses the ratio of the limited and unlimited values, that is:

$$(\nabla\phi)_{r,0} = \alpha\,(\nabla\phi)^{u}_{r,0} \tag{2576}$$

For each cell-0, the quantities are defined:

$$\phi_0^{\max} = \max\,(\phi_0, \phi_{\text{neighbors}}) \tag{2577}$$

$$\phi_0^{\min} = \min\,(\phi_0, \phi_{\text{neighbors}}) \tag{2578}$$

where $\phi_{\text{neighbors}}$ represents the cell value in each neighbor that has a common face with cell-0.

These quantities can also be defined:

$$\Delta_{\max} = \phi_0^{\max} - \phi_0 \tag{2579}$$

$$\Delta_{\min} = \phi_0^{\min} - \phi_0 \tag{2580}$$

For each face *f* of cell-0, is defined:

$$\Delta_f = \phi_{f,0} - \phi_0 = \mathbf{s}_0 \cdot (\nabla\phi)^{u}_{r,0} \tag{2581}$$

Now defining:

$$r_f = \begin{cases} \dfrac{\Delta_f}{\Delta_{\max}} & \text{for } \Delta_f > 0 \\ \dfrac{\Delta_f}{\Delta_{\min}} & \text{for } \Delta_f \le 0 \end{cases} \tag{2582}$$

The Venkatakrishnan [439] limiter gives for the face:

$$\alpha_f = \frac{2r_f + 1}{r_f(2r_f + 1) + 1} \tag{2583}$$

The Min-Mod limiter gives for the face:

$$\alpha_f = \min(1, 1/r_f) \tag{2584}$$

The Modified Venkatakrishnan limiter gives for the face:

$$\alpha_f = \frac{r_f(2r_f + 1) + 1}{r_f(r_f(2r_f + 1) + 1) + 1} \tag{2585}$$

The Venkatakrishnan limiter is the default option. In certain cases, the Min-Mod and the modified Venkatakrishnan limiter can be less dissipative than the default option, particular in unsteady simulations for aeroacoustics or when using LES/DES for turbulence modeling.

The cell value of α is given by:

$$\alpha = \min(\alpha_f) \tag{2586}$$

Total Variation Bounded Gradient Limiting

If gradients are excessively limited, convergence can slow or stall. This problem arises most often with the gradients of near-constant fields, or when the limiter cycles between 0 and 1 without settling on a convergence value.

To avoid this problem, activate the model property TVB Gradient Limiting (Total Variation Bounded gradient limiting). This property lets reconstructed values vary by an additional difference of

$$\delta = \psi \cdot \max(\Delta_{max} - \Delta_{min}) \tag{2587}$$

where ψ is a fraction between 0 and 1, and $\max(\Delta_{max} - \Delta_{min})$ is the largest difference between local maximum and local minimum anywhere in the simulation.

Set the value of ψ with the Acceptable Field Variation property. The default value is 0.05. Larger values mean gradients are limited less. This can improve accuracy, especially for smooth flow field gradient reconstruction. However, excessively large values can lead to loss of solver robustness and lower accuracy, by producing large over- or under-shoots at flow field discontinuities.

TVB Gradient Limiting is useful for high-fidelity simulations, such as aeroacoustics and turbulence simulations using Large Eddy Simulation (LES) or Detached Eddy Simulation (DES).

Accuracy Level

When computing steady-state solutions of complex flow problems using high-order discretization, you can sometimes encounter convergence stall or even lack of convergence. These issues are due to flow instabilities, sometimes local ones, that are generally caused by numerical artifacts.

One solution for these situations is to tune the numerical scheme dissipation to a user-specified level, between the first order and second order of accuracy.

In the Gradients model, the Custom Accuracy Level Selector property allows you to set an accuracy level. When necessary, use this property to introduce more damping in the numerical scheme and so improve convergence.

The accuracy level selector only applies when the second-order scheme is selected. It is implemented as a global scaling applied to the reconstruction gradients, thus affecting only convection.

Note: Gradients used for diffusion are not affected by this scaling.

Cell Gradients

Cell gradients are computed using the reconstruction gradients. When using the Hybrid Gauss-LSQ method, the cell gradients are a copy of the reconstruction gradients. For the Green-Gauss Method, the improved estimates of the face values are obtained from the reconstruction gradients. They can in turn be used to obtain better estimates of the cell gradients using the divergence theorem of Gauss:

$$\nabla\phi = \frac{1}{V_0}\sum_f \phi_f \mathbf{a}_f \tag{2588}$$

where the face value is approximated by the arithmetic average of the face values reconstructed from the adjacent cell values:

$$\phi_f = \frac{\phi_{f,0} + \phi_{f,1}}{2} \tag{2589}$$

where, as before:

$$\phi_{f,0} = \phi_0 + \mathbf{s}_0 \cdot (\nabla\phi)_{r,0} \tag{2590}$$

$$\phi_{f,1} = \phi_1 + \mathbf{s}_1 \cdot (\nabla\phi)_{r,1} \tag{2591}$$

Implicit Iteration

The discretization approach results in the coefficients of a linear equation system being obtained. This system is solved implicitly, in an iterative fashion. The following sections describe:

- methods that are used for implicit under-relaxation;
- methods that are used for the delta form of the equations; and
- the need for iterative solution.

Implicit Under-Relaxation

The algebraic system for the transported variable ϕ at iteration *k+1* is written implicitly as:

$$a_p \phi_p^{k+1} + \sum_n a_n \phi_n^{k+1} = b \tag{2592}$$

where the summation is over all the neighbors n of cell p. The right-hand side, b, represents the explicit (that is, evaluated with the results from iteration k) contributions to the discretized equation. The coefficients a_p and a_n are obtained directly from the discretized terms.

An under-relaxation factor ω is introduced implicitly as follows:

$$\frac{a_p}{\omega} \phi_p^{k+1} + \sum_n a_n \phi_n^{k+1} = b + \frac{a_p}{\omega}(1-\omega)\phi_p^k \tag{2593}$$

where the superscript $k + 1$ implies the value after the solution is produced, and the source term on the right-hand side is evaluated at the previous iteration (k).

Delta Form

Rather than solving Eqn. (2593) for the unknowns ϕ^{k+1}, it is instead cast into "delta" form. Defining $\Delta\phi_p = \phi_p^{k+1} - \phi_p^k$, the system becomes:

$$\frac{a_p}{\omega} \Delta\phi_p + \sum_n a_n \Delta\phi_n = b - a_p \phi_p^k - \sum_n a_n \phi_n^k \tag{2594}$$

The right-hand side:

$$r = b - a_p \phi_p^k - \sum a_n \phi_n^k \tag{2595}$$

is termed the residual, and represents the discretized form of the original equation (Eqn. (2531)) at iteration k. By definition, then, the residual becomes zero when the discretized equation is satisfied exactly.

Iterative Solution

For linear phenomena such as constant-property solid conduction, the linear system needs only be constructed and solved only once. In most situations, however, the system is non-linear. For example, the source term S_ϕ or transport coefficient Γ in Eqn. (2530) could themselves be functions of ϕ. In this case, an iterative solution is required. There are two levels of iteration; an outer iteration loop controlling the solution update and an inner loop governing the iterative solution of the linear system. Since the outer iterations are repeated multiple times, it is sufficient to solve the linear system only approximately at each iteration. The iterative solution of linear equations is covered in the next section.

Algebraic Multigrid

The algebraic multigrid methods solve the discrete linear system iteratively. The AMG Linear Solver node, which has its own properties and subject to some common menu items, contains the objects that you would use for multigrid cycling strategies.



The result of the discretization approach that is described above is a linear system:

$$\mathbf{A}\,\mathbf{x} = \mathbf{b} \tag{2596}$$

representing the algebraic equations for each computational cell. The matrix $\mathbf{A}$ represents the coefficients of the linear system (for example, the coefficients a_p and a_n on the left-hand side of Eqn. (2594), the vector $\mathbf{x}$ represents the unknowns ($\Delta\phi$ in Eqn. (2594)) in each cell, and the vector $\mathbf{b}$ represents the residuals (see Eqn. (2595)) from each cell.

Typically, the matrix $\mathbf{A}$ is sparse. Direct methods such as Gauss elimination or LU decomposition on such systems are costly, since the triangular factors of sparse matrices are not sparse. Therefore, such methods are untenable for practical problems involving large grids. It is therefore preferable to use an efficient iterative method, such as the algebraic multigrid method in STAR-CCM+.

More detail is available on:

- Basic iterative methods
- Multigrid methods
- Multigrid cycling strategies
- The preconditioned conjugate-gradient method
- The bi-conjugate gradient stabilized method

- AMG Linear Solver Troubleshooting

AMG Linear Solver Properties

Max. cycles	Specifies the maximum number of multigrid cycles to run in case the residuals of the linear system remain above Epsilon and the Convergence Tolerance not be met.
Verbosity	Specifies the amount of feedback (directed to the console window) by the AMG Linear Solver.
	None — The AMG Linear Solver runs silently. There is no feedback.
	Low — Residuals of the linear system are printed at the completion of each multigrid cycle.
	High — Same as Low, plus the number of multigrid levels and number of equations at each level are printed once all multigrid cycles are complete.
	Diagnostics — Provides the highest degree of information regarding performance of the AMG system. This option can affect performance, and therefore is not recommended for normal iteration.
Convergence Tolerance	Factor by which residuals of the linear system (that are above the value of Epsilon) must be reduced before multigrid cycling stops (provided Max. cycles have not been reached). For example, a convergence tolerance of 0.01 requires that cycling continue until residuals of the linear system have dropped by two orders of magnitude.
Cycle Type	Specifies the multigrid cycling strategy.

F Cycle	Adds an F cycle sub-node to this node. After restriction one W-cycle is applied on the next coarsest level, followed by one V-cycle (also on the next coarsest level).
V Cycle	Adds an V Cycle sub-node to this node. After restriction one cycle (also a V-cycle) is applied on the next coarsest level.
W Cycle	Adds an W Cycle sub-node to this node. After restriction two cycles (also W-Cycles) are applied on the next coarsest level.
Flex Cycle	Adds a Flex Cycle sub-node to this node.

AMG Linear Solver Expert Properties

Enable Direct Solver	Controls whether to use the direct solver for AMG coarse levels.	
	Activated	Activates this feature.
	Deactivated	Leaves this feature deactivated.
Maximum Direct Solver Equations	Sets the maximum number of equations for the direct solver.	
Epsilon	Cutoff value for residuals of the linear system below which convergence is considered complete and multigrid cycling stops, regardless of the Convergence Tolerance setting. For example, an Epsilon value of 1.e-08 implies that cycling stops when residuals of the linear system drop below 1.e-08.	
Group Size Control	Specifies how the group size must be set.	
	Auto	The group size is set automatically to 2 in two-dimensional cases and 4 in three-dimensional cases.
	Specified	Allows a user-specified group size. Adds a Group Size sub-node to this node.
Group Size	The number of equations on a given multigrid level that are agglomerated to form one equation on the next coarsest level. This value is read-only and reports the current group size being used, as determined by the `Auto` setting or from the Group Size sub-node when `Specified`.	
Relaxation Scheme	Specifies method for iteratively correcting (relaxing) the linear equation that is set during multigrid cycling.	
	Jacobi	A slow, simple method.
	Gauss-Seidel	Provides better convergence. The default option.

	ILU	A more efficient but more expensive method.
Acceleration Method	Specifies the acceleration type for the AMG solver.	
	None	No acceleration method used.
	Conjugate Gradient	Uses the Conjugate-Gradient method with the AMG solver as the preconditioner. This method is available with the Constant Density and Segregated Flow models only. It is accessed through the Expert Properties of the AMG Linear Solver sub-node of the Segregated Flow > Pressure solver node.
	Bi-Conjugate Gradient Stabilized	Uses the Bi-Conjugate Gradient Stabilized method to solve non-symmetric (general) linear systems.
Scaling	Controls the manner in which corrections are prolongated from the next coarsest level. Scaling is only appropriate when the linear system matrix $\mathbf{A}$ (Eqn. (2603)) is symmetric positive-definite.	
	Auto	Scaling is automatically applied to the coarse level corrections to achieve optimal convergence.
	Disabled	Coarse level corrections are prolongated without any scaling applied. (The default.)

Basic Iterative Methods

The general principle behind iterative methods is that, given an approximate solution $\mathbf{x}^k$, a better approximation $\mathbf{x}^{k+1}$ is sought, and then the process is repeated.

If the error at iteration k is defined as:

$$\mathbf{e}^k = \mathbf{x} - \mathbf{x}^k \tag{2597}$$

where $\mathbf{x}$ represents the exact solution and the residual as:

$$\mathbf{r}^k = \mathbf{b} - \mathbf{A}\mathbf{x}^k \tag{2598}$$

it follows that:

$$\mathbf{A}\mathbf{e}^k = \mathbf{r}^k \tag{2599}$$

Therefore, continuing the iteration until the residual is driven to a small value, drives the error a small value.

The most basic iterative methods are Jacobi and Gauss-Seidel iteration. These methods involve visiting each cell in sequence, and updating the value of x_i in each cell i using the coefficients of its n neighbor cells as follows:

$$x_i = \frac{1}{A_{i,i}}\left(b - \sum_{\text{neighbors } n} A_{i,n}\, x_n \right) \tag{2600}$$

The difference between Jacobi and Gauss-Seidel iteration is subtle: Jacobi uses the "old" values of x_n, while Gauss-Seidel uses the available values that have been updated, resulting in better convergence.

The ILU (Incomplete Lower-Upper) method, Saad [438], solves iteratively for $\boldsymbol{x}^{k+1}$ using:

$$\tilde{\boldsymbol{A}}(\boldsymbol{x}^{k+1} - \boldsymbol{x}^k) = b - \boldsymbol{A}\boldsymbol{x}^k \tag{2601}$$

where

$$\tilde{\boldsymbol{A}} = \boldsymbol{L}\boldsymbol{U} = (\boldsymbol{D} + \boldsymbol{L}_A)(\boldsymbol{I} + \boldsymbol{D}^{-1}\boldsymbol{U}_A) \tag{2602}$$

where $\boldsymbol{I}$ is the identity matrix, $\boldsymbol{L}_A$ and $\boldsymbol{U}_A$ are the lower and upper triangular matrices of the original matrix $\boldsymbol{A}$, and $\boldsymbol{D}$ is a diagonal matrix that is calculated using the MILU(0) method.

The ILU method is more computationally expensive but more robust than the Gauss-Seidel method. The robustness of the ILU method means that it can handle group sizes of 8 or more in three dimensions; the Gauss-Seidel method can only handle group sizes of 4. Use the ILU method for complex problems, and to improve convergence and parallel scalability. For three-dimensional problems, use it with a group size of 8.

Multigrid Methods

The primitive iteration methods that are described above, while relatively simple to implement, exhibit relatively slow convergence characteristics. They tend to be only effective at removing high-frequency (rapidly varying)

components of the error. This effectivity suggests that some of the work could be done on a coarse grid. Computations on coarse grids are much less costly and the Gauss-Seidel method converges four times faster on a grid half as fine. Multigrid algorithms do this using the following steps:

1. Agglomerate cells to form coarse grid levels.
2. Transfer the residual from a fine level to a a coarser level (known as restriction).
3. Transfer the correction from a coarse level back to a finer level (known as prolongation).

Multigrid algorithms can be divided into two types: geometric and algebraic.

- Geometric multigrid uses the grid geometry and the discrete equation at the coarse level to arrive at the linear system that is to be solved on that level.
- Algebraic multigrid derives a coarse level system without reference to the underlying grid geometry or discrete equations. The coarse-grid equations are derived from arithmetic combinations of the fine-grid coefficients.

Since it is not always straightforward to obtain suitable discrete equations on the coarse levels, algebraic multigrid (AMG) is clearly at an advantage. Therefore, it is used for the solution of all linear systems in STAR-CCM+.

The Group Size object facilitates agglomeration.

Accelerating the Linear Solver

Use one of two optional methods for accelerating the linear solver, Preconditioned Conjugate-Gradient (PCG) for incompressible flows and Bi-Conjugate Gradient Stabilized (BiCGStab) for all others, in either of the following situations:

- AMG is struggling to converge (taking many cycles).
- Many cores in parallel require improvements in scalability. In parallel, the AMG solver loses its efficiency somewhat, taking longer to converge. In addition, it has poor scalability on a high number of cores. By contrast, AMG with acceleration takes fewer cycles to converge, and both PCG and BiCGStab have better scalability than the AMG solver itself.

Group Size

The Group Size node is made available when the Group Size Control property of the AMG Linear Solver is set to `Specified`. Manually set this value in the Properties window. It is also subject to common menu items.



Group Size Properties

Group Size	Sets the number of equations on a given multi-grid level that are agglomerated to form one equation on the next coarsest level.

Multigrid Cycling Strategies

Multigrid techniques enable a basic iteration scheme such as Gauss-Seidel to be greatly accelerated by using simple correction sweeps on a sequence of coarse meshes. The strategy for visiting the coarse levels can have a significant impact on the efficiency of the algorithm. Two cycling strategies are available in the AMG solver that is used in STAR-CCM+ -- fixed and flexible:

- Fixed cycles -- A complete multigrid cycle represents the recursive application of a single cycle that is composed of the following steps:
  a. (Pre)smooth
  b. Restrict
  c. Cycle anew
  d. Prolongate
  e. (Post)smooth

These steps are applied to a sequence of successively coarser grids or equation sets. Smooth means to apply any number of iterative relaxation sweeps to the equations on the current fine level to compute a new set of corrections. Restrict is to transfer the existing residuals down to the next coarsest level where a new cycle is applied. Following that the resultant corrections are prolongated--transferred back up to the current fine level--where smoothing is again applied.

The three types of fixed cycles are:

- F fixed cycles
- V fixed cycles
- W fixed cycles

- Flexible cycles -- This type of cycle is a more economical cycling strategy for linear systems that are not stiff. Rather than using all multigrid levels in a regular pattern, the residuals are monitored after every sweep on a given grid level. If the ratio of the residuals exceeds a given threshold, the solution continues on a coarser level. If the residual on a given level is reduced more than a specified tolerance, the solution moves to a finer level. A limit is further imposed on the number of sweeps that are allowed at any level.

F Fixed Cycle

The F Cycle node represents the AMG F Fixed Cycle. This object shares properties with the other cycle types and is subject to common menu items.



The F cycle is a variant of the W cycle. Illustrated below, this cycle involves less coarse-level sweeps than the W cycle. but still more than the V cycle.

- 

V Fixed Cycle

The V Cycle node represents the AMG V Fixed Cycle. This object shares properties with the other cycle types and is subject to common menu items.



The V cycle is the simplest type of fixed cycle, and only has two legs. In the first leg, one performs a number of relaxation sweeps on the finest level and transfers the residuals to the next level. Then, the operation is repeated on successively coarse levels until the coarsest level is reached. A coarse "grid" generally contains only a few "cells". After finishing the sweeps on the coarsest level, the solution is used to correct the solution on the next finer

level. Some relaxation sweeps are performed on that level before repeating the process until the finest level is reached. This procedure is illustrated below.

- 

W Fixed Cycle

The W Cycle node represents the AMG W Fixed Cycle. This object shares properties with the other cycle types and is subject to common menu items.



For stiff systems, the V cycle sometimes is not sufficient, and more coarse iterations are advantageous. The W cycle increases the number of coarse relaxation sweeps as illustrated below.

- 

Fixed Cycle Properties

The three fixed-cycle objects share these properties.

Pre-Sweeps	Number of smoothing sweeps at the start of each cycle.
Post-Sweeps	Number of smoothing sweeps at the end of each cycle.
Max Levels	Maximum number of coarse levels in the sequence.

Flex Cycle

The Flex Cycle node represents the AMG Flex Cycle, which has its own properties and is subject to common menu items.



Flex Cycle Properties

Restriction Tolerance	Ratio of two successive residuals of the linear system below which progression to the next coarsest level is suspended.
Prolongation Tolerance	Specifies the amount by which the residuals of the linear system must be reduced before the solution is stopped on any given level.
Sweeps	Number of smoothing sweeps at the start of each cycle.

Preconditioned Conjugate-Gradient Method

To accelerate the convergence of the iterative solution of the linear system in the Pressure solver, a conjugate-gradient method is available that uses the AMG solver as the preconditioner. Use this method for incompressible flows (Constant Density model) with the Segregated Flow model.

To solve the linear system $\mathbf{A}\,\mathbf{x} = \mathbf{b}$, the preconditioned conjugate-gradient (PCG) takes the following form:

$$\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0 \tag{2603}$$

$$\mathbf{z}_0 \cong \text{AMG fixed cycle } (\mathbf{A}, \mathbf{r}_0) \tag{2604}$$

$$\mathbf{p}_0 = \mathbf{z}_0 \tag{2605}$$

$$\mathbf{k} = 0 \tag{2606}$$

Each iteration involves the calculations in Eqns. (2607) through (2613):

$$\alpha_k = \frac{(\mathbf{r}_k^T \mathbf{z}_k)}{(\mathbf{p}_k^T \mathbf{A}\mathbf{p}_k)} \tag{2607}$$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k \tag{2608}$$

$$\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k \mathbf{A}\mathbf{p}_k \tag{2609}$$

If $|\mathbf{r}_{k+1}| < \varepsilon$, then the iteration ends.

$$\mathbf{z}_{k+1} \cong \text{AMG fixed cycle } (\mathbf{A}, \mathbf{r}_{k+1}) \tag{2610}$$

$$\beta_k = \frac{(\mathbf{r}_{k+1}^T \mathbf{z}_{k+1})}{(\mathbf{r}_k \mathbf{z}_k)} \tag{2611}$$

$$\mathbf{p}_{k+1} = \mathbf{z}_{k+1} + \beta_k \mathbf{r}_k \tag{2612}$$

$$k = k + 1 \tag{2613}$$

The solution of the linear system is $\mathbf{x}_{k+1}$.

A necessary condition for the convergence of the PCG algorithm is that the matrix $\mathbf{A}$ (which here is also the preconditioner) is symmetric positive-definite (SPD). The matrix of the linear system of the Pressure solver in pure incompressible cases satisfies this condition.

Equations (2604) and (2610) in the iteration process are preconditioner solver calculations, which involve one AMG fixed-cycle loop to approximate the auxiliary linear systems:

$$\mathbf{A}\mathbf{z}_k = \mathbf{r}_k \qquad k = 0, 1, \ldots \tag{2614}$$

As implemented, the PCG accelerates convergence and improves the parallel scalability of the Pressure solver. For normal cases (no non-linear instability), it is expected that the PCG does not change significantly the overall non-linear convergence of STAR-CCM+. The PCG can provide robust convergence for those cases where numerical instability results from a lack of incomplete convergence of the linear system of the Pressure solver.

Bi-Conjugate Gradient Stabilized Method

The Bi-Conjugate Gradient Stabilized (BiCGStab) method has been developed to solve nonsymmetric (general) linear systems. STAR-CCM+ implements an AMG-preconditioned BiCGStab algorithm to improve the robustness and the speed of convergence of linear systems that the solvers produce.

To solve the general linear system $\mathbf{A}\,\mathbf{x} = \mathbf{b}$, the preconditioned BiCGStab takes the following steps:

$$\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0 \tag{2615}$$

$$\mathbf{r} = \mathbf{r}_0 \tag{2616}$$

$$i = 0 \tag{2617}$$

Each iteration i involves calculations of Eqns. (2618) through (2631):

$$\rho_{i-1} = \mathbf{r}^T \mathbf{r}_{i-1} \tag{2618}$$

If $\rho_{i-1} = 0$, the method fails.

If $i = 1$:

$$\mathbf{p}_i = \mathbf{r}_{i-1} \tag{2619}$$

Otherwise, as per Eqns. (2620) and (2621):

$$\beta_{i-1} = \left(\frac{\rho_{i-1}}{\rho_{i-2}}\right)\left(\frac{\alpha_{i-1}}{\omega_{i-1}}\right) \tag{2620}$$

$$\mathbf{p}_i = \mathbf{r}_{i-1} + \beta_{i-1}(\mathbf{p}_{i-1} - \omega_{i-1}\mathbf{v}_{i-1}) \tag{2621}$$

$$\mathbf{p} = \text{AMG fixed cycle } (\mathbf{A}, \mathbf{p}_i) \tag{2622}$$

$$\mathbf{v}_i = \mathbf{A} \cdot \mathbf{p} \tag{2623}$$

$$\alpha_i = \frac{\rho_{i-1}}{(\mathbf{r}^T \cdot \mathbf{v}_i)} \tag{2624}$$

$$\mathbf{S} = \mathbf{r}_{i-1} - (\alpha_i \cdot \mathbf{v}_i) \tag{2625}$$

If $(|\mathbf{S}| < \varepsilon)$, then as per Eqn. (2626):

$$\mathbf{x}_i = \mathbf{x}_{i-1} + \alpha_i \mathbf{p} \tag{2626}$$

and the iteration ends.

Otherwise the iteration continues as follows:

$$\mathbf{\hat{S}} = \text{AMG fixed cycle } (\mathbf{A}, \mathbf{S}) \tag{2627}$$

$$\mathbf{t} = \mathbf{A} \cdot \mathbf{\hat{S}} \tag{2628}$$

$$\omega_i = \frac{\mathbf{t}^T \cdot \mathbf{S}}{\mathbf{t}^T \cdot \mathbf{t}} \tag{2629}$$

$$\mathbf{x}_i = \mathbf{x}_{i-1} + \alpha_i \mathbf{\hat{p}} + \omega_i \mathbf{S} \tag{2630}$$

$$\mathbf{r}_i = \mathbf{S} - \omega_i \mathbf{t} \tag{2631}$$

If $|\mathbf{r}_i| < \varepsilon$, then the iteration ends.

If $\omega_i \neq 0$, then the iteration continues.

The solution of the linear system is $\mathbf{x}_i$.

Eqns. (2622) and (2627) in the iteration process are the preconditioner solver calculations, which involve *one* AMG fixed-cycle loop to obtain an approximate solution of:

$$\mathbf{A} \cdot \mathbf{\hat{p}} = \mathbf{p}_i \quad \text{and} \quad \mathbf{A} \cdot \mathbf{\hat{S}} = \mathbf{S} \quad \text{respectively} \tag{2632}$$

As with the Preconditioned Conjugate-Gradient method, BiCGStab can improve parallel scalability and improve robustness for those cases where numerical instability was due to lack of sufficient convergence of the linear system.

Troubleshooting the AMG Linear Solver

Follow the steps below when you suspect that the AMG Linear Solver causes convergence problems:

1. In the *AMG Linear Solver Properties* window, set the Verbosity to `Diagnostics`.



2. Continue to iterate, and in the *Output* window, identify the solver for which AMG is crashing.
3. For the solver in which the AMG Linear Solver is crashing, select the cycle node and, in the *Properties* window, set the Max Levels property to `0`. This setting deactivates the multigrid procedure.



4. Also change the Max Cycles property to a larger value, around `50 - 100`. Try running the simulation again.



5. If divergence still occurs, it indicates an error in the coefficient of the linear system coming from the discretization. The problem is not located in the AMG Linear Solver -- it could come from a meshing issue, for example.
6. If Step 2 resolves the problem, it could indicate a technical issue inside the AMG Linear Solver. Contact your CD-adapco support engineer.

Transport Equation Nomenclature

The following sections provide a guide to the terms used for:

- scalars;
- tensors or column vectors;
- Cartesian vectors;
- subscripts;
- superscripts; and
- field operators.

Scalars

a	coefficient arising from finite volume discretization
b	explicit contribution to discretized equation
D_f	face diffusion term
$\dot{m}_f$	face mass flow rate
V	cell volume
r	residual
t	time
α	gradient limiting factor
α_i	volume fraction of phase i
Γ	diffusion coefficient
Δ	difference between max or min allowable value and cell value
ϕ	scalar quantity
ω	under-relaxation factor
ρ	density
χ	void fraction

Tensors or Column Vectors

$\mathbf{A}$	coefficient matrix
$\mathbf{b}$	residual vector
$\mathbf{e}$	error vector
$\mathbf{x}$	solution vector

Cartesian Vectors

$\mathbf{a}$	face area vector
$\mathbf{ds}$	vector between cell centroids
$\mathbf{v}$	velocity
$\mathbf{v}_g$	grid velocity

$\mathbf{s}$	vector between cell face and cell centroids
$\mathbf{x}$	position vector
$\vec{\alpha}$	face metric quantity

Subscripts

0	cell-0 quantity
1	cell-1 quantity
c	cell quantity
f	face quantity
n	neighbor cells or coefficients in matrix
p	current cell or center coefficient matrix
r	reconstructed value

Superscripts

k	iteration k
n	time-step
u	unlimited (gradients)

Field Operators

- Gradient of a scalar quantity:

$$\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$$

- Scalar (inner, dot) product of two vectors, whose result is a scalar:

$$\mathbf{a}\cdot\mathbf{b} = a_i b_i$$

Transport Equation Bibliography

[434] Darwish M.S. and Moukalled F. 1994. "Normalized Variable and Space Formulation Methodology for High-Resolution Schemes", *Num. Heat Trans.*, Part B, **26** pp. 79-107.

[435] Ferziger, J.H. and Peric, M. 2002. "Computational Methods for Fluid Dynamics", 3rd rev. ed., Springer-Verlag, Berlin.

[436] Leonard, B. P. 1991. "The ULTIMATE conservative difference scheme applied to unsteady one-dimensional advection", *Comp. Methods in Applied Mechanics and Engineering*, **88** pp. 17-74.

[437] Murthy, J. Y. 2002. "Numerical Methods in Heat, Mass and Momentum Transfer", Draft Notes, School of Mechanical Engineering, Purdue University.

[438] Saad, Y. 2003. "Iterative Methods for Sparse Linear Systems", 2nd ed., Society for Industrial and Applied Mathematics

[439] Venkatakrishnan, V. 1994."On the Convergence of Limiters and Convergence to Steady State Solutions", AIAA-93-0880.

Transport Equation Solvers Reference

This section details the properties available to all solvers that control the transport equation models.

It also has descriptions of the URF ramping and the linear ramp.

Properties

There are no (non-expert) properties that are common to solvers controlling transport equation models.

Expert Properties

Unless you are thoroughly familiar with the discretization techniques that are used in STAR-CCM+, do not change any properties within the *Expert* category. The values in that category have been optimized for accuracy and performance.

The following expert properties are available with all solvers that control transport equation models.

Reconstruction Frozen	Sets whether reconstruction gradients are updated with each iteration.	
	Activated	Use reconstruction gradients from the last iteration in which they were updated. Temporary Storage Retained must be activated.
	Deactivated	Update reconstruction gradients with each iteration.
Reconstruction Zeroed	Allows you to set reconstruction gradients to zero, so that face values used for upwinding (Eqns. (2533) and (2537)) and for computing cell gradients (Eqns. (2586) and (2588)) are first-order estimates.	
	Activated	Set reconstruction gradients to zero at the next iteration. Gradients remain zeroed until you deactivate the checkbox again.
	Deactivated	Do not reset reconstruction gradients. If they had been zeroed, they are recomputed at the next iteration.

Temporary Storage Retained

Retains or clears temporary storage at the end of the iteration. Examples include the *ap* coefficients, the residual, the corrections, the reconstruction gradients, and the cell gradients. These quantities become available as field functions during subsequent iterations.

Activated — Retain temporary storage.

Deactivated — Do not retain temporary storage.

Solver Frozen

Prevents or allows the solver from updating the solution while iterating. In general, to freeze flow or turbulence, you should enable Temporary Storage Retained for both flow and turbulence for one iteration after freezing any solver.

Activated — The solver does not update the solution.

Deactivated — The solver updates the solution.

Setting URF (Under-Relaxation Factor) Ramping

The Under-Relaxation Factor Ramp node is a subnode of various solvers. You can activate or deactivate linear URF ramping in its properties.



Variations exist for this object, depending on the type of model you activate. For example, the Segregated Energy solver has separate ramps for the Fluid and Solid URFs.



Under-Relaxation Factor Ramp Properties

Ramp Method	Activates or deactivates the ramp for calculation.	
	No Ramp	Leaves off ramp calculation.
	Linear Ramp	Activates linear ramp value calculation. The Linear Ramp node becomes available in the Under-Relaxation Factor Ramp node.

Working with the Linear Ramp

The Linear Ramp node allows you to control linear URF ramp calculations with its properties.



Linear Ramp Properties

Start Iteration	Specifies the iteration on which to begin ramping, typically 1.
End Iteration	Specifies the iteration on which to stop ramping, typically within the first 10–100 iterations.
Initial Value	Uses a text entry to specify the value before ramping starts.

Using Tools

The Tools manager node, which is subject to common menu items, contains various objects used in support of the STAR-CCM+ CFD process.



This part of the documentation describes tools in STAR-CCM+ that facilitate your work. It includes:

- Using coordinate systems
- Using field functions
- Using reference frames
- Using tables
- Using engineering units

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Modeling physics -- the concepts, theory and implementation details of the models

- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format
- A guide to using STAR-View+, a viewer for three-dimensional STAR-CCM+ scenes

Using Coordinate Systems

The Coordinate Systems node, which has a pop-up menu, is the manager object for working with coordinate systems in the simulation tree. The coordinate systems manager can hold any number of coordinate systems up to the memory limits of the computer.



This section defines coordinate systems and more specifically explains the concepts of the laboratory coordinate system, local coordinate systems, and the block-mapped coordinate system.

Instructions are provided for:

- Creating a Local Coordinate System
- Creating a local coordinate system interactively
- Editing a local coordinate system interactively
- Transforming a local coordinate system
- Creating a block-mapped coordinate system

- Displaying a coordinate system in a scene
- Copying properties of a local coordinate system

What Are Coordinate Systems?

Coordinate systems are used to specify spatial vectors, such as position and direction for boundary and region values, initial conditions, and properties of location and direction for derived parts. Coordinate systems are also used to define the position of sketch planes in 3D-CAD. These coordinate systems can be exported for use with STAR-CCM+. For simple applications, the default laboratory coordinate system is sufficient. However, more complex situations can benefit from the definition and use of a local coordinate system or a block-mapped coordinate system.

A coordinate system can be renamed using the standard techniques. If a user-defined coordinate system is not being used, it can be deleted from an existing simulation.

What Is the Laboratory Coordinate System?

The laboratory coordinate system is the base system in which all other local coordinate systems are defined. The laboratory coordinate system is a Cartesian system. It exists in every STAR-CCM+ simulation by default, and is subject to common menu items. It is always defined and cannot be deleted.



The Laboratory node represents the laboratory coordinate system within the coordinate systems manager. It contains a child node, Local Coordinate Systems, that is used to manage all local coordinate systems (and their derived systems) within the simulation.

What Is a Local Coordinate System?

A local coordinate system can be defined in terms of cartesian, cylindrical, or spherical coordinates, and can take either the base laboratory coordinate system, or a previously defined cartesian local coordinate system, as the reference system. A local coordinate system can also be exported from a sketch plane in 3D-CAD for use in STAR-CCM+.

In terms of the object tree, the reference system acts as a parent node. Several layers of nesting are permissible, up to the memory limit of the computer. The Local Coordinate Systems manager node manages local coordinate systems and is a child node of the relevant reference coordinate system.



Local coordinate systems can be accessed by derived parts and reports, and also serve as scalar and vector field functions. Examples of how they can be used are:

- Initializing the velocity in a region to solid body rotation
- Providing a swirling inlet velocity profile

What Is a Cartesian Coordinate System?

A Cartesian coordinate system is one of three types of user-defined local coordinate systems in STAR-CCM+. It is represented in the simulation tree by a node that has its own properties and uses a common pop-up menu.



This type of coordinate system is defined by an origin and three right-handed orthogonal direction vectors for the X Axis, Y Axis, and Z Axis. Points are described using x, y, and z coordinates in the reference coordinate system (identified as the coordinate system containing the

Local Coordinate System manager node to which the current system belongs). Vectors are described using basis vectors aligned with the X, Y, and Z Axes of the reference coordinate system.



A variation of this coordinate system is available with 6-DOF motion: the body-fitted coordinate system.

What Is a Body-Fitted Coordinate System?

A body-fitted coordinate system is a variation of the Cartesian coordinate system. It is represented in the simulation tree by a node that has its own properties and is subject to common menu items.



Unlike the other local coordinate systems, the body-fitted one is only created automatically when you create a 6-DOF body with using 6-DOF motion in the simulation. A set of axes that stays "fixed" to a rigid body as it undergoes motion helps with reporting various quantities, especially forces, and moments in a body reference frame.

Properties of the Cartesian and Body-Fitted Local Coordinate Systems

You can use the properties that are listed below to adjust the specifications of this coordinate system.

X Axis Direction	A Cartesian unit vector that, along with the Origin, defines the x-axis.
Y Axis Direction	A Cartesian unit vector that, along with the Origin, defines the y-axis.
Z Axis Direction	A Cartesian unit vector that, along with the Origin, defines the z-axis (read-only).
Origin	The Cartesian position vector defining the origin of the x-, y-, and z-axes.
Reference System	The system to which the current coordinate system refers (read-only).

What Is a Cylindrical Coordinate System?

A cylindrical coordinate system is one of three types of user-defined local coordinate systems in STAR-CCM+. It is represented in the simulation tree by a node that has its own properties and uses a common pop-up menu.



This type of coordinate system is defined by a reference Cartesian coordinate system, with an origin and three right-handed orthogonal direction vectors for the radial, tangential, and axial axes. It therefore belongs within the Local Coordinate Systems manager node of either the Laboratory coordinate system, or a cartesian local coordinate system.

Points are described using r, θ and z coordinates. Vectors are described as the radial, tangential, and axial components of the vector in the local basis of that vector. This coordinate system is common in simulations that are rotating about an axis, or where a swirling velocity profile is used.



Cylindrical Coordinate System Properties

Radial Axis Direction	A Cartesian unit vector that, along with the Origin, defines the radial axis.
Tangential Axis Direction	A Cartesian unit vector that, along with the Origin, defines the tangential axis.
Axial Axis Direction	A Cartesian unit vector that, along with the Origin, defines the axial axis (read-only).
Origin	The Cartesian position vector defining the origin of the radial, tangential, and axial axes.
Reference System	The system to which the current coordinate system refers (read-only).

What Is a Spherical Coordinate System?

A spherical coordinate system is one of three types of user-defined local coordinate systems in STAR-CCM+. It is represented in the simulation tree by a node that has its own properties and uses a common pop-up menu.



This type of coordinate system is defined by an underlying Cartesian coordinate system, with an origin and three right-handed orthogonal direction vectors for the X Axis, Y Axis, and Z Axis. Points are described

using *r*, theta, and phi coordinates. Vectors are described as the radial, polar, and azimuthal components of the vector in the local basis of that vector. This coordinate system is useful for simulations in which there is spherical symmetry or rotation about a point.



Spherical Coordinate System Properties

Radial Axis Direction	A Cartesian unit vector that, along with the Origin, defines the radial axis.
Polar Axis Direction	A Cartesian unit vector that, along with the Origin, defines the polar axis.
Azimuthal Axis Direction	A Cartesian unit vector that, along with the Origin, defines the azimuthal axis (read-only).
Origin	The Cartesian position vector defining the origin of the radial, polar, and azimuthal axes.
Reference System	The system to which the current coordinate system refers (read-only).

Pop-Up Menu of Local Coordinate Systems

Transform System	Transforms the coordinate system in a manner selected in the submenu.	
	Translate	Activates a dialog for translating the coordinate system.
	Rotate	Activates a dialog for rotating the coordinate system.

See also: common menu items.

Creating a Local Coordinate System

To create a new local coordinate system, right-click on the Local Coordinate Systems manager node of either the base laboratory system, or a previously defined cartesian local coordinate system. Choose New followed by the type of local coordinate system you wish to create (Cartesian, Cylindrical, or Spherical).



If you have a scene display open, you are able to create the system interactively. Otherwise the new coordinate system appears in the Local Coordinate Systems manager node belonging to the parent system.



If you want to create a duplicate of an existing Cartesian coordinate system, you can do it using the technique for copying properties.

You can rename the coordinate system and set its axis and origin properties. You can also edit it interactively in a scene.

Creating a Local Coordinate System Interactively

When you have a scene display active, visual tools are activated when you choose to create a new local coordinate system. An in-place dialog appears in the edit tab of the simulation window.



The text boxes allow you to specify coordinates as text entries. Click the Renormalize button to apply your entries. When you click this button, STAR-CCM+ makes sure that the vectors are normalized and orthogonal, which typically results in some of the entered values changing.

Alternatively, you can click one of two mouse icons in the dialog.



The icon on the left activates the tool for specifying three points that define the XY plane.

1. The first click defines the origin of the coordinate system.
2. The vector from the first click to the second click defines the X axis.
3. The third click defines the direction of the Y axis. The actual Y axis needs to be perpendicular to the X axis.

The Z axis is defined as orthogonal to the X and Y axes. These mouse clicks are illustrated below.



The icon on the right activates the tool for specifying three points on a circle. You should click three points that are not co-linear, and that uniquely define a circle. The origin of the new coordinate system is at the center of this circle. The X axis points towards the first click, and the Z axis is perpendicular to

the plane of the circle with its direction determined by the order in which the points are defined. In the example that is shown below, the three clicks were made clockwise beginning from the first click.



You can adjust the display while this tool is active.

When you finish setting the parameters, click the Renormalize button to apply your entries. Click Create to add the coordinate system to the simulation tree, then click Close. Alternatively, click Close to cancel the creation process and discard all your settings.

Adjusting the Display While Creating a Coordinate System

To change the view in the scene display while the coordinate system creation tool is active, use normal mouse movements for changing a 3D display or changing a 2D display.

- For example, select the first of three points on the XY plane.



- To get access to an area that is behind the view, move the view as needed:
  - To rotate, click and hold the mouse button, then drag.

- To zoom, use the standard techniques with the middle mouse button.



- After changing the view as desired, release the mouse button.

The vector of the creation tool reappears.



- Continue defining the coordinate system, or make other adjustments to the view.

Editing a Local Coordinate System Interactively

If you have a scene display active, you can edit a local coordinate system interactively by right-clicking on it and selecting Edit in Current Scene.



This pop-up menu item is disabled (grayed out) if a scene display is not active.

The subsequent steps for modifying a local coordinate system in a scene are the same as those used for creating one.

Transforming a Local Coordinate System

STAR-CCM+ allows you to translate or rotate a local coordinate system to a new location. You access these transformations through a right-click menu on the local coordinate system node.

Translate

To translate the coordinate system, right-click on the node and select Transform System > Translate... in the pop-up menu.



In the *Translate Coordinate System* dialog, enter the component distances by which you wish to translate the local coordinate system. Distances can be either positive or negative, and can be entered in any valid length unit. Click the *Reference System* drop-down list for a choice of existing coordinate systems. Click Apply to complete your changes. Click Close to exit the dialog.



Rotate

To rotate the coordinate system, right-click on the node and select Transform System > Rotate... in the pop-up menu.



In the *Rotate Coordinate System* dialog, enter the angle by which you wish to rotate the local coordinate system, and the components of the vector to rotate it around. The *Reference System* drop-down list offers a choice of coordinate systems in the simulation. Click Apply to complete your changes. Click Close to exit the dialog.



What Is a Block-Mapped Coordinate System?

A block-mapped coordinate system is a type of user-defined coordinate system in STAR-CCM+ that defines points based on evenly distributed gradations between selected end surfaces in a topological cube. Each coordinate direction has a range of 0 to 1, where [0, 0, 0] defines one corner of the cube, and [1, 1, 1] defines the diagonally opposite corner of the cube. This type of coordinate system makes it possible to create iso surfaces that

follow the contours of the bounding surfaces, as would be desired, for example, in turbomachinery post-processing for generating blade-to-blade or meridional views.



The system is represented in the simulation tree by a node that has its own properties and is subject to common menu items. This coordinate system permits the creation of section surfaces along periodic directions, and has its own field functions.



The six subnodes in the Block-mapped node, known as block-mapped boundaries, represent the six sides of the topological cube. Each side should consist of one or more boundaries that you select when creating this coordinate system. All boundary edges in the group should be connected to one other boundary to form a closed cube.



Example Mapping

Umin = `Inlet`
Umax = `Outlet`
Vmin = `Shroud`
Vmax = `Hub`
Wmin = `Bottom`
Wmax = `Top`

Block-Mapped Coordinate System Properties

You can use the properties that are listed below to adjust the specifications of this coordinate system.

Region	A drop-down list of the regions in the simulation, from which you select a region to specify boundary groups.
Interpolation Points	A set of points along the three dimensions, representing the number of cells in the background grid along the Umin/Umax, Vmin/Vmax, and Wmin/Wmax directions. Increasing the number of points improves the accuracy with which the coordinate system follows the contours of the mapped surfaces.

Working with Field Functions of the Block-Mapped Coordinate System

Three field functions are used with the block-mapped coordinate system. When they are calculated, their nodes are added under the Tools > Field Functions manager node: Block-mapped: Umin - Umax, Block-mapped: Vmin - Vmax and Block-mapped: Wmin - Wmax. The names of these field functions can be altered by renaming the coordinate system and/or any of its six subnodes.



The values in the field functions are calculated when they are needed. A section surface can be defined by creating an iso surface of one of these field functions.

The field functions' values vary from about 0 to 1, where 0 is the first edge of a boundary group pair and 1 is the other edge. The maximum and minimum are not exactly 0.0 and 1.0 due to approximation errors.

The field functions are computed by creating a background grid and interpolating where each point in the actual mesh would be located. If this background mesh is too coarse, the resulting section surfaces will not be smooth. In this case, increase the resolution of the background grid by changing the coordinate system's Interpolation Points property.

What Are Block-Mapped Boundaries?

Block-mapped boundaries nodes represent grouped boundaries within a block-mapped coordinate system. These nodes have their own properties.



Block-Mapped Properties

Boundaries	Lists boundaries in this boundary group. Editing the property accesses the in-place object selector to add or remove boundaries.

Creating a Block-Mapped Coordinate System

To create a new block-mapped coordinate system, open the Tools node, right-click on the Coordinate Systems node, and select the New Block-mapped menu item.



The new coordinate system appears in the Coordinate Systems node. It contains six subnodes representing the six sides of the model.



You can then rename the coordinate system and set its properties.

Setting Properties of a Block-Mapped Coordinate System

Specify a region from which the boundaries are taken. To do this, make sure that you have the node of the coordinate system selected, and then choose the region from the drop-down list of the Region property.



Set up the boundary groups by selecting each of the block-mapped boundaries nodes and clicking the right half of its Boundaries property. Then specify the boundaries in the object selector.

Displaying a Coordinate System in a Scene

When a coordinate system is in a scene, it functions as an alternate set of axes. The original set of axes is itself a laboratory coordinate system, though it is relative to the screen rather than to the origin in the space being modeled.

You can display a coordinate system in a scene using one of two techniques: drag-and-drop or object selection.

Displaying a Coordinate System Using Drag-and-Drop

To display a coordinate system in a scene with the drag-and-drop technique, drag the node of the particular coordinate system onto the scene display, and release the mouse button when your pointer is within the display.

If you do not have the scene open, you can also drag the coordinate system node onto the scene node.

The axes of the coordinate system appear in the display. In the following screenshot, a Cartesian coordinate system has been added to the scene.



Displaying a Coordinate System Using Object Selection

To display a coordinate system in a scene by selecting it as an object, select the node of the particular scene.



In the *Properties* window, click in the right half of the Coordinate Systems property.



This activates a standard object selection dialog for selecting coordinate systems for the scene.

The axes of the coordinate system appears in the display. In the following screenshot, a Cartesian coordinate system has been added to the scene.



Copying Properties of Local Coordinate Systems

When properties are copied from one local coordinate system to another, all of the editable properties are applied. The read-only Reference System remains the same.

It is not possible to copy properties between a user-defined coordinate system and the laboratory coordinate system, which is pre-defined.

You can use these copying techniques to create a duplicate of a local coordinate system. Simply apply the Paste command to the Coordinate Systems manager node (or make it the destination of your drag-and-drop operation).

A copy of the coordinate system appears in the tree.

Pop-Up Menu of the Coordinate Systems Manager

These pop-up menu items are available when you right-click the Coordinate Systems manager node.

New Block-Mapped	Executes the create block-mapped coordinate system action to add a new block-mapped coordinate system to the Coordinate Systems node.

See also: common menu items.

Using Field Functions

The Field Functions node, which has its own context menu, is the manager object for working with field functions in the simulation tree. The field functions manager can hold any number of field functions up to the memory limits of the computer.



This section describes field functions, explains what the uses of field functions are, and specifies how they are used in the GUI.

Instructions for working with user field functions are provided, specifically the following actions:

- Creating a user field function
- Copying a user field function
- Ignoring boundary values for field function data
- Interpolating field functions using inverse distance weighting

Reference sections are also provided for:

- The syntax of user field functions
- Some examples of user field functions

What Are Field Functions?

Field functions provide a mechanism by which fields (scalar or vector data that are evaluated at cells, vertices, or boundary faces) can be viewed and defined in STAR-CCM+. Available field function nodes are viewed in the Field Functions manager node under Tools in the object tree.



All field function nodes are subject to common menu items, but their properties vary by type. The most common set of properties is for the primitive type.

Once the field function has been made available, it is provided in scalar and/or vector drop-down lists for any operation that requires it: notably visualization, XY plotting, reporting, monitoring, and creating table data. In addition, the following objects are typically included in these lists:

- Field function monitors
- Field functions corresponding to reports
- Cell sets

A field function remains active in the GUI as long as any other object is still accessing it.

The field functions are classified based on the following:

- Distinctions between system field functions and user field functions
- Distinctions among scalar, vector, array, and position field functions

Field functions can only be accessed on a part if the data used to define the field function exists on that part. For the data to exist it must be allocated on the specific part and must have been initialized. For example, wall shear stress is a quantity that is stored only on walls in viscous-flow simulations.

Therefore the wall shear stress field function, and related field function quantities such as Wall y^+, can only be displayed on wall boundaries, and only after the solution has been initialized.

Primitive Field Function Properties

Type	The type of field function: scalar, vector, array, or position (read-only).
Function Name	The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.
Dimensions	The physical dimensions of this field function (read-only).

Primitive Field Function Expert Properties

Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values.	
	Activated	Interpolate node values using inverse distance weighting.
	Deactivated	Use nearest neighbor as node values.

What Are System Field Functions?

System field functions are predefined field functions that active models create. System field functions share a particular set of properties. Every system field function has a unique name and can be of a scalar, vector, array, or position type.

System Field Function Properties

Type	The type of field function: scalar, vector, array, or position (read-only).
Dimensions	The physical dimensions of this field function (read-only).
Function Name	The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.
Definition	The definition of the field function in the STAR-CCM+ field function syntax (read-only).

Expert Properties

Assembly Code	An unambiguous tree-based (but humanly readable) form of the field function formula (read-only). This property is empty until the field is used, for example in a contour plot.
Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values. Activated: Interpolate node values using inverse distance weighting. Deactivated: Use nearest neighbor as node values.

What Are User Field Functions?

User field functions are field functions that you create to access field data, and share a particular set of properties. They can be of a scalar, vector, array, or position type. They are created manually and are defined in terms of other field functions.

Similar to other objects, a user field function can be renamed within STAR-CCM+. User field functions can be deleted from an existing simulation as long as they are not in use.

User-Defined Field Function Properties

Type	The type of field function: scalar, vector, array, or position (read-only).
Dimensions	The physical dimensions of this field function, which you can customize.
Function Name	The name of this field function. Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree. If you inadvertently use a function name that another field function uses, STAR-CCM+ appends an underscore (_) and a digit to what you typed. The result is a name that is still unique.
Definition	The formulation of this field function using the special STAR-CCM+ field function syntax. Entering an invalid definition causes the error: `Error compiling Field Function: parse error` to be displayed in the *Output* window, and the Assembly Code property to be populated with `error`.

Expert Properties

Ignore Boundary Values	Determines whether the field functions are computed based on the data that are stored in the cells that are next to the boundaries rather than the field data stored on boundaries (if available). Activated — Use adjacent cell values. Deactivated — Use boundary values.

Assembly Code	An unambiguous tree-based (but humanly readable) form of the field function formula (read-only). Empty until the Definition is entered or until the field is used, for example in a contour plot. The assembly code is updated each time that you change the Definition and press `<Enter>`. This property can be useful to make sure that STAR-CCM+ interprets your definition as expected (for example, operator precedence).
Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values. Activated: Interpolate node values using inverse distance weighting. Deactivated: Use nearest neighbor as node values.

What Are Scalar Field Functions?

Scalar field functions are field functions that correspond to just one number (zero, positive or negative) at a specific location. Six of them in particular have different types of properties:

- Heat transfer coefficient
- Nusselt number
- Pressure coefficient
- Skin friction coefficient
- Temperature coefficient
- Total pressure coefficient

Heat Transfer Coefficient Field Function Properties

Reference Temperature	The reference temperature, which can be specified with a text entry.
Type	The type of field function (scalar; read-only).
Dimensions	The physical dimensions of this field function (read-only).
Function Name	The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.

Expert Properties

Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values.	
	Activated	Interpolate node values using inverse distance weighting.
	Deactivated	Use nearest neighbor as node values.

Nusselt Number Field Function Properties

Dimensions	The physical dimensions of this field function (read-only).
Reference Length	The reference length, which can be specified with a text entry.
Reference Temperature	The reference temperature, which can be specified with a text entry.
Reference Thermal Conductivity	The reference thermal conductivity, which can be specified with a text entry.
Type	The type of field function (scalar; read-only).

Function Name — The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.

Expert Properties

Inverse Distance Weight — Whether to use inverse distance weighting to interpolate node values.

- Activated — Interpolate node values using inverse distance weighting.
- Deactivated — Use nearest neighbor as node values.

Pressure Coefficient Field Function Properties

Dimensions — The physical dimensions of this field function (read-only).

Reference Density — The reference density, which can be specified with a text entry.

Reference Pressure — The reference pressure, which can be specified with a text entry.

Reference Velocity — The reference velocity, which can be specified with a text entry.

Type — The type of field function (scalar; read-only).

Function Name — The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.

Expert Properties

Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values.	
	Activated	Interpolate node values using inverse distance weighting.
	Deactivated	Use nearest neighbor as node values.

Skin Friction Coefficient Field Function Properties

Dimensions	The physical dimensions of this field function (read-only).
Reference Density	The reference density, which can be specified with a text entry.
Reference Velocity	The reference velocity, which can be specified with a text entry.
Type	The type of field function (scalar; read-only).
Function Name	The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.

Expert Properties

Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values.	
	Activated	Interpolate node values using inverse distance weighting.
	Deactivated	Use nearest neighbor as node values.

Temperature Coefficient Field Function Properties

Dimensions	The physical dimensions of this field function (read-only).
Reference Temperature Delta	The reference temperature difference, which can be specified with a text entry.
Reference Temperature	The reference temperature, which can be specified with a text entry.
Type	The type of field function (scalar; read-only).
Function Name	The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.

Expert Properties

Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values.	
	Activated	Interpolate node values using inverse distance weighting.
	Deactivated	Use nearest neighbor as node values.

Total Pressure Coefficient Field Function Properties

Dimensions	The physical dimensions of this field function (read-only).
Reference Density	The reference density, which can be specified with a text entry.
Reference Pressure	The reference pressure, which can be specified with a text entry.

Reference Velocity	The reference velocity, which can be specified with a text entry.
Type	The type of field function (scalar; read-only).
Function Name	The name of this field function (read-only). Use this property in the definition of other field function definitions, not the node name. Use this property to remind you of the actual field function if you have renamed the node in the tree.

Expert Properties

Inverse Distance Weight	Whether to use inverse distance weighting to interpolate node values.	
	Activated	Interpolate node values using inverse distance weighting.
	Deactivated	Use nearest neighbor as node values.

What Are Vector Field Functions?

Vector field functions define a vector that can be expressed in different coordinate systems. They have the following components when presented in the scalar selection list, depending on the coordinate system with which they are associated:

- `Magnitude`
- `Cartesian` or `Laboratory`:
  - `i`
  - `j`
  - `k`
- `Cylindrical`:
  - `Radial`
  - `Tangential`
  - `Axial`
- `Spherical`:
  - `Radial`
  - `Polar`

- Azimuthal

What Are Array Field Functions?

Array field functions are field functions that are defined as three numbers at each location. No coordinate system transformations are possible. They have the following components when presented in the scalar selection list:

- Magnitude
- [X-component value]
- [Y-component value]
- [Z-component value]

What Are Position Field Functions?

Position field functions are field functions that are defined as a position at each location. They can be transformed into different coordinate systems using position transformations. They have the following components when presented in the scalar selection list, depending on the coordinate system with which they are associated:

- Magnitude
- Cartesian or Laboratory:
  - X
  - Y
  - Z
- Cylindrical:
  - r
  - theta
  - z
- Spherical:
  - r
  - theta
  - phi

What Are the Uses of Field Functions?

The most common use of field functions is examining the solution. You can, for example, display pressure on the wall of a body. In this case, `Pressure` is a field function that the system defines.

A more advanced use of field functions is the specification of boundary and region values as well as initial conditions. In this case, you typically define a new field function and apply it. An example would be writing a function for a parabolic velocity profile that is a function of the y-coordinate and using it for the inlet velocity at a boundary.

How Do Field Functions Work in the GUI?

Field functions are typically activated in STAR-CCM+ by continuum models. However, when objects such as reports, monitors, or displayers access field functions, that connection forces the field function to remain present in the simulation even after you deactivate the related model. This is because STAR-CCM+ considers a field function to be in use as long as any other object in the simulation accesses the field function. At the same time, the field function cannot be expected to be computed correctly if the activating model is no longer active.

Therefore when you are changing continuum models and/or choosing to delete a user field function, take care to examine your simulation for objects that keep an obsolete field function active. Failure to do so can lead to the inability to delete user field functions, as well as confusion as to which visible field functions are actually being calculated.

Creating a User Field Function

To create a user field function, right-click the Tools > Field Functions node and select New followed by the type of field function you wish to create.



A new field function appears in the object tree. The type that you selected in the context menu is in the read-only Type property of the field function.



Initially the node has a default name such as User Field Function 1. However, you can rename the field function.

A new field function also appears automatically with the creation of a report. You cannot rename such field functions directly; their names contain the names of their respective reports.

To define the user field function, do either of the following:

- In the Definition expert property, enter the definition directly.
- For longer definitions, use the definition dialog. To activate the dialog, click the property customizer (...).

An example definition, using the dialog, is shown below.



In the bottom right panel, a pencil icon appears alongside the name of the field function that you are defining.

- To see the definition of another user field function in the simulation, hover over it with the pointer. The definition appears in a tooltip.



- In the expression text panel at the top of the dialog, define the field function using any combination of the following techniques:
  - Typing the definition directly.
  - Inserting expressions interactively from the two bottom panels.
  - Using common editing techniques, such as copying and pasting:
    - Apply standard keyboard commands such as `<Ctrl><Z>` (Undo), `<Ctrl><Y>` (Redo, also with `<Ctrl><Shift><Z>`), and `<Ctrl><C>` (Copy).
    - To access the commands from a context menu, right-click in the expression text panel.



Note: The syntax that is used in the definition for user field functions is a subset of the C programming language.

Space between tokens is not significant, so you can use spacing to make the definition easier to read. You can even include line breaks. If you indent on one line, subsequent lines are automatically indented. To maintain readability of deeply nested parenthetical expressions, use this indenting capability.

The expression text panel uses colors to help guide you through the editing process:

- Green indicates a syntactically correct reference to an existing field function:



- Parentheses appear in yellow to indicate the matching parenthesis:



This color technique also applies to square brackets.

- Red indicates an incorrect entry:



However, if a field function reference has red text (but no black text or red background) the color represents a syntactically correct reference to a field function that does not yet exist.

- If you click OK when the definition includes an error, dialogs appear with error messages. The messages are repeated in the *Output* window.



When such an error occurs, examine and debug the definition until it is accepted without any error notifications.

Note: Field function definitions are always done in SI units, independent of the units that you use for input or display.

Inserting Expressions Interactively

To add complete expressions (field functions or system functions) into the expression text panel, make selections from the two bottom panels of the dialog.

To insert a field function expression:

- Right-click on one of the names in the bottom right panel and select Insert.



- Alternatively you can double-click the row, or highlight the row and press the spacebar.

  When you insert a field function, the field function reference replaces any selected text in the expression text panel.

  For details on including field functions within a field function, see Referencing Field Functions.

To insert a system function expression:

- Open any of the objects with the handle (⊶) in the bottom left panel. Right-click one of the child objects and select Insert.



- Alternatively you can double-click the row, or highlight the row and press the spacebar.

  For detailed information about system functions, see the following sections:

  - Trigonometric Functions
  - Vector Functions
  - Miscellaneous Functions
  - Working with Table Values Through Field Functions

  When you insert a system function, any selected text becomes enclosed by the function. For example, if you have `$$Velocity` selected and you insert the system function `mag($$u)`, `mag($$Velocity)` appears in the expression text panel. If there is no selection, the system function gets inserted at the cursor location with no parameters, for example `mag()`.

Previewing Expressions

You can preview the calculations of your user field functions without actually running the simulation. Scalar values appear in a plot which gets updated instantly with your changes. The output plot is calculated by substituting the values of each of the independent variables and sweeping through the range of values provided for the dependent variable.

An example of a plot preview is shown below. In this example, the output is `acos(1) = 0`.



Note: Plot previews are available for scalar functions only.

To preview a user field function:

- In the definition dialog, click the Preview tab.
- Select the Dependent Variable and enter the Range and Values to use in the preview.

  For a vector component, as you select the radio button, the right column value appears for that component. You can plot a range for only one of the vector components at a time.

- Within the table of Independent Variables, enter values for the field function to use in generating the preview plot.



Note: If you make errors, such as dividing by zero, STAR-CCM+ informs you by stating the nature of the error in the status field at the bottom of the definition dialog.

Copying a User Field Function

There are several ways to copy a user field function:

- Copy the user field function to the clipboard and paste it into the Field Functions node.
- Drag and drop the user field function onto the Field Functions node.
- Copy the properties of the user field function onto another.

Copying Properties of a User Field Function

When properties are copied from one user field function to another, the following properties and expert properties are changed:

- Dimensions
- Type
- Assembly code
- Definition
- Ignore Boundary Values

Change the Function Name expert property by entering text or renaming it.

It is not possible to copy properties between other types of field functions and user field functions.

Ignoring Boundary Values

Field functions access data on specific parts that are defined in the simulation. When field functions are accessed on boundary parts, two rules are used by default:

- If the data is stored on the boundary, the boundary data are accessed.
- If the data is not stored on the boundary, the cell data from the cells in the adjacent region are accessed.

In specific situations where data is stored on the boundary and in the adjacent cells, access the adjacent cell data:

1. Create a user field function.
2. Define the new user field function such that it is defined in terms of an existing field function that is stored in both cells and on boundaries, but the values in each location differ. An example is velocity magnitude, which has the wall velocity values stored in wall boundaries, and the wall-adjacent fluid velocity in the cells. In this case, the field function is defined as:
   `$$Velocity.mag()`

3. Activate Ignore Boundary Values in the Expert category.



If this user field function is accessed on a boundary, it contains a scalar value corresponding to the velocity magnitude in the cell that is next to the boundary.

If the field function is applied on an interface, and the Ignore Boundary Values option is activated, the field function is computed using data from the cell next to the boundary. This data is identified using the Boundary-0 property of the interface.

Interpolating Field Functions Using Inverse Distance Weighting

In STAR-CCM+, most variables are solved and stored at cell centers. To smooth these values, STAR-CCM+ uses a first order method by default to interpolate the data that are stored in cell centers to vertices. When using a field function in a derived part, the value of the field function at the extraction points is interpolated from these vertex values.

You can interpolate field functions from cells centers to vertices using a higher-order method:

1. Create a user field function.
2. Define the new user field function.
3. Activate Inverse Distance Weight in the Expert category.



When Inverse Distance Weight is activated, STAR-CCM+ interpolates vertex values from the nearest cell centers using inverse distance weighting.



The field function ϕ_v at a node is calculated as

$$\phi_v = \sum_i \frac{1}{\delta_i}\phi_i \tag{2633}$$

where δ_i is the distance between the node and the centroid i (cell center for a cell value and face center for a face value) and ϕ_i is the variable that is stored at the centroid.

Field Function Programming Reference

This section details the syntax that must be followed when writing your own user field functions. Specific topics include:

- Working with table values through field functions
- Referencing field functions

- Arithmetic operators
- Logical operators
- Trigonometric functions
- Vector functions
- Other miscellaneous functions
- Conditional expression
- Coordinate system averaging
- Alternate value functions
- Working with units in field functions

Working with Table Values Through Field Functions

You can access table data from user field functions within STAR-CCM+ and use these values for post-processing. Column table data can be interpolated as functions of other scalar field functions (for example `$Time`), and position tables can be spatially interpolated and accessed as fields. Both scalar and vector tables can be interpolated.

Interpolating Data as a Function of Another Table Column

You can use field functions to access table data as a function of another table column. The syntax for doing this involves the use of a table file's name and the names of its columns. For this technique to work properly, column names must not include any units associated with that column.

Using the field function table interpolation syntax to generate a scalar expression:

```
interpolateTable(@Table("<tablename>"), "<x-column>",
{LINEAR | STEP | SPLINE}, "<phi-column>",
<scalar_expression>)
```

Using the field function table interpolation syntax to generate a vector expression:

```
interpolateTable(@Table("<tablename>"), "<x-column>",
{LINEAR | STEP}, "<phi_0-column>", "<phi_1-column>",
"<phi_2-column>", <scalar_expression>)
```

where:

- `<tablename>` is a table name from the Tools > Tables node in the object tree.

- `<x-column>` is the name of a table column to be used as the `x` data for the interpolation.
- `<scalar_expression>` is a scalar field function that is used as the `x` value for the interpolation of `y`.
- For scalar functions:
  - `<phi-column>` is the table column of scalar values to be interpolated as a function of `x`.
- For vector functions:
  - `<phi_0-column>` is the table column of 0th components to be interpolated as a function of `x`.
  - `<phi_1-column>` is the table column of first components to be interpolated as a function of `x`.
  - `<phi_2-column>` is the table column of second components to be interpolated as a function of `x`.

  The value that is returned from a vector expression is an array of the interpolated `0`, `1`, `2` component columns.

The following interpolation methods are supported:

- `LINEAR` -- specifies a linear interpolation between `y` table values.
- `STEP` -- uses `y` table values at the largest `x` table value that is less than the value of `$scalar_expression`.
- `SPLINE` -- uses B-spline interpolation between `y` table values.

For example, assuming a table `"DataTable"` with columns `"Time"` and `"Phi"` the following:

```
interpolateTable(@Table("DataTable"), "Time", LINEAR,
"Phi", $Time)
```

linearly interpolates `"Phi"` as a function of `"Time"`, and provides the interpolated value at the simulation time.

Interpolating Data as a Function of Position

You can use field functions to access table data as a function of position. The syntax for doing this involves the use of a table file's name and the names of its columns. For this technique to work properly, column names must not include any units associated with that column.

Using the field function table position interpolation syntax to generate a scalar expression:

```
interpolatePositionTable(@Table("<tablename>"),
"<phi-column>")
```

or, with a coordinate system:

```
interpolatePositionTable(@Table("<tablename>"),
@CoordinateSystem("<coordinatesystem>"), "<phi-column>")
```

or, with a nested coordinate system:

```
interpolatePositionTable(@Table("<tablename>"),
@CoordinateSystem("<coordinatesystem>.<coordinatesystem>.
<coordinatesystem>"), "<phi-column>")
```

Using the field function table position interpolation syntax to generate a vector expression:

```
interpolatePositionTable(@Table("<tablename>"),
"<phi_0-column>","<phi_1-column>","<phi_2-column>")
```

or, with a coordinate system:

```
interpolatePositionTable(@Table("<tablename>"),
@CoordinateSystem("<coordinatesystem>"),
"<phi_0-column>","<phi_1-column>","<phi_2-column>")
```

or, with a nested coordinate system:

```
interpolatePositionTable(@Table("<tablename>"),
@CoordinateSystem("<coordinatesystem>.<coordinatesystem>.
<coordinatesystem>"),
"<phi_0-column>","<phi_1-column>","<phi_2-column>")
```

where:

- `<coordinatesystem>` is the name of a coordinate system in the Tools > Coordinate Systems node in the object tree. Multiple instances of this name, which dots separate, represent layers of nesting. When a coordinate system is not specified, the laboratory system is used and assumed to be Cartesian.
- `<tablename>` is a table in the Tools > Tables node, which contains three position columns, one for each of the coordinates corresponding to the type of `<coordinatesystem>` specified. Column names must be those that the coordinate system dictates, that is, Cartesian (`"X"`, `"Y"`, `"Z"`), cylindrical (`"r"`, `"theta"`, `"z"`) or spherical (`"r"`, `"theta"`, `"phi"`).
- For scalar functions:
  - `<phi-column>` is the table column for spatially interpolation using the nearest neighbor, corresponding to the position coordinate columns.
- For vector functions:
  - `<phi_0-column>`, `<phi_1-column>`, and `<phi_2-column>` represent the column names of the three values of the *components of the vector* in `<coordinatesystem>`.

The points are considered unordered in the file; each face or cell is assigned the value from the closest point in the table.

Interpolating Direction from Parts

The interpolateDirection field function generates a field of direction vectors based on line curves or surfaces that you specify. Having a field of direction vectors can be useful when you model anisotropic behavior in materials. The field function returns a field of orientation vectors, interpolated from the lines or surfaces to every cell.

The syntax for this field function is:

```
interpolateDirection(@Part("<partname>"),
@Representation("<representation>"),FORWARD)
```

or, with a coordinate system:

```
interpolateDirection(@Part("<partname>"),
@Representation("<representation>"),
@CoordinateSystem("<coordinatesystem>"),FORWARD)
```

where:

- `<coordinatesystem>` is the name of a coordinate system in the Tools > Coordinate Systems node in the object tree. The direction vectors are returned based on this coordinate system. When a coordinate system is not specified, the laboratory system is used and assumed to be Cartesian.
- `<partname>` is a surface or a line curve. For a surface, the direction field is given using the normals of that part. For a line curve, the direction field is given using the line segment directions. If using a surface or a curve name of a part, the syntax is `<partname.surface/curve name>`. For a region, it is `<regionname.boundary/feature curve name>`.
- `<representation>` is the representation that is used to extract the part. It is set in the Representation node in the object tree.
- `FORWARD` is used to set the orientation of the vectors based on the direction of the normals of the input surface or the direction of the line segment of the input line. Use `BACKWARD` to reverse the direction vectors.

The syntax for this field function with multiple parts is:

```
interpolateDirection(@Part("<partname1>::<partname2>::<pa
rtname3>"), @Representation("<representation>"),FORWARD)
```

or, with a coordinate system:

```
interpolateDirection(@Part("<partname1>::<partname2>::<pa
rtname3>"), @Representation("<representation>"),
@CoordinateSystem("<coordinatesystem>"),FORWARD)
```

The parts are delimited using `::`

The following is an example of using this field function. The left image is a tangential field function that uses a line curve input. The right image shows a radial field function that uses a surface input for the given geometry.



Referencing Field Functions

When defining an expression or a field function, you can reference other field functions within it by using the name from the FunctionName property prefixed with `$`, such as `$MachNumber`.

If the name contains any non-alphanumeric characters apart from '_', Enclose the field function name in curly braces, '{}'. Also use curly braces if the name begins with a number, or if you are referencing object names that include characters from non-Western alphabets. Some examples of referencing such names are:

- ${Five Things}
- ${Temp@X}

- ${3rdInletMassflow}
- ${test-Temperature}

Field functions used to define expressions cannot be a function of space; they have to evaluate to a single number as the expressions are constants in space. Variables such as `$Time`, `$TimeStep`, and `$Iteration`, as well as the field functions associated with reports, are valid. Variables such as `$Temperature` or `$Pressure` have different values at every different point in space and cannot be used in expressions.

This section includes more specific information as follows:

- Referencing scalar field functions
- Referencing vector field functions
- Referencing components of a vector field function
- Referencing local coordinate systems
- Referencing reference frames

Referencing Scalar Field Functions

To reference an existing scalar field function, the variable name in the new definition is the existing function name (as defined in the Function Name property of the *Expert* category) prepended by a dollar ($) symbol.

For example, to reference the field function for temperature, first look at the field function name:



In this case, the name is Temperature. The variable name is thus:

```
$Temperature
```

Since the names are case-sensitive, they must be entered in the field function definition exactly as they are shown in the function name.

Referencing Vector Field Functions

To reference an existing vector, array, or position field function, the variable name in the new definition is the existing function name (as defined in the Function Name property of the *Expert* category) prepended by two dollar symbols (`$$`).

For example, to reference the field function for velocity, first look at the field function name:



In this case, the name is Velocity. The variable name is thus:

```
$$Velocity
```

Since the names are case-sensitive, they must be entered in the field function definition exactly as they are shown in the function name.

Referencing Components of a Vector Field Function

If you are defining a new user field function that uses components of a vector, array, or position field function, your function is clearer and executes a bit more quickly if you write the function directly in terms of the original vector function, rather than going through the derived scalar.

For example, suppose that you want to define a new scalar function that is equal to the sum of the Cartesian velocity components. You could define it directly in terms of the velocity vector:

```
$$Velocity[0] + $$Velocity[1] + $$Velocity[2]
```

Referencing Local Coordinate Systems

For position and vector field functions, user field functions support a syntax to represent Position and Centroid in different local coordinate systems. For example:

```
$$Centroid(@CoordinateSystem("Laboratory.Coordinate
System 1"))
```

This yields a vector of `Centroid` in `Coordinate System 1`. For positions, components of the vector are in the units of the coordinate system specified:

- If the coordinate system is Cartesian, `$$Centroid(@CoordinateSystem("Laboratory.CartesianSystem"))` is `[X,Y,Z]`.
- If the coordinate system is cylindrical, `$$Centroid(@CoordinateSystem("Laboratory.CylindricalSystem"))` is `[radius, theta, z]`.
- If the coordinate system is spherical, `$$Centroid(@CoordinateSystem("Laboratory.SphericalSystem"))` is `[radius, theta, phi]`.

For cases where local coordinate systems have one or more levels of nesting, use the qualifier `@CoordinateSystem` and separate each coordinate system from its parent using the '.' character. As an example, consider a field function that gives the centroid in the Spherical 1 local coordinate system for the simulation tree below:



The field function would be written as:

```
$$Centroid(@CoordinateSystem("Laboratory.Initial ship
orientation.Spherical 1"))
```

The Boat-CSys coordinate system would be referred to as:

```
$$Centroid(@CoordinateSystem("Laboratory.Boat-CSys"))
```

Referencing Reference Frames

For position and vector field functions, user field functions support a syntax to represent Position and Centroid in different reference frames. For example:

```
$$Velocity(@ReferenceFrame("Rotating.Rotating Relative
Frame"))
```

This yields a vector of `Velocity` in `Rotating Relative Frame`. For positions, components of the vector are in the units of the coordinate system specified in the reference frame.

Arithmetic Operators

Operator	Function
+	plus
-	minus
*	multiply
/	divide

Logical Operators

Operator	Function
==	equal to
!=	not equal to
>	greater than
<	less than
>=	greater than or equal to
<=	less than or equal to
&&	and
||	or

Trigonometric Functions

Operator	Function
acos($x)	arccosine of x
asin($x)	arcsine of x
atan($x)	arctangent of x
atan2($y,$x)	arctangent of y/x
cos($x)	cosine of x

cosh($x)	hyperbolic cosine of x
sin($x)	sine of x
sinh($x)	hyperbolic sine of x
tan($x)	tangent of x
tanh($x)	hyperbolic tangent of x

Vector Functions

These operators work with vector, array, and position field functions.

Operator	Function
mag($$u)	vector magnitude
mag2($$u)	vector magnitude squared
unit($$u)	unit vector
dot($$u,$$v)	dot product of u and v
cross($$u,$$v)	cross product of u and v
curl($$v)	curl of vector v In vector field notation, the curl of vector $\mathbf{v}$ is defined as: $\vec{\nabla} \times \mathbf{v} = \left(\frac{\partial v_z}{\partial y} - \frac{\partial v_y}{\partial z}\right)\mathbf{i} + \left(\frac{\partial v_x}{\partial z} - \frac{\partial v_z}{\partial x}\right)\mathbf{j} + \left(\frac{\partial v_y}{\partial x} - \frac{\partial v_x}{\partial y}\right)\mathbf{k}$
div($$v)	divergence of vector v returned as a scalar In vector field notation, the divergence of vector $\mathbf{v}$ is defined as: $\nabla \cdot \mathbf{v} = \frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} + \frac{\partial v_z}{\partial z}$

Miscellaneous Functions

Operator	Function

ceil($x)	smallest integer not less than x, as a real
exp($x)	exponential function, e to the power x
abs($x)	absolute value of x
floor($x)	largest integer not greater than x, as a real
log($x)	base e logarithm of x, natural log
log10($x)	base 10 logarithm of x
sqrt($x)	square root of x
pow($x,$y)	x raised to the power of y
fmod($x,$y)	x modulus y; the floating point remainder of x/y
mod($x,$y)	equivalent to fmod($x,$y)
min($x,$y)	the minimum of x and y
max($x,$y)	the maximum of x and y
grad($phi)	gradient of scalar phi returned as a vector The gradient of scalar ϕ is defined as: $\nabla\phi = \frac{\partial\phi}{\partial x}\mathbf{i} + \frac{\partial\phi}{\partial y}\mathbf{j} + \frac{\partial\phi}{\partial z}\mathbf{k}$ Values calculated using this approach are not necessarily the same as gradients returned from the solvers. This is due to different approaches used in calculating the gradients. In the solvers, a two-pass Gauss algorithm is used, whereas the function provided here uses a least squares approach. Solvers also have more information on boundary values than is available to this function.

Conditional Expression

The ternary operator "?:" defines a condition operation similar to the C language. For example:

```
(a > b) ? a : b
```

is equivalent to

```
z = max(a, b)
```

You can also use parentheses to nest conditional expressions. For example, to specify the following conditions:

```
if position > p1 then...
else if position < p2 then...
else...
```

You can use the following syntax:

```
($$Centroid[1] > p1) ? 3000 : (($$Centroid[1] < p2) ? -30:
0)
```

Coordinate System Averaging

Coordinate system averaging combines the three-dimensional STAR-CCM+ solution results with a typical two-dimensional meridional view. While the meridional view is commonly used in turbomachinery, it can also be applied to certain other applications that involve axisymmetric regions.

This technique can provide an average of results in the θ direction, although you can choose any direction for projecting this average. The calculation is defined as:

$$\overline{\phi_\theta} = \frac{\int_{\theta_1}^{\theta_2} \phi \partial\theta}{\int_{\theta_1}^{\theta_2} \partial\theta} = \frac{1}{(\theta_2 - \theta_1)} \int_{\theta_1}^{\theta_2} \phi \partial\theta \tag{2634}$$

This definition can be implemented in STAR-CCM+ with a `csavg()` field function, in a coordinate system and direction that you specify:

```
csavg($Pressure, @CoordinateSystem("Block-Mapped 1"), 2,
0, 20, 20)
```

where, in this example:

- `$Pressure` is the scalar being defined. (A vector can also be defined in this type of field function.)
- `CoordinateSystem("Block-Mapped 1")` is the (block-mapped) coordinate system being used for the averaging.
- `2` is the averaging direction.

- `0` is the slice direction.
- `20, 20` are the number of slices and the integrations per slice, respectively.

Using a two-dimensional grid, the approach involves projecting region vertices to the grid and interpolating the average to the vertices.

Examples in this section are from the Rotor 37 case.

This scalar scene illustrates projection of vertices to the grid:



- The blue-gray plane shows a constant z slice.
- Black lines are averaging lines of constant r.
- The results from multiple slices appear in a two-dimensional average mesh.

The following screenshot shows these values in a plot:



- Red dots are the resulting average mesh.
- Blue "x" symbols are the (*z*, *r*) vertices from the region projected onto this mesh.
- Green dots are extrapolated (halo) nodes from the average mesh.

For comparison, look at the scalar-scene representation of the average mesh.



The scalar scene in the following screenshot shows the Rotor 37 case using the following user-defined field function:

```
csavg($AbsoluteTotalPressure,
@CoordinateSystem("Cylindrical 1"), 1, 2, 20, 60)
```



This field function takes the average of Absolute Total Pressure, using the Cylindrical 1 coordinate system in the θ direction, and slicing in the *z* direction. It uses a *20 x 60* averaging mesh -- that is, this mesh has 20 *z*-slices and, for each *z*-slice, 60 averaging lines (or *r*-slices).

Alternate Value Functions

If the value for a scalar or vector expression on a part is not available or computable, you can provide an alternate value for that expression.

Normally STAR-CCM+ requires all terms in an expression to exist for the expression to be computable. For example, in the following expression:

```
$Pressure * $Volume
```

if `Pressure`, `Volume`, or both do not exist for a part, this expression is not computable on that part.

To provide an alternate value, use the `alternateValue()` operator:

```
alternateValue($Pressure, 0) * $Volume
```

In this example, STAR-CCM+ substitutes zero for `Pressure` if `Pressure` is not available.

The alternate value operator is not restricted to constants; it can also work with another expression. The syntax for this use of the operator is `alternateValue(<scalar/vector expression 1>, <scalar/vector expression 2>)`, for example:

```
alternateValue($Pressure1, $Pressure2) * $Volume
```

In this example, if both `Pressure1` and `Pressure2` are not computable, then the `alternateValue` expression is not computable.

Working with Units in Field Functions

You can include units when defining a field function. STAR-CCM+ assumes all physical quantities to be in SI units by default. However, literal constants can have a unit specification in a different system. The syntax for the quantities is `<number units>`, for example:

```
<32 ft/s^2>
$Pressure - <32 psi>
<27 in> + <3 ft>
```

The units are parsed at run time, so they do not use the Preferred units setting and do not necessarily need to match any of the predefined units that the objects in the Tools > Units manager node represent.

Additionally, "-" binds tighter than "/", so `kg/m-s^2` equals `(kg)/(m-s^2)` not `(kg/m)*(s^2)`.

User Field Function Examples

This section provides some examples of user field functions. The examples include:

- Scaling an existing field function
- Setting a transient boundary condition
- Specifying components of a vector field function
- Calculating and visualizing face normals
- Creating a velocity field function that is based on grid flux
- Calculating derivatives with vector field functions

Scaling an Existing Field Function

Suppose that you wish to create a user field function that defines the distance in the x direction in inches, rather than in the SI units of the computation. You could create a field function that is called, say, "X inches" with the following definition:

```
$$Centroid[0]/0.0254
```

Setting a Transient Boundary Condition

Suppose that you want to set a transient boundary condition, such that temperature ramps up linearly from 300K to a value of 1000K at a time of 0.01 seconds. Your user field function definition (to be used in the definition of the boundary temperature) would be:

```
($Time >= 0.01)  ?  1000 : 300+70000*$Time
```

Specifying Components of a Vector Field Function

To define a vector, array or position field function in terms of its components, specify the three scalar components inside square brackets separated using commas. For example, to set up a field function to initialize the velocity in a reducing pipe, where the velocity is 2 m/s in the larger diameter and 10 m/s in the smaller diameter, and the diameter changed at x=4. The field function definition would be:

```
[($$Centroid[0] < 4) ? 2 : 10, 0, 0]
```

Calculating and Visualizing Face Normals

In STAR-CCM+, the Area field function is a vector field function of the area components, that is the area magnitude multiplied by the unit normal. Therefore, to calculate the unit face normals, a field function can be created to normalize the Area field function by the area magnitude:

- Create a field function that is called, for example, `Face Normals`.
- Set its type to `Vector` and the definition to:

  ```
  $$Area / $$Area.mag()
  ```

This field function can be used to visualize the face normals by creating a vector scene, and setting the vector quantity to your Face Normals field function. You can also add a vector displayer to an existing scene. An example scene is shown below.



Creating a Velocity Field Function Based on Grid Flux

To create a field function that is based on Grid Flux that has the same units as velocity, create a field function that is called, for example, `Grid Flux Velocity`. Set its dimensions to Length = `1` and Time = `-1`, and set its type to `Vector`. Set its definition to:

```
$$GridFlux * $$Area / $$Area.mag2()
```

In this expression, `$$GridFlux / $$Area.mag()` is the magnitude of the velocity, and `$$Area / $$Area.mag()` is the direction of the flux, normal to the face area.

Calculating Derivatives with Vector Field Functions

Typical user field functions involving derivatives are:

- Divergence of velocity $\nabla \cdot \mathbf{v}$ is specified as:

  `div($$Velocity)`

- Divergence of density, velocity $\nabla \cdot (\rho \mathbf{v})$ is specified as:

  `div($Density * $$Velocity)`

- Gradient of pressure ∇p is specified as:

  `grad($Pressure)`

- Curl of velocity $\nabla \times \mathbf{v}$ is specified as:

  `curl($$Velocity)`

A cell-based least squares gradient algorithm is used for the computation of the field gradients. At boundaries, zero-order extrapolation is used. This algorithm differs from the gradients that the various solvers compute. Therefore, the gradients computed for post-processing are likely to differ slightly from solver gradients that are viewed by using the Temporary Storage Retained option in specific solvers. However, they are sufficiently close for most practical purposes.

Context Menu of the Field Functions Manager

These context menu items are available when you right-click the Field Functions manager node.

New	Creates a field function of the type you select in the sub-menu.	
	Scalar	Creates a scalar field function.
	Array	Creates a array field function.
	Vector	Creates a vector field function.
	Position	Creates a position field function.

See also: common menu items.

Using Reference Frames

In STAR-CCM+, it is possible to specify multiple frames of reference for use in applying a constant grid flux to the governing equations of fluid motion, or for post-processing the results of a simulation. Reference frame objects represent these frames of reference. Reference frames are also added when particular types of motion are added to the Motions manager. The Reference Frames node, which has a pop-up menu, is the manager object for working with reference frames in the simulation tree.



The default type is the Laboratory Reference Frame, and the two other general types are Rotating and Rotating and Translating. Depending on how they are added, the latter two types may be either read-only or user-defined, and can vary in their labeling.



You can add more than one of any types other than the Laboratory frame.

Pop-Up Menu of the Reference Frames Manager

These pop-up menu items are available when you right-click the Reference Frames manager node.

New	Creates a new user-defined reference frame.	
	Rotating	Adds a user-defined Rotating reference frame to the Reference Frames node.
	Rotating and Translating	Adds a user-defined Rotating and Translating reference frame to the Reference Frames node.

See also: common menu items.

What Is the Lab Reference Frame?

The Lab Reference Frame node, which is subject to common menu items, is the default stationary reference frame in a STAR-CCM+ simulation.



Lab is the one type of reference frame present in the reference frames manager when a simulation is created, and is the default selection for all regions in the simulation. It cannot be deleted, and more reference frames of this type cannot be created.

What Is a Rotating Reference Frame?

There are two types of Rotating reference frames: read-only and user-defined. They use the same properties, although only the user-defined type can be added or modified.

What Is a Read-Only Rotating Reference Frame?

A read-only Rotating reference frame is created automatically for each of the following objects that you add to the simulation:

- Rotation motion
- DFBI Superposed Rotation motion

The label of the read-only Rotating reference frame includes the name of the object with which it was created.



If you rename the related object, the label of the Rotating reference frame changes accordingly.

What Is a User-Defined Rotating Reference Frame?

A user-defined rotating reference frame is a rotating frame of reference that can be applied to regions to generate a constant grid flux, or used in post-processing as an alternative frame of reference for certain quantities. Add this type of frame using the pop-up menu of the Reference Frames manager, and define it using its properties.



User-defined rotating reference frames can have further relative rotating frames defined within them. This creates a nested hierarchy of rotating reference frames, where child frames are defined relative to the parent frame.

Rotating Reference Frame Properties

These properties are shared by both the read-only and user-defined rotating reference frames, but are editable only for the user-defined type.

Axis Direction	The Cartesian direction vector specifying the axis of rotation.
Axis Origin	The Cartesian position vector specifying the origin of the axis of rotation.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.
Rotation Rate	The rotation rate around the axis specified by Direction and Origin.

Adding a User-Defined Rotating Reference Frame

To create a user-defined Rotating reference frame, right-click on the Reference Frames node and select New > Rotating from the pop-up menu.



What Is a Rotating and Translating Reference Frame?

A Rotating and Translating Reference frame provides an input for constant translation motion in addition to constant rotation. There are two types of Rotating and Translating reference frames: read-only and user-defined. They use the same properties, although only the user-defined type can be added or modified.

What Is a Read-Only Rotating and Translating Reference Frame?

A read-only Rotating and Translating reference frame is created automatically for each of the following objects that you add to the simulation:

- Rotation and Translation motion
- 6-DOF body

The label of the read-only Rotating and Translating reference frame includes the name of the object with which it was created.



If you rename the related object, the label of the Rotating and Translating reference frame changes accordingly.

What Is a User-Defined Rotating and Translating Reference Frame?

A user-defined rotating and translating reference frame is a moving frame of reference that can be applied to regions to generate a constant grid flux, or used in post-processing as an alternative frame of reference for certain quantities. It has inputs for both rotation and translation.

Add this type of reference frame using the pop-up menu of the Reference Frames manager, and define it using its properties.



User-defined rotating and translating reference frames can have further relative rotating frames defined within them. This creates a nested hierarchy of reference frames, where child frames are defined relative to the parent frame. The child reference frames have rotation only.

Rotating and Translating Reference Frame Properties

These properties are shared by both the read-only and user-defined Rotating and Translating reference frames, but are editable only for the user-defined type.

Axis Direction	The Cartesian direction vector specifying the axis of rotation.
Axis Origin	The Cartesian position vector specifying the origin of the axis of rotation.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.

Rotation Rate	The rotation rate around the axis specified by Direction and Origin.
Translation Velocity	The translation velocity of the region. This is additive with Rotation Rate.

Adding a User-Defined Rotating and Translating Reference Frame

To create a user-defined Rotating and Translating reference frame, right-click on the Reference Frames node and select New > Rotating And Translating from the pop-up menu.



What Are Relative Reference Frames?

Relative reference frames are rotating reference frames defined as children of other reference frames. They are contained within the Relative Reference Frames manager node, which is a child of both Rotating and Rotating and Translating reference frames. This node has its own pop-up menu that includes the command to add a new relative reference frame.



Relative Reference Frames Manager Popup Menu

New	Creates a new rotating relative reference frame.

See also: common menu items.

Creating a New Rotating Relative Reference Frame

Create a new relative reference frame by right-clicking on the Relative Reference Frames manager node and selecting New.



This adds a Relative Reference Frame node as a child of the manager node. Select this node and use its properties to define the relative reference frame.

Defining a Rotating Relative Reference Frame

Relative Reference Frames are defined by the properties of their node in the object tree.



These nodes can be renamed, copied, and deleted as required. Nodes of relative reference frames which are assigned to the Motion Specification of regions cannot be deleted. They must be unassigned before the option to delete becomes available.

Relative Reference Frame Properties

Axis Direction	The Cartesian direction vector specifying the axis of rotation.
Axis Origin	The Cartesian position vector specifying the origin of the axis of rotation.
Coordinate System	Specifies the coordinate system (laboratory by default) for defining the origin and axis. The drop-down list of this property includes any local coordinate systems that exist in the simulation.
Rotation Rate	The rotation rate around the axis specified by Direction and Origin.

Working with Reference Frames

One use of reference frames is to mimic the effect of a constant motion without actually moving the vertices of the mesh. This approach means that a steady solution only is required, without the need for a transient method. This is a simplification, however, and an accurate analysis of the fluid flow in such a system is likely to require use of an appropriate motion in transient flow.

To apply a reference frame to a region:

- Create the required reference frame (rotating, or rotating and translating).
- Add any further relative reference frames.
- Select the Motion Specification node of the relevant region and set its Reference Frame property to the newly created frame.

An example is shown below of a relative reference frame being applied to the Motion Specification of a region. The Motion property remains at its Stationary default value.



Currently, only the morphing and solid displacement motions can be used with user-defined reference frames. All other motions require either the Laboratory reference frame, or add their own reference frames automatically.

You are ready to use reference frames for analysis since they can be accessed in scalar and vector menus. You can display them in scalar or vector scenes, add them to reports, or plot them in monitor and XY plots.

Reference frames can be renamed, and user-defined reference frames can also be copied and deleted.

Accessing Reference Frames as Scalar Options

Once a reference frame is added to a STAR-CCM+ simulation, it is available in any drop-down list of scalar or vector values. It is listed as a suboption of relative values, such as relative total temperature in the following screenshot.



Displaying Reference Frames in Scalar Scenes

To display a value in a scalar scene with a particular reference frame, select the value and the reference frame in the scalar list. The scalar display uses that reference frame.



Displaying Reference Frames in Vector Scenes

To display a value in a vector scene with a particular reference frame, select the value and the reference frame in the list. The vector display uses that reference frame.



Using Reference Frames in Reports

Reference frames can be added to reports, and the relative values with the reference frames can be processed accordingly, such as when you run a report.



It is also possible to generate monitors and plots from such reports for plotting.

Plotting Reference Frames

As scalar or vector values, reference frames can be included in the following:

- Monitor plots



- XY plots



Copying User-Defined Reference Frames

A user-defined reference frame can be copied. When such an object is copied, all properties are copied, resulting in an exact duplicate of the original.

You can also copy the properties of one user-defined reference frame onto another, provided they are of the same type.

Copying Properties of a User-Defined Reference Frame

It is possible to copy properties among similar user-defined reference frames.

Among reference frames of the same type, this operation causes all of the properties of the source reference frame to be applied.

It is not possible to copy properties between different types of reference frames.

Renaming Reference Frames

Any reference frames may be renamed using the standard mechanisms, but be careful that the name you apply still reflects the type of reference frame. If it does not (as in the example below), you need to identify the type of the reference frame from the description area at the bottom of the *Properties* window.



Deleting User-Defined Reference Frames

User-defined reference frames may be deleted using the standard methods, unless the reference frame is being used by a region. In this case the Delete pop-up menu item is grayed and the `<Delete>` key has no effect (no warning is issued).

Read-only reference frames are deleted only when the object associated with them is deleted. This is possible provided they are not being used by any region.

Governing Equations in Moving Reference Frames

Moving Reference Frames (MRF) are reference frames that can rotate and translate with respect to the laboratory reference frame. MRF models in STAR-CCM+ assume that:

- The angular velocity of the body is constant.
- The mesh is rigid.

The governing equations can be written in terms of:

- the velocity in the laboratory reference frame (absolute velocity), or
- the velocity relative to the moving reference frame (relative velocity).

Moving Reference Frames are represented by nodes under the Tools > Reference Frames manager, except for Lab Reference Frame which represents the laboratory reference frame.



Incompressible Navier-Stokes Equations in the Laboratory Reference Frame

Note: This section presents the incompressible Navier-Stokes equations for the sake of simplicity. The extra terms that appear in the compressible Navier-Stokes equations receive no contribution from the motion of the reference frame.

In the laboratory reference frame, the incompressible Navier-Stokes equations are written as:

$$\nabla \cdot \boldsymbol{u}_0 = 0 \tag{2635}$$

$$\frac{D\boldsymbol{u}_0}{Dt} = -\nabla(p/\rho) + \nu\nabla\cdot\nabla\boldsymbol{u}_0 \tag{2636}$$

Expanding the substantial derivative in the momentum equation:

$$\frac{\partial\boldsymbol{u}_0}{\partial t} + \nabla\cdot(\boldsymbol{u}_0\otimes\boldsymbol{u}_0) = -\nabla(p/\rho) + \nu\nabla\cdot\nabla\boldsymbol{u}_0 \tag{2637}$$

The following relation is used to transform the convection term in the conservative form:

$$\nabla\cdot(\boldsymbol{u}_0\otimes\boldsymbol{u}_0) = \boldsymbol{u}_0\cdot\nabla\boldsymbol{u}_0 + \underbrace{\left(\nabla\cdot\boldsymbol{u}_0\right)}_{=\,0}\boldsymbol{u}_0 \tag{2638}$$

Using Eqn. (2638), the momentum equation becomes:

$$\frac{\partial\boldsymbol{u}_0}{\partial t} + \boldsymbol{u}_0\cdot\nabla\boldsymbol{u}_0 = -\nabla(p/\rho) + \nu\nabla\cdot\nabla\boldsymbol{u}_0 \tag{2639}$$

Incompressible Navier-Stokes Equations in a Moving Reference Frame with Relative Velocity

Consider a moving reference frame p rotating with a constant angular velocity $\boldsymbol{\omega}_p$ as measured in the laboratory frame 0. The velocity $\boldsymbol{u}_p$ relative to the moving reference frame is related to the absolute velocity $\boldsymbol{u}_0$:

$$\boldsymbol{u}_0 = \boldsymbol{u}_p + \boldsymbol{\omega}_p\times\boldsymbol{r} \tag{2640}$$

where $\boldsymbol{r}$ denotes the position vector from the origin of the moving reference frame p to the center of a control volume.

Eqn. (2640) is used to obtain the Navier-Stokes equations in terms of the relative velocity. First, the left-hand side of Eqn. (2636) is transformed:

$$\frac{D\boldsymbol{u}_0}{Dt} = \frac{D\boldsymbol{u}_p}{Dt} + \frac{d\boldsymbol{\omega}_p}{dt}\times\boldsymbol{r} + 2\boldsymbol{\omega}_p\times\boldsymbol{u}_p + \boldsymbol{\omega}_p\times\boldsymbol{\omega}_p\times\boldsymbol{r} \tag{2641}$$

Using Eqn. (2638),

$$\frac{D\boldsymbol{u}_0}{Dt} = \frac{\partial\boldsymbol{u}_p}{\partial t} + \nabla\cdot(\boldsymbol{u}_p\otimes\boldsymbol{u}_p) + \frac{d\boldsymbol{\omega}_p}{dt}\times\boldsymbol{r} + 2\boldsymbol{\omega}_p\times\boldsymbol{u}_p + \boldsymbol{\omega}_p\times\boldsymbol{\omega}_p\times\boldsymbol{r} \tag{2642}$$

The diffusive term is:

$$\nabla \cdot \nabla \boldsymbol{u}_0 = \nabla \cdot \nabla [\boldsymbol{u}_p + \boldsymbol{\omega}_p \times \boldsymbol{r}] \tag{2643}$$

$$= \nabla \cdot \nabla \boldsymbol{u}_p + \underbrace{\nabla \cdot \nabla (\boldsymbol{\omega}_p \times \boldsymbol{r})}_{=\ 0}$$

$$= \nabla \cdot \nabla \boldsymbol{u}_p$$

Applying all transformations, the momentum equation becomes:

$$\frac{\partial \boldsymbol{u}_p}{\partial t} + \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_p) + \frac{d\boldsymbol{\omega}_p}{dt} \times \boldsymbol{r} + 2\boldsymbol{\omega}_p \times \boldsymbol{u}_p + \boldsymbol{\omega}_p \times \boldsymbol{\omega}_p \times \boldsymbol{r} = -\nabla(p/\rho) + \nu \nabla \cdot \nabla \boldsymbol{u}_p \tag{2644}$$

Both the convection and convected velocities are in terms of velocity relative to the rotating frame. This formulation is also known as the relative velocity formulation.

Similarly, the continuity equation is transformed using:

$$\nabla \cdot \boldsymbol{u}_0 = \nabla \cdot [\boldsymbol{u}_p + \boldsymbol{\omega}_p \times \boldsymbol{r}] \tag{2645}$$

$$= \nabla \cdot \boldsymbol{u}_p + \underbrace{\nabla \cdot (\boldsymbol{\omega}_p \times \boldsymbol{r})}_{=\ 0}$$

$$= \nabla \cdot \boldsymbol{u}_p$$

giving

$$\nabla \cdot \boldsymbol{u}_p = 0 \tag{2646}$$

Eqn. (2646) and Eqn. (2644) form the incompressible Navier-Stokes equations in the moving frame, with relative velocity.

Incompressible Navier-Stokes Equations in a Moving Reference Frame with Absolute Velocity

In this form, the convection velocity is the velocity $\boldsymbol{u}_p$ relative to the moving reference frame and the convected velocity is in terms of the velocity $\boldsymbol{u}_0$ in the laboratory frame (that is, the absolute velocity). The convection term is of the form $\nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0)$.

Starting from Eqn. (2644), the term corresponding to the convective fluxes $\nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_p)$ is written as:

$$\nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_p) = \nabla \cdot [\boldsymbol{u}_p \otimes (\boldsymbol{u}_0 - \boldsymbol{\omega}_p \times \boldsymbol{r})] \tag{2647}$$

$$= \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0) - \underbrace{\nabla \cdot \boldsymbol{u}_p}_{=\,0} (\boldsymbol{\omega}_p \times \boldsymbol{r}) - \boldsymbol{u}_p \cdot \nabla(\boldsymbol{\omega}_p \times \boldsymbol{r})$$

The term $\boldsymbol{u}_p \cdot \nabla(\boldsymbol{\omega}_p \times \boldsymbol{r})$ is expanded as follows:

$$\boldsymbol{u}_p \cdot \nabla(\boldsymbol{\omega}_p \times \boldsymbol{r}) = \left(\boldsymbol{u}_p \cdot \underbrace{\nabla \boldsymbol{\omega}_p}_{=\,0} \right) \times \boldsymbol{r} + \boldsymbol{\omega}_p \times (\boldsymbol{u}_p \cdot \nabla \boldsymbol{r}) \tag{2648}$$

$$= \boldsymbol{\omega}_p \times \boldsymbol{u}_p \quad \text{with} \quad \nabla \boldsymbol{r} = \boldsymbol{I}$$

Therefore, Eqn. (2647) becomes:

$$\nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_p) = \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0) - \boldsymbol{\omega}_p \times \boldsymbol{u}_p \tag{2649}$$

Simplifying the steady terms in the left-hand side of Eqn. (2644):

$$\begin{aligned} \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_p) + 2\boldsymbol{\omega}_p \times \boldsymbol{u}_p + \boldsymbol{\omega}_p \times \boldsymbol{\omega}_p \times \boldsymbol{r} \\ &= \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0) + \boldsymbol{\omega}_p \times \boldsymbol{u}_p + \boldsymbol{\omega}_p \times \boldsymbol{\omega}_p \times \boldsymbol{r} \\ &= \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0) + \boldsymbol{\omega}_p \times (\boldsymbol{u}_p + \boldsymbol{\omega}_p \times \boldsymbol{r}) \\ &= \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0) + \boldsymbol{\omega}_p \times \boldsymbol{u}_0 \end{aligned} \tag{2650}$$

The unsteady term of Eqn. (2644) is transformed as:

$$\frac{\partial \boldsymbol{u}_p}{\partial t} = \frac{\partial \boldsymbol{u}_0}{\partial t} - \frac{\partial(\boldsymbol{\omega}_p \times \boldsymbol{r})}{\partial t} \tag{2651}$$

$$= \frac{\partial \boldsymbol{u}_0}{\partial t} - \underbrace{\frac{\partial \boldsymbol{\omega}_p}{\partial t}}_{=\,0} \times \boldsymbol{r} - \boldsymbol{\omega}_p \times \underbrace{\frac{\partial \boldsymbol{r}}{\partial t}}_{=\,0}$$

Using all the transformed terms, the momentum equation in terms of the absolute velocity is:

$$\frac{\partial \boldsymbol{u}_0}{\partial t} + \frac{d\boldsymbol{\omega}_p}{dt} \times \boldsymbol{r} + \nabla \cdot (\boldsymbol{u}_p \otimes \boldsymbol{u}_0) + \boldsymbol{\omega}_p \times \boldsymbol{u}_0 = -\nabla(p/\rho) + \nu \nabla \cdot \nabla \boldsymbol{u}_0 \tag{2652}$$

Eqn. (2652) is the momentum equation in the moving reference frame with the absolute velocity formulation. In this formulation, Eqn. (2635) is the continuity equation. This formulation is solved in STAR-CCM+.

Embedded Reference Frames

An embedded reference frame e allows you to specify rotational motion with respect to a parent frame p that can be both rotating and translating. The diagram below shows an example of an embedded, rotating frame that

is defined with respect to a moving reference frame and the laboratory reference frame. **X**, **Y**, and **Z** represent the axes for each frame. **O** represents the origin in each frame.



The Moving Reference Frame model handles motion in moving reference frames. The Relative Moving Reference Frame model handles motion in embedded reference frames. Both models are formulated in terms of the absolute velocity formulation that is given by Eqn. (2646) and Eqn. (2652).

Embedded reference frames are represented by nodes under the [Moving Reference Frame] > Relative Reference Frames manager.



Both models require a single initial evaluation of grid flux throughout the domain. The source term, being a function of the velocity field, is evaluated at each iteration.

Moving Reference Frame

In the Moving Reference Frame, the grid flux is calculated as:

$$G_f = (\boldsymbol{a} \cdot \boldsymbol{v}_g)_f \tag{2653}$$

with

$$\boldsymbol{v}_g = \boldsymbol{\omega}_p \times \boldsymbol{r} + \boldsymbol{V}_{tp}$$

where

- the subscript f corresponds to a face value
- the subscript c corresponds to a cell value
- $\boldsymbol{v}_g$ is the grid velocity

- $\boldsymbol{\omega}_p$ is the angular velocity of the moving frame p. This angular velocity is measured in the laboratory reference frame.
- $\boldsymbol{r} = \boldsymbol{R} - \boldsymbol{O}_p$ is the position vector with respect to the parent frame. $\boldsymbol{O}_p$ is the origin of the parent frame.
- $\boldsymbol{V}_{tp}$ is the translational velocity of the parent frame. This translational velocity is measured in the laboratory reference frame.
- $\boldsymbol{R}$ is the position vector of a control volume in the parent frame, as measured in the laboratory reference frame.

The source term is calculated as $\boldsymbol{\omega}_p \times \boldsymbol{v}_c$, where $\boldsymbol{v}_c$ is the velocity at the cell center as measured in the laboratory reference frame.

Embedded Reference Frame

In the embedded Moving Reference Frame, the grid flux is calculated as:

$$G_f = (\boldsymbol{a} \cdot \boldsymbol{v}_{total})_f \tag{2654}$$

where the total grid velocity $\boldsymbol{v}_{total}$ is calculated as the sum of grid velocities of the parent frame p and the embedded frame e:

$$\boldsymbol{v}_{total} = \boldsymbol{v}_p + \boldsymbol{v}_e = \boldsymbol{\omega}_p \times \boldsymbol{r}_p + \boldsymbol{V}_{tp} + \boldsymbol{\omega}_e \times \boldsymbol{r}_e \tag{2655}$$

where

- $\boldsymbol{\omega}_e$ is the angular velocity of the embedded moving frame.
- $\boldsymbol{r}_p = \boldsymbol{R}_0 - \boldsymbol{O}_p$ is the position vector with respect to the parent frame.
- $\boldsymbol{r}_e = \boldsymbol{R}_0 - \boldsymbol{O}_e$ is the position vector with respect to the embedded frame. $\boldsymbol{O}_e$ is the origin of the embedded frame.
- $\boldsymbol{R}_0$ is the position vector of a control volume in the embedded frame, as measured in the laboratory reference frame.

The source term is calculated as $\boldsymbol{\omega}_{total} \times \boldsymbol{v}_c$, where the angular velocity is the vector sum of the angular velocities of the parent and embedded frames:

$$\boldsymbol{\omega}_{total} = \boldsymbol{\omega}_p + \boldsymbol{\omega}_e \tag{2656}$$

Using Tables

The Tables node, which has a pop-up menu, is the main manager object for all tables that have been created in the simulation. It exists even when empty to allow creation of the first table. The tables manager can hold any number of tables up to the memory limits of the computer.



This section describes tables and how to use them. Specifically, instructions are provided for:

- Reading a table data file
- Creating and filling a table
- Exporting a table
- Copying part data between simulations
- Displaying a table
- Automatically extracting and exporting tables
- Copying properties of a table
- Deleting a table
- Renaming a table

A reference is also provided for the format of table files.

What Are Tables?

Tables refer to objects in STAR-CCM+ containing tabular data. These objects are multi-variable entities akin to spreadsheets. They can be used for:

- Specifying tabular data for the simulation, particularly for boundary conditions and/or initial conditions
- Extracting tabular data from the simulation, for instance to use as boundary conditions for other simulations

The table contains a combination of location and scalar data. Normally, a file table has been created using external data such as from experiments or an analysis, but it could also be extracted and exported from a simulation. There are several types:

- File tables
- Xyz internal tables
- Radial internal tables
- Histogram tables
- Accumulated force tables

Table Pop-Up Menu

Export...	Executes the export table action and saves the contents of the table as a separate file in tabular data set (.tab) format.
Extract	Executes the extract tabular data action and fills a table with data.
Reload file...	Executes the reload table file action and updates the table data from a modified file (file tables only).
Tabulate...	Opens a dialog displaying tabular data.

See also: common menu items.

What Are File Tables?

A file table contains data that STAR-CCM+ gets by reading a table data file, which is generated in an experiment or analysis. The file formats are .csv, .tab, and .xy. It is represented in the object tree by a node that has properties and a pop-up menu.



It is possible to access the table data from a field function.

File Table Properties

Extracted	A read-only list of the names of the scalar field in the file. This property is blank if a new tabular data set has been created but no scalars have yet been selected.
File Path	The path to the table file from which this table object gets its data. This path is relative to the location of the simulation file.

What Are Xyz Internal Tables?

An xyz internal table specifies location in Cartesian coordinates corresponding to cell or face centroid positions in the Laboratory coordinate system. The data set must contain an identifier specifying that it is an xyz table. It also needs a list of identifiers corresponding to the "columnar data" in the file. The columnar data must contain the Cartesian coordinates (x and y are only required for two-dimensional simulations) as well as the scalar

data. An example of the file format is given in Xyz Internal Table File Format. It is represented in the object tree by a node that has properties and a pop-up menu.



The interpolation method for xyz tabular data is zeroth order. When interpolating the data to a specific cell or face location, STAR-CCM+ finds the coordinate location in the tabular data that is closest to the specific cell or face centroid. STAR-CCM+ then applies the appropriate scalar data from the table to that cell or face.

Internal xyz Table Properties

Coordinate System	Specifies the coordinate system (laboratory by default). The drop-down list of this property includes any local coordinate systems that exist in the simulation (xyz internal table only).
Parts	Editing the property activates the Select Objects dialog for parts.
Scalars	The property customizer activates a list box functionally similar to the Select Objects dialog, with which you can select scalars.
Extracted	A read-only list of the names of the scalar field in the file. This property is blank if a new tabular data set has been created but no scalars have yet been selected.

What Are Radial Internal Tables?

A radial internal table specifies location in terms of radius, and the data set must contain an identifier specifying that it is a radial table. It also needs a list of identifiers corresponding to the "columnar data" in the file. The columnar data must contain the radial coordinate as well as the scalar data.

An example of the file format is given in Radial Internal Table File Format. It is represented in the object tree by a node that has properties and a pop-up menu.



The interpolation method for radial tabular data is a cubic spline. Hence, it is assumed that the radial data are sorted and single-valued.

You can customize how this table is updated as the simulation iterates.

Radial Internal Table Properties

Number of Bins	Number of radial stations.
Parts	Editing the property activates the Select Objects dialog for parts.
Scalars	The property customizer activates a list box functionally similar to the Select Objects dialog, with which you can select scalars.
Extracted	A read-only list of the names of the scalar field in the file. This property is blank if a new tabular data set has been created but no scalars have yet been selected.

What Are Histogram Tables?

Histogram tables allow you to create a table showing the distribution of a selected quantity with optional weighting. The parameters that are required for STAR-CCM+ to create the table are set in the properties of the table node and its child Histogram node.



You can customize how this table is updated as the simulation iterates.

Histogram Table Properties

Parts	Parts from which the table data is extracted.
Extracted	A read-only list of the names of the data sets extracted for the table.

Defining the Histogram

The properties of the Histogram node, and its sub-node (some of which are optional) control the definition of the histogram table. The properties of the Bin sub-node are used to split the scalar field function into intervals (or bins).



Histogram Properties

Number of bins

Number of discrete intervals in the histogram.

Smooth Values

Node values can be used instead of cell values

- **Activated** — Node values are used in populating the histogram.
- **Deactivated** — Cell center values are used.

Generate Profile

For each set of points that is allocated to a bin, STAR-CCM+ evaluates the field function that is chosen in the properties of the Profile node. STAR-CCM+ then returns the maximum and minimum values of the function.

- **Activated** — Adds the Profile sub-node. Select the field function that is used for generating the profile data.
- **Deactivated** — No profile data is provided.

Weighting

STAR-CCM+ weights the count of each point that is allocated to a bin by the amount that is selected here.

- **Frequency** — Equivalent to a weighting of one. The sum for each bin is the frequency with which points fall into the bin.
- **Specified Scalar** — Adds the Weighting sub-node. Select the field function that is used to weight each data point allocated to a bin.
- **Specified Scalar with Accumulation** — Same as Specified Scalar, but with the additional property that values for successive bins are summed cumulatively along the bin direction.

Range Mode

Select how to find the maximum and minimum range of bin values.

- **Auto** — STAR-CCM+ automatically calculates the maximum and minimum range of values for the bin field function, and sets the range from this calculation.

Manual	Adds the Range sub-node. Specify the range of values within which bins are formed.
Range	Minimum and maximum binning values (read only).
Spacing	Interval width for each bin (read only)

Histogram Range Properties

Range	Minimum and maximum values of the binning quantity. Set as a scalar array profile.

Histogram Weighting Properties

Field Function	Selects the scalar field function that is used to weight the data points allocated to each bin.
Units	Choose appropriate units from the drop-down menu.

Histogram Profile Properties

Field Function	Selects the scalar field function from which to extract minimum and maximum values that are calculated using the set of data points in each bin.
Units	Choose appropriate units from the drop-down menu.

Histogram Bin Properties

Field Function	Selects the scalar field function to split into separate intervals (or bins).
Units	Choose appropriate units from the drop-down menu.

What Are Accumulated Force Tables?

Accumulated force tables provide a convenient way to plot accumulated pressure and shear forces (or both) over the surface of selected parts. The parameters that are required for the calculation are set through the properties of the table node, and also through the properties of its child Force Histogram node.



You can customize how this table is updated as the simulation iterates.

Accumulated Force Table Properties

Parts	Editing the property activates the Select Objects dialog for parts.
Extracted	A read-only list of the names of the data sets extracted for the table.

Defining the Force Histogram

To set the parameters that are required for building the histogram of forces and accumulated force across the chosen parts use the properties of the Force Histogram node. The forces that are extracted are the same as the forces given for the Force report. When force normalization is activated, the force coefficients, and accumulated force coefficients, are returned.



Force Histogram Properties

Number of bins	Number of discrete intervals (or bands) across the parts.
Coordinate System	Coordinate system in which the bin, force, and profile directions are defined. The drop-down list of this property includes any local coordinate systems that exist in the simulation. The default is laboratory. This property does not control the Position values given in the displayed table. The Position values given in the table are always for the centroids of the bins, which are given in the laboratory coordinate system along the specified bin direction.
Bin Direction	Direction in which the intervals (or bands) are distributed.
Force Direction	Direction in which the force is computed
Profile Direction	Direction to use in extracting the profile from the parts.
Force Option	Choose which force components are included in the reported data.
Pressure	Only the force due to pressure, $\mathbf{f}_f^{\text{pressure}}$ is considered. See Eqn. (2654) for the Force report.
Shear	Only the force due to shear, $\mathbf{f}_f^{\text{shear}}$ in (2654) is considered.
Pressure and Shear	Both pressure and shear forces, $\mathbf{f}_f^{\text{pressure}}$ and $\mathbf{f}_f^{\text{shear}}$ in (2654) is considered.
Normalization Mode	Choose whether to normalize the force as a force coefficient, or to return the forces directly. See Eqn. (2660).
None	No normalization. The table contains the forces.

	Standard	Adds the Force Normalization child node. The table contains the force coefficients.
Range	Maximum and minimum band values (read only).	
Reference Pressure	The reference pressure for the computation of Eqn. (2655).	
Range Mode	Sets how band positions are calculated.	
	Auto	Use this option if you want STAR-CCM+ to calculate the start and end positions of the bands.
	Manual	Use this option if you want to specify where the first band begins and the last band ends.
Spacing	The width of the bands across the selected parts (read only)	

Force Normalization Properties

Reference Density	The reference density that is used to calculate the force coefficients (see Eqn. (2660)).
Reference Velocity	The reference velocity that is used to calculate the force coefficients (see Eqn. (2660)).
Reference Area	The reference area that is used to calculate the force coefficients (see Eqn. (2660)).

Reading a Table Data File

To read a table data file, right-click the Tables node under Tools in the object tree, then select New Table > File Table as shown in the following screenshot.



You can also read multiple table data files with this one menu action.

Select an available tabular data file from the *Open* dialog, then click Open. A table is created using the data from the specified file.



The path to the table file is relative to the location of the simulation file.

The new table appears in the Tables node. The table is initially assigned a name corresponding to the file name (without the `.csv` extension), but can be renamed.



If you modify a table file that has been already read into the simulation, you can reread the file easily.

It is also possible to change the file path of the table.

Reading Multiple Table Data Files

Once you execute the action for reading a table data file, you can read multiple table files at one time.

In the *Open* dialog, use the `<Ctrl><Click>` approach to select multiple tabular data files. The names of the selected files are listed in the File Name text box. When finished selecting files, click Open. One table is created for each specified file.



All the table files that you select must be in the same directory. For more table files in another directory, repeat this action.

The new table objects appear in the Tables node.



Rereading a Table Data File

To reread a table data file, right-click its table node and select Reload file...



This action updates the table data for this object in the simulation.

Changing the File Path of the Table

If you want to use a different set of data in the same table object, you can simply change the file path. Select the table node.



In the Path property, enter the path for the new file.



Alternatively, you can click the property customizer (...) to open the same dialog that is used for reading a table file. This dialog allows you to navigate to the new file.

Once you have changed the path, reload the table file.

For best results, ensure that your new table file has the same titles, numbers of columns, and rows as the old one.

Creating and Filling a Table

Create an xyz internal table or a radial internal table that is filled with the data within an existing simulation using the following four-part process:

1. Create a new (empty) table.
2. Specify the parts from which the data are extracted.
3. Specify the scalar field function values to extract.
4. Extract the data.

A table that is created in this way can either be used in the current simulation or exported to a file.

Creating a New (Empty) Table

To create the new (empty) table, right-click on the Tables node under Tools in the object tree, then select New Table > XYZ Internal Table... as shown below, or alternatively select New Table > R Internal Table... if necessary.



The default name for the table is XyzInternalTable 1, and it can be renamed if necessary.

Specifying the Parts

To select the parts from which the data are extracted, select the node for the new table.



In the *Properties* window, click the right column of the Parts property.



The *Select Objects* dialog displays and allows you to select the required parts.



Specifying the Scalars

To specify the scalars that must be extracted, select the table node in the object tree. In the properties for this table, click the property customizer of the Scalars property. The *Scalars* selection dialog displays. Using this dialog, select the required scalars.



Extracting the Data

Once the parts and the scalars are specified, right-click on the table name and select Extract. The table is then filled with data.



Use this feature to update the table object when the external table file has been updated with new data.

Exporting a Table

To export an existing table to a file, right-click the table name and select the Export... menu item.



A *Save Table to File* dialog opens. Enter the file name (with the extension `.csv`) in the File Name text box, and click Save.



Copying Part Data Between Simulations

You can use tables to copy solution data of a single part from one simulation to another.

Begin by creating a table in the simulation that contains the solution data. In that new table, specify the part from which you want to copy the data. Next, select the scalar that you want to copy. Extract the data and export the table to a file. Finally, read the table into the simulation file where you want to add the data.

In the destination file, set up a scalar scene with the part on which you want the table data applied.



Suppose that the destination part is a boundary. Open the Regions node, the node of the region containing the boundary, the Boundaries node, the node of the boundary and the Physics Values node. Select the profile node that corresponds to the table values that you are copying.



Select a tabular method for the profile. Then set the tabular method properties to use the values in the table file.



Set the value in the scalar scene to the one that corresponds to the table values that you are copying.



Next, step the solution. (When the *Residuals* display appears, click the tab of the scalar scene display to bring it back into focus.) The table values are displayed on the part in the scalar scene.



Displaying a Table

To display the data contained in a table, right-click the table node and select Tabulate...



The data from the plot are presented in the *Plot Tabular Data* dialog.



Units are displayed in the preferred unit for the selected scalar.

Automatically Extracting and Exporting Tables

You can update, extract, and export tables based on a trigger event. To select the trigger event, expand the Tools > Tables node, and choose the table (internal table or histogram table). The Update node, which has its own properties, customizes how tables are updated as the simulation iterates. See Trigger Types.



The Update node has a sub-node with which to control the frequency of the action (update, extraction, and export). The properties of the Update Frequency child node are used to change the frequency of updates. See Trigger Properties.



For more information on selecting the trigger type and defining the update frequency, go to Setting the Trigger for an Update.

Update Properties

Trigger	Determines how the table is updated. Select the desired trigger type from the drop-down list. See Trigger Types.
Auto Extract	Controls automatic extraction. Activated — Extracts automatically. Deactivated — Does not use this feature.
Save to File	Controls automatic export. Activated — Exports automatically. Deactivated — Does not use this feature.
Output Directory	Path to where the files specified in the Base File-name Append Tag property are saved.
Base Filename Append Tag	The prefix of the file name to use when saving tables to file.

Copying, Deleting, and Renaming Tables

Copying Tables

When properties are copied from one table to another, the Dimensions and binsNumber properties are changed. (The binsNumber property appears in radial tables.) The Type property remains unchanged. You would still specify the parts and extract the data of the destination table.

It is not possible to copy properties between different types of tables.

Deleting Tables

Tables can be deleted from an existing simulation as long as they are not in use. To prevent tables that are in use from being deleted, the Delete menu item for that table are grayed out.

Renaming Tables

Similar to other objects, a table can be renamed within STAR-CCM+.

File Format Reference

This section details the formats of files when tables are saved to disk:

- Xyz internal table
- Radial internal table

It also describes the formats of files from which you can read a table:

- .csv
- .tab
- .xy

The following principles apply to using variables in either internal table format:

- You can use any variable name that you like in the file as a heading row, and then assign it to the scalar value in the boundary condition.
- The file can contain any number of variables, and they can be used or unused.
- However, the xyz file must have the fields X, Y and (if three-dimensional) Z, while the radial file needs the field R.
- A variable name can contain a units specification. The units are specified within a set of parentheses () as part of the variable name.

  ```
  "Volume (in^3)", "Density (lb/ft^3)"
  ```

  Imported units are converted to the preferred units for a dimension that is based on the Preferred System property, which is specified in the Units node. If the provided units are not known to the units manager, the *Output* window displays that unit conversion is not possible.

Xyz Internal Table

An example of a two-dimensional xyz internal table file is shown below, with columns for 'X' and 'Y' and the field functions 'U' and 'SaNut'. Critical aspects of the file format are:

- The specific parentheses used
- The use of quotation marks to identify text strings

- The literal identifiers including 'Type', 'xyz' and 'DataSets'.

The columnar data inside the brackets are free format, and delimiters are used to separate the individual numbers. Each column of data must have the same number of elements. There must be columns of data that are identified using 'X', 'Y' (and 'Z' in three-dimensional space).

```
{'Type': 'xyz'
'DataSets': [ 'X'  'Y'  'U'  'SaNut']
'X':
[-4 -4 -4 -4 -4 -4 -4 -4 -4 -4 -4]
'Y':
[1.0 1.1 1.2 1.3 1.4 1.5 1.6 1.7 1.8 1.9 2.0]
'U':
[0.04 0.123 0.209 0.291 0.359 0.409 0.445 0.471 0.491 0.509
0.523]
'SaNut':
[1.14e-05 3.51e-05 6.078e-05 8.84e-05 0.00012 0.00015 0.0
0019 0.00022 0.00026 0.00030 0.00035] }
```

Radial Internal Table

An example of a radial internal table file is shown below. Critical aspects of the file are the specific parentheses used, the use of quotation marks to identify text strings, and the literal identifiers such as 'Type', 'radial' and 'DataSets'. The columnar data inside the brackets are free format, and any delimiters can be used to separate the individual numbers. Each column of data must have the same number of elements. There must be a column of data identified as 'R'.

```
{'Type': 'radial'
'DataSets': [ 'R'  'axial-vel']
'R': [0 0.01 0.02 0.03 0.05 0.07 0.09 0.1 0.12 0.13 0.14
0.15]
'axial-vel': [0.9700 0.9756 0.9697 0.9878 0.9955 1.0223
1.0404 1.0519 1.059543 1.065043 1.064233 1.071543] }
```

.csv Table File

A simplified example of a `.csv` (comma-separated values) table file is shown below. The individual numbers can be separated by commas or tabs. Each column of data must have the same number of elements, as well as a label at the top.

```
Vol   Press
0.01  458.58
0.10  417.85
0.20  286.26
```

```
0.30  278.52
0.40  267.68
0.50  254.59
0.60  239.05
 ...   ...
2.20 -204.01
```

This format is the best way to bring in user data (`.usr`) files from pro-STAR that were written in ASCII format using the `SAVU` command.

.tab Table File

An example of a `.tab` table file is shown below. Each column of data must have the same number of elements.

Critical aspects of the file format are:

- The specific parentheses used
- The use of quotation marks to identify text strings
- The literal identifiers such as 'DataSets', 'Modified Diffusivity' and 'Velocity: X-Component'.

```
{'DataSets': ['Modified Diffusivity', 'Velocity:
X-Component', 'Velocity: Y-Component', 'X', 'Y', 'Z'],
'Modified Diffusivity': [0.00050444749649614095688,
0.001555675407871603965 8, ...], 'Type': 'xyz', 'Velocity:
X-Component': [0.55402863025665283203,
0.68478977680206298828,...], 'Velocity: Y-Component':
[-7.3806604632409289479e-06, -3.4772736398736014962e-05,
...], 'X': [-4.0000000000000000000,
-4.0000000000000000000,...], 'Y': [8.9650001525878906250,
8.8903274536132812500,...], 'Z': [0.0000000000000000000,
0.0000000000000000000,...]}
```

Although .csv is the standard format for exporting table data from STAR-CCM+, `.tab` is still supported as a legacy format for reading into a simulation.

.xy Table File

An example of a `.xy` table file is shown below. Each column of data must have the same number of elements. There must be columns of data identified as 'X', 'Y' (and 'Z' in three-dimensional space).

```
Y X Z U
0.0 0.0 0.0 0.00000
0.0005 0.0 0.0 0.11408
0.001 0.0 0.0 0.22230
```

```
0.002 0.0 0.0 0.42120
0.003 0.0 0.0 0.59670
0.004 0.0 0.0 0.74880
0.005 0.0 0.0 0.87750
 ...   ...   ...   ...
0.3 0.0 0.0 0.00000
```

Although .csv is the standard format for exporting table data from STAR-CCM+, `.xy` is still supported as a legacy format for reading into a simulation.

Pop-Up Menu of the Tables Manager

These pop-up menu items are available when you right-click the Tables manager node.

New Table	Creates a new tabular data set of the type that is selected in the submenu:	
	File...	Executes the open file action and reads in an existing tabular data file.
	Accumulated Force	
	Histogram Table	
	xyz Internal Table	Executes the create internal tabular data set action and creates an empty xyz tabular data set.
	R Internal Table	Executes the create internal tabular data set action and creates an empty radial tabular data set.

Using Engineering Units

STAR-CCM+ has flexible and customizable support for engineering units. The standard SI and USCS systems are built in, and you can add new units as desired. Physical quantities can be specified and displayed in any of these units. Individual quantities can be specified with new units on the fly by simply typing something like “80 mph”.

This section explains the basic concepts associated with units. In addition, details are given about working with units and using units while setting values.

What Is a Unit?

A unit is a reference amount of a physical quantity used to measure quantities of the same dimensions. Each unit is represented by a tree node with properties; its pop-up menu consists of common menu items.



Examples of some standard units are meter, foot, second, kilogram, Newton, pound and Watt.

A unit has an associated set of dimensions. For example a meter has dimensions of length. A foot, inch and millimeter also have dimensions of length. Any quantity with dimensions of length can therefore be expressed in meters, feet, inches and millimeters. It is also possible to set preferred units.

You can delete a unit that you have defined if it is not being used.

Properties of a Unit

Conversion	The conversion factor that converts the quantity from this unit to SI units. Read-only except for user-defined units.
Description	The full name of the unit. Read-only except for user-defined units.
Dimensions	The physical dimensions of the units. Read-only except for user-defined units.
Offset	The offset factor used to convert the quantity from this unit to SI units. Used for Celsius and Fahrenheit temperatures. Read-only except for user-defined units.

Preferred

For each dimensional combination in use in the simulation, STAR-CCM+ marks exactly one unit to be the preferred unit for that dimension. For example the preferred unit for quantities with the dimension of length might be meters and the preferred unit for quantities with dimensions of length per time might be kph (kilometers per hour).
The preferred units are the ones that are used by default for any new physical quantity. You can change the preferred unit for an individual dimensional combination, and you can change the preferred units for all quantities to a standard system. Although the preferred unit is used by default, you can override it at any time you like in specific situation. For example, you might have the preferred unit for dimensions of velocity to be the SI standard of meters per second, but for a specific velocity vector display you could change the units to kilometers per hour.
Note: Field functions default to SI units, without using to the Preferred property. See Working with Units in Field Functions.

Ticked — Makes this unit the preferred one for these dimensions.

Cleared — Causes the preferred unit for this dimension to revert back to the default for the current preferred system (SI or USCS).

System

The systems (SI, USCS, or Other) to which this unit belongs. "s" (seconds), for example, is part of both SI and USCS. User-defined units belong to "Other".

What Is a Dimension?

An understanding of the concept of dimensions is critical for the manipulation of units. Physical quantities are characterized by their dimensions.

This dialog is activated by the property customizer button (...) in the Dimensions property of a unit or a field function. It remains in read-only mode when you activate it in the properties of a default unit. However, it can be in edit mode to let you define your own units or specify the dimensions of a user field function.



The fundamental dimensions used in STAR-CCM+ are:

- Length
- Time
- Current
- Mass
- Temperature
- Luminosity

- Quantity
- Angle

and the derived dimensions are:

- Force
- Energy
- Power
- Pressure
- Velocity
- Angular velocity
- Volume
- Moment
- Stress

Every physical quantity has a set of dimensions which are one of these dimensions, or a combination of these basic dimensions. For example gravity has dimensions of acceleration, which is dimensions of length per time squared.

Working with Units

The Units manager node, which has its own properties and pop-up menu, contains all of the unit nodes in the object tree.



This section explains the following:

- selecting a standard system;

- specifying the default units;
- defining your own units; and
- copying the properties of units.

Selecting a Standard System

There are two primary units systems available for selection in STAR-CCM+:

- Systeme Internationale (SI), which includes meter, kilogram and second; and
- United States Common System (USCS), which includes foot, pound and second.

To set the SI or USCS systems as your preferred system, simply select it in the Units node in the Tools node:



What happens when you change the system of units is that STAR-CCM+ runs through all the units and sets the preferred ones to be the SI or USCS selection.

Keep in mind that any time a new preferred unit is selected it will reset the units everywhere in the system to that new value. Thus, it will reset any manual overrides you may have set in specific quantities.

Specifying the Default Units

You can specify which system and unit STAR-CCM+ uses as its default for every new simulation.

To do this, first select a standard system and/or set a preferred unit. Then right-click the Units manager node in the Tools node and select Save as Defaults.



Your defaults will be applied when a new simulation is created.

Defining Your Own Units

It is easy to create a new custom unit. You will give this unit a certain set of dimensions and conversion information, and then you can use the unit anywhere a quantity with those dimensions is used.

First create the units with the right-click menu of the Units node in the Tools node:



This will result in a new blank unit:



You would then rename the unit.



Then click the property customizer button (...) of the Dimensions property to set the dimensions. In this case the dimensions are length, so the length dimension exponent should be set to 1.



Press OK, and then set the Conversion and Description properties for the new unit:



You can now use this unit anywhere a quantity with dimensions of length is required.

Copying Properties of Units

When properties are copied from one unit to another, all of the properties other than System are applied.

It is not possible to copy properties between built-in and user-defined units.

Properties of the Units Manager

Preferred System	Selects a preferred system of units.	
	Systeme Internationale	Includes kilogram, meter, Pascal, Newton, Joule, Watt, second.
	United States Customary System	Includes pound, foot, second.

Pop-Up Menu of the Units Manager

New Units	Executes the create units action to add a user-defined unit type to the Units node.
Save as Default	Executes the save as default action so that all new simulations use the specified unit settings.

See also: common menu items.

Using Units While Setting Values

This section explains techniques for adjusting units at the same time that you change values:

- using the keyboard to specify different units;
- using the keyboard to convert units; and
- using the editor to convert units.

Using the Keyboard to Specify Different Units

You can specify a constant in non-default units simply by typing the units along with the number. For example, if you use the default SI unit system, the units for velocity will be meters per second:



If you want to change it to a velocity of 100 kilometers per hour, simply type `100 kph` into the Value field:



and press `<Enter>`:



Using the Keyboard to Convert Units

You can convert units in an entry by typing the new units without typing a number. Following from the example above, let's say we wanted to convert the 100 kilometers per hour to miles per hour. In this case simply type `mph` in the Value field:



and hit `<Return>`:



The velocity value is converted from 100 kph to 62.137 mph.

Using the Editor to Convert Units

You can also convert units using the graphics editor. Continuing with the example above, let's convert back feet per second. Select the Value field and press the property customizer (...):



This will open up the editor for the quantity:



For the Units field, just select feet per second:



The velocity value will change from 62 miles per hour to 91 feet per second:



You can then press OK to accept or Cancel if you do not want to save the change.

Interacting with CAD/CAE Products

This part of the documentation describes how STAR-CCM+ can integrate a wide range of other computer-aided engineering (CAE) tools into a simulation.

The basic import and export features let you use external geometry, surface and volume mesh data, as well as export such data from a STAR-CCM+ simulation. Therefore you can incorporate a host of external file formats and CAE models into your simulation.

A more advanced integration of CAE with STAR-CCM+ is coupling. This allows you to use STAR-CCM+ to obtain part of the model solution, and a second code to obtain another part of the model solution. By coupling the two codes together, solution data is passed between them at specified intervals during the simulation, and you can obtain a solution for the whole model. STAR-CCM+ offers two methods of coupling: file-based coupling, whereby data is exchanged between codes via files stored in the computer disk, and co-simulation, in which data is exchanged automatically while the simulation is running. File-based coupling is very versatile and can be used for most types of coupled simulation. It relies on the user to manually execute the coupling steps via the STAR-CCM+ interface, although the process can be automated via scripts. Co-simulation is intended for two-way closely coupled cases, and uses automatic data exchange between the following codes:

- GT-POWER
- OLGA
- AMESim
- WAVE
- RELAP
- STAR-CCM+ (co-simulation between two STAR-CCM+ sessions)
- Reacting Channel (1D Plug Flow Reactor in DARS basic CFD)
- Abaqus
- Isight

The methods available for each coupling technique depend on which third party software you are using.

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Modeling physics -- the concepts, theory and implementation details of the models
- Using tools in STAR-CCM+ to facilitate your work
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format
- A guide to using STAR-View+, a viewer for three-dimensional STAR-CCM+ scenes

Importing into STAR-CCM+

STAR-CCM+ can import geometry or mesh files in a wide variety of commonly used file formats. This lets you use material from your CAD or meshing tools in a STAR-CCM+ simulation.

Additionally, if you have files from STAR-CD, you can post-process their data in STAR-CCM+. This allows you to take advantage of powerful visualization and plotting tools to review your STAR-CD material.

STAR-CCM+ also allows you to import models from a number of other CAE packages so that you can use data associated with the model in your STAR-CCM+ simulation, or export data from STAR-CCM+ for use in third party software. This technique can be used for coupling STAR-CCM+ with other codes.

Importing Geometry and Meshes

This section provides details on the following:

- You can import a volume mesh from a range of external file formats, and ensure its quality using various STAR-CCM+ features, including diagnostics and reordering.
- If you choose to create a volume mesh within STAR-CCM+, it is possible to work with surface formats, for importing into 3D-CAD and/or importing a surface mesh.
- When you import a surface mesh or volume mesh, STAR-CCM+ automatically assigns boundary types.

Importing a Volume Mesh

A number of different file formats are available in STAR-CCM+ to import the volume mesh. A progress report and, optionally, mesh diagnostics feedback, can be produced during the import process. Multiple mesh imports, either simultaneously or consecutively, are possible in the same simulation.

Once the volume mesh import is complete, you can review new simulation objects that were created. If the volume mesh is modified in any way thereafter, it can be reordered to optimize the solver performance.

If you choose to use an external mesh generator, it may create a mesh that is not consistent with the convention for face orientation in STAR-CCM+. Although mesh files exported from pro-STAR and other CD-adapco products conform to the expected (right-handed) face orientation, Gridgen can write meshes with either left- or right-handed face orientations. Incorrect face orientations can result in negative cell volumes and hence an invalid mesh. For this reason, make sure that your mesh was written with the correct face orientation.

Exporting from pro-STAR

Volume meshes can be exported from pro-STAR v4.0x to STAR-CCM+ using the `.ccm/.ccmg` file format (the .ccmg format replaces the .ccm format from pro-STAR v4.04 onwards). If you are using an older version of pro-STAR (say v3.2x), use the `.ngeom` file format instead. The boundary types in STAR-CCM+ will, as far as possible, be inferred from the pro-STAR model for both formats. Any unrecognized types are interpreted as walls. Boundaries that are defined as cyclic in pro-STAR are also imported as walls in STAR-CCM+, but they can then be used to create a periodic interface.

Only mesh information is contained in the `.ngeom` and `.ccmg` files but the `.ccm` file format supports both mesh and solution data. Any simulation parameters (such as under-relaxation values or choice of differencing schemes) or boundary conditions that are defined in pro-STAR are ignored when exporting the mesh data. These parameters must be redefined in STAR-CCM+ once the mesh or meshes have been imported.

To export a mesh from pro-STAR as an `.ccm` file:

1. Start pro-STAR and resume an existing pro-STAR model or create a new one.
2. Define and/or merge boundary regions. There is a one-to-one correspondence between the boundary regions defined in pro-STAR and the boundaries in STAR-CCM+. This step is not essential as boundaries can be manipulated in STAR-CCM+ after the import, but it may require less effort to have the boundaries demarcated properly before the mesh is imported.
3. Define and/or merge cell type properties. There is a one-to-one correspondence between all the cells with the same properties (material label, material index, porosity index, spin index and group index) defined in pro-STAR and the regions in STAR-CCM+. The use of different cell types has no effect on the resultant region definition unless one of the above properties is used as well.
4. Enter the `GEOM` command in the command line to specify the file and scale factor for the export.

The steps that are described above can be used to export a `.ngeom` file. Use the `NGEOM` command instead of `GEOM`.

In both the above instances, the mesh can either be scaled during export, or it can be scaled within STAR-CCM+ after the import has been completed.

Scaling During Export

Both the `GEOM` and `NGEOM` commands in pro-STAR allow a scale factor to be supplied during the export process in order to correctly size the model for STAR-CCM+. The formats of the commands are as follows:

```
GEOM,[filename.ccm/.ccmg],[scale factor],CCM
```

```
NGEOM,[filename.ngeom],[scale factor],NGEOM
```

However, the two file formats that are produced by the commands differ in the use of the scale factor. In the case of the `.ccm/.ccmg` file, the scale factor is not used to scale the vertices in the mesh but is stored on the file and used by STAR-CCM+ to scale the vertices during the import process. The mesh is therefore stored in the same units as used in the original pro-STAR model.

In the case of `.ngeom` files, the scale factor is applied during the export process and is used to scale the vertices directly. The scale factor is also stored in the file but is not used by STAR-CCM+ during the import process.

Therefore, in both the above cases, the application of the correct scaling factor results in the mesh being scaled appropriately once it is imported into STAR-CCM+ and no further scaling is necessary.

Importing a Volume Mesh

Before you can import a volume mesh, a new or existing simulation must be open. You can then import a volume mesh by selecting the Import Volume Mesh... menu item of the File > Import submenu.



The *Open* dialog appears.



Select one of the supported format types. You can select more than one file to import if desired. You can filter files according to the mesh file types in the Files of Type drop-down list. If you have numerous files of the same type, you can further restrict the selection by typing the common name and an asterisk with the extension of the file type that you want. For example:



This action changes two attributes of the *Open* dialog:

- Only files of the type you specified are displayed in the dialog.
- That type alone is listed in the Files of Type drop-down list.

Select the mesh file that you want to import and click the Open button.

When you attempt to import a file that contains more than one solution state, the following dialog appears:



To start the import, select the solution state and click Close.

Once you start the import, an import progress report is displayed in the *Output* window and new regions and boundaries are added to the simulation tree. STAR-CCM+ automatically assigns boundary types. If you activate the appropriate option, a mesh diagnostics report and/or geometry scene is generated.

Reviewing the Import Progress Report

When importing volume meshes into STAR-CCM+, the progress of the import is reported in the *Output* window. The example below shows the full output from a successful `.ccm` file import.

```
Output - Star 1

Importing /home/star_import/files/hx_unit.ccm

Reading /home/star_import/files/hx_unit.ccm...
CCM file version 20601
Cell type[1] label: "Fluid_body", material: "fluid" (2) porosity: 0 spin: 0 group: 1
Cell type[2] label: "Solid_body", material: "solid" (2) porosity: 0 spin: 0 group: 2
Boundary type[0] id: 1, name: "Interface", label: "Interface", type: "wall"
Boundary type[1] id: 2, name: "Inlet", label: "Inlet", type: "inlet"
Boundary type[2] id: 3, name: "Outlet", label: "Outlet", type: "pressure"
Boundary type[3] id: 4, name: "Symmetry", label: "Symmetry", type: "symmetry"
Boundary type[4] id: 5, name: "Interface", label: "Interface", type: "wall"
Boundary type[5] id: 6, name: "Symmetry", label: "Symmetry", type: "symmetry"
Boundary type[6] id: 7, name: "Inner_cylinder_wall", label: "Inner cylinder wall", type: "wall"
Boundary type[7] id: 8, name: "Outer_cylinder_wall", label: "Outer cylinder wall", type: "wall"

Processing mesh
Finished reading /home/star_import/files/hx_unit.ccm
Reordering Reordering cells in region Fluid_body... done
Reordering faces in region Fluid_body... done
Reordering cells in region Solid_body... done
Reordering faces in region Solid_body... done

Finished importing meshes
Imported coordinate extents (m):
   X extents: [-0.0254, 0.4318]
   Y extents: [-0.1269488478, 1.555250069E-17]
   Z extents: [-0.127, 0.127]
```

Importing files from other formats generates a slightly different output, but in general, the feedback consists of:

- the names or identifiers for regions and boundaries;
- the boundary types;
- output from the mesh reordering process;
- the coordinate extents of the imported mesh or meshes; and
- any error or warning messages if any problems are encountered

When exporting a mesh from pro-STAR using the `.ccm` file format, any material and/or boundary names that are set are used for naming regions and boundaries in STAR-CCM+ respectively. If no names are specified, the boundary names are inferred from the boundary types where possible.

The number of vertices in the file may not match the number of vertices in the resulting regions if any vertices were not referenced by cells or were referenced by cells from more than one region. STAR-CCM+ requires that all vertices belong to one region only.

The final section in the report displays the extents of the mesh in meters. Check these numbers to make sure that they are what you expect, especially if your pre-processor model was not in meters and some scaling was required during the export process. If the dimensions of the mesh are not correct, the entire mesh can be easily scaled.

Importing More than One Volume Mesh

More than one volume mesh can be imported into the same simulation simultaneously using any combination of the allowable formats. The process is the same as when importing a single file using the *Open* dialog, but using multiple file selections with the Ctrl key depressed. Alternatively, each mesh file can be imported one at a time in a consecutive manner. The result will be the same regardless of which method is used.

This capability is useful if you split up the pre-processing effort into multiple files, or if you want to experiment with different geometries of the same part. For example, if you are evaluating the effect of different wing mirrors on the flow over a car body, the body could be in one mesh file, and several alternative mirror designs in other files. Each simulation would then import the same mesh for the car body, and a different mesh for each mirror.

If the same file is imported twice, two copies of the mesh from that file are created, resulting in two different region definitions. You can transform one or both of the meshes so they do not occupy the same space.

Reviewing New Simulation Objects

Once the volume mesh import is complete, new region and boundary objects are added to the simulation tree. Review region names, boundary names, and region properties against the import progress report, and rename any objects that require more meaningful names.

Reviewing the Region Names

The region names that are inferred from the imported volume mesh can be seen by opening the Region node in the simulation tree:



These names correspond with the names displayed in the volume mesh import progress report. If no region name or names were provided in the input file, the material type identifier is used (for example, fluid or solid).

Reviewing the Boundary Names

The boundary names that are inferred from the imported volume mesh can be seen by opening the Boundaries node for each imported region:



Comparing the boundaries in the simulation tree with the import progress report in the output window confirms that the required boundaries were imported.

Where duplicate boundary name types are imported, the names are numbered sequentially.

Renaming Regions and Boundaries

The regions and boundaries (except for interface-type boundaries) can be renamed in the object tree. The key rule to remember when editing these names is that region names must be unique within a single simulation. Boundary names need only be unique within a single region. The uniqueness rule is case-sensitive so, for example, `Wall-1` is different from `wall-1`.

You can also rename boundaries in a scene.

Reviewing the Region Properties

The properties of each region can be viewed from the volume mesh representation in the Representations node. The read-only properties of Cells, Interior Faces, and Vertices are listed in the *Properties* window.



Reordering a Mesh

The volume mesh in STAR-CCM+ can be reordered to perform a "bandwidth reduction", where the memory references for cells and faces are relocated to be closer together, resulting in an optimized solver performance. The reordering process also attempts to provide an ordering that aligns somewhat with the expected flow pattern, which can further optimize the solver performance due to the way it updates solution variables.

Mesh reordering is done automatically whenever a volume mesh is imported or generated within STAR-CCM+. There are two circumstances when you will likely need to reorder the mesh manually:

- if the mesh topology should change, for example: boundary fusing, combining regions, splitting regions, and convert to 2D
- if the boundary types are changed from the defaults that are used during the import or mesh generation process

To reorder a given mesh region or regions, right-click on the region name or names and select Reorder Mesh from the pop-up menu.



The mesh is reordered as described above. The reordering progress for each region can be monitored from the *Output* window.

Mesh reordering is not a requirement and only minimal gains in solver speed are observed for most cases.

Initial Assignment of Boundary Types

When a surface mesh or volume mesh is imported, STAR-CCM+ checks the names of each boundary for the text that is listed in the following table. If there is a match within the imported boundary name, the boundary is assigned the boundary type given in the second column. If there is no match, the boundary is assigned a Wall type by default.

If an imported boundary contains	It is given the following type
`massflow`	Mass Flow Boundary
`inle`	Velocity Inlet Boundary
`outl`	Flow-Split Outlet Boundary
`pres`	Pressure Boundary
`stag`	Stagnation Boundary
`sym`	Symmetry Boundary
`free`	Free Stream Boundary
`axis`	Axis Boundary
(Default)	Wall Boundary

Surface Data Format and Import Reference

This section provides additional format and import details for:

- surface meshes;
- neutral format data; and
- native CAD data

Specific details on the import dialogs are provided for each format, along with background information. This information can help you determine the best format to use for your situation.

Version numbers for supported formats (where relevant) are given in the table below:

CAD Import Versions

CAD import is available on all platforms. The supported file formats and corresponding version numbers are given below.

File Format	File Extensions	Versions Supported	Add-on Required	Type of Import
CATIA Graphics Representation	.cgr	R4 to R22 V5-6R2013	CAD Exchange	B-rep
CATIA V4	.model, .exp, .session	Up to 4.2.5	CAD Exchange	B-rep
CATIA V5	.CATPart, .CATProduct	R4 to R21	CAD Exchange	B-rep
CATIA V5-6 2012	.CATPart, .CATProduct	R22, V5-6R2013	CAD Exchange	B-rep
3D XML	.3DXML	2011 to 2013	CAD Exchange	Triangulated Data
SolidWorks	.sldprt, .sldasm	Up to 2014	CAD Exchange	B-rep
JT Open	.jt	Up to 9.5	JTOpen Reader & CAD Exchange	Triangulated Data & B-rep
NX	.prt	v11 to NX 9.0	CAD Exchange	Triangulated Data & B-rep
Parasolid	.x_t, .xmt_txt, .x_b, .xmt_bin	Up to v26	None	B-rep
Pro/E	.asm, .prt	Up to Wildfire 5	CAD Exchange	Triangulated Data & B-rep
Creo	.asm, .prt	Elements/Pro 5.0, Parametric 2.0	CAD Exchange	Triangulated Data & B-rep
Inventor	.ipt, .iam	Up to 2014	CAD Exchange	Triangulated Data
IGES	.igs, .iges	5.1, 5.2, 5.3	None	B-rep
STEP	.stp, .step	AP 203 E1/E2, AP 214	None	B-rep

File Format	File Extensions	Versions Supported	Add-on Required	Type of Import
Stereo Lithography	.stl		None	Triangulated Data
VDA-FS	.vda	1.0 & 2.0	CAD Exchange	B-rep
Rhino 3D	.3dm	4, 5	CAD Exchange	Triangulated Data & B-rep

CAD Export Versions

CAD export is available on all platforms. 3D-CAD supports exporting to the following file formats.

File Format	File Extensions	Versions Supported	Add-on Required	Type of Export
Parasolid	.x_t, .x_b	v26	None	B-rep
IGES	.igs, .iges	5.3	CAD Exchange	B-rep
STEP	.stp, .step	AP 203 E1/E2, AP 214	CAD Exchange	B-rep

Supported Surface Mesh Input Formats

The currently allowable import formats that support surfaces are:

- .dbs - pro-STAR surface database mesh file
- .inp - pro-STAR cell/vertex shell input file
- .nas - NASTRAN shell file
- .pat - PATRAN shell file
- .stl - Stereolithography file
- .fro - FELISA front surface file
- .jt - JTOpen surface file
- .bif - Medina binary input file
- .3dm - Rhino 3D Model

All of these formats support single or multiple region definitions but always result in a single region definition when geometries are imported from a single file. When importing multiple files, you can choose to generate either a single region or multiple regions. For more information on the various import options, see Importing Surface Data into STAR-CCM+.

All formats other than STL and PATRAN support feature curves (or line cells). If feature curves exist in the original surfacing package, they are useful to maintain.

A summary of the surface import formats is provided below:

Format Name	File Extension	Binary ASCII	Feature Curves	Multi-Region	Boundary Names
Database	.dbs	Binary	Yes	Yes	Yes
Cell/Vert/Inp	.cell/.vrt/.inp	Both	Yes	Yes	Yes
NASTRAN	.nas, .bdf, .dat	ASCII	Yes	Yes	Yes
PATRAN	.pat, .ntl, .neu	ASCII	No	Yes	Yes
STL	.stl, .STL	Both	No	Yes	No
FELISA	.fro	Both	No	Yes	Yes
JTOpen	.jt	Binary	No	Yes	Yes
MEDINA Binary Input	.bif	Binary	No	Yes	Yes
Rhinoceros 3D Model	.3dm	Binary	No	Yes	No

pro-STAR Database Format

The pro-STAR database format is native to the pro-STAR and pro-surf pre-processing software by CD-adapco and can store either surface type or volume type meshes (or in combination if desired). STAR-CCM+, however, can only read surface meshes from this file. More than one geometry can be stored in the file, with each being identified from a numerical database ID value and an optional descriptive name. Feature curves that are supported and boundary names that are used in the pre-processing software are preserved when read into STAR-CCM+. Cell table indexes (in pro-STAR) and surface groups (in pro-surf) both produce the equivalent patch definitions when read into STAR-CCM+. The file is binary to preserve file size and is the recommended file format to use when transferring surface data from other CD-adapco pre-processing software programs to STAR-CCM+.

Surfaces that are imported from a database file are always represented as surface mesh parts.

pro-STAR Cell and Vertex Shell Input Format

The pro-STAR cell/vertex shell input format is a combination of three files which together reconstruct the surface geometry details. The three files are:

- .cel- Cell file
- .vrt- Vertex file
- .inp- pro-STAR command input file

The input (.inp) file is always ASCII and contains pro-STAR commands that define the cell table entries for the mesh, as well as commands to read in the other two files with the appropriate ASCII or binary format option. Feature curves are supported and boundary names are read from the cell table entries when available. Different cell table indexes in the cell file result in the equivalent patch definitions when read into STAR-CCM+.

Surfaces that are imported from pro-STAR cell/vertex/input files are always represented as surface mesh parts. Additionally, when exporting a surface mesh to the .inp format, any spaces within boundary names are replaced with underscores.

NASTRAN Format

The NASTRAN format is an ASCII-based file that is commonly used by a number of CAD packages and pre-processing software. This format supports feature curves and patch definitions. The current importer only recognizes shell (two-dimensional face) type elements, so volume meshes cannot be imported using this format. When preparing the NASTRAN file, it is recommended that the Free Field Format is used to export bulk data.

Surfaces that are imported from Nastran files are always represented as surface mesh parts. Surface names are preserved on import.

PATRAN Format

The PATRAN format is an ASCII-based file that is commonly used by a number of CAD packages and pre-processing software. It does not support feature curves. The current importer only recognizes shell (two-dimensional face) type elements so volume meshes cannot be imported using this format. Patch definitions are supported.

Surfaces that are imported from PATRAN files are always represented as surface mesh parts.

Stereolithography (STL) Format

STL files can be binary or ASCII and are a universal method of transferring surface data using the least possible number of triangles. As such, the quality of the surface data is usually not good and some additional work is

required before the geometry is usable. Feature curves are also not supported and all surface patches are always placed into one boundary definition.

Surfaces that are imported from STL files are always represented as surface mesh parts.

FELISA Front Surface Format

The Front Surface .fro file contains 3D surface triangulation data that is produced by the SURFACE module from the FELISA collection of programs. The SURFACE module accepts geometric curve and surface components as its input (generated by hand or from CAD software). The surface triangulation data can be used as a starting point for surface and volume meshing within STAR-CCM+.

Surfaces that are imported from Front Surface files are always represented as surface mesh parts.

JTOpen Surface Format

JTOpen is a 3D data format provided by Siemens PLM Software. It is commonly used for product visualization and data exchange. It can contain a combination of tessellated and B-rep data to represent surfaces. The STAR-CCM+ import process recognizes units specifications in the JTOpen file.

If the JTOpen file contains both a tessellation representation and B-rep representation, the *JT Import Options* box appears. From here, you can choose what representation you want to import into STAR-CCM+. The import dialog presents options that are relevant to the type of data that is contained in the file.

To import tessellated data from a JTOpen file, you require a JTOpen feature license. To import B-rep data from a JTOpen file, you require a CAD Exchange license.

When importing in Part Mode, if the JTOpen file does not contain units, you can select the desired units from the drop-down menu.

When importing in Regions Mode, you can select the units from the drop down menu. The selected units are applied to all the JTOpen files, even if they contain different units. If a hierarchy of parts of parts is being imported in the JT format, and one or more of the parts have units specified, it is essential for all units to be the same.

Tessellated surfaces that are imported from JTOpen files are represented as surface mesh parts. B-rep surfaces are represented as CAD parts.

MEDINA Binary Input File

The MEDINA binary input file format is a binary file that is used in the automotive and aerospace industries. Patch definitions, surface, and volume meshes are all supported by this format. XML (MEDINA format) files are not supported with version of STAR-CCM+ running on Windows operating systems.

The table below summarizes the element types that can be imported to, and exported from, STAR-CCM+ using the MEDINA format.

Element Type	Import Support	Export Support
Element ID	Yes	No
NODE 1 (Vertex)	Yes	Yes
PROP 464 (Property Data)	Yes	Yes
TRIA 31	Yes	Yes
QUAD 41	Yes	Yes
TETRA 61	Yes	No
PENTA 63	Yes	No
HEXA 65	Yes	No
PYRA 71	Yes	No

Surfaces that are imported from MEDINA files are always represented as surface mesh parts. Surface names are preserved on import.

Rhinoceros 3D Model File

Rhino 3D Model file format was originally designed to accurately transmit NURBS geometries between applications and is commonly used by the marine industry.

Rhino files can contain a combination of tessellated and B-rep data. If the geometry contains both types of data, you must import the geometry twice, changing the Import Type each time.

Surfaces that are imported from Rhino 3D Model files are always represented as surface mesh parts.

Supported Neutral Format Input Formats

The currently supported neutral format data files are:

- .igs/.iges- Initial Graphics Exchange Specification file

- .stp/.step- Standardized Exchange of Product file
- .vda - Verband der Automobilindustrie surface data exchange file
- .x_t/.x_b- Parasolid Transmit file

All of these formats support single or multiple region definitions in a single file. When importing multiple files as regions, the choice of generating either a single region per file or multiple regions is provided. For more information on the various import options, see Importing Surface Data into STAR-CCM+.

None of the neutral file formats support feature curves (line cells). Feature curves can be automatically created based on an angle (sharp edges) or all feature edges during the import process if desired.

A summary of the neutral format import types is given below:

Format Name	File Extension	Feature Curves	Multi-Region	Boundary Names
IGES	.igs/.iges	No	Yes	No
STEP	.stp/.step	No	Yes	Yes
VDA	.vda	No	Yes	Yes
Parasolid Transmit	.x_t/.x_b	No	Yes	Yes

IGES Format

Initial Graphics Exchange Specification (IGES) is an ASCII-based neutral data format that allows the exchange of digital information among CAD systems. It can contain models in the form of solid models as well as other types of product data. Although it has been around for almost 30 years, and has been superseded by other formats such as STEP, it is still the most widely used standard for translating CAD data.

The demarcation of parts and surfaces is determined automatically based on the bodies and patches that are contained in the file. Surfaces that are imported from IGES files are always represented as CAD parts initially.

No units specification option is provided as this information is automatically preserved by the file.

STEP Format

Standard for the Exchange of Product model data (STEP) is an international standard for the exchange of industrial product data. It was developed to provide a method for describing product data through its life cycle but is typically used to exchange data between various CAD, CAM, and CAE applications. STEP files are saved in ASCII. STAR-CCM+ currently allows the import of AP203 and AP214 (geometry only) STEP file formats only.

Surfaces that are imported from STEP files are always represented as CAD parts initially. No units specification option is provided as this information is automatically preserved by the file.

STAR-CCM+ attempts to use entity names (particularly body and face names) contained in the STEP file as part of the naming process for regions/boundaries and parts/surfaces where possible.

VDA Format

VDA is a CAD data exchange format that allows surface data to be transferred from one CAD program to another. The VDA name stands for Verband der Automobilindustrie and originated from the German automobile industry. It is also called VDA-FS in some CAD programs.

No units specification option is provided as this information is automatically preserved by the file. Surfaces that are imported from VDA files are always represented as CAD parts initially.

Unclosed surfaces in the original `.vda` file remain unclosed when imported into STAR-CCM+. Use the surface wrapper to obtain a closed surface suitable for volume meshing.

Parasolid Transmit Format

Parasolid is a geometric modeling kernel that is owned by Siemens PLM. It is used as the basis of many CAD, CAM, and CAE packages for modeling three-dimensional solid objects. Parasolid Transmit files can be binary (.x_b) or ASCII (.x_t) - the selection of the file format type is made during the export process.

No units specification option is provided as this information is automatically preserved by the file. Surfaces that are imported from Parasolid Transmit files are always represented as CAD parts initially.

STAR-CCM+ attempts to use entity names (particularly body and face names) contained in the Parasolid Transmit file as part of the naming process for regions and boundaries where possible. To allow precision to be handled correctly, all parts of the Parasolid body must fit in a box no larger

than 1000 x 1000 x 1000 units and is centered at the origin. For example, if the default units in the CAD application are set to meters, the maximum distance in any one direction that can be modeled is 1 km.

Supported Native CAD Import Formats

The type of native CAD file you can import is dependent on the hardware that you are running the STAR-CCM+ client on. The following table lists the supported options:

Native CAD Program
Creo Elements/Pro
CATIA v4
CATIA v5
CATIA v5-6 2012
NX
SolidWorks

The CAD Exchange Reader add-on license is required to import native CAD formats. All the above formats support single or multiple region definitions in a single file. When importing multiple files as regions, the choice of generating either a single region per file or multiple regions is provided. None of the native CAD formats support feature curves (line cells). You can automatically create feature curves that are based on an angle or all feature edges during the import process if desired. For more information on the various import options, see Importing Surface Data into STAR-CCM+

Surfaces that are imported from native CAD files are always represented as CAD parts initially.

Creo Elements/Pro 5.0

Creo Elements/Pro 5.0 (formerly Pro/ENGINEER) is a parametric and fully integrated 3D CAD/CAM/CAE software package from Parametric Technology Corporation (PTC). Solid models that are produced using the CAD modeler are saved using files with the extensions `.prt` (for parts) and `.asm` (for assemblies). These files can be read into STAR-CCM+ while still retaining individual part and assembly names, as well as the association between each. If your geometry models contain accelerator (`.xpr`/`.xas`) files, then these files are required to import the geometry into STAR-CCM+.

CATIA V4 Format

CATIA (Computer Aided Three Dimensional Interactive Application) is a multi-platform 3D CAD/CAM/CAE software from Dassault Systemes. CATIA V4 is principally based on piece wise polynomial surfaces. CAD models that are saved using files with extensions `.model`, `.exp` and `.session` can be read into STAR-CCM+.

CATIA V5 Format

CATIA (Computer Aided Three Dimensional Interactive Application) is a multi-platform 3D CAD/CAM/CAE PLM software from Dassault Systemes. CATIA V5 features a parametric NURBS-based modeler. CAD models that are saved using files with extensions `.CATPart` and `.CATProduct` can be read into STAR-CCM+.

CATIA V5-6 2012

CATIA (Computer Aided Three Dimensional Interactive Application) is a multi-platform 3D CAD/CAM/CAE PLM software from Dassault Systemes. CATIA V5-6 2012 features a parametric NURBS-based modeler and offers enhanced compatibility between CATIA V5 and CATIA V6 files. CAD models that are saved using files with extensions `.CATPart` and `.CATProduct` can be read into STAR-CCM+.

UniGraphics NX Format

NX is a 3D CAD/CAM/CAE PLM software from Siemens PLM Software. NX was originally called UniGraphics. NX is a parametric solid/surface feature based software that utilizes the Parasolid modeling kernel. CAD models that are saved using files with extensions `.prt` can be read into STAR-CCM+.

SolidWorks Format

SolidWorks is a 3D mechanical CAD program from the SolidWorks Corporation, a subsidiary of Dassault Systemes. It is available only for the Windows OS. SolidWorks is a parametric feature-based software that utilizes the Parasolid modeling kernel. CAD models that are saved using files with extensions `.sldprt` and `.sldasm` can be read into STAR-CCM+.

Volume Mesh Import Formats Reference

This section provides details on importing volume mesh files with the following formats:

- CD-adapco CCM/CCMG;
- CD-adapco NGEOM;
- Fluent;
- Plot3D; and
- Gridgen

In addition, the volume mesh importer can read model files relating to the following packages for further manipulation within STAR-CCM+:

- Abaqus (.inp);
- Ensight (.case);
- NASTRAN (.nas, .bdf);
- RadTherm (.ntl, .neu);
- STAR-CCM+ Boundary Data (.sbd) ; and
- Visual Toolkit (.vtk)

Currently, STAR-CCM+ is limited to importing surface descriptions for NASTRAN, RadTherm, and Visual Toolkit files.

CD-adapco CCM Format

The CCM file type is native to the STAR-CD solver (STAR) and pre-processor (pro-STAR). Since STAR-CD v4.04, the following extensions are used:

- `.ccmg` signifies a CCM file from pro-STAR containing geometry data.
- `.ccmp` signifies a CCM file from STAR (the solver) containing solution post-processing data.
- `.ccmt` signifies a CCM file from STAR containing transient post-processing data.

STAR-CCM+ can read the `.ccmg` file directly to retrieve the mesh. When reading a `.ccmp`, the corresponding `.ccmg` is also read in at the same time so that the post data can be displayed. Only the data corresponding to the first state in the `.ccmp` file is available. In the case of transient `.ccmt` files, rename them as `.ccmp` files, which allows the first time step saved on the file to be read in. Other time step values in the same file are available.

Mesh data is exported to the file by pro-STAR using the command GEOM. In the case of the `.ccmp` and/or `.ccmt` files, solution data is saved by the STAR solver directly (as part of an analysis run). pro-STAR can also write solution mapped data (command SMAP) to the `.ccmp` file as part of a restart.

A scale factor for the mesh is saved to the file during export from pro-STAR but the vertices that are contained in the file are not scaled. The scale factor is used to redimension the mesh during the import process into STAR-CCM+.

Each collection of cells in pro-STAR with the same properties (material label, material index, porosity index, spin index and group index) is considered to be one region. Different cell types do not result in different region definitions unless one of the above properties differs for the cell type.

The CCM file format is the recommended way of transferring meshes and post data from STAR-CD to STAR-CCM+. STAR-CCM+ can also export both mesh and solution data to STAR-CD using the CCM file format, although not all solution variables are transferable.

CD-adapco NGEOM Format

The `.ngeom` file type was originally designed for use with the COMET solver, but is now replaced by the CCM format. The `.ngeom` format is exported by pro-STAR using the NGEOM command and can only contain mesh data.

If the appropriate scale factor was applied during the export process, the vertices that are contained in the file conform to SI and no further action is required on importing the file into STAR-CCM+. The scale factor that is used is also saved to the file but is not used. Boundary types are inferred from boundary names, and if no inference can be made, then a boundary is assigned the type wall.

Each collection of cells with the same properties (material label, material index, porosity index, spin index and group index) in pro-STAR is considered to be one region. Cell types do not affect the region definition unless one of the above properties differs.

Fluent Format

Fluent format files can have the following extensions:

- `.cas` - Fluent case file
- `.msh` - Gambit
- `.grd` - TGrid

These files have the same format internally, but can be written in text or binary form. STAR-CCM+ automatically recognizes a file as being binary or text and imports it without requiring further action.

Also see the guidelines for migrating from Fluent.

Plot3D Format

Plot3D files have various formats and extensions. STAR-CCM+ currently recognizes this file type only with the following extensions:

- `.grd`;
- `.p3d`;
- `.xyz`; or
- `.x`

and with the following characteristics:

- three-dimensional;
- single-block or multi-block; and
- FORTRAN-unformatted or C-binary

There is no need to specify the characteristics of the file; the importer recognizes the file type and also determines whether byte-swapping is necessary. STAR-CCM+ currently does not support two-dimensional, unstructured (tetrahedral), or formatted (ASCII) Plot3D files. Double precision files are now supported.

When you import a Plot3D file, STAR-CCM+ can read boundary definitions from a `.fvbnd` file at the same time, if the `.fvbnd` file has the same base file name and is in the same directory as the Plot3D file.

Gridgen Meshes

Gridgen can export Fluent format files with the `.cas` extension and STAR-CCM+ files with the `.inp` extension. Files with the `.cas` extension have the same format as a Gambit `.msh` file and as such can be renamed to have this extension. These formats are the best route from Gridgen to STAR-CCM+. It is important to ensure that the mesh is specified as right-handed in Gridgen before it is exported.

Although Gridgen can export STAR-CD cell/vertex/boundary/input files also, they are not compatible with volume mesh import in STAR-CCM+. If you have access to pro-STAR, the files can be read in using the IFILE command and then resaved as a CCM file.

Importing STAR-CD Mesh and Solution Data

It is possible to import mesh and solution data from STAR-CD in order to post-process the solution data using the STAR-CCM+ visualization tools. This section describes the process.

Since the STAR-CD v4.04 release, the solution data provided by STAR (the solver) has been stored in two files: the `.ccmp` file for standard (restart) solution data, and the `.ccmt` file for transient solution data. Both of these files contain a link to the corresponding `.ccmg` file for the mesh data; if this `.ccmg` file is missing, it is not possible for STAR-CCM+ to read either the `.ccmp` or the `.ccmt` file. For a transient file, STAR-CCM+ only reads the first set of solution data in the file.

An example of importing solution data is provided in the following section.

Example of Importing STAR-CD Solution Data

The procedure for importing both mesh and solution data from the STAR-CD solver is described as follows (assuming that the necessary STAR-CD files have been copied into the working directory):

- Launch STAR-CCM+ and start a new simulation in the manner appropriate to your environment.

- Select File > Import > Import Volume Mesh... as shown.



- In the *Open* dialog, select the desired `.ccmp` or `.ccmt` file. Ensure that the corresponding `.ccmg` file is also available.



- Click Open.

If solution data is successfully read from the file, information on the process is displayed in the *Output* window. This information includes details of the solution data found in the file, and details of the scalar fields in STAR-CCM+ to which the solution has been mapped. An example of this output is shown below.

```
CCM file version 20617

Cell type[1] label: "Default_Fluid", material: "fluid" (1)
porosity: 0 spin: 1 group: 0
Boundary type[0] id: 0, name: "Default_Boundary_Region",
label: "", type: "wall"
Boundary type[1] id: 1, name: "Inner_inlet", label: "",
type: "inlet"
Boundary type[2] id: 2, name: "Outer_inlet", label: "",
type: "inlet"
Boundary type[3] id: 3, name: "", label: "", type:
"pressure"
Boundary type[4] id: 4, name: "", label: "", type:
"symmetry plane"
Reading solution from state Restart_1
Read 'Velocity Component U' (SU) m/s  at 2000 cells
Read 'Velocity Component U' (SU) m/s  at 960 faces
```

```
Read 'Velocity Component V' (SV) m/s  at 2000 cells
Read 'Velocity Component V' (SV) m/s  at 960 faces
Read 'Velocity Component W' (SW) m/s  at 2000 cells
Read 'Velocity Component W' (SW) m/s  at 960 faces
Read 'Mass Flux' (MassFlux) kg/s  at 5500 faces
Read 'Mass Flux' (MassFlux) kg/s  at 960 faces
```

... and so on ...

```
Processing mesh:
Mapping imported data DENS to Density
Mapping imported data ED to Turbulent Dissipation Rate
Mapping imported data H to Enthalpy
Mapping imported data P to Pressure
Mapping imported data SU to Velocity-X
Mapping imported data SV to Velocity-Y
Mapping imported data SW to Velocity-Z
Mapping imported data T to Temperature
Mapping imported data TE to Turbulent Kinetic Energy
Mapping imported data VIS to Turbulent Viscosity
```

It is observed that solution data is read for both cells and faces during the import process. This corresponds to cell and wall data in STAR-CD.

Importing CAE Models

Imported models are CAE models that have been prepared in other CAE products, for example an FEA code, and imported into STAR-CCM+. The model can be formed from surfaces, shells, beams and/or volume elements. When a CAE model is imported, STAR-CCM+ recognizes geometrical elements that are defined in the input file, and creates a new region from each volume that is defined in the CAE model. If the CAE model contains any surface definitions, these create new boundaries.

From version 9.02 of STAR-CCM+, a CAE model can be imported as a Region, or as a Region And Imported Model. The first option allows you to import a CAE model into STAR-CCM+ as a region that can then be used in the same way as any other region in the simulation. The second option allows you to couple STAR-CCM+ to other codes using the file-based coupling method: the CAE model is still imported as a region, but an extra node called Imported Models is added to the simulation tree. From this node, you can map solution data from the STAR-CCM+ model onto the CAE model, and export it in a format that can be read into the third-party software. Similarly, you can import solution data that is associated with the CAE model into STAR-CCM+, and use it to define a boundary condition in the simulation.

Simulation files that contain imported CAE models, and were created with previous versions of STAR-CCM+, are compatible with the import process introduced in 9.02.

An imported CAE model can be visualized in a scene, in the same way as regions, boundaries, and parts.

STAR-CCM+ supports the import of CAE models containing shells, surfaces and/or volumes from the following codes:

- Abaqus
- ANSYS
- Nastran
- RadTherm (shells only)
- es-ice or pro-STAR

It is also possible to import a CAE model in the STAR-CCM+ boundary format (`.sbd`). These files are used in the same way as CAE models imported from the codes that are listed above, but you cannot import them as regions. STAR-CCM+ also supports the import of Ensight input files, which can be used to visualize and compare multiple solutions in one simulation.

For Nastran models, you can import data that was defined in local coordinate systems. STAR-CCM+ transforms the data into a global coordinate system during the import process. All other formats are imported in a global coordinate system, as importing from local coordinate systems is not supported.

What Is the Imported Models Node?

When the Import Mode is set to Create New Region(s) and Imported Model(s), STAR-CCM+ creates a new region for every imported volume, shell, or beam, and boundaries for the imported surfaces. STAR-CCM+ also adds a new node called Imported Models to the simulation tree.

Every CAE model that is imported in the simulation is represented by a node within the Imported Models node. These imported CAE model nodes, located within the Imported Models node, provide an interface for file-based coupling.

The node layout for an imported CAE model is as follows:

- Imported CAE model node
  - Imported Surfaces node
    - Imported surfaces
  - Imported Elements node
    - Imported elements (includes shells, beams, and volumes)

The Imported Models node has no properties but is subject to common menu items.



What Is the Imported CAE Model Node?

The imported CAE model node, located within the Imported Models node, provides an interface for file-based coupling. From an imported CAE model node, it is possible to map solution data from the STAR-CCM+ model onto the CAE model, and export it in a format that can be read into the third-party software. You can also import solution data for a CAE model, computed by the third-party code, map the data onto the STAR-CCM+ model, and use the data to define boundary conditions, or as an input to the morpher.

STAR-CCM+ automatically names the imported CAE model node according to the following convention:

```
<File Type>: <File Name>
```

For example, when the Abaqus input file `exhaust.inp` is imported, the corresponding imported CAE model node is named `Abaqus: exhaust`.



When a CAE model is imported, STAR-CCM+ recognizes geometrical elements that are defined in the input file, and creates corresponding volume, shell, or beam regions. The imported surfaces, volumes, shells, and beams are also listed under the relevant imported CAE model node.



The imported CAE model features are fully stored in STAR-CCM+ under the Regions and Representations nodes, whereas the imported CAE model node simply provides an interface for file-based coupling operations. After

the import process, the imported CAE model nodes are connected to the corresponding regions. Deleting a region associated with an imported CAE model deletes the imported model itself. The regions created for the imported CAE models are subject to all the operations that apply to regions in STAR-CCM+, for example, transform operations such as rotation and translation, mesh reordering, and so on. When these operations are performed, the imported CAE model updates accordingly.

The imported CAE model node is used for file-based coupling, which is carried out using the pop-up menu. The node has no properties.

Imported CAE Model Pop-Up Menu

Import CAE Solution Data...		Imports data from an external file that is associated with the imported CAE model, such as an Abaqus `.odb` file.
Map Data	Surface Data...	Maps surface data between CAE model surfaces or vertices and STAR-CCM+ boundary faces or vertices.
	Volume Data...	Maps volume data between CAE model vertices and STAR-CCM+ region cells or vertices.
	Beam Data...	Maps beam data between CAE model vertices (beam nodes) and STAR-CCM+ boundaries.
Export Mapped Data...		Exports mapped data to a separate file for use in an analysis with the other software.
Deform Using Imported Surface Displacements		Deforms the CAE model mesh based on imported surface vertex displacements.

See also: common menu items.

What Is the Imported Surfaces Node?

When a CAE model is imported, STAR-CCM+ automatically recognizes the surfaces that have been defined in the CAE model, and defines corresponding boundaries. The imported surfaces are also listed under the Imported Surfaces node.



The Imported Surfaces node has only one property called Imported Surfaces. This property is read-only, and its value specifies the number of surfaces imported into STAR-CCM+.



Each imported surface is represented by a node within the Imported Surfaces node. The nodes representing each imported surface have two read-only properties: Index, which identifies the surface with a unique number, and Boundaries, which specifies the region boundaries that are associated with the imported surface.



The imported surface nodes can be selected in the relevant dialogs when mapping/exporting solution data from STAR-CCM+ to the third-party CAE package, or when importing/mapping CAE model solution data to the STAR-CCM+ model. These processes are used for file-based coupling and are carried out using the pop-up menu on the imported CAE model node.

What Is the Imported Elements Node?

When a CAE model is imported, STAR-CCM+ automatically recognizes the shell, beam, and volume elements that have been defined in the CAE model, and creates corresponding regions and boundaries. Shell, beam and volume elements are also listed under the Imported Elements node.

The Imported Elements node has only one property called Imported Elements. This property is read-only, and its value specifies the number of elements imported into STAR-CCM+.



Imported Shell Elements

Shell elements have a positive and a negative side. When a CAE model is imported into STAR-CCM+, a shell element defines a region, and the surfaces of the shell element are stored as boundaries within that region. Each shell element is also represented by a node within the Imported Elements node.



The nodes representing each imported shell have two read-only properties: Index, which identifies the shell with a unique number, and Regions, which specifies the region that is associated with the imported shell.



These shell element nodes contain a Boundaries sub-node, in which the two surfaces of the shell are listed. The surface that corresponds to the positive facing side of the shell are suffixed with POS, SPOS, or Front, depending on the input format. Similarly, the surface that corresponds to the negative facing side of the shell are suffixed with NEG, SNEG or Back. The positive direction is found by applying the right-hand rule to the vertex numbering on a face. If the fingers of your right hand are wrapped around the face in the direction in which vertex numbers increase, then your thumb points in the positive direction. The nodes representing the shell surfaces listed under the [shell element] > Boundaries node, have two properties: Index, which

identifies the shell surface with a unique number, and Boundaries, which specifies the region boundaries that are associated with the imported shell surfaces.



The shell surfaces can be used for file-based coupling in the same way as imported surfaces. When mapping onto shell surfaces, it is essential to ensure that you map to the correct side of the shell. If in doubt, activate the Check Source/Target Face Normals in the *Nearest Neighbor Search Constraints* group box, that is located on the *Map Imported Model Data* dialog. You can also visualize the direction of normals on the different surfaces. The normals of the source and target faces face each other.

Imported Beam Elements

Currently, STAR-CCM+ supports the import of linear beam elements from Abaqus (.inp), Nastran (.bdf and .dat), and ANSYS (.cdb).

Beam elements are one-dimensional objects in three-dimensional space, and are often used to model 3D structures in which one dimension is greater than the other two dimensions. This approximation is based on the fact that the relevant physical quantities are mainly functions of the position along the dominant dimension, that is, the length of the structure. The cross section of a beam element defines the beam profile and its dimension is negligible compared to the beam length.

Beam elements have a two or three node representation, depending on the order of interpolation used for its finite element formulation. A beam element deforms under the application of either a bending moment, axial force, torque or a combination of these loads.

When a CAE model is imported into STAR-CCM+, the body and ends of a beam element are stored as two individual boundaries within the relevant region node. The name of the boundary that corresponds to the beam body is suffixed with LINE; the name of the boundary that corresponds to the beam ends is suffixed with ENDS.

Each beam element is also represented by a node within the Imported Elements node.





The nodes representing imported beams, within the Imported Elements node, have two read-only properties: Index, which identifies the beam with a unique number, and Regions, which specifies the region that is associated with the imported beam.



These beam element nodes contain a Boundaries sub-node, which contains a single boundary. This boundary represents both the ends and body boundaries that are located under the associated region node.

The boundary node under the [beam element] > Boundaries node, has two properties: Index, which identifies the beam boundary with a unique number, and Boundaries, which specifies the region boundaries that are associated with the imported beam.



It is possible to map fluid traction data (pressure and shear forces) from the wetted surface in STAR-CCM+ onto the imported beam element as forces, and export the mapped field functions to an external file that can be read into the third party software. It is also possible to map nodal displacement and rotation data from the imported beam to the wetted surface in STAR-CCM+. Displacement and rotation data can then be used by the mesh morpher to deform the model. These file-based coupling operations are carried out using the pop-up menu for the imported CAE model node.

Imported Volume Elements

When a CAE model is imported into STAR-CCM+, volume elements are stored as regions under the Regions node. The extracted surface of the volume element is stored as a boundary under the relevant region Boundaries node.



Each volume element is also represented by a node within the Imported Elements node. The name of the node depends on the names that are specified in the input file, for example manifold-1 in the image below.



The nodes representing each imported volume have two read-only properties: Index, which identifies the volume element with a unique number, and Regions, which specifies the region that is associated with the imported volume.

Volume element nodes also contain a Boundaries node, under which the extracted surface of the volume element appears. The extracted surface node has two properties: Index, which identifies the volume surface with a unique number, and Boundaries, which specifies the region boundaries that are associated with the extracted surface.



It is possible to map region temperature data from STAR-CCM+ onto an imported volume via the imported CAE model node, and export this temperature to an external file. This file can be read into the third-party software and used to define a temperature field throughout the volume. This operation is used for file-based coupling and carried out via the pop-up menu for the imported CAE model node. Upon import, STAR-CCM+ automatically extracts the free surface of a volume element. The free surface can be used in the same way as an imported surface to import, export, and map solution data, and also for visualization.

Importing a CAE Model

To import a CAE model:

1. Select File > Import > Import CAE Model...



2. In the *Open* dialog, select the file to import. If necessary, adjust the *Files of Type* drop down list so that it is set to the format of your file.



3. Click Open.
4. In the *Import Options* dialog, set the Import Mode. You can choose to import the CAE model as a region and imported model (default), or as a region only.



5. The Units menu in the *Import Options* group box can be used to specify the units of the CAE model.
6. Click OK.

The file is imported after a brief period. The progress is shown in the *Output* window.



Note: The import procedure is different when importing a CAE model from pro-STAR or es-ice.

Importing a CAE Model as a Region

Importing a CAE model as a region causes STAR-CCM+ to create a new region from each volume that is defined in the CAE model. If the CAE model contains any surface definitions, these create new boundaries. Surfaces that are not included in any surfaces definitions form a default boundary.

To import a model as a region, follow the import procedure for a CAE model. In the *Import Options* dialog, set Import Mode to Create New Region(s).



When a CAE model is imported as a region, it can be used in the same way as any other region in the simulation.

Importing a CAE Model as a Region and Imported Model

When importing a CAE model as a region and imported model, STAR-CCM+ creates a new region from each volume that is defined in the CAE model, and adds the Imported Models node, containing the relevant imported CAE model node, to the simulation tree. This import option allows you to access the Imported Models framework for file-based coupling operations.

To import a model as a region and imported model, follow the import procedure for a CAE model. In the *Import Options* dialog, set Import Mode to Create New Region(s) and Imported Model(s).



A new top-level node, Imported Models, appears in the simulation tree.



Open the Imported Models node to visualize the imported CAE model node. The following image shows the imported CAE model node that is created when importing an Abaqus input file called `Manifold.inp`.



The imported CAE model node typically shows two other manager nodes, Imported surfaces and Imported elements. These nodes show the imported surfaces and elements that are defined in the CAE model file.

Note: When you import a CAE model containing a volume, STAR-CCM+ automatically extracts the free surface of the imported volume. When you import a model containing shells, STAR-CCM+ automatically locates the positive and negative side of each shell.

Visualizing an Imported CAE Model

To display the surface mesh geometry of the imported CAE model, right-click the Scenes node and select New Scene > Mesh.



A new scene appears with the mesh geometry appearing.



Select the Mesh Scene 1 > Displayers > Mesh 1 > Parts node.



Depending on the type of data you have imported, there may be parts in addition to the surfaces that you want to display. In the *Properties* window, click in the right half of the Parts property. In the selection dialog, remove all the parts that may currently be selected by clicking Select None. Then expand the Imported Models node and tick the CAE model elements that you wish to display (EXTERIOR and INNER in the following screenshot). Click OK.



The imported surface appears in the mesh scene.



Visualizing Shell Normals

When shells are imported in a CAE model, STAR-CCM+ automatically locates the positive and negative sides, or front and back, of each shell. It is often useful to be able to check which side of the shell is which, and this can be achieved by visualizing the face normals of each shell side.

First, create a new vector scene, and add one of the shell sides to the Vector displayer. One method for doing this is to click-and-drag the Imported Elements > shell > Boundaries > shell side node onto the scene. When

you release the mouse button, select Add to Vector 1 in the pop-up menu. In the screenshot below, the negative side of the manifold shell, Manifold: NEG, is added to the scene.



In the scene, set the Vector Field to Area. Vector glyphs are displayed on the shell surface. The direction of the vectors shows which side of the shell the surface is on. For example, in the following screenshot, the negative surface

is on the inside of the manifold shell. Had you displayed face normals on the positive surface of the shell, the arrows would be facing outward, indicating that the positive surface is on the outside of the manifold shell.



In the above screenshot, the gray surface was activated by adding the same negative shell surface, Manifold:NEG, to a geometry displayer and enabling surfaces on the displayer.

Exporting from STAR-CCM+

What you create in STAR-CCM+ can also be used in other codes. It is possible to export geometry after drawing 3D-CAD designs or generating mesh surfaces. After you run a simulation, the mesh and solution data that is produced can be exported to a variety of formats. Additionally, you can control which data gets exported based on scalar or vector variables, regions and boundaries, and/or derived parts.

Exporting Geometry

Separate sections provide instructions and format information for exporting the following:

- 3D-CAD geometry
- Mesh surfaces

Exporting Mesh and Solution Data

STAR-CCM+ supports exporting mesh and solution data to other packages via the File > Export ... menu item. This menu item opens a modified form of the operating system-specific standard *Save To File* dialog. In this dialog you can specify the file type, the name of the new file, and which types of data to write to the file.

Use the lists in the Vector and Scalar tabs in the right portion of the dialog to select solution variables for export. If your simulation has multiple regions, you can also select the regions for data export in the Regions tab. (A Boundaries tab replaces the Regions tab when you select any of the following formats from the *Files of Type* drop-down list: Boundary Shell, STAR-CCM+ Boundary Data, and RadTherm Patran Neutral). For certain formats, if your simulation has one or more derived parts, you can export their data by selecting them in the Derived Parts tab. The Derived Parts tab is available only when the `Ensight Gold files (*.case)` or `Surface fft (.trn)` file type is selected.

More than one variable can be selected using `<Ctrl><Click>` to select discontinuous items and `<Shift><Click>` to select a continuous range of items. To deselect an item in a list, use `<Ctrl><Click>`.

You can choose the scope of the data to select from the drop-down list:

- Mesh only — this option causes the lists in the Vector and Scalar tabs to be unavailable.
- Mesh and solution data
- Solution data only

Use the *Files of Type* drop-down list to choose the output format. The currently supported export formats include:

- CCM — If you select the CCM format, you can activate the Append to existing file checkbox to add the exported information to a `.ccm` file that you select. For example ESI/VA One, FFT/Actran, and LMS/Virtual.Lab, which can be used in noise propagation cases, can read this format.
- Ensight Gold — To export vertex data, activate the Export Results on Vertices checkbox. This option is only available when your export includes solution data.
- FIELDVIEW — When exporting to this format from a parallel session, STAR-CCM+ includes grid partition information in the exported file. A parallel session of FIELDVIEW can use this grid partition information to distribute the data among the processes available to it.
- Tecplot 360 (Version 2008 and newer)

- Boundary Shell — The Boundary Shell option actually creates multiple files of different formats: `.cel` and `.vrt`. A third format, `.daten`, appears when you specify data by selecting scalars, vectors and/or boundaries.
- STAR-CCM+ Boundary Data (.sbd)
- RadTherm Patran Neutral
- Surface FFT files

Finally, enter the output file name and click Save.



It is also possible to export automatically at regular intervals.

Automatically Exporting Mesh and Solution Data at Regular Intervals

When running a long simulation (either interactively or in batch mode) it is useful to enable the auto export feature. This feature causes STAR-CCM+ to export the solution data automatically at regular intervals during the run, when the Export... command is disabled. This section includes a description of the properties of the *Auto Export* dialog.

To set up the auto-export of solution data:

- Select File > Auto Export... from the menu.



The *Auto Export* dialog appears.



- Activate the auto-export feature:
  - Use the node handle to expand the Update property.
  - Activate the Enabled checkbox.
- To specify the file format that you want to export to, choose from the drop-down list in the Values column of the File Format property.
- Specify the Base Name. The base name is the prefix of the name of the file that you want to export your data to.

- Specify the Export Directory if necessary. This directory is the location of the file that you want to export your data to. The default path is the same as the path of your simulation file. Clicking the property customizer (...) activates a dialog similar to the *Save* dialog. You can use this dialog to navigate to a new location for the file.
- Enter the type of separator that should be used in the Separator property. By default, when STAR-CCM+ saves copies of the simulation file, it places a @ character to separate the base file name and the number of the trigger event. If the use of this character in this manner causes any problems with your system, you can enter a different one as the separator. If you attempt to leave this field blank, STAR-CCM+ puts the @ character there by default.
- If you want to add the exported data to an existing file, activate the Append To File checkbox.
- If you prefer not to export the mesh data, activate the Export Solution Data Only checkbox.
- To export vertex data, activate the Export Data on Vertices checkbox. This option is only available when your export includes solution data, and when you are exporting to an Ensight Gold file.
- Select the scalar values that you want to export: click the property customizer (...) next to the Scalar Functions property.

This action opens an object selection dialog (the dialog in the screenshot below is called *Scalar Functions*).



- Select the desired scalar functions and click OK.

- Using the same method as defining the scalar functions, select the vector functions, regions, boundaries, and derived parts that you want.
- Select the Trigger Type from the drop down menu. This property specifies how the auto export operation is triggered. See Trigger Types. Unless you are running a transient analysis, leave the setting at `Iteration`.
- To control the number of intervals that pass between auto exports, expand the <Trigger> Frequency node and enter the value that you want. See Trigger Properties.
- Click Close.

STAR-CCM+ automatically exports during the next run at the intervals you specified. To deactivate this feature, simply reopen the *Auto Export* dialog and deactivate the checkbox of the Enabled property.

Auto Export Properties

You can use the following properties to activate and customize this feature:

File Format	Selects the format of the file from a drop-down list.
Base Name	The prefix of the file name to use when exporting your data.
Export Directory	The directory that contains your file with exported data.
Separator	Specifies the character that appears between the base file name and the number of the trigger event (`@` by default).
Format Width	Width of the iteration, time, or step number in the formatted file name.
Append to File	This option, when its checkbox is activated, adds the exported information to an existing file. This file would be the one you specify in the Base Name, Export Directory and File Format properties. This option works for `.ccm`, Ensight and `.trn` file formats only.

Export Solution Data Only	This option, when its checkbox is activated, exports solution data without any mesh data.
Export Data on Vertices	This option, when its checkbox is activated, exports vertex data. This option is only available when your export includes solution data, and when you are exporting to an Ensight Gold file.
Scalar Functions	Determines which scalar values to export.
Vector Functions	Determines which vectors to export.
Regions	Selects regions to export. See Exporting Region and Boundary Data.
Boundaries	Selects boundaries to export. See Exporting Region and Boundary Data.
Derived Parts	Selects derived parts to export.
Update	When expanded, this property makes available the controls listed below.
Enabled	Enables or disables the auto export feature. Activated — Enabled Deactivated — Disabled
Trigger	Provides choices for triggering the automatic export. See Trigger Types.
Iteration Frequency	This property controls the number of *trigger events* between exports. See Trigger Properties.

The settings are saved in the simulation database.

Exporting Region and Boundary Data

The following rules apply when exporting region and/or boundary data:

- If no regions and no boundaries are selected, data for all the regions (including the boundaries) in the simulation is exported.
- Selecting a region automatically selects the boundaries that are associated with that region. Data for the selected region (including the associated boundaries) is exported. This action has the same effect as manually selecting a region and the boundaries in that region.
- Selecting a boundary does not automatically select the region that is associated with that boundary. Data for the selected boundary is exported.
- Selecting boundaries without also selecting regions results in a `.ccm` file that contains only boundary data. This option is not compatible with pro-STAR.

Using File-Based Coupling

STAR-CCM+ can be coupled with external packages using the file-based coupling method. This allows you to use the STAR-CCM+ solvers to obtain part of the solution, and another CAE code to obtain another part of the solution. This technique can be used in cases such as fluid-structure interaction (FSI), where STAR-CCM+ would solve for the fluid, a third party code would solve for the structure, and data would be exchanged between the two codes a the solid-fluid interface. This section includes the following:

- Understanding file-based coupling
- Mapping STAR-CCM+ solution data to a CAE model
- Exporting mapped data
- Mapping and Exporting Transient Surface Data
- Importing CAE model data
- Mapping CAE solution data to a STAR-CCM+ model
- Using data mappers
- Details for specific CAE products

Understanding File-Based Coupling

File-based coupling is a method used to couple STAR-CCM+ with another CAE code via external files stored in the computer memory. The method allows you to obtain a solution using separate CAE codes to solve for each part of the model, such as the fluid and solid domains in an FSI problem. The powerful set of tools available in STAR-CCM+ allow you to import a CAE model into STAR-CCM+ and carry out a one-way or two-way coupled simulation.

For one-way coupling, solution data from STAR-CCM+ can be mapped onto the CAE model and exported in a format that can be read into the third-party application. The exported data can then be used to specify boundary or operating conditions in the CAE model. An example of this is using STAR-CCM+ to compute the conjugate heat transfer solution in the fluid and solid domains of a model, then exporting the solid region temperature solution to an external file. This file can then be read into a third party code, and used to calculate the thermal stresses in the solid. An example of the typical workflow for one-way coupling is shown in the diagram below (solid blue arrows).

For two-way coupling, you can continue by importing solution data from the third-party code into STAR-CCM+. This data can then be mapped onto the STAR-CCM+ model and used to define a boundary condition in the simulation, or to update the mesh. An example of two-way coupling is using STAR-CCM+ to compute the pressure and wall shear stress at a fluid-solid interface, using this data to define a traction load on the FSI surface in the CAE model, and calculating the resulting displacements using the third-party code. The nodal displacements in the solid model can be imported into STAR-CCM+ and used as an input to the mesh morpher. The STAR-CCM+ simulation can then be run using the updated geometry, and the whole process can be repeated until the solution has converged. The

typical workflow for two-way coupling is an extension of that for one-way coupling. The additional steps for two-way coupling are shown in the diagram below (dashed gray arrows).



An important part of file-based coupling is the mapping process, and this is discuss in greater detail in the Understanding Data Mapping section.

Understanding Data Mapping

Mapping data is an important part of the file-based coupling method, as it allows you to use data computed in STAR-CCM+ to specify a boundary or field condition in a third-party CAE code. Similarly, you can use data computed in a third-party code to set boundary conditions in the STAR-CCM+ model. Mapping is the process of interpolating a set of data between two separate models, for example pressure data from a fluid domain can be mapped onto the corresponding surface of a solid domain

and used as a traction load. Typically data is mapped at an interface between two models, and this may involve interpolating data between two surfaces, a surface and a set of vertices, or between a region and a volume. Mapping presents several challenges to the user, such as how to overcome differences in mesh density or topology at the interface between the two models, or how to preserve the accuracy of the solution. These challenges, along with the type of mapping algorithms used in STAR-CCM+ and the typical workflow, are discussed below. For details on how to carry out mapping operations, refer to the following chapters:

- Mapping STAR-CCM+ solution data to a CAE model
- Mapping CAE solution data to a STAR-CCM+ model

All of the formats in STAR-CCM+ file-based coupling use the same core interpolation algorithms to map data from one surface to another. In the current version of STAR-CCM+ the accuracy of the mapping techniques has been improved from zeroth-order to at least first-order. Mapping in STAR-CCM+ is carried out by performing operations via the object tree as described below.

Mapping Operations

A key operation in coupling with imported models is to map data between STAR-CCM+ boundaries and/or regions and imported surfaces and/or volumes. It is also possible to map data between STAR-CCM+ boundaries and imported beam curves. The pop-up menus for imported surfaces, beams and volumes present several map actions. The mapping functionality is encapsulated into Data Mapper objects. Each of these map actions creates a Data Mapper object that is listed under the Tools > Data Mappers node for you to modify or reuse. Depending on whether the map action is associated with surfaces, beams, or volumes, a Surface Data Mapper, Beam Data Mapper or Volume Data Mapper is created.

The map action dialog (*Map Imported Model Data*) allows you to choose a STAR-CCM+ boundary/region, an imported surface/volume, and a field function to map between the source and target boundaries/regions. In addition, you can specify certain parameters to control the underlying data mapper. These options include proximity tolerance and normalcy tolerance (only for the surface data mapper). A description of these parameters can be found under Data Mapper Expert Properties. Your selections from the map action dialog are automatically populated in the created data mapper object, thereby allowing you to achieve the same mapping either by creating mappers as described here or by manually creating them. Moreover, in case you perform subsequent mappings between the same surface or volume pairings, the interpolation stencil will be reused, thereby optimizing the computational time.

Interpolation Techniques

The interpolation techniques are based on either the least squares method or shape function method. The choice is based on whether the source data is face-centric or node-centric. For instance, consider a conjugate heat transfer (CHT) application in which the thermal analysis of the solid is being done using ABAQUS and the flow/thermal analysis of the fluid is being performed using STAR-CCM+. The coupling of these analyses takes place in two steps:

1. Mapping the heat transfer coefficient and reference temperature from the STAR-CCM+ fluid boundary to the ABAQUS solid surface.
2. Mapping the surface temperature from the ABAQUS solid surface to the STAR-CCM+ fluid boundary.

In Step 1, the STAR-CCM+ boundary is the source surface and the source data is face-centric. The target points are the ABAQUS face centers. In this case, we use least squares interpolation to map boundary face-centered data to ABAQUS face centers. In Step 2, the solid surface temperatures are imported from ABAQUS. In the above-mentioned CHT case, the ABAQUS surface is the source surface and the source data resides on the nodes. The shape functions for the ABAQUS elements are used to interpolate functions to map data from the ABAQUS nodes to STAR-CCM+ boundary face centers. The shape functions are implemented for linear/quadratic tri/quad elements only.

Meeting Challenges

One of the challenges in mapping data is to identify the matching faces or surfaces on the source and target sides. In cases of complex applications like the under-hood of an automobile, heat transfer between solid and fluid takes place across many common surfaces attached to many different parts. These surfaces and parts vary in size, shape and thickness. Depending on the analytic requirement, the parts are grouped together differently. For instance, the fluid region can group many of the exposed surfaces together in doing the CFD analysis. However, the solid region needs to treat each of the parts separately to account for different material properties and to analyze the thermal stresses in each of them. Moreover, the discrete representation of the common surfaces can be drastically different in the solid and fluid regions. The mesh topology and mesh density on the solid and fluid surfaces can also be very different. For instance, STAR-CCM+'s mesh engine generally creates polygonal faces and cells, while most of the finite element packages operate on tri/quad elements of various orders. In addition, the CFD analysis generally requires higher mesh densities (than solids) to capture boundary layers. Depending on the quality of mesh

generated, the geometric detail of common surfaces can be compromised. These challenges can be addressed by using the additional mapping options.

Mapping STAR-CCM+ Solution Data to a CAE Model

Solution data from a STAR-CCM+ model can be mapped onto a CAE model for use with third-party CAE software. This could involve using STAR-CCM+ to compute the flow solution in the fluid domain, then mapping and exporting traction data to the third party software to compute strains in the solid. Another example would be to use STAR-CCM+ to compute a full CHT (conjugate heat transfer) solution across the fluid and solid domains in the model, then map and export the solid region temperature data to a third-party code, and hence compute the thermal stress solution.

To map STAR-CCM+ solution data to a CAE model, first you need to import the CAE model into STAR-CCM+. It is then possible to map solution data from STAR-CCM+ onto faces, vertices, or volumes in the CAE model. Following this, you can export the mapped data, and use it in a third-party code.

Mapping Boundary Data to Imported Surfaces

To map boundary data to the faces or vertices of an imported surface, right-click the imported CAE model node under the Imported Models node and select Map data > Surface Data... This menu option is also used to map

imported CAE data to the STAR-CCM+ model.



This will activate the *Map Imported Model Data* dialog. Use the dialog as follows:

- Ensure the *Map Direction* is set to STAR-CCM+ to Imported Model.
- In the *Mapping Inputs* box, select the STAR-CCM+ boundaries you wish to map from, and the imported model surfaces you wish to map to.
- The *Data to map* menu allows you to select which field data will be mapped. This menu displays field functions that are available for mapping in the simulation. When you select the field data, the *Imported Data Type* and *Boundary Data Type* will be set automatically. If you are mapping local heat transfer coefficients, you will also be able to select a reference temperature evaluation option.
- In some cases the mapping algorithm may fail to find appropriate neighbors for each face on the source boundary. If this happens, use the additional constraints provided in the Nearest Neighbor Search Constraints box at the bottom of the dialog. Turning on these constraints also enables you to check the success of the mapping operation using the mapped flag field function.

In the following screenshot, the settings specify that static pressure data will be mapped from the Fluid Wall boundary of the STAR-CCM+ model to the INNER surface of the imported Abaqus model.



When the settings in the dialog are complete click Apply, then Close.

The time taken to carry out the mapping operation will be displayed in the *Output* window, followed by `Done`. A field function will be exposed corresponding to the mapped data, for example the `Mapped Pressure` field function will be exposed if you map `Pressure` data. You can use this field function to export the mapped data.

Mapping Boundary Data to Imported Beams

It is possible to map pressure and shear stress data, that are defined on the faces of the wetted surface in STAR-CCM+, onto the nodal forces and moments for each node on the beam.

To map boundary data to the nodes of an imported beam, right-click the imported CAE model node under the Imported Models node and select Map data > Beam Data... This menu option is also used to map imported CAE data to the STAR-CCM+ model.



This will activate the *Map Imported Model Data* dialog. Use the dialog as follows:

- Ensure the *Map Direction* is set to STAR-CCM+ to Imported Model.
- In the *Mapping Inputs* box, select the STAR-CCM+ boundaries you wish to map from, and the imported beam you wish to map to.
- The *Data to map* menu allows you to select which field data will be mapped. This menu displays field functions that are available for mapping in the simulation. When you select the field data, the *Imported Data Type* and *Boundary Data Type* will be set automatically.
- Activate the Include Wall Shear Stress option if you wish to map wall shear stress data to the imported beam.

In the following image, the settings specify that fluid traction data will be mapped from the Fluid-Wall boundary of the STAR-CCM+ model to the imported beam element, PART-1-1.



When the settings in the dialog are complete click Apply, then Close.

The time taken to carry out the mapping operation will be displayed in the *Output* window, followed by `Done`. A field function will be exposed corresponding to the mapped data, for example the `Mapped Traction` field function will be exposed if you map `Traction` data. You can use this field function to export the mapped data.

Mapping Region Data to Imported Volumes

To map region cell data to the volume vertices of an imported model, right-click the imported CAE model node under the Imported Models node and select Map data > Volume Data...



This will activate the *Map Imported Model Data* dialog. Use the dialog as follows:

- Ensure the *Map Direction* is set to STAR-CCM+ to Imported Model.
- In the *Mapping Inputs* box, select the STAR-CCM+ regions you wish to map from, and the imported model volumes you wish to map to.
- The *Data to map* menu allows you to select which field data will be mapped. Currently, it is only possible to map temperature data from a region to an imported volume. When you select the field data, the *Imported Data Type* and *Boundary Data Type* will be set automatically.
- Use the additional constraints in the Nearest Neighbor Search Constraints box at the bottom of the dialog to ensure the success of the mapping operation.

- The settings in the following screenshot indicate that temperature field data will be mapped from the Fluid region to the Manifold Volume.



When the settings in the dialog are complete click Apply, then Close.

The time taken to carry out the mapping operation will be displayed in the *Output* window, followed by `Done`. A field function will be exposed corresponding to the mapped data, called `Mapped Vertex Temperature`. You can use this field function to export the mapped data.

Nearest Neighbor Search Constraints

In some cases you will need to use additional checks to ensure that the matching entities of the source and target objects are correctly identified. These checks, or constraints, are provided in the *Nearest Neighbor Search*

Constraints group box of the *Map Imported Model Data* dialog. Using these constraints also allows you to use the mapped flag field function to assess the result of the mapping.



- Check Source/Target Proximity -- The source face from which data is being mapped should lie within a certain distance of the target face. This ensures that data is not mapped from surfaces which are too far away. The tolerance is set by specifying a fraction of the maximum of the mesh sizes on the source/target surfaces. The default value is 1.0.
- Check Source/Target Face Normals -- This check enforces the rule that the source face normal and target face normal be facing each other. The Minimum angle option specifies the minimum angle permissible between opposing face normals.

  This constraint is essential when a single surface wraps around edges and corners of the geometry. In the example below, the constraint ensures that data on each side of the source box is mapped to the correct side of the target box.



Without face normals check



With face normals check

This check should also be activated when you are mapping to shell elements--particularly when the imported shell does not have a consistent distribution of positive and negative faces. If the geometries of the source and target surfaces match closely, you can choose to map to both sides of the shell element and rely on this constraint to choose the correct side (positive or negative) on which data is mapped.

If some faces are not intended to be mapped, and will fail either of the criteria above, you can choose to assign reasonable alternate values for the unmapped faces. This can be done by specifying a value in the alternate value text boxes. The default alternate values are 0.0.

Using the Mapped Flag Field Function

The faces that satisfy the specified nearest neighbor search constraints are successfully mapped using an appropriate interpolation technique. The faces that fail these tests are not mapped at all. To allow you to troubleshoot and choose appropriate values for the tolerances, a field function named `Mapped Flag` is created upon each mapping operation. This field has the following values:

- 0 -- mapping was successful
- 1 -- proximity check failed
- 2 -- normalcy check failed

You can visualize the unmapped faces by plotting contours of `Mapped Flag`.

Selecting a Reference Temperature Evaluation Option

When you map heat transfer coefficients from STAR-CCM+ to an imported CAE model, you can select the method by which STAR-CCM+ evaluates the heat transfer coefficient and reference temperatures. Further discussion is provided in the section, Exchanging Heat Transfer Coefficients.

When *Data to map* is set to Local Heat Transfer Coefficient, the *Reference Temperature Evaluation Option* box appears within the *Map Imported Model Data* dialog. You can choose the method by which STAR-CCM+ evaluates the reference temperature by selecting one of the four options available.



The particular form of heat transfer coefficient that is mapped will depend on the reference temperature chosen. For the first three options, the mapped heat transfer data will be stored in the field function `Mapped Local Heat Transfer Coefficient`.

- **Local heat transfer reference temperature**

This option allows you to map the local heat transfer reference temperature, which results in the creation of the `Mapped Local Heat Transfer Reference Temperature` field function. With this option, STAR-CCM+ will base its mapping on the Local Heat Transfer Coefficient.

- **Specified bulk temperature**

This option allows you to specify a single temperature to use as the reference temperature. Selecting this option results in the creation of the `Mapped Bulk Temperature` field function. STAR-CCM+ will base its mapping on the Heat Transfer Coefficient.

- **Near wall cell temperature**

This option allows you to map the cell-based reference temperature values. It results in the creation of the `Mapped Cell Based Temperature` field function. STAR-CCM+ will base its mapping on the Local Heat Transfer Coefficient.

- **Specified y+ temperature**

This option is for turbulent flows and allows you to map reference temperatures at a specified y+ value. This option is preferred over the local heat transfer reference temperature. Selecting this option results in the creation of the `Mapped Heat Transfer Reference Temperature User Y+`. Note that the specified y+ heat transfer coefficient and reference temperature do not take account of boiling heat transfer and radiation. For

cases involving boiling and radiation, you should map and export the boundary heat flux, or alternatively the local heat transfer coefficient and reference temperature.

With this option, STAR-CCM+ will base its mapping on the Specified y+ Heat Transfer Coefficient, and will store the mapped heat transfer data in the field function `Mapped Specified y+ Heat Transfer Coefficient`.

Exporting Mapped Data

STAR-CCM+ solution data that has been mapped onto an imported CAE model can be exported from STAR-CCM+ so that it can be used in an external CAE package. Data is exported from STAR-CCM+ in text files that can be read by the respective CAE package.

The following table shows the export options supported by STAR-CCM+:

	CAE Package			
Field Function	**NASTRAN**	**Abaqus**	**ANSYS**	**RadTherm**
Pressure	✓	✓	✓	-
Traction	✓	✓	✓	-
Heat Transfer Coefficient and Reference Temperature	✓	✓	✓	✓
Heat Flux	✓	✓	✓	✓
Temperature	✓	✓	✓	-

Before exporting data, ensure you have mapped the surface, beam or volume data from the STAR-CCM+ model to the imported model.

To export mapped data, right-click on the imported models node and select Export Mapped Data...



The *Export File* dialog will open. To use this dialog, do the following:

- Navigate to the location you want to save the file and enter a file name, or select an existing file. If you activate Append to existing file, the exported data will be added to the end of the selected file.
- Use the *Export Mode* menu at the bottom-right of the dialog to select either the Surface Data or Volume Data export option. This selection will affect the options shown in the dialog.

- In the Imported Surfaces or Imported Volumes tab (depending on the selected export mode), select the surfaces or volumes for which you wish to export data. In the following screenshot, the settings specify that surface data will be exported for the imported surface INNER.



- Select the FieldFunctions tab.
- In the *Export data as* menu, select the field data you wish to export.
- From the *Data to export* menu, select the corresponding field function.

  Note that to export heat transfer coefficient data, select the mapped heat transfer coefficient field function in the *Data to export* menu and the mapped reference temperature field function in the *Tref to export* menu. To export traction data, select the mapped wall shear stress field function in the *Data to export* menu and the mapped pressure field function in the *Pressure to export* menu.

  If you are exporting Temperature data for surfaces, you should select `Temperature BC`. If you are exporting Temperature data for shell surfaces, you can choose either temperature option.
- Ensure the correct units are selected in the *Export units* menu.

- A summary of the selected surfaces/volumes and field functions will be displayed in the box underneath the tabbed panel. In the screenshot below, the settings specify that the `Mapped Local Heat Transfer Coefficient` and `Mapped Local Heat Transfer Reference Temperature` field functions will be exported as heat transfer coefficient data.
- To export the data, click Save.



If no mapped field functions are available for the selected surface or volume, a message will be displayed in the summary box in the *Export File* dialog. Further information on exporting data to NASTRAN, Abaqus, ANSYS, and RadTherm are provided in a separate section.

Exporting Temperature Data on Shells

When exporting temperature data on shell surfaces you can choose to export the data as a temperature field, or as a boundary condition. When the data is exported as a field, it can be used to specify a constant temperature across the thickness of the shell in the external CAE package. This is used for very thin shells, where any temperature gradient across the shell thickness would be negligible.

Temperature data exported as a boundary condition specifies the temperature on one side of the shell. The data can then be used in an external CAE package to compute the temperature gradients across the shell thickness.



To do this, map the temperature data to the appropriate shell surface (for field data export, you can export the data to either shell surface). Select the export option for the imported CAE model, and select one of the following options in the *Export data as* menu of the FieldFunctions tab:

- Temperature Field -- to export temperature data as a field
- Temperature BC -- to export temperature data as a boundary condition

In the *Data to export* menu select the Mapped Vertex Temperature field function.

Mapping and Exporting Transient Surface Data

STAR-CCM+ allows you to map transient surface data onto the surfaces of an imported transient model, and to export the mapped data to an external `.trn` or `.mat` (Matlab) file.

Map and export operations are carried out through the imported transient model node trn:fileName (trn:turbine in the image below) and its subnode, Export mapped transient data.



Imported transient model [trn:fileName] Properties

Transient Display Time	Specifies the current display time. The slider can vary from the maximum to the minimum time of the current signal.	
Transient End Time	Displays the maximum time in the file.	
Export Mapped Transient Data	Allows to activate the Export mapped transient data node, for map and export operations.	
	Activated	Enables map and export of transient surface data through the Export mapped transient data node.
	Deactivated	Map and export of transient surface data is not enabled.
Update Data Set Functions	Allows automatic update of the data set functions.	
	Activated	The data set functions are automatically updated after any change to the imported transient model node.
	Deactivated	The data set functions are not updated automatically.

Export Mapped Transient Data Properties

Export Directory	Specifies the directory to which the file containing the transient surface data is exported.
Mapped File Format	Specifies the format of the file containing the transient surface data. Two options are available, `.trn` and `.mat`.
Source Surfaces	Allows to select the transient model surfaces from which the transient surface data is mapped.
Target surfaces	Allows to select the transient model surfaces onto which the transient surface data is mapped.

Scalar Functions	Allows to select the scalar fields to be mapped from the source surfaces to the target surfaces.
Vector Functions	Allows to select the vector fields to be mapped from the source surfaces to the target surfaces.
Interpolation Method	Specifies the interpolation technique that is used for mapping.
Export Data Location	Specifies whether the transient surface data is mapped to the target mesh vertices or surfaces.
Export Frequency	Specifies the number of time-steps between exports.
Start Export At Time Step	Specifies at what time-step the transient surface data export starts.
End Export At Time Step	Specifies at what time-step the transient surface data export ends.

To map and export transient surface data to target surfaces of an imported transient model:

1. Import the transient model into STAR-CCM+.
2. Expand the Imported Models node.
3. Select the relevant imported transient model node.
4. Activate Export Mapped Transient Data.

   A new node, Export mapped transient data, is added within the imported transient model node.
5. Select the Export mapped transient data node and set its properties to specify the map and export settings.

6. Specify mapping constraints by setting the Proximity Constraint node properties and/or the Normal Direction Constraint node properties.



7. Right-click the Export mapped transient data node and select Map & Export Data.



An alternative map and export procedure that uses data set functions is described under Using Data Set Functions.

Importing CAE Solution Data

In two-way coupling, a solution is obtained by sequentially running STAR-CCM+ and a third-party code to obtain a complete solution. This may involve using STAR-CCM+ to compute the traction load on a flexible structure caused by fluid flow, and a third-party code to compute the deflection of the structure and associated displacements and stresses. In order to carry out a two-way coupled simulation, not only do you need to map and export solution data from STAR-CCM+, but you will also need to

import solution data from the third-party software, and map it onto the STAR-CCM+ model. Further information on file-based coupling is available in the Understanding File-Based Coupling section.

Importing CAE solution data from es-ice or pro-STAR is different to other CAE products in terms of limitations and available data. This information is covered in another section of the User Guide.

The vertex data associated with an imported CAE model can be imported into STAR-CCM+ via the imported CAE model node. The following table shows the import options currently supported by STAR-CCM+, not including es-ice or pro-STAR:

	CAE Package			
Field Function	**NASTRAN**	**Abaqus**	**ANSYS**	**RadTherm**
Temperature	✓	✓	✓	✓
Displacement	✓	✓	✓	-
Eigenmodes	✓	✓	✓	-
Qvol	✓	-	-	-
Temperature (volume import)	✓	✓	✓	-

For Abaqus data import, in the *Open* dialog file selection pathway, it is also necessary to enter the command used to launch Abaqus, and the Load Step Name for the data. Additional settings are available when you are importing eigenmodes. When you import Nastran data using the Qvol card, in the *Import Data Options* dialog, set the Data Type to Cell as the values are present at the cell centers.

When vertex data is imported into the simulation it is stored in the simulation ready to be mapped onto a boundary. The data cannot be used in the simulation until it has been mapped onto the STAR-CCM+ model, but it can be visualized as described below. If displacement data has been imported into the simulation you can deform the CAE model based on the displacements.

To import CAE solution data, right-click on the imported CAE model node under the Imported Models manager node and select Import CAE Solution Data...



In the *Open* dialog select the output file containing the solution data. In the *Files of Type* drop-down list you can select the appropriate format for the CAE model.



This will activate the *Import Data Options* dialog, which is used to specify the properties of the imported data. Use the *Data Type* menu to select the type of data being imported, for example Vertex. In the *Field Name* menu, select the field data, then select the corresponding units from the *Units* menu.



After you click OK, the import process will result in the data being represented as one or more field functions in STAR-CCM+.



This field function can be displayed just like any other scalar value.



Importing Eigenmode Data

When importing eigenmode data for Nastran and Abaqus models, you can select the number of modes to import. To do this, open the *Import Data Options* dialog and set the *Field Name* to EigenModes. This will activate the Num modes text box. You can enter the following options to specify the number of modes:

- `ALLMODES` -- imports data for all modes in the data file.

- `X` -- to import the first *X* modes, for example `3, 5, 10`.



Deforming the CAE Model Using Imported Displacements

The CAE model can be deformed based on imported displacement data. This process deforms the CAE model mesh by applying the displacement data at each surface vertex. This allows you to view the deformed shape of the CAE model within STAR-CCM+. It is also possible to display scalar quantities on the deformed model within a scalar scene.

To deform a CAE model, you must first import displacement data for the model. This reveals the deform option in the CAE model pop-up menu. Right-click on the CAE model node and select Deform Using Imported Surface Displacements.

The following example shows an imported CAE model of a plate fixed at the base. In the external CAE package, a traction load was placed on the front surface of the plate.



In the following screenshot the CAE model has been deformed based on the imported surface displacements. Note that it is often necessary to change the parts displayed in the scene in order to view the deformed mesh. This is because the displacement data is typically associated with a particular

surface of the model, such as an imported surface, or volume surface. Therefore, you may need to add one of the other CAE model surfaces to the scene.



The deformed model can also be used to display scalar data. In the example below, the `Imported Displacement` field function is displayed on the surface. A geometry displayer has also been used to show the original position of the plate.



Mapping CAE Solution Data to a STAR-CCM+ Model

In two-way coupling, data from the third-party CAE solution is used to specify boundary conditions in STAR-CCM+. To do this, the data must first be imported into STAR-CCM+, then it can be mapped onto the corresponding boundary. For example, nodal displacements computed in an external FEA package may be mapped onto an FSI boundary in STAR-CCM+, and used to define mesh motion via the morpher. The mapping operation is described below. Currently, it is possible to map surface data and beam data from a CAE model to a STAR-CCM+ model.

Mapping Imported Surface Data to STAR-CCM+ Boundaries

To map surface data from the CAE model to STAR-CCM+, right-click on the imported CAE model node and select Map Data > Surface Data...



This will activate the *Map Imported Model Data* dialog. Use the dialog as follows:

- Set the *Map Direction* to Imported Model to STAR-CCM+.
- In the *Mapping Inputs* box, select the imported model surfaces you wish to map from, and the STAR-CCM+ boundaries you wish to map to.
- The *Data to map* menu allows you to select which field data will be mapped. This menu displays field functions that are available for mapping in the simulation. When you select the field data, the *Imported Data Type* and *Boundary Data Type* will be set automatically. If you are mapping local heat transfer coefficients, you will also be able to select a reference temperature evaluation option.
- In some cases the mapping algorithm may fail to find appropriate neighbors for each face on the source boundary. If this happens, use the additional constraints provided in the Nearest Neighbor Search Constraints box at the bottom of the dialog. Turning on these constraints also enables you to check the success of the mapping operation using the mapped flag field function.

In the following screenshot, the settings specify that imported temperature data will be mapped from the Manifold INNER surface of the imported Abaqus model to the Fluid Wall boundary of the STAR-CCM+ model.



When the settings in the dialog are complete click Apply, then Close.

The time taken to map the data will be shown in the *Output* window, and the message `Done.` will be displayed when the mapping operation is finished.

To display the mapped data results, create a scalar scene. Select the scalar function and the mapped surface part to display. The scalar values will appear.



When the mapping operation is carried out for the first time, a new field function will be created, according to the type of data that has been mapped from the CAE model. For example, if temperature data was imported in the previous step, an `Imported Temperature` field function will be created automatically by STAR-CCM+. When this data is mapped onto a boundary, the field function `Mapped Imported Temperature` is created. This field function can then be used when specifying a boundary condition, for example by setting the Thermal Specification type to Field Function, then selecting `Mapped Imported Temperature` from the list.

Mapping Imported Beam Data to STAR-CCM+ Boundaries

It is possible to map nodal displacement and rotation data, defined for the imported beam, onto the wetted surface in STAR-CCM+. Imported displacement and rotation data can be used as input for the morpher to deform the model.

To map beam data from the CAE model to STAR-CCM+, right-click on the imported CAE model node and select Map Data > Beam Data...



This will activate the *Map Imported Model Data* dialog. Use the dialog as follows:

- Set the *Map Direction* to Imported Model to STAR-CCM+.
- In the *Mapping Inputs* box, select the imported model beams you wish to map from, and the STAR-CCM+ boundaries you wish to map to.
- The *Data to map* menu allows you to select which field data will be mapped. This menu displays field functions that are available for mapping in the simulation. When you select the field data, the *Imported Data Type* and *Boundary Data Type* will be set automatically. Beams only support Vertex data.

In the following image, the settings specify that imported displacement data will be mapped from the imported Abaqus model to multiple boundaries of the fluid STAR-CCM+ model.



When the settings in the dialog are complete click Apply, then Close.

Using Data Mappers

In STAR-CCM+ data mapping is the transfer of solver data from one CAE mesh location to another. The data mapper feature allows you to create mappers, choose parts and fields for mapping, and specify various properties of the mapper. It is located under Tools > Data Mappers. This node comes into existence upon importing a CAE model.



This section describes data mappers and examples of how to use them. Specifically, instructions are provided for the following:

- Creating a data mapper
- Setting data mapper properties
- Setting CAE-specific properties
- Mapping data
- Data mapper formulation

What Are Data Mappers?

Data Mappers refer to STAR-CCM+ tools that are used to map field data between different mesh representations of a part. The source parts can be tables, surfaces, or volumes, and the discrete representations can be:

- finite-volume meshes representing STAR-CCM+ boundaries and regions, or
- finite-element meshes imported from third-party CAE packages.

STAR-CCM+ offers the following types of mappers:

- **Surface Data Mapper:** a mapper that interpolates data from a set of source surfaces to locations on a set of target surfaces, or volumes.
- **Volume Data Mapper:** a mapper that interpolates data from a set of source volumes to locations on a set of target surfaces, or volumes.
- **Beam Data Mapper:** a mapper that interpolates data from/to a single beam or line to/from a set of surfaces. Multiple beams require the creation of multiple mappers. Data can be mapped from/to beam vertices only.
- **Tabular Data Mapper:** a mapper that interpolates data from an XYZ table you import to locations on a set of target surfaces, or volumes.

For all types of mapper, the target volumes can be either STAR-CCM+ regions with a finite-volume mesh, imported CAE volumes with a finite-element mesh or a combination of both. For Surface, Beam and Volume Data Mappers, the sources can be either STAR-CCM+ regions with a finite-volume mesh, imported CAE volumes with a finite-element mesh or a combination of both. For a Tabular Data Mapper, the source is a table containing X, Y, and Z co-ordinates and source data.



There are two ways to create a data mapper in an existing simulation:

- Invoking mapping actions on imported models -- A data mapper is automatically created and listed under Tools > Data Mappers whenever a map action is invoked on the imported CAE model node. The properties of the mapper are automatically filled based on the selections in the Map dialog. The created mapper can be modified or reused as desired.
- Manual setup -- A surface/volume data mapper can be manually set up by following these steps:
  - Creating a data mapper
  - Setting data mapper properties
  - Specifying CAE-specific properties
  - Mapping the data

Data mappers can be deleted from an existing simulation if they are not required.

Caution: Deleting a data mapper does not delete the mapped data. To delete mapped data, first use the Clear Mapped Fields action.

Data Mappers Pop-Up Menu

These options become available, along with common pop-up menu items, when you right-click the Data Mappers node.

New Data Mapper	Creates a data mapper of the type that is selected in the submenu.	
	Surface Data Mapper	Executes the create data mapper action to create a surface data mapper.
	Volume Data Mapper	Executes the create data mapper action to create a volume data mapper.
	Tabular Data Mapper	Executes the create data mapper action to create a tabular data mapper.

	Beam Data Mapper	Executes the create data mapper action to create a beam data mapper.
Clear Mapped Fields	This action deletes any fields that were populated using data mappers and also removes any mapped data field function. If a mapped data field function is used by other objects in the simulation, such as scenes or reports, make sure that you remove any dependencies before using the Clear Mapped Fields option. If you do not remove the dependencies, any data associated with the mapped field function will be deleted but the mapped field function will not be removed from the simulation tree.	

Creating a Data Mapper

To create a data mapper:

- From the *File* menu, select File > Import > Import CAE Model and select a model for import.

  After the model is imported, the Data Mappers node appears under Tools in the object tree.

- Right-click the Data Mappers node, select New Data Mapper and then one of the data mapper options. See Data Mappers Pop-Up Menu.



Setting Data Mapper Properties

The data mappers are versatile tools that let you map any available data between any number of like parts irrespective of their discrete representation. Set data mappers properties at the source and targets, remembering that you can have several targets for each source. The types of source properties are:

- Surface source properties
- Volume source properties
- Beam source properties
- Tabular source properties

The target properties are:

- Surface target properties
- Volume target properties
- Beam target properties

Guidelines for setting up source properties:

- Using surface and volume sources
- Using field functions for mapping from surfaces or volumes
- Using a representation for mapping from surfaces or volumes
- Using tabular sources

Guidelines for setting up target properties:

- Using targets
- Using interpolation method
- Projecting targets on sources
- Using the proximity constraint check
- Using the normal direction constraint check
- Specifying alternate values

Surface Data Mapper Properties

Source Surfaces
A list of source parts that are used for this data mapper.

Source Stencil
Selects the type of source for the data, depending on where the source data is available.

- **Face** Selects faces as the source.
- **Vertex** Selects vertices as the source.

Scalar Field Functions
A list of scalar field functions that are used for this data mapper.

Vector Field Functions
A list of vector field functions that are used for this data mapper.

Representation
Sets the finite volume representation to get source data from.

Use original mesh
Controls whether to use the original mesh or the current mesh as the reference mesh.

- **Activated** The mapper uses the original mesh for reference mesh.
- **Deactivated** The mapper uses the current mesh for reference mesh.

Verbosity
Controls the level of verbosity of the selected data mapper.

- **Activated** Additional details are printed to the output window, such as showing the progress of various steps in the mapping process and the percentage of cells that is mapped successfully.
- **Deactivated** Only the time that is taken to map the data and a message to indicate that it has completed are printed.

Volume Data Mapper Properties

Source Volumes
A list of source parts that are used for this data mapper.

Source Stencil
Selects the type of source for the data, depending on where the source data is available.

- **Cell** Selects cells as the source.

Scalar Field Functions
A list of scalar field functions that are used for this data mapper.

Vector Field Functions
A list of vector field functions that are used for this data mapper.

Representation
Sets the finite volume representation to get source data from.

Use original mesh
Controls whether to use the original mesh or the current mesh as the reference mesh.

- **Activated** The mapper uses the original mesh for reference mesh.
- **Deactivated** The mapper uses the current mesh for reference mesh.

Verbosity
Controls the level of verbosity of the selected data mapper.

- **Activated** Additional details are printed to the output window, such as showing the progress of various steps in the mapping process and the percentage of cells that is mapped successfully.
- **Deactivated** Only the time that is taken to map the data and a message to indicate that it has completed are printed.

Beam Data Mapper Properties

Source Beams

A list of source parts that are used for this data mapper.

Source Stencil

Selects the type of source for the data, depending on where the source data is available. When mapping beam data, set Source Stencil to Vertex: Beams do not support face data; only vertex data are supported.

- **Face** Selects faces as the source.
- **Vertex** Selects vertices as the source.

Scalar Field Functions

A list of scalar field functions that are used for this data mapper.

Vector Field Functions

A list of vector field functions that are used for this data mapper.

Representation

Sets the finite volume representation to get source data from.

Use original mesh

Controls whether to use the original mesh or the current mesh as the reference mesh.

- **Activated** The mapper uses the original mesh for reference mesh.
- **Deactivated** The mapper uses the current mesh for reference mesh.

Verbosity

Controls the level of verbosity of the selected data mapper.

- **Activated** Additional details are printed to the output window, such as showing the progress of various steps in the mapping process and the percentage of cells that is mapped successfully.
- **Deactivated** Only the time that is taken to map the data and a message to indicate that it has completed are printed.

Tabular Data Mapper Properties

Source Table	Select the imported source table to use. The table must be imported into the Tools > Tables node before it can be imported here.
Source: X-Coordinate	Select the column in the source table for the X-coordinate of each node.
Source: Y-Coordinate	Select the column in the source table for the Y-coordinate of each node.
Source: Z-Coordinate	Select the column in the source table for the Z-coordinate of each node.
Source: Data	Select from a list of columns in the source table that are potential scalar field functions for this data mapper.
Verbosity	Controls the level of verbosity of the selected data mapper. **Activated**: Additional details are printed to the output window, such as showing the progress of various steps in the mapping process and the percentage of cells that is mapped successfully. **Deactivated**: Only the time that is taken to map the data and a message to indicate that it has completed are printed.

Surface Target Properties

Target Surfaces

A list of parts that are used as the target for this data mapper.

Target Stencil

Selects the type of destination for the data, depending on where the data for mapping.

- **Face** Selects faces as the target.
- **Vertex** Selects vertices as the target.

Use Original Mesh

Controls whether to use the original mesh or the deformed/morphed mesh for mapping.

- **Activated** The mapper uses the original mesh for mapping.
- **Deactivated** The mapper uses the deformed/morphed mesh for mapping.

Target Source Transform

Specifies the transform to project targets on the sources (surface-to-surface only).

- **Identity Transform** Uses original target points.
- **Simple Transform** Rotates, translates, or scales the target points.

Limiter

If this option is beyond the range of the contributing stencil values, it specifies how the mapped value is limited.

- **Min/Max of Contributing Stencil** Limits the mapped value based on the values of the points chosen for interpolation.
- **Min/Max of Contributing Sources** Limits the mapped values to the minimum/maximum of the source surface/volume. This option allows more extrapolation as compared to the first option.

No Limiting. Allow Extrapolation	Allows extrapolation and does not limit the mapped value.

Interpolation Method

Specifies the interpolation scheme for evaluation interpolation stencil and weights. Each method uses reciprocity between connected faces to ensure two-way flows of data and enhances stability and robustness.

Shape Function	Uses shape function interpolation (for Vertex data location only).
Least Squares	An averaging technique where the mapped data is weighted by the distance between source and multiple target data locations. If the source and target are coincident, the resultant weight at that location is one.
Nearest Neighbor	Maps the data to the nearest target location only.
Exact Imprint	Includes the intersecting areas between source and target faces in the interpolation stencil (for Face source and target data stencil only).
Approximate Imprint	Includes all source faces that lie within a maximum radius of a target face in the interpolation stencil (for Face source and target data stencils only).

Volume Target Properties

Target Volumes

A list of parts that are used as the target for this data mapper.

Target Stencil

Selects the type of destination for the data, depending on where the data for mapping.

- **Cell** — Selects cells as the target.
- **Vertex** — Selects vertices as the target.

Use Original Mesh

Controls whether to use the original mesh or the deformed/morphed mesh for mapping.

- **Activated** — The mapper uses the original mesh for mapping.
- **Deactivated** — The mapper uses the deformed/morphed mesh for mapping.

Boundary Map Option

Determines how boundary faces of the target volume are mapped from the source volumes.

- **Map from Boundaries of Source Volumes** — Use if the boundaries of the target and source volume coincide (even though the meshes are different).
- **Map from Interior of Source Volumes** — Use if the target volume is only a portion of a larger source volume.

Target Source Transform

Specifies the transform to project targets on the sources.

- **Identity Transform** — Uses original target points.
- **Simple Transform** — Rotates, translates, or scales the target points (for volume-to-volume mapping only).
- **Normal Projection Transform** — Projects the points in the normal direction onto the source surface (for surface-to-volume mapping only).

	Helical Projection Transform	Projects the points in a helical fashion onto the source surface (for surface-to-volume mapping only).
Limiter	If this option is beyond the range of the contributing stencil values, it specifies how the mapped value is limited.	
	Min/Max of Contributing Stencil	Limits the mapped value based on the values of the points chosen for interpolation.
	Min/Max of Contributing Sources	Limits the mapped values to the minimum/maximum of the source surface/volume. This option allows more extrapolation as compared to the first option.
	No Limiting. Allow extrapolation	Allows extrapolation and does not limit the mapped value.
Interpolation Method	Specifies the interpolation scheme for evaluation interpolation stencil and weights. Each method uses reciprocity between connected faces to ensure two-way flows of data and enhances stability and robustness.	
	Shape Function	Uses shape function interpolation (for Vertex data stencil only).
	Least Squares	An averaging technique where the mapped data is weighted by the distance between source and multiple target data locations. If the source and target are coincident, the resultant weight at that location is one.
	Nearest Neighbor	Maps the data to the nearest target location only.

Beam Target Properties

Target Beams	Specifies the beam element that is used as the target for this data mapper.

Using Surface and Volume Sources

Choose any number of parts in the source and target set. These parts can have different discrete representations or they can be grouped in various ways into regions and/or boundaries.

The source stencil is the location where the solution data is available on the source mesh. If the source parts are STAR-CCM+ regions and/or boundaries, the source stencil is most likely either a cell or a face. If the source parts are imported volumes and/or surfaces, the source stencil is typically a vertex. While you can choose any source stencil, choosing the wrong combination of stencils reduces the integrity of the mapping operation.

- Set the source surfaces and stencil in the properties of the Surface Data Mapper or Volume Data Mapper node.



Using Field Functions for Mapping from Surface or Volume Sources

Each data mapping requires you to select at least one scalar or vector field.

- Select the Scalar or Vector Field Functions expert property of the Surface Data Mapper or Volume Data Mapper node.

- Click the property customizer (...) button next to these properties to select from the list of available field functions for mapping.



Using Representation for Mapping from Surface or Volume Sources

In the Representation property, you can specify the mesh and solution from which the source data is picked. The source data can include any representation contained in a solution history file. This property is not relevant in cases where the data is from an imported model.



Using Tabular Sources

In STAR-CCM+, it is possible to apply an imported table as a boundary condition or as an initial condition. If the table data is much more coarse than the associated volume mesh, use a data mapper to apply an interpolation between the table and mesh.

The table can be an XYZ table created in STAR-CCM+ or a file table you created elsewhere. For a table to be useful, it must have at least four named columns. Three of the columns specify the x, y, and z location of a value and the fourth column specifies the scalar value to be mapped at that location. It is possible to have more than one variable of interest at each location (for example, pressure and temperature) in the table.

To use a tabular data mapper:

- Import the table into Tools > Tables node.
- Right click the Data Mappers node and select New Data Mapper > Tabular Data Mapper.

Note: The Data Mappers node is available after you import a CAE model

- In the source properties for the data mapper select the table, the scalar data, and the x, y, and z co-ordinates that define each data point.



- Specify the target, and interpolation type.
- Map the data.

  This creates a field function named Mapped <Field Function Name>.
- Use the Mapped <Field Function Name> on the target boundary or region settings.

Using Targets

- Set the target surfaces and stencil in the properties of the Target Specifications > Surface or Volume node.



Choose a volume target specification for the surface data mapper and a surface target specification for the volume data mapper.

- Right-click the Target Specifications node and select New Group > Surface or Volume.



Using Mapper Interpolation

The following interpolation methods are provided in STAR-CCM+:

- Nearest Neighbor -- the interpolation stencil includes only the nearest neighbor. The weight of the nearest neighbor in the stencil is 1, and is therefore the only value included in the interpolation. This method is available for any source and target surface and volume data locations.

- Least Squares -- the interpolation stencil includes the nearest neighbor and its immediate neighbors, weighted by the distance between the source and target data locations. This method is available for any source and target surface and volume data locations.
- Shape Function -- the interpolation stencil is determined using shape functions, which are special relations between vertices. This method is available for vertex source surface data location.
- Exact Imprinting -- uses a maximum radius of the target face to find the source faces to include in the interpolation stencil. The maximum radius is the distance from the target face centroid to one of its vertices. The target and source faces are intersected and the weights are the areas of the subfaces that match up between the target and sources. This method is available when mapping between surfaces. Since exact imprinting is computationally intensive to find exact intersections, this method relies on a proximity check to determine the stencil.
- Approximate Imprinting -- as with the exact imprint method except that the weights include all source faces that are found within the maximum radius, and do not depend on the intersecting process. This method is available when mapping between surfaces.

The Exact and Approximate imprinting schemes are globally conservative methods, but do not strictly satisfy local conservation. They result in more accurate solutions than the least squares methods when mapping from a fine to a coarser mesh. The least squares method is more accurate locally and is a better choice in cases where the target and source surfaces have comparable mesh densities. In terms of performance, the exact imprinting method is computationally more expensive than the least squares scheme, while the approximate method is of comparable efficiency to the least squares.

Set the interpolation method in the properties of the Target Specifications > Surface or Volume node.



Projecting Targets on Sources

The target source transform is set up to define how the target entity is transformed into the source entity. Four transforms are possible -- identity, simple, normal, and helical projection.

- The identity transform is available when mapping from any source surface or volume to any target surface or volume. It does not change the target points.
- Simple transform is available for surface-to-surface or volume-to-volume mapping. Define the transform in terms of rotation, scaling, or translation. A Simple Transform subnode is added.
- The Normal Projection Transform projects a cell centroid in the normal direction onto the source surface.
- For the Helical Projection Transform, the cell centroid is projected onto the surface using a helical path. The path is defined using the axis, origin, and pitch of the helix. A Helical Projection Transform subnode is added in which the transform properties are defined.

For each type of transform:

- Set the target source transform in the expert properties of the Target Specifications > Surface or Volume node.



Helical Projection Transform Properties

Pitch	Specifies the pitch of the helix, which is the distance along the axis of one complete turn of the helix.
Helical axis	Specifies the axis of rotation of the helix.
Helical Axis Origin	Specifies the origin of the helix.

Simple Transform Properties

Rotation Angle (deg)	Specifies the angle of rotation about the rotation axis in degrees.
Scale	Specifies the scaling vector.
Rotation Origin	Specifies the axis of rotation origin.
Rotation Axis	Specifies the axis of rotation.
Translation	Specifies the translation vector.

Using the Proximity Constraint Check

The proximity check is relevant if the surfaces are physically separated from each other or have some mismatch to avoid mapping data. Such a situation is possible if the source parts do not completely represent some of the target parts. This scenario is often encountered in solid/fluid coupling, especially with external packages. For co-located surfaces, a proximity check is not required.

When the proximity constraint is activated, the distance between parts is compared to a tolerance you define. If the distance is greater than this tolerance, an alternate value is used.

Note: Be aware of the general mesh sizes on parts while choosing this tolerance. If the mesh sizes vary extensively among the various selected parts, choose appropriate subsets of the parts for mapping.

Activate and set the proximity constraint in the properties of the Target Specifications > Surface or Volume > Proximity Constraint node.



Proximity Constraint Properties

Enable	If activated, the mapper checks if the target point and its nearest source point are within a certain distance.	
	Activated	The mapper checks if the target point and its nearest source point are within a certain distance. An Alternate Value Specification sibling node is added.
	Deactivated	The check is not performed. This is the default setting.
Use Absolute Tolerance	If activated, the mapper checks if the target point and its nearest source point are within the Absolute Proximity Tolerance.	

Activated	The mapper checks if the target point and its nearest source point are within the Absolute Proximity Tolerance. If this is deactivated, the mapper uses a value based on Relative Proximity Tolerance multiplied by the longest element edge-length of the target surface mesh.
Deactivated	The check is not performed. This is the default setting.

Relative Proximity Tolerance	Specifies a multiplication factor for checking proximity. The tolerance is calculated by multiplying this factor by the longest element edge-length of the target surface mesh. For example, if the maximum target edge length is 0.05 mm and the source and target meshes are 0.1 mm apart, then a factor of 2 is required to enable mapping. This value is used when the Enable checkbox is activated and the Use Absolute Tolerance checkbox is deactivated.
Absolute Proximity Tolerance	Specifies a value for the tolerance for checking proximity. This value is used when both the Enable and the Use Absolute Tolerance checkboxes are activated.

Using the Normal Direction Constraint Check

The normalcy check is relevant for surface-to-surface mapping. This check is intended to ensure that the surface normal at the target point faces the surface normal at the closest source point. Such a condition is true for solid-fluid interfaces, for instance. Also, for thin solids, this check avoids wrongly identifying the matching surface.

The normalcy tolerance is the lower limit of the angle between surface normals at the target point and the corresponding source point. While ideally this angle is 180 degrees, account for differences in surface meshes on either side and also for differences in the modeling itself. A tolerance of 90 degrees or greater means that the normals face each other.

Enable and set the proximity constraint in the properties of the Target Specifications > Surface > Normal Direction Constraint node.



Normal Direction Constraint Properties

Enable	Controls whether to check the angle between surface normals (Surface Data Mapper only).	
	Activated	The mapper checks if the angle between surface normals at the target point and the corresponding source point is greater than a certain value.
	Deactivated	Leaves this option deactivated.
Minimum angle (degrees)	Specifies the lower limit for the angle between source and target normals. This value is registered only if the Normals Check option is activated (Surface Data Mapper only).	

Specifying Alternate Values

Data is not mapped to target points that fail the proximity or normalcy checks. Choose to assign constant alternate values for each field to these points. Specify the alternate value in the properties of the Target Specifications > Surface or Volume > Alternate Value Specification node.



The field functions you selected for mapping populate the relevant property fields of the Alternate Values Specification node.

Note: Be careful while specifying alternate values. In STAR-CCM+, the order of the values is tied to the order in which the selected fields are listed. This rule applies to both scalar and vector alternate values.

Setting Alternate Values for a Selected Scalar Function

For scalar field functions, enter the single value for each function as follows:

- Click to the right of the Scalar Alternate Values.
- Enter the constant values for all the functions as a comma-separated list in the same order as the selected field functions. Place the entire list in square brackets.

For example, to set the alternate value of Local Heat Transfer Coefficient to 0 and the Local Heat Transfer Reference Temperature to 300:

- If the Selected Scalar Fields are listed as [Local Heat Transfer Coefficient, Local Heat Transfer Reference Temperature], then in the Scalar Alternate Values type `[0.0,300]`

Setting Alternate Values for a Selected Vector Function

For vector field functions, enter the three components for each function as follows:

- Click to the right of the Vector Alternate Values.
- Enter the three components of each function as a comma-separated list contained in round brackets. List all the functions as a comma-separated list in the same order as the selected field functions. Place the entire list in square brackets.

  For example, to set the alternate value of Velocity to (1,0,0) and the Wall Shear Stress to (0,0,0):

  - If the Selected Vector Fields are listed as [Velocity, Wall Shear Stress], then in the Scalar Alternate Values type `[(1,0,0),(0,0,0)]`

Alternate Value Specification Properties

Selected Scalar Fields	This read-only property lists the scalar fields for mapping.
Scalar Alternate Values	Enter an alternate value for each of the chosen scalar fields to be assigned to target points that are not mapped. This property is relevant only if the proximity check (or normalcy check for the surface mapper) option is activated. See setting a scalar alternate value for more explanation and an example.
Selected Vector Fields	This read-only property lists the scalar fields for mapping.
Vector Alternate Values	Enter an alternate value in vector notation for each of the chosen vector fields to be assigned to target points that are not mapped. This property is relevant only if the Proximity Check (or Normalcy Check for the surface mapper) option is activated. See setting a vector alternate value for more explanation and an example.
Estimate Alternate Values	Controls whether to estimate an alternate value for unmapped targets. These options only appear for target surfaces.
	Do Not Estimate Alternate Values: The specified alternate values are used.
	Estimate For All Unmapped Patches: An alternate value is estimated by averaging the surrounding mapped targets. If this fails, the specified alternate values are used.

Max. Patch Face Area	An alternate value is estimated by averaging the surrounding mapped targets, limited to a maximum patch area. A Maximum Unmapped Patch Area subnode is created. If this mapping fails, the specified alternate values are used.
Max. Patch Face Count	An alternate value is estimated by averaging the surrounding mapped targets, limited to a specified number of surrounding faces. A Maximum Unmapped Face Count subnode is created. If this fails, the specified alternate values are used.

Setting CAE-Specific Properties

As mentioned earlier, the data mappers are created manually or by invoking a map-action on surfaces or volumes of imported models. The map-action dialog has special properties depending on the type of surfaces and the data type chosen for mapping. CAE-specific properties exist for the surface data mappers in the Heat Transfer Export Options and Shell Side subnodes.

What Is the Heat Transfer Export Options Object?

In coupling with CAE packages for Conjugate Heat Transfer applications, the CAE model is imported into the STAR-CCM+ simulation. Typically, the Local Heat Transfer Coefficient (HTC) is mapped from the STAR-CCM+ fluid boundary to the corresponding solid surface from the imported CAE model. As a fluid temperature is required to calculate the heat flux, the Local Heat Transfer Reference Temperature is mapped and exported along with HTC as well. The map-action dialog also allows you to specify a constant bulk temperature as the reference temperature. The dialog also

allows you to choose the cell-center temperature in case radiation is being solved in both the solid and the fluid. The surface data mapper displays a Heat Transfer Export Options subnode when appropriate.



This node is displayed if you choose all of the following:

- STAR-CCM+ boundaries as a source surfaces
- Imported Surfaces as target surfaces
- The source and target stencils set at Face
- Local Heat Transfer Coefficient included among the Scalar field functions for mapping.

The HTC export options can be specified by selecting the node in the object tree and working with the properties.

Heat Transfer Export Options Properties

You cannot activate both the Use Near Wall Cell Temperature and Use Specified Bulk Temperature options at the same time. However, if neither one is activated, the field function for Local Heat Transfer Reference Temperature is automatically appended to the list of scalar field functions for mapping.

Use Near Wall Cell Temperature	Accounts for the inclusion of radiation in the simulation.	
	Activated	Exports cell center temperature as the reference temperature. This is applicable in case radiation is being solved in both the STAR-CCM+ and CAE packages.

	Deactivated	Leaves this option deactivated.
Use Specified Bulk Temperature	Controls the export of the bulk temperature.	
	Activated	Exports a constant bulk temperature as the reference temperature.
	Deactivated	Leaves this option deactivated.
Bulk Temperature	Specifies a value for bulk temperature.	

What Is the Shell Side?

All the imported RadTherm models consist of shell surfaces. Because the shell can potentially be exposed to different fluids on either side, choose the side of the shell for mapping the data. In a mapping dialog for the imported RadTherm model, radio buttons let you choose the shell side. For shell targets, select the mapping target-side from the properties of the Shell Side subnode that appears under the surface data mapper. This node is also subject to common menu items.



Shell Side Properties

Shell Side	Specifies mapping to a particular side of a shell.	
	Front	Selects the front face of the shell.
	Back	Selects the back face of the shell.

Mapping the Data

After you have selected the source and target parts, data fields for mapping and other data mapper properties, right-click the data mapper node and select Map Data from the pop-up menu.



If the Verbosity expert property is enabled, statistics of the neighborhood search and mapping are presented in the *Output Window* at the end of the mapping process.

If the solution on the source parts has been updated, you can also use this feature to update the mapped fields.

Using the Mapped Flag Field Function

The faces that satisfy the specified mapping search constraints are successfully mapped using an appropriate interpolation technique. The faces that fail these tests are not mapped at all. To allow you to troubleshoot and choose appropriate values for the tolerances, a field function named `Mapped Flag` is created upon each mapping operation. This field has the following values:

- 0 -- mapping was successful
- 1 -- proximity check failed
- 2 -- normalcy check failed

Visualize the unmapped faces by plotting contours of `Mapped Flag`.

Examples of Using Data Mappers

Data mapping is a general concept which entails interpolating field data from a set of source points to a set of target points. STAR-CCM+ encompasses this functionality into data mappers and employs them for explicit coupling in multiphysics applications. Currently STAR-CCM+ supports coupling with Abaqus, RadTherm, and Nastran. This section provides two examples of applications where explicit coupling can be used for analysis:

- Explicit coupling within STAR-CCM+ (Conjugate Heat Transfer) -- Thermal analysis of a solid with flow-thermal analysis of fluid, for instance in under-hood analysis. The CFD in the fluid region can be solved in STAR-CCM+ while the solid region is analyzed either in STAR-CCM+ or in an external CAE package.
- Explicit coupling between STAR-CCM+ and Abaqus (Fluid Structure Interaction) -- Stress analysis of a solid coupled with CFD in fluid, for instance analyzing the flow around an F1 wing. In this case, STAR-CCM+ solves the CFD and the structural analysis is either done using STAR-CCM+'s finite volume stress solver or using an external CAE package.

The coupling between the different physics can take place either at a surface level or at a volume level. An example of surface-to-surface coupling is a fluid structure interaction problem, where the solid and fluid models share an interface on which the fluid forces and solid displacements are exchanged. An example of a volume-to-volume coupling is a thermo-structural analysis. The temperature field from a conjugate heat transfer analysis is provided as an initial condition for performing thermo-structural analysis in the solid region. In both cases, the data is mapped between different discrete representations of the same physical space (either surface or volume). The rest of the section describes two coupling scenarios to highlight the utility of data mappers for explicit coupling.

Explicit Coupling Within STAR-CCM+

Consider a conjugate heat transfer (CHT) application in which both the solid and fluid regions are created and meshed in STAR-CCM+. Instead of creating an interface as is traditionally done for implicit coupling, both the regions are solved separately as follows:

1. The physics continua are set up by selecting all the appropriate models to solve for flow and energy in the fluid and heat transfer in the solid.

The idea is to solve solid and fluid separately, but couple them by exchanging boundary data at the common boundary.

2. Now two surface data mappers are created from the Tools > Data Mappers manager node, one for mapping from solid to fluid (S2F) and the other for mapping from fluid to solid (F2S).
3. In CHT applications, the data for exchange includes:
   - Local Heat Transfer Coefficient, and Local Heat Transfer Reference Temperature from fluid to solid
   - Surface Temperature from solid to fluid
4. Now in the S2F mapper choose the solid boundary as the source surface and the fluid boundary as the target surface, and select face stencils for both source and targets. In addition, choose Temperature as a scalar field for mapping.
5. In the F2S mapper, choose the fluid boundary to be the source and the solid boundary to be the target, and again do face-to-face mapping. The scalar fields in this case are htc and Tref.
6. Now the mappers are set up to map data. Map once to create Mapped field functions.
7. The next step is to apply these Mapped field functions as boundary conditions:
   - For the solid boundary, choose a Convection as the Thermal Specification and choose Field Function as the method to apply heat transfer coefficient and ambient temperature. Then select the `MappedLocalHeatTransferCoefficient` and `MappedLocalHeatTransferReferenceTemperature` as the respective field functions.
   - For the fluid boundary, choose to apply the Dirichlet boundary condition using the Field Function method. In this case choose `MappedTemperature` as the applied field function.
8. Now when the solver iterates, these fields are applied as boundary conditions.
9. A java macro can be created to take the appropriate number of steps before performing a map operation. **Perform a map operation to update the Mapped field functions.** The coupling frequency can be determined based on the nature of the problem.

Explicit Coupling Between STAR-CCM+ and Abaqus

Consider an FSI application in which the solid region is modeled in Abaqus and the fluid region is modeled in STAR-CCM+. The fluid model is created and meshed in STAR-CCM+ as usual. The solid model is created in Abaqus.

1. The fluid physics continua and solver properties are specified appropriately. An FSI application requires enabling the morpher model to account for solid displacement.
2. Once the Abaqus model is set up, create an `.inp` file as well as a database file (`.odb`) with initial nodal displacements.
3. The Abaqus `.inp` file is imported into the simulation. The interfacing surface pairs are identified.
4. Now two surface data mappers are created from the Tools > Data Mappers manager node, one for mapping from solid to fluid (S2F) and the other for mapping from fluid to solid (F2S).
5. In FSI applications, the data are:
   - Pressure and Wall Shear Stress from fluid to solid
   - Surface Displacement from solid to fluid
6. Now in the S2F mapper choose the Abaqus surface as the source surface and the fluid boundary as the target surface, and choose the vertex source stencil and the face target stencil. In addition, choose Incremental Surface Displacement as a mapped vector field (this field is created when nodal displacements are imported from Abaqus).
7. In the F2S mapper, choose the fluid boundary to be the source and the solid boundary to be the target, and choose face stencils for both source and targets. Choose Pressure as a scalar field and Wall Shear Stress as a vector field for mapping.
8. Now the mappers are set up to map data. Map using the S2F mapper to create Mapped field functions.
9. The next step is to apply these Mapped field functions as boundary conditions.
   - For the fluid boundary, choose Displacement as the Morpher specification. The displacement is applied using the mapped field function `MappedIncrementalSurfaceDisplacement`.
   - For the Abaqus imported surface, `MappedPressure` and `MappedWallShearStress` are exported to a separate `.inp` file and applied as loads for the Abaqus model. The details of using Abaqus are beyond the scope of this section.
10. Now you can create a java macro to control the frequency of coupling depending on the nature of the problem. However, each time the coupling happens, perform the following steps:
    - Import new displacement data from Abaqus.
    - Map data to the STAR-CCM+ boundary.
    - Map data from the fluid boundary to the imported surface.

- Export the mapped data to the `.inp` file.
- Apply the load on the Abaqus model.

Data Mapper Formulation

Data mappers use interpolation schemes to transfer data from source to target meshes. It is possible to map between surfaces and between volumes. STAR-CCM+, like other finite volume codes, stores the solution at each cell or face centroid of the mesh. Finite element codes store the solution at vertices of the mesh.

STAR-CCM+ uses various interpolation schemes for data mapping. The schemes vary depending on the choice of source stencil. For surface data mappers, the available methods are:

- Nearest Neighbor
- Least Squares
- Shape Function
- Exact Imprinting
- Approximate Imprinting

For volume data mappers, the available methods are:

- Nearest Neighbor
- Least Squares
- Shape Function

The following sections describe the implementation of the various interpolation schemes. The shape function scheme is illustrated in the context of vertex to face mapping and all the other schemes are illustrated in the context of face to face mapping.

Nearest Neighbor

In the figure below, consider mapping data from faces of the red background mesh to faces of the blue mesh using a nearest neighbor scheme. **Face k** is the closest point to the target **Face n**. The data from **Face k** is assigned to **Face n** as the mapped value.



Least Squares Interpolation

In the figure below, consider mapping data from faces of the blue background mesh to faces of the red mesh using a least squares scheme. **Face 0** is the closest source to the target **Face a**. The neighbors of **Face 0** are

defined as any face that shares at least one vertex with **Face 0**. These neighbors are included as part of the interpolation stencil. In the example below the neighbors include faces **0, 2, 3, 4, 5, 6**, and **8**.



STAR-CCM+ uses the second-order terms of a Taylor series expansion as a "cost function" to approximate the error of the temperature distribution at point **a**. It then minimizes the cost function in Eqn. (2657) to solve for the temperature at the target point **a**.

$$J = \sum_{n \in F(c)} \left(T + (\boldsymbol{x}_n - \boldsymbol{x}) \bullet \nabla T - T_n\right)^2 \tag{2657}$$

$$\nabla T = \begin{cases} \dfrac{\partial T}{\partial t_1}\boldsymbol{t}_1 + \dfrac{\partial T}{\partial t_2}\boldsymbol{t}_2 & \text{for surfaces} \\ \dfrac{\partial T}{\partial x}t_x + \dfrac{\partial T}{\partial y}t_y + \dfrac{\partial T}{\partial z}t_z & \text{for volumes} \end{cases} \tag{2658}$$

where:

$\boldsymbol{t}_1$ and $\boldsymbol{t}_2$ are surface tangents.

The drawback of a neighbor-based least squares scheme is that some of the blue faces (for example **1, 10, 11, 7**) that the target face imprints upon are not included in the stencil.

Note: The least squares method is applicable for all element types and for source and vertex source stencils.

Exact Imprinting Interpolation

In the figure below, consider mapping data from faces of the blue background mesh to faces of the red mesh using the exact imprinting scheme. In this method, the target red face is imprinted on the background blue mesh. This results in several facelets that are the intersection of the background mesh (faces **0, 1, 2, 3, 4, 5, 6, 7, 8, 10, 11**) with the target face.



The normalized areas of these facelets are used as interpolation weights in evaluating the mapped value at face-center **a**. This method is preferable when:

- a flux-preserving mapping is desired
- the source mesh is finer than the target mesh and the exact imprint does not fail to map.

Unlike the least squares scheme, all the source faces that the target face imprints upon are included in the stencil.

Mapping with the exact imprint method can fail when surface curvature is inadequately resolved in either the source or target mesh.

Approximate Imprinting Method

In the figure below, consider mapping data from faces of the blue background mesh to the red face using an approximate imprinting method. This method is a variation of exact imprinting. Since exact imprinting is computationally intensive to find exact intersections, this method relies on a proximity check to determine the stencil. The stencil includes all source faces that lie within a circle that is centered at the target face-centroid and

passes through the farthest vertex of the target face. In the figure, the circle includes faces **0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13** and **14** as part of the stencil.



Exact imprinting would have eliminated faces **9, 13** and **14** from the stencil because they do not intersect the target face. As for exact imprinting, the normalized face areas are used as interpolation weights in this case.

Although the approximate imprinting method is more robust than the exact imprinting method, it does not preserve flux quantities. The approximate method has comparable accuracy to the exact method when the source mesh is much finer than the target mesh.

Shape Function Interpolation

In the figure, consider mapping data from vertices of the blue source mesh to faces of the red target mesh using shape function interpolation. This scenario is typical when mapping data from a finite-element mesh to a finite-volume mesh. The face centroid **n** lies within the element **k**. The shape functions of element **k** are used to interpolate data to the point **n**.



The general form of shape function interpolant is:

$$T_n = \sum_{l \in N(k)} T_1 N_1(\xi_n, \eta_n, \chi_n) \tag{2659}$$

where:

N is the shape function for the specific element type

(ξ_n, η_n, χ_n) are the computational coordinates of the point **n** in the element **k**

For a four-node isoparametric quadrilateral in 2D, the shape function for each node is:

$$N_1 = \frac{(1-\xi)(1-\eta)}{4} \tag{2660}$$

$$N_2 = \frac{(1+\xi)(1-\eta)}{4} \tag{2661}$$

$$N_3 = \frac{(1+\xi)(1+\eta)}{4} \tag{2662}$$

$$N_4 = \frac{(1-\xi)(1+\eta)}{4} \tag{2663}$$

For a three-node isoparametric triangle in 3D, the shape function at each node is:

$$N_1(\xi, \eta) = 1 - \xi - \eta \tag{2664}$$

$$N_2(\xi, \eta) = \xi \tag{2665}$$

$$N_3(\xi, \eta) = \eta \tag{2666}$$

Note: The shape function scheme is applicable only for triangle and quadrilateral elements in 2D and tetrahedral and hexahedral elements in 3D.

When mapping with shape functions from a FEA mesh to a STAR-CCM+ mesh, STAR-CCM+ recognizes the element type and uses the relevant shape function for interpolation. STAR-CCM+ supports higher order elements with mid-side nodes for Abaqus and other FEA codes.

When projecting points from STAR-CCM+ to the FEA mesh, however, STAR-CCM+ only uses corner nodes for the parametric map. Thus the mapper projects to Finite Element triangle faces as if they are flat. Projecting to Finite Element quad faces accounts for warping if the corner nodes do not lie in one plane.

Exchanging Heat Transfer Coefficients

In convective heat transfer studies, the local heat flux at the wall is given by Newton's law of cooling, that is,

$$\dot{q}''_w = h(T_w - T_r) \tag{2667}$$

where h is the heat transfer coefficient, T_w is the wall temperature, and T_r is the fluid reference temperature. Examples of fluid reference temperatures are the near-wall temperature, or the mean fluid temperature. An important point is that the heat transfer coefficient and reference temperature are paired; there are many pairs that give the same wall heat flux.

In conjugate heat transfer calculations, where all calculations are part of the same simulation, there is no concern over the choice of heat transfer coefficient and reference temperature. In conjugate simulations the heat flux and wall temperature are calculated by the implicitly coupled system of equations. In co-simulation calculations, where the fluid and solid parts of the geometry are treated separately, you must choose the information to exchange at the interface between the simulations.

This section of the documentation is intended to help you choose the right approach for exchanging heat transfer coefficients. The following topics are discussed:

- What is a general methodology for a coupled thermal analysis?
- What methods are available for exchanging heat transfer coefficients?
- What are recommended best practices for exchanging heat transfer coefficients?

What Is a General Methodology For an Explicitly Coupled Heat Transfer Analysis?

An explicitly coupled thermal analysis is one in which the non-isothermal flow problem is solved in one STAR-CCM+ simulation, and the solid thermal problem is solved in a second simulation either using STAR-CCM+ or another CAE code. This is typically done for analyses in which solid components exchange heat with a fluid whose properties vary in a cyclic manner, or cases in which the fluid and solid have disparate time scales.

An example for which this approach would be appropriate is an analysis of the heat transfer between the exhaust gases of an internal combustion engine (ICE) and its exhaust manifold. In such an analysis, the flow field leaving the engine will fluctuate in response to the exhaust valve dynamics. This means that the temperature of the fluid entering the exhaust manifold

will experience rapid changes to the order of several hundred degrees. The solid body of the manifold will not respond so quickly to these temperature changes due to its greater thermal inertia relative to the fluid, so this results in a conjugate heat transfer problem where the time step that is appropriate for the fluid analysis is smaller than the time step that is appropriate for the solid thermal analysis -- hence the need to have two distinct simulations explicitly coupled by an exchange of thermal information on their common boundary. The solid thermal analysis will, in fact, tend towards a steady state solid conduction problem, as small fluctuations in solid temperatures during an engine cycle are not usually of interest.

A suggested methodology for conducting this analysis is illustrated in the figure below. Two separate simulations are set up, one for the fluid analysis and another for the solid thermal analysis.



In the first step, an assumed wall temperature is applied to the wall boundaries of the transient flow problem. This flow problem is solved for a total time equal to one engine cycle. It is important that this initial assumed wall temperature be reasonable in this step, as otherwise the solution may not converge. Following the simulation, a time-averaged, spatially varying, heat transfer coefficient and its corresponding reference temperature (both of which are calculated using field mean monitors and field functions) are calculated on the wall boundaries and then passed to the solid thermal simulation which is being run in steady state. To ensure that the time averaged heat transfer at the interface is conserved, the reference temperature (which is not the time averaged heat transfer reference temperature) that must be used is defined as:

$$\hat{T}_r = T_w - \frac{\overline{\dot{q}''_w}}{\overline{h}} \tag{2668}$$

where $\overline{\dot{q}''_w}$ is the time averaged interface heat flux and $\bar{h}$ is the time averaged heat transfer coefficient. The solid thermal simulation is then run to convergence.

In the second step, the wall temperatures predicted by the solid thermal simulation are passed back to the transient flow simulation and applied to the wall boundaries. The flow simulation is solved for a second time to obtain a second set of heat transfer coefficients and reference temperatures. This new set of data is passed to the solid thermal simulation for a second time.

This iterative process continues until successive steps show little or no meaningful change in the values of the predicted heat transfer coefficients or temperatures. Some situations may require many co-simulation cycles until the overall converged state is achieved.

A key question in this approach is: what heat transfer coefficient and reference temperature should be exchanged between the simulations? STAR-CCM+ provides several choices.

What Methods Are Available for Exchanging Heat Transfer Coefficients?

For turbulent flows, the central concept for modeling convective heat transfer at the wall is given by the standard wall function, which can be written as:

$$\dot{q}''_w = \frac{\rho_f(y_c)C_{p,f}(y_c)u^*(y^+(y_c))}{T^+(y^+(y_c))}(T_w - T_r) \tag{2669}$$

where y_c is the distance to the near wall cell center, y^+ is the non-dimensional wall distance, u^* is the reference velocity, and T^+ is the non-dimensional temperature -- both of which are determined by the chosen wall treatment. ρ_f and $C_{p,f}$ are the density and specific heat of the fluid. Note that ρ_f, $C_{p,f}$, u^*, and T^+ are functions of the distance y_c from the wall.

Comparing Eqn. (2669) with Eqn. (2667), the heat transfer coefficient can be written as:

$$h(y_c) = \frac{\rho_f(y_c)C_{p,f}(y_c)u^*(y^+(y_c))}{T^+(y^+(y_c))} \tag{2670}$$

This expression is the foundation for various definitions of heat transfer coefficient in STAR-CCM+. These coefficients, and their corresponding fluid reference temperatures, are listed in the table below.

Field Function	Reference Temperature for Data Exchange
Local Heat Transfer Coefficient	Local Heat Transfer Reference Temperature
Specified Y+ Heat Transfer Coefficient	Specified Y+ Heat Transfer Reference Temperature
Heat Transfer Coefficient	User-defined reference temperature
Virtual Local Heat Transfer Coefficient	N/A

Local Heat Transfer Coefficient

This is the heat transfer coefficient evaluated at the centroid of the cell next to the boundary, that is $h(y = y_c)$. It is given by:

$$h_{local} = \frac{\rho_f(y_c)C_{p,f}(y_c)u^*}{T^+(y^+(y_c))} \tag{2671}$$

The corresponding fluid reference temperature is given by:

$$T_{local} = T_w + \frac{\dot{q}''_w}{h} \tag{2672}$$

Specified Y+ Heat Transfer Coefficient

This is the heat transfer coefficient evaluated at a user specified value of y^+, such that $y^+(y) = Y^+_{user}$. It follows that:

$$h_{y^+} = \frac{\rho_f(y_c)C_{p,f}(y_c)u^*}{T^+(Y^+_{user})} \tag{2673}$$

The corresponding fluid reference temperature, the Specified Y+ Heat Transfer Reference Temperature is given by:

$$T_{f,y^+} = T_w + \frac{\dot{q}''_w}{h_{y^+}} \tag{2674}$$

As this definition does not guarantee that T_{f,y^+} lies within the temperature range of the local region, the following conditions are enforced to give realistic values:

$$T_{f,y^+} = \begin{cases} T_w + \dfrac{\dot{q}''_w}{h_{y^+}} & T_{min} \le T_w + \dfrac{\dot{q}''_w}{h_{y^+}} \le T_{max} \\ T_{min} & T_w + \dfrac{\dot{q}''_w}{h_{y^+}} < T_{min} \\ T_{max} & T_w + \dfrac{\dot{q}''_w}{h_{y^+}} > T_{max} \end{cases} \tag{2675}$$

In the latter two cases, h_{y^+} will be re-evaluated such that

$$h_{y^+} = \frac{\dot{q}''_w}{(T_{f,y^+} - T_w)} \tag{2676}$$

For these field functions, the y^+ value is set on the Specified Y+ Heat Transfer Coefficient field function node within the Tools > Field Functions manager node.



Virtual Local Heat Transfer Coefficient

This is an approximation to the true Local Heat Transfer Coefficient that does not require an energy model to be activated. It is defined as the heat transfer coefficient evaluated at $y = y_c$, where y_c is the distance of the near wall cell center from the wall.

$$h_{virtual} = \frac{\rho_f(y_c)C_{p,f}(y_c)u^*}{T^+(y^+(y_c))} \tag{2677}$$

Since thermal material properties are not available, the specific heat, $C_{p,f}$ is specified as a reference value, as well as the molecular and turbulent Prandtl numbers. These are set on the Virtual Local Heat Transfer Coefficient field function node, as shown below. The fluid density, ρ_f is obtained locally from the simulation.



Heat Transfer Coefficient

This is yet another choice for defining heat transfer coefficients. It is evaluated based on a specified bulk temperature T_{bulk}.

$$h_{bulk} = \frac{\dot{q}''_w}{(T_{bulk} - T_w)} \tag{2678}$$

This definition does not guarantee that the evaluated heat transfer coefficient is positive. When exchanging data with other CAE codes, if the evaluated h_{bulk} is not positive, then the local heat transfer coefficient, h_{local}, and its corresponding reference temperature, T_{local}, are exchanged instead.

The bulk temperature is set as a property of the Heat Transfer Coefficient field function node.



What Are Recommended Best Practices for Exchanging Thermal Data?

The following best practices should govern the exchange of thermal data in a file-based coupling of STAR-CCM+ with another CAE code:

1. When explicitly coupling a fluid and a solid simulation, the suggested practice is for:
   - the fluid model to provide spatially varying heat transfer coefficients and their corresponding reference temperatures to the solid model, and,
   - the solid model to return the local wall temperatures to the fluid model.

   In order to avoid issues with numerical stability, you are advised to avoid alternative techniques of exchanging thermal boundary data between STAR-CCM+ and other applications.

   Two specific scenarios are:

   a. For the case of fluid flow with, for example, no radiative transport or phase-change, use the Specified y+ Heat Transfer Coefficient together with the Specified y+ Heat Transfer Reference temperature. A relatively large value of y^+ such as 100 should be specified, in

which case, the local temperature difference $(T_w - T_r)$ in Eqn. (2667) will be relatively large due to the relatively large wall distance. This in turn should allow the solid temperatures room to change significantly, and so allow the models to converge to a steady state in fewer exchanges between them.

b. For the case of fluid flow with radiative transport and/or phase-change, use the Local Heat Transfer Coefficient together with the Local Heat Transfer Reference Temperature.

2. Exchanging local heat fluxes in one direction (whether fluid to solid, or solid to fluid), and receiving local wall temperatures in return, is not a recommended practice for reasons of numerical stability.
3. More than a single round of data exchange between the fluid and solid models will be required. A minimum of two rounds is recommended. This is because the initial guess of wall temperature for the fluid model may have been significantly different from the final solution.
4. During the first round, the initial temperature distribution in the solid model should be the same as the initial wall temperature used in the fluid model.
5. As either the fluid or the solid model must be solved first in the initial round, it is suggested that the fluid model be solved first. An initial guess for the value of wall temperature is required.
6. When comparing the heat transfer coefficient predictions with benchmarks (either measurements or model predictions), it is important to ensure that the heat transfer reference temperature is the same as in the benchmark. Otherwise the comparisons of heat transfer coefficient can be significantly different.
7. In a thermo-fluids CFD analysis, the energy equation should be solved (which is usually relatively cheap, and should always be done when circumstances allow, especially when temperature-dependent properties are involved as in buoyant flows).
8. When the energy equation is not solved, the local heat transfer coefficient is called the Virtual Local Heat Transfer Coefficient.
9. When using a low-y+ mesh for turbulent flows, you should reduce the thickness of the prism cells adjacent to walls until the fluxes at the walls become grid independent.
10. For laminar flows, you should reduce the thickness of the prism cells adjacent to walls until the fluxes at the walls become grid independent.
11. As the near wall region is coupled to the flow in the interior of the domain, it is important that the interior solution is also grid independent within practical constraints.

Working with Specific CAE Products

This section discusses the specific aspects of file-based coupling with each of the following software packages:

- Abaqus
- ANSYS
- NASTRAN
- RadTherm
- es-ice or pro-STAR

Also covered are the following topics:

- Use of Ensight input files for comparing multiple solutions in the same simulation
- The special format available in STAR-CCM+, .sbd (STAR-CCM+ Boundary Data)
- The Visualization Toolkit (VTK) format

The supported versions of these software packages are listed in the Release Notes.

Working with Abaqus

This section provides notes on the following topics that are related to importing from and mapping to Abaqus with STAR-CCM+:

- Preparing data in Abaqus
- Importing Abaqus solution data
- Exporting data to Abaqus

Preparing Data in Abaqus

The Abaqus geometry data is imported from a standard ASCII-format Abaqus keyword input file (`.inp`). This data is for surfaces and volume elements that are defined in Abaqus. Surfaces are defined as a set of faces of 3D continuum elements and/or faces of shell elements. Volumes are defined as sets of 3D continuum elements. STAR-CCM+ supports surface-to-surface and volume-to-volume mapping, and the meshes do not need to be conformal. Before proceeding in STAR-CCM+, first prepare the input file in Abaqus as follows:

- Define the Abaqus surfaces for data to be mapped, using the `*Surface` keyword command. Generate this surface using the Tools > Surfaces > Create dialog in Abaqus/CAE.
- For volume import and mapping, define nodes and elements in the `.inp` file (using the `*NODE` and `*ELEM` commands). STAR-CCM+ automatically imports them into the simulation and create imported volume objects.

  Data can then be mapped to the surface and/or volume element from selected parts of the STAR-CCM+ mesh.

Importing Abaqus Solution Data

To import data from Abaqus, you must have access to an Abaqus installation. In the *Import Data Options* dialog for an Abaqus model, additional options are required in order to set up the import operation. The dialog shown in the example below is for importing Temperature data.



To set the additional options, do the following:

- In the Abaqus Command text box enter the command you use to launch Abaqus, for example `abq6102`. Typically this is based on your installation path for Abaqus. On Windows the command is `abaqus.bat`
- In the Load Step Name text box, enter the name of the step in the Abaqus analysis job for which you would like to import data. The default name for the first load step in an Abaqus input file is `Step-1`.

When you click OK the data is imported into STAR-CCM+. During this process STAR-CCM+ launches Abaqus in the background. If you do not have access to an Abaqus installation, an error message is displayed.

Exporting Data to Abaqus

When the export action is executed, an Abaqus `.inp` file is produced that contains mapped data from the STAR-CCM+ solution. This file is a standard ASCII-format keyword-driven Abaqus file. STAR-CCM+ automatically writes the corresponding keyword for the type of field data that is exported from the CFD model at the top of the data set:

- Pressure loads are defined via the `*DLOAD` keyword.
- The heat flux is defined via the `*DFLUX` keyword.
- Film coefficient and temperatures are defined via the `*FILM` keyword.
- Temperatures on surfaces are defined via the `*BOUNDARY` keyword for use in a heat transfer analysis.
- Temperatures on a volume are defined via the `*TEMPERATURE` keyword, which can only be used in an uncoupled temperature-displacement analysis. It cannot be used for heat transfer or coupled temperature-displacement analyses.

You can use the exported field data from STAR-CCM+ in an analysis job by adding an extra line in the job file (`.inp`) that instructs Abaqus to read in the load data. This line uses the `*INCLUDE` keyword, and is added in the history section of the input file in the appropriate load step. For example, the following line would be added between the `*STEP` and `*End Step` keywords for the step in which the load data is applied to the Abaqus model:

```
*INCLUDE, input=mappedHTC.inp
```

Where `mappedHTC.inp` is the name of a data file containing mapped heat transfer coefficients from STAR-CCM+.

Working with ANSYS

This section provides information on the following topics:

- Preparing data in ANSYS
- Importing ANSYS solution data
- Mapping STAR-CCM+ to ANSYS
- Exporting to ANSYS

Preparing Data in ANSYS

An imported model from ANSYS contains the full FEA mesh as well as FSI interfaces for mapping data to and from STAR-CCM+ and ANSYS. The wetted FSI interfaces are defined through *Components* of *Nodes* and *Elements* in ANSYS, which are named lists of nodes or elements. Use either surface or volume FSI interfaces.

Defining Surface and Volume Interfaces using Nodes and Elements

Node components can only be used to specify the FSI surface interfaces. The nodes must define an external (free) surface on the FEA model (use the **EXT** item of the ANSYS command **NSEL**), and only element faces for which all the nodes are selected are included in the surface definition. Node components are the only method for defining FSI surfaces on ANSYS SOLID elements.

It is recommended that element components are used to define FSI surface interfaces on SHELL elements. Element components are also used to define FSI volume interfaces for collections of SOLID elements. Both wall shear stress and pressures can be used with this component type.

For a Node or Element component to be imported for mapping results, include the string "FSI" or "fsi" somewhere in its name. When the ANSYS model is imported into the simulation, STAR-CCM+ recognizes these components and display them in the object tree; other components are automatically filtered out.

ANSYS Element Types

The following are the list of all supported ANSYS elements:

Type	Element Name	Description
3D Shells	SHELL43	4-Node Plastic Large Strain Shell
	SHELL63	4-Node Elastic Shell
	SHELL93	8-Node Structural Shell
	SHELL131	4-Node Layered Thermal Shell
	SHELL132	8-Node Layered Thermal Shell
	SHELL143	4-Node Plastic Small Strain Shell
	SHELL150	8-Node Structural Shell p-Element
	SHELL163	Explicit 4-Node Thin Structural Shell
	SHELL181	4-Node Finite Strain Shell

	SHELL281	8-Node Finite Strain Shell
3D Continuum Solid	SOLID5	8-Node Coupled-Field Solid
	SOLID45	8-Node Structural Solid
	SOLID46	8-Node Layered Structural Solid
	SOLID62	8-Node Magneto-Structural Solid
	SOLID64	8-Node Anisotropic Structural Solid
	SOLID65	8-Node Reinforced Concrete Solid
	SOLID69	8-Node Coupled Thermal-Electric Solid
	SOLID70	8-Node Thermal Solid
	SOLID72	4-Node Linear Tet, obsolete
	SOLID73	8-Node Linear Brick, obsolete
	SOLID87	10-Node Tetrahedral Thermal Solid
	SOLID90	20-Node Thermal Solid
	SOLID92	10-Node Tetrahedral Structural Solid
	SOLID95	20-Node Structural Solid
	SOLID98	Tetrahedral Coupled-Field Solid
	SOLID147	Brick Structural Solid p-Element
	SOLID148	Tetrahedral Structural Solid p-Element
	SOLID164	Explicit 3D Structural Solid
	SOLID168	Explicit 3D 10-Node Tet Structural Solid
	SOLID185	8 Structural Solid or Layered Solid
	SOLID186	20 Structural Solid or Layered Solid
	SOLID187	10-Node Tetrahedral Structural Solid
	SOLSH190	8-Node Layered Solid Shell
	SOLID191	20-Node Layered Structural Solid
	SOLID226	20-Node Coupled-Field Solid
	SOLID227	10-Node Coupled-Field Solid
	VISCO89	20-Node Viscoelastic Solid
	VISCO107	8-Node Viscoplastic Solid

Surface Elements 3D	SURF152_8_N	See notes below.
	SURF154_8_N	See notes below.

In ANSYS, tractions are applied to special "Surface" element types: SURF154. This is the only element type that can have a traction other than pressures in a three-dimensional analysis.

Thermal loads are applied directly to solid element faces but can also be applied to special "Thermal" element types: SURF152.

There are two methods of defining a fluid-structure interaction (FSI) surface of these SURFACE elements:

1. As an ANSYS component of Elements, with the letters "FSI" in the name.
2. As a surface node set like all of the other surface definitions for solid elements, with the letters "FSI" in the name.

The first method is the most straightforward. With method (2), using a node set, there are two surfaces created: one from the faces of the solid elements, and one from faces of the surface elements. In this case, "-SOLID" are appended to the name of the surface comprised of solid element faces, and "-SURF" are appended to the name of the surface that is comprised of surface elements.

Exporting using the .cdb file

ANSYS models are imported into STAR-CCM+ in the `.cdb` format. When the file is generated via the CDWRITE command in ANSYS, only some of the options write enough information so that STAR-CCM+ can reliably read the file.

Use the following options:

CDWRITE, **DB**, <*filename*>.cdb

CDWRITE, **COMB**, <*filename*>.cdb

CDWRITE, **ALL**, <*filename*>.cdb

Do not use the following options:

CDWRITE, **GEOM**, <*filename*>.cdb

CDWRITE, **CM**, <*filename*>.cdb

CDWRITE, **LOAD**, <*filename*>.cdb

CDWRITE, **SOLID**, <*filename*>.cdb

Further Considerations for Model Preparation

The Imported Model FSI Surface for a single ANSYS component is divided into multiple surfaces by element style: SOLID, SHELL, and/or SURF. If the ANSYS component name is "fsi_surface", then the FSI surface names in the STAR-CCM+ imported model are:

Element Style	Example Surface Name
SOLID	`fsi_surface-SOLID_FACE`
SURF	`fsi_surface-SURF-EL`
SHELL	`fsi_surface-SHELL:POS` `fsi_surface-SHELL:NEG`

This is useful for solid models with Node components, where SURF or SHELL elements are overlaid on the SOLID free surface. This allows you to apply traction loads to the SURF element set and heat flux loads to the SOLID element set in the same simulation.

For each SHELL element, note that there are two sides: a positive and a negative side. These are determined by the right-hand rule applied to the nodal sequence of the element definition. If you know that the side of a SHELL element is either positive or negative, you can include "POS" or "NEG" in the component name, and the Imported Model only includes one surface for the SHELL elements, for example either "fsi_surface-SHELL:POS" or "fsi_surface-SHELL:NEG".

Importing ANSYS Solution Data

Displacement, eigenmode and eigenvector, and temperature results can be imported from ANSYS listing files for an ANSYS model into a STAR-CCM+ simulation. Write the results to a `.lis` file from ANSYS in order for STAR-CCM+ to import them.

Before listing results in ANSYS, load the results using the following commands in ANSYS:

`/POST1`	Enter the ANSYS Post Processor.
`SET,LS,SS`	Set load step and substep to load.
`INRES,NSOL`	Load the nodal solution for the selected load step, substep combination.

To create the listing file, the following commands are needed:

`/OUTPUT,fname,ext`	Open the file `fname.ext` for output.
`PRNSOL,NAME,OPT`	Print the output to `fname.ext`.
`PRNSOL,TEMP`	Print temperature to the file.
`PRNSOL,U,COMP`	Print the three displacement or eigenvector components to the file, depending on whether analysis is structural or modal.
`/OUTPUT`	An additional /OUTPUT command with no parameters to close the file and send further output back to the screen.

The above commands are general instructions for obtaining output data from ANSYS. Before importing the data into the STAR-CCM+ simulation, export the data to a `.lis` file (the default format for the /OUTPUT command in ANSYS).

Note: While STAR-CCM+ reads displacements and temperatures for a particular load step from a `.lis` file, it reads *all* eigenmode results in the file.

Example Temperature Listing creation:

`/POST1`	
`SET,1`	! set load step 1, (substep 1 is default)
`INRES,NSOL`	
`/OUTPUT,ansys_temperature,lis`	! open new file `ansys_temperature.lis`
`PRNSOL,TEMP`	
`/OUTPUT`	! turn off output to file

Example Displacement Listing creation:

`/POST1`	
`SET,1`	! set load step 1, (substep 1 is default)
`INRES,NSOL`	
`/OUTPUT,ansys_disp,lis`	! open new file `ansys_disp.lis`
`PRNSOL,U,COMP`	! print all displacement components
`/OUTPUT`	! turn off output to file

Example Eigenmode Listing creation for the first six modes of a modal analysis:

```
/POST1
SET,1                          ! set load step 1, (substep 1 is default)
INRES,NSOL
/OUTPUT,ansys_modes,lis        ! open new file ansys_modes.lis
SET,FIRST
PRNSOL,U,COMP
SET,NEXT
PRNSOL,U,COMP
SET,NEXT
PRNSOL,U,COMP
SET,NEXT
PRNSOL,U,COMP
SET,NEXT
PRNSOL,U,COMP
SET,NEXT
PRNSOL,U,COMP
/OUTPUT                        ! turn off output to file
```

Mapping the STAR-CCM+ solution to the ANSYS mesh

In the STAR-CCM+ CAE mapping tools, map each fluid boundary to the associated ANSYS exterior surface detected using the source/target proximity checks. Depending on the conformality of the grids at the edges of the wetted area and the proximity tolerance, there may be some stretching of the wetted area onto non-wetted faces by the amount of the proximity tolerance.

Exporting to ANSYS

Data is exported from STAR-CCM+ to ANSYS via ANSYS macro scripts. Load the data onto the ANSYS model with the appropriate load step activated. The ANSYS command to input a macro script is "/INPUT,<*filename*><*extension*>", or select File > Read Input From.

In the ANSYS macro scripts, data exported from STAR-CCM+ set values in the ANSYS model, as follows:

- For Surface interfaces:
  - Vertex temperatures are used to set ANSYS boundary conditions (ANSYS command: **D**).
  - Face values, including pressures, local heat transfer coefficient and reference temperature, and heat flux, are used to specify ANSYS surface loads (ANSYS command: **SFE**).
- For Volume interfaces:
  - Vertex temperatures are used as ANSYS loads or field variables (ANSYS command: **BF**).

Working with NASTRAN

This section provides notes on the following topics related to importing from and mapping to NASTRAN with STAR-CCM+:

- Preparing data in NASTRAN
- Importing NASTRAN solution data
- Exporting Data to NASTRAN
- Exporting a Model to NASTRAN

Preparing Data in NASTRAN

STAR-CCM+ supports the import of shells, volumes (solids) and surface elements from NASTRAN models text files with the `.nas`, `.bdf` or `.dat` file extensions.

Importing NASTRAN Solution Data

Once a NASTRAN model has been imported, its associated solution data can be imported into STAR-CCM+ from the OUTPUT2 (`.op2`) binary file format.

Exporting Data to NASTRAN

Data that has been mapped onto an imported NASTRAN model can be exported from STAR-CCM+ to external files.

To export data from STAR-CCM+ to NASTRAN:

1. Import the NASTRAN model into STAR-CCM+ (see Importing a CAE Model).
2. Map the data that you want to export to NASTRAN onto the imported NASTRAN model (see Mapping STAR-CCM+ Solution Data to a CAE Model).

3. Export the mapped data to external files (see Exporting Mapped Data), choosing `.nas` as file extension.

The export files, saved using the `.nas` file extension, are text-based files that contain the Bulk Data keywords for the selected export quantity. To use the exported data in a NASTRAN simulation, you can append it to the NASTRAN simulation file by using the `INCLUDE` keyword, for example `INCLUDE 'starccm_pressure.nas'`.

Exporting a Model to NASTRAN

When you export a model to NASTRAN you can choose if patches with the same name are merged or not. Activating Merge Cell Types by Name exports the boundary/part surfaces after merging the patches with the same name. This typically leads to one Property ID (PID) for the whole boundary/part surface instead of one PID per patch.



Working with RadTherm

This section provides notes on the following topics related to importing from and mapping to RadTherm with STAR-CCM+:

- Preparing a Model in RadTherm
- Importing a RadTherm model
- Importing and Mapping RadTherm data
- Analyzing with RadTherm
- Known limitations

It is also possible to export simulation data to RadTherm.

Preparing a Model in RadTherm

To be imported into STAR-CCM+, export a RadTherm file to the Patran Neutral file format. A common technique is the RadTherm command-line utility, `tdf2patran` (`tdf2patran.exe` on Windows) that is bundled with RadTherm. In this utility you would execute the following at the command prompt:

```
tdf2patran -temp -v -f back <RadThermFileName>.tdf
```

The `-temp` and `-v` options above are required to ensure that nodal temperatures are exported, rather than the temperature at the surface cell centroids. The `-f back` (or `front`) option can be omitted to export both the front/back nodal temperatures.

A quick inspection of the generated file is recommended so as to ensure that at a minimum, the generated file contains the following Patran "packets", which are required for successful import into STAR-CCM+:

- Packets 1 & 2 for the nodal coordinates and element connectivity/type respectively
- Packet 4 for the element/nodal properties that are associated with the parts
- Packet 10 for the exported nodal temperatures at both the front/back faces of each shell
- Packet 21 for the parts definition, which consists of listings of the element type, ID for each exported part or component, and the part name

By default the RadTherm Patran utility exports thermal quantities at both the front/back surfaces of a shell or component. A negative temperature value of `-990` is assigned to the nodes of a face that does not have a user-specified temperature during model export. However, you can select either the front or back shell surface when you map the surface data.

Importing the RadTherm Model

After you launch STAR-CCM+ and start a simulation, you can import the RadTherm model (`.ntl` or `.neu`) in one of two ways:

- If the CFD model is being developed from scratch, then select File > Import Surface Mesh... to import the RadTherm model as a surface mesh or geometry part. This import option is useful if you wish to generate a volume mesh from the RadTherm model, for example to compute a fluid solution for a volume that is adjacent to the RadTherm shells.

- If the STAR-CCM+ simulation model is already loaded, then select File > Import CAE Model... to load the RadTherm model file. This option allows you to import and map solution data for the RadTherm model. Note that the Create new Region option in the *Import Options* dialog cannot be used for RadThem model import, as it does not apply to shells.

To set up a fluid-structure interaction case, it may be necessary to use the first import method to create the fluid volume, then use the second method to import the structural shells. File-based coupling techniques can then be used to couple the STAR-CCM+ and RadTherm models.

Mapping and Exporting to RadTherm

Solution data from STAR-CCM+ can be mapped and exported to the imported RadTherm model.

When you map surface data from STAR-CCM+ to the imported model, you can select which shell surface to export the data for. To do this, in the Imported Surface section of the dialog expand all the nodes for the RadTherm model and select one of the shell surfaces, referred to as Front and Back. Currently it is only possible to map heat transfer coefficient data to the RadTherm model using the mapping dialog. However, if you wish to map a different field function, the full list of field functions can be made available by manually creating data mappers via the following steps:

- Setting data mapper properties
- Mapping the data
- Specifying CAE-specific properties

Exporting to RadTherm

You can export solution data to Radtherm using similar techniques as for other CAE models.

You can also export mesh and solution data from STAR-CCM+ to RadTherm.

Select File > Export ... to open the *Export File d*ialog.



In this dialog, set up the export as follows:

- Select the `RadTherm Patran Neutral Files (*.ntl)` option from the *Files of Type* drop-down list.
- Select the interface boundaries of interest from the list in the Boundaries tab.
- Switch to the Scalars tab and select `Local Heat Transfer Coefficient` and `Local Heat Transfer Coefficient Reference Temperature` from the list. Only these two quantities are currently supported for this direct type of export to RadTherm.
- Enter a name for the output file and click Save.

The generated file is ready for import into RadTherm.

Analyzing with RadTherm

When you have finished the mapping work in STAR-CCM+, export the data to the Patran Neutral file format (`.ntl`). Then do as follows:

- Import the file into RadTherm.
- Set up the boundary conditions and thermal parameters in RadTherm and run the analysis.
- Prepare and export the results from the RadTherm file to the Patran Neutral file format.
- Import the data into STAR-CCM+.
- Export the STAR-CCM+ fluid analysis results into the Patran Neutral file format.
- Back in the RadTherm file, import the Patran file created by STAR-CCM+.
- If there are no error messages after importing the geometry, use RadTherm's Import CFD Results feature to map the STAR-CCM+ results to the geometry of the RadTherm file.
- Use RadTherm's Model Parameter Display feature to view the fluid reference temperature values, and make sure they are consistent with those generated by STAR-CCM+.

Known Limitations

The following are the known limitations of the current implementation:

- The exchange of data is limited to steady-state, or data at a single time-step. The import/export of thermal data at multiple time steps, or time-levels is not supported.
- The current implementation does not support the conversion of units of measurement. It is the user's responsibility to ensure consistency of units in both RadTherm and STAR-CCM+.
- It is currently not possible to import into STAR-CCM+, "dummy" Patran shells, which are defined as Packet-21 entries that have no element type, and ID attributes. Delete dummy entries from the Patran Neutral file after export from RadTherm, since they are not required for surface definition in STAR-CCM+.

Working with es-ice or pro-STAR

The `.dbs` file format is commonly used to import a surface mesh. It can also be used to map solution data from pro-STAR or es-ice to STAR-CCM+ by importing the file as a CAE model, rather than a surface mesh. When used in this manner, store the solution data in a `.usr` file that was created by the same session of pro-STAR or es-ice that created the `.dbs` file.

For pro-STAR, up to six sets of wall solution data can be stored in the six Registers available in the `.usr` file. STAR-CCM+ can import solution data from these six Registers to the CAE model, and then map the solution data on to boundaries.

For es-ice, two `.usr` files are used to export cycle-averaged wall heat transfer data.

This section provides notes on the following topics related to obtaining solution data from es-ice or pro-STAR:

- Preparing wall heat transfer data in es-ice.
- Preparing wall data in pro-STAR.

As for mapping data from es-ice or pro-STAR, there are three steps in the process:

- Import the model as a CAE model (`.dbs` file).
- Imported the associated solution data (`.usr` file).
- Map the imported solution data from the surface of the CAE model to the STAR-CCM+ boundaries.

An example of the es-ice cycle-averaged heat transfer coefficient displayed as a scalar on the imported CAE model is shown below.



The following image shows the same imported heat transfer coefficient mapped onto the boundaries of a STAR-CCM+ volume mesh.



Preparing Wall Heat Transfer Data in es-ice

Calculating spatially varying, cycle-averaged wall heat transfer data in es-ice produces CAE data, which can be imported into STAR-CCM+. This es-ice post-processing facility generates a `.dbs` file, containing the CAE model and calculates eight data-sets. As the `.usr` file can only contain six data-sets at most, es-ice writes two .usr files, with the data described below:

- Data file 1:

Register 1: Average Heat Transfer Coefficient (Watts/m2-K)

$$\bar{h} = \frac{1}{\theta_2 - \theta_1}\int_{\theta_1}^{\theta_2} h(\theta)d\theta$$

Register 2: Average Near-Wall Gas Temperature (K)

$$\overline{T_g} = \frac{1}{(\theta_2 - \theta_1)\bar{h}} \int_{\theta_1}^{\theta_2} h(\theta) T_g(\theta) d\theta$$

Register 3: Average Heat Flux (Watts/m2)

$$\overline{q''} = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} q''(\theta) d\theta$$

Register 4: Average Wall Boundary Temperature (K)

$$\overline{T_w} = \frac{1}{(\theta_2 - \theta_1)\bar{h}} \int_{\theta_1}^{\theta_2} h(\theta) T_w(\theta) d\theta$$

Register 5: Average y^+

$$\overline{y^+} = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} y^+(\theta) d\theta$$

Register 6: Average Distance from Boundary to y^+=100 (m)

$$\overline{D} = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} D(\theta) d\theta$$

Register 2: Average Near-Wall Gas Temperature (K)

- Data file 2:

Register 1: Average Heat Transfer Coefficient at y^+=100 (Watts/m2-K)

$$\overline{h_y} = \frac{1}{\theta_2 - \theta_1}\int_{\theta_1}^{\theta_2} h_y(\theta)d\theta$$

Register 2: Average Near-Wall Gas Temperature at y^+=100 (K)

$$\overline{T_y} = \frac{1}{(\theta_2 - \theta_1)\overline{h_y}}\int_{\theta_1}^{\theta_2} h_y(\theta)T_y(\theta)d\theta$$

The nomenclature for the above equations is provided below:

Notation	Description
θ	Crank Angle (degrees)
θ_1	Start Crank Angle (degrees)
θ_2	End Crank Angle (degrees)
h	Heat Transfer (W/m^2-K)
h_y	Heat Transfer at y^+=100 (W/m^2-K)
T_g	Temperature of near-wall gas (K)
T_w	Temperature of wall boundary (K)
T_y	Temperature at y^+=100 (K)
q''	Heat Flux (W/m^2)
y^+	Non-dimensional wall co-ordinate, see Eqn. (449)
D	Distance from boundary to y^+=100 (m)

The solution data is defined with respect to cell IDs, which must match exactly to the cell IDs in the `.dbs` file. To ensure the cell IDs correlate, follow the procedure below to generate this solution data.

- Run an es-ice analysis with the Heat Transfer option selected in the *Star Controls* > *Post-processing* > *Global Settings* panel. This generates heat transfer data (`es-ice_htx.pos`) while STAR is obtaining a solution.



- When the solver is complete, open the *Post-process* panel and select Heat Transfer from the drop-down menu at the top of the panel.



- Click the ellipsis to select `es-ice_htx.pos` from the file browser and then click Add after selected. STAR has generated this file when running the analysis.



- Click Load post data and enter an appropriate Crank angle range to cover the last full engine cycle. Then click Cycle average to calculate the cycle-averaged wall heat transfer data. The start crank angle (θ_1 in the equations above) is entered in the first text field, and the end crank angle (θ_2) in the second text field



- Specify the file names of the `.dbs` and two `.usr` files, set the database ID to 1 and click Write cycle avg. The CAE model (`.dbs`) and the solution data (`.usr`) is located in the current working directory.



Currently, STAR-CCM+ can only import coded (ASCII) format `.usr` files, whereas es-ice generates binary format files. Therefore, an extra step is required to convert the binary format file to a coded format file in pro-STAR.

1. Import the `.dbs` file into pro-STAR using the following commands:

   ```
   dbase, open, <filename>.dbs
   dbase, get, 1
   ```

   This ensures the cell IDs in the solution data match the CAE model.

2. Read the binary format `.usr` file into pro-STAR using the following commands:

   ```
   getu, <filename>.usr, cell, all, binary
   ```

3. Save the solution data to a coded format `.usr` file with the following commands:

   ```
   savu, <filename>.usr, all, coded, all
   ```

4. Repeat steps 2 and 3 for the second `.usr` file.

Preparing Wall Data in pro-STAR

pro-STAR can export the surface of a mesh into a `.dbs` file. As discussed previously, this surface can be imported into STAR-CCM+ as either a surface mesh or a CAE model. pro-STAR can also export solution data (steady-state or transient) to a .usr file, which can be imported into STAR-CCM+ and mapped onto boundaries.

Only surface meshes from pro-STAR can be imported into STAR-CCM+ as a CAE model. Therefore, you can only export wall (surface) post-processing data from pro-STAR, as opposed to cell (volume) data. Note: STAR-CCM+ can only import solution data from a coded (ASCII) format `.usr` file. Remember these limitations when preparing data in pro-STAR.

The solution data is defined with respect to cell IDs, which must match exactly to the cell IDs in the `.dbs` file. Otherwise, the solution data are imported incorrectly. To ensure the cell IDs correlate, follow the procedure below to generate the CAE model and solution data.

1. Load the pro-STAR model and post-processing data. If you are running a transient analysis, you must also store the data for a given time-step.
2. Read the required solution data from the stored post-processing data. The command shown below reads the wall temperature data:

   ```
   getw, t
   ```

3. Plot and export the surface mesh to a `.dbs` file using the following commands:

   ```
   wplot
   dbase, open, <filename>.dbs
   dbase, put, <id>, star
   ```

4. Use the post-processing Registers to store up to six data sets. By default, scalar data is always read into Register 4 and the x, y, and z components of vector data are read into Registers 1, 2, and 3 respectively. The commands shown below swap the wall temperature data currently stored in Register 4 to Register 6, and reads the wall pressure data into the now empty Register 4:

   ```
   operate, swap, 4, 6
   getw, p
   ```

5. Save the solution data to a coded format `.usr` file using the following commands:

   ```
   savu, <filename>.usr, all, coded, all
   ```

Importing the es-ice or pro-STAR CAE Model

The procedure for importing the CAE model from es-ice or pro-STAR unlike importing models from other CAE products.

To import a CAE model from es-ice or pro-STAR:

- Select File > Import > Import CAE model...



- The *Open* dialog appears. Select the `.dbs` file of a surface mesh (shells) from es-ice or pro-STAR.
- Selecting a file opens the *Import Options* dialog.



Note: There is no option to import a CAE model from es-ice or pro-STAR as a region. Also note that the *Units* must correspond to the length unit used to generate the model in es-ice or pro-STAR. For es-ice, is typically millimeters.

- The CAE model appears in the simulation tree as a sub-node of the Imported Models manager.



Importing the es-ice or pro-STAR Solution Data

To import solution data from es-ice or pro-STAR:

- Right-click the CAE model node and select Import CAE Solution Data.



- The *Open* dialog appears. Select the `.usr` file in a coded (ASCII) format.
- Selecting a file opens the *Import Data Options* dialog.



The options available in the panel are described below:

- *Data Type* – Currently, only cell data is available.
- *Registers* – Selects the Registers that contain the required post-processing data. If you choose to import a Register that does not contain any data, the import fails. The following options are available:
  - Scalar (Register 4) – Selects the scalar data that is stored in Register 4.
  - Vector (Registers 1-3) – Selects the vector data that is stored in Registers 1-3.

  - Both – Selects both the vector and scalar data that is stored in Registers 1-3 and 4 respectively.
  - All – Selects all Registers. This is recommended when importing data from es-ice.
  - Register *n* – Selects Register *n*.
- *Data Format* – Currently, only coded format files are permitted.
- *Cell Offset* – to the cell ID numbers in the solution data. Typically this option is not required, especially you have followed the guidelines for preparing data in es-ice or pro-STAR.
- *Heading* – Label used to identify the Register in STAR-CCM+. For example, using the default heading (Register) stores the imported data in a field function called *Imported Register n*, where *n* is the Register number.

Mapping data from es-ice or pro-STAR to STAR-CCM+

Note that data can only be mapped from es-ice or pro-STAR to STAR-CCM+, not the opposite. This means that only one-way coupling is possible.

To map imported solution data from es-ice or pro-STAR to STAR-CCM+:

- Right-click the CAE model node and select Map Data between STAR-CCM+ and Imported Models > Map Surface Data this opens the *Map Data* dialog
- In the *Map Direction* group box, select Imported Model to STAR-CCM+.
- In the *Imported Surface* tree, tick the surfaces of the imported model, from which data is to be taken.
- In the *Boundaries* tree, tick the equivalent boundaries to which the imported data is to be mapped.
- In some cases, you might need to select additional options in the *Nearest Neighbor Search Constraints* group box. Use these options if your initial attempt at mapping fails.

- Select the Register to be mapped from the *Data to map* drop-down menu. An example of the dialog is shown below:



- Click Apply to map the imported solution data to the STAR-CCM+ mesh.

This generates a field function with the same name as the imported data, prefixed by `Mapped` (e.g Mapped Imported Register 4). This field function can then be used in the same way as any other scalar field function.

Working with Ensight

You can use the export capabilities of STAR-CCM+ to compare multiple solutions in one simulation. This approach involves saving the results of each run in a separate Ensight file and is only possible for boundaries and parts, not volumes.

Note: If results are to be imported directly into Ensight, boundary and region names must not exceed 80 characters.

To compare results:

- Complete the first run.
- Export the results to an Ensight `.case` file. When your export includes solution data, you have the following options:
  - To export vertex data, activate the Export Results on Vertices checkbox in the *Export File* dialog. This option is only available when your export includes solution data.
  - To export cell data, leave the Export Results on Vertices checkbox deactivated.
- Clear the solution and make your modifications to the simulation. Run the simulation again.
- Import the data from the `.case` file. The mesh surfaces appear in the object tree.



The scalars that you exported appear in the simulation as field functions, with the name of your Ensight file preceding the scalar name. You can select them just like any other scalar values, as shown in the following screenshot.



You can compare the results side-by-side, selecting the simulation's scalars in one scene and the imported ones in the other.

Note that the minus operator ("`-`") can appear in the field functions generated by an import of an Ensight file. A temperature function imported from a file called `test.case` would have the name `test-Temperature`. In order for such a function to be referenced by another field function, the following syntax would have to be used:

${test-Temperature}

Working with the STAR-CCM+ Boundary Data (.sbd) Format

This section provides notes on the following topics related to using the `.sbd` file format in STAR-CCM+:

- What is the .sbd file format?
- Importing and mapping the .sbd format
- Working with field functions
- Examples of post-processing

What Is the .sbd File Format?

The STAR-CCM+ Boundary Data (`.sbd`) file format is an internal STAR-CCM+ format that allows you to exchange boundaries across different simulations. For boundaries that you choose, this binary file contains mesh information as well as selected scalar or vector solution data. The advantages of this format for exchanging boundaries include:

- It enables coupling of simulations with different physics.

- Each of the analyses can be performed separately and the coupling can take place at desired intervals.
- It is especially suitable when each of the analyses requires significant computer resources.
- You can isolate boundaries alone for coupling, rather than importing a whole region. The emphasis on boundaries is shown in the Export File dialog, which replaces the Regions tab with one for boundaries when you select the `.sbd` format in the *Files of Type* drop-down list.



An example of using this file format is a coupled conjugate heat transfer (CHT) analysis to optimize the design of an automobile. This optimization involves exchanging the heat transfer coefficient and reference temperature with vehicle thermal analysis and retrieving surface temperatures. Given the distinct nature of physics to be captured in the under-hood fluid analysis and in the vehicle thermal analysis, the models created are drastically different. Differences include the mesh densities and topologies,

the grouping of parts and boundary definitions, and the preservation of geometric features such as sharp corners and edges. However, the `.sbd` file makes this possible with the following procedure:

1. Choosing the boundaries representing the under-hood surface from the fluid analysis and selecting the local heat transfer coefficient and reference temperature for export
2. Exporting the chosen boundaries and the selected fields to an external `.sbd` file
3. Importing the `.sbd` file into the solid analysis, and mapping the imported surface data to the corresponding under-hood surface in the vehicle thermal model
4. Performing thermal analysis
5. Choosing the boundaries comprising the solid under-hood surface and selecting the surface temperature for export
6. Exporting the chosen boundaries and the selected fields to an external `.sbd` file
7. Importing the `.sbd` file into the fluid analysis and mapping data onto the surface

Importing and Mapping the .sbd Format

The sequence of operations is similar to those for coupling with external CAE packages such as Abaqus or RadTherm. The only difference is that the imported data is node-centric in case of the finite element packages and face-centric in case of the `.sbd` format. The pop-up menus of the imported surface objects indicate this, as they refer to faces rather than vertices:

- Import Face Data for Imported Surface



- Map Imported Surface Face Data to Boundary Faces



Working with Field Functions

Field functions work with the `.sbd` format as follows:

- Imported Field Functions -- When an `.sbd` file is imported, field functions are automatically created for all the imported data fields. These field functions have storage only on the imported surfaces.
- Mapped Field Functions -- When data is mapped between the imported surfaces and STAR-CCM+ boundaries, a new set of field functions are automatically generated. These field functions have storage only on the target surfaces. Apart from this mapped field, a field function named `Mapped Flag` is created for all mappings. This field has a value of 0, 1 or 2. This field function also has storage only on the target surface. On any face, a value of 0 indicates that mapping was successful, a value of 1 indicates that the mapping failed the proximity check and a value of 2 indicates failure of the normalcy check. The proximity and normalcy checks are additional mapping options.

Examples of Post-Processing

The following screenshots show how the imported quantities are displayed in STAR-CCM+:

- Imported local heat transfer coefficient



- Mapped imported local heat transfer coefficient



Working with the Visualization Toolkit (VTK) Format

The VTK format is part of a system of visualization and computer graphics programming known as the Visualization Toolkit. The `.vtk` format itself is designed to present different types of data consistently. Two styles of this file format exist: legacy (serial) and XML-based. STAR-CCM+ works with the first style, either binary or ASCII, to import surface data. VTK files are also used by the arbitrary probe derived part.

Creating a VTK File

It is possible to create a VTK file by exporting data from a derived part.

For creating a file manually to import into STAR-CCM+, the following example shows how to input 7 data points.

```
# vtk DataFile Version 2.0
Example Data for STAR-CCM+ Arbitrary Probe
ASCII
DATASET POLYDATA
POINTS 7 float
0.0 0.0 0.0
1.0 0.0 0.0
2.0 0.0 0.0
0.0 1.0 0.0
1.0 1.0 0.0
2.0 1.0 0.0
0.0 0.0 1.0
```

The sensor and/or probe locations *must* lie inside of the geometry in order for a successful import to take place.

Coupling with RCAS

RCAS (Rotorcraft Comprehensive Analysis System) is a Computational Structural Dynamics (CSD) tool for modeling rotor blades.

STAR-CCM+ allows you to import RCAS models for fluid-structure interaction analyses using file-based coupling. The RCAS model mesh and motion data are stored within the same file, with the standard CSD `*.in` format.

This section describes how to:

- import RCAS models and solution data into STAR-CCM+
- perform mapping operations from the STAR-CCM+ user interface
- export STAR-CCM+ data to RCAS

Importing the RCAS Beam Model

Rotor blades are modeled in RCAS as 6-DOF beams and can be imported into STAR-CCM+ as beam regions. For more information on beam elements in STAR-CCM+, see Imported Beam Elements.

To import the rotor blade mesh in a STAR-CCM+ simulation:

1. Select File > Import > Import CAE Model...
2. Select the RCAS `*.in` file that contains the model mesh and motion data.
3. In the *Import Options* dialog:
   a. Set Import Mode to Create new Region(s) and Imported Model(s).
   b. Specify the RCAS Import Options, that is, the rotor Angular Velocity, Origin, and Radius.

These data items are stored within the imported beam model node, for later use.



c. Click OK.

Importing Displacement and Rotation Data

The RCAS mesh and motion data are stored in the same `*.in` file. In STAR-CCM+, you can import the RCAS displacement and rotation data from the RCAS Imported Model node. The imported RCAS displacement and rotation data become available in STAR-CCM+ as the Imported Displacement and Imported Rotation field functions.

To import the RCAS motion data:

1. Right-click the imported RCAS model node and select Import CAE Solution Data...



2. In the *Open* dialog, select the RCAS model `*.in` file that contains the mesh and motion data.

3. In the *Import Data Options* dialog, specify the Angle Units and Import Time, then click OK.



You can view the imported displacement and rotation data in a scene by selecting the corresponding field functions.



Mapping Field Data

You can map pressure and shear stress data, defined on the faces of the wetted surface in STAR-CCM+, onto the nodal forces and moments for each node on the RCAS beam model. Similarly, you can map the nodal displacement and rotation data, defined on the RCAS beam model, onto the wetted surface in STAR-CCM+.

Instructions for performing mapping operations from/to beam models are given in the sections Mapping Boundary Data to Imported Beams and Mapping Imported Beam Data to STAR-CCM+ Boundaries of the STAR-CCM+ User Guide.

Exporting to RCAS

Data that has been mapped onto an imported RCAS model can be exported from STAR-CCM+ to external `*.out` files.

To export data from STAR-CCM+ to the RCAS model:

1. Import the RCAS beam model into STAR-CCM+. For more information, see Importing the RCAS Beam Model.
2. Map the data that you want to export to RCAS onto the imported RCAS model.
3. Export the mapped data to RCAS following the instructions outlined in the section, Exporting Mapped Data. Choose `.in` as file extension.

As the export of CFD loads is transient, the export file is appended for discrete values of the azimuth angle (rotor plane) for each revolution.

Co-Simulation with One-Dimensional Codes

It is possible to combine STAR-CCM+'s capabilities with analysis in a one-dimensional code. The set up procedure in STAR-CCM+ is specific for GT-POWER, WAVE, OLGA, RELAP, and Reacting Channel coupled simulations. This section includes the following:

- What are the 1D Coupling Models?
- Activating a 1D Coupling Model
- Using the 1D Coupling Solver

What are the 1D Coupling Models?

The 1D coupling models allow for the transfer of physics values and species data between a one-dimensional code and STAR-CCM+.

The 1D coupling models available for co-simulation with STAR-CCM+ are AMESim Coupling, GT-POWER Coupling, OLGA Coupling, OLGA EMP Coupling, RELAP Coupling, WAVE Coupling, and Reacting Channel Coupling.

One-dimensional codes are fast, and good for situations where the flow is broadly one-dimensional or zero-dimensional, or can be modeled as such (with correlations and so on). Three-dimensional codes are much slower, but can handle requirements such as new geometries and separated flows.

Coupling works best if the interfaces are located where the flow is largely one-dimensional, as this will reduce errors due to averaging. The worst place for a coupling zone would be where there is partial reverse flow, because the mean flow is not representative of either the forward or reverse flows.

Before beginning a new simulation with one-dimensional coupling, make the following preparations with the one-dimensional software:

- The one-dimensional software should be installed with the correct environment variables.
- For convenience, make sure the one-dimensional data files are in a single directory.

A 1D Coupling model is represented with a node among the physics continuum models. Its node is subject to common menu items, and appears when you activate the model.



Activating a 1D Coupling Model

The 1D Coupling models are selected from the Physics Model Selection dialog on the relevant continuum. The 1D Coupling models are available once the Coupled or Segregated flow models have been activated.

For information on how to activate a specific 1D Coupling model, see:

- Activating the GT-POWER Model
- Activating the OLGA Coupling Model
- Activating the OLGA EMP Coupling Model
- Activating the WAVE Coupling Model
- Activating the RELAP Coupling Model
- Activating the AMESim Coupling Model
- Activating the Reacting Channel Coupling Model

Activating a 1D Coupling model (other than the Reacting Channel Coupling model) will activate the 1D Coupling Solver under the Solvers manager.

Using the 1D Coupling Solver

The 1D coupling solver is represented by the 1D Coupling Solver node, which is subject to common menu items.



Solver Frozen

The solver frozen option allows the user to determine whether the 1D coupling solver will be active whilst obtaining a solution. This is generally not recommended, and should only be used as a debugging tool.

To set the solver as frozen, select the 1D coupling solver, located in the feature tree under the Solver manager.



Then tick the Solver Frozen checkbox in the *Properties* window.



GT-POWER Coupling

The GT-POWER Coupling feature handles two-way co-simulation between STAR-CCM+ and GT-POWER, the 1D engine simulation tool from Gamma Technologies. This coupling enables the creation of advanced simulations where GT-POWER handles the bulk of the flow in an engine powertrain, but where a portion of the flow, say in an intake manifold, is calculated in 3D within STAR-CCM+. The STAR-CCM+ simulation therefore has realistic time-accurate boundary conditions with which to calculate the flow within the component being modeled. This coupling also reduces the computational overhead predicting the flow throughout the rest of the system.

The following section provides the definition of common terms:

- GT-POWER Coupling
- GT-POWER Species
- GT-POWER Zones

The documentation also includes a guide to using GT-POWER Coupling, which details:

- Preparing the GT-POWER Case
- Setting Up the STAR-CCM+ Simulation
- Coupling STAR-CCM+ to GT-POWER

The material is specific to GT-POWER v7.3.

What Is GT-POWER Coupling?

GT-POWER coupling provides data exchange between three-dimensional STAR-CCM+ and one-dimensional GT-POWER. This coupling provides a compromise between the speed of one-dimensional code, with the detail of three-dimensional code.

Note: The feature provides optimum results when coupling occurs in a zone with one-dimensional flow.

When running the solver, the physics values and species mass fractions that GT-POWER sends are applied uniformly at the STAR-CCM+ boundary. If there is outflow at the boundary, STAR-CCM+ averages the values and mass fractions across the zone, and sends data to GT-POWER. Data exchange occurs at each time step or inner iteration depending on the simulation setup.

In the GT-POWER simulation, GT-POWER flow components represent the STAR-CCM+ regions. Similarly, GT-POWER flow connections represent STAR-CCM+ boundaries.



This method provides volume-averaged data, and data transition from three-dimensional to one-dimensional simulations.

The Co-Simulations > GT-POWER 1 node is available when you activate the GT-POWER model. It has subnodes that contain species and zones, when STAR-CCM+ is coupled to GT-POWER. Selecting this node provides access to the properties and expert properties.



GT-POWER Properties

Directory Name

Specifies the directory to the required GT-POWER case files. The path to the directory can be either relative or absolute. If the path is relative, two interpretations apply:

- if GT-SUITE is automatically started by STAR-CCM+, the path is relative to the directory containing the STAR-CCM+ simulation file.
- if GT-SUITE is started manually using the gtlink command, the path is relative to the directory containing the gtlink command. Manually starting GT-SUITE is required when coupling across different machines.

Case Name

Specifies the name of the GT-SUITE model input files that are used in the coupling simulation. These files must correspond to the major release version of GT-SUITE, as indicated in the GT-SUITE Version (v#.#) property.
When using GT-SUITE v7.3 or higher, specify the full name of the GT-ISE file including the extension `*.gtm`.
With older GT-SUITE versions, the Case Name property specifies the GT-SUITE solver input file ending with extension `.dat`, which could be indicated either by including the `.dat` extension or excluding the extension altogether.

GT-POWER Expert Properties

GT-SUITE Version (v#.#)

Specifies the GT-SUITE version number to which STAR-CCM+ couples. This property must be in the format, `v#.#`, where '`#`' are integers (for example, v7.3)

GT-SUITE Root

Specifies the root directory where STAR-CCM+ searches for the library link file, `libgtlink_dp.so`. If set to `default`, STAR-CCM+ uses the `GTIHOME` environment variable.

Coupling Data Precision	Adjusts the precision of the data transfer between STAR-CCM+ and GT-POWER. A value of `1` denotes single precision and a value of `2` denotes double precision.
Couple From Step	Specifies the point at which two-way coupling between the STAR-CCM+ and GT-POWER begins. Until this time-step, data is only transferred from GT-POWER to STAR-CCM+. Typically, you do not use this property, unless the solver initially diverges.

What Is a GT-POWER Species?

A GT-POWER species defines components of the fluid mixture in the simulation. The physics values and mass fractions of these species are exchanged at zones, between STAR-CCM+ and GT-POWER.

Within STAR-CCM+, you assign all species from GT-POWER to a mixture component or passive scalar. You can assign multiple species to individual gas components and passive scalars, which reduces process times, but compromises detail and accuracy.

The GT-POWER Species node is available when you activate and start GT-POWER Coupling.



Unburned Species

GT-POWER defines unburned species with the *FluidGas*, *FluidLiqIncompress*, *FluidLiqCompressible*, or *FluidNASA-LiqGas* reference objects, as illustrated below:



Burned Species

The GT-POWER combustion models determine the burned species from the products of reaction between the unburned species. In GT-POWER v7.3, the combustion products consist of 13 components with varying composition.

Species Name	Description
prod_ar	Argon
prod_co	Carbon Monoxide
prod_co2	Carbon Dioxide
prod_h	Monatomic Hydrogen
prod_h2	Hydrogen

prod_h2o	Water
prod_n	Monatomic Nitrogen
prod_n2	Nitrogen
prod_no	Nitric Oxide
prod_o	Monatomic Oxygen
prod_o2	Oxygen
prod_oh	Hydroxide
prod_so2	Sulphur Dioxide

You can define how GT-POWER coupling tracks and exchanges the burned species between GT-POWER and STAR-CCM+ using CFD Species Maps.

GT-POWER Species Properties

GT Species Name	The species name that is specified within the GT-POWER case (read-only).
Mixture Component / Passive Scalar	Specifies the gas component or passive scalar which represents the species in the STAR-CCM+ simulation.

What Is a GT-POWER Zone?

A GT-POWER zone is a specified section of the CFD region, which provides data averaging and exchange between STAR-CCM+ and GT-POWER. You define individual zones in GT-POWER, and equivalent boundaries in the STAR-CCM+ simulation.

STAR-CCM+ volume-averages the physics values and species mass fractions in the zones, which provides a single value to determine the equivalent boundary conditions in GT-POWER. You assign each GT-POWER zone must to a boundary in the STAR-CCM+ simulation.

The GT-POWER simulation defines zone length. To provide the optimum transition between STAR-CCM+ and GT-POWER, make the zone represent a section where there is one-dimensional flow. You can visualize the

GT-POWER zones as a scalar displayer using the AvgZoneFlag field function. The AvgZoneFlag field function values have no significance and simply distinguish the individual zones



The GT-POWER Zones node is available when you activate and start GT-POWER Coupling.



GT-POWER Zone Properties

GT Zone Name	The zone name that is specified within the GT-POWER case (read-only).
GT Zone Type	The CFDConnection boundary type that is specified within the GT-POWER case, that is inlet or pressure (read only).
GT Zone Id	The ID number that is given to the zone within the GT-POWER case (read only).
GT Zone Len	The CFDConnection zone length that is specified within the GT-POWER case (read only).
Boundary	Specifies the boundary at which coupling occurs in the STAR-CCM+ simulation. Only boundaries that are specified as coupled are available for selection.

Preparing the GT-POWER Case

You define settings in the GT-POWER case before coupling, which determine certain properties and characteristics within the STAR-CCM+ simulation:

- GT-POWER Boundary Types
- Setting the Zone Diameter
- Setting the Zone Length
- Setting the Number of Pre-Cycles
- Defining the CFD Species Map

GT-POWER Boundary Types

The boundary type determines the physics values exchanged between GT-POWER and STAR-CCM+. There are two available boundary types in a GT-POWER simulation, which behave in different ways when coupled to STAR-CCM+.

- GT-POWER inlet boundary, which sends the following data to STAR-CCM+:
  - Density
  - Mass Flow
  - Normal Velocity
  - Temperature
- GT-POWER pressure boundary, which sends the following data to STAR-CCM+:
  - Mass Flow
  - Pressure
  - Temperature

Turbulence values on the coupled boundaries are not exchanged between GT-POWER and STAR-CCM+. Turbulence values are specified in STAR-CCM+ at the boundary level in the usual manner, through the Physics Conditions > Turbulence Specification and the Physics Values nodes.

Setting a GT-POWER Inlet Boundary

A GT-POWER inlet boundary only couples to a STAR-CCM+ mass flow or velocity inlet.

To set an inlet boundary type in GT-POWER:

1. Right-click a CFD connection, and select Edit Properties...



2. In the *Edit Object* dialog, set *Boundary Type* to inlet.



Setting a GT-POWER Pressure Boundary

A GT-POWER pressure boundary only couples to a STAR-CCM+ pressure outlet.

To set a pressure boundary type in GT-POWER:

1. Right-click a CFD connection, and select Edit Properties...



2. In the *Edit Object* dialog, set *Boundary Type* to pressure.



Setting the Zone Diameter

GT-POWER coupling requires the diameter of the zone in STAR-CCM+ to be equal to the GT-POWER CFD connection.

If the zone cross-section is non-circular, then use the effective diameter. The effective diameter being the diameter of an equivalent circle with the same cross-sectional area as the zone.

Using the example below, the diameter at the manifold intake is 50 mm.



To set the CFD connection diameter in GT-POWER:

1. Right-click a CFD connection, and select Edit Properties...



2. In the *Edit Object* dialog, set CFD Domain Diameter to the zone diameter in STAR-CCM+.



Setting the Zone Length

The zone length defines the extent of the volume where STAR-CCM+ averages values and sends them to GT-POWER. Ensure that the zone length covers an area of the simulation where there is one-dimensional flow.

To set the zone length in GT-POWER:

1. Right-click a CFD connection, and select Edit Properties...



2. In the *Edit Object* dialog, set CFD Domain DX to the zone length.



When STAR-CCM+ couples to GT-POWER, the zone length is rounded up or down to cover the nearest complete mesh cell. You can visualize the zones using the AvgZoneFlag field function as a scalar displayer.



Setting the Number of Pre-Cycles

Pre-cycles improve stability in GT-POWER, before transferring data to STAR-CCM+. This facility provides more accurate boundary conditions from GT-POWER, which reduces the time to reach a converged solution.

To use pre-cycles in a coupled simulation:

1. Right-click a CFD component, and select Edit Properties...



2. Set Number of Cycles Before CFD Started to the number of pre-cycles.



When obtaining a solution, the GT-POWER solver runs independently, until the number of pre-cycles, at which point data is sent to STAR-CCM+.

Defining a CFD Species Map

CFD species maps determine how the gas components are transferred between STAR-CCM+ and GT-POWER.

To define a CFD species map:

1. Right-click a CFD component, and select Edit Properties...



2. In the *Edit Object* dialog, click the ellipsis (...) next to the Species Mapping Object attribute.
3. In the *Value Selector* dialog, select CFDSpeciesMap from the *Create New Object* group box and click OK.
4. In the *Create Object* dialog, set Object to the name of the species map.
5. Set User Defined Species Map to one of the following options:

- `def` — Uniquely associates all GT-POWER species with a STAR-CCM+ gas component or passive scalar.
- CFDDefinedSpeciesMap — The user-defined CFDDefinedSpeciesMap reference object defines the association of one or more GT-POWER species with a STAR-CCM+ gas component or passive scalar. For example, you can map all harmful gases to a STAR-CCM+ passive scalar called Pollutants.

6. Set each CFD species to one of the following options:
   - `ign` — Ignores the CFD species and does not send it to STAR-CCM+. The mass fraction of this species is redistributed between the other species to maintain mass conservation. This option is only recommended for trace components.
   - CFDSpecies — Associates the GT-POWER species with the user-defined CFDSpecies in STAR-CCM+.



Defining a CFDDefinedSpeciesMap

A CFDDefinedSpeciesMap is similar to a CFDSpeciesMap, but you can map multiple GT-POWER species to a singe STAR-CCM+ species. You can also define the mapping for the intake charge species.

The following steps assume that you are defining a CFD species map, as detailed in the previous section.

To define a CFDDefinedSpeciesMap:

1. In the *Edit Object* dialog of the CFDSpeciesMap, click the ellipsis (...) next to the User Defined Species Map attribute.
2. In the *Value Selector* dialog, select CFDDefinedSpeciesMap from the *Create New Object* group box and click OK.
3. In the *Create Object* dialog, set Object to the name of the CFDDefinedSpeciesMap.
4. Set each entry in the GT-SUITE Species Name column to the GT-POWER species that you want to include in the map.
5. Set each entry in the CFD Species Name column to one of the following options:
   - `ign` — Ignores the CFD species and does not send it to STAR-CCM+. The mass fraction of this species is redistributed between the other species to maintain mass conservation. This option is only recommended for trace components or species that are not used in the GT-POWER case that you are analyzing.

- CFDSpecies — Associates the GT-POWER species with the user-defined CFDSpecies in STAR-CCM+.



Defining a CFDSpecies

A CFDSpecies is the species that GT-POWER passed to STAR-CCM+ during a coupled simulation. GT-POWER writes the material properties for these species when you run the simulation preprocessing.

The following steps assume that you are defining a CFD species map, as detailed in the previous section.

To define a CFDSpecies:

1. In the *Edit Object* dialog of the CFDSpeciesMap, click the ellipsis next to a CFD species.

2. In the *Value Selector*, select the relevant species from the *GT-SUITE Library* tab and click OK.



Setting Up the STAR-CCM+ Simulation

This section details the process of setting up a STAR-CCM+ simulation for coupling with GT-POWER:

- Activating the GT-POWER Coupling Model
- Setting the GT-POWER Species Representation
- Setting the Reference Pressure
- Setting the Initial Conditions
- Setting the STAR-CCM+ Boundary

Activating the GT-POWER Coupling Model

The GT-POWER Coupling model is available from the Physics Model Selection dialog on the relevant continuum.

To select the GT-POWER Coupling model, the following models are required:

1. Select Three Dimensional from the *Space* group box.
   - The Gradients model is automatically selected.
2. Select Multi-Component Gas from the *Material* group.
3. Select Non-Reacting from the *Reaction Regime* group.
4. Select either Coupled or Segregated from the *Flow* group, depending on the required solver type. The Segregated solver is recommended in most cases.
5. Select Implicit Unsteady from the *Time* group.
6. Select Co-Simulation from the *Optional Models* group.

7. Select GT-POWER Coupling from the *Co-Simulation* group.



Complete the physics model selection by activating the following models:

- Select Ideal Gas from the *Equation of State* group box.
- Select one of the Segregated Fluid Energy Models from the *Energy* group box.
- Select Turbulent from the *Viscous Regime* group box.
- If you intend to specify turbulence values in STAR-CCM+, select any turbulence model.
- Select Passive Scalar from the *Optional Models* group, which can be neglected if you represent all species with gas components.

You can select further models from the *Optional Models* group to meet individual simulation requirements.

GT-POWER Species Representation

A GT-POWER species defines components of the fluid mixture in the coupled simulation. You assign all GT-POWER species to either a mixture component or passive scalar, which defines the behavior as active or passive respectively.

The choice between using a mixture component or passive scalar to represent a GT-POWER species is dependent on user requirements. A mixture component influences the physical properties of the simulation, thus providing the most accurate and realistic solution. A passive scalar has no effect on the physical properties of the simulation, which reduces processing times. If you use passive scalars for all species, specify a background fluid in the region.

Note: Although the "fuel-combust" species is a liquid, you can represent it with a gas in the Multi-Component Gas model. This modeling approach does not affect results, as the gas only represents the species in STAR-CCM+, with the physics values received from GT-POWER.

To avoid inconsistencies when transferring species mass fractions to GT-POWER, assign each GT-POWER species to a unique mixture component or passive scalar in STAR-CCM+. If you assign two GT-POWER species to a single gas component, the solver stops with an error message before the first iteration. This error check was not in STAR-CCM+ version 6.04 and earlier. Therefore, change the setup in older simulation files to fix this error.

Importing the GT-POWER Material Properties

Simulation preprocessing in GT-POWER generates a material database file, `<case_name>.chemkin_GT.dbs`. To use the same material properties as the GT-POWER model, import this file into STAR-CCM+

To import the GT-POWER material properties:

1. Right-click the Tools > Material Databases node and select New Material Database.
2. Right-click the new material database and select Import Properties.
3. In the *Open* dialog, select the `<case_name>.chemkin_GT.dbs` file and click OK.
4. In the *Import Properties* dialog, select Import properties for all materials and click OK.



Defining Gas Components for GT-POWER Species

Gas components represent GT-POWER species that influence the bulk flow. That is, you associate gas components with GT-POWER species that have non-trace quantities.

To add a mixture component to the simulation:

1. Right-click the Gas Components node and select Select Mixture Components.



2. In the object selection dialog, select gas components from the tree.



3. Click Apply to add the gases to the simulation.

Defining Passive Scalars for GT-POWER Species

Passive scalars represent GT-POWER species that do not influence the bulk flow. That is, you associate passive with GT-POWER species that have trace quantities.

To add a passive scalar to the simulation:

- Right-click the Passive Scalars manager, and select New.



Setting the Reference Pressure

The reference pressure reduces the error due to numerical round-off. All other pressures in a STAR-CCM+ simulation are defined relative to the reference pressure.

However, in a GT-POWER simulation, all pressures are absolute. Take these differences into consideration when defining pressure values between the two codes.

To reduce the numerical round-off, set the reference pressure to the average pressure in GT-POWER. You can identify the average pressure in the region by running the GT-POWER simulation and using GT-POST to analyze the results.

To set the reference pressure:

- Select the Reference Pressure node and set the Value property.



Initial Conditions

This section details the process for setting the initial conditions for a GT-POWER coupled simulation. The following examples use constant values—while field functions and tabular data are also available.

The initial conditions in a STAR-CCM+ simulation improve the simulation convergence. For best initialization, set the STAR-CCM+ initial conditions to the average values from GT-POWER, which you can obtain using the following steps:

- Run the GT-POWER simulation until converged.
- Analyze the results in GT-POST.

This section covers the following initial conditions:

- Initial Pressure
- Initial Temperature
- Mass Fractions

Setting Initial Pressure

The initial pressure in STAR-CCM+ is relative to the reference pressure. Therefore, set the value to the difference between the average pressure from GT-POWER and the STAR-CCM+ reference pressure.

To set the initial pressure:

- Select the Pressure > Constant node and set the Value property.



Setting Initial Temperature

The initial temperature in STAR-CCM+ is an absolute value. Therefore, set the value to be the same as the average temperature from GT-POWER.

To set the initial temperature:

- Select the Constant node within the Static Temperature manager and set the Value property.



Setting Initial Mass Fractions

The species mass fraction varies depending on simulation requirements. As an example, if you are coupling to a component on the intake side of an engine, it is intuitive to initialize the region with air. Hence, the mass fractions are:

Mixture Component	Value
fuel-comb	0.0
fuel-vap	0.0
n2-vap	0.767
o2-vap	0.233

To set the initial mass fractions:

- Select the Species Mass Fraction > Constant node and set the Value property.



- Alternatively, click the ... (ellipsis) button to enter the mass fractions in an array.



STAR-CCM+ Boundary Conditions

You can couple a GT-POWER inlet boundary to a mass flow inlet or velocity inlet boundary in STAR-CCM+.

When a STAR-CCM+ boundary is coupled to GT-POWER, it is no longer a boundary in the traditional sense, but an interface between the two codes. For this reason, you do not generally set specific quantities for boundaries that are coupled to GT-POWER. An exception to this rule is for turbulence quantities, which must be set manually for the STAR-CCM+ boundary.

The recommended practice is to set the boundary type to mass flow inlet for all coupled boundaries—regardless of whether they are inlets or outlets. If you use a mass flow inlet, STAR-CCM+ uses the boundary area, normal velocity, and density from GT-POWER to determine the mass flow rate at the boundary. Small changes in density are reflected in the mass flow rate across the boundary.

For coupled boundaries that have large negative flows (such as flows that are leaving the STAR-CCM+ domain), use coupled velocity inlet boundaries. Using coupled mass flow inlet boundaries can sometimes lead to instabilities for some models. If you use a velocity inlet, STAR-CCM+ takes the normal velocity from GT-POWER and applies it uniformly at the coupled boundaries.

Note: Setting all coupling interfaces to mass flow or velocity inlets does not have the same negative numerical repercussions that would occur in a standalone CFD model. The absolute pressure of the CFD model drives the GT-SUITE solution and sets the mass flux or velocity at the boundary.

Mass Flow Inlet Data Exchange

You couple a STAR-CCM+ mass flow inlet to a GT-POWER inlet boundary. The following table summarize the data that STAR-CCM+ sends and receives at mass flow inlets.

Sent to GT-POWER	Received from GT-POWER
If flow is into STAR-CCM+:	
Pressure	Density
	Normal Velocity
	Temperature
	Mass Fractions
If flow is into GT-POWER:	
Pressure	Density
Temperature	Normal Velocity

Mass Fractions

Passive Scalars

Pressure Outlet Data Exchange

You couple a STAR-CCM+ pressure outlet to a GT-POWER pressure boundary. The following table summarize the data that STAR-CCM+ sends and receives at pressure outlets.

Sent to GT-POWER	Received from GT-POWER
If flow is into STAR-CCM+:	
Normal Velocity	Pressure
	Temperature
	Mass Fractions
If flow is into GT-POWER:	
Normal Velocity	Pressure
Temperature	
Mass Fractions	
Passive Scalars	

Setting the Boundary Conditions

The boundaries that couple to GT-POWER require a coupled boundary condition. These coupled boundaries require no other conditions as GT-POWER supplies all the flow information.

To set the boundary conditions:

1. Select the coupled boundary.



2. Set the Type property to Mass Flow Inlet, Velocity Inlet, or Pressure Outlet.



3. Select the Physics Conditions > 1D Coupling Specification node and set 1D Coupling Specification to Coupled.



Coupling STAR-CCM+ to GT-POWER

The following section details the process for coupling to GT-POWER.

- Defining the GT-POWER Case File
- Starting the GT-POWER Solver
- Assigning the GT-POWER Species
- Assigning the GT-POWER Zones
- Setting the Couple From Step Option
- Coupling to Alternative Versions of GT-POWER
- Coupling to GT-POWER on a Remote Host
- Running a GT-POWER Coupled Simulation on Parallel Hosts
- Setting up STAR-CCM+ and GT-POWER for a Restart Run

With the above points addressed as necessary, a solution can be obtained.

Defining the GT-POWER Case File

Specifying a GT-POWER case file determines the case that GT-POWER runs during the coupled simulation.

To define the GT-POWER case file:

- Select the Co-Simulations > GT-POWER 1 node and set the *Directory Name* and *Case Name* properties.



You can enter the directory name in an absolute or relative form. STAR-CCM+ does not change an absolute directory, but interprets a relative directory in one of two ways:

- If you launch STAR-CCM+ from a directory and supply the simulation file on the command line, the path is relative to the **starting** directory.
- If you launch STAR-CCM+ and load the simulation file with the file browser, the path is relative to the **simulation file** directory.

Starting the GT-POWER Solver

Once you have defined the case file, you can couple the STAR-CCM+ simulation to GT-POWER. When coupled, STAR-CCM+ reads the species and zones from the GT-POWER case file. You can then assign the species to the previously defined gas components and passive scalars. You can also assign the zones to boundaries, with the correct boundary type and conditions.

The following example couples STAR-CCM+ to GT-POWER when both codes are on the local host. For example, when running both the three-dimensional and one-dimensional analysis on a single machine or on a cluster.

To run the codes on independent machines, there are some additional steps in the Coupling to GT-POWER on a Remote Host section. You require these additional steps when running STAR-CCM+ on a cluster, but GT-POWER is on an independent machine with a separate installation.

To start GT-POWER coupling:

- Right-click the GT-POWER 1 node and select Start GT-POWER.



When coupled, the *Output* window displays

```
GT-POWER Initialized ...
```

You only start GT-POWER in this manner the first time you couple to GT-POWER. Thereafter, clicking Run (Ctrl+R) starts GT-POWER coupling and begins obtaining a solution.

You can stop the GT-POWER solver by clicking Solution > Clear Solution from the menu bar, which breaks the connection between STAR-CCM+ and GT-POWER. You then start GT-POWER using the previous method, which resets the coupling and updates changes to the GT-POWER case. This facility only applies if the STAR-CCM+ session remains active.

If you close STAR-CCM+, perform a restart-run by modifying the GT-POWER case file to restart from a user-defined time step and restarting the STAR-CCM+ simulation from an autosave file.

Assigning the GT-POWER Species

You assign all species in GT-POWER coupling to previously defined gas components or passive scalars. This assignment provides a means of representing the gas components in the STAR-CCM+ simulation. The GT-POWER species are only available for assignment when you define a STAR-CCM+ boundary as coupled to GT-POWER.

To avoid inconsistencies when transferring species concentrations to GT-POWER, assign each GT-POWER species to a unique mixture component or passive scalar in STAR-CCM+. If you inadvertently assign two GT-POWER species to a single gas component, the solver stops with an error message before the first iteration. This error check was not in STAR-CCM+ version 6.04 and earlier. Therefore, to fix this error, change the setup in older simulation files before continuing the run.

To assign a GT-POWER species:

- Select a GT-POWER Species and set the *Mixture Component / Passive Scalar* property to a mixture component.



Assigning the GT-POWER Zones

You assign all zones in GT-POWER coupling to a STAR-CCM+ boundary, with the correct boundary type and conditions. This assignment identifies the location of data exchange between the one-dimensional and three-dimensional codes.

To assign a GT-POWER zone:

- Select a GT-POWER Zone and set the Boundary property to a STAR-CCM+ boundary.



Setting the Couple from Step Option

The Couple From Step option delays the data transfer to GT-POWER until a defined time step. This facility allows the STAR-CCM+ solver to reach a level of stability while receiving data from GT-POWER, before sending data to GT-POWER. Although STAR-CCM+ does not transfer data to GT-POWER, the codes both advance at the same time. During this period of one-way transfer, the GT-POWER analysis runs with constant boundary conditions at the coupled boundary.

Note: Using the Couple From Step option is only required if the solver diverges during the initial time-steps.

To set the Couple From Step option:

- Select the Co-Simulations > GT-POWER 1 node and set the Couple From Step property.



Coupling to Alternative Versions of GT-POWER

By default, STAR-CCM+ couples to GT-POWER v7.3, double precision, within the GTIHOME environment variable directory. However, you can modify the GT-POWER node properties so that STAR-CCM+ couples to an alternative version or installation of GT-POWER. This section details the following three properties:

- GT-SUITE Version (v#.#)
- GT-SUITE Root
- Coupling Data Precision

GT-SUITE Version (v#.#)

The GT-SUITE Version (v#.#) property specifies the version of GT-POWER to which you are coupling.

The default GT-POWER version depends on the STAR-CCM+version:

- If you create a simulation file in STAR-CCM+ v7.02 or newer, the default GT-POWER version is 7.2.
- If you create a simulation file in STAR-CCM+ v8.04 or newer, the default GT-POWER version is 7.3.

- If you resume a simulation file from STAR-CCM+ v8.04 or newer, the GT-POWER version comes from the simulation file.

To couple to an alternative GT-POWER version number:

- Select the Co-Simulations > GT-POWER 1 node and set GT-SUITE Version (v#.#) to the required version number.
  - The version number must be in the form v#.#, where the # is a number (for example, v6.2 or v7.3).



GT-SUITE Root

The GT-SUITE Root property specifies the directory in which STAR-CCM+ loads the GT-SUITE coupling library file, `libgtlink*.so` (or `gtlink*.dll` on Windows). This file is specific to the STAR-CCM+ platform as well as the major release version of GT-SUITE. A GT-SUITE major release version is indicated by the first two numbers in the GT-SUITE software designation (v#.#).

The machine on which STAR-CCM+ runs does not need to have GT-POWER installed. The library from the GT-POWER installation, however, is required in order to initiate the communication between GT-POWER and STAR-CCM+.

If GT-POWER is installed on the same machine as STAR-CCM+, do not modify the GT-SUITE Root property. If GT-SUITE Root is not set, or set to the default value, STAR-CCM+ tries to load the library from:

```
$GTIHOME/[GT-SUITE_version].0/GTsuite/bin/[platform_archi
```

```
tecture]/coupling/
```

where:

- `[GT-SUITE_version]` is the GT-SUITE version number, for example: `v7.3`.
- `[platform_architecture]` is the platform architecture, for example: `linux_x86_64`.

If you installed GT-POWER in a non-default directory, modify GT-SUITE Root and set it to your GT-SUITE root directory. STAR-CCM+ tries to load the library file from the same path as above, but with `$GTIHOME` substituted with your GT-SUITE root directory:

```
[GT_SUITE_ROOT]/[GT-SUITE_version].0/GTsuite/bin/[platfor
m_architecture]/coupling/
```

If no library is found in the above path, STAR-CCM+ tries to load the library from the GT-SUITE root directory:

```
[GT_SUITE_ROOT]/
```

Setting the GT-SUITE Root Directory

To set the GT-SUITE Root directory:

- Select the Co-Simulations > GT-POWER 1 node and set GT-SUITE Root to the root directory.



Coupling Data Precision

The Coupling Data Precision property specifies the precision of the data transfer between STAR-CCM+ and GT-POWER. A value of 1 specifies single precision and a value of 2 specifies double precision.

Double precision coupling data is only available with GT-POWER v7.1 onwards. Therefore, to maintain compatibility with older versions, switch to single precision.

To change the coupling data precision:

- Select the Co-Simulations > GT-POWER 1 node and set the Coupling Data Precision property.



Coupling to GT-POWER on a Remote Host

By default, GT-POWER coupling assumes that the one-dimensional and three-dimensional codes are on the same machine, with an automatically assigned port. This approach is typical when running both STAR-CCM+ and GT-POWER on the same machine or cluster; providing you have installed both programs on the local host.

However, you can manually modify this procedure to allow for STAR-CCM+ to link to GT-POWER on a remote host, with a user-defined port. This modification provides greater flexibility and can allocate specific system resources to the GT-POWER solver.

GT-POWER coupling can also run remotely across different operating systems.

To couple to a GT-POWER simulation on a different machine:

1. On the STAR-CCM+ machine, create a file that is called `gtlink.prm` in the same directory as the STAR-CCM+ simulation file, which contains the following:

   ```
   HOST=<host name>
   PORT=<port number>
   ```

2. Launch STAR-CCM+ from the same directory as the `gtlink.prm` file and STAR-CCM+ simulation file.
   - On UNIX/Linux, in a terminal window:

     ```
     cd path/to/directory/
     starccm+ <filename.sim>
     ```

   - On Windows, in the command prompt:

     ```
     cd path\to\directory\
     starccm+ <filename.sim>
     ```

3. For the GT-POWER 1 node properties, leave the Directory Name property **blank**, and set the Case Name property to the absolute path of the GT-POWER case file. For example, with a case file called `4cyl-cfd`:
   - On UNIX/Linux systems:

     ```
     /home/coupling/GT-POWER/4cyl-cfd
     ```

   - On Windows systems:

     ```
     C:\coupling\GT-POWER\4cyl-cfd
     ```

4. On the GT-POWER machine, start the GT-POWER server in the correct directory using the following commands:
   - On UNIX/Linux, in a terminal window:

     ```
     cd path/to/directory/
     $GTIHOME/bin/gtlink -s:<port number>
     ```

   - On Windows, in the command prompt:

     ```
     cd path\to\directory\
     gtlink -s:<port number>
     ```

   The gtlink command initiates the GT-POWER server on the independent machine. The -s: argument defines the port number, which must match the port number in the gtlink.prm file.

5. Copy the relevant gtlink library from the installation on the GT-POWER machine to the STAR-CCM+ machine.

- The following table summarizes the required libraries appropriate for the STAR-CCM+ machine architecture.

Platform	Precision	Library File Location
Linux 32-bit	Single	`$GTIHOME/v7.3.0/GTsuite/bin/linux_x86/coupling/libgtlink.so`
	Double	`$GTIHOME/v7.3.0/GTsuite/bin/linux_x86/coupling/libgtlink_dp.so`
Linux 64-bit	Single	`$GTIHOME/v7.3.0/GTsuite/bin/linux_x86_64/coupling/libgtlink.so`
	Double	`$GTIHOME/v7.3.0/GTsuite/bin/linux_x86_64/coupling/libgtlink_dp.so`
Windows 32-bit	Single	`%GTIHOME%/v7.3.0/GTsuite/bin/win32/coupling/gtlink.dll`
	Double	`%GTIHOME%/v7.3.0/GTsuite/bin/win32/coupling/gtlink_dp.dll`
Windows 64-bit	Single	`%GTIHOME%/v7.3.0/GTsuite/bin/win64/coupling/gtlink.dll`
	Double	`%GTIHOME%/v7.3.0/GTsuite/bin/win64/coupling/gtlink_dp.dll`

6. In the GT-POWER 1 *Properties* window, set the GT-SUITE Root property to the location where you copied the gtlink library file in the previous step.
   - This location can be a relative or absolute path.
7. Start GT-POWER coupling.

When coupling to GT-POWER, STAR-CCM+ finds the gtlink.prm file and uses the contents to identify the GT-POWER server host machine. The GT-POWER server (which must be started before attempting to connect from STAR-CCM+) provides access to the GT-POWER solver. When you start GT-POWER coupling, STAR-CCM+ connects to this GT-POWER server, and you can obtain a solution exactly as if GT-POWER was available locally.

Running a GT-POWER Coupled Simulation on Parallel Hosts

If you intend to run on parallel hosts (a multiple node cluster or several machines over a network), the procedure for running a GT-POWER coupled simulation must meet certain criteria.

The following examples demonstrate coupling over parallel hosts using two machines or nodes; A is the master host and B is the slave host:

- Running GT-POWER on the Master Host
- Running GT-POWER on a Slave Host

Running GT-POWER on the Master Host

The easiest way to run a coupled simulation is to allow STAR-CCM+ to launch the GT-POWER server automatically on the master host. However, the master host must have access to the GT-SUITE installation.

To run the GT-POWER server on the master host, A:

- Launch a parallel simulation using the procedure in Using a Remote Multi-Core Workstation, ensuring that the master host is the first entry in the table. Alternatively, from the command line:

  ```
  starccm+ -server -on A:x,B:y -np x+y
  ```

  where:

  - `x` – is the number of processes on A
  - `y` – is the number of processes on B

Thereafter, you can either allow STAR-CCM+ to start the GT-POWER server automatically, or you can start the server manually using the remote hosts procedure—both approaches are valid.

Running GT-POWER on a Slave Host

If the master host does not have access to a GT-SUITE installation, run the GT-POWER server on one of the slave hosts. This scenario requires a different coupling procedure.

To run the GT-POWER server on the slave host, B:

1. Launch a parallel simulation using the procedure in Using a Remote Multi-Core Workstation, ensuring that the master host is the first entry in the table. Alternatively, from the command line:

   ```
   starccm+ -server -on A:x,B:y -np x+y
   ```

   where:

   - `x` – is the number of processes on A
   - `y` – is the number of processes on B

2. In the *GT-POWER 1 - Properties* window, set the GT-SUITE Root property to the absolute path for the GT-SUITE installation on the slave host, for example:



3. Start the GT-POWER server manually using the remote hosts procedure.

Setting Up STAR-CCM+ and GT-POWER for a Restart Run

On occasion, it is necessary to stop the solver prematurely, or continue running an existing analysis for more time steps. To meet these requirements, GT-POWER coupling provides facilities to perform a restart run. This section details this procedure for a case that has been stopped before completion.

To continue an analysis for more time steps, the following steps can also be followed, but increase the STAR-CCM+ stopping criteria to extend the duration of the simulation.

Before running the solver, enable the Autosave feature in STAR-CCM+ to save the simulation at regular intervals. The autosave files provide an exact time step from which you can restart the solver, thus ensuring the two codes begin at the same point in the analysis. The autosave file has a name that corresponds to the current time step, for example, `casename@01000.sim` is the saved simulation at time step 1000.

To enable autosave:

1. From the menu bar, select File > Auto Save.



2. In the *Auto Save* dialog, set Trigger Type to Time Step, as the restart time step must be specified in the GT-POWER case.



3. Set Frequency and Max Autosaved Files as required.
4. Start the GT-POWER solver and obtain a solution in the usual manner.

To restart the simulation in GT-POWER, modify the case setup so that the solver restarts from a user-specified time step.

1. Right-click a CFD component, and select Edit Properties...



2. In the *Edit Object* dialog, select the *Advanced* tab.
3. Set Read to CFD Time Step Number to the required time step.
   - The time step must match the STAR-CCM+ autosave file. The example below uses a `casename@00500.sim` file, so the required time step is 500.



4. Save the GT-POWER case file and create a `.dat` file, by selecting Run > Run-Simulation Utilities > Create .dat file from the menu bar.

To restart the simulation in STAR-CCM+:

1. Loaded the autosave file into STAR-CCM+
2. Run the solvers in the usual manner.

The STAR-CCM+ solver restarts from the last time step and the GT-POWER solver restarts from the user-specified time step.

OLGA Coupling

The OLGA Coupling feature provides the ability for two-way transfer of data between three-dimensional STAR-CCM+, and one-dimensional OLGA. For example, the effect on three-phase flow behavior due to a surge of oil can be analyzed in detail in STAR-CCM+, with OLGA providing the input data.

This simulation method provides increased detail in specified areas of the OLGA case, by adding a three-dimensional element to a one-dimensional simulation. Similarly, processing times can be reduced in a STAR-CCM+ simulation, by utilizing the speed of one-dimensional codes.

Currently, STAR-CCM+ supports coupling with OLGA 5.3.2 and OLGA 7.2.2. OLGA runs on Windows platforms only.

Note: When coupling to OLGA 5, STAR-CCM+ can run on any platform and connect to an OLGA session running on Windows. When coupling to OLGA 7, STAR-CCM+ must run on a Linux platform and connect to an OLGA session running on Windows. Additionally, OLGA 7 requires APIS, a third party tool that allows STAR-CCM+ to communicate with OLGA.

This section provides the definition of the following concepts:

- OLGA Coupling
- OLGA Species
- OLGA Zones

This section also includes a guide to using OLGA Coupling, which details:

- Setting Up the STAR-CCM+ Simulation for coupling with OLGA
- Coupling STAR-CCM+ to OLGA 5
- Coupling STAR-CCM+ to OLGA 7

What is OLGA Coupling?

OLGA coupling provides a means of data exchange between three-dimensional STAR-CCM+ and one-dimensional OLGA. This simulation method allows for a compromise between the speed of one-dimensional code, with the detail of three-dimensional code.

Coupling provides optimum results when data transfer occurs at a zone where the flow is expected to be mostly one dimensional.

You can use either the Multiphase Segregated Flow (EMP) or Volume of Fluid (VOF) models to simulate the multiphase fluid in STAR-CCM+. The OLGA EMP Coupling and the OLGA Coupling (for VOF) models handle the transfer of phase volume fractions from OLGA to their respective multiphase fluid models. There is no difference in setting up the simulation within OLGA for coupling to EMP or VOF.

When running the solver, the physics values and species volume fractions received from OLGA are applied stratified at the STAR-CCM+ boundary. If there is outflow at the boundary, the values and mass fractions are averaged at the zone, and sent back to OLGA.

The STAR-CCM+ solver uses a smaller time step than OLGA, due to OLGA being more efficient with longer time steps. Data is exchanged at each OLGA time step, with the received values stored and used at the next STAR-CCM+ time step.

In the OLGA simulation, a STAR-CCM+ velocity inlet connects to an OLGA *boundary*, and a STAR-CCM+ pressure outlet to an OLGA *source*.



Using this method provides transition from three-dimensional to one-dimensional simulation with area-averaged data.

In STAR-CCM+, a co-simulation session for coupling with OLGA is automatically created, within the Co-Simulations node, when you select an OLGA coupling model. By default, the co-simulation session node is called

OLGA5, and is set up for coupling with OLGA 5



To couple STAR-CCM+ to OLGA 7, you must set the OLGA Version property, for the OLGA Coupling model, to 7. This adds an extra node called OLGA7 under the Co-Simulations node. STAR-CCM+ uses the OLGA5 or OLGA7 node for co-simulation, depending on the OLGA version that is specified for the OLGA Coupling model.

The OLGA5 node holds two subnodes, OLGA Species and OLGA Zones, which contain species and zones, once you couple the simulation to OLGA. The same subnodes are available within the OLGA7 node. The OLGA7 node also contains the OLGA Locations subnode, which populates automatically when the simulation is coupled to OLGA.



What is an OLGA Species?

An OLGA species defines a component of the multiphase mixture used in the simulation. You can model these species in STAR-CCM+ with either the Volume of Fluid (VOF) or the Multiphase Segregated Flow (EMP) models, by selecting the correct OLGA coupling model.

The physics values and volume fractions of these species are exchanged at zones, between STAR-CCM+ and OLGA. All species from OLGA must be assigned to an Eulerian phase, within the STAR-CCM+ simulation.

There are three species exchanged in all OLGA coupled simulations, which are detailed below:

- Gas — A generic species representing the mixture of gas components from OLGA. OLGA contains tables, which are used to specify the material properties of the species as a result of the mixture.
- Oil — A generic species representing the mixture of oil components from OLGA. OLGA contains tables, which are used to specify the material properties of the species as a result of the mixture.
- Water — A specific species representing the water component from OLGA.

The OLGA Species node holds the individual species nodes, which become available when you select an OLGA coupling model. The following image shows the OLGA Species node for co-simulation with OLGA 5. The same node is available within the OLGA7 node, when coupling to OLGA 7.



What is an OLGA Zone?

An OLGA zone is a representation of the boundary of an OLGA model that is coupled to the CFD solution. Because of the nature of the OLGA representation, some field variables are available at the OLGA model **Boundary** or **Source**, and others at the adjacent control volume or **Position**. These variables are grouped together into the OLGA zone.

Each OLGA zone must be associated with an appropriate boundary in STAR-CCM+. The physics values and species volume fractions at the associated boundary are area-averaged, providing a single value which is applied to the OLGA boundary/source/position.

You can create new OLGA zones in the STAR-CCM+ simulation by right-clicking the OLGA Zones node. The OLGA boundary/source and position names are used to identify the coupling points in OLGA. Setting the name of the OLGA zone automatically sets the name of the OLGA coupled boundary/source, while the OLGA position name is set by filling in the property value.

The OLGA Zones node holds the individual zone nodes, which become available when you select an OLGA coupling model.The following image shows the OLGA Zones node for co-simulation with OLGA 5. The same node is available within the OLGA7 node, when coupling to OLGA 7.



Setting Up the STAR-CCM+ Simulation

This section details the procedure for setting up an OLGA coupling simulation.

The procedure described in this section is the same for OLGA 5 and OLGA 7. When coupling to OLGA 7, an extra step is required to change the OLGA version from 5 (default) to 7.

Also remember that, when coupling to OLGA 7, STAR-CCM+ must run on a Linux platform and connect to an OLGA session running on Windows.

- OLGA Coupling Models
  - Activating the OLGA Coupling Model
  - Activating the OLGA EMP Coupling Model

- Creating an OLGA Species Representation
  - Activating the Phase Models for Volume of Fluid
  - Activating the OLGA EMP Coupling Model
- Setting the OLGA Species Material Properties
- Identifying and Setting the Reference Values
- Identifying the Initial Conditions
  - Setting the Initial Pressure
  - Setting Initial Temperature for Volume of Fluid
  - Setting Initial Temperature for Eulerian Multiphase
  - Setting the Initial Volume Fractions
- Setting the STAR-CCM+ Boundary Type
  - Setting a Velocity Inlet
  - Setting a Pressure Outlet
  - Setting Coupled Boundary Conditions
- Setting the OLGA Version

OLGA Coupling Models

When coupling to OLGA, you can simulate the OLGA species by using the Multiphase Segregated Flow (EMP) or Volume of Fluid (VOF) models. The choice between each modeling approach depends on individual simulation requirements.

The EMP model provides a more detailed solution, and is suitable for cases where the phases are not in equilibrium. It solves the physical quantities, such as velocity, temperature, and turbulence, for each phase individually. However, it does assume that pressure is constant for all phases. It can also model interactions between phases, such as drag forces and energy transfer.

The VOF model provides a more numerically efficient solution, and is suitable for cases where the interface between each phase covers a large area and you can assume that the phases are in equilibrium. It assumes that the physical quantities, such as velocity, temperature, and turbulence, of all phases are identical.

See What Is the Multiphase Segregated Flow Model? and What Is the VOF Model? for more information on these multiphase flow models.

Activating the OLGA Coupling Model

You use the OLGA Coupling model to simulate the multiphase fluid with the Volume of Fluid model in STAR-CCM+. You select this coupling model from within the Physics Model Selection dialog of the relevant continuum.

To activate the OLGA Coupling model with VOF:

- Select Three Dimensional from the *Space* group box.
  - The Gradients model is selected automatically.
- Select Implicit Unsteady from the *Time* group box.
- Select Eulerian Multiphase from the *Material* group box.
  - The Multiphase Interaction model is selected automatically.
- Select Volume of Fluid (VOF) from the *Multiphase Model* group box.
  - The Multiphase Equation of State and Segregated Flow models are selected automatically.
- Select Co-Simulation from the *Optional Models* group box.
- Select OLGA Coupling from the *Co-Simulation* group box.
  - The Segregated Multi-Phase Temperature and Gravity models are selected automatically.

To complete the physics continua definition for an OLGA coupled simulation, activate the following models:

- Select Turbulent from the *Viscous Regime* group box.
  - The Reynolds-Averaged Navier-Stokes model is selected automatically.

- Select any turbulence model from the *Reynolds-Averaged Turbulence* group box, depending on simulation requirements.



Further models can be selected from the *Optional Physics Models* group, to meet individual simulation requirements.

Activating the OLGA EMP Coupling Model

You use the OLGA EMP Coupling model to simulate the multiphase fluid with the Multiphase Segregated Flow (EMP) model in STAR-CCM+. You select this coupling model from within the Physics Model Selection dialog of the relevant continuum.

To activate the OLGA EMP Coupling model:

- Select Three Dimensional from the *Space* group box.
  - The Gradients model is selected automatically.
- Select Implicit Unsteady from the *Time* group box.

- Select Eulerian Multiphase from the *Material* group box.
  - The Multiphase Interaction model is selected automatically.
- Select Multiphase Segregated Flow from the *Multiphase Flow Model* group box.
  - The Multiphase Equation of State model is selected automatically.
- Select Co-Simulation from the *Optional Models* group box.
- Select OLGA EMP Coupling from the *Co-Simulation* group box.
  - The Gravity model is selected automatically.



To complete the physics continua definition for an OLGA coupled simulation, activate the following models:

- Select Phase Coupled Fluid Energy from the *Energy* group box.
- Select Turbulent from the *Viscous Regime* group box.

Further models can be selected from the *Optional Physics Models* group, to meet individual simulation requirements.

Creating an OLGA Species Representation

An OLGA species defines a component of the multi-phase mixture used in the coupled simulation. All OLGA species must be assigned to an Eulerian phases for coupling to be successful.

Note: Creating a phase applies the OLGA EOS model to the phase automatically, as no other equation of state models are applicable for OLGA coupling.

To add an Eulerian phase to the simulation:

- Right-click the Eulerian Phases node, and select New.



This action adds a phase within the Eulerian Phases manager, which requires specific phase models to define its behavior. Rename each phase according to the OLGA species, for clarity when coupling the species to the phases.

Activating the Phase Models for Volume of Fluid

The VOF phase models define the material properties, equation of state, and basic turbulence behavior for each OLGA phase.

To define the Eulerian phase models for VOF

- Right-click the Eulerian Phases > [Phase] > Models manager and select Select models.



- In the *[Phase] Model Selection* dialog:
  - For the oil and water species, select Liquid from the *Material* group box.
  - For the gas species, select Gas from the *Material* group box.



Activating the Phase Models for Eulerian Multiphase

The EMP phase models define the material properties, equation of state, and detailed turbulence behavior for each OLGA phase.

To define the Eulerian phase models for EMP:

- Right-click the Eulerian Phases > [Phase] > Models manager and select Select models.



- In the *[Phase] Model Selection* dialog:
  - For the oil and water species, select Liquid from the *Material* group box.
  - For the gas species, select Gas from the *Material* group box.
- Select any turbulence model from the *Reynolds-Averaged Turbulence* group box, depending on simulation requirements.

- Select the Segregated Fluid Temperature model from the *Energy Group* box.



Setting the OLGA Species Material Properties

Each of the OLGA phases must be defined with a material, which closely matches the OLGA species properties.

To define the phase material properties:

- Right-click the Models > Gas/Liquid > Material node and select Replace with.



- In the object selection dialog, select a material that closely matches the OLGA species.



- Modify the material properties so that they match the species properties from OLGA.
  - OLGA automatically defines the Density, Density pressure derivative, and Density temperature derivative material properties, which you cannot change.



- Set all other material properties to the average values from the OLGA species property table.



Defining Phase Interactions for Eulerian Multiphase

When using the OLGA EMP coupling model, define phase interactions between phases to determine how the phases behave when there is an interface between them.

The choice of which phases interact with one another and the relevant phase interaction models depends on the flow regime that you are modeling. See Modeling Phase Interactions in Multiphase Segregated Flow for more information on phase interactions.

Note: For simulations where the gas and liquid phases can change from continuous to dispersed (that is from bubbly to droplet flows), use the symmetric drag force method.

To define a phase interaction:

- Right-click the Continua > Multiphase Interaction > Phase Interactions node and select New.



- Right-click the Phase Interaction > Models node and select Select models.



- In the *Phase Interaction Model Selection* dialog, select the required phase interaction models to model your flow regime.

- Select the Phase Interaction > Models > Eulerian Continuous-Dispersed node and set the Continuous Phase and Dispersed Phase properties.



Identifying and Setting Reference Values

Reference values determine the value to which some physical quantities are relative. For OLGA coupled simulations, set the gravity orientation and reference pressure to appropriate values.

Gravity Orientation

To ensure that the simulation runs with the oil, water, and gas species separated as expected, set the reference gravity correctly. The direction of gravity is relative to the global coordinate system.

In the example below, gravity must be defined as acting in the negative y-direction.



To set the gravity orientation as in the above example:

- Select the Gravity node and set the Value to `[0.0, -9.81, 0.0] m/s^2`.



Reference Pressure

The reference pressure is used in STAR-CCM+ to reduce the error due to numerical round-off. All other pressures in a STAR-CCM+ simulation are defined relative to the reference pressure.

However, in an OLGA simulation, all pressures are absolute. Take this fact into consideration when defining pressure values between the two codes.

To reduce round-off error, set the reference pressure to a value that is close to the expected average pressure from OLGA. You can identify the average pressure in the region by running the OLGA simulation independently until converged. Upon completion, OLGA reports the average pressure in the CFD component in its .out file.

```
2. EXTERNAL BOUNDARY CONDITIONS
*******************************


 INLET BRANCH              DOWNSTREAM

 BOUNDARY COND TYPE -    : CLOSED


 OUTLET BRANCH             PIPELINE

 BOUNDARY COND TYPE -    : PRESSURE
 VOID FRACTION      -    :    0.866
 TEMPERATURE     (C)     :     20.0
 PRESSURE        (PA)    :    0.800E+07
```

To set the reference pressure:

- Select the Reference Pressure node and set Value to the average pressure in OLGA.



Identifying the Initial Conditions

This section details the process of setting the initial conditions for OLGA coupled simulations. The initial conditions in a STAR-CCM+ simulation are used to improve the time taken for convergence.

To identify appropriate initial conditions:

- Run the OLGA simulation independently from STAR-CCM+ until converged.
- Open the OLGA .out file.
  - You can typically find this file in the same directory as the OLGA simulation file.

- Note the pressure and temperature for the branch connected to the STAR-CCM+ simulation. For example:

```
2. EXTERNAL BOUNDARY CONDITIONS
*******************************

 INLET BRANCH              DOWNSTREAM

 BOUNDARY COND TYPE -   : CLOSED

 OUTLET BRANCH             PIPELINE

 BOUNDARY COND TYPE -   : PRESSURE
 VOID FRACTION      -   :    0.866
 TEMPERATURE     (C)    :     20.0
 PRESSURE        (PA)   :    0.800E+07
```

Setting the Initial Pressure

The initial pressure in STAR-CCM+ is defined relative to the reference pressure. Therefore, the value is set to the difference between the average pressure from OLGA, and the STAR-CCM+ reference pressure.

This example uses a constant value, but you can also use field functions and tabular data.

To set the initial pressure:

- Select the Initial Conditions > Pressure > Constant node and set the Value property.



Setting Initial Temperature for Volume of Fluid

For VOF simulations, you set the initial temperature at the continuum level. The initial temperature in STAR-CCM+ is an absolute value; so the value is the same as the average temperature from OLGA.

This example uses a constant value, but you can also use field functions and tabular data.

To set the initial temperature for a VOF simulation:

- Select the Initial Conditions > Static Temperature > Constant node and set the Value property.



Setting Initial Temperature for Eulerian Multiphase

For EMP simulations, you set the initial temperature at the phase level for each species. The initial temperature in STAR-CCM+ is an absolute value; so the value is the same as the average temperature from OLGA.

This example uses a constant value, but you can also use field functions and tabular data.

To set the initial temperature for an EMP simulation:

- Select the Eulerian Phases > [Phase] > Initial Conditions > Static Temperature > Constant node and set the Value Property.

Setting Initial Volume Fractions

If you do not have accurate initial volume fractions, initialize the region with a gas volume fraction of 1. The gas is then pushed out of the region as the solver runs, providing a faster solution than initializing with oil or water.

However, if you have more accurate volume fractions, you can use these values as an alternative.

To set the initial volume fractions:

- Select the Initial Conditions > Volume Fraction > Constant node and click the ellipsis (...) for the Value property.



- In the array, enter the mass fractions for each phase.



STAR-CCM+ Boundary Type

The boundary type determines the physics values exchanged between STAR-CCM+ and OLGA. The flow in an OLGA coupled simulation is typically one-directional, although temporary reverse flow can be modeled. Hence, OLGA coupling uses a velocity inlet boundary for flow from OLGA, and a pressure outlet for flow to OLGA.

The following table summarizes of the data that STAR-CCM+ receives from OLGA:

Boundary Type	Received Data
Velocity Inlet	Mass flux of each species
	Eulerian component velocities, applied stratified at zone
	Eulerian component densities, applied stratified at zone
Pressure Outlet	Mass flux of each species, used for reverse flow
	Pressure, applied uniformly at zone
	Temperature, applied uniformly at zone

OLGA defines most of boundary conditions in an OLGA coupled simulation, except turbulence.

Setting a Velocity Inlet

Note: You can only couple a STAR-CCM+ velocity inlet an OLGA *boundary*.

To set a velocity inlet:

- Select the coupled boundary and set Type to Velocity Inlet.



Setting a Pressure Outlet

Note: You can only couple a STAR-CCM+ pressure outlet to an OLGA *source*.

To set a pressure outlet:

- Select the coupled boundary and set the Type to Pressure Outlet.



Setting the Coupled Boundary Conditions

To set a boundary as coupled:

- Select the Physics Conditions > 1D Coupling Specification node and set 1D Coupling Specification to Coupled.



Setting the OLGA Version

Currently, STAR-CCM+ supports coupling with OLGA 5.3.2 and OLGA 7.2.2. When activating an OLGA Coupling model, STAR-CCM+ automatically sets the OLGA Version property for the OLGA Coupling model to 5. A co-simulation session node, OLGA5, is added within the Co-Simulations node. If you wish to couple STAR-CCM+ to OLGA 7, you need to change the setting for the OLGA Version property.

To change the OLGA version for coupling to OLGA 7:

- Expand the Continua > [Physics Continuum] > Models node.

- Select the OLGA Coupling node and set OLGA Version to 7.



When the OLGA version is set to 7, an extra node, OLGA7, is added under the Co-Simulations node. When coupling to OLGA7, make sure that you are working within the relevant co-simulation session node, that is, OLGA7.



Coupling STAR-CCM+ to OLGA 5

This section details the procedure for coupling STAR-CCM+ to OLGA 5.

- Assigning the OLGA Species
- Creating OLGA Zones
  - Setting the OLGA Name
  - Setting the OLGA Position Name
  - Assigning a Boundary to an OLGA Zone
- Starting OLGA Coupling
  - Starting the OLGA Server
  - Setting the Port Number and Host Name
  - Defining the OLGA Input File
  - Defining the OLGA Property File
  - Connecting to the OLGA Server
- Setting the Couple From Step Option
- Setting the OLGA Resume File

With the above points addressed as necessary, a solution can be obtained.

Assigning the OLGA Species

All species in OLGA coupling must be assigned to a previously defined Eulerian phase. This assignment provides a means of representing the multiphase mixture in the STAR-CCM+ simulation.

To assign an OLGA species:

- Select an OLGA Species > Species and set Eulerian Phase to the relevant user-defined phase.



Creating an OLGA Zone

All zones in OLGA coupling must be assigned to a STAR-CCM+ boundary, with the correct boundary type and conditions. This assignment provides a means of identifying the location of data exchange between the one-dimensional and three-dimensional code.

To create and assign an OLGA zone

- Right-click the OLGA > OLGA Zones node and select Add New OLGA Zone.



The following properties are then required to define the zone:

- OLGA Name
- OLGA Position Name
- Boundary

Setting the OLGA Name

The OLGA Name is set by renaming the OLGA Zone node in the simulation tree, with the same name as the corresponding entity in OLGA. An OLGA boundary is required when coupling to a STAR-CCM+ velocity inlet, and an OLGA source when coupling to STAR-CCM+ pressure boundary.

Note: This name must match the OLGA name exactly for coupling to be successful, and is case-sensitive.

To set the OLGA name:

- Select an OLGA Zone, press <F2> on the keyboard, and rename the node.



The OLGA Name also appears as a read-only property in the *Properties* window.

Setting the OLGA Position Name

The *OLGA Pos Name* is set in the *Properties* window, with the same name as the corresponding position in OLGA.

Note: This entry must match exactly for coupling to be successful, and is case-sensitive.

To set the OLGA position name:

- Select an OLGA Zone and set OLGA Pos Name to the OLGA position name.



Assigning a Boundary to an OLGA Zone

The zone is assigned to a boundary that you have specified as available for coupling.

To assign a boundary to an OLGA zone:

- Select an OLGA Zone and set Boundary to a coupled boundary.



Starting OLGA Coupling

Before coupling for the first time, perform the following steps on the machine on which OLGA will run:

- Navigate to `C:\WINDOWS\system32\drivers\etc\` and open the file, `services`, with a pure text editor.
- Add at least one service for the OLGA server at the end of the file, making sure to match the format of the existing entries:

  ```
  [olga_server_0_name] 12340/tcp #For STAR-OLGA coupling 0
  ```

  Where:

  - `[olga_server_0_name]` is the name of the OLGA server.
  - `12340/tcp` is the port number of the OLGA server.
  - `#For STAR-OLGA coupling 0` is a description for the service.

Note: You can add any number of OLGA servers to the service file.

The following steps are then required to start OLGA coupling:

- Start the OLGA server
- Set the port number and host name
- Define the OLGA input file
- Define the OLGA property file
- Connect to the OLGA server

Starting the OLGA Server

To couple STAR-CCM+ to OLGA, first start an OLGA server. This server runs OLGA in batch mode and sends data to STAR-CCM+.

An OLGA server must be started manually in a command prompt. The easiest way to do this is to create a batch file as follows:

- Create a batch file, named `olgaServer.bat`, on the machine that is used to run OLGA. The batch file must contain the following:

  ```
  C:\"Program Files"\"SPT Group"\"OLGA 5.3.2"\OlgaExecutables_5_3_2\olga-5.3.2.exe -server olga
  ```

  `olgaServer.bat`

  Where:

  - The text in red is the path to your OLGA executable.
  - The text marked in blue is the OLGA service name. This name must match the name you specified in the `services` file.
  - `olgaServer.bat` keeps the server open, waiting for input from STAR-CCM+.

- Double-click the `olgaServer.bat` file.

  This action opens a command prompt, executes OLGA as a server, and waits for input from STAR-CCM+.



When running the coupled simulation, the OLGA solver prints data in the command prompt.

Setting the Port Number and Host Name

The port number and host name identify the location of the server.

To define the port number and host name:

- Select the Co-Simulations > OLGA5 node and set Port Number to the port that you identified above.



- Set Host Name to the host that you identified above



Defining the OLGA Input File

STAR-CCM+ requires you to define the OLGA input file used in the coupled simulation. This file must be on the machine that runs OLGA during the coupled simulation.

Note: You define this file with its absolute directory and file name, which has a .geninp extension.

Currently, you can only use a maximum of 64 characters to define the input file location, which includes the directory path and file name. To overcome this issue, move input file to a location on the hard disk that meets this criteria.

To define the OLGA input file:

- Select the Co-Simulations > OLGA5 node and set Input File to the absolute path and file name.



Defining the OLGA Property File

STAR-CCM+ requires you to define the OLGA property file used in the coupled simulation. The property file must be in the working directory of both the STAR-CCM+ simulation and the `olgaServer.bat` file. On the machine running OLGA, the property file must be in the same directory as the `.geninp` file.

Note: You define this file with its file name only, as STAR-CCM+ assumes that the file is in the current working directory of both the STAR-CCM+ simulation file and OLGA server.

To define the OLGA property file:

- Select the Co-Simulations > OLGA5 node and set Property File to the property file name.



Connecting to the OLGA Server

Once the simulation is fully defined, you can couple STAR-CCM+ to OLGA. STAR-CCM+ connects to the OLGA server, allowing for the exchange of species volume fractions and physics values between the one-dimensional and three-dimensional code.

To start OLGA coupling:

- Right-click the OLGA5 node, and select Start OLGA.



- Alternatively, click Run (Ctrl+R) in the toolbar.

  When complete, the *Output* window displays:

  ```
  OLGA Initialized
  ```

Use the right-click method the first time that you couple STAR-CCM+ to OLGA or if you want to verify that all the settings are correct before beginning the run.

You can stop the OLGA solver at any time by clicking Solution > Clear Solution from the menu bar, which breaks the connection between STAR-CCM+ and OLGA. You can then restart OLGA using the methods mentioned above, which resets the coupling and takes into account any changes made to the OLGA case.

Setting the Couple From Step Option

The Couple From Step option delays the transfer of data to the OLGA solver until a defined time step. This feature allows the STAR-CCM+ solver to reach a level of stability while receiving data from OLGA, before sending data back to OLGA.

To set the Couple From Step option:

- Select the Co-Simulations > OLGA5 node and set Couple From Step to the amount of time steps that you want to delay data transfer.



Setting the OLGA Resume File

The OLGA resume file restarts the coupled simulation from a specific time step. This restart facility is useful in simulations with long processing times, as you can stop and restart the solver when necessary.

OLGA creates the resume file automatically when you specify the simulation as a restart run. Resume File0 contains the solution data the previous OLGA time step, and Resume File1 contains the solution data the next OLGA time step. Together these two files fully define the point at which coupling resumes on a restart, as the three-dimensional time step is less than the one-dimensional time step.

To set the resume files:

- Select the Co-Simulations > OLGA5 node and set the Resume File0 and Resume File1 properties to the resume file names.



Coupling STAR-CCM+ to OLGA 7

This section explains the procedure for coupling STAR-CCM+ to OLGA 7.

- Setting Up the OLGA Case for Coupling
- Configuring APIS
- Assigning the OLGA Species
- Creating OLGA Zones
  - Assigning a Boundary to an OLGA Zone
- Starting OLGA Coupling
  - Setting the Port Number
  - Setting the Host Name
  - Defining the OLGA Property File
  - Setting the Certificates Root Directory
  - Setting the Couple From Step Option
  - Connecting to the OLGA Server
  - Setting the OLGA zone Position Locations and Boundary/Source Names

With the above points addressed as necessary, a solution can be obtained.

Setting Up the OLGA Case for Coupling

Modify your case in OLGA 7, to set it up for coupling to STAR-CCM+. In OLGA 7, the variables that you want to set using STAR-CCM+ incoming data, must be *Exposed*. The variables that you want to send to STAR-CCM+ must be set as *Output* variables.

Select the data that will be exchanged between STAR-CCM+ and OLGA 7:

- Launch OLGA 7.2.2.
- Open the OLGA case that you want to couple to your STAR-CCM+ simulation.
- *Expose* the variables that you want STAR-CCM+to set.



- Set the variables that you want to send to STAR-CCM+ as *Output*.

- In the Toolbar, click Run Batch:

A message similar to the following is displayed in the command prompt:.



Configuring APIS

To couple STAR-CCM+ to OLGA 7, you need to install the third party software, APIS (Prediktor Advanced Plant Information System) from Prediktor, which allows STAR-CCM+ to communicate with OLGA 7.

APIS must be installed on the Windows platform running OLGA 7.

For information about APIS installation, please refer to chapter 2 of the APIS Installation Guide, which is provided in your APIS installation folder.

To configure APIS:

- Start APIS Buddy.
- In the top menu, click Apishive > Start.
- In the top menu, click Apishive > Configure.

  The *ApisHive - [Console Root\ApisHive\Namespace Root]* window appears.
- Right-click the ApisHive node on the left, and select New > Module.
- In the *Add Apis Module* dialog:
  - Set Type to ApisOPC.
  - Set Name to `OPCmodule`.

    - click OK.
- In the *Initial OPCmodule Properties* window:
    - Select Server and set Property value to SPT.[servername].7.
    - If OLGA is running on a different machine, select Computer and set Property value to the ip address of the machine running OLGA.
    - Select ReadMode and set Property value to Synch I/O - caller thread.
    - Press Expert.
    - Select ReconnectTime and set Property value to 10 sec.
    - Select ForceRefreshOnConnect and set Property value to true.
    - Select ReconnectSrvShutdown and set Property value to true.
    - Click Ok.
- Right-click the ApisHive > OPCmodule node on the left, and select New > OPC Item.
- In the *Add OPC Item item(s) to module OPCmodule* window:
    - Click Click to Browse for items.
    - Click Check... at the bottom of the window.
    - In the *check candidates* window, click All and then OK.
- Click OK.
- Click the ApisHive > Namespace Root node.

  A list of items, representing the data that is exchanged between STAR-CCM+ and OLGA, is displayed in the right panel. If nothing is displayed, right-click the ApisHive > Namespace Root node and make sure that Auto Refresh is activated.

Assigning the OLGA Species

All species in OLGA coupling must be assigned to a previously defined Eulerian phase. This assignment provides a means of representing the multiphase mixture in the STAR-CCM+ simulation.

To assign an OLGA species:

- Expand the Co-Simulations > OLGA7 > OLGA Species node.

- Select an OLGA species and set Eulerian Phase to the relevant user-defined phase.



Creating an OLGA Zone

All zones in OLGA coupling must be assigned to a STAR-CCM+ boundary, with the correct boundary type and conditions. This assignment provides a means of identifying the location of data exchange between the one-dimensional and three-dimensional code.

To create and assign an OLGA zone:

- Right-click the Co-Simulations > OLGA7 > OLGA Zones node and select Add New OLGA Zone.



Assigning a Boundary to an OLGA Zone

The zone is assigned to a boundary that you have specified as available for coupling.

To assign a boundary to an OLGA zone:

- Select an OLGA Zone and set Boundary to a coupled boundary.



The Boundary/Source Name and Position Location properties can be set once the simulation is coupled to OLGA.

Starting OLGA Coupling

The following steps are then required to start OLGA coupling:

- Setting the Port Number
- Setting the Host Name
- Defining the OLGA Property File
- Setting the Certificates Root Directory
- Setting the Couple From Step Option
- Connecting to the OLGA server
- Setting the OLGA zone Position Locations and Boundary/Source Names

Setting the Port Number

The port number and host name identify the location of the server.

The port number needs to match the APIS port number, which is specified in the `ApisHive_UA.config` file, contained in your APIS installation folder.

To check the port number:

- Open the `ApisHive_UA.config` file with a text editor of your choice.
- Locate the following text:

  ```
  <!-- Specify ua server url and port -->
  ```

  the port number is specified after localhost:

  ```
  <core url="opc.tcp://localhost:[portnumber]">
  ```
- Close the file.

To set the port number in STAR-CCM+:

- Select the Co-Simulations > OLGA7 node and set Port Number to match the APIS port number:



Setting the Host Name

To set the host name:

- Select the Co-Simulations > OLGA7 node.
- Set Host Name to the name of the machine running OLGA:



Defining the OLGA Property File

STAR-CCM+ requires you to define the OLGA property file used in the coupled simulation. The property file must be in the working directory of the STAR-CCM+ simulation. On the machine running OLGA, the property file must be in the same directory as the OLGA project file.

Note: You define this file with its file name only, as STAR-CCM+ assumes that the file is in the current working directory of both the STAR-CCM+ simulation file and OLGA server.

To define the OLGA property file:

- Select the Co-Simulations > OLGA7 node and set Property File to the name of the property file.



Setting the Certificates Root Directory

In order to couple STAR-CCM+ to OLGA 7, APIS requires STAR-CCM+ to provide a security certificate. To obtain the certificate, contact SPT Group.

In STAR-CCM+, you must specify the location of this certificate:

- Select the Co-Simulations > OLGA7 node.

- Set Certificates Root Directory to the absolute or relative path to the folder containing your certificate file. If you enter a relative path, the path must be relative to the directory from which STAR-CCM+ is launched.



Setting the Couple From Step Option

The Couple From Step option delays the transfer of data to the OLGA solver until a defined time step. This feature allows the STAR-CCM+ solver to reach a level of stability while receiving data from OLGA, before sending data back to OLGA.

To set the Couple From Step option:

- Select the Co-Simulations > OLGA7 node and set Couple From Step to the amount of time steps that you want to delay data transfer.

Connecting to the OLGA Server

Once the simulation is fully defined, you can couple STAR-CCM+ to OLGA. STAR-CCM+ connects to the OLGA server, allowing for the exchange of species volume fractions and physics values between the one-dimensional and three-dimensional code.

To start OLGA coupling:

- Right-click the OLGA7 node, and select Start OLGA 7.



- Alternatively, click Run (Ctrl+R) in the toolbar.

  When complete, the *Output* window displays:

  ```
  OLGA Initialized
  ```

Use the right-click method the first time that you couple STAR-CCM+ to OLGA or if you want to verify that all the settings are correct before beginning the run.

You can stop the OLGA solver at any time by clicking Solution > Clear Solution from the menu bar, which breaks the connection between STAR-CCM+ and OLGA. You can then restart OLGA using the methods mentioned above, which resets the coupling and takes into account any changes made to the OLGA case.

Setting the OLGA Zone Position Location and Boundary/Source Name

Once the simulation is coupled to OLGA 7, you can assign the correct OLGA Position and OLGA Boundary/Source to the OLGA Zones.

To set the OLGA Position for an OLGA zone:

- Select the Co-Simulations > OLGA7 > OLGA Zones > [OLGA zone].
- Set Position Location to the OLGA Position corresponding to this zone.



To set the OLGA Boundary/Source for an OLGA zone:

- Select the Co-Simulations > OLGA7 > OLGA Zones > [OLGA zone].

- Set Boundary/Source Name to the OLGA Boundary corresponding to this zone.



Area Mismatch Error in OLGA Coupling

Coupling requires the assigned boundary to have a circular cross-section, with an area, and hence diameter, equal to the OLGA zone. Currently, the area of the zone cannot be read directly from OLGA. Instead, the area is calculated using the received physics values from OLGA (mass flux, density, and velocity). A large discrepancy between these physics values typically indicates one of two issues that require attention:

- An area mismatch exists between the two codes. To resolve these types of issues, check the units and sizes of the pipe in STAR-CCM+ and in OLGA, and correct as necessary. If the mesh in STAR-CCM+ is excessively coarse, discretization errors might alter the calculated area. In this case, you may require to adjust the nominal radius to get the correct flow area.
- The flow is unstable, changing rapidly, or is diverging. In such cases, you will have to adjust the run settings (for example, reduce the time-step).

In cases of area mismatch, if the calculated flow area is larger than 10% of the geometric area, the following warning message is displayed in the *Output* window:

```
Warning: Area mismatch between OLGA data and STAR-CCM+
geometry,check for geometry errors or solution divergence
```

If the calculated flow area is negative and less than -10% of the geometric area, the following warning message is displayed in the *Output* window:

```
Warning: Apparent negative area implied by OLGA boundary
data, solution may be diverging
```

Errors of less than 10% are ignored.

WAVE Coupling

The WAVE Coupling feature provides the ability for two-way transfer of data between three-dimensional STAR-CCM+, and one-dimensional WAVE. For example, STAR-CCM+ can provide the pressure-drop calculations for tubes and manifolds, with this information fed back to WAVE for a full engine-power simulation.

This facility provides increased detail in specified areas of the WAVE case, by adding a three-dimensional element to a one-dimensional simulation. Similarly, processing times can be reduced in a STAR-CCM+ simulation, by utilizing the speed of one-dimensional codes.

The following section provides the definition of common terms:

- WAVE Coupling
- WAVE Species
- WAVE Zones

The documentation also includes a guide to using WAVE Coupling, which details:

- Preparing the WAVE Case
- Setting Up the STAR-CCM+ Simulation
- Coupling STAR-CCM+ to WAVE

What is WAVE Coupling?

WAVE coupling provides a means of data exchange between three-dimensional STAR-CCM+ and one-dimensional WAVE. This allows for a compromise between the speed of one-dimensional code, with the detail of three-dimensional code. The feature provides optimum results when coupling occurs at a zone where the flow is expected to be largely one-dimensional.

When running the solver, the physics values and species mass fractions that are received from WAVE are applied uniformly at the STAR-CCM+ boundary. If there is outflow at the boundary, the values and mass fractions are averaged at the zone, and sent back to WAVE. Data is exchanged at either each time step or inner iteration, depending on the simulation setup.

In the WAVE simulation, the STAR-CCM+ boundary connects to a WAVE CFD orifice.



Using this method provides averaged data, and a means of transition from three-dimensional to one-dimensional simulation.

The WAVE node, which is located within the Co-Simulations node, becomes available when the WAVE Coupling model is activated. It holds sub-nodes, which are populated with species and zones, once the simulation is coupled to WAVE.



What is a WAVE Species?

A WAVE species defines a component or components of the fluid mixture that is used in the simulation. The physics values and mass fractions of these species are exchanged at zones, between STAR-CCM+ and WAVE. All species from WAVE must be assigned to a mixture component or passive scalar, within the STAR-CCM+ simulation. Multiple species can be assigned to individual mixture components and passive scalars, which reduces process times, but compromises detail and accuracy.

WAVE uses combustion models to determine the products of reaction between fuel and air, which is given the engine operating conditions and environment. There are five species that are exchanged in all WAVE coupled simulations, which are detailed below:

- air -- Air, with gas composition aN_2, bO_2, cCO_2, dH_2O, defined in WAVE pre-processing.
- burnedAir -- Products of combustion, with gas composition, aN_2, bO_2, cCO_2, dH_2O.
- burnedFuel -- Remaining products of combustion, with gas composition, eCO, fNO, gOH, hN, iO, jH kH_2.

- fuelLiquid -- Fuel in liquid form, with molecular composition $C_C H_H O_O N_N$, specified in WAVE pre-processing.
- fuelVapour -- Fuel in vapor form, with molecular composition $C_C H_H O_O N_N$, specified in WAVE pre-processing.

The WAVE Species manager holds the individual species nodes, which become available when WAVE Coupling is activated and started.



What is a WAVE Zone?

A WAVE zone is a specified boundary in the CFD region, which provides a means of data exchange between STAR-CCM+ and WAVE. Individual zones are defined in WAVE, with equivalent boundaries required in the STAR-CCM+ simulation. The physics values and species mass fractions at the zones are averaged, providing a single value that is used to determine the equivalent boundary conditions in the one-dimensional code. Each WAVE zone must be assigned to a boundary in the STAR-CCM+ simulation.

Zones are labeled according to the *Interface Number* defined in the WAVE simulation. Rename the corresponding boundary with the same name, for clarity when coupling the zone to the boundary.

The WAVE Zones manager holds the individual zone nodes, which become available when WAVE Coupling is activated and started.



Preparing the WAVE Case

The settings in the WAVE case are defined before coupling, and determine certain properties and characteristics within the STAR-CCM+ simulation:

- Setting the Zone Diameter
- Setting WAVE as the Child Process
- Setting WAVE Pre-Cycles
- Delaying Data Transfer to WAVE

Setting the Zone Diameter

WAVE coupling requires the diameter of the zone in STAR-CCM+ to be equal to the WAVE CFD orifice. If the zone cross-section is non-circular, then the effective diameter is used. The effective diameter being the diameter of an equivalent circle with the same cross-sectional area as the zone.

Using the example below, the diameter at the manifold intake is 50 mm.



To set the CFD orifice diameter in WAVE, double-click External CFD Orifice.



The *External CFD Orifice* dialog appears. Here, set the value in the Diameter textbox to match the zone diameter in STAR-CCM+.

Setting WAVE as the Child Process

In order for STAR-CCM+ to couple to WAVE, the *External CFD Model* settings must define WAVE as the child process. This parameter allows STAR-CCM+ to start the WAVE solver when obtaining a solution.

To set WAVE as the child process, from the WAVE menu bar, select Model > External > CFD..., and select Child from the *Parameters* group.



Setting WAVE Pre-Cycles

Pre-cycles are used to improve stability in WAVE, before data is transferred to STAR-CCM+. This facility provides more accurate boundary conditions from WAVE, reducing the time that is required to reach a converged solution.

To use pre-cycles in a coupled simulation, from the WAVE menu bar, select Model > External > CFD..., and select Postponed from the *Process Control* group.

The number of coupled cycles is then entered in the textbox labeled Coupled Duration.

This value is subtracted from the total number of simulation cycles, to determine the number of pre-cycles.

When obtaining a solution, the WAVE solver is run independently, until the number of pre-cycles is met, at which point data is sent to STAR-CCM+.

Delaying Data Transfer to WAVE

The transfer of data to WAVE from STAR-CCM+ can be delayed until a defined time step. This facility allows the STAR-CCM+ solver to reach a level of stability while receiving data from WAVE, before sending data back to WAVE.

To set a transfer delay, from the WAVE menu bar, select Model > External > CFD..., and enter the number of time-steps in the One Way Steps textbox.

Setting Up the STAR-CCM+ Simulation

The following section details the process of setting up a STAR-CCM+ simulation for coupling with WAVE:

- Activating the WAVE Coupling Model
- WAVE Species Representation
  - Defining Mixture Components for WAVE Species
  - Defining Passive Scalars for WAVE Species
- Setting the Reference Pressure
- Identifying Initial Conditions
  - Setting the Initial Pressure
  - Setting the Initial Temperature
  - Setting Mass Fractions
- Setting the STAR-CCM+ Boundary

Activating the WAVE Coupling Model

The WAVE Coupling model is selected from within the Physics Model Selection dialog on the relevant continuum.

To select the WAVE Coupling model, the following models are required:

- Select Three Dimensional from the *Space* group box.
- Select Multi-Component Gas from the *Material* group box.
- Select Non-Reacting from the *Reaction Regime* group box.

- Select either Coupled or Segregated from the *Flow* group, depending on the required solver type. The Segregated solver is recommended in most cases.
- Implicit Unsteady from the *Time* group box.
- Select Co-Simulation from the *Optional Physics Models* group box.
- Select WAVE Coupling from the *Co-Simulation* group box.

To define the physics continua that are used in a WAVE coupled simulation, activate the following additional models:

- Ideal Gas from the *Equation of State* group
- Turbulent from the *Viscous Regime* group
- Any turbulence model from the *Reynolds-Averaged Turbulence* group, depending on simulation requirements. K-Epsilon Turbulence is recommended for most simulations.

- Passive Scalar from the *Optional Physics Models* group, which may be neglected if all species are represented by mixture components.



Further models can be selected from the *Optional Physics Models* group to meet individual simulation requirements.

WAVE Species Representations

A WAVE species defines a component or components of the fluid mixture that is used in the coupled simulation. All WAVE species must be assigned to either a mixture component or passive scalar, which defines the behavior as active or passive.

The choice between using a mixture component or passive scalar to represent a WAVE species is dependent on user requirements. A mixture component influences the physical properties of the simulation, providing the most accurate and realistic solution. A passive scalar has no effect on the

physical properties of the simulation, which reduces processing times. If passive scalars are used for all species, a background fluid must be present in the region.

Although the fuelLiquid species is a liquid, it can be represented by a gas in the Multi-Component Gas model. This will not affect results, as the gas simply represents the species in STAR-CCM+, with the physics values received from WAVE.

Defining Mixture Components for WAVE Species

To add a mixture component to the simulation, right-click the Gas Components manager, and select Select Mixture Components...



The object selection dialog appears, where you can select gas components from the tree.



Clicking Apply adds the gases to the feature tree, within the Gas Components node.

Defining Passive Scalars for WAVE Species

To add a passive scalar to the simulation, right-click the Passive Scalars node, and select New.



A new passive scalar is added to the feature tree, within the Passive Scalars node.

Setting the Reference Pressure

The reference pressure is used in STAR-CCM+ to reduce the error due to numerical round-off. All other pressures in a STAR-CCM+ simulation are defined relative to the reference pressure. However, in a WAVE simulation, all pressures are absolute. Take this information into consideration when defining pressure values between the two codes.

In order to reduce the risk of round-off error, set the reference pressure to a value close to the expected average pressure from WAVE. A useful means of identifying the average pressure in the region is to run the WAVE simulation independently until converged. The average pressure in the CFD component can then be found in the relevant WAVE `.out` file.

```
-----------------------------------------------------------------------------
  F I N A L   O U T P U T   O F   D U C T S
-----------------------------------------------------------------------------
        Duct        Junction  TWALL   TAV    PAV    PMAX    PMIN    UMAX
                               [K]    [K]   [bar]   [bar]   [bar]   [m/s]

       duct1            amb1  300.0  300.0  1.000   1.000   1.000     0.1
                      Intake  300.0  300.0  1.000   1.000   1.000     0.1

      duct10       cylinder2  400.0  595.8  1.000   1.052   0.947    57.2
                        orif6  400.0  564.5  1.001   1.049   0.956    15.4

      duct11            cyl1  300.0  329.9  0.172   0.368   0.146     0.0
                        orif1  300.0  388.5  0.172   0.365   0.146     4.9
```

To set the reference pressure, select the Reference Pressure node, and enter the Value in the *Properties* window.



Setting the Initial Conditions

This section details the process for setting the initial conditions in STAR-CCM+. The following examples use constant values, with field functions and tabular data also available:

- Setting the Initial Pressure
- Setting the Initial Temperature

- Setting Initial Mass Fractions

The initial conditions in a STAR-CCM+ simulation are used to reduce the time that it takes for the solution to converge. Set the STAR-CCM+ initial conditions to the average values from WAVE, which can be obtained using the following steps:

- Run the WAVE simulation independently until converged.
- Using a text editor, open the `.out` file that is associated with the WAVE simulation. This file is typically found in the same directory as the WAVE simulation file.
- Note the average pressure and temperature for the adjacent duct to the CFD inlet. For example:

```
-----------------------------------------------------------------------------
  F I N A L   O U T P U T   O F   D U C T S
-----------------------------------------------------------------------------
        Duct        Junction  TWALL   TAV    PAV     PMAX    PMIN     UMAX
                                [K]    [K]   [bar]    [bar]   [bar]    [m/s]

       duct1            amb1  300.0  300.0  1.000   1.000   1.000      0.1
                      Intake  300.0  300.0  1.000   1.000   1.000      0.1

      duct10        cylinder2  400.0  595.8  1.000   1.052   0.947     57.2
                        orif6  400.0  564.5  1.001   1.049   0.956     15.4

      duct11            cyl1  300.0  329.9  0.172   0.368   0.146      0.0
                        orif1  300.0  388.5  0.172   0.365   0.146      4.9
```

Initial conditions are located in the feature tree under the Physics manager.



Setting the Initial Pressure

The initial pressure in STAR-CCM+ is defined relative to the reference pressure. Therefore, the value is set to the difference between the average pressure from WAVE, and the STAR-CCM+ reference pressure.

To set the initial pressure, select the Constant node within the Pressure manager, and enter the Value in the *Properties* window.



Setting the Initial Temperature

The initial temperature in STAR-CCM+ is defined as an absolute value. Therefore, the value is set the same as the average temperature from WAVE.

To set the initial temperature, select the Constant node within the Static Temperature manager, and enter the Value in the *Properties* window.



Setting Mass Fractions

The species mass fraction varies depending on the simulation requirements. Initial mass fractions must sum to the value of 1.

As an example, if coupling to a component on the intake side of an engine, it would be intuitive to initialize the region with air. Hence, the mass fractions would be as follows:

Mixture Component	Value
air	1.0
fuelVapour	0.0
fuelLiquid	0.0

To set the initial mass fractions in STAR-CCM+, select the Constant node within the Species Mass Fraction manager. The Value for each respective mixture component can then be entered in the *Properties* window.



Alternatively, to enter the mass fractions in an array, click the ... (ellipsis) button, which opens the *Value* dialog.



Setting the STAR-CCM+ Boundary

The only available STAR-CCM+ boundary type in a WAVE coupled simulation is a mass flow inlet. The following list summarizes the data that is sent and received between STAR-CCM+ and WAVE:

- Density
- Temperature
- Pressure
- Mass Flux
- Mass Fractions

Most boundary conditions in a WAVE coupled simulation are governed by WAVE, except turbulence. Therefore, the boundary is set as coupled, with turbulence specification required.

Boundary Type

To set a mass flow inlet, select the boundary where coupling to WAVE is required.



The boundary type can be changed to Mass Flow Inlet in the *Properties* window, using the drop-down menu labeled *Type*.



Boundary Conditions

To set a boundary as coupled, select the 1D Coupling Specification node, under the Physics Conditions manager.



The boundary can then be set as coupled in the *Properties* window, by selecting Coupled from the *1D Coupling Specification* drop-down menu.



Coupling STAR-CCM+ to WAVE

The following section details the process for coupling to WAVE.

- Defining the WAVE Case File
- Starting the WAVE Solver
- Assigning the WAVE Species
- Assigning the WAVE Zones

With the above points addressed as necessary, a solution can be obtained.

Defining the WAVE Case File

STAR-CCM+ requires you to identify the WAVE case file that is used in the coupled simulation.

To define the WAVE case file, select the WAVE node, which is located in the feature tree under the Co-Simulations node.



The Directory Name and Case Name are then identified in the *Properties* window.



Starting the WAVE Solver

Once the case file is correctly defined, the STAR-CCM+ simulation can be coupled to WAVE. When coupled, the species and zones are read from the WAVE case file and added under the managers.

The species can then be assigned to the previously defined mixture components and passive scalars. The zones can also be assigned to boundaries, with the correct boundary type and conditions.

To start WAVE coupling, right-click the WAVE node, and select Start WAVE.



When complete, the *Output* window displays, "*** WAVE STARTED ***".

This method is only required the first time WAVE is coupled, and the species and zones are correctly assigned. There after, clicking Run (Ctrl+R) in the toolbar will start WAVE coupling and begin obtaining a solution.

The WAVE solver can be stopped at any time by clicking Solution > Clear Solution from the menu bar, which will break the connection between STAR-CCM+ and WAVE. WAVE can then be restarted using the methods that are mentioned above, which will reset the coupling and take into account any changes that are made to the WAVE case.

Assigning the WAVE Species

All species in WAVE coupling must be assigned to previously defined mixture components or passive scalars. This provides a means of representing the gas components in the STAR-CCM+ simulation. The WAVE species will only be available for assignment when at least one STAR-CCM+ boundary is defined as coupled to WAVE.

To assign a WAVE species, select a Species from the WAVE Species manager.



The species can then be assigned to a mixture component or passive scalar in the *Properties* window, using the *Mixture Component / Passive Scalar* drop-down menu.



Assigning the WAVE Zones

All zones in WAVE coupling must be assigned to a STAR-CCM+ boundary, with the correct boundary type and conditions. This provides a means of identifying the location of data exchange between the one-dimensional and three-dimensional code.

To assign a WAVE zone, select a Zone from the WAVE Zones manager.



The zone can then be assigned to a boundary in the *Properties* window, using the drop-down menu labeled *Boundary*.



RELAP Coupling

The RELAP Coupling feature provides the ability for two-way transfer of data between three-dimensional STAR-CCM+, and one-dimensional RELAP. For example, the natural convection currents in coolant fluid can be analyzed in detail in STAR-CCM+, with RELAP providing the boundary conditions.

This provides increased detail in specified areas of the RELAP case, by adding a three-dimensional element to a one-dimensional simulation. Similarly, processing times can be reduced in a STAR-CCM+ simulation, by utilizing the speed of one-dimensional codes.

The following section provides the definition of common terms:

- RELAP Coupling
- RELAP Zones or Flow Zones

The documentation also includes a guide to using RELAP Coupling, which details:

- Preparing the RELAP Case
- Setting Up the STAR-CCM+ Simulation
- Coupling STAR-CCM+ to RELAP

Note that RELAP is only available on Linux. STAR-CCM+ can run on Windows or Linux and connect to RELAP running on Linux.

What is RELAP Coupling?

RELAP coupling provides a means of data exchange between three-dimensional STAR-CCM+ and one-dimensional RELAP. This allows for a compromise between the speed of one-dimensional code, with the detail of three-dimensional code. The feature provides optimum results

when coupling occurs at a flow zone where the flow is expected to be largely one-dimensional. This recommendation does not apply to zones, as they apply thermal data to a wall boundary.

When running the solver, the flow or thermal data received from RELAP are applied at the coupled STAR-CCM+ boundary. If there is outflow at the boundary, the temperature is averaged at the flow zone, and sent back to RELAP. The STAR-CCM+ solver typically uses a smaller time step than RELAP. Data is exchanged at each RELAP time step, and the boundary conditions for STAR-CCM+ are interpolated between these time steps.

The following table summarizes how STAR-CCM+ and RELAP boundaries couple to one another:

STAR-CCM+	RELAP
Wall	Heat Structure
Mass Flow	Time Dependent Volume and Junction

The RELAP node, located within the Co-Simulations node, becomes available when the RELAP Coupling model is activated. It holds subnodes, which are populated with zones and flow zones, once the simulation is coupled to RELAP.



What is a RELAP Zone or Flow Zone?

A RELAP zone or flow zone represents the equivalent of a STAR-CCM+ boundary in RELAP. Each RELAP zone and flow zone must be assigned to a coupled boundary in the STAR-CCM+ simulation.

A RELAP zone is a specified wall boundary in the CFD region, which provides a means of exchanging thermal data between STAR-CCM+ and RELAP. Similarly, a RELAP flow zone is a mass flow boundary, which provides a means of exchanging flow data between STAR-CCM+ and RELAP. The physics values at the zones or flow zones are averaged, providing a single value used to determine the equivalent boundary conditions in the one-dimensional code.

RELAP zones transfer thermal data based upon the axial discretization in RELAP. This means that multiple constant values from RELAP are applied at the wall boundary. The direction and distribution of the one-dimensional discretization is determined by the Orientation property. The RELAP discretization on the coupled boundary can be visualized as a scalar displayer using the AvgZoneFlag field function.



RELAP flow zones transfer flow data, with a uniform value applied at the STAR-CCM+ boundary. This value can be multiplied by a flow multiplier to modify the received value. A field function or table can also be used to

multiply the received value by a variable scale factor. For example, this can be used with cell positions or wall distance to model the effect of a boundary layer at the flow boundary.



The RELAP Zones and RELAP Flow Zones nodes hold the individual zone nodes, which become available when RELAP Coupling is activated and started.



Preparing the RELAP Case

In order to couple RELAP to STAR-CCM+, you must create a `.control` file that specifies the port number and which parts of the RELAP model are to be coupled to STAR-CCM+. This file must have the same case name as the RELAP case input file and be located in the same directory. For example, if the RELAP case input file is `casename.i`, then the file required by STAR-CCM+ must be called `casename.i.control`.

The `.control` file must follow a certain format, with an example and explanation below:

```
PORT 5555
NUMCOUPLED,5,1,0
DEBUGOUTPUT
PRESBC TDV 980 JUN 979
FLUXBC HS 9000 SIDE 1
FLUXBC HS 9001 SIDE 1
TEMPBC HS 9700 SIDE 2
TEMPBC HS 9701 SIDE 2
```

`TEMPBC HS 9702 SIDE 2`	
`PORT 5555`	Set the port number for communication, which must match the RELAP node property.
`NUMCOUPLED,5,1,0`	Specifies the number of zones, followed by the number of flow zones. In the example, five thermal boundaries (zones) and one flow boundary (flow zone).
`DEBUGOUTPUT`	Makes RELAP write debug output when running
`PRESBC TDV 980 JUN 979`	Define a pressure boundary condition at a specified time dependant volume ID (e.g. 980) and a junction ID (e.g. 979). This will be represented by a flow zone in STAR-CCM+.
`FLUXBC HS 9000 SIDE 1`	Define a heat flux boundary condition at a specified heat structure ID (e.g. 9000) and side (e.g. 1). This will be represented by a zone in STAR-CCM+ that sends temperature data and receives heat flux.
`TEMPBC HS 9700 SIDE 2`	Define a temperature boundary condition at a specified heat structure ID (e.g. 9702) and side (e.g. 2). This will be represented by a zone in STAR-CCM+ that sends heat flux data and receives temperature.

Setting Up the STAR-CCM+ Simulation

The following section details the process of setting up a STAR-CCM+ simulation for coupling with RELAP:

- Activating the RELAP Coupling Model
- Setting the STAR-CCM+ Boundary

Activating the RELAP Coupling Model

The RELAP Coupling model is selected from within the Physics Model Selection dialog on the relevant continuum.

To select the RELAP Coupling model, the following models are required:

- Select Three Dimensional from the *Space* group box.
- Select Implicit Unsteady from the *Time* group box.
- Select any material from the *Material* group box, except Solid.
  - If you select Multi-Component Gas or Multi-Component Liquid, select Non-Reacting from the *Reaction Regime* group box.
  - If you select Eulerian Multiphase, select Volume of Fluid (VOF) from the *Multiphase Flow Model* group box.
- Select either Segregated or Coupled from the *Flow* group box, depending on the required solver type. The Segregated solver is recommended in most cases.
- Select Co-Simulation from the *Optional Physics Models* group box.
- Select RELAP Coupling from the *Co-Simulation* group box.

To fully define the physics continua used in a RELAP coupled simulation, the following additional models must also be activated:

- Either a Segregated Energy or Coupled Energy model from the *Optional Physics Models* group.
- Any model from the *Equation of State* group box.
- Turbulent from the *Viscous Regime* group.

- Any turbulence model from the *Reynolds-Averaged Turbulence* group, depending on simulation requirements.



Further models can be selected from the *Optional Physics Models* group, to meet individual simulation requirements.

Setting the STAR-CCM+ Boundary

The boundary type determines the physics values exchanged between STAR-CCM+ and RELAP. RELAP coupling is capable of transferring thermal data at walls and flow data at inlets and outlets.

The first step in setting the STAR-CCM+ boundary is to specify the boundary as coupled. For thermal data, you must specify whether the STAR-CCM+ wall boundary will send temperature and receive heat flux, or vice versa. For flow data, there is an option to multiply the received data by a scale factor.

1D Coupling Specification

The 1D Coupling Specification property determines whether the boundary in question is coupled to RELAP.

To set a boundary as coupled, select the 1D Coupling Specification node, under the Physics Conditions manager.



The boundary can then be set as coupled in the *Properties* window, by selecting Coupled from the *1D Coupling Specification* drop-down menu.



Coupling a STAR-CCM+ Wall Boundary to a RELAP Heat Structure

To transfer thermal data, a STAR-CCM+ wall must be assigned to a RELAP heat structure at a zone. There are then two options available:

- STAR-CCM+ receives temperature and sends heat flux:

- To perform this type of coupling, set the STAR-CCM+ boundary Physics Conditions > Thermal Specification option to Temperature, and specify the `FLUXBC` boundary condition in the `casename.i.control` file.



- STAR-CCM+ receives heat flux and sends temperature:

- To perform this type of coupling, set the STAR-CCM+ boundary Physics Conditions > Thermal Specification option to Heat Flux, and specify the `TEMPBC` boundary condition in the `casename.i.control` file.



Coupling a STAR-CCM+ Mass Flow Boundary to a RELAP Time Dependant Volume and Junction

To transfer flow data, a STAR-CCM+ mass flow boundary must be assigned to both a RELAP time dependant volume and a junction, at a flow zone.

A summary of the data transferred at a flow zone is included below:

- Pressure sent from STAR-CCM+ to the RELAP time dependant volume
- Mass flux sent from the RELAP junction to STAR-CCM+
- If there is inflow into STAR-CCM+, temperature is sent from the RELAP time dependant volume to STAR-CCM+
- If there is outflow from STAR-CCM+, temperature is sent from STAR-CCM+ to the RELAP time dependant volume

The majority of boundary conditions at a flow zone are governed by RELAP, except turbulence. Therefore, the boundary is set as coupled, with turbulence specification required.

Flow Multiplier

The Flow Multiplier allows you to multiply the data received from RELAP at the coupled boundary with a scale factor. For example, if the RELAP simulation modeled one tube, but the STAR-CCM+ simulation models two tubes, you could use a multiplier of 2 to double the received flow data at the coupled boundary.

To set a flow multiplier, select the Flow Multiplier node, under the Physics Values node. The *Method* menu allows you to select Constant, Field Function, Tabular or User Code input. Then, select the Flow Multiplier sub-node and enter the required input in the *Properties* window. The example below shows use of the Constant option to double the received data



The main use of this property is to take the one-dimensional, uniform flow data, and provide a parabolic profile at the coupled boundary, which takes the boundary layer at the wall into consideration. As an example, you can use a field function to multiply the data by a varying scale factor dependant on cell position.

The following field function uses the cell centroid in the radial direction, relative to a cylindrical coordinate system located at the centre of a pipe. This value is put to the power of 25 to give a parabolic shape, representative of an arbitrary boundary layer. The profile is then flipped and transposed so that the majority of the flow is given a flow multiplier of approximately 1.



This example profile is illustrated below by displaying the flow multiplier field function on the inlet of a cylindrical pipe:



A similar effect can also be achieved on a per-cell basis by using tabular data with reference to cell IDs.

Coupling STAR-CCM+ to RELAP

The following section details the process for coupling to RELAP.

- Starting the RELAP Solver

- Setting the Couple From Step Option
- Assigning the RELAP Zones and Flow Zones
- Setting RELAP Zone Properties

With the above points addressed as necessary, a solution can be obtained.

Starting the RELAP Solver

The following steps are required to start RELAP coupling:

- Set the port number
- Set the RELAP command
- Set the host name
- Set the case name
- Start and couple to RELAP

Setting the Port Number

The Port Number property must be set to the same port number specified in the `casename.i.control` file. This is the port at which RELAP and STAR-CCM+ will communicate.

To set the port number, select the Co-Simulations > RELAP node. Then, in the *Properties* window, set the Port Number.



Setting the RELAP command

The Relap Command property must be set to the command that is used to execute RELAP in a Linux shell. This can also be set as the absolute path to the RELAP executable, or as a relative path from the current working directory.

To set the RELAP command, select the Co-Simulations > RELAP node. Then, in the *Properties* window, set the Relap Command.

Setting the Host Name

The Host Name property must be set to the name of the machine where RELAP will be launched. If this is the same machine as the STAR-CCM+ client, the property should be set to `localhost`. This will allow STAR-CCM+ to launch RELAP automatically. If the RELAP machine is different to the STAR-CCM+ machine, the exact name of the machine on the network should be used. In this case, RELAP must be started on the host manually.

To set the host name, select the Co-Simulations > RELAP node. Then, in the *Properties* window, set the Host Name.



Setting the Case Name

The Case Name property must be set to the case name of the RELAP case file, (i.e. the same as `casename` in the `casename.i.control` file).

To set the case name, select the Co-Simulations > RELAP node. Then, in the *Properties* window, set the Case Name.



Starting and Coupling to RELAP

Once the simulation is fully defined, the STAR-CCM+ simulation can be coupled to RELAP. When coupling is started, RELAP reads the `casename.i.control` file and sends the zones and flow zones to STAR-CCM+. The zones and flow zones can then be assigned to coupled boundaries.

To start and couple to RELAP, right-click the Co-Simulations > RELAP node and select Start RELAP.



When complete, the *Output* window will display `RELAP Connection Initialized`. There will also be a summary of the RELAP zones added to the simulation.

RELAP can be stopped at any time by clicking Solution > Clear Solution from the menu bar, which will break the connection between STAR-CCM+ and RELAP. RELAP can then be restarted using the methods mentioned above, which will reset the coupling and take into account any changes made to the RELAP case. Closing STAR-CCM+ will also close RELAP.

Setting the Couple From Step Option

The Couple From Step option allows the transfer of data from STAR-CCM+ to RELAP to be delayed until a defined time step. This allows the STAR-CCM+ solver to reach a level of stability whilst receiving data from RELAP, before sending data back to RELAP. Note that, although data is not transferred to RELAP, the codes are both advancing at the same time. During this period of one-way transfer, the RELAP analysis is running with constant boundary conditions at the coupled boundary.

Using the Couple From Step option is not typically recommended. It is only required if the solver diverges during the initial stages, and you are unable to obtain a solution.

To set the Couple From Step option, select the Co-Simulations > RELAP node. Then, in the *Properties* window, set Couple From Step.

Assigning the RELAP Zones and Flow Zones

All zones and flow zones in RELAP coupling must be assigned to a STAR-CCM+ boundary, with the correct boundary type and conditions. This provides a means of identifying the location of data exchange between the one-dimensional and three-dimensional code.

The node names are based upon data received from RELAP. Zones are named "Heat Structure ID : Side ID" (e.g. 9000:1) and flow zones are named "Time Dependant Volume ID : Junction ID" (e.g. 980:979). This allows you to identify the part of the RELAP model to which the node refers, so that the appropriate boundary can be assigned.

To assign a zone or flow zone, select a sub-node of the RELAP Zones or RELAP Flow Zones manager.



In the *Properties* window, assign the boundary using the *Boundary* drop-down menu.



For comparison, the image below shows boundary assignment in the RELAP Flow Zone properties window.



Setting RELAP Zone Properties

The RELAP zone properties provide two user-defined settings:

- Orientation
- Face Area is Relative

Orientation

The Orientation property determines the direction of the axial discretization from RELAP. The axial discretization allows STAR-CCM+ to receive a different temperature or heat flux value from each part of the discretization in RELAP, which provides a gradient along the wall in the specified direction. For example, the following image shows a temperature distribution on the coupled boundary as a result of the RELAP discretization.



To set the orientation of the axial discretization, select a RELAP Zone node, and set the Orientation with the x, y and z components of a direction vector.



Face Area Is Relative

In some cases, the surface area of the STAR-CCM+ wall boundary may not match the area of the RELAP heat structure. In this instance, you should activate the Face Area is Relative option so that the STAR-CCM+ wall boundary is considered relative to the RELAP heat structure. This will apply the necessary scale factor to the values received at the coupled boundary.

To activate the Face Area is Relative option, select a RELAP Zone, and activate Face Area is Relative.



AMESim Coupling

The AMESim Coupling feature makes it possible for STAR-CCM+, (a three-dimensional code), to communicate with AMESim, (a one-dimensional code), in order to simulate coupled systems. For example, the displacement of a check valve due to fluid forces acting on it can be calculated by AMESim, while STAR-CCM+ calculates the force acting on the surface of the valve.

AMESim is a simulation software created by LMS Imagine, and is used for modelling and analyzing multi-domain systems. For more information on AMESim, visit the official website.

Data is transferred from AMESim to STAR-CCM+ using a TCP element in the AMESim system. In STAR-CCM+, data is transferred between AMESim and STAR-CCM+ through zones. There are two types of zone:

- AMESim Zones—these are pre-defined zones that exchange the most common type of data across boundaries: pressure and mass flux. The flow in these zones is expected to be one-dimensional.
- AMESim Field Function Zones—these zones allow you to send any scalar quantity to AMESim, such as the integrated force acting on a valve. These scalar quantities come from field functions, which can be either user-defined or system field functions.

Coupling STAR-CCM+ with AMESim is a two stage process. The first stage involves adding a TCP element to the AMESim system and setting it up to act as a server to allow co-simulation. You then define the simulation parameters and run the AMESim solver. The second stage involves setting up STAR-CCM+ to couple with AMESim.

Note: Currently, STAR-CCM+ supports coupling with AMESim Rev 12.

This section contains the following topics:

- Setting Up the AMESim System
- Setting Up the STAR-CCM+ Simulation

Setting Up the AMESim System

To couple AMESim with STAR-CCM+, you add a TCP element (the `dynamic_cosim_socket` element) to the AMESim system. You then configure this TCP element to act as a server. After configuring the simulation parameters, you run the AMESim solver. The AMESim solver must be already running for STAR-CCM+ to couple with it. When an AMESim system contains the `dynamic_cosim_socket` element, the AMESim solver waits for the STAR-CCM+ solver to run before doing any of its own calculations.

This section contains the following topics:

- The AMESim TCP Element
- Configuring AMESim to Run as a Server
- Setting AMESim Simulation Parameters
- Starting the AMESim Solver

The AMESim TCP Element

AMESim communicates with STAR-CCM+ through TCP/IP sockets using a TCP element (the `dynamic_cosim_socket` element).



This element contains input pins and output pins. When adding this element to an AMESim system, you choose how many input and output pins to create.



The pins that are on the left side of the element correspond to the number of data inputs into STAR-CCM+ from AMESim. The pins that are on the right side of the element correspond to the number of data outputs from STAR-CCM+ to AMESim. For example, if you want to transfer two scalar values from STAR-CCM+ to AMESim, and transfer three scalar values from AMESim to STAR-CCM+, you require a total of five pins on the TCP element: three pins on the left side and two pins on the right side.

Note: In AMESim, the data coming into and going out of the TCP element is a dimensionless signal.

AMESim TCP Pin Field Functions

When STAR-CCM+ couples with AMESim, two field functions are generated for each TCP pin. These field functions are stored in the Tools > Field Functions node. The `AMESimTransfer_x` field functions correspond to data from the current time-step, while the `AMESimTransfer_oldx` field functions correspond to data from the previous time-step.



Configuring AMESim to Run as a Server

For STAR-CCM+ to be able to communicate with AMESim, set the TCP element to act as a server. When the TCP element is set as a server, you can set the time-step for the AMESim solver.

To configure the TCP element:

1. Click to switch to parameter mode. The appropriate license for your AMESim system is checked out.
2. Double-click the `dynamic_cosim_socket` element. The *Change Parameters* dialog appears.
3. Set mode to server. A new property, sample time, appears in the dialog.
4. Set sample time to the desired time-step for your case.
5. Click Close.

Setting AMESim Simulation Parameters

After configuring the TCP element to act as a server, you can modify the run parameters for the AMESim solver. You can adjust the time-step and total time of the simulation. It is good practice to set the total run time of the AMESim solver to a larger value than the total run time set in STAR-CCM+. When coupled to AMESim, STAR-CCM+ keeps two steps of AMESim data and interpolates between them.

To modify the run parameters of the AMESim solver:

1. Click to switch to simulation mode.
2. Click to set the run parameters. The *Run Parameters* dialog appears.
3. Set Start time to correspond with the coupling time in STAR-CCM+. See Setting Up the STAR-CCM+ Simulation Parameters.
4. Set Final time to a value greater than the total simulation time in STAR-CCM+. See Setting Up the STAR-CCM+ Simulation Parameters. This action is required to make sure that AMESim does not finish solving before STAR-CCM+, due to numerical precision issues.
5. Set Print interval to match the time-step you set in the TCP element. See Configuring AMESim to Run as a Server. This property is used for post-processing in AMESim.
6. Click OK.

Starting the AMESim Solver

Once you have set up the TCP element to run as a server, with the correct run parameters, you can start the AMESim solver. When an AMESim system contains a TCP element, the AMESim solver waits until the STAR-CCM+ solver runs before it begins solving the model.

To start the AMESim solver:

- Click (Start a simulation).

Setting Up the STAR-CCM+ Simulation

Before coupling STAR-CCM+ with AMESim, decide which method you want to use for exchanging data between the two codes.

AMESim Zones allow you to exchange common types of data between STAR-CCM+ and AMESim on a limited number of boundary types. AMESim Zones are essentially a short-cut method for coupling STAR-CCM+ with AMESim.

AMESim Field Function Zones allow you to send any scalar field function data from STAR-CCM+ to a specific AMESim TCP pin. By using this method, you can also send data using reports.

In addition to these zone types, STAR-CCM+ generates two field functions for each pin on the AMESim TCP element. See The AMESim TCP Element. You can use these field functions as boundary conditions on any boundary in STAR-CCM+.

Suggested Workflow

To couple STAR-CCM+ with AMESim, follow the suggested workflow outlined below:

1. Activate the AMESim Coupling model. See Activating the AMESim Coupling Model
2. Couple the boundaries:
   - For coupling boundaries with AMESim Zones, see Coupling Boundaries Using AMESim Zones
3. Create the zones:
   - For AMESim Zones, see AMESim Zones
   - For AMESim Field Function Zones, see AMESim Field Function Zones

4. Couple STAR-CCM+ to AMESim. See Coupling STAR-CCM+ to AMESim
5. Couple STAR-CCM+ boundaries with AMESim TCP pin field functions. See Coupling Boundaries Using AMESim Pin Field Functions
6. Set the STAR-CCM+ simulation parameters. See Setting STAR-CCM+ Simulation Parameters

Activating the AMESim Coupling Model

The AMESim Coupling model in STAR-CCM+ allows you to exchange data between AMESim and STAR-CCM+. To use this model, first select the Co-Simulation model. The AMESim coupling model is only available for transient simulations.

To activate the AMESim coupling model:

1. Open the Physics Model Selection dialog. See Using the Physics Model Selection Dialog.
2. Select the models that are suitable for your case. The simulation must be transient.
3. Select Co-Simulation from the *Optional Models* group.

4. Select AMESim Coupling from the *Co-Simulation* group.



5. Click Close.
6. Set up the STAR-CCM+ Simulation parameters.
7. Couple STAR-CCM+ with AMESim. See Coupling STAR-CCM+ to AMESim.

AMESim Zones

AMESim Zones are boundaries in the CFD region where data is transferred between AMESim and STAR-CCM+. The following table demonstrates what data can be transferred between STAR-CCM+ and AMESim.

Boundary	Data Transferred	Direction
Mass Flow Inlet	Mass Flux	AMESim to STAR-CCM+
	Pressure	STAR-CCM+ to AMESim
Velocity Inlet	Mass Flux	AMESim to STAR-CCM+
	Pressure	STAR-CCM+ to AMESim
Pressure	Mass Flux	STAR-CCM+ to AMESim
	Pressure	AMESim to STAR-CCM+

If you want to exchange any other type of data between the two codes, or if you want to exchange data on other types of boundaries, use AMESim FF Zones. See AMESim Field Function Zones.

Creating AMESim Zones

You can only create AMESim zones in STAR-CCM+ after activating the AMESim Coupling model in STAR-CCM+. See Activating the AMESim Coupling Model.

To create an AMESim zone:

1. Right-click the Co-Simulations > AMESim > AMESim Zones node and select Add New AMESim Zone.



A new AMESim zone appears in the STAR-CCM+ object tree. You can rename new AMESim zones. See Renaming Objects.

2. Set up the AMESim Zone. See AMESim Zone Properties Reference.

An example of an AMESim Zone coupled to a Pressure Outlet boundary is shown in the following image.



AMESim Zone Properties Reference

Each AMESim Zone that you create has the following properties:

AMESim Zone Type (read only)	The type of zone, as defined in AMESim. A value of 1 indicates that the zone is an inlet boundary. A value of 2 indicates that the zone is a pressure outlet boundary.
Pressure Direction (read only)	Indicates in which direction Pressure data is being transferred. This property is automatically updated when you associate a boundary with this AMESim Zone.
Mass Flux direction (read only)	Indicates in which direction Mass Flux data is being transferred. This property is automatically updated when you associate a boundary with this AMESim Zone.
AMESim Name (read only)	The name of the zone as it appears in AMESim.

Boundary	Specifies which STAR-CCM+ boundary is associated with this zone. Only boundaries that are marked as coupled appear here. See Marking Boundaries for Coupling.
AMESim I/O number, Pressure	The pin number (in AMESim) that is associated with pressure data.
AMESim I/O number, Mass	The pin number (in AMESim) that is associated with mass flux data.

AMESim Field Function Zones

AMESim Field Function Zones determine the data that STAR-CCM+ sends to particular AMESim TCP pin. In these zones, you can associate a specific AMESim TCP pin number with any scalar field function in STAR-CCM+.

In STAR-CCM+, all reports automatically generate a field function that returns the value of the report on each iteration of the simulation. This facility means that in addition to the default field functions, you can send report data to AMESim, such as the sum of pressure and shear forces on a surface. Create a report with the correct force, and use the associated field function to send the data to the desired AMESim Pin.

Creating AMESim Field Function Zones

You can only create AMESim Field Function Zones in STAR-CCM+ after activating the AMESim Coupling model in STAR-CCM+. See Activating the AMESim Coupling Model.

To create an AMESim FF Zone:

1. Right-click the Co-Simulations > AMESim > AMESim FF Zones node and select Add New AMESim Zone.



A new AMESim zone appears in the STAR-CCM+ object tree. You can rename these new zones. See Renaming Objects.



2. Set up the AMESim Field Function Zone. See AMESim Field Function Zone Properties Reference.

AMESim Field Function Zone Properties Reference

Each AMESim Field Function Zone that you create has the following properties:

AMESim I/O number	The pin number (in AMESim) that you want to send field function data to.
Scalar Field Function	The field function that you want to associate with the AMESim Field Function zone.

Setting Up STAR-CCM+ Boundaries for Coupling

There are two methods that you can use to send AMESim data to STAR-CCM+ boundaries from AMESim:

- Using AMESim Zones
- Using AMESim Pin Field Functions

Coupling Boundaries Using AMESim Zones

When using AMESim Zones to couple STAR-CCM+ with AMESim, you must mark which of the supported boundaries you want to couple to.

To mark a boundary as a coupled boundary:

1. Navigate to the boundary you want to couple. Only Velocity Inlet, Mass Flow Inlet, and Pressure Outlet boundaries support this method of coupling.
2. Select Physics Conditions > 1D Coupling Specification and set 1D Coupling Specification to Coupled.



The selected boundary is marked as a coupling boundary and you can use it with an AMESim Zone. See AMESim Zone Properties Reference. The physics values that are relevant to the boundary type will be exchanged between AMESim and STAR-CCM+ through the AMESim Zone. For a Pressure Outlet, for example, STAR-CCM+ send the mass flux data to AMESim, and AMESim sends the pressure data to STAR-CCM+.

Coupling Boundaries Using AMESim Pin Field Functions

STAR-CCM+ automatically creates field functions corresponding to the pins on the AMESim TCP element when you couple to and AMESim server. You can use these field functions to set values on boundaries. This method allows you to transfer any type of data from AMESim to the STAR-CCM+ boundary.

Note: For the AMESim TCP pins to appear in STAR-CCM+, STAR-CCM+ must be already coupled to AMESim. See Coupling STAR-CCM+ to AMESim.

To couple a boundary using the AMESim TCP pin field functions:

1. Navigate to the boundary you want to couple.

2. Select Physics Values > [scalar quantity] and set Method to Field Function.



3. Select [scalar quantity] > Field Function and set Scalar Function to the AMESim TCP pin field function that corresponds to the desired scalar quantity in AMESim.



When the coupled simulation runs, the boundary receives data from the selected AMESim TCP pin.

Coupling STAR-CCM+ to AMESim

The AMESim solver must be already running before you can couple STAR-CCM+ to it.

Note: Some reports, for example Force reports that include Shear Force, are only available after the flow solution is stepped at least once. If such reports are used in AMESim Field Function Zones, set the coupling to start from Step 1 instead of Step 0. See the Couple From Step option.

To set up the coupling parameters:

1. Click the Co-Simulations > AMESim node.



2. Configure the coupling properties. See AMESim Coupling Properties Reference.
3. Right-click the Co-Simulations > AMESim node and select Start AMESim.



STAR-CCM+ couples to AMESim. A message, summarizing the coupling procedure appears in the *Output* window.



AMESim Coupling Properties Reference

The following AMESim coupling parameters are available in the AMESim node.

AMESim Host	Specifies the name of the machine where AMESim is running. If AMESim is on the same machine as the STAR-CCM+ installation, you can leave this property as `localhost`.
AMESim Port	Specifies the port to use when communicating with AMESim. The default port for AMESim co-simulation is `60000`.
AMESim Pins	Specifies the total number of pins on the `dynamic_cosim_socket` element. STAR-CCM+ determines how many TCP pins are inputs and how many are outputs.
Couple From Step	Specifies the time-step from which STAR-CCM+ starts sending data to AMESim. This option delays the transfer of data to the AMESim solver until a defined time-step. This feature allows the STAR-CCM+ solver to reach a level of stability while receiving data from AMESim, before sending data back to AMESim. See Coupling STAR-CCM+ to AMESim.

Setting STAR-CCM+ Simulation Parameters

To set the time-step in STAR-CCM+, see Setting the Time-Step For Implicit Unsteady Simulations.

It is also good practice to set the total run time in AMESim to a larger value than the value set in STAR-CCM+. To set the total run time in STAR-CCM+, see Setting the Maximum Physical Time.

Reacting Channel Coupling

The Reacting Channel Coupling feature provides the ability for exchange of data between three-dimensional STAR-CCM+ and the one-dimensional Plug Flow Reactor in Dars-CFD.

This section describes:

- What is the Reacting Channel Coupling Model?
  - What are Reacting Channels?
  - Activating the Reacting Channel Coupling Model
  - Setting up Reacting Channels
- Reacting Channel Coupling Formulation

What is the Reacting Channel Coupling Model?

You can use the Reacting Channel Coupling model in STAR-CCM+ to solve reacting flow simulations in tubular heat exchangers with channels that are long and thin. When the Reacting Channel Coupling model is selected along with the Co-Simulation model, a Co-simulation > Reacting Channel CoSimulation node appears.

Several geometrical channels can represent one reacting channel as long as all of the geometrical channels have identical geometries of the same orientation. For geometrical channels that differ in length, diameter, or orientation, define these as a separate reacting channel.



This functionality is useful for simulating chemical processes such as steam re-forming and firebox combustion.

STAR-CCM+ meshes and solves the volume of the outer flow. The reacting channels are solved using a Plug Flow Reactor (PFR) approximation. The PFR is part of the Dars-CFD package. The reacting channels are not meshed and are considered 1D. However, the outer flow volume must account for the presence of the reacting channels in its geometry.

The coupling happens on the reacting channel wall boundary between the outer flow in STAR-CCM+ and the reacting channel in the PFR. The temperature is exchanged from the outer flow to the reacting channel and heat flux is returned back from the reacting channel to the outer flow. Since the reacting channels are not meshed, the exchange of data occurs at axial divisions, whose frequency you specify, along the channel.

What are Reacting Channels?

Although reacting channels appear as 3D features in a geometry, they become 1D representations when defined as reacting channels in the Reacting Channel Coupling model.

Since the reacting channels are one-dimensional and have no mesh, data is exchanged between axial points along a vector through the centre of the reacting channel and the surface boundary of the outer flow, as represented by the diagram below:



Note: It is imperative that all geometrical channels that are defined within a single reacting channel, are the same length and diameter with the same properties.

Reactions that occur in the reacting channel can consume or create heat energy. At specified time intervals, STAR-CCM+ sends the temperature along the geometrical channel wall to the Plug Flow Reactor (PFR). The PFR then solves for the reacting flow, determines the temperature within the reacting flow, and calculates the heat flux at each axial point in the reacting flow channel. The PFR then sends this data back to the outer flow in STAR-CCM+ . See Eqn. (2686).

See also:

Reacting Channel Coupling Formulation

What is the Reacting Channel Coupling Model?

Activating the Reacting Channel Coupling Model

The Reacting Channel Coupling model in STAR-CCM+ allows you to exchange data between reacting channels that use the Dars-CFD one-dimensional Plug Flow Reactor and three-dimensional STAR-CCM+ .

To activate the Reacting Channel Coupling model:

1. Open the Physics Model Selection Dialog. See Using the Physics Model Selection Dialog.

2. Select the specific models that you require for the simulation, including the following models:

Group Box	Model
Space	Three Dimensional
	Gradients (Selected Automatically)
Optional Models	Co-Simulation
Co-Simulation	Reacting Channel Coupling

3. Click Close.

Setting up Reacting Channels

When using the reacting channel coupling model, set up at least one reacting channel. Each reacting channel can represent several geometrical channels.



To define a reacting channel:

1. Expand the Co-Simulations > Reacting Channel CoSimulation node
2. Right-click the Reacting Channels node
3. Select New Reacting Channel

A Reacting Channel 1 node appears.



4. Specify properties for the Reacting Channel sub-nodes:
   - Reacting Channel Inlet
   - Reacting Channel Wall
   - Reacting Channel Outlet
   - Reacting Channel Solver Settings
   - Reacting Channel Geometries
   - Reacting Channel XY Plots
5. Import the Reacting Channel mechanism.

See also:

What is the Reacting Channel Coupling Model?

Reacting Channel Inlet

The Reacting Channel Inlet node contains a set of properties and a sub-node, Inlet Viscosity, which allow you to define the parameters of the reacting channel inlet.



Reacting Channel Inlet Properties

Initial Mass Fraction	This array is available to populate after the mechanism is imported. Allows you to specify the composition of the initial mass fraction at the inlet. See Setting Up a Complex Chemistry Definition in Dars-CFD Format for Reacting Channel Coupling.
Inlet Area	Specifies the surface area of the inlet for each individual geometrical channel that is part of the same reacting channel.
Inlet Temperature	Specifies the temperature of the initial reacting flow at the inlet.
Inlet Velocity	Specifies the velocity of the initial reacting flow at the inlet.
Inlet Pressure	Specifies the pressure of the initial reacting flow at the inlet.

Inlet Viscosity Properties

Method	Allows you to select a method for specifying viscosity.	
	Calculate	STAR-CCM+ calculates the viscosity.
	Constant	Defines the Method as a single constant for which you can specify a Value.

Reacting Channel Wall

The Reacting Channel Wall node contains a set of properties and three sub-nodes which allow you to define the parameters of the reacting channel

wall.



Reacting Channel Wall Properties

Number of Axial Divisions	Specifies the number of axial points at regular intervals along each geometrical channel where data is exchanged from the Plug Flow Reactor.
Underrelaxation Factor	Represented in Eqn. (2686) as, a. Specify a value between one and zero. The lower the value, the less effect the most recently exchanged heat flux has on the solution at each time step.
Orientation	Defines the vector of the reacting channel orientation through which the axial points are determined.
Pipe Length	Specify the length of each geometrical channel in the reacting channel. All geometrical channels in the same reacting channel must be this length.

Wall area — Represented in Eqn. (2681) as, A_c. Specifies the surface area of each geometrical channel, determined automatically from the area of one representative geometrical channel that this reacting channel is associated with.

Heat Transfer Coefficient

Method — Allows you to specify the method for determining the heat transfer coefficient, h_f, that is used in Eqn. (2682).

- Calculate Nonuniform — Determines the heat transfer coefficient using data from each axial point.
- Calculate Uniform — Determines the heat transfer coefficient using values for all parameters that are taken directly from the reacting channel inlet.
- Constant — Allows you to specify the heat transfer coefficient as a constant array value.
- Tabular — Allows you to specify the heat transfer coefficient using tabular data that you can import.

Wall Temperature

Method — Allows you to specify how temperature, T_W, is determined at the wall boundary.

- Nonuniform — Determines the temperature using data from each axial point.
- Uniform — Determines the temperature using data that is averaged from all axial points.

Pipe Friction Coefficient

Method	Allows you to specify a multiplying factor for the pipe friction, f.	
	Constant	Specifies a constant value for the pipe friction coefficient.
	Tabular	Allows you to specify the pipe friction coefficient using tabular data that you can import.

Reacting Channel Outlet

The Reacting Channel Outlet node represents the end of the channel.



Reacting Channel Outlet Properties

Outlet area — Specifies the surface area of the outlet which is automatically determined from the Inlet Area value, A_c, that is set for the Reacting Channel Inlet. Read only.

Reacting Channel Solver Settings

The Reacting Channel Solver Settings node contains a set of properties which allow you to control the solution method that is used by the Plug Flow Reactor.



The Plug Flow Reactor (PFR) employs an adaptive time step scheme which adjusts the time step size according to the convergence of the previous step. To control the accuracy of the solution, improve robustness, and efficiency, you can set limits for the Minimum Time Step Size and Maximum Time Step Size.

Reacting Channel Solver Settings Properties

Absolute Tolerance	Absolute tolerance for the Dars-CFD ODE solver.
Relative Tolerance	Relative tolerance for the Dars-CFD ODE solver.
Minimum Time Step Size	Defines a limit for the minimum time step size.
Maximum Time Step Size	Defines a limit for the maximum time step size.
Limit Rates	Allows you to limit the reaction rates to make sure that a physical solution is achieved. Activated — Reaction rates are limited. Deactivated — Reaction rates are not limited.

Reacting Channel Solver Settings Expert Properties

Maximum BDF Order	The number of previous time steps that are considered when calculating the predictor for the next step.
Maximum Newton Steps	Defines how many times a gradient is re-used before performing a new full iteration step.

Reacting Channel Geometries

The Reacting Channel Geometries node contains a set of properties which allow you to define the geometrical channels that are part of the reacting

channel.



Reacting Channel Geometries Properties

Reacting channels	Select the boundaries of the geometrical channels that represent the reacting channel. The boundaries must be from long, thin, cylindrical channels that are geometrically identical. Sub-nodes appear for each of the boundaries that are selected.

Reacting Channel XY Plots

The Reacting Channel XY Plots node allows you to create XY plots for a variety of variables that are involved in the simulation, including:

- Temperature
- Wall Temperature
- Heat Flux
- Density
- Velocity
- Heat Transfer Coefficient



When there are chemical species defined in the initial mass fraction, each species and a variety of possible products are also available to plot.

The Reacting Channel XY Plots node has no properties, although each [XY Plot] sub-node contains a set of properties and further sub-nodes which allow you to define the parameters for each XY Plot.



XY Plot Properties

Title	Defines the title of the XY plot.
Variable Name	Displays the title of the variable that is selected for the XY plot.
X-Axis Variable	Defines the variable that is displayed on the X-axis.

XY Plot Expert Properties

See XY Plot Expert Properties

Setting Up a Complex Chemistry Definition in Dars-CFD Format for Reacting Channel Coupling

The Dars-CFD library requires a specific form of input files for the complex chemistry definition. To create these input files, the DARS tool must be available.

Note: For reacting channels, *f90 files are required regardless of if run time compile is activated or not.

Two ASCII input files are needed for Dars-CFD, `InputRedKinMec.txt` and `InputRedKinTherm.txt`.

Importing the Complex Chemistry Definition

To import the complex chemistry definition in Dars-CFD Format when using the Reacting Channel Coupling model:

1. Expand the Co-Simulation > Reacting Channel CoSimulation > Reacting Channels node.
2. Right-click the [Reacting Channel] node for which you are importing the complex chemistry definition.

3. Select Import Complex Chemistry Definition (Dars-CFD Format).



4. Locate and select the input files, `InputRedKinMec.txt` and `InputRedKinTherm.txt`.

To replace the definition:

1. From the pop-up menu of the [Reacting Channel] node, select Delete Complex Chemistry Case... > Dars-CFD Reaction Model node.
2. Repeat the import step.

If you try importing without deleting, the following error message appears:



Reacting Channel Coupling Formulation

The Reacting Channel Coupling feature couples a 3D STAR-CCM+ simulation with the 1D Plug Flow Reactor in Dars-CFD.

Reacting Channel Plug Flow Reactor Formulation

The flow in the reacting channel is simulated using the one-dimensional Plug Flow Reactor (PFR) from Dars-CFD. At any axial point within the reacting channel, the flow variables are radially uniform. For a given reaction, the axial diffusion of heat and species are consistent and are therefore ignored. The evolution of species and temperatures in the PFR are represented as follows:

- Species:

$$\frac{dY_k}{dt} = \frac{\dot{\omega}_k}{\rho} \tag{2679}$$

- Temperature:

$$\frac{dT}{dt} = \frac{\sum_k \dot{\omega} h_k + \dot{Q}_C}{\rho c_p} \tag{2680}$$

where,

$\dot{\omega}$ = net source term, $(\mathrm{kg/m^3\text{-}s})$

Y_k = k^{th} species mass fraction

W_k = k^{th} molecular weight, $(\mathrm{kg/kmol})$

$\dot{\omega}_k$ = k^{th} net molar reaction rate, $(\mathrm{kg/m^3\text{-}s})$

h_k = k^{th} total enthalpy, $(\mathrm{J/kg})$

ρ = mixture density, $(\mathrm{kg/m^3})$

T = reacting channel bulk temperature, (K), as specified in the Reacting Channel Inlet Properties

c_p = specific heat, $(\mathrm{J/kg\text{-}K})$

$\dot{Q}_C$ = convective heat transfer source, $(\mathrm{W/m^2})$

The convective heat transfer source, $\dot{Q}_C$, in the reacting channel is determined using:

$$\dot{Q}_C = h_f(T_W - T) \tag{2681}$$

where,

T = reacting channel bulk temperature, (K)

T_W = reacting channel wall temperature which is averaged from the 3D outer flow temperature field on the resolved reacting channel wall, (K).

h_f = convective heat transfer coefficient of the reacting channel

The convective heat transfer coefficient of the reacting channel, h_f, is determined using:

$$h_f = \frac{NuK_C}{D_C} \tag{2682}$$

where,

K_c = gas phase thermal conductivity, (W/m -K)

D_c = channel diameter, (m)

Nu = Nusselt number

The Nusselt number, Nu, is determined according to fully developed pipe flow empirical correlations:

- For laminar flow where Re < 3000:

$$Nu = 3.66 \tag{2683}$$

- For turbulent flow where Re > 3000, Nu is determined using Gneilinski's correlation for turbulent flow:

$$Nu = \frac{f(\mathrm{Re}\text{-}1000)\mathrm{Pr}}{1 + 12.7\sqrt{f}(\mathrm{Pr}^{2/3} - 1)} \tag{2684}$$

where:

f = pipe friction

For smooth, cylindrical, reacting channels, f, is determined using:

$$f = \frac{0.125}{(0.79\log(\mathrm{Re}) - 1.64)^2} \tag{2685}$$

where the pipe friction coefficient is specified in the reacting channel wall properties.

The equations in the PFR are solved using a stiff Ordinary Differential Equation (ODE) solver with time-steps that are based on the grid size of the reacting channel and the local reacting channel velocity.

Outer Flow STAR-CCM+ Formulation

The outer flow, which you can define as reacting, or non-reacting, is simulated using three-dimensional STAR-CCM+. You can mesh the three-dimensional outer flow geometry as long as the geometrical channels within the geometry that represent the reacting channel remain unmeshed.

The heat flux boundary condition at the reacting channel wall, $\dot{Q}_{3D}$, is calculated as an under-relaxation of the heat that the outer flow gains from the reacting channel. This heat flux is determined using:

$$\dot{Q}_{3D} = -a\dot{Q}_C - (1-a)\dot{Q}_C^o \tag{2686}$$

where,

$\dot{Q}_C$ = reacting channel heat flux, $(\mathrm{W/m^2})$

a = under-relaxation factor as specified in the reacting channel wall properties

$\dot{Q}_C^o$ = heat flux from previous iteration, $(\mathrm{W/m^2})$

When the solution for the simulation converges, the heat flux from the reacting channel is equal to the flux gain from the outer flow.

STAR-CCM+ to STAR-CCM+ Co-Simulation

STAR-CCM+ allows you to couple two running simulations together in a co-simulation analysis. The STAR-CCM+ to STAR-CCM+ co-simulation capability targets simulations involving conjugate heat transfer (CHT), fluid-structure interaction (FSI) or exchange of scalar/vector fields across a coupling boundary. The co-simulation is between 3D geometries that have one or more shared interfaces but different mesh densities and time scales. Thermal and mechanical fields can be exchanged across the fluid-solid interface; auxiliary fields can be exchanged across fluid-solid, fluid-fluid, or solid-solid interfaces. One instance of STAR-CCM+ solves the solid domain and the other solves the fluid domain. This technique is similar to the technique described in Explicit Coupling within STAR-CCM+ but without the requirement to manually create data mappers in each simulation or to create any Java macros that automate the exchange between two instances of STAR-CCM+.

Conjugate heat transfer simulations are commonly used to model both solid and fluid domains in a single instance of STAR-CCM+. In these simulations, however, it is not possible to account for the different time scales present in the solid and fluid domains. For scenarios where the time scales of the thermal conduction in the solid are much longer than the time scales in the fluid, STAR-CCM+ to STAR-CCM+ co-simulation offers a particular advantage. For example, the exhaust gases escape the tail pipe almost immediately after the engine is started, but the engine itself may take minutes to reach its maximum temperature.

For simulations with widely varying time scales, STAR-CCM+ to STAR-CCM+ co-simulation allows you to treat the fluid and solid domains separately. The fluid is treated in a steady state simulation and the solid in a transient simulation. The solid and fluid simulations must exchange data at a frequency that captures the slow change in the solid temperature.

A solid thermal simulation typically requires less iterations per time step to reach convergence than is required by the fluid simulation. By allocating a separate simulation to the solid thermal analysis, the compute time required to solve the overall solid-fluid analysis can be considerably smaller when the co-simulation technique is used. And although the computer CPU time needed to reach the final time in a STAR-CCM+ to STAR-CCM+ CHT co-simulation over a single STAR-CCM+ CHT simulation can be considerably smaller, with a properly configured co-simulation the same level of accuracy is achieved as in a fully-implicit conjugate simulation.

There are two types of co-simulation:

- One-way co-simulation
- Two-way co-simulation

This section describes how to use the co-simulation technique. Specifically, instructions are provided for the following:

- STAR-CCM+ Co-Simulation Workflow
- Activating the STAR-CCM+ Co-Simulation models
- Specifying the Co-Simulation global settings
- Setting up the coupling zones
- Specifying the co-simulation exchange conditions
- Running the Co-Simulation

What is One-Way STAR-CCM+ to STAR-CCM+ Co-Simulation?

The diagram below demonstrates a one-way STAR-CCM+ to STAR-CCM+ co-simulation process. In the diagram, the leading partner provides data to the lagging partner but does not receive any data from it. You can invert the relationship so the lagging partner provides information to the leading partner but not receive information from it. After the leading simulation has run for a specified number of iterations, a mapping operation occurs and then the data in the shared zone is transferred to the waiting lagging simulation. The lagging simulation uses the data as part of its solution and runs. However, the leading simulation does not use the data in the shared zone before recommencing. The process of transferring data to the lagging

partner simulation, running, and waiting continues until both solutions have satisfied their stopping criteria. The setup process is similar to two-way coupling, except deactivate thermal or mechanical import.



What is Two-Way STAR-CCM+ to STAR-CCM+ Co-Simulation?

The diagram below demonstrates the two-way STAR-CCM+ to STAR-CCM+ co-simulation process. After the leading simulation has run for a specified number of iterations, a mapping operation occurs and then the data in the shared zone is transferred to the waiting lagging simulation.

The lagging simulation uses the data as part of its solution and runs. The process of transferring data to the partner simulation and waiting continues until both solutions have satisfied their stopping criteria.



STAR-CCM+ Co-Simulation Workflow

For co-simulation cases, the following workflow is recommended:

1. Develop one STAR-CCM+ simulation for the solid region, and another STAR-CCM+ simulation for the fluid region, without any co-simulation options. Use provisional boundary conditions that are close to the expected values from the partner simulation.

   Run the models to convergence before attempting co-simulation.

2. Activate a STAR-CCM+ co-simulation model in each simulation.
3. Specify global co-simulation conditions:
   a. Specify a leading and a lagging simulation. The lead simulation runs first, connects, and determines the coupling time-steps and mapping mode.
   b. In each simulation, specify when the data mapping occurs.
   c. Specify the server identity of the lagging simulation in the Connect to Partner Simulation node of the co-simulation settings for the lead simulation.

- To select the partner co-simulation, click ... (Custom Editor). A selection dialog opens and displays a list of all currently active STAR-CCM+ simulations. Select the partner co-simulation.
- Note that the server list in the selection dialog contains only servers that are visible in the *Servers* window. To access the *Servers* window, select *Window* > *Servers* or click on the Servers window tab next to the simulation window tab.

4. Set up one or more co-simulation coupling zones. In each simulation, a coupling zone defines the coupled boundaries and the field import/export specification.
   a. Select the coupled boundaries in each simulation.
   b. Specify the field export and import conditions in each simulation.
   c. If the current simulation requires non-default control of the data mapping, create a Surface Data Mapper and select it for each of the mapper nodes under the zone Values.
5. Define the frequency of exchange so that the cumulative number of iterations / physical time is less than the maximum set in the stopping criteria.
6. Set stopping criteria so that each simulation provides the correct number of exchanges.
7. Run the leading simulation. The coupling surfaces are exchanged as an imported region in each simulation. At the end of each exchange the data is mapped automatically to the surface and then sent to the partner simulation.

   The simulations continue to run and send data until their individual stopping criteria are satisfied.

Note: If either stand-alone model is improperly defined or has numerical problems, you are unlikely to obtain a solution for a co-simulation.

Activating the STAR-CCM+ Co-Simulation Models

To enable coupling between two STAR-CCM+ simulations, activate a STAR-CCM+ co-simulation model in each partner simulation.

There are three co-simulation models for coupling STAR-CCM+ simulations:

- *STAR-CCM+ Steady Co-Simulation* model: for steady simulations, this model allows you to specify the number of iterations per data exchange.
- *STAR-CCM+ Explicit Co-Simulation* model: for transient simulations, this model allows data exchange at a coupling time-step that you specify.
- *STAR-CCM+ Implicit Co-Simulation* model: for transient simulations, this model allows to exchange data more than once per time step.

In general, for CHT and FSI problems, the *STAR-CCM+ Implicit Co-Simulation* model allows you to reach a stable solution with larger time-steps than those required by the *STAR-CCM+ Explicit Co-Simulation* model. Using the *STAR-CCM+ Implicit Co-Simulation* model is also recommended for FSI problems that require multiple data exchanges per time-step to reach a stable solution.

Activate a STAR-CCM+ co-simulation model in each partner simulation:

1. Expand the Continua node.
2. Right-click the relevant physics continuum node and select Select Models...
3. In the Physics Model Selection dialog:
   a. Select appropriate models from the *Space*, *Material*, and *Time* group boxes.
   b. Select Co-Simulation from the *Optional Models* group box.

c. Select a STAR-CCM+ co-simulation model of your choice.



When you activate a STAR-CCM+ co-simulation model, STAR-CCM+ adds a new physics continuum under the Continua node to represent the partner simulation. For the partner continuum, the External Application model is automatically selected.

The following image shows an example of model selection for two partner simulations coupled in a STAR-CCM+ to STAR-CCM+ co-simulation.



Once a STAR-CCM+ Co-Simulation model is activated, you can specify the co-simulation settings. Complete the co-simulation settings for both the leading and the lagging simulation.

Specifying Co-Simulation Conditions and Values

In STAR-CCM+, you specify the settings for co-simulation in a co-simulation session under the Co-Simulations node.

The Co-Simulations node becomes available in the simulation tree when you activate a STAR-CCM+ Co-Simulation model. This node is automatically populated with a co-simulation session that has the default name, STAR-CCM+ Co-Simulation 1. It is useful to rename the session using a name that is associated with the partner simulation in the co-simulation.

The [STAR-CCM+ Co-Simulation 1] node allows you to specify co-simulation global settings, as well as to define and set up co-simulation coupling zones.



In STAR-CCM+, a coupling zone defines a set of boundaries where fields are exchanged with the partner simulation. STAR-CCM+ Co-Simulation supports multiple coupling zones, that allow you to specify different conditions on different sets of coupled boundaries.

Co-simulation global settings, such as the method for connecting to the partner simulation and the concurrency mode, are general conditions and values that apply to all the zones defined in the simulation.

Instructions for specifying co-simulation global settings and zone settings are outlined in this section.

For the [STAR-CCM+ Co-Simulation 1] node, the following properties and pop-up menu items are available.

STAR-CCM+ Co-Simulation Properties

State	Displays the status of the Co-Simulation. (Read-only)
Co-Simulation Type	Displays the type of co-simulation (Read-only). For a STAR-CCM+ to STAR-CCM+ co-simulation, Co-Simulation Type is set to STAR-CCM+ Co-Simulation.

STAR-CCM+ Co-Simulation Pop-Up Menu

The following pop-up menu items are available when you right-click the [STAR-CCM+ Co-Simulation 1] node.

Disconnect from Partner Co-Simulation	Closes the connection between the simulations.

See also: common menu items.

Specifying Co-Simulation Global Settings

In STAR-CCM+, you set global co-simulation conditions under the [STAR-CCM+ Co-Simulation 1] > Conditions node.



In each partner simulation, specify the Concurrency Mode, the Connect Method, and the Mapping Mode. The available properties are listed below.

Concurrency Mode Properties

Method — Defines the initial roles of the partner simulations. The co-simulation is run from the leading simulation as the lagging simulation waits. When the exchange criteria have been reached the data is transferred, the leading simulation waits and the lagging simulation runs. This process continues until the stopping criteria have been satisfied.

- Lead — Defines the simulation that runs before transferring initial data. An empty child node, Lead, is added.
- Lag — Defines the simulation that waits for initial data before running. An empty child node, Lag, is added.

Connect Method Properties

Method	Specify the method to connect the simulations.	
	Connect to Partner Simulation	Connects to a running STAR-CCM+ server visible from the current client. To select the partner co-simulation, click the custom editor button ... of the Connect to Partner Simulation child node. Choose the partner co-simulation from the STAR-CCM+ servers list in the selection dialog.
	Launch Application and Connect	Launches an external application and connects to it, providing that the external application is associated with a STAR-CCM+ server (STAR-CCM+ will echo a string in the form `Server::start -host <host_name>:<port_number>`). Selecting this option when coupling partner code to STAR-CCM+ through the Co-Simulation API allows you to use a single shared license for both the STAR-CCM+ and the STAR-CCM+ API server. For more information, see Running a Co-Simulation with Partner Code. This option enables the child node, Launch Application and Connect. Use this node to specify the command to launch the external application.
	Launch STAR-CCM+ and Connect	Launches a new STAR-CCM+ session and loads the specified simulation. Use the Launch STAR-CCM+ and Connect child node to specify the Launch command by clicking the custom editor button ... and selecting the simulation file and run mode.
	Wait for Connection	Waits for a lead simulation to connect. An empty child node, Wait for Connection, is added.

Mapping Mode Properties

Method — Determines which simulation maps the co-simulated fields.

- **Map Before Export** — The current simulation maps the exported solution data. The fluid-only simulation maps its solution (heat transfer coefficient and reference temperature) onto the solid-only mesh before exchanging data with the solid-only simulation. The solid-only simulation maps its solution (wall temperature) onto the fluid-only mesh before exchanging data with the fluid-only simulation.
  A Map Before Export child node is added. Mapper controls for fields that are exported from this simulation are added to the Values node. This option is typically optimal when the solver is run in parallel. When this option is used, both of the co-simulations have a copy of both meshes.
  Both the leading and lagging simulation use this mode if it is selected for the leading simulation.
- **Map After Import** — The current simulation maps the imported solution data before the solver run starts. The fluid-only simulation sends its mesh and solution to the solid-only simulation, and then the solid simulation maps the fluid solution onto its mesh. Meanwhile, the solid-only simulation sends its mesh and solution to the fluid-only simulation, and the fluid simulation maps the solid solution onto its mesh.
  A Map After Import child node is added. Mapper controls for fields this simulation requires are added to the Values node.
  When this option is used, both of the co-simulations have a copy of both meshes.
  Both the leading and lagging simulation use this mode if it is selected for the lead.
  For STAR-CCM+ to STAR-CCM+ co-simulations, this option is effectively the same as the Map Before Export option.

Map Both Directions	The current simulation maps the imported solution data before the solver run starts and the exported solution data after the solver stops. The fluid-only simulation has both of the meshes, and maps its solution onto the solid mesh before exchanging data with the solid-only simulation. The solid-only simulation exchanges its solution data with the fluid-only simulation, which then maps the solid solution onto the fluid mesh. When this option is used, only one of the co-simulations (typically the fluid-only simulation) has a copy of both meshes. A Map Both Directions child node is added. Mapper controls for all values the co-simulation uses are added to the Values node. If the leading simulation uses this mode, set the lagging simulation to use Partner maps both directions. You can use this option when one of the simulations has many more CPUs than the other (that is, allow the simulation with the most processors to do all of the mapping work). You can also use this option when co-simulating with one-dimensional codes, or co-simulating with codes that do not have mapping capabilities.
Partner Maps Both Directions	The partner simulation maps the imported solution data before the solver run starts and the exported solution data after the solver completes. A Partner Maps Both Directions child node is added. If the leading simulation uses this mode, set the lagging simulation to use Map Both Directions. You can use this option when STAR-CCM+ is coupled to a third-party code (for example, Abaqus) and you think the third-party mapper is more accurate or efficient.

Specifying Launch Conditions for Batch Mode

If you launch a simulation using a batch queue and the command line option `-np N,M` (where `N` and `M` are the number of processes for the lead and lag simulations respectively), be sure to observe these conditions:

- The simulation that is submitted to the batch queue must have the Connect Method set to Launch STAR-CCM+ and Connect.
- The Launch STAR-CCM+ and Connect method must have its Launch command pointing to the second simulation on the batch queue file system.

- The STAR-CCM+ version used in the launch command must be the same as the one used for the first simulation.
- Depending on the user run environment, you must specify all the command line options that are required for running a regular STAR-CCM+ session in parallel, such as `-mpi` or `-ssh`.
- In case of trouble with this option, ensure beforehand that a standalone STAR-CCM+ session can be run in parallel on the same environment.

Example:

```
% starccm+ -np 16,8 model01.sim
```

This command launches the lead simulation in `model01.sim` with 16 processes in the lead simulation. The lag simulation is then launched on 8 processes, following which the connection between the two simulations is established.

Specifying Co-Simulation Zone Settings

In STAR-CCM+ to STAR-CCM+ co-simulation, a coupling zone allows you to specify the coupled boundaries and the fields that are exchanged with the partner simulation.

As STAR-CCM+ co-simulation supports multiple coupling zones, each zone is identified with an index and connects to the zone with the same index that is defined in the partner simulation. The number of coupling zones defined in a simulation must match the number of zones in the partner simulation.



In a zone, you specify coupled boundaries that exchange data with target boundaries in the equivalent zone of the partner simulation. The zone import/export field specification defines this exchange.

Note: A coupling zone connects to a zone in the partner simulation based on the index, not on the node name. Renaming a zone does not affect its index.

For a coupling zone, various import/export field conditions are available depending on the type of co-simulation. STAR-CCM+ offers a co-simulation framework for thermal and mechanical analyses, as well as the option to use auxiliary fields, which allow you to select any field functions defined on the coupled boundaries.

- In a thermal co-simulation analysis, one simulation solves for temperature and the other solves for convective flux (heat transfer coefficient + reference temperature).
- In a mechanical co-simulation analysis, one simulation solves for pressure or traction (pressure + wall shear stress) and the other solves for displacements.

- In general, the partner simulations can exchange any field functions defined on the coupled boundaries, using auxiliary fields.

Zone Properties

Index	Specifies the unique index that identifies a zone. A zone is connected with the zone of same index in the partner simulation. (Read-only)
Zone Type	Specifies the type of coupling zone. The available option is Surface Zone.

Selecting the Coupled Boundaries

The coupled boundaries selected for a zone couple to the boundaries selected for the zone of same index in the partner simulation.

For a coupling zone, you specify the coupled boundaries from the boundary nodes within the relevant [Region] > Boundaries node.

To assign a coupled boundary to a zone:

1. Expand the Regions > [Region] > Boundaries node.
2. Expand the node of the boundary that you want to assign to a coupling zone.

3. Select the Physics Conditions > External Code Coupling Specification node and for the External Code Coupling Specification property select the coupling zone of interest.



STAR-CCM+ automatically updates the zone Coupled Boundaries condition with the selected boundary.



The External Boundaries condition displays the partner simulation coupled boundaries. When you start the co-simulation, the External Boundaries condition is updated automatically based on the coupled boundaries that you selected for the partner zone.

When you run the co-simulation, the partner simulation coupled boundaries are imported in the current simulation as a region. The physics continuum that represents the partner simulation is automatically assigned to the region representing the partner simulation.

Specifying Co-Simulation Thermal Conditions

Thermal co-simulation conditions become available when you select an energy model. To create a one-way coupling, activate an export option. To create a two-way coupling, also activate an import option. There are separate properties for fluid export, fluid import, solid export, and solid import.



When exporting Convective Flux, you choose between two forms of heat transfer coefficient and their corresponding reference temperatures:

Heat Transfer Coefficient	Reference Temperature
Local Heat Transfer Coefficient	Local Heat Transfer Reference Temperature
Specified Y+ Heat Transfer Coefficient	Specified Y+ Heat Transfer Reference Temperature

For most FSI applications, use the specified y+ heat transfer reference temperature. If the simulation involves radiation, use the local heat transfer reference temperature. You can define the y+ value for the specified y+ heat transfer reference temperature using the `Specified Y+ Heat Transfer Coefficient` field function. These can be found under the Tools > Field Functions node in STAR-CCM+.

When exporting Temperature, choose the Temperature field function.

Fluid Thermal Export Specification Properties

Method	Controls the export of thermal data.	
	No Thermal Export	Deactivates thermal export. A No Thermal Export field empty child node is added.
	Convective Flux Export Field	Select this option for the fluid simulation. A Convective flux export fields child node is added where you can select the appropriate Heat Transfer Coefficient and Reference Temperature field functions. A Heat Transfer Coefficient Export Mapper node and a Reference Temperature Export Mapper node are added to the Values

Fluid Thermal Import Specification Properties

Method	Controls the import of thermal data.	
	No Thermal Import	Deactivates thermal import.
	Temperature	Activates import of temperature from the partner simulation.



Solid Thermal Export Specification Properties

Method	Controls the export of thermal data.	
	No Thermal Export	Deactivates thermal export. A No Thermal Export field empty child node is added.
	Temperature Export Field	A Temperature Export Field child node is added where you can select the appropriate Temperature field function from which to transfer the temperature at the wall. A Temperature Export Mapper node is added to the Values

Solid Thermal Import Specification Properties

The convective fluxes exported by the fluid simulation can be imported by the solid simulation as convective fluxes or environmental fluxes. Both the Convective Flux and Environmental Flux import methods allow the solid simulation to import the same thermal fields, but under different conditions. The Convective Flux import option enables the convection thermal condition on the coupled boundaries, whereas the Environmental Flux import option enables the environmental thermal condition.

Both import methods can be used to model radiation and convection in the solid simulation. However, when the Convective Flux import method is used, radiation modeling between meshless regions requires the outward normal option on boundaries, which makes the solid bodies opaque by definition. Alternatively, the Environmental Flux import method allows radiation modeling within a semitransparent solid and across meshless regions without any restriction.

The Environmental Flux import method also allows you to model thermal conduction and convection using the environmental thermal condition on the boundaries. Use the Environmental Flux import method when the solid simulation accounts for radiation in addition to convection and conduction. Otherwise the Convection Flux import method is sufficient.

Method	Controls the import of thermal data.	
	No Thermal Import	Deactivates thermal import.

Convective Flux	Activates import of convective fluxes from the partner simulation, enabling the Convection condition on the co-simulation coupled boundaries.
Environmental Flux	Activates import of environmental fluxes from the partner simulation, enabling the Environment condition on the co-simulation coupled boundaries. For information on how to set the Environment Condition see Setting the Environment Condition.

Specifying Co-Simulation Mechanical Conditions

In a solid simulation, mechanical co-simulation conditions become available when you select the *Finite Volume Solid Stress* model.

In a fluid simulation, the conditions become available when you assign a Morphing motion to the region.

To activate the Morpher for the fluid region:

1. Right-click the Tools > Motions node and select New > Morphing.
2. Expand the Regions > [Region] > Physics Values node.
3. Select the Motion Specification node and set Motion to Morphing.

   STAR-CCM+ automatically updates the coupled boundary nodes to use the displacement data from the partner simulation. The Morpher physics condition for the coupled boundaries is automatically set to STAR-CCM+ Co-Simulation.

To create a one-way coupling, activate an export option. To create a two-way coupling, also activate an import option There are separate properties for fluid export, fluid import, solid export, and solid import.



Fluid Mechanical Export Specification Properties

Method	Controls the export of mechanical data.	
	No Mechanical Export	Deactivates mechanical export. A No Mechanical Export field empty child node is added.
	Pressure Export Field	Activates export of pressure data in the fluid simulation. A Pressure Export Field child node is added where you can select the Pressure field function. A Pressure Export Mapper node is added to the Values
	Traction Export Fields	Activates export of pressure and wall shear stress data in the fluid simulation. A Traction Export Fields child node is added where you can select the Pressure and Wall Shear Stress field functions. Two nodes, Pressure Export Mapper and Wall Shear Stress Export Mapper are added to the Values.

Fluid Mechanical Import Specification Properties

Method	Controls the import of mechanical data.	
	No Mechanical Import	Deactivates thermal import.
	Displacements	Activates import of displacements from the partner simulation.


Servers
Fluid
Solid
Simulation
Co-Simulations
STAR-CCM+ Co-Simulation 1
Zones
Zone 1
Conditions
Auxiliary Fields Export Specification
Auxiliary Fields Import Specification
Coupled Boundaries
External Boundaries
Solid Mechanical Export Specification
Displacement Export Field
Solid Mechanical Import Specification
Values
Displacement Export Mapper

Solid Mechanical Export Specification Properties

Method	Controls the export of mechanical data.	
	No Mechanical Export	Deactivates mechanical export. A No Mechanical Export field empty child node is added.
	Displacement Export Field	A Displacement Export Field child node is added where you can select the appropriate Displacement field function from which to transfer the displacements. A Displacement Export Mapper node is added to the Values

Solid Mechanical Import Specification Properties

Method	Controls the import of mechanical data.	
	No Mechanical Import	Deactivates mechanical import.
	Pressure	Activates import of pressure from the partner simulation.

Note: In a mechanical co-simulation analysis, you can access the settings for the coupling algorithm between the fluid and the solid solvers by activating the Partitioned Coupling model in each of the partner simulations. For more information, see Using Partitioned Coupling with STAR-CCM+ to STAR-CCM+ Co-Simulation.

Specifying Co-Simulation Auxiliary Fields

In STAR-CCM+ co-simulation, auxiliary fields allow you to exchange any scalar or vector fields defined on the coupled boundaries. For a coupling zone, you specify auxiliary fields exchange conditions through the Auxiliary Fields Export Specification and Auxiliary Fields Import Specification nodes.

To export an auxiliary field to the partner simulation:

1. Right-click the Auxiliary Fields Export Specification node and select New Auxiliary Field.
2. Set the Field Function property of the newly created Auxiliary Field node through the *Auxiliary Field - Properties* panel: you can specify any scalar or vector field defined on the coupling boundary.

On adding an auxiliary field, STAR-CCM+ adds an extra field function whose name is composed of the field function selected for export, followed by the text “Exported by Co-Simulation”. For example, “Velocity Exported by Co-Simulation”. In a similar way, when the co-simulation starts running, any auxiliary field imported from the partner co-simulation is identified by a field function whose name ends with the text "Imported by Co-Simulation".



Auxiliary Fields Export Specification Properties

Number of Auxiliary Fields Exported	Displays the number of Auxiliary Fields that are exported to the partner co-simulation.

Auxiliary Fields Import Specification Properties

Number of Auxiliary Fields Imported	Allows you to specify the number of Auxiliary Fields that are imported from the partner co-simulation.

Auxiliary Field Properties

Index	Displays the unique number assigned to the Auxiliary Field in order to distinguish it from others. (Read-only).
Field Function	Specifies the scalar or vector field function to be exported to the partner co-simulation.

Specifying Co-Simulation Values for Mapping

In order to exchange data between partner simulations, STAR-CCM+ must map data across the coupled boundaries between simulations.

For each quantity mapped by the current simulation, STAR-CCM+ adds a mapper control node under the relevant [Zone] > Values node. These mapper nodes reference the data mappers that are defined within Tools > Data Mappers.

You can either create data mapper nodes manually within Tools > Data Mappers, or rely on STAR-CCM+ to create them automatically. If you do not create a surface data mapper for each of the mappers under the [Zone] > Values node, then one is created using default values when the co-simulation starts.

The screenshot below shows a Mapping Mode of Map both directions. As a result, the Heat Transfer Coefficient and Reference Temperature fields each require a mapper at import and the Temperature field requires a mapper before it is exported.



If you create your own mapper, part and field properties are specified for you in both source and target. You must still specify proximity constraints, normal direction constraints, and interpolation method and you may wish to specify other mapper properties.

Note: You can inspect properties of a mapper you have created at any time, but you can inspect the properties of default mappers only after the co-simulation run has started.

Mapper Properties

Mapper	Select the data mapper to use for the variable.	
	Default	The default setting for each mapper before running the simulation. If you do not specify a new mapper, then a mapper with default values is created automatically when the co-simulation starts.
	Mappers from the Tools > Data Mappers node	STAR-CCM+ makes other mappers available here, if they are available in the Tools > Data Mappers node. See Creating a Data Mapper.

Specifying Co-Simulation Exchange Conditions and Solver Settings

The frequency with which data is exchanged between the two partner simulations is controlled by specifying the number of iterations between exchanges in steady simulations, or the coupling time-step in transient simulations. If you are using the *STAR-CCM+ Implicit Co-Simulation* model, you can also set coupling to occur more than once per time-step by specifying the number of inner iterations between exchanges within a time-step.

You access the frequency of exchange properties and menu items by selecting the relevant STAR-CCM+ co-simulation solver under the Solvers node.

Co-Simulation Steady Solver Controls

Set the data exchange frequency in steady simulations through the Co-simulation Steady > Iterations per Exchange node.

Iterations per Exchange Properties

Number of Iterations per Exchange	Specifies the number of iterations to run before transferring data. Only available for Steady State simulations.

Co-Simulation Explicit/Implicit Unsteady Controls

Specify the coupling method in unsteady simulations through the child node, Coupling Step Size Negotiation Method node.

Coupling Step Size Negotiation Method Properties

Method	Controls the method to determine the number of time-steps to run before transferring data. To prevent conflict, the first simulation to start will dictate the method that is used in both simulations.	
	Constant Negotiation Method	Allows you to specify a constant coupling time-step at which data is exchanged. You set the coupling time-step size through the child node, Constant Negotiation Method. Make sure that you enter a value that is consistent with your simulation time-step (that is, the time-step that you set for the Implicit Unsteady, or Explicit Unsteady, solver).
	Export Negotiation Method	Coupling occurs at the simulation time-step specified in the current simulation. A child node, Export Negotiation Method, is added.
	Import Negotiation Method	Coupling occurs at the simulation time-step specified in the partner simulation. A child node, Import Negotiation Method, is added.

Maximum Negotiation Method	Coupling occurs at the larger simulation time-step specified in the two simulations. A child node, Maximum Negotiation Method, is added.
Independent Negotiation Method	You set the coupling time step independently in the two simulations. This method is particularly useful when the time scales in the two simulations are very different, as it allows the two simulations to transfer data at convenient, independent frequencies. A child node, Independent Negotiation Method, is added. From this node, you can set the coupling time-step size. Make sure that the value that you specify for the coupling time-step is consistent with the time-step of your simulation.

For the Co-simulation Implicit Unsteady solver you can also specify the number of inner iterations between exchanges within a time-step through the Co-simulation Implicit Unsteady > Inner Iterations Per Field Exchange node.

Inner Iterations Per Fields Exchange Properties

This property only applies to the Co-simulation Implicit Unsteady solver.

Number of Inner Iterations Per Fields Exchange	Controls how many times within a time-step data is exchanged. In a time-step, coupling occurs every time the Number of Inner Iterations Per Fields Exchange that you specify is run.

Co-Simulation Solver Properties

The Co-Simulation Steady, Co-Simulation Explicit Unsteady, and the Co-Simulation Implicit Unsteady solver nodes share common properties:

Co-Simulation Solver Properties

Verbosity	Provides more output for this value while the simulation is running. This option can be useful for debugging problems when they occur.	
	None	Leaves output at its default.
	Low	Displays additional information.
	High	Displays detailed information about the co-simulation run.
Solver Frozen	Prevents or allows the solver from updating the solution while iterating.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.
Temporary Storage Retained	Specifies whether to retain temporary storage at the end of the iteration, or not.	
	Activated	Retains temporary storage.
	Deactivated	Does not retain temporary storage.

Running the Co-Simulation

Before running a co-simulation interactively, load the partner simulations and specify the co-simulation settings as described in the previous sections.

To run a co-simulation:

1. Select the leading simulation.
2. Click (Run) or select Solution > Run.

The co-simulation progress can be monitored in the *Output* window, and STAR-CCM+ reports output for each data exchange process.

Running on a Cluster

To run both simulations on the same cluster, use the command for each case:

```
starccm+ <options> -load <simulation>
```

Once each case has opened, enter the master server process in the lead case under the Connect to Server node as described previously.

Stopping the Co-Simulation:

To stop a co-simulation:

1. Press (Stop Iterating) in any of the participating simulations.
2. Once the simulation has stopped running, right-click the relevant node within the Co-Simulations node, and select Disconnect from Partner Co-simulation.

Both simulations are now disconnected and the co-simulation has stopped.

Note: Do not press (Stop Iterating) in both simulations as this option is not supported.

Restarting the Co-Simulation

In certain cases, it might be necessary to change some properties and restart the co-simulation.

To do this:

1. Stop one of the participating simulations. The participating simulations can either be in Running/Computing mode or in Waiting mode.

   You need to stop only one of the participating simulations. If you stop the Running/Computing simulation, the simulation stops immediately. If you stop the Waiting simulation, the simulation stops in the immediately following iteration update, that is, after the first iteration is run.
2. Make the intended changes in both or one of the simulations.

To restart the co-simulation:

- Press (Run) in the simulation that you stopped.

Partitioned Coupling Module

In STAR-CCM+, the Partitioned Coupling Module (PCM) provides a framework for partitioned coupling algorithms.

In simulations involving multiple physics disciplines, for example fluid flow interacting with structural mechanics, the two domains can be coupled using either monolithic or partitioned methods. In the monolithic approach, the equations for the fluid flow and the solid structure displacement are solved at the same time as a single linear system. In the partitioned approach, independent solvers solve for the fluid and the solid domains separately and a coupling algorithm ensures convergence of the Fluid-Structure Interaction (FSI) problem. This approach preserves each solver specialization and capabilities, but the accuracy of the solution relies on the stability of the coupling algorithm.

In STAR-CCM+, multiple physics are coupled with the partitioned coupling method. For Fluid-Structure Interaction analyses, the FSI coupling algorithm that is currently available is based on the fixed-point iteration method.

Using Partitioned Coupling with STAR-CCM+ to STAR-CCM+ Co-Simulation

STAR-CCM+ to STAR-CCM+ Co-Simulation allows two running simulations to exchange field data across their coupled boundaries.

In a two-way mechanical co-simulation analysis, the solid simulation solves for the structural displacement in response to the traction fields that the fluid simulation provides at the coupled boundaries.

The fluid simulation, in turn, uses the structural displacements provided at the coupled boundaries by the solid simulation as inputs to the mesh morpher.

For more information, see STAR-CCM+ to STAR-CCM+ Co-Simulation.

The Partitioned Coupling Module provides a coupling algorithm based on the fixed-point iterative method. This method uses the field data exchanged at the fluid-structure interface during the co-simulation to ensure convergence of the coupled problem.

This section describes how to use the Partitioned Coupling Module in conjunction with STAR-CCM+ to STAR-CCM+ Co-Simulation.

Activating the Partitioned Coupling Model

In a mechanical STAR-CCM+ to STAR-CCM+ co-simulation analysis, you access the *Partitioned Coupling Module (PCM)* framework by activating the Partitioned Coupling physics model in each of the partner simulations.

The Partitioned Coupling model becomes available for selection after you select a STAR-CCM+ Co-Simulation model. In the fluid simulation, activate the Partitioned Coupling model after you set up the morpher for the fluid region. For more information, see Setting Up the Morpher.

To activate the Partitioned Coupling model:

1. Expand the Continua node.
2. Right-click the relevant physics continuum node and select Select models...
3. In the *Physics Model Selection* dialog:
   a. Select appropriate physics models based on your analysis requirements.
   b. Select Co-Simulation from the *Optional Models* group box.
   c. From the *Co-Simulation* group box, select a STAR-CCM+ co-simulation model. For more information, see Activating the STAR-CCM+ Co-Simulation Models.
   d. From the *Optional Models* group box, select Partitioned Coupling.

When you activate the Partitioned Coupling model, STAR-CCM+ automatically selects the Fixed-Point Iteration model from the *Partitioned Coupling* group box.

Fixed-Point Iterative Solver

In the fluid simulation, selecting the Fixed-Point Iteration model enables the Fixed-Point Iterative solver. This solver uses the fixed-point iterative method to under-relax the solid displacement at the fluid-structure interface.

Fixed-Point Iterative Solver Properties

Verbosity	Provides additional output for this value while the simulation is running. This property is useful for debugging problems when they occur. The available options are None, Low, and High.	
Solver Frozen	Prevents or allows the solver from updating the solution while iterating. This option is primarily a debugging tool.	
	Activated	The solver does not update the solution.
	Deactivated	The solver updates the solution.
Temporary Storage Retained	Specifies whether to retain temporary storage at the end of the iteration, or not.	
	Activated	Retains temporary storage.
	Deactivated	Does not retain temporary storage.

Setting Up the Morpher

In the fluid simulation, set up the morpher for the fluid region before you activate the Partitioned Coupling model.

Since the Partitioned Coupling Module controls the mesh morpher in the fluid simulation, make sure that the Morph at Inner Iterations property for the Mesh Morpher Solver is deactivated when using PCM. For more information on how to set up the morpher for STAR-CCM+ Co-Simulation, see Specifying Co-Simulation Mechanical Conditions.

Partitioned Coupling Controls in the Fluid Simulation

When you activate the Partitioned Coupling module in the fluid simulation, the FSI Coupling Algorithm and the FSI Displacement Under-Relaxation nodes become available in the simulation tree. The settings that you specify for

these nodes are general settings that apply to all the coupling zones defined in the simulation. For more information, see Specifying Co-Simulation Global Settings.



The FSI Coupling Algorithm node allows you to select the implicit or explicit coupling scheme. The Explicit coupling method results in weak coupling, appropriate for problems in which the effect of the structure on the fluid is substantially larger than the effect of the fluid on the structure. To reach a stable solution when using the Explicit coupling scheme, use a sufficiently small simulation time-step.

The Implicit coupling method results in strong coupling and allows you to use larger time-steps than the Explicit scheme. This method also ensures equilibrium at the fluid-structure interface, which is not guaranteed with

explicit coupling. The Implicit coupling scheme also enables additional settings, the FSI Under Relaxation Option and the FSI Displacement Convergence Tolerance.



FSI Coupling Algorithm Properties

Method — Defines the type of algorithm used for coupling at the fluid-structure interface.

- Explicit — The coupling algorithm only works at the beginning of each time-step. Under-relaxation of solid displacement and fluid load control is applied only at the beginning of the time-step. No equilibrium is ensured at the fluid-structure interface during the time-step evaluation.
- Implicit — The coupling algorithm works at all inner iterations within each time-step. This ensures kinematic and dynamic equilibrium at the fluid-structure interface in each time-step. Selecting the Implicit method enables the FSI Under Relaxation Option and the FSI Displacement Convergence Tolerance nodes in the simulation tree.

FSI Displacement Under-Relaxation Properties

Under-Relaxation Factor

Specifies the factor used to under-relax the displacement of the structure. When using the Explicit FSI Coupling algorithm, this factor is applied to the solid displacement field at the beginning of each time-step. When using the Implicit FSI Coupling algorithm with the Aitkens FSI Under-Relaxation Option, this factor is applied to the solid displacement field at the second iteration of each time-step. When using the Implicit FSI Coupling algorithm with the Constant FSI Under-Relaxation Option, this factor is applied at all iterations starting at the second iteration of each time-step.

FSI Under-Relaxation Option Properties (Only available for Implicit FSI Coupling Algorithm)

Method

Specifies the method that is used to calculate/apply the under-relaxation factor.

Constant

A constant under-relaxation factor is applied to the solid displacements, starting from the second iteration of each time-step. At the first iteration of each time-step, no under-relaxation factor is applied.

Aitkens

The under-relaxation factor is calculated adaptively based on the displacement residuals, according to Aitken's method. The under-relaxation factor is applied at each iteration within a time-step. This method helps to accelerate convergence of the fixed-point algorithm.

FSI Displacement Convergence Tolerance Properties (Only available for Implicit FSI Coupling Algorithm)

Displacement Convergence Tolerance	Sets the convergence tolerance to evaluate convergence during each time-step. When convergence is achieved in a time-step, a message is displayed in the *Output* window. Specify a value in the range [0, 1].

Aitkens Urf Bounds Properties (Only available for Implicit FSI Coupling Algorithm)

The following properties are available only when the FSI Under Relaxation Option Method is set to Aitkens.

Urf Lower Bound	Lower bound for under-relaxation factor.
Urf Upper Bound	Upper bound for under-relaxation factor.

Partitioned Coupling Controls in the Solid Simulation

When you activate the Partitioned Coupling module in the solid simulation, the FSI Fluid Load Control Option node becomes available in the simulation tree as a global co-simulation condition.

For this node, the following properties are available.

FSI Fluid Load Control Option Properties

Method	Specifies the method for applying the fluid load on the solid structure.	
	None	No fluid load control is used.
	Linear Ramp & Clipping	Allows you to control how the fluid load is applied on the solid at the fluid-structure interface. Enables the FSI Fluid Load Linear Ramp & Clipping node under [STAR-CCM+ Co-Simulation] > Values.

FSI Fluid Load Linear Ramp & Clipping Properties

Time Ramping Parameters	You can specify an initial and ending time to increase the imported fluid traction load progressively, from zero to the actual load. Enter the Time Ramping Parameters as [t_start, t_end].
Traction Clipping Parameters	To reflect the actual operating conditions, you can limit the traction values imported from the fluid simulation. For the imported traction load, you can specify the maximum and minimum pressure, as well as the peak shear stress magnitude. Enter the Traction Clipping Parameters as [p_max, p_min, shear_max].

Field Functions and Residual Monitor

When you activate the Partitioned Coupling Module, the following field functions become available

MappedMorpherDisplacement	Displacement that the morpher uses to morph the fluid mesh.
MappedDisplacementAtIteration	Imported solid displacement, before the under-relaxation is applied.
FsiUnderRelaxationFactor	Under-relaxation factor calculated during the time-step.

Also, STAR-CCM+ adds the FSI Displacement monitor to the Residual Monitors.

Co-Simulation API

STAR-CCM+ provides an Application Programming Interface (API) that allows you to write partner code that exchanges data with a running STAR-CCM+ simulation. As the partner code can be part of any other application that supports linking to a flat C library, the co-simulation API extends the range of co-simulation scenarios possible with STAR-CCM+.



When initialized, the Co-Simulation API launches a companion STAR-CCM+ API Server alongside the partner software. During co-simulation, the companion STAR-CCM+ API Server communicates with other STAR-CCM+ servers using the same protocols as a co-simulation between two STAR-CCM+ servers. Partner software makes calls into the Co-Simulation API library to receive data from, and transfer data to, the STAR-CCM+ API Server.

Co-Simulation API Files

The STAR-CCM+ installation contains a header file that describes the functions available for partner code to call. This header file is located at:

- Windows:

  `[INSTALLDIR]\star\include\StarccmplusCoSimulationApi.h`

- Linux:

  `[INSTALLDIR]/star/include/StarccmplusCoSimulationApi.h`

You can access an HTML version of the header file from the STAR-CCM+ interface by selecting Help > Co-simulation API.

The library to which partner code must link is located at:

- Windows:

  `[INSTALLDIR]\star\[OS]\[COMPILER]\lib\StarccmplusCoSimulationApi.dll`

- Linux:
  `[INSTALLDIR]/star/[OS]/[COMPILER]/lib/StarccmplusCoSimulationApi.so`

An example partner program is included at:

- Windows:

  `[INSTALLDIR]\star\examples\api\SpringMass`

- Linux:

  `[INSTALLDIR]/star/examples/api/SpringMass`

You are advised to study the example code before attempting to create your own partner program. A tutorial, Co-Simulation API: Spindle Valve, is also available within the Coupling with CAE Codes section of the Tutorial Guide.

The STAR-CCM+ installation contains two helper scripts for Windows, and two helper scripts for Linux:

- Windows:
  - buildapi: `[INSTALLDIR]\star\bin\buildapi.bat`
  - starapi : `[INSTALLDIR]\star\bin\starapi.bat`
- Linux:
  - buildapi : `[INSTALLDIR]/star/bin/buildapi`
  - starapi : `[INSTALLDIR]/star/bin/starapi`

buildapi

The `buildapi` helper script compiles the code that uses the co-simulation API. The script works with the following compilers:

- gcc for C and Intel Fortran 11 on Linux
- Intel C++ Compiler and Intel Fortran Compiler on Windows

To use the script to compile the code, enter the following command:

- On Windows:

```
> call "[INTEL_C++_INSTALLDIR]\bin\compilervars.bat"
intel64
```

```
> [INSTALLDIR]\star\bin\buildapi.bat [<verbosity>]
<executable name> [<source files>]
```

- On Linux:

```
% [INSTALLDIR]/star/bin/buildapi [<verbosity>] <executable
name> [<source files>]
```

where:

- `[<verbosity>]` specifies the verbose mode.

  The available options are:

  - `-v` for low verbosity
  - `-vv` for medium verbosity
  - `-vvv` for high verbosity
- `<executable name>` specifies the target name of the final executable.
- `[<source files>]` specifies the source code files or includes source file directories. If you do not specify this argument, the compiler will compile all the files with extension `.cpp`, `.c`, or `.f`, `.f90` that are located in the current folder.

Note: If you experience problems running the script, open it in a text editor and check that all required paths are resolved by the environment you are working with.

starapi

Use the `starapi` script for launching the partner program. The script sets the environment that allows the partner program to load the co-simulation library and launch the STAR-CCM+ API Server.

To launch the partner program from a command console:

- On Windows:

```
> [INSTALLADIR]\star\bin\starapi.bat [<verbosity>]
<executable> [<executable arguments>]
```

- On Linux:

```
% [INSTALLDIR]/star/bin/starapi [<verbosity>] <executable>
[<executable arguments>]
```

where:

- `[<verbosity>]` specifies the verbose mode.

  The available options are:

  - `-v` for low verbosity
  - `-vv` for medium verbosity
  - `-vvv` for high verbosity
- `<executable>` specifies the executable path and name.
- `[<executable arguments>]` contains additional parameters required by the partner program.

Interacting with the STAR-CCM+ API Server

Partner code accesses the STAR-CCM+ API Server through the Co-Simulation API library. The partner code is responsible for initializing the companion STAR-CCM+ API Server.

To prepare the STAR-CCM+ API Server in partner code, include the following actions:

1. Load the Co-Simulation API library
2. Initialize the library
3. Initialize the STAR-CCM+ API Server

   e. Create a physics continuum on the STAR-CCM+ API Server

   f. Create a region on the STAR-CCM+ API Server

   g. Register callback functions with message handlers

   h. Provide a surface mesh to the STAR-CCM+ API Server that conforms to the STAR-CCM+ Surface Polyhedral Mesh format. For more details, see Surface Polyhedral Mesh Format.

   i. Register data fields that the client code sends to STAR-CCM+.

   j. Register data fields that the client code receives from STAR-CCM+.

4. Notify the API that the surface mesh is ready for use.
5. Co-simulation loop:
   a. Calculate the desired time-step size and pass to the Co-Simulation API. The API returns the target time-step.
   b. Wait for incoming fields. On arrival, retrieve the registered incoming fields within the fillContainer callback function.
   c. Solve for partner solution.
   d. Prepare solution data for export.
   e. Notify the Co-Simulation API that outgoing fields are ready.
   f. Repeat the co-simulation loop till done.
6. When complete, finalize the Co-Simulation API.

Data Containers in Co-Simulation

The Co-Simulation API provides you with containers for accessing data on the STAR-CCM+ API Server.

A container is an intermediate data store used to transfer data between the partner program and the STAR-CCM+ API Server. Each container has its own ID of type `signed int`. API calls that create containers return the container ID. Partner code must retain container IDs for all future operations that involve the same container.

Continuum Container

Use the continuum container to instantiate a partner physics continuum in the STAR-CCM+ API Server. The continuum container requires you to set a material model and a time model when you first create it. For STAR-CCM+ v9.04, you must choose a single component solid for the material model, and the implicit unsteady for the time model.

Example:

```
int continuumId = lib.createContinuum("SpringMass",
StarccmplusSingleComponentSolid,
StarccmplusImplicitUnsteady,lib.createPropertiesContainer
());
```

Region container

Identifies the single region supported by the STAR-CCM+ API Server.

Example:

```
int regionId = lib.createRegion(continuumId, "Surface1");
```

Properties container

Use a properties container to transfer simulation properties between programs. Also use this container type to pass additional options to other methods.

- Properties are stored as name/value pairs.
- Property types can be int, double, or string.
- Property data can be single values or arrays.

Mesh Container

Use a mesh container to provide a surface mesh to the STAR-CCM+ API Server. This surface mesh must provide values and arrays as defined in the specification for a STAR-CCM+ Polyhedral Surface Mesh.

- Use a container registration functions to register the mesh containers
- Before you can register the mesh, you must already have an ID for the region container.
- Fill the container with values and arrays when the API calls the `fillContainer` callback function with the mesh container ID.

Field Container

Use field containers to exchange solution data with the STAR-CCM+ API Server.

- Use container registration functions to specify the fields that the partner program works with.
- Send or retrieve data when the STAR-CCM+ API Server calls the `fillContainer` callback function with a valid ID.

STAR-CCM+ Polyhedral Surface Mesh

When the STAR-CCM+ API requests that the partner code fill the surface mesh, the partner code should provide the data identified in the following table:

Data Type	String Identifier	Length
Integer	`NumberOfFaces`	Single value
Integer	`NumberOfVertices`	Single value
Integer	`NumberOfDimensions`	Single value
Double Array	`Coordinates`	`NumberOfVertices * NumberOfDimensions`
Integer Array	`VertexIds`	`NumberOfVertices`
Integer Array	`FaceVertexList`	`Sum over all faces([Face Vertex Count + 1])`

The FaceVertexList is a ragged array whose sequence follows the pattern:

```
[<Num Vertices on Face 0>, V0, V1, V2,..., <Num Vertices on
Face 1>, V0, V1, V2, ...]
```

The following diagram illustrates the relationship between the arrays:



After constructing the surface mesh, the partner program should call API function `notifyOutgoingMeshesReady()`.

STAR-CCM+ Field Types for Exchange

The Co-Simulation API provides functions for outgoing and incoming fields. Registration functions instruct the API to prepare for exchanging particular fields.

Example registration for outgoing fields:

```
int dispId = lib.registerOutgoingField(
meshId,
StarccmplusDisplacementSpecification,
StarccmplusDisplacements,
lib.createPropertiesContainer());
```

Example registration for incoming fields:

```
int htcId = lib.registerIncomingField(
meshId,
StarccmplusConvectiveFluxSpecification,
StarccmplusHeatTransferCoefficient,
lib.createPropertiesContainer());

int trefId = lib.registerIncomingField(
meshId,
StarccmplusConvectiveFluxSpecification,
StarccmplusReferenceTemperature,
lib.createPropertiesContainer());
```

Use the following table in registration functions for data fields:

StarccmplusSpecification	StarccmplusFieldType
StarccmplusPressureSpecification	StarccmplusPressure
StarccmplusTractionSpecification	StarccmplusPressure AND StarccmplusWallShearStress
StarccmplusDisplacement-Specification	StarccmplusDisplacements
StarccmplusTemperature-Specification	StarccmplusTemperature
StarccmplusConvectiveFlux-Specification	StarccmplusHeatTransferCoefficient AND StarccmplusReferenceTemperature
StarccmplusEnvironmentalFlux-Specification	tarccmplusHeatTransferCoefficient AND StarccmplusReferenceTemperature
StarccmplusAuxiliarySpecification	StarccmplusAuxiliaryField

Use the `StarccmplusSpecification` to say what quantity the partner program is going to exchange (incoming or outgoing). Use the `StarccmplusFieldType` to define the type of data that the codes exchange for that quantity.

StarccmplusFieldType	Array Type	Array Length	Dimension	Units
`StarccmplusPressure`	1 double per face	`NumberOfFaces`	Pressure	Pa
`StarccmplusWallShearStress`	3 doubles per face	`NumberOfFaces * 3`	Pressure	Pa

StarccmplusFieldType	Array Type	Array Length	Dimension	Units
`StarccmplusDis-placements`	3 doubles per vertex	`NumberOfVertices * 3`	Length	m
`StarccmplusTem-perature`	1 double per face	`NumberOfFaces`	Temperature	K
`StarccmplusHeat-TransferCoeffi-cient`	1 double per face	`NumberOfFaces`	Power/Length^2-Temperature	W/m^2K
`StarccmplusRefer-enceTemperature`	1 double per face	`NumberOfFaces`	Temperature	K
`StarccmplusAuxil-iaryField`	1 double per face	`NumberOfFaces`		

Loading the Co-Simulation API Library

`LibraryLoader.h` and `LibraryLoader.cpp` in the SpringMass example demonstrate how to wrap the Co-Simulation API for dynamic loading. The LibraryLoader class also forwards function calls to the library. This forwarding behavior means that the SpringMass program makes calls on the LibraryLoader object rather than the API directly.

Checking the Co-Simulation API Version

Each release of the STAR-CCM+ Co-Simulation API is identified with a unique number, that is, the API version. The STAR-CCM+ Co-Simulation API and STAR-CCM+ have independent life cycles. Typically, new versions of the STAR-CCM+ Co-Simulation API will be released less frequently than new versions of STAR-CCM+.

The STAR-CCM+ Co-Simulation API version includes three sets of integers, separated by a period, identifying the major version, minor version, and patch version respectively. Example code for checking the STAR-CCM+ Co-Simulation API version is included in the SpringMass example.

Registering API Callback Functions

Partner code must register callback functions as follows:

Handling messages:

```
typedef void (*StarccmplusMessageHandler)
(StarccmplusMessageType const msgType, char const* msg);
```

Handling calls to send data:

```
typedef void (*StarccmplusFillContainerHandler) (int
containerId, char const* requirement);
```

See the example file, `SpringMass.cpp`, for further details.

Note: Register callback functions only after creating the region container.

Compiling Partner Code

Before coupling your code to STAR-CCM+, compile the code using the `buildapi` helper script that is provided in the STAR-CCM+ installation.

To compile partner code:

1. In a terminal/command prompt, go to the folder that contains your code files.
2. From this location, run the `buildapi` helper script:
   - On Windows, set the Intel C++ Compiler environment before running the buildapi.bat script:

     ```
     > call "[INTEL_C++_INSTALLDIR]\bin\compilervars.bat"
     intel64
     ```

     ```
     > [INSTALLDIR]\star\bin\buildapi.bat <executable name>
     ```

   - On Linux:

     ```
     % [INSTALLDIR]/star/bin/buildapi <executable name>
     ```

   where `<executable name>` specifies a name for the output executable file.

   As the argument for the source files is not specified, this command compiles all the `.cpp` files that are located in the project folder and creates an executable output file that is used to launch the partner program. See buildapi for more information as well as additional arguments that you can use with this script.

Running a Co-Simulation with Partner Code

To run a co-simulation with previously compiled partner code:

1. Launch the STAR-CCM+ fluid simulation.
2. Set up the simulation for STAR-CCM+ Co-Simulation. For more information, see Specifying Co-Simulation Conditions and Values.
3. In the fluid simulation, select the Co-Simulations > STAR-CCM+ Co-Simulation 1 > Conditions > Connect Method node and set Method to Launch Application and Connect.
4. Select the Connect Method > Launch Application and Connect node. For the Launch Command property, enter:
   - On Windows:

     ```
     > [INSTALLDIR]\star\bin\starapi.bat <executable>
     ```

   - On Linux:

     ```
     % [INSTALLDIR]/star/bin/starapi <executable>
     ```

   where `<executable>` specifies the path to the executable file created when compiling the code.
5. Run the co-simulation from the fluid simulation.

   When you run the co-simulation, the `starapi` helper script launches the partner program and the associated STAR-CCM+ API server. The STAR-CCM+ fluid simulation connects to the running program through the STAR-CCM+ API server and the co-simulation analysis starts.

   See starapi for more information about this script as well as additional options that you can use with this script.

   When using the Launch Application and Connect method, the STAR-CCM+ fluid simulation and the STAR-CCM+ API server share the same STAR-CCM+ license, meaning that you require a single license to run the co-simulation. On the other hand, if you launch the partner program separately (for example running the starapi script in a command console) and then connect the STAR-CCM+ fluid simulation to the running program using the Connect to Partner Simulation Method, you will require two STAR-CCM+ licenses.

Co-Simulation with Abaqus

The co-simulation technique can be used to solve complex fluid-structure interaction (FSI) problems by coupling Abaqus and STAR-CCM+. Abaqus solves the structural domain and STAR-CCM+ solves the fluid domain. The two domains are interconnected using the SIMULIA Co-Simulation Engine (CSE). This chapter describes how to use the co-simulation technique with STAR-CCM+ and Abaqus to solve FSI problems. For general information on using Abaqus, refer to the Abaqus Analysis User's Guide, and for more information on co-simulation in Abaqus, see "Co-simulation", Section 17.1 of the Abaqus Analysis User's Guide.

Target Applications

The STAR-CCM+ and Abaqus Co-Simulation capability targets simulations involving compliant structural response to fluid flow and their complex interrelationship. Mechanical and/or thermal properties can be exchanged across the fluid-structure interface.

Most of the features in STAR-CCM+ can be used in a co-simulation, including compressible and incompressible flow, laminar and turbulent flow, heat transfer, a rich set of material options, and free-surface, multi-species, and multiphase driven flows. The implicit unsteady solver is enabled for use with Abaqus co-simulation. For mechanical moving boundary problems, the co-simulation technique works with the morpher motion model to solve the transport equations. For further details, refer to the Modeling section of the User Guide.

Most of the features in Abaqus are also available for use in a co-simulation, including nonlinear materials, nonlinear geometric effects, and contact. The following Abaqus procedures support co-simulation: static stress/displacement, dynamic (implicit and explicit), heat transfer

(steady-state and transient), coupled temperature-displacement, coupled thermal-electrical, and piezoelectric analysis. The interface domain can be modeled with most continuum, shell, and membrane elements.

Co-simulation with Abaqus supports wall boundaries, baffle boundaries, and porous baffle boundaries as coupled boundaries. In most cases, you can use wall boundaries.

The use of baffle boundaries or porous baffle boundaries is recommended for modeling the thermal resistance or porosity of a thin structure. In both cases, thermal resistance and porosity are modeled within STAR-CCM+, whereas Abaqus models the structure mechanical response to the pressure field sent by STAR-CCM+.

Applications that this FSI capability does not target include:

- Frequency domain FSI problems, including vibro-acoustic and aero-acoustic.

Note: Such problems are treated more effectively using frequency domain methods and are not suited for the partitioned solution approach. The Abaqus acoustic capability is applicable for such problems (see "Acoustic, shock, and coupled acoustic-structural analysis" Section 6.10.1 of the Abaqus Analysis User's Guide, for more information). The Harmonic Balance method (a frequency domain analysis) in STAR-CCM+ is not supported with co-simulation.

- Structures that are modeled with truss, beam, and any special-purpose elements that are wetted by the fluid.

  Such discretization has an inconsistent geometric idealization with the fluid domain. If such structural elements exist in the model, they must be linked to other structural elements that unambiguously define the actual wetted surface. Imported and exported data is then mapped to/from this wetted surface to the corresponding surface in STAR-CCM+.

- Problems involving injection molding, casting, material solidification, and superplastic forming.

  If the fluid-structure interface changes during the simulation, these applications are not suited for the FSI approach. The Abaqus Interface for Moldflow can be applicable (see the Abaqus Interface for Moldflow User's Guide, for more information).

- Problems involving rupture, penetration, and fragmentation.

  In these applications, the structural deformation can cause a single fluid domain to separate into multiple domains.

- Problems involving flow in porous media where the intention is to account for coupled effects on the fine-scale internal structure of the media (specifically, where both the fluid and solid domain overlap).

  There is no prohibition on the use of the lumped parameter porous resistance model available for regions in STAR-CCM+, but this model only has an impact on the fluid domain that STAR-CCM+ computes. Similarly, you can use the coupled pore fluid diffusion/stress procedure in Abaqus/Standard to model the porous media for the solid domain.

Preparing Your System for Co-Simulation

This section describes the key requirements for running a co-simulation case:

1. Abaqus and STAR-CCM+ must be installed.
2. You must have access to the SIMULIA Co-Simulation Engine.

Installing Abaqus and STAR-CCM+

Before you start, ensure that the following versions of Abaqus and STAR-CCM+ have been installed:

- STAR-CCM+ 9.06
- Abaqus 6.14-1

It is important that you verify that each product is properly installed before continuing. Consult the respective product Installation Guides for installation instructions and product verification.

Licensing Information

STAR-CCM+ and Abaqus are separately licensed products, and licenses are required from both SIMULIA and CD-adapco. In addition, a separate license is required for the SIMULIA Co-Simulation Engine. Information on installing and managing Abaqus licenses can be found under Licensing on the Support page at www.3ds.com/simulia. Information on installing and managing STAR-CCM+ licenses can be found in the Setting Up Licensing section of the STAR-CCM+ Installation Guide. Contact your local SIMULIA or CD-adapco support office for assistance in determining licensing requirements.

Obtaining the SIMULIA Co-Simulation Engine

The SIMULIA Co-Simulation Engine is a software component responsible for communication between Abaqus and STAR-CCM+. It is distributed and installed with the Abaqus software as dynamic link libraries. For STAR-CCM+ to communicate with Abaqus during the co-simulation, it must have access to these dynamic link libraries. To find the libraries, STAR-CCM+ works through the following three methods, in order, until it has found them:

1. STAR-CCM+ uses the path that the SIMULIA_CSE_LIBPATH environment variable sets. If this environment variable exists, and contains a valid path to dynamic link libraries, STAR-CCM+ uses these libraries. Set this path when working with libraries that have been downloaded separately to the main Abaqus installer.
2. STAR-CCM+ locates the libraries relative to the executable path or launch command that is given for Abaqus. The Executable Name property of the Co-Simulations > [Abaqus Co-Simulation 1] > Values > Abaqus Execution node takes either an Abaqus command or full path to the Abaqus executable. STAR-CCM+ can use this setting to locate the libraries. This setting is sufficient when the default dynamic link libraries contained within the Abaqus installation are used.
3. STAR-CCM+ searches the paths that are set in the LD_LIBRARY_PATH environment variable on Linux, or the PATH environment variable on Windows. Including any paths added to the program environment using the `-ldlibpath` option on the `starccm+` command.

Downloading Libraries Separately

On rare occasions, it is necessary to obtain updated versions of the dynamic link libraries from the SIMULIA download area. Instructions on downloading libraries separately can be found in the document, “SIMULIA Co-Simulation Engine dynamic libraries for fluid-structure interaction analyses using Abaqus and STAR-CCM+”. This document is located in the Dassault Systemes Knowledge Base at: www.3ds.com/support/knowledge-base.

On installing these libraries, set the environment variable SIMULIA_CSE_LIBPATH to the path where the dynamic link libraries were installed:

- On Windows platforms:

  `> set SIMULIA_CSE_LIBPATH=<`*dynamic_libs_dir*`>\lib`

- On Unix/Linux platforms:
  - For Korn and Bash shells:

```
% export SIMULIA_CSE_LIBPATH=<dynamic_libs_dir>/lib
```

- For C and Tcsh shells:

```
% setenv SIMULIA_CSE_LIBPATH <dynamic_libs_dir>/lib
```

where `dynamic_libs_dir` is the path to the directory containing the SIMULIA Co-Simulation Engine dynamic link libraries. Add these settings to your environment or profile.

Solving Fluid-Structure Interaction Problems

For co-simulation cases, the following workflow is recommended strongly:

1. Develop and test stand-alone models; that is, an Abaqus structural model and a STAR-CCM+ fluid model without any co-simulation options. Run each stand-alone simulation to convergence before attempting co-simulation.
2. Prepare the Abaqus model for co-simulation.
3. Prepare the STAR-CCM+ model for co-simulation.
4. Launch and execute the co-simulation from STAR-CCM+.

This workflow ensures that the Abaqus and STAR-CCM+ models are defined properly before you perform a co-simulation. Some additional considerations for setting up co-simulation analyses are provided at the end of this section.

Note: If either stand-alone model is improperly defined or has numerical problems, you are unlikely to obtain a solution for a co-simulation.

Developing Standalone Models

The following steps describe the general procedure for preparing the standalone Abaqus and STAR-CCM+ models.

1. Develop a standalone STAR-CCM+ fluid simulation with the following considerations:

- For mechanical co-simulation involving surface displacements, apply provisional fixed wall boundary conditions to the FSI boundary of interest. Alternatively, use a morphing motion model with a prescribed displacement or grid velocity that mimics the actual displacement.
- For thermal co-simulation analyses, apply provisional temperature boundary conditions to the wall boundary of interest.

2. Develop a standalone Abaqus structural model with the following considerations:
   - Identify the fluid-structure interfaces using surface definitions as described in "Preparing an Abaqus analysis for co-simulation", Section 17.2.1 of the Abaqus Analysis User's Guide.
   - Verify the model using provisional pressure/heat flux loads at the fluid-structure interface. Apply pressure/heat flux load magnitudes that are reasonable and similar to the expected fluid loads. You can specify the assumed pressure/heat flux loads using surface-based loading options (`*DSLOAD`, `*DSFLUX`, and `*SFILM`).
3. When you are satisfied that the standalone models are defined correctly, remove the provisional loads from each standalone model and replace the prescribed morphing with the co-simulation.

Preparing the Model in Abaqus

This section outlines how to prepare an Abaqus model for a co-simulation. You can either manually modify the Abaqus input (*.inp) files to define the co-simulation parameters, or choose to have STAR-CCM+ write them for you. For further details on co-simulation parameters in Abaqus, refer to "Preparing an Abaqus analysis for co-simulation," Section 17.2.1, of the Abaqus User's Guide.

The steps that are covered in this section are:

1. Identifying the Abaqus step for the co-simulation
2. Identifying the interface surface
3. Identifying the fields exchanged across the fluid-structure interface
4. Requesting Output of Restart Frames

Note: Steps 1, 2, and 3 can be skipped if the STAR-CCM+ property Write Abaqus *CO-SIMULATION Keywords and accompanying properties are set. For more information, see Selecting the Coupled Boundaries.

Identifying the Abaqus Step for the Co-Simulation

In Abaqus, you must identify the analysis step as a co-simulation step. To specify a co-simulation step, add the following options in the history section of the Abaqus input file:

```
*CO-SIMULATION, NAME=<name>, PROGRAM=MULTIPHYSICS
```

For further information regarding the different sections of the Abaqus input file, refer to "Defining a model in Abaqus", Section 1.3.1 of the Abaqus Analysis User's Guide.

Communication between Abaqus and STAR-CCM+ is initiated during the co-simulation step and is terminated when either program ends the co-simulation. Abaqus terminates the co-simulation event when the end of the analysis step is reached or when the analysis cannot proceed any further; for example, due to convergence problems.

Identifying the Interface Surface

In Abaqus, identify the co-simulation interface surface and the fields that are exchanged on this surface. For fields imported into Abaqus from STAR-CCM+, use

```
*CO-SIMULATION REGION, TYPE=SURFACE, IMPORT
<surface>, <field_identifier>
<surface>, <field_identifier>
```

For fields exported from Abaqus to STAR-CCM+, use

```
*CO-SIMULATION REGION, TYPE=SURFACE, EXPORT
<surface>, <field_identifier>
<surface>, <field_identifier>
```

where `<surface>` is an element-based surface (see "Element-based surface definition," Section 2.3.2 of the Abaqus Analysis User's Guide). The `<field_identifier>` designates the field that is exchanged during the co-simulation, as described in the next section.

Note: For each coupling zone defined in the STAR-CCM+ simulation, specify a single unique surface name. This surface can be defined as the union of multiple disconnected surfaces. The same surface name can be specified on multiple lines when defining the imported/exported fields.

Identifying the Fields Exchanged Across the Fluid-Structure Interface

The tables below identify fields that can be exchanged between Abaqus and STAR-CCM+ during a co-simulation. The first table lists the fields which are imported (incoming) or exported (outgoing) for each available Abaqus analysis procedure type.

Abaqus Procedure	Incoming Fields			Outgoing Fields		
	CF	CFL	CFILM	U	V	NT
Static	●	-	-	●	-	-
Dynamic (Implicit)	●	-	-	●	●	-
Dynamic (Explicit)	●	-	-	●	●	-
Heat Transfer	-	●	●	-	-	●
Coupled Temperature-Displacement	●	●	●	●	-	●

The following table defines the correspondence between the Abaqus field and the STAR-CCM+ field function as well as the dimensionality of the fields.

Abaqus			STAR-CCM+	
Field Name	Abaqus Field ID	Dimension-ality	Field Function Names	Dimension-ality
Displacement	U	L	Nodal Displacement	L
Velocity	V	LT^{-1}	Velocity	LT^{-1}
Traction Vector	CF	F	Wall Shear Stress, Static Pressure, or Absolute Pressure	FL^{-2}
Temperature	NT	θ	Temperature	θ

Abaqus			STAR-CCM+	
Field Name	Abaqus Field ID	Dimension-ality	Field Function Names	Dimension-ality
Heat Flux	CFL	JT^{-1}	Boundary Heat Flux	$JT^{-1}L^{-2}$
Film Properties (heat transfer coef-ficient, ambient temperature)	CFILM	$JT^{-1}L^{-2}\theta^{-1}$ θ	Boundary Heat Flux, Local Heat Transfer Ref-erence Temperature or Specified Y+ Heat Trans-fer Reference Tempera-ture	$JT^{-1}L^{-2}$ θ

Requesting Output of Restart Frames

Optionally, you are advised to write restart frames during the Abaqus analysis so that the co-simulation can be restarted in case of termination of the analysis due to a power outage, etc. The time points in which Abaqus saves a restart frame must be synchronized with the Auto Save settings in STAR-CCM+ to provide a consistent state from which to restart the co-simulation.

In the co-simulation step definition, add the following keyword to write `n` equally spaced (in time) restart frames to the Abaqus restart file during the co-simulation:

```
*RESTART, WRITE, NUMBER INTERVAL=<n>
```

Additionally, to save only the latest restart frame on the restart file (analogous to setting Max Autosaved Files to 1 in STAR-CCM+) use:

```
RESTART, WRITE, NUMBER INTERVAL=<n>, OVERLAY
```

Preparing the Model in STAR-CCM+

This section outlines how to prepare the STAR-CCM+ simulation for co-simulation with Abaqus.

1. Activating the Abaqus Co-Simulation Model
2. Selecting the Coupled Boundaries
3. Specifying Exported and Imported Fields
4. Specifying the Mapper Settings
5. Setting the Abaqus Execution Properties
6. Setting the Coupling Algorithm Properties

7. Selecting the Units for Imported and Exported Fields

Activating the Abaqus Co-Simulation Model

In STAR-CCM+, you enable co-simulation with Abaqus by activating the Abaqus Co-Simulation model for the physics continuum. In a simulation containing multiple physics continua, activate the Abaqus Co-Simulation model for the physics continuum that the coupled region uses.

This procedure assumes that you have already selected appropriate material and flow models for the simulation.

To activate the Abaqus Co-Simulation model:

1. Right-click the relevant physics continuum node under Continua and select Select Models...
2. In the *Physics Model Selection* dialog:
   a. Select Co-Simulation from the *Optional Physics Models* group box.

b. Select Abaqus Co-Simulation from the *Co-Simulation* group box.



Specifying Co-Simulation Conditions and Values

In STAR-CCM+, you specify the settings for co-simulation in a co-simulation session within the Co-Simulations node.

The Co-Simulations node becomes available in the simulation tree when you activate the Abaqus Co-Simulation model. This node is automatically populated with a co-simulation session that has the default name, Abaqus Co-Simulation 1.

The [Abaqus Co-Simulation 1] node allows you to specify co-simulation global settings, as well as to define and set up co-simulation coupling zones.



In STAR-CCM+, a coupling zone defines a set of boundaries where fields are exchanged with the Abaqus model. Co-Simulation with Abaqus supports multiple coupling zones, that allow you to apply different conditions to different sets of coupled boundaries.

Co-simulation global settings, such as the Abaqus Execution properties, the units used by Abaqus, and the coupling algorithm, are general conditions and values that apply to all the coupling zones.

Instructions for specifying co-simulation global settings and zone settings are outlined in the following sections.

For the [Abaqus Co-Simulation 1] node, the following properties and pop-up menu items are available.

Abaqus Co-Simulation Properties

Co-Simulation Type	Displays the type of co-simulation (Read-only). For co-simulation with Abaqus, Co-Simulation Type is set to Abaqus Co-Simulation.

Abaqus Co-Simulation Pop-Up Menu

Terminate Abaqus	Closes the connection between STAR-CCM+ and Abaqus.

Specifying Co-Simulation Zone Settings

In STAR-CCM+, a coupling zone defines a set of coupled boundaries that exchange data with a specified Abaqus surface.

Co-Simulation with Abaqus supports multiple coupling zones, that allow you to apply different conditions on different sets of coupled boundaries. The fields that are exchanged between the codes, and the mapper settings, are set independently for each coupling zone. Other general settings, such as the Abaqus execution options and the coupling algorithm, apply to all the zones defined in the STAR-CCM+ simulation.

This section outlines the recommended workflow for setting up a coupling zone. For general co-simulation parameters, see Specifying Co-Simulation Global Settings.

Zone Properties

Index	Specifies the unique index that identifies a zone. (Read-only)
Zone Type	Specifies the type of coupling zone. The available option is Surface Zone.

Selecting the Coupled Boundaries

In a co-simulation with Abaqus, STAR-CCM+ boundaries exchange data with the corresponding surfaces of the Abaqus model.

You can specify both the STAR-CCM+ coupled boundaries and the corresponding Abaqus surfaces in your STAR-CCM+ simulation. Alternatively, you can define the Abaqus coupled surfaces directly in the Abaqus input file.

For each coupling zone in STAR-CCM+, select the coupled boundaries and define the corresponding Abaqus surface.

To assign a coupled boundary to a zone:

1. Expand the Regions > [Region] > Boundaries node.
2. Expand the boundary that you want to assign to a coupling zone.
3. Select the Physics Conditions > External Code Coupling Specification node and set External Code Coupling Specification to the coupling zone of interest.



STAR-CCM+ automatically updates the zone Coupled Boundaries condition with the selected boundary.



To define the Abaqus surface that couples to the selected boundaries:

1. Expand the Co-Simulations > [Abaqus Co-Simulation 1] > Values node.
2. Select the Abaqus Execution node and activate Write Abaqus *CO-SIMULATION Keywords.
3. Expand the relevant [Abaqus Co-Simulation 1] > Zones > [Zone] node.
4. Select the Values > Abaqus *CO-SIMULATION REGION Name and use the Set Name property to specify a unique name for the Abaqus surface.

If you are not using the Write Abaqus *CO-SIMULATION Keywords option, specify the surface name directly in the Abaqus .inp file, under *CO-SIMULATION REGION (see Identifying the Interface Surface). Specify a single unique surface name for each coupling zone defined in the STAR-CCM+ simulation. In any case, when you are using multiple coupling zones, set the [Zone] > Values > Abaqus *CO-SIMULATION REGION Name property in STAR-CCM+, to assign each Abaqus surface to the corresponding STAR-CCM+ zone.

To verify that the STAR-CCM+ coupled boundaries and Abaqus co-simulation surface are co-located, import the Abaqus model and create a scene to visually inspect the fluid-structure interface.

Specifying Exported and Imported Fields

Specify the fields that are exchanged between Abaqus and STAR-CCM+ during the co-simulation.

If you activated the Write Abaqus *CO-SIMULATION Keywords option, STAR-CCM+ will automatically write the fields definition keywords in the Abaqus .inp file, according to the import/export conditions specified for the coupling zones.

If you are not using the Write Abaqus *CO-SIMULATION Keywords option, specify the fields imported/exported by Abaqus in the Abaqus input file. The field functions that you specify for import/export in STAR-CCM+ must match with the exported and imported field types that you set in the Abaqus input file. Fields that are set for export in Abaqus must be set for import in STAR-CCM+. Likewise, fields that are set for export in STAR-CCM+ must be set for import in Abaqus. See section, Identifying the Fields Exchanged Across the Fluid-Structure Interface.

Thermal Coupling

In co-simulation analyses involving thermal coupling at the fluid-solid interface, STAR-CCM+ exports boundary heat flux to Abaqus and imports temperature data.

To specify the imported/exported fields in STAR-CCM+:

1. Expand the relevant coupling zone Values node (Co-Simulations > [Abaqus Co-Simulation 1] > Zones > [Zone] > Values).
2. Select the Fields Exported To Abaqus node and set Exported Field Function to Boundary Heat Flux.

   When exporting Boundary Heat Flux with reference temperature data, you can choose from the following reference temperature field functions:

   - Local Heat Transfer Reference Temperature
   - Specified Y+ Heat Transfer Reference Temperature

   For most FSI applications, use the specified y+ heat transfer reference temperature. If the simulation involves radiation, use the local heat transfer reference temperature instead. You can define the y+ value that is used for the specified y+ heat transfer reference temperature using the `Specified Y+ Heat Transfer Coefficient` field function, which can be found under the Tools > Field Functions node in STAR-CCM+.
3. Select the Fields Imported From Abaqus node and set Imported Field Function to Temperature.

   STAR-CCM+ automatically updates the coupled boundary nodes to use the thermal data from Abaqus. The Thermal Specification physics condition for the coupled boundaries is automatically set to Temperature; the Method for the Static temperature physics value is automatically set to Abaqus Co-Simulation.

Mechanical Coupling

In co-simulation analyses involving mechanical coupling at the fluid-solid interface, STAR-CCM+ exports traction data (pressure and wall shear stress) to Abaqus and imports nodal displacements.

STAR-CCM+ uses morphing motion to move the boundaries in response to the imported nodal displacements.

Activate the morpher and assign it to the relevant region:

1. Right-click the Tools > Motions node and select New > Morphing.
2. Select the Regions > [Co-simulation region] > Physics Values > Motion Specification node and set Motion to Morphing.

To specify the imported/exported fields:

1. Expand the relevant coupling zone Values node (Co-Simulations > [Abaqus Co-Simulation 1] > Zones > [Zone] > Values).
2. Select the Fields Exported To Abaqus node and set the Exported Field Function property.
3. Select the Fields Imported From Abaqus node and set Imported Field Function to Nodal Displacements.

   STAR-CCM+ automatically updates the coupled boundary nodes to use the displacement data from Abaqus. The Morpher physics condition for the coupled boundaries is automatically set to Abaqus Co-Simulation.

Specifying the Mapper Settings

To control the data mapping between STAR-CCM+ and Abaqus, use the Mapper Settings node.

Fields are mapped between a boundary on the fluid grid and a surface on the structural mesh. A separate mapping is done for each type of exchanged field (displacement and temperature versus traction and heat flux) and each direction of mapping (imported versus exported). For fields exported to Abaqus, the fluid boundary acts as a source and the structure surface acts as a target. For fields imported from Abaqus, the structure surface acts as a source and the fluid boundary acts as a target.

The position and extent of the coupled fluid boundary and structure surface must match at the beginning of the co-simulation. Possible causes of a mismatch in position are: discretization differences, an error identifying the units that Abaqus uses, or the use of an incorrect reference configuration in the initial search in the mapper. Each code reports initial (target) search failures for its coupled boundary/surface separately (in the *Output* window in STAR-CCM+ and the `.msg`/`.sta` files in Abaqus).

Initial Search and Reference Configuration

The mappers make an initial search using a reference configuration. The reference configuration can be set to either the current or original coordinates. If you mesh both the fluid and the structure using the same geometry, their original configurations match, and you can set the reference configuration to the original coordinates for the first and subsequent (restart) co-simulations. If you mesh either the fluid or the structure using a different configuration from the other, use the current coordinates for the initial search in the mapper. An example of this latter scenario is when the structure undergoes some deformation for pre-loading, and the fluid mesh is generated based on the deformed structure.

Current coordinates can be used in all cases. However, when restarting a co-simulation, it is best to use the original coordinates when possible. Mapping using updated (current) coordinates perturbs the initial search and can cause small jumps in the solution in the first time step.

Some mismatch in the coupled boundaries is expected due to discretization differences or errors when the coupled boundaries are not conformal. To account for this mismatch automatically, the mappers search a small distance from the source surface to find targets during the initial search. (The STAR-CCM+ mappers search one-half the maximum source boundary/surface face size; the Abaqus mappers search one-half the average source boundary/surface face size). Remesh the fluid and/or structure to resolve gross mismatches of the coupled boundaries; increase the proximity tolerances to resolve small mismatches. Arbitrarily increasing the proximity tolerance can mask errors in the model setup by allowing the search to find neighbors that are not in close proximity. Matching neighbors that are not in close proximity leads to incorrect results. The normals of coupled boundaries are also checked to ensure that surfaces face each other. In cases where initial search failures occur due to discretization differences in areas of large curvature variation, you can deactivate the normal check.

Remeshing in STAR-CCM+

If the deformation of the mesh during the analysis requires the fluid to be remeshed before continuing, the choice of the Mapper Code will define the workflow:

- If Mapper Code is set to STAR-CCM+, then the analysis can be continued without any intervention.
- If Mapper Code is set to Abaqus, then the co-simulation will be terminated and must be restarted using the Current Coordinates option for the mapper Reference Configuration. See Restarting From a Previous Abaqus Simulation

Mapping Rotations

When mapping displacements, an error is introduced that is a function of the distance between the source and target and the amount of rotation of the source. Activate the Map Rotations option on the corresponding STAR-CCM+ mapper to eliminate this error by adding a rotational correction to the mapped displacement.

In cases where the source undergoes large rotations, the Map Rotations option should be activated to preserve the shape of the fluid surface as the structure rotates. The Map Rotations option should be deactivated in cases where the structure is modeled using shell elements and the deformation causes the physical shell surface (the surface of the solid body represented by the shell) to self-intersect during the analysis.

Mapper Settings Properties

Mapper Code	Chooses whether STAR-CCM+ or Abaqus maps the exchanged data.	
Reference Configuration	Chooses the configuration to use for the initial search in the mapper.	
	Current Coordinates	Use the current (deformed) mesh configuration to perform the initial search.
	Original Coordinates	Use the original (undeformed) mesh configuration to perform the initial search.
Proximity Tolerance	The relative search distance to use in the initial search in the mapper. If STAR-CCM+ mappers are used the distance is relative to the maximum face size on the coupled boundary source surface. If Abaqus mappers are used the distance is relative to the average face size on the coupled boundary source surface.	
Check Normals	When set, checks that the STAR-CCM+ boundary and Abaqus surface are oriented toward one another in the initial search in the mapper.	
	Activated	Checks orientations.
	Deactivated	No check is made.
Map Rotations	When set, the mappers correct displacements for target points that are offset from the source.	
	Activated	Accounts for rotations when mapping displacements.
	Deactivated	Ignores rotations.

Specifying Co-Simulation Global Settings

In STAR-CCM+, you specify the general co-simulation parameters, such as the Abaqus execution options, the units used by Abaqus for the imported/exported fields, and the coupling algorithm, within the co-simulation session [Abaqus Co-Simulation 1].

This section outlines the recommended workflow for specifying general co-simulation parameters, that are independent from the coupling zones set up. For specifying co-simulation zone settings, see Specifying Co-Simulation Zone Settings.

Setting the Abaqus Execution Properties

The Abaqus Execution node allows you to set the properties for launching and running the Abaqus job, and to set STAR-CCM+ to overwrite the Abaqus co-simulation keywords in the Abaqus input file.



Specifying the Abaqus Input File

Specify the (`*.inp`) input file that is prepared for the co-simulation by clicking the [...] button next to the Input File property. Alternatively, specify the file name and path by typing directly into the property box. You can enter either the full path to the file or a path relative to the STAR-CCM+ simulation file.

Restarting from a Previous Abaqus Simulation

An Abaqus restart must be performed when continuing from a previous Abaqus analysis, whether the analysis:

- terminated due to an abort, for example in case of convergence failure in the Abaqus solver or a power failure
- completed successfully (either a co-simulation analysis or an Abaqus stand-alone analysis to pre-load the structure)
- terminated under STAR-CCM+ control when the fluid mesh was remeshed in the middle of an analysis. (Remeshing the fluid analysis only forces termination of Abaqus when Abaqus mappers are selected under Mapper Settings > Mapper Code).

In order to restart, both the STAR-CCM+ and Abaqus models must be in the same configuration at the point in time of the restart. Significant differences in configuration (such as a mismatch in the FSI boundary mesh positions) results in mapping errors (reported in the STAR-CCM+ log window) or mapping inaccuracies. A good practice is to do both of the following steps when running an analysis:

- Save restart frames at evenly spaced points in the Abaqus analysis. Do this using the `*RESTART, WRITE, NUMBER INTERVAL=<n>` command in the original Abaqus `.inp` file.
- Synchronize the auto-save feature in STAR-CCM+ to save at the same time intervals.

For analyses that were terminated before the end of the co-simulation analysis this practice ensures that there is a matching restart frame in Abaqus corresponding to a saved `.sim` file in STAR-CCM+.

The mapping configuration must be chosen to suit the type of restart. For co-simulations where both the fluid and the structure are run together in all preceding analyses, set the mapper to use the original coordinates of the coupled FSI boundary. Set this option under Mapper Settings > Mapper Configuration. In certain cases, there is a mismatch in the configuration that is represented by the original coordinates of each mesh. Examples of such cases are:

- where the fluid analysis is remeshed
- the Abaqus analysis is restarted from a stand-alone pre-loading step in Abaqus

In these cases, the current coordinates option must be selected under Mapper Settings > Mapper Configuration.

Define a restart from a point at which the models are in equilibrium. This condition is automatically satisfied when both analyses are started at the same point in time for the restart. If Abaqus is first run stand-alone to pre-load the structure and/or STAR-CCM+ is run with a fixed FSI boundary to develop an initial flow field, tractions exported from STAR-CCM+, and loads/boundary conditions applied in Abaqus, may need to be ramped to allow a smooth transition to an equilibrium state at the fluid-structure interface.

To set up a restart analysis:

1. Create an Abaqus `.inp` file with `*RESTART, READ, ...`, corresponding to the restart frame you are restarting from. If you wish to restart from a time point other than the time of completion of the previous step, include the `END STEP` parameter in the `*RESTART` definition. Also, include a new `*STEP` definition with the time duration of the new analysis and any additional structural loading or boundary conditions. Check the Abaqus `.msg` (Abaqus/Standard) or `.sta` (Abaqus/Explicit) files of the previous analysis for information about the restart frames that were written.
2. Modify the Abaqus Execution properties in STAR-CCM+ (Current Job Name, Restart > New Step, Old Job Name, Input File). The Abaqus restart files from the analysis defined by Old Job Name are read at the start of the co-simulation and must be present in the same directory as the Input File of the restart analysis.
3. Set the Mapper Settings > Reference Configuration to Current Coordinates if the original configurations of the structure and fluid do not match.
4. Set the STAR-CCM+ stopping time to the current time in STAR-CCM+ plus the Abaqus restart step time. The co-simulation time corresponds to the Abaqus step time and starts from zero each time a co-simulation is restarted.
5. Save the `.sim` file to a new name and click Run in STAR-CCM+.

Specifying the Executable Name and Command Line Options

For the Executable Name property, you can enter the full path to the Abaqus executable, or else enter the command that is used to launch Abaqus. For the latter option, include the directory containing the Abaqus command in the system path.

Some examples of these entries are:

- Linux
  - `abq6141` -- Abaqus command without the path prepended; the `Commands` directory in the Abaqus install directory must be in the system path

    - `/apps/abaqus/6.14-1/code/bin/abq6141` -- full path to the Abaqus executable
- Windows
    - `abq6141` or `abq6141.bat` -- Abaqus command without the path prepended; the `Commands` directory in the Abaqus install directory must be in the system path. If `.BAT` is included in the `PATHEXT` environment variable, you can omit the `.bat` extension
    - `C:\SIMULIA\Abaqus\code\bin\abq6141.exe` -- full path to the executable.

If you wish to use more command line options, these options can be entered in the Additional Command Line Options property. For example, entering `-double both -output_precision full -user usersub.f` would run Abaqus/Explicit in double precision, output results to the output database (`*.odb`) file in double precision, and specify a user subroutine to be used in the analysis.

Note: Separate multiple options by empty spaces.

In addition, some command line options are not compatible with the co-simulation analysis mode.

Abaqus Execution Properties

Write Abaqus *CO-SIMULATION Keywords	Instructs STAR-CCM+ to overwrite any existing *CO-SIMULATION keywords in the `*.inp` file. You are advised to let STAR-CCM+ write the keywords to the Abaqus input file to maintain consistent settings in both models.	
	Activated	Overwrites existing keywords.
	Deactivated	Does not overwrite existing keywords.
Abaqus *CO-SIMULATION Step Number	Specifies the step in the Abaqus input file where the co-simulation keywords are inserted.	
Current Job Name	Defines the current arbitrary job name for the Abaqus Co-Simulation output files. Must be different from Old Job Name.	
Restart	Selects the restart method.	
	None	No restart.
	New Step	The Abaqus Input File contains a new step definition that is used for the current simulation. The Abaqus restart file, Old Job Name, is read at the start of the co-simulation.
Old Job Name	Specifies the job name from a previous co-simulation for restart.	
Input File	Selects the Abaqus input file for the co-simulation.	
Executable Name	Specifies the name of the Abaqus executable or command.	
Additional Command Line Options	Specifies more command line options for use with the Abaqus executable.	
Number of CPUs	Specifies the number of processors that Abaqus uses when the co-simulation is run. Must be entered as an integer. If the number of processes for Abaqus is entered using the command-line argument, `-np`, this property is ignored.	

Host List	Specifies the hosts and CPUs per host on which Abaqus is run (for example, host1:2,host2:4).
Remote Shell	Specifies the type of shell that is used to log in to the remote machine. It is recommended that you use the default, `rsh`, for unsecured networks (usually within a firewall). The other supported option is `ssh`, and must be used when all connections are secure (for example, when `rsh` is blocked). If using `ssh`, make sure that you set it up so that it does not prompt for a password.
Remote User	Specifies the username for the remote shell login. If this option is left blank, the local username is used.

Setting Up the Coupling Algorithm

To control the coupling method that is used in the co-simulation, use the Coupling Algorithm node.



Selecting the Coupling Scheme

The Coupling Scheme property controls the sequence of exchanges in the co-simulation. For one-per-time-step exchange, choose the Explicit scheme. This method results in weak coupling, appropriate for problems in which

the effect of the structure on the fluid is substantially larger than the effect of the fluid on the structure. A sufficiently small time-step must be chosen to maintain stability with this scheme.

To pass data between Abaqus and STAR-CCM+ more than once per time-step, choose the Implicit scheme. Doing so ensures strong coupling (with several coupled computations performed within one time-step), and allows you to use larger time-steps than the Explicit scheme. STAR-CCM+ adds a Field Exchange Controls sub-node under the Abaqus Co-Simulation solver node, where you set expert control properties for the coupling solver.

Selecting the Time-Marching Sequence

The Time-Marching Sequence property defines whether STAR-CCM+ and Abaqus are run simultaneously or sequentially. For implicit coupling, only sequential runs are possible.

To run both codes simultaneously, select Concurrent from the drop-down menu. This method is the least stable method. However, if the computation time per time-step is equal and the partition of processors is balanced, then running the codes simultaneously is faster than running sequentially if resources are available to run each code on separate hosts.

To run the codes sequentially during the co-simulation, select either Abaqus Leads or STAR-CCM+ Leads from the drop-down menu. The stability and computational costs of both the Abaqus Leads and STAR-CCM+ Leads are identical.

Selecting the Coupling Negotiation Method

The Coupling Negotiation Method in STAR-CCM+ sets the method by which the coupling step size is determined. If the Constant method is selected, the Constant Coupling Time Step property must be defined.

Selecting an Appropriate Time-Step Size

The time-step size determines the time that it takes to solve a case, the stability of the solution and its ability to capture all physical effects in the case. The following are a list of recommendations to help you select an appropriate time-step size:

- For steady-state cases with static loads, such as a deformed wing under fluid loads, the effect of time-step on accuracy is minimal. Use time-steps larger than the fluid or structural time scale for the best results.
- For cases with impact loading, such as fluid loads from sloshing, choose a time-step to resolve the peak load.

- For cases with oscillations or vibrations, such as releasing a pre-stressed cantilever, set a time-step of one hundredth (1/100) of the expected period of oscillations to capture the behavior of the system adequately.

To test if the time-step is sufficiently small for accuracy, run a simulation with a time-step of half the original value, and observe if there is a change in results.

Coupling Algorithm Properties

Coupling Scheme	This coupling control defines whether data is passed more than once per time-step in the co-simulation.	
	Explicit	Sets a one per time-step exchange between Abaqus and STAR-CCM+.
	Implicit	Sets a more than once per time-step exchange between Abaqus and STAR-CCM+.
Time-Marching Sequence	This coupling control defines serial or parallel coupling in the co-simulation.	
	Abaqus Leads	Sets Abaqus and STAR-CCM+ to run sequentially with Abaqus leading the co-simulation.
	STAR-CCM+ Leads	Sets Abaqus and STAR-CCM+ to run sequentially with STAR-CCM+ leading the co-simulation.
	Concurrent	Sets Abaqus and STAR-CCM+ to run simultaneously during the co-simulation. Only available if the coupling scheme is set to explicit.

Coupling Negotiation Method	This coupling control defines how the coupling step size is determined during the co-simulation.	
	Constant	Sets a constant user-defined coupling step size.
	Min	Sets the minimum of the coupling step sizes suggested by Abaqus and STAR-CCM+.
	Max	Sets the maximum of the coupling step sizes suggested by Abaqus and STAR-CCM+.
	Import	Allows Abaqus to export a coupling step size for STAR-CCM+ to use.
	Export	Allows Abaqus to import a coupling step size that STAR-CCM+ suggests.
	User Defined	Allows STAR-CCM+ and Abaqus to use different coupling step sizes. This feature is useful if the analyses are running at different time scales. When selected, an additional node appears under Solvers > Abaqus Co-Simulation. See User-Defined Negotiation Method Controls.
Constant Coupling Time Step	This coupling control sets the value of the coupling step size for a Constant Coupling Negotiation Method.	

Selecting the Units for Imported and Exported Fields

To specify the units for data that is imported from and exported to Abaqus, use the Abaqus Units node. For imported fields, the unit specifications tell STAR-CCM+ the units of the incoming data. For exported fields, the unit specifications ensure that data sent to Abaqus is in the correct units. In each

case, these properties depend on which consistent unit system was used when defining the Abaqus model. These settings do not affect the units that are used in STAR-CCM+ calculations.

If a unit used in the Abaqus model is not a pre-defined unit in the STAR-CCM+ units manager, you must create the unit before it appears and can be selected in the Abaqus Units node.



Abaqus Units Properties

Length Unit	Selects the unit of length that is used in the Abaqus model.
Stress Load Unit	Selects the unit of stress that is used in the Abaqus model.
Temperature Unit	Selects the unit of temperature that is used in the Abaqus model.
Heat Flux Unit	Selects the unit of heat flux that is used in the Abaqus model.

Specifying Exchange Conditions and Solver Settings

To set the properties for controlling the coupling method that is used in the co-simulation, use the Solvers > Abaqus Co-Simulation sub-nodes. These properties can be changed during the analysis. You can control:

- field exchange settings for implicit coupling
- exported traction field settings for implicit and explicit coupling
- grid flux settings for implicit and explicit coupling

- user-defined negotiation method for explicit coupling



Set the Abaqus Co-Simulation solver Verbosity to High to output information about the exchange of data, co-simulation time, and solver timing in the STAR-CCM+ output window.

Controlling Field Exchange

To modify the field exchange settings for implicit coupling, use the Field Exchange Controls node. STAR-CCM+ adds this node to the Abaqus Co-Simulation solver when you select the implicit coupling scheme for the coupling algorithm.



The Inner Iterations Per Exchange property controls the number of inner iterations STAR-CCM+ performs for each field data exchange within one coupling time-step. The data exchange for each time-step is terminated once either the Field Convergence Tolerance or the maximum number of inner iterations, set in the stopping criteria properties, is reached.

To accelerate convergence of the implicit loop, you can modify the under-relaxation factor for the imported data, which affects the solution update at each iteration. Choosing a Constant imported field URF applies a fixed parameter in computing each iterate while an Adaptive type optimizes the solution update by controlling the under-relaxation factor based on the residual. Limits on the adaptive URF are set using the Minimum Adaptive URF and Maximum Adaptive URF values. In cases where you set the URF type to Adaptive, the Imported Field URF represents the initial URF value. Otherwise, it represents the constant URF value.

Field Exchange Controls Expert Properties

Inner Iterations Per Exchange	Sets the number of inner iterations per field exchange in implicit coupling.	
Minimum Number of Exchanges	Sets the lower bound of the total number of exchanges in a time step.	
Field Convergence Tolerance	Sets the convergence tolerance for terminating field exchanges in a time-step.	
Imported Field URF Type	Sets the method to under-relax imported data.	
	Adaptive	Automatically scales the URF value to optimize solution update at each iteration.
	Constant	Uses a constant URF value to update the solution at each iteration.
Minimum Adaptive URF	Sets the minimum adaptive under-relaxation value for the imported data.	
Maximum Adaptive URF	Sets the maximum adaptive under-relaxation value for the imported data.	
Imported Field URF	Sets the constant under-relaxation value or initial value for the Adaptive under-relaxation type.	

Controlling Exported Traction Fields

To modify the exported traction field settings, use the Exported Traction Field

Control node.



Some co-simulation problems exhibit numerical instabilities resulting from high inertia of added-mass, high compliance of the structure, mismatch in initial conditions of the structure and fluid, or other sources. In such cases, elect to suppress or control the destabilizing loading conditions by progressively increasing and/or limiting the exported traction load. Although doing so would generally allow for the use of a much larger coupling step size while maintaining numerical stability, the stability would be achieved at the expense of accuracy. The Pressure Ramping Parameters, Traction Clipping Parameters, and Traction Field Multiplier can be used to modify the traction field sent to Abaqus in a difficult co-simulation problem.

Specifying Pressure Ramping Parameters

The Pressure Ramping Parameters are values that can be employed to increase the pressure component of the exported traction load to Abaqus progressively. This is particularly useful at the start of a co-simulation where the structure is initially unloaded and an initial flow field in the fluid was calculated in a standalone analysis with a rigid structure. Two numerical values are required for the pressure ramping parameters: **t_zero** and **t_couple**. During **t<t_zero,** the STAR-CCM+ exported pressure and shear stress are zero. Between **t_zero<t<t_couple** the exported pressure value is smoothly blended so that by **t ≥ t_couple** the exported pressure equals the actual pressure STAR-CCM+ computed. In the same period, the exported shear stress vector is the same direction as the shear stress vector

STAR-CCM+ computed, but the magnitude is smoothly blended. Similar to the exported pressure, at **t ≥ t_couple** the exported shear stress equals the shear stress STAR-CCM+ computed.

The pressure ramping parameters are entered into the property box as `[t_zero, t_couple]`. If the co-simulation time is chosen, set **t_zero** to `0.0`. If STAR-CCM+ physical time is chosen, set **t_zero** to the physical time where the pressure is zero.



Specifying Traction Clipping Parameters

Optionally, limit values in the fluid domain to reflect the actual operating conditions. Values can be enforced for the maximum and minimum pressures that are exported for the pressure component of the traction load, as well as the peak shear stress magnitude. These values are referred to as **p_max**, **p_min** and **shear_max**, respectively, and are constant throughout

the co-simulation. When setting the pressure parameters, as shown in the following diagram, **p_max** must be greater than **p_min** and both values can be negative.



For the shear stress clipping parameter, **shear_max**, the direction of the vector is kept the same and just the magnitude is clipped, as shown in the following.



An example of when these limiting values are useful is in simulating multiphase flow of gas and liquid in a pressurized tank. Consider a case where the known peak operating pressure in the tank is 70 bar. Rather than export a computed pressure that is potentially unrealistic or can result in an unstable co-simulation (especially in the early stages of the co-simulation), you can enforce the peak operating pressure (**p_max**).

The traction clipping parameters are entered into the property box as `[p_max, p_min, shear_max]`.

Specifying a Traction Field Multiplier

A Traction Field Multiplier can be specified to define a spatially and/or time varying function to modify the traction exported to Abaqus.

The multiplier is applied before the Pressure Ramping Parameters or Traction Clipping Parameters are applied. An existing or user-defined scalar field function can be used as the multiplier. Some uses of a multiplier include applying a factor of safety to the load in a one-way coupled analysis or masking pressures in regions deemed to be undergoing contact.

Exported Traction Field Controls Expert Properties

Pressure Ramping Time Scale	Sets the time scale for the pressure ramping time parameters.	
	Co-Simulation Time	Employs pressure ramping that is based on co-simulation time.
	Star-ccm+ Time	Employs pressure ramping that is based on STAR-CCM+ time.
Pressure Ramping Parameters	Sets the pressure ramping values in the early stages of the co-simulation.	
Traction Clipping Parameters	Specifies values for traction clipping.	
Apply Traction Field Multiplier	Activates/deactivates the traction field multiplier.	
	Activated	STAR-CCM+ applies the traction field multiplier for the Traction Field Multiplier property.
	Deactivated	The traction field multiplier option is not used.
Traction Field Multiplier	Specifies a scalar field function to be multiplied with the traction field before export to Abaqus.	

Controlling Grid Flux Terms at the Fluid-Structure Interface

To modify the grid flux under-relaxation settings, use the Grid Flux Controls node.



The grid flux is the volume swept out by the movement of a face in one time-step, which is divided by the time-step. This property is effectively the normal velocity of the face that is multiplied by its area. The grid flux triggers the response of the fluid to the motion of the structure. If you are searching for a steady-state solution ("static" solution in FEA), the velocity of the structure within the fluid simulation can be assumed to be zero. The history of how the structure moved to get to the steady-state location is not important, just the location at which it eventually arrives. Therefore, the grid flux can be neglected and you can still arrive at the correct solution.

In transient problems ("dynamic" in FEA), it is advantageous to neglect the grid flux terms in the initial steps of the simulation. The initial conditions are far from any realistic conditions and the grid flux terms can be unrealistic and destabilizing. Ignoring the grid flux allows the fluid and structure solutions to evolve into a more realistic state. You can then choose to switch on the grid flux after the solution has settled down.

The grid flux terms can be ignored by ticking the Ignore Grid Flux Terms checkbox in the *Properties* window.

Another way to mitigate the destabilizing effects of the grid flux is using the Grid Flux Under-Relaxation Factor. A value of zero for the under-relaxation factor is effectively the same as ignoring the grid flux terms. A value of 1.0 (the default value) causes the grid flux terms, dictated by the motion of the structure, to be fully utilized in the simulation. Entering an intermediate value for the under-relaxation factor causes STAR-CCM+ to use an applied

grid flux. The applied grid flux is a blend of the grid flux at the current time-step and the grid flux at the previous time-step. An under-relaxation factor of less than 0.5 is not recommended for problems requiring dynamic accuracy, as it is not time-accurate. However, to allow the fluid and structure simulations to settle down, it can be useful to under-relax with values much less than 0.5 in the initial stages of the co-simulation. If you observe large oscillations in the pressure at the fluid-structure interface from time-step to time-step, but the coupling time-step size cannot be reduced due to CPU considerations, then using the grid flux under-relaxation factor is recommended.

Grid Flux Controls Expert Properties

Ignore Grid Flux Terms	Sets whether STAR-CCM+ ignores grid flux terms.	
	Activated	Ignore grid flux terms from mesh motion that correspond to imported Abaqus displacements.
	Deacti-vated	Do not ignore grid flux terms from mesh motion that corresponds to imported Abaqus displacements.
Grid Flux Under-Relax-ation Factor	Governs the extent to which the old grid flux terms are supplanted by the newly computed grid flux terms at each iteration. The default value is 1.0, which causes the new grid flux terms to supplant the old terms at each itera-tion.	

User-Defined Negotiation Method Controls

In some co-simulations cases, it can be expensive to run STAR-CCM+ and Abaqus at the smallest time scale of the two codes. The user-defined coupling negotiation method allows you to specify different coupling step sizes for each code. This coupling negotiation method is useful for loosely coupled problems.

Setting the Coupling Negotiation Method to User-Defined activates the Solvers > Abaqus Co-Simulation > User-Defined Negotiation Method Controls node. See Setting Up the Coupling Algorithm.



The coupling step sizes for each code, and the method that is used to calculate the exported fields, is controlled from the node properties.

User-Defined Negotiation Method Controls Properties

STAR-CCM+ Coupling Time Step	Specifies the coupling step size for the STAR-CCM+ analysis.	
Abaqus Coupling Time Step	Specifies the coupling step size for the Abaqus analysis.	
Exported Field Calculation	Specifies the method that is used to calculate the field values that are exported to Abaqus.	
	Average	Averages the field values over all time-steps between exchanges.
	Instantaneous	Exports the field values as they appear at the time of exchange.

STAR-CCM+ Properties and Equivalent Abaqus Keywords

The following table summarizes the STAR-CCM+ field functions and equivalent Abaqus keywords for defining the physical quantities that are transferred between the two codes.

STAR-CCM+		**Abaqus**
Property	**Value**	**Keyword**
Exported Field Functions	Pressure and Wall Shear Stress	`*CO-SIMULATION REGION, TYPE=SURFACE, IMPORT <Abaqus surface>, CF`
	Boundary Heat Flux	`*CO-SIMULATION REGION, TYPE=SURFACE, IMPORT <Abaqus surface>, CFL`
	Boundary Heat Flux and either Specified Y+ Heat Transfer Reference Temperature or Local Heat Transfer Reference Temperature	`*CO-SIMULATION REGION, TYPE=SURFACE, IMPORT <Abaqus surface>, CFILM`

Imported Field Functions	Nodal Displacement	`*CO-SIMULATION REGION, TYPE=SURFACE, EXPORT <Abaqus surface>, U`
	Velocity	`*CO-SIMULATION REGION, TYPE=SURFACE, EXPORT <Abaqus surface>, V`
	Temperature	`*CO-SIMULATION REGION, TYPE=SURFACE, EXPORT <Abaqus surface>, NT`
Coupled Boundaries	STAR-CCM+ boundary	`Specify the co-located Abaqus surface in the *CO-SIMULATION REGION keyword (as above).`

Running the Co-Simulation

This section describes the steps for running the co-simulation case.

- Running the co-simulation
- Terminating Abaqus from STAR-CCM+

It is assumed that the model is loaded and prepared for a co-simulation as described in previous sections.

Running the Co-Simulation

Run the co-simulation by clicking the (Run/Restart) button in the *Solution* toolbar, or by selecting Solution > Run. STAR-CCM+ launches Abaqus automatically. The co-simulation progress can be monitored in the *Output* window, and STAR-CCM+ reports output for each data exchange process.

Terminating Abaqus from STAR-CCM+

If Abaqus fails to terminate automatically at the end of the co-simulation, or if you discover that the co-simulation is being conducted using the wrong set of parameters, you can manually terminate the co-simulation. To manually terminate the co-simulation, right-click the Co-Simulations > Abaqus Co-Simulation 1 node and select Terminate Abaqus.

Post-Processing the Co-Simulation Results

The structural response and interface response for FSI simulations can be post-processed in Abaqus/CAE. The fluid behavior and interface solution can be viewed in STAR-CCM+. Combined fluid and structural solution post-processing can be done using third-party software.

Analyzing with STAR-CCM+

In STAR-CCM+ you can view imported data using the corresponding field function. For example, `Nodal Displacement`, `Velocity`, and `Temperature`, in scenes and reports for the coupled boundary. Generated images from the scenes can be exported for use in creating animations using the methods that are described in Creating Images for an Animation. For general information on post-processing in STAR-CCM+, refer to the Analyzing section of the User Guide.

Output Analysis in Abaqus

The structural response and the co-simulation traction vectors (CF), as well as the co-simulation heat flux vectors (CFL) for thermal coupling, can be visualized in Abaqus/CAE. The output file for the co-simulation is located in the directory containing the Abaqus input file. The output file is named `<current job name>.odb`, where `<current job name>` corresponds to the Current Job Name property specified in the co-simulation session properties. Details on how to choose the data that is included in the `*.odb` file can be found in "Output to the output database", Section 4.1.3 of the Abaqus Analysis User's Guide.

Imported co-simulation traction and heat flux vectors can be displayed by requesting CF and CFL as nodal field output. The following example illustrates an output request for nodal displacements, nodal traction vector, nodal temperature (NT), and heat flux vectors imported from STAR-CCM+, and computed material point stresses.

Input File Usage:

```
*OUTPUT, FIELD
*NODE OUPUT
 U, CF, NT, CFL
*ELEMENT OUTPUT
 S
```

Allocating Cores to Abaqus in Batch Queuing Systems

When running both STAR-CCM+ and Abaqus as parallel jobs on batch queuing systems, STAR-CCM+ allows you to set the number of cores to allocate to Abaqus in the STAR-CCM+ launch command. It does this through a second parameter on the `-np` option, that is:

```
-np [starccm+_cores],[abaqus_cores]
```

To run an Abaqus Co-Simulation on a batch queuing system:

1. Edit the launch command in the run script and allocate the desired number of cores for each process. In the following example, STAR-CCM+ is allocated 24 cores, and Abaqus is allocated 8 cores:

   ```
   [INSTALL_DIR]/star/bin/starccm+ -np 24,8 -batchsystem
   <system> -batch <macro_file>.java
   ```

   The Abaqus value (`8`) that is defined in the above argument automatically overwrites the value of Number of CPUs in the Execution Properties.

2. Submit the run script to the batch management system.

   Once the run script starts, STAR-CCM+ generates a machine file from the hosts that the queuing system allocates. STAR-CCM+ uses the number of cores defined by the first value in the `-np` argument. When STAR-CCM+ subsequently launches Abaqus, it allocates to Abaqus the number of cores defined by the second value in the `-np` argument using hosts not already allocated to run STAR-CCM+.

   The total number of cores requested from the queuing system determines whether STAR-CCM+ and Abaqus are run on separate cores or if they run on the same cores.

   - To run each code on separate cores, request the sum of the cores designated on the -`np` argument from the queuing system.
   - You can run each code on the same cores for cases where the codes are not running concurrently, (that is, the Coupling Algorithm > Time-Marching Sequence is set to STAR-CCM+ Leads or Abaqus Leads). The memory footprint of each code must fit within the local memory of each host. If these conditions apply, then ask the queuing system for the largest of the two numbers appended to the `-np` argument. The two codes then use the same cores but run in sequence (so that STAR-CCM+ sleeps while Abaqus runs, and vice versa).

Additional Considerations

This chapter includes more technical information for performing co-simulations using Abaqus and STAR-CCM+.

- Exporting pressure loads
- Model dimensions and coordinate systems
- Co-location of the Abaqus and STAR-CCM+ models
- Supported Abaqus elements
- Supported STAR-CCM+ cells
- Known limitations

Exporting Pressure Loads

For co-simulation, export the pressure load on a wall boundary using one of two field functions: `Static Pressure` or `Absolute Pressure`. The `Absolute Pressure` in STAR-CCM+ is the absolute static pressure, which is the sum of the static pressure and the reference pressure. The `Static Pressure` is equal to the absolute pressure minus the reference pressure, and is often termed the "gauge" pressure in literature. The following diagram illustrates the relationship between these values.



Note: Use consistent applications of pressure in both Abaqus and STAR-CCM+. In Abaqus, the reference pressure is the pressure at which the structure is considered to be unloaded, with no displacements. It is the variation of the absolute static pressure with respect to the reference pressure that causes deformations and strains in the structure.

As an example of how to ensure consistency, consider a cylinder that is loaded on the inside diameter (ID) and the outside diameter (OD). STAR-CCM+ models the fluid on the inside of the cylinder. Assume that the OD of the cylinder is subject to 1 bar (atmospheric conditions), and that the STAR-CCM+ reference pressure is 1 bar (which is the default value). If STAR-CCM+ is set to export the `Static Pressure` as the load on the ID, then, in Abaqus, load the OD of the cylinder with zero load to be consistent. In effect, this means that no load is required in Abaqus at the OD. If, however, STAR-CCM+ is set to export the `Absolute Pressure` as the load on the ID of the cylinder, then, in Abaqus, load the OD with 1 bar (that is, the same value as the reference pressure in STAR-CCM+).

If the absolute static pressure within the cylinder is 10 bar when the fluid velocity is zero, it is convenient for the STAR-CCM+ reference pressure to be set to 10 bar. With the 10-bar reference pressure, if you export `Static Pressure` as the load on the ID, in Abaqus the OD must be loaded with a pressure of -9 bar (that is, a tension load) to be consistent. If, however, you choose to export `Absolute Pressure` from STAR-CCM+, load the OD in the Abaqus with a pressure of 1 bar.

This procedure of loading the OD of the cylinder with the pressure that will give the correct net loading on the structure is correct only for elements where through-thickness behavior is not considered (for example, shell elements). If solid elements are used to model the cylinder, then applying fictitious loads to the OD will produce the correct bulk loading, but incorrect behavior in the walls of the cylinder. In this case, absolute pressure should be exported to Abaqus for the ID and the ambient pressure should be applied on the OD within the Abaqus model.

Model Dimensions and Coordinate Systems

Co-simulation between two-dimensional and axisymmetric models is not supported. Represent two-dimensional models as a three-dimensional slice of unit thickness, or wedge element, with the appropriate symmetry boundary conditions applied.

Vector fields are transferred according to Abaqus and STAR-CCM+ conventions. The first component represents the quantity along the x-axis, the second component represents the quantity along the y-axis, and the third component represents the quantity along the z-axis. All fields are

transferred in the global coordinate system (the Laboratory Coordinate System in STAR-CCM+). If a local coordinate system is applied in the Abaqus model, the incoming and outgoing fields are transformed appropriately.

Co-Location of the Abaqus and STAR-CCM+ Models

Co-locate the Abaqus and STAR-CCM+ models. Place the fluid-structure interface regions of both models in the same geometric positions in space. Ensure that the coupled wall boundaries in STAR-CCM+ are not larger or smaller in extent than the corresponding interface in the Abaqus model.

Supported Abaqus Elements

Abaqus supports first-order and second-order continuum, shell, and membrane elements. Modified tetrahedral elements are not supported due to mapping limitations.

Supported STAR-CCM+ Cells

STAR-CCM+ supports tetrahedral and trimmed mesh cell types, as well as cells of arbitrary polyhedral type in 3D for co-simulation with Abaqus.

Known Limitations

This section lists the known limitations of co-simulation using Abaqus and STAR-CCM+.

- Run Abaqus and STAR-CCM+ on binary compatible platforms.
- The coupled boundaries must be of the same spatial extents - the coupled boundaries cannot be larger or smaller in extent than their counterpart in either code.
- Under *CO-SIMULATION REGION in the Abaqus .inp file, specify a single unique surface name per co-simulation zone defined in the STAR-CCM+ model; the same surface name can be specified on multiple lines when defining the imported/exported variables.
- If coupling both sides of a double-sided structure, a separate surface must be created for each side of the structure and assigned to a separate coupling zone in the fluid domain.

The STAR-CCM+ Component for Isight

Your STAR-CCM+ installation contains a Java component that can be used within the Isight process automation suite from SIMULIA. The component is located:

Linux:

```
[INSTALL_DIR]/isight/component/6.02/starccm.jar
```

Windows:

```
[INSTALL_DIR]\isight\component\6.02\starccm.jar
```

Add this component to your Isight Design Gateway using the Isight Library. It becomes available for selection in the Isight interface when it is successfully published.



Documentation for the component can be accessed by clicking Help on the *Component Editor* dialog.



Best Practices

This part of the documentation provides best-practice guidelines on applying STAR-CCM+ models to your applications.

Guidelines for Fluid Flow and Energy

Some guidelines apply to a wide range of simulations that involve fluid flow and energy.

This section contains some guidelines that apply in the following categories:

Mesh

The construction of the volume mesh has a direct influence on how well STAR-CCM+ simulates fluid flow and energy.

The construction of the volume mesh influences the rate of convergence and the accuracy of the final solution. In general, construct a volume mesh according to the following guidelines:

- Design a mesh that gives adequate resolution in regions where spatial gradients are high. Such regions occur when:
  - The mean flow changes rapidly. Examples are: irrigation canals, vortex shedding, forced mixing, and atmospheric flows.
  - There are strong shear layers. Examples are: atmospheric flows, fluid jets, flows past solids, and flows with strong vorticity.
- Whenever possible, align the mesh with the flow to improve accuracy and increase the rate of convergence.
- In order to make sure that simulation results are independent of the mesh, perform a grid sensitivity study with two or more meshes.

Note: Turbulence has a dominant role in the transport of momentum and other scalars for most complex turbulent flows. For this reason, simulations that involve turbulent flow tend to be more dependent on the mesh resolution than simulations that involve laminar flow.

- For external flows, perform domain convergence studies to make sure that what is happening at the flow boundaries does not affect the solution. For example, if the simulation has an outflow boundary close to a bluff body, the flow cannot recover from the effect of the wake before it meets the boundary.

Near-Wall Prism Layers

- In order to minimize numerical diffusion in the near-wall region, create layers of prism cells on walls and the fluid side of solid-fluid interfaces.
- The thickness of prism layers next to the wall have a direct impact on fluid y+ values. Check the y+ values in the solution, and make sure that they come within the range of validity recommended for the turbulence model and wall treatment that you have selected.
- For high accuracy modeling of boundary layers, create 5 - 10 cell layers within the boundary layer when using wall functions for the

turbulence model. Create around 40 cells in the boundary layer when using a low-Reynolds number turbulence model.

Cell Aspect Ratio

- Avoid creating a mesh that has excessive stretching in the direction normal to a wall.
- Keep cell aspect ratios as small as possible (for example, below 1000) to avoid convergence problems in the solver. Some problems require lower cell aspect ratios, such as simulations that involve solid stress calculations.

Thermophysical Properties

This section contains guidelines for defining thermophysical properties.

The following guidelines apply only when you use the energy equation in the simulation:

- Include the effects of temperature on thermophysical properties such as viscosity, specific heat, and thermal conductivity, when you have the information.
- When using temperature-dependent thermophysical properties, make sure that relevant dimensionless numbers maintain the values that you expect them to have. Examples of dimensionless numbers are the Reynolds number and the Prandtl number. As the Prandtl number appears explicitly in temperature wall functions, wall heat transfer predictions are sensitive to the value of thermophysical properties.
- When you use constant thermophysical properties, that is, independent of temperature, make sure that you specify all physical properties at the same temperature.
- When the Prandtl number is small, (for example, less than 0.1), use a mesh suitable for low values of y+ at the wall.
- When you specify surface properties for radiation heat transfer, be aware that in some circumstances the properties of real surfaces change significantly during service. For example, it is not unusual for the emissivity of an aluminized shield to increase several times through oxidation, dust, or the presence of moisture. This change of emissivity is not automatically captured by STAR-CCM+: if you wish to simulate such effects you must alter the surface properties.

Boundary Conditions

This section contains guidelines for defining the boundary conditions.

- Choose the shape and location of boundaries so that flow passes through a boundary in one direction only, that is, entering or leaving the domain. Avoid any recirculation across a flow boundary.
- As far as possible, choose a static pressure boundary for the outlet boundary.
- When possible, extend the computational domain to have flow regions upstream and downstream of the main region of interest. Extending the computational domain in this way minimizes the effect that boundary conditions have on the solution within the main region of interest. You can specify a coarser mesh in extended regions, so that the computational requirements are not increased unduly.
- In some simulations, the solution does not converge with the boundary conditions you require. When this situation occurs, start with a different set of boundary conditions and change them as the simulation proceeds. For example, when solving high-Mach number flows over bluff bodies, gradually increase the free-stream Mach number instead of starting with the final value.
- Choose boundary types that you know are compatible with the type of flow you are simulating. For example, a velocity inlet boundary is not a well-posed type for a highly compressible flow.

Conjugate Heat Transfer (CHT)

For applications that involve simultaneous heat transfer in both a solid and a fluid, there are two options for obtaining thermal solutions in the solid and the fluid.

These options are:

- STAR-CCM+ can obtain solutions in both the solid and the fluid.
- STAR-CCM+ can obtain the flow and thermal solution in the fluid only, and pass thermal information to another application that solves for the thermal solution in the solid. An example of this arrangement is where STAR-CCM+ exchanges thermal data with Abaqus in an explicitly coupled simulation.

Modeling Conjugate Heat Transfer

This section describes the conditions that apply to applications that involve simultaneous heat transfer in both a solid and a fluid.

For applications that involve simultaneous heat transfer in both a solid and a fluid, such as modeling the fluid and solid in a heat exchanger, the following apply:

- Fluid and solid equations are implicitly coupled and are solved simultaneously.
- The following surface conditions (temperature and energy) are enforced at the fluid-solid interface:

$$T_{solid}\big|_{surface} = T_{fluid}\big|_{surface}$$

$$q_{solid}\big|_{surface} = q_{fluid}\big|_{surface}$$

$$-K\frac{\partial T_{solic}}{\partial n}\bigg|_{surface} = h_f\left(T_{surface} - T_{ref}\right) + q_{rad} + q_c + \cdots$$

  where:

  - n is the surface outward normal
  - K_s is the solid conductivity
  - h_f is the fluid heat transfer coefficient
  - T_{ref} is the characteristic fluid temperature (for example, the domain mean or near wall cell temperature)

- q_{rad} and q_{pc} correspond to net radiative heat flux and added heat flux due to phase change, respectively.

Using Co-simulation

You can use a single STAR-CCM+ simulation to obtain thermal solutions in both the solid and the fluid, where the fluid and the solid are implicitly coupled at the interface. Alternatively, you can use co-simulation to obtain each solution in a separate simulation, where the fluid and the solid are explicitly coupled at the interface and exchange data back and forth as the simulation progresses.

You can use co-simulation when the fluid and solid time scales are significantly different. For example, if the flow is strongly turbulent and well-mixed, and solid conduction is slow, you can assume that the flow adjusts to the changes in the wall temperature quickly. This adjustment is essentially instantaneous with respect to the solid time scale, so the fluid-only simulation can be run steady and the solid-only simulation can be run unsteady. This arrangement allows you to update the fluid solution only when the solid wall temperature has changed significantly.

Co-simulation can provide significant run-time savings, with recent results showing significant speed-ups.

If you use co-simulation, be aware of the following issues:

- In co-simulation, convergence problems can arise because the fluid and solid are explicitly coupled. You should ensure that:
  - The correct information is passed between the simulations.

    For example, the heat transfer coefficient and its corresponding reference temperature are passed from the fluid-only simulation to the solid-only simulation. The wall temperature is passed from the solid-only simulation to the fluid-only simulation.

    Although there are other choices of boundary conditions to pass between the simulations (for example, you could pass the wall heat flux from the fluid-only simulation), for numerical stability reasons their use is not recommended.

    The recommended heat transfer coefficient and corresponding reference temperature are:

    - Fluid flow with no radiative transport or phase change:

Use the Specified Y+ Heat Transfer Coefficient and Specified Y+ Heat Transfer Reference Temperature. Start with a y+ of 100, but if needed for better convergence behavior, you can try a different y+. Increasing y+ decreases the Specified Y+ Heat Transfer Coefficient while decreasing y+ increases the Specified Y+ Heat Transfer Coefficient.

- Fluid flow with radiative transport and/or phase change:

  Use the Local Heat Transfer Coefficient and Local Heat Transfer Reference Temperature. If necessary, you can linearly transform this pair to another pair that has better convergence behavior.

- The correct co-simulation parameters are used. For example, the number of iterations, time-step, and update frequency.

- Loss of accuracy. The amount of accuracy loss strongly depends on the co-simulation parameters that are used and the amount of speed-up desired.

Co-simulation Time Scales

In a co-simulation, the fluid and solid systems are solved independently: a separate system of equations is solved for each of the fluid and solid regions.

Some of the most common time scales for the fluid region are as follows:

- Diffusion time scale:

$$t_{f1} = \frac{L_f^2}{v}$$

- Inertial time scale:

$$t_{f2} = \frac{L_f}{U_f}$$

- Internal/external forcing time scale:

$$t_{f3} = \frac{1}{\omega_f}$$

The time scale for the solid region is:

- Diffusion time scale:

$$t_s = \frac{L_s^2}{a}$$

The time scale ratios are:

$$\left[\frac{t_s}{t_{f1}}\frac{t_s}{t_{f2}}\frac{t_s}{t_{f3}}\right]=\left[\left(\frac{L_s{}^2 v}{L_f{}^2 a_s}\right)\left(\frac{L_s{}^2 U_f}{L_f a_s}\right)\left(t_s\omega_f\right)\right]$$

where:

- L_f and L_s are characteristic length scales (such as domain or grid size) for the fluid and solid, respectively
- v is the fluid kinematic viscosity
- U_f is a characteristic velocity scale (for example, domain mean or cell velocity)
- ω_f is a characteristic forcing frequency
- $a_s=\frac{K_s}{\rho C_p}$ is the solid diffusivity.

Run Time Speed-up

This section gives the equations for the run time of an implicitly coupled CHT simulation and an explicitly coupled CHT co-simulation.

The run time of an implicitly coupled (unsteady) CHT simulation is:

$$t_{total}^{IC}=\frac{t_{final}}{\Delta t}\left(N_{it}\tau_{it}\right)$$

where:

- t_{final} is the simulation ending time
- Δt is the time step
- N_{it} is the number of inner iterations per time step
- τ_{it} is the CPU time per inner iteration.

The run time of an explicitly coupled CHT co-simulation is:

- Fluid (steady): short time scale
- Solid (unsteady): long time scale

$$t_{total}^{EC}=\left(\frac{t_{final}^s}{t_{ex}^s}\right)\left[\left(N_{it}^f\tau_{it}^f\right)+\left(\frac{t_{ex}^s}{\Delta t^s}\right)\left(N_{it}^s\tau_{it}^s\right)\right]$$

where:

- t_{final}^s is the solid simulation ending time
- t_{ex}^s is the exchange period between the solid and fluid simulations

- N_{it}^{f} and N_{it}^{s} are the fluid and solid number of inner iterations, respectively
- τ_{it}^{f} and τ_{it}^{s} are the fluid and solid run times per inner iteration, respectively
- Δt^{s} is the solid simulation time-step.

In general, the speed-up in run time of an explicitly coupled CHT co-simulation compared to an implicitly coupled CHT simulation is:

$$\frac{t_{total}^{IC}}{t_{total}^{EC}} = \frac{N_{it}\tau_{it}t_{ex}^{s}}{\frac{\Delta t}{\Delta t^{s}}\left[\Delta t^{s}N_{it}^{f}\tau_{it}^{f} + t_{ex}^{s}N_{it}^{s}\tau_{it}^{s}\right]}$$

Setting the Co-simulation Parameters

When you set up a CHT co-simulation, you specify various co-simulation parameters.

The co-simulation parameters include:

- Mapping Mode
- Iteration Count
- Update Frequency
- Time Step

For the best results, use a conformal mesh at the solid to fluid interface. A conformal mesh is one whose vertices and faces match exactly on either side of an interface.

Setting the Initial Parameter Values

After setting the initial settings for these parameters, run the co-simulation.

If necessary, you can perturb near these initial settings and look for combinations that give better results (that is, a set of parameters that give good accuracy and short run times).

Fluid-only Simulation (Steady)

- Iterations

Set the number of iterations, which in general can be a strong function of time, to whatever it takes for convergence (that is, for the residuals to stop changing).

Solid-only Simulation (Unsteady)

- Update Frequency

  Set the initial update frequency to:

$$f_s = \frac{1}{\beta t_s} = \frac{a_s}{\beta L^2} = \frac{K_s}{\rho_s C_{p,s} L^s \beta}$$

  where:

  - L is a characteristic length scale for the solid (for example, overall length, half length, grid size)
  - β is a scaling factor less than one (for example, 0.2 is a good first choice).

- Time Step

  Set the solid-only simulation time step Δt^s to be some fraction of the update period. For example, $\beta t_s / 10$ can be appropriate if the solution is smooth and not too steep.

  Alternatively, you can slave the update period $\beta t_s$to a particular amount of change in wall temperature (for example, the maximum change in one cycle). This allows the update period to adjust itself and the time-step as the simulation progresses.

- Iterations

  Set the number of inner iterations to whatever it takes for convergence.

Cyclone Separators

This section gives some guidelines for modeling cyclone separators.

When you are modeling cyclone separators, consider the following:

Meshing

In order to minimize unwanted numerical diffusion effects, the best type of mesh is a hexahedral type with the mesh elements aligned with the circumference of the cyclone.

You can use the Directed Meshing feature to construct this mesh within STAR-CCM+.

Turbulence Model

An anisotropic turbulence model is required to model accurately the free to forced vortex transition that occurs in cyclonic flows. Standard K-Epsilon models and other models that are based on assumptions of isotropic turbulence are not suitable. These models tend to over-predict the turbulent viscosity and exaggerate the forced vortex.

Boundary Layer Resolution

Wall resolution is not critical as the turbulence is generated in the main flow.

Freezing the Solvers

If used correctly, freezing the solvers can be of great use in promoting initial simulation stability, computational efficiency, and total run time.

There are two primary situations in which freezing solvers can help:

- Early in the simulation, convergence issues can arise because of the lack of a good physical initial condition for transient simulations or initial iteration for steady simulations. At this point, you can activate solver verbosity so that you see exactly what each solver is doing and where problems occur. Freezing the solvers associated with more complex physics can allow you to obtain approximate converged solutions. Thereafter you can unfreeze the solvers one by one, continuing the solution each time, to see if any convergence issues remain.
- Later in the simulation, because of differing physical time scales or for purely numerical reasons, some solvers reach their asymptotic states much faster than others and are no longer significantly changing the solution.

  For instance, in a typical conjugate heat transfer problem, the flow time scale can be small and the solid conduction time scale can be large. Consider, for example, the flow of hot water in a thick-walled pipe. The temperature of the inner pipe wall quickly reaches the temperature of the hot water; the flow solution stabilizes to a steady flow condition. If the only remaining solution concerns the thermal distribution within the pipe wall, then the flow solvers can be frozen while the energy solver runs.

  For quantities that have reached their asymptotic states in steady simulations, you can freeze the solvers for the remainder of the simulation. Solvers can be frozen for the remainder of steady simulations. For unsteady simulations where transient effects occur, do not keep solvers frozen for a time longer than the time-scale of the transient effects. (That is, freeze the solvers for a short time and then unfreeze them to see if their asymptotic states have changed. Repeat the procedure as needed). This workflow is similar in nature to co-simulation but is not as flexible.

Guidelines for Aeroacoustics Calculations

This section provides best-practice guidelines for conducting aeroacoustics analyses in the following applications.

Note: Use the double precision version of STAR-CCM+ for all aeroacoustic and hydroacoustic analyses.

Automobile Aerodynamics

At higher speeds, the aeroacoustic noise dominates noise from the engine, exhaust, and tires.

The aeroacoustic noise comes from air flows over pillars, wing-mirrors, and the underbody, and by flow through ducts (including fan noise). This section presents the following recommended practices to model automobile aerodynamics:

- Recommended Process
- Recommended Mesh Settings
- Recommended Time-Step
- Recommended Solver Settings

Recommended Process

1. Apply best practices for mesh-resolution for a RANS-based steady state aerodynamics case. It is vital to get the mean flow correct. Compromise the wake refinement only if rear-vehicle aeroacoustics are not important.
2. Plot the mesh cut-off frequency to judge if the refinement in the acoustic source regions resolves frequencies in the region of 3–5 kHz. In practice, a mesh cell size of 2 mm or less is required. Refine appropriately and map the solution to the refined mesh.

   You can identify the acoustic source regions from the regions where the acoustic correlations are large. Use Proudman for volume sources:



   and Curle for surface sources:



3. Ensure that the mesh in the near field is adequate for local sound propagation. For example, use 3.5 mm or less in the region of the A-pillar, side-view mirror, and side glass to propagate a 5 kHz signal.
4. Choose a time-step size suitable for the most demanding (that is, smallest) of the following constraints: frequency resolution, explicit CFL, or implicit CFL (as appropriate) and Courant number.
5. Restart a transient case from the initial steady state simulation that you used to judge the previous requirements. Switch the turbulence model from RANS to DES or LES. Use second-order time discretization, and a spatial discretization that is based on central differencing.
6. Select and fine-tune the numerical parameters to ensure convergence within each time-step. You can conclude that the solution has converged within a time-step when all (not a subset) of the following criteria are satisfied:
   a. Monotonically reducing or minimized residuals at the end of the time-step.
   b. An unchanging monitor. The monitor can be: pressure at a point, lift, or drag.

Recommended Mesh Settings

Select cell size, Δ, based on criteria:

Criterion	Comments / Description	Formula	Typical Values
Mean Flow	Use best practices for external steady aerodynamics. Refine in the		Δ_{min} = 2 mm y+ = 1 for low-Re wall treatment

Criterion	Comments / Description	Formula	Typical Values
	near-wall region to capture separation from the body.		
Turbulence integral scales, $L_{integral}$	A measure of the "typical" size of dissipative eddies. Approximated from steady state, $\Delta \leq 0.5\ L_{integral}$	$L_{integral} = \frac{C_\mu^{0.75} \cdot k^{1.5}}{\epsilon}$	$L_{integral} > 10$ mm $\Delta \leq 2$ mm
Frequencies that are associated with turbulent structures	Mesh cut-off frequency, that is based on local turbulent fluctuations compared to best-practice mesh size. Approximated from steady state.	$F_{MCF} = \frac{\sqrt{\frac{2}{3}k}}{2\Delta}$	$\Delta = 4$ mm captures frequencies in the range 1000–1500 Hz.
Propagation of pressure waves	Based on second order space and time discretization, 20 cells per acoustic wavelength, λ (less for higher-order spatial schemes). a, is the speed of sound.	$\lambda = \frac{a}{F}$ $\Delta \leq \frac{\lambda}{20}$	$\Delta \leq 3.5$ mm. Based on a requirement to propagate a pressure wave of $F = 5000$ Hz in air.

Recommended Time-Step

Select the time-step size Δt, based on criteria:

Criterion	Comments / Description	Formula	Typical Values
A signal of frequency F	The Nyquist sampling theorem states that a signal with frequency, F, can be resolved by sampling at a rate greater than $2F$. For second-order spatial schemes, tests show that around 15 points are required to resolve the wave amplitude.	$\Delta t_F \leq \frac{1}{15F}$, recommended	Δt_F = 2.22E-5 s. Based on requirement to capture F = 5000 Hz.
Convection and acoustics	Based on the convection speed, U, and speed of sound (acoustic pressure waves propagate at the local speed of sound, a). The Courant (CFL) number is based on $(U+a)$	Explicit: $\Delta t_{exp} = \frac{\Delta}{(U+a)}$ Implicit: $\Delta t_{imp} = CFL \cdot \Delta t_{exp}$	Δt_{exp} = 5.0E-6 s. Based on Δ = 2.0 mm, U = 50 m/s and a = 350 m/s. Δt_{imp} = 5.0E-5 s. Based on CFL = 10.
Convection only	Based on convection speed, U. Convective Courant Number is based on U only.	$\Delta t_{Courant} = \frac{\Delta}{U}$	$\Delta t_{Courant}$ = 4.0E-5 s. Based on Δ = 2.0 mm and U = 50 m/s.
Diffusive time-scales	Most important in the boundary layer where the local (eddy)	$\Delta t_{diff} = \frac{\Delta_{normal}^2}{V_{eff}}$	Δt_{diff} = 1.0E-6 s. Based on Δ_{normal}= 5.0E-6 m,

Criterion	Comments / Description	Formula	Typical Values
	viscosity, V_{eff}, is largest. Unimportant to acoustic phenomena		V_{eff} = 2.0E-5 m^2/s.

Recommended Solver Settings

You can select solver settings for a slower ("regular") mesh and time-step setup, a faster ("aggressive") setup, or the default settings.

	Velocity URF	Pressure URF	# Iterations Per Time step
Default	0.8	0.2	5
Regular Δt = 1.0E-4 to 1.0E-5 s. Δ_{min} = 1–4 mm.	0.7	0.7	4–10
Aggressive Δt = 2.0E-5 s. Δ_{min} = 2 mm.	1.0	0.9	3

Sunroof Buffeting

At low speeds, flow resonance can occur when the shear layer coming off the front of a vehicle passes over its open sunroof and generates periodic pressure fluctuations within the cabin.

This phenomenon is known as sunroof buffeting, and can be modeled:

- Recommended Process
- Recommended Mesh Settings
- Recommended Time-Step
- Recommended Solver Settings

Recommended Process

The recommended process is:

1. Assess the resonance of the modeled cabin volume by performing a FRET (**F**requency **RE**sponse **T**est):
   a. Apply a pressure disturbance to the sunroof aperture with no flow using an initial condition and field function.
   b. Run the compressible solver with the same settings as a flow case to set up a compressible standing wave in the cavity.
   c. Perform an FFT of the internal pressure-time signal, to identify the cavity resonance frequencies.
2. The resonance of the real vehicle cavity (from experiment) can be different from the FRET for various reasons, namely: panel impedance, leakages, or surface attenuation. Use an artificial compressibility factor, ϕ, based on the ratio of experimental CFD frequencies to modify the response of the modeled cavity as follows:

$$a' = \phi a_0 \tag{2687}$$

$$T'_{in} = \phi^2 T_{in} \tag{2688}$$

$$M' = \frac{M_0}{\phi} \tag{2689}$$

$$p'_{ref} = \phi^2 p_{ref} \tag{2690}$$

w$a_0$here:

- is the speed of sound in the fluid.
- T_{in} is the inlet temperature.
- M_0 is the mass of the air.

- p_{ref} is the reference pressure.

Recommended Mesh Settings

Select cell size, $\Delta_{shear-layer}$, based on criteria:

Criterion	Comments / Description	Formula	Typical Values
Shear layer	Full length of aperture and height of shear-layer oscillations.		$\Delta_{shear-layer}$ = 4 mm

Recommended Time-Step

Select the time-step size Δt, based on criteria:

Criterion	Comments / Description	Formula	Typical Values
Buffeting frequency F_B	Shear-layer oscillates over the opening. Buffeting frequency F_B, is directly proportional to the free-stream speed, U_∞, and is in the range 20-30 Hz.	$F_B = \frac{U_\infty}{3L}\left(N - \frac{1}{4}\right)$ where: N = oscillation mode L = length of sunroof aperture	Δt_F= 1.0E-4 s.

Recommended Solver Settings

The recommended solver settings are:

	Velocity URF	Pressure URF	# Iterations Per Time-step
Default	0.8	0.2	5
Regular Δt = 1.0E-4 s. Δ_{min} = 1–4 mm.	0.7	0.5	5–10
Aggressive Δt = 1.0E-4 s. Δ_{min} = 1–4 mm.	0.9	0.7	5

Radial Compressors

Compressors are usually open to the air side of engines. The generated noise is audible outside the engine. Noise from compressors can be associated with blade passing, blade stall, and other dynamic flow mechanisms.

This section presents recommended practices to model acoustic sources in radial compressors:

- Recommended Process
- Recommended Time-Step
- Recommended Solver Settings

Recommended Process

Qualify Using Steady State

Before running an unsteady analysis, it is vital to first qualify the model by ensuring that the fan performance curve is predicted well from steady state simulations. Use the Realizable K-Epsilon or the K-Omega SST turbulence models for consistency with the DES model in the unsteady analysis.

To avoid pressure reflections, set a free-stream boundary at either the inlet or the outlet, depending on which side is most important. Having a free-stream boundary on both inlet and outlet is unstable in a steady analysis. If it is necessary to control the mass flow, then set the mass flow on the downstream boundary, since acoustics on the inlet side are often more important.

Use Compressible Flow

It is essential to use compressible flow modeling because of the operating conditions of these machines. The peak Mach number is generally in the supersonic range, and a 'choke' regime is present due to a passage shock. Choke occurs when a shock wave spans the full space between the blade pressure and suction sides.

Recommendations for Obtaining Performance Curves in Steady State

1. Initialize the flow field using GSI (Grid-Sequencing Initialization).
2. For high speed radial compressors, validations show that the mixing plane interface is superior to the frozen rotor.
3. The segregated solver is preferred to the coupled solver due to the follow-on step to transient where it is definitely faster. Prediction accuracy between the two is marginal.

4. Marginal differences exist between high and low Reynolds number resolution on the blade. A high Reynolds number wall treatment is preferred as the mesh can be coarser and the grid sequencing initialization is more effective. Practice shows that low Reynolds number meshes produce more accurate predictions. You can refine the mesh and interpolate between high and low Reynolds number predictions.
5. The preferred boundary condition combination is to use stagnation inlet and static pressure outlet. Start with a low pressure ratio between the inlet and outlet. To get points at higher pressure ratios, increase the back pressure. As the pressure ratio reaches the maximum for the device, activate the target mass flow setting in the pressure boundary properties. This setting allows you to fine-control the mass flow while STAR-CCM+ still applies an outlet pressure through successive corrections. (Alternative boundary combinations are also possible - you can choose either an inlet mass flow boundary or an outlet mass flow boundary in steady state. An outlet mass flow boundary is preferred when moving to an unsteady analysis.)
6. The SST K-Omega Turbulence model seems to be marginally more accurate than Spalart Allmaras turbulence model as conditions approach choked flow. Otherwise both perform with a similar level of accuracy.

A performance curve of 10 points can be generated in approximately a day when running a 4–6 million-cell mesh on 48 processors.

Note: Some over-prediction of mass flow at the choked condition is observed with the recommended settings. The over-prediction could be due to the inflow profile, or geometric simplifications like neglecting the under-wheel leakage.

Eliminating Reflective Boundary Effects

1. Extract reports of the mass-flow averaged Mach number, pressure, and temperature from the steady-state run at inlet and outlet boundaries.
2. Change the steady state case inlet boundary type to *Freestream* and set the Mach number, pressure, and temperature from the steady state reports.
3. Run the simulation to convergence with these new boundary settings.
4. Ensure that the pressure ratio and mass flow remain unchanged. Alternatively, wait until the mass flow balances across the inlet and outlet.

Recommendations for the Unsteady Simulation

1. Start the unsteady simulation from a converged steady solution, as described in the previous section.
2. Convert to transient.
3. Change the turbulence model to Detached Eddy Simulation (DES) using the SST k-omega model.
4. Apply the recommended Segregated Flow settings.
5. Use the Segregated Flow Model setting Hybrid-Bounded Central Differencing (BCD) or full BCD, if cell quality allows with blend=0.3. Use second-order time.
6. Freeze the Wall Distance Solver.
7. Run the case.
8. Run for between 100 and 200 full rotations to get a good statistical sample for digital signal processing.

Using these settings approximately 15 full rotations a day are possible when running a 4–6 million-cell mesh on 48 processors.

Choosing Transient or Harmonic Balance

Radial compressors are usually close to a downstream scroll volute, diffuser vanes, or other geometries which are not rotating. The correct way to simulate this system is to account for the periodic interaction of the rotating and stationary parts, namely downstream wake, and upstream pressure gusts of turning vanes, blades, and supports.

There are two ways to do model radial compressors in STAR-CCM+: using transient Rigid Body Motion (sliding mesh) or the Harmonic Balance method. Both are more time-consuming, but more accurate, than steady state approximations that use frozen-rotor or mixing-plane stator-rotor interfaces. The harmonic balance method runs faster than transient. However, the harmonic balance method only resolves tones that are based on the modes solved (blade passing frequency, harmonics, and Strouhal shedding). It does not resolve broadband noise which is associated with resolved turbulence.

Recommended Time-Step

Select the time-step size Δt, based on criteria:

Criterion	Comments / Description	Formula	Typical Values
Rotation speed ω	Aim for 1° rotation per	Minimum of	Δt = 8.3E-7 s.

Criterion	Comments / Description	Formula	Typical Values
	time-step, or less if the number of blades exceeds 24. Then select Δt to give at least 15 time-steps per blade passage.	$\Delta t = \frac{1}{6\omega}$ (1° rotation per time-step) and $\Delta t = \frac{4}{N_b \omega}$ (15 time-steps per blade passage) where: N_b = number of blades.	Based on ω = 200,000 RPM.

Recommended Solver Settings

The recommended solver settings are:

	Velocity URF	Pressure URF	# Iterations Per Time-step
Default	0.8	0.2	5
Regular	0.8	0.7	3–4
Aggressive	1.0	0.9	3

Axial Fans

Axial fans, commonly used for cooling and ventilation, are placed in locations where they are highly audible.

This section presents recommended practices to model acoustic sources in axial fans:

- Recommended Practices
- Recommended Time-Step
- Recommended Solver Settings

Recommended Practices

Qualify Using Steady State

Before running a transient analysis, it is vital to first qualify the model by ensuring that the fan performance curve is predicted well using steady state simulations. Use the Realizable K-Epsilon or the K-Omega SST turbulence models for consistency with the DES model in the unsteady analysis.

To avoid pressure reflections, set the inlet or outlet boundaries to be freestream depending on which is the important side. If it is necessary to control the mass flow, then set the mass flow at a boundary farther away from the fan. Aim to place boundaries at least ten fan diameters away from the fan.

Use Compressible Flow

Rotors passing relative to stators create flow and pressure pulsations which are strictly a compressible phenomenon. The pulsations can result in accumulated mass (monopoles) within the system, which an incompressible calculation does not permit. Even at low Mach-numbers, compressibility is preferred.

Recommendations for Obtaining Steady Performance Curves

1. No special initialization procedures are required. To accelerate convergence, set the initial axial velocity to give the expected mass flow through the fan.
2. No preferences exist between mixing plane and frozen rotor.
3. The segregated solver is preferred to coupled due to the existence of a follow-on step to transient where it is definitely faster. Prediction accuracy between the two is marginal. Use the same solver for both steady state and transient calculations to minimize flow-wide disturbances.

4. Low Reynolds number resolution on the blade is preferred for fan performance prediction.
5. Start by running with the known boundary conditions (pressure-pressure or mass flow-pressure) obtained from the experimental conditions.
6. Sample the Mach number at the boundary and reset the steady boundary conditions to freestream with the sampled pressure, Mach number, and temperature balance.
7. Run with freestream boundaries, and the system reregulates itself to find the correct mass-flow / pressure balance.

Using the recommended settings, a performance curve of 10 points can be generated in approximately a day when running a 4–6 million-cell mesh on 48 processors.

Recommendations for the Unsteady Simulation

1. Start the acoustics simulation from the steady solution with non-reflective boundaries as described.
2. Change the turbulence model to Detached Eddy Simulation (DES) using the SST k-omega model.
3. Use the recommended Segregated Flow settings.
4. Use the Segregated Flow Model setting Hybrid-Bounded Central Differencing (BCD) or full BCD, if cell quality allows with blend = 0.3. This higher blending factor can be necessary due to non-orthogonality across the sliding mesh. Use second-order time.
5. Freeze the wall-distance calculation. The wall distances are needed mainly near to the surfaces to calculate wall damping characteristics. Since the near-wall distances are unlikely to change, it is inefficient to calculate this distance at every time-step.
6. Convert to transient and run.
7. Run for between 100–200 full rotations to get a good statistical sample for digital signal processing.

Using the settings approximately 15 full rotations a day are possible when running a 4–6 million-cell mesh on 48 processors.

Choosing Transient or Harmonic Balance

Axial fans, that are used for under-hood cooling or building HVAC, are placed close to casing, mounting struts or other geometries which are not rotating. The correct way to simulate this system is to account for the periodic interaction of the rotating and stationary parts, namely downstream wake, and upstream pressure gusts of turning vanes, blades, and supports.

There are two ways to model axial fans in STAR-CCM+: using transient Rigid Body Motion (sliding mesh) or the Harmonic Balance method. Both are more time-consuming, but more accurate, than steady state approximations that use frozen-rotor or mixing-plane stator-rotor interfaces. The Harmonic Balance method runs faster than transient. However, the Harmonic Balance method resolves tones that are based on the modes solved (blade passing frequency, harmonics, and Strouhal shedding). It does not resolve broadband noise which is associated with resolved turbulence.

Recommended Time-Step

Select the time-step size Δt, based on criteria:

Criterion	**Comments / Description**	**Formula**	**Typical Values**
Rotation speed ω	Aim for 1° rotation per time-step, or less if the number of blades exceeds 24. Then select Δt to give at least 15 time-steps per blade passage.	Minimum of $\Delta t = \frac{1}{6\omega}$ (1° rotation per time-step) and $\Delta t = \frac{4}{N_b\omega}$ (15 time-steps per blade passage) where: N_b = number of blades.	Δt = 5.0E-5 s. Based on ω = 3,333 RPM.

Recommended Solver Settings

The recommended solver settings are:

	Velocity URF	**Pressure URF**	**# Iterations Per Time-step**
Default	0.8	0.2	5
Regular	0.7	0.5	5–10
Aggressive	1.0	0.8	5

Choosing DES, LES, or URANS for Aeroacoustics

Where boundary layer structures do not directly affect the aeroacoustics, DES and LES using Bounded Central Differencing perform similarly in terms of pressure spectra.

LES is quicker as it solves fewer equations. When the hybrid BCD discretization scheme is used in DES, some key regions, such as an A-pillar or separated shear layers, suffer from numerical damping of the pressure fluctuations. The best practice (at this time) is to resolve the boundary-layer down to y+ equal to 1.

In general, Unsteady Reynolds Averaged Navier Stokes (URANS) is not recommended, unless the flow features are guaranteed to be tonal (such as obtained from blade passing or coherent shedding). Never use URANS when there is broadband content arising from turbulence. DES uses RANS only in the boundary layer. RANS-type meshes are sufficient in the bulk flow (away from boundary layers). You can investigate the mesh suitability by interrogating the spatial distribution of the turbulence integral scales arising from a steady RANS calculation. Use a mesh size less than 0.5x(local integral turbulence length scale).

Mesh Issues

Highly Skewed Cells at Sliding or Rotating Interfaces or Small Gaps

In circumstances where the cell orthogonality or connectivity is compromised, the solver can experience instability with Bounded Central Differencing (for LES) or Hybrid-Bounded Central Differencing (for DES). Under these circumstances, second-order upwind spatial discretization is recommended; it is more robust than Bounded Central Differencing. However, expect the flow structures to suffer smearing from the upwind contribution of the second-order scheme. The main effect is the reduction of the spectral magnitudes at higher frequencies, and, possibly, the redistribution of energy to lower frequencies.

Refinement Interfaces on a Trimmed Mesh

There is some evidence that (acoustic) pressure reflections or spurious (acoustic) convective disturbances can occur at mesh coarsening interfaces for trimmed meshes. As a result, the local flow fluctuating structures swamp the acoustic pressure effects. The likely cause is local orthogonality errors that reduce the spatial accuracy toward first-order accuracy. Using 20 cells per acoustic wavelength is recommended. But, where the local discretization due to refinement is reduced to first order, use 40 cells per acoustic wavelength.

Another solution to the issue is to use an equivalent polyhedral mesh with gradual coarsening/refinement. The polyhedral mesh eliminates reflections, but a speed penalty factor of around 1.5 exists for an equivalent mesh size—directly proportional to the increased number of faces per cell.

Freestream Boundary Condition Issues

The freestream boundary, usable with an ideal-gas formulation only, is non-reflective to one-dimensional disturbances incident (that is, normal) to the boundary.

It is based on the one-dimensional method of characteristics, calculating the boundary face fluxes for all the (inviscid) transported properties from the classical "compression", "expansion" and "convective" waves. These waves have eigenvalues $U+a$, $U-a$ and U. However, the following limitations exist when using a freestream boundary:

- This boundary treatment is found to be ill-posed when flow structures cross the boundary. When the flow structure crosses a boundary, spurious compressible pressure waves can emanate from the interaction and propagate upstream to influence the source field.
- Where the boundary is not normal to the incident compressible waves, errors can occur in the non-normal directions.

Acoustics in Liquid Media

In STAR-CCM+, hydroacoustic analysis is an extension of aeroacoustic analysis.

The ability to analyze sounds in liquids is possible under the following conditions:

- Use a double precision version of STAR-CCM+ to improve the numerics in calculating the speed of sound in liquids.
- Select a compressible EOS (Equation of State) for the fluid medium.
- Select an appropriate boundary condition. There are no non-reflective boundary treatments for liquid compressible flows, so use treatments to damp out the pressure waves from the source region before they reach the inflow and outflow boundaries.

Post-Processing the Acoustic Pressures and Pressure Perturbations

Although there are several ways to visualize the acoustical pressures and pressure perturbations, two useful methods are:

Using Wavenumber Decomposition

Using a line spatial Fourier transform, you can visualize the separate acoustical and convective energy content. The screenshot below shows the upstream (u^+) and downstream (u^-) convecting waves and the upstream (a^+) and downstream (a^-) acoustically propagating waves for a generic side-view mirror case:



For more information, see What Are Line Spatial Fourier Transforms?.

Using More Field Functions

A useful method of visualizing the fluctuating flow field is to subtract the average flow from the instantaneous flow, that is:

$$p' = p - \bar{p} \tag{2691}$$

In aeroacoustics, use the technique below to display the fluctuating pressure p':

1. Run the transient calculation so that it is statistically steady. That is, the initial state has been removed from the domain and the fluctuations in the flow are unchanging.
2. Create the field mean value for pressure. To create this value, right-click **Monitors** and choose **New Monitor** > **Field Mean**. For the *Field Function*, select `Pressure`.
3. In order to obtain a statistical mean for each point in the field, run the simulation for several further cycles.
4. Create a field function, **Pprime**, defined as `${Pressure} - ${PressureMeanMonitor}`
5. Create a plot of **Pprime** to show the hydrodynamic and acoustical pressures.

An example plot is shown below:



Index	Description
1	Turbulent fluctuation in the near field
2	Acoustical fluctuation in the far field

Guidelines for Aerodynamics Calculations

The following sections provide best practices guidelines when using a steady state RANS approach for:

Incompressible External Aerodynamics: Steady State RANS Approach

The following sections describe best practices for estimation of external aerodynamic/hydrodynamic loads for incompressible flow using a steady-state RANS approach.

This approach is valid when the maximum Mach number in the domain is less than 0.2–0.3, and when unsteady wake effects are not important.

Note: If unsteady wake effects are important, use a DES, LES, or URANS approach.

This page contains the following sections:

- Recommended Domain Settings
- Recommended Physics Setup
- Recommended Mesh Settings

Recommended Domain Settings

- Use a hexahedral or "bullet-shaped" far-field boundary, with domain extents around 8–10 body lengths or wing spans, whichever is larger, from the body. The mesh size on this boundary can be large: typically of the order of the body length.
- Use a velocity inlet boundary condition type for the upstream and lateral boundaries
- Use a pressure outlet boundary condition for the downstream boundary.
- Using a half-body with a symmetry plane is appropriate for cases with no sideslip. Take care to match the symmetry plane mesh growth with the mesh growth of the volume mesh.
- Optional: Angles of attack/sideslip are typically defined by changing the far-field boundary condition flow direction, not the orientation of the body.

Recommended Physics Setup

To set up an incompressible external aerodynamics model using the steady state RANS approach, choose the following physics models in order:

Group Box	Model
Space	**Three Dimensional**
Time	**Steady**
Material	Select a material model that is appropriate for your simulation.
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**
Turbulence	**Reynolds-Averaged Navier-Stokes** (Selected automatically)
	Realizable K-Epsilon or **Spalart-Allmaras**

To select a specific turbulence model, choose from the following:

- The **Realizable K-Epsilon** model is preferred, especially if there is a blunt boundary at the end. Either the **low y+** or **all y+** boundary treatments can be used with this model, depending on the mesh.

  This model works best when complex geometric protrusions exist, bluff body separation occurs, or both.
- The **Spalart-Allmaras** turbulence model can be used in cases of streamlined geometries without large base separation regions.

  This model works best for attached boundary layers or mildly separated flows (that is, flow past a wing at or below stall).

Recommended Mesh Settings

Both polyhedral and trimmed meshes have proven successful on this class of problem.

In either case, typical rules for good meshing practices apply and are important. The mesh should give proper resolution of geometry and flow field features, along with good cell quality.

For instance, refine the mesh:

- Around leading and trailing edges of lifting surfaces, control surfaces, and other protrusions.
- At tips of lifting surfaces and control surfaces in order to capture vortices that are shed in these locations.
- At body/lifting-surface junctions.
- To resolve any bluff-body wakes.

- On airfoil leading edges. Use aligned surface meshes in these regions when possible.
- If feasible, use at least two cells to resolve thin blunt leading edges, or other thin step-like protrusions.

Note: External aerodynamics simulations can be sensitive to the quality of the surface geometry representation. Care must be taken when using the surface wrapper such that junctions/intersections result in smooth, well-defined edges. Whenever possible, the preferred approach is to start with clean CAD geometry and bypass surface wrapping.

Polyhedral Meshes

- Polyhedra have the advantage of smooth growth away from the body.
- Due to the smooth growth behavior, polyhedral are typically preferred when aerodynamic coefficients are of key importance.
- Due to the pseudo-random orientation of the faces, polyhedra also generally give better results than trimmed meshes when large sweeps in flow direction (such as angle of attack or sideslip) are used.
- Typical polyhedral meshes contain 2–20 million cells, depending on the complexity of the geometry and flow field, and the near-wall treatment.
- Use Volume Growth Rate, with the expansion rate set between 1.05 and 1.15.

Caution: Be careful when using volumetric controls with polyhedral meshes: abrupt changes in cell size can occur if there is a significant difference in the sizes that are specified on the surface and in the volumetric control.

Trimmed Meshes

- If you are modeling a single flow direction, or only small variation in flow, a trimmed mesh with the grid lines aligned with the flow often provides the most efficient approach to getting good answers.
- Typical trimmed meshes contain 4–40 million cells, depending on the complexity of the geometry and flow field, and the near-wall treatment.
- In order to capture the near field properly, set the volume growth rate to **Slow** or **Very Slow**.

Prism Layers

- Proper resolution of the boundary layer is critical. Choose the thickness of the prism layer such that the entire boundary layer is contained within it.
- The use of wall functions rather than integrating to the wall can be appropriate depending on several factors, including:
  - Desired accuracy
  - Relative importance of skin friction drag
  - Importance of transition
  - Existence/importance of separation/reattachment on smooth boundaries
  - For integrating to the wall, 20–30 prism layers are typically used, with the near-wall y+ being of the order 1.
  - For wall functions, 5–8 prism layers are typically used, with the near-wall y+ being of the order 50–150.
  - When using the Spalart-Allmaras Turbulence model, integrate to the wall (do not use wall functions).

Transonic/Supersonic External Aerodynamics: Steady State RANS Approach

The following sections describe best practices for estimating external aerodynamic loads for compressible flow, with freestream Mach numbers ranging from 0.5 to 2.5.

If unsteady phenomena such as vortex shedding are not important, and angles of attack are not extreme, steady state RANS solutions are acceptable for this type of analysis.

Recommended Domain Settings

- Use a spherical far-field boundary, with domain extents around 8–10 body lengths or wing spans, whichever is larger, from the body. The mesh size on this boundary can be large: typically of the order of the body length.
    - When replicating a wind-tunnel test, modeling the wind-tunnel walls is appropriate if blockage effects are important. In this case, slip walls are often defined at the tunnel walls, and no prism layers are constructed.
- Use a freestream boundary condition type at the far-field.
- Using a half-body with a symmetry plane is appropriate for cases with no sideslip. Take care to match the symmetry plane mesh growth with the volume mesh growth.
- Optional: Angles of attack/sideslip are typically defined by changing the far-field boundary condition flow direction, not the orientation of the body.

Recommended Physics Setup

To set up a transonic/supersonic external aerodynamics model using the steady state RANS approach, choose the following physics models in order:

Group Box	Model
Space	**Three Dimensional**
Time	**Steady**
Material	**Gas**
Flow	**Coupled Flow**

Group Box	Model
Equation of State	**Ideal Gas**
Optional Physics Model	**Coupled Energy** (Selected automatically)
Viscous Regime	**Turbulent**
Turbulence	**Reynolds-Averaged Navier-Stokes** (Selected automatically)
	SST K-Omega or **Spalart-Allmaras**

To select the specific turbulence model, choose from the following:

- The **SST K-Omega** model is preferred. This model is robust, is well-behaved near walls, and does well in recirculation regions.

  Either the low y+ or all y+ boundary treatments can be used with this model, depending on the mesh. This model works best when complex geometric protrusions exist, bluff body separation occurs, and at supersonic speeds.
- The **Spalart-Allmaras** turbulence model can be used in cases of streamlined geometries without large base separation regions. This model works best for attached boundary layers or mildly separated flows (that is, flow past a wing at or below stall), and at subsonic speeds.

Recommended Mesh Settings

Both polyhedral and trimmed meshes have proven successful on this class of problem. In either case, typical rules for good meshing practices apply and are important. The mesh should give proper resolution of geometry and flow field features, along with good cell quality.

For instance, refine the mesh:

- Around leading and trailing edges of lifting surfaces, control surfaces, and other protrusions.
- At tips of lifting surfaces and control surfaces in order to capture vortices that are shed in these locations.
- At body/lifting-surface junctions.
- Where shocks exist.
- To resolve any bluff-body wakes.
- On airfoil leading edges. Use aligned surface meshes in these regions when possible.

- If feasible, use at least two cells to resolve thin blunt leading edges, or other thin step-like protrusions.

Note: External aerodynamics simulations can be sensitive to the quality of the surface geometry representation. Care must be taken when using the surface wrapper such that junctions/intersections result in smooth, well-defined edges. Whenever possible, the preferred approach is to start with clean CAD geometry and bypass surface wrapping.

Polyhedral Meshes

- Polyhedra have the advantage of smooth growth away from the body.
- Due to the smooth growth behavior, polyhedral are typically preferred when aerodynamic coefficients are of key importance.
- Due to the pseudo-random orientation of the faces, polyhedra also generally give better results than trimmed meshes when large sweeps in flow direction (such as angle of attack or sideslip) are used.
- Typical polyhedral meshes contain 2–20 million cells, depending on the complexity of the geometry and flow field, and the near-wall treatment.
- Use Volume Growth Rate, with the expansion rate set between 1.05 and 1.15.

Note: Be careful when using volumetric controls with polyhedral meshes: abrupt changes in cell size can occur if there is a significant difference in the sizes that are specified on the surface and in the volumetric control.

Trimmed Meshes

- If you are modeling a single flow direction, or only small variation in flow, a trimmed mesh with the grid lines aligned with the flow often provides the most efficient approach to getting good answers.
- Typical trimmed meshes contain 4–40 million cells, depending on the complexity of the geometry and flow field, and the near-wall treatment.
- In order to capture the near field properly, set the volume growth rate to **Slow** or **Very Slow**.

Prism Layers

- Proper resolution of the boundary layer is critical. Choose the thickness of the prism layer such that the entire boundary layer is contained within it.

- The use of wall functions rather than integrating to the wall can be appropriate for the SST K-Omega turbulence model depending on several factors, including
  - Desired accuracy
  - Relative importance of skin friction drag
  - Importance of transition
  - Existence/importance of separation/reattachment on smooth boundaries
  - For integrating to the wall, 20–30 prism layers are typically used, with the near-wall y+ being of the order 1.
  - For wall functions, 5–8 prism layers are typically used, with the near-wall y+ being of the order 50–150.

When using the Spalart-Allmaras Turbulence model, integrate to the wall (do not use wall functions).

Solution Procedure

To improve convergence, especially as the Mach number increases, order the cells in the general direction of the flow using directional reordering.

Initialization

- Grid sequencing initialization with default values typically works well for this type of problem, and is recommended up to around Mach 2.
- The flow field can be initialized to uniform conditions matching the far-field boundary—recommended for Mach numbers greater than 2. However, an arbitrary linear boundary layer profile in the initialization (using a field function) may be required.

Solution

- The Expert Driver with default values works well.
- Appropriate CFL values are dependent on mesh quality and Mach number.

  For a high-quality mesh, a general rule-of-thumb is $CFL \approx 20 / Mach$ for this regime. For example $CFL \approx 40$ for Mach 0.5, and $CFL \approx 10$ for Mach 2.

If the Expert Driver is used, CFL ramping is necessary only if poor-quality cells are present or if the flow field is especially complex.

- Residuals may drop by only 2–3 orders of magnitude for this type of problem, especially if grid sequencing initialization is used (because grid sequencing initialization provides a better initial estimate of the flow field).
- Monitor the forces and moments on the object to determine convergence.

Vehicle Simulations

This section presents the following recommended practices to model external automobile aerodynamics:

- Choosing Between an Unsteady or Steady State Model
- Recommended Turbulence Model
- Recommended Mesh Settings
- Recommended Number of Inner Iterations
- Recommended Time-Step
- Recommended Monitors and Reports

Choosing Between an Unsteady or Steady State Model

Steady State

Steady state simulations are excellent tools for optimization studies to reduce drag. The advantages of running a simulation with steady state include:

- Quick turnaround times.
- Accuracy in absolute values and trends.

Unsteady

Unsteady simulations must be used when looking at cases such as vehicle handling, overtaking, or fleet modeling. In these cases, understanding the unsteady flow characteristics can improve the aerodynamics. Optionally, unsteady simulations can also be used in cases where transient eddy effects are important, such as aero-acoustics or wheel well modeling.

Unsteady simulations take longer to run. However, the overall turnaround time of an unsteady simulation can be reduced as follows:

- The time step can be increased using transient SIMPLE. However, if the time step is raised too high, the solution becomes non-physical.
- Run transient simulations with higher under-relaxation factors than steady simulations. The higher the relaxation faction, the fewer the inner sweeps that are required.
- Initialize transient simulations using a steady flow field. Use the K-Omega SST model to initialize the flow field and reduce the jump between steady and DES simulations.

- You can use either coupled or segregated solvers to initialize the flow field. The coupled solver has the advantage of reduced run time for larger case studies.

Recommended Turbulence Model

- Either SST (Menter) K-Omega or Spalart-Allmaras Detached Eddy Simulation model is recommended for vehicle aerodynamics. Both seem to give similar solutions.
- For vehicle aerodynamics, DES offers improved information on large eddy structures in the wake. For wake flow, DES is more accurate than URANS. So, for all studies of vehicle passing or wake flow, DES is recommended.
- Aero-acoustics studies have shown some cases where Spalart-Allmaras offers slightly better trends. Typically, for passenger cars, the SST K-Omega turbulence model is used.
- Large time periods are more critical than local flow structure for thermal transient simulations. Therefore, URANS is appropriate for this class of problem.

Recommended Mesh Settings

- The grid density must be refined enough to capture local physics. However, over-refinement slows down performance. Use volumetric sources to refine the mesh only in the required areas.
- The grid size must be fine enough to capture the turbulent length scale.
- Resolve the wake of vehicles and tires when modeling passenger cars. Smaller length scales are less critical. The wake is generally refined for half of the car length to the rear.

  A more realistic estimation is to calculate the frontal area. Based upon the frontal area of the vehicle, extend the same distance in the wake.
- Examining the results of a steady state analysis can help determine grid resolution for DES simulation using the following field functions:

Name	Function Name	Definition	Purpose
Turbulent Length Scale	**TurbulentLengthScale**	`pow(0.09,-1)*sqrt($TurbulentKineticEnergy)/`	Determines Turbulence Length Scale for SST

Name	Function Name	Definition	Purpose
		`$SpecificDissipationRate`	(Menter) K-Omega model
Grid Size	**GridSize**	`pow($Volume, 1.0/3.0)`	Estimate Edge Length of cell (based upon cubes)
Length Ratio	**LengthRatio**	`$TurbulentLengthScale/(2*$GridSize)`	To resolve turbulence properly, length ratio > 1 would be needed

1. Run a preliminary steady simulation.
2. Use the `TurbulentLengthScale` and `GridSize` field functions to calculate the Length Ratio, as defined in the table.

   The turbulent length scale is based upon the turbulent dissipation. The grid size is estimated based upon the volume of each cell
3. Use a mesh that yields a length ratio greater than one in the preliminary steady simulation for the transient analysis.

Recommended Number of Inner Iterations

- Lowering inner iterations to speed up the solution is a better strategy than increasing the time step. Large time steps adversely affect results.
- In most cases, 8 inner iterations are enough to converge a time step. Lowering inner iterations between 3–5 helps in reducing turnaround time.
- The overall drag value is not sensitive to the number of inner iterations. The grid density and refinement locations are more critical in getting accurate drag values.

Recommended Time-Step

- Time-step size depends on the physics that is being resolved.
  - For aero-acoustics, the time step must be small enough to resolve the desired frequency.
  - For aerodynamics, larger time steps are used to capture the large eddies in the flow field.

- For general drag values, resolve the large wake structures behind the tires/vehicle. A good estimate for a time step is approximately 100 time steps for a particle to reach the span of the vehicle:

$$\Delta t = \frac{\text{Length of Vehicle}}{100 \times \text{Vehicle Speed}}$$

- The courant number must be below 1 for most of the flow field. Small local peaks from flow accelerating around a sharp corner are fine if they spike over 1.

Recommended Monitors and Reports

Use the following monitors to check for convergence in inner iterations:

- Force on vehicle: this monitor reaches a steady value at convergence.
- Pressure probes in tire and vehicle wake

  80% convergence to final value is fine for drag prediction
- Velocity magnitude

  The peak velocity must be below 150 m/s. If the peak velocity is higher, raise the time scale coefficient for upwind blending (*CdesTimeLim*) expert property of the SST (Menter) K-Omega model to improve stability.

Modeling Fluid-Structure Interaction

This section describes the various techniques and considerations for fluid-structure interaction (FSI) cases, and how to deal with FSI cases in STAR-CCM+.

FSI can be modeled entirely within STAR-CCM+, without resorting to a third-party code. You can model FSI as a fluid interacting with a rigid body using dynamic fluid-body interaction (DFBI), or as a fluid interacting with a deformable body using the finite volume stress capabilities. Guidelines for the latter are provided in this chapter.

You can also model FSI problems by coupling STAR-CCM+ to a third-party code, whereby STAR-CCM+ solves for the fluid domain and the third-party code solves for the structure. STAR-CCM+ can be coupled to NASTRAN, ANSYS, Abaqus and Radtherm using the File-Based Coupling technique, or to Abaqus using Co-Simulation.

See Also

Modeling FSI Using Finite Volume Stress

What Is Fluid-Structure Interaction?

Fluid-structure interaction (FSI), broadly speaking, is the thermo-mechanical interaction of a fluid and a solid structure.

If one asked one hundred engineers what FSI means to him/her, one would likely get one hundred different answers. This section attempts to make some classifications of the interactions and the capabilities that are required to address FSI problems.

Classifications of FSI Coupling

This section describes the various types of fluid-structure interactions.

FSI can be grouped into two distinct categories:

- "One-way" interactions
- "Two-way" interactions

Each type of interaction is discussed below.

One-Way Interactions

In "one-way" interactions, the fluid can impart some action on the structure but the response of the structure to the fluid loading does little to affect the fluid motion. For example, a fluid can heat or cool a stiff, supported structure, which produces thermal stress loads and deformations in the solid material. However, these thermal deformations do not necessarily lead to any significant change in the flow patterns of the fluid. Such a problem involves a two-way exchange of heat, which is referred to as Conjugate Heat Transfer. Nonetheless, the mechanical exchange is only one way, so it is sufficient to compute the temperature in the fluid and solid domains in separate simulations. Then you can compute the stresses in another simulation using the results of the temperature that is computed in the previous simulation.

The one-way interaction can also be the other way. For example, the motion of a piston and valves in an internal combustion engine certainly affects the fluid flow patterns in the combustion chamber. However, the motion of the fluid has little effect on the deformation of the piston and valves (other than, or course, the thermal effects).

Two-Way Interactions

In "two-way" interactions, the fluid motion and pressure affect the displacement and deformation in the structure. The response of the

structure has a significant effect on the fluid flow. A prime example is the Dynamic Fluid Body Interaction capability of 6DoF to model the coupled interaction between a rigid body and the fluid.

Traditionally, FSI implies two-way coupling of a fluid and a deformable structure, such that the deformation and rigid body DOF are coupled with the fluid. Even under this definition, there are classifications of the degrees of coupling, ranging from “weak” or “loose” coupling to “strong” coupling.

In two-way interactions, it is clear that the solid structure responds to the pressure or viscous shear traction in the fluid. On the other hand, the fluid responds to the structure in various ways. The shape of the structure obviously affects the flow field, but it is the velocity of the structure (or change in displacement) that has the greatest effect. In fact, it can be shown that the acceleration of the structure causes proportional pressure responses much like how the acceleration of the fluid leads to a hydrostatic pressure gradient. For example, striking a hammer on a constrained pipe. In the fluid, the response would be a pressure wave that could be heard but not seen, and the small deformations in the pipe would be visually imperceptible. In this case, the response is determined by how fast the structure moved, or more accurately, how the structure accelerated. It is not determined by how much the structure moved.

Aero (Hydro) Elastic Equilibrium

In most cases, FSI implies “dynamic” simulations, or what would be referred to in CFD as “transient” simulations. However, in some situations, “static” solutions can be important. For example, the static deformation of a wing structure due to a “steady-state” airflow around the deformed wing. In such cases, a transient analysis in the fluid and structure can still be used to reach the static steady-state solution. However, the time-step can be thought of as a pseudo time-step, since the simulation does not need to be “time accurate” in its search for the steady solution. The fluid and solid time integration can be reduced to first order. Also, by definition, “static” implies that the velocity of the structure is zero. Therefore in the static solution you can ignore the grid flux terms in the fluid, since the grid flux terms are zero. On the structural side, you can increase the material damping or Rayleigh damping for the same reason. That is, when the velocity of the structure goes to zero, the damping forces are zero. More precisely, setting the damping to “critical” damping is the quickest way to drive the structure to the steady-state. Thus, ignoring the grid flux and adding structural damping promotes stability.

Weak Coupling

In some two-way FSI problems, you can regard the coupling to be "weak", and so employ a "loose" coupling algorithm to find the solution.

Consider a structure that bends in steady-state flow. This bending can significantly alter the flow patterns around the structure, which, in turn, affect the bending. When searching for a steady-state condition, one is interested in the shape the structure takes under the steady condition. As the structure and fluid reach the steady-state, the structure material point velocities approach zero. In this situation, the coupling is "weak" since in the steady state the structure imparts no motion on the fluid. In fact, it is almost "one-way", since the primary influence of the structure on the fluid is from its velocity, which in steady state is zero. In this situation, a "loose" coupling algorithm between the solid and fluid is sufficient to determine the end steady condition.

For example, the fluid can initially assume the structure is rigid and constrained, and a flow field can be computed independently. The pressure that is generated can be passed to a separate structural solver and the deformations in the structure can be computed from this initial guess of the fluid loading. Once the deformations in the structure are known, they can be provided to the CFD analysis to compute the steady flow with the new deformed geometry. This iteration can be repeated until a steady-condition is met. The important thing to notice is that this analysis has no consideration for the time scales. Since the point of interest is not how it reaches the steady-state solution, but simply what steady-state solution the fluid/structure reaches. The fluid solver only requires to know where the fluid/structure interface lies, not how fast it is moving or accelerating.

In a general sense, weak coupling implies that the response times of the structure to a disturbance in the fluid are slow, compared to the fluid. The opposite is also true, the response times in the fluid to a disturbance in the solid are slow, compared to the solid.

A distinction must be made between "loose coupling" and "loose coupling algorithms". Loose coupling implies the physical strength of the coupling, whereas loose coupling algorithms imply the degree to which the fluid and structure codes are coupled. The loose coupling algorithm can mean simply passing data back and forth between different codes using files on the disk. It can also mean the fluid and structure solvers are not necessarily in the processor memory at the same time. Also, most certainly it means a superficial coupling strategy that does not go deep into the solution algorithms of either code. In

particular, loose coupling algorithms are often explicit, in that the response of the solid depends on the conditions of the fluid at the previous time step. The opposite can also be said. This coupling strategy is often referred to as “staggered” coupling.

Strong Coupling

In “strong” coupling, the physical coupling is two way and the coupling between the codes is pronounced.

Often “strong” coupling is associated with “dynamic” problems where the hydrodynamic loads and the structural velocities change dramatically. In strong coupling algorithms, the fluid and structure solvers can be resident in processor memory at the same time. Data is passed at regular intervals (often referred to as coupling time or coupling step) from the memory that the structure solver uses to/from the memory that the fluid solver uses. Depending on the actual strength of the coupling it can be necessary to exchange data between the solvers more than once per time step for the simulation to remain stable. This type of coupling implies a deeper communication between the structure and fluid codes and is known as: “implicit”, “iterative staggered”, “iterative successive substitution”, or “multiple iterative coupling”.

Implicit coupling is necessary when a relatively light or compliant structure interacts with a relatively heavy, almost incompressible fluid. For example, a steel structure interacting with air is not coupled strongly in the physical sense. However, a steel structure in water can be coupled strongly as long as the structure is relatively light compared to the fluid it displaces or moves. This phenomenon is true even although the structure is relatively stiff compared to the bulk modulus of the water. A ship only floats because it is lighter than the water it displaces, and thus the interaction is coupled strongly. Indeed the Dynamic Fluid Body Interaction (DFBI) capabilities in STAR-CCM+ couple the rigid body dynamics of a structure with the fluid flow in an implicit manner.

Strong coupling can occur in situations that are not immediately evident. For example, flow of water through a stiff, vibrating pipe. In this case, the mass of the pipe can be smaller than the mass of the water that is contained within the pipe. Unless the coupling time step is small, explicit coupling is not stable in this situation. Such cases can be dealt with in STAR-CCM+ by taking advantage of the implicit coupling between the CFD and FV Stress solvers. This coupling algorithm is strong and deep, and the exchange of data occurs at each iteration within a time step. This coupling also considers the load balancing and the partitioning, and minimizes the cross-talk among different processors between the fluid and the solid.

Other Categories of Coupling Algorithms

As mentioned in the previous section, the DFBI capability in STAR-CCM+ is implicitly coupled. You can take advantage of this stable implicit coupling even when you are interested in the deformations and resulting stresses in the ship. If actual deformations about the rigid body modes are relatively small, then the fluid motion is not sensitive to the deformations, only the rigid body motions. Thus, for example, the CFD and rigid body modes can be solved together in this implicit fashion. Then the transient loads can then be applied to the structural solver, which accounts for deformations and stresses. This process, in effect, is a one-way coupling, since the deformations computed in the structure code are not passed back to the CFD solver. STAR-CCM+ can be controlled using Java macros to map and export fluid dynamic loads for use in a structural solver such as Abaqus.

Dynamic Mesh Evolution

The motion of a structure in a fluid requires the CFD solver to account for changes in the shape and position of the solid structure. In STAR-CCM+, various strategies are used for this problem.

The strategies are discussed in the following sections:

Mesh Rotation and Translation in FSI

The simplest type of mesh motion strategy is to displace the fluid mesh as if it were a rigid body. In this case, all the cells maintain their shape and the description of the mesh motion is from a displacement vector and Euler angles.

In DFBI, the position of the grid is determined by solving the equations of the assumed rigid structure, and the fluid transport equations automatically account for motion of the grid. This type of mesh motion is suitable when the rigid structure does not approach other solid structures. For example, this technique could be used for a ship that is in relatively deep water with no other structures or ships nearby. If other structures were nearby, the mesh is sometimes required to morph to account for the relative motion between the structure and other structures.

STAR-CCM+ also allows the fluid mesh to rotate with a rigid structure, and the rotating mesh to be embedded within another stationary mesh in a different region. The two mesh regions can then slide with respect

to each other at an interface. The solver automatically calculates this interface projection and imprinting at each time step.

Morphing in FSI

Morphing, in this context, is the deformation of the fluid grid. Morphing occurs by moving the fluid vertices in such a manner as to conform to the solid structure and maintain a reasonable quality fluid grid.

This process is often referred to as a "topologically constant" operation since all the cells maintain the same neighbors but the shapes of the cells can change over time. The fluid transport equations are solved to account for this arbitrary motion of the mesh. This technique is sometimes referred to as the Arbitrary Lagrangian Eulerian (ALE) technique, but often ALE is associated with FEA, rather than FVA. STAR-CCM+ uses a "space conservation law" to conservatively and accurately express the transport motion. The key word in ALE is arbitrary, since the motion of the vertices is arbitrary except that it must conform to the moving boundaries and it must provide a reasonable quality mesh.

The morpher in STAR-CCM+ currently uses a multi-quadric morphing model that is based on radial basis functions to define the motion of interior vertices. This model originates from the motion of the vertices on the structural surface. When the structure moves in relation to other solid structures, the fluid mesh is distorted to conform to the structures. However, when two structures come into contact or move nearby, a morphing strategy can lead to poor cells.

It is not possible to morph between topologies that are dramatically different. In this case, you can:

- use the overset mesh capability and morph the overset mesh around the deforming structure, or
- use the advanced automatic meshing capabilities in STAR-CCM+ to remesh the model completely, using the most recent boundaries as a starting point.

Overset Mesh for FSI

When two structures approach each other, the fluid mesh between them must be squeezed in such a way as to retain the boundary shape that is formed by the structures. When two structures move apart, the fluid mesh must expand. In both cases, there is a threshold beyond which the morpher cannot avoid creating poor quality cells. For such situations,

the overset mesh technology in STAR-CCM+ provides an effective solution.

Overset meshes that wrap around deformable structures can still be morphed using the structural deformation for the control points. If structures drift apart, however, most of the distance between the structures is treated using the background mesh.

Overset interpolation requires that the overset mesh always overlaps the background mesh by at least 4–5 cells rows. For this reason, it is not possible to close a gap completely when using the overset technique. However, you can still satisfy the overall interpolation requirement for a small gap size by adding more prism cell layers on the walls of both approaching structures.

Gap Closure

To mimic gap closure in STAR-CCM+, make the final gap sufficiently small that it gives a high resistance to fluid moving through it. If you intend to stop the flow (as in a stop valve), then drive the flow using pressure boundary conditions at both the outlet and inlet. In these circumstances, the maximum flow rate is proportional to the pressure difference across the gap, and inversely proportional to the viscous resistance of the flow within the gap. Ultimately the gap does leak fluid, but potentially at a small enough rate for the simulation to be an acceptable engineering approximation for the needs of FSI analysis.

Mapping Between Structure and Fluid Grids

Another challenge that must be overcome in FSI cases is the difference in the resolution between the fluid and structure grids. Often the grids are different due to the difference in physical processes in the fluid and the structure.

If the structural mesh and the fluid mesh are both constructed in STAR-CCM+, the two meshes can be made conformal at the fluid-structure interface. That is, the vertex positions on the fluid surface match the vertex positions on the surface of the structure. In this case, the mapping of fluid loads and structural displacements is more or less trivial.

When the opposing meshes are not conformal, STAR-CCM+ offers various strategies to map the data efficiently and accurately to and from meshes that are imported from other CAE products. This mapping is accurate, continuous, bounded, and as conservative as possible. When the opposing surface is an FEA mesh, STAR-CCM+ uses the innate

shape functions of the finite element topology to map between the structure and fluid meshes.

Reference Configuration

Mapping requires that fluid and solid meshes must be similar at some reference configuration. The meshes must be sufficiently similar that the neighbors of faces on one side are matched to faces on the opposite side. The reference configuration is usually the geometry at the initial time. The fluid and solid mesh are assumed to move together so that their neighbors remain the same. For this reason, the weights of the mapping are defined entirely by the geometry of the mesh at the reference configuration—the mapping does not depend on the current deformed geometry.

In the data mappers, you can choose whether the reference configuration is set by the original mesh or by the current mesh, which could be deformed as a result of morphing. When the reference configuration is set to use the current mesh, STAR-CCM+ performs a neighbor search each time the mesh is moved or morphed. As this neighbor search is an expensive operation, you are recommended to set the reference configuration to the original mesh. STAR-CCM+ generally uses the original mesh configuration when it automatically sets the data mapping schemes appropriate for the physics in co-simulation. It also uses this configuration for mapped fluid/solid interfaces within a single simulation.

The assumption that neighbors at the reference configuration will always be neighbors is violated if the fluid mesh slides with respect to the solid mesh. In this case, the reference configuration must always be the current deformed mesh and neighbors must be re-established at each new time step.

The reference configuration need not always be at the initial time or at a re-meshing event. For example a structure can be pre-loaded to deform the shape. The fluid mesh can be initially constructed to conform around the deformed shaped of the structure (for example, the deformed tire loaded by the weight of the car and contact with the pavement). The reference configuration is now the deformed shape at this point in time.

HPC and Grid Partitioning on Multiple Processors

Modeling FSI also presents challenges regarding the partition of the fluid and structural grids.

In STAR-CCM+, when one employs the FV Stress solver for the structure, the grid is automatically partitioned for load balancing and to minimize processor-to-processor communications. In this case, the partition is not required to respect the boundaries between the structure and the fluid. In other words, any particular partition can have part of the fluid grid and part of the structural grid. It is this unique feature that makes STAR-CCM+ efficient for FSI applications on multiple processors.

Modeling FSI Using Finite Volume Stress

The **Finite Volume Solid Stress** model (or "FV stress") makes it possible to simulate fluid-structure interaction (FSI) entirely within in STAR-CCM+.

Depending on the case, large or small displacements of the solid domain can be involved, and guidelines for setting up each case are covered in the following topics:

Using Large Displacements for FSI

The **Large Displacement** model can be used in a fluid-structure interaction (FSI) case to simulate movement in the solid mesh.

Guidelines describing the settings that are required for preparing such a case using implicit coupling and the large displacement model are described in the following pages:

Selecting Physics Models

This section describes the physics models that you select to use large displacements.

1. For the solid continuum, select the **Finite Volume Solid Stress** model and enable Large Displacements by activating the **Solid Displacement** motion.
2. For the fluid continuum, activate the **Morphing** motion so that the fluid mesh can move in response to motion in the solid mesh.

If you are modeling a liquid, it can be useful to employ the **User-Defined Density** model instead of the default **Constant Density** (incompressible) model for the purposes of both numerical stability and accuracy. When you are working with a user-defined density in an FSI case, it can be useful to specify the liquid as compressible. To do so:

3. Activate the **Compressible** option for the **User-Defined Density** model under the **Continua** manager node.



The interaction of a liquid with a solid is often more numerically stable when the liquid is allowed to be slightly compressible, particularly when the flow is internal to the structure. To provide this numerical stability, you can define a user-specified density using a field function. For example, Density `rho = 1000 + $Pressure/($c*$c)`, Density pressure derivative `drhodp = 1/($c*$c)` and Speed of Sound `c = 600`. In this case, it is often beneficial to use a slower speed of sound than the "true" value for the liquid. This setting leads to better numerical stability and faster convergence without a loss of accuracy.

The choice of using a compressible liquid or not depends on the physics and appropriate time scales. When the fluid sound transient times are short relative to other time scales (such as natural frequency of the structure or vortex shedding frequencies), it is possible to reduce the sound speed of the fluid. Reducing this value has little influence on the results, while improving stability and robustness. If the physical relevant time scales are of the same order as the sound transit times (such as an impact or underwater

explosion), then it is necessary to account for the exact nature of the fluid compressibility accurately and employ more accurate user-specified density laws for the liquid.

For external flows, it is possible to use an incompressible fluid. If the fluid is a gas, then using a perfect gas law is appropriate. These models can be selected from the *Equation of State* box in the *Physics Model Selection* dialog.

Setting the Fluid Morpher Boundary Condition

Activate the morpher and select a morpher physics condition for the interface boundary.

1. For the fluid region, activate the mesh morphing motion so that the interface boundary in the fluid domain morphs according to the solid stress displacement.
2. Once morphing has been activated for the region, select the morpher physics condition under the boundary interface node and specify **Solid Stress** for the *Method* property.



Preparing the Fluid-Structure Interface

At the interface between the solid and fluid domains, there are two nodes available that affect the FSI solution. The **FSI Coupling Method** node allows you to control the frequency of data exchange across the fluid-solid interface. The **FSI Under-Relaxation** node can be used to set the mesh motion under-relaxation factor at the FSI interface.

For an implicit case, set the **FSI Coupling Method** to **Implicit**, which causes data to be exchanged after each inner iteration. The other option is **Explicit**, which causes data to be exchanged after each time-step.

For Explicit, the fluid morphing is only applied once per time-step. For Implicit, there are two options:

- The fluid morpher is executed only on the first iteration of the time-step. Thereafter a correction is made only on the boundary to account for the acceleration of the structure. (Default)
- The fluid morpher is executed for each iteration within a time-step. This function is activated using the *Morph At Inner Iterations* property on the Morpher solver.



Using Small Displacements for FSI

The small displacement theory assumes that the displacements are relatively small compared to the cell size. In this case, the solid mesh is not moved and the adjacent fluid mesh does not need morphed. Instead, the velocity of the solid is imparted as a boundary condition to the fluid.

Generally, the assumption of small displacement FSI is only applicable to vibration problems where the maximum displacements are small. However, the frequency of the oscillations can be large - implying that

the velocity and the acceleration of the solid structure are significant. Discussions regarding the compressibility of liquid fluids (see Using Large Displacements for FSI) are still relevant to the small displacement theory.

Modeling Atmospheric Boundary Layers

This section describes a recommended technique for ABL (Atmospheric Boundary Layer) modeling in STAR-CCM+.

The ABL is an important part of simulating wind and is used in modeling weather, wind-borne pollution, and designing wind-farms.

This chapter contains the following sections:

- What is the ABL?
- Setting the parameters
- Recommended values
- Bibliography

What is the ABL?

The ABL (Atmospheric Boundary Layer) is the bottom layer of the atmosphere which is in contact with the surface of the Earth.

Within this layer, the presence of the surface of the Earth directly affects flow: frictional effects and temperature create turbulence and vertical mixing. Above this layer, rotational effects dominate the flow rather than turbulence. The depth of the ABL can vary in both time and space.

Garratt decomposes the boundary layer in the atmosphere of the Earth into two sublayers:

- An outer Ekman sublayer where there is little dependence on the nature of the surface.
- An inner sublayer where surface effects such as friction and heating are important.

There is an interfacial sublayer within this inner layer that defines the height and structure (together termed “roughness”) of obstacles to the flow on the ground.



Index	Description
1	Ekman sublayer
2 & 3	Inner sublayer
3	Interfacial sublayer

Formulation for ABL in STAR-CCM+

In the literature, the preferred method of modeling the roughness is to use the Richard and Hoxey Law for the velocity distribution, u^+_{RH}:

$$u^+_{RH} = \frac{1}{\kappa}\ln\left(\frac{y_0 + y}{y_0}\right) \tag{2692}$$

where:

- κ is the von Karman constant.
- y_0 is the physical roughness length
- y is normal distance from the wall to the wall-cell centroid.

For ABL simulations, the flows are predominantly turbulent, and the velocity distribution maps to Eqn.(305) in STAR-CCM+. Following [435], let Eqn.(2607) equal Eqn.(305) and cancel the common terms to give:

$$\frac{y_0 + y}{y_0} = E' y^+ \tag{2693}$$

Substitute E'using Eqn.(453) and Eqn.(468) for rough walls ($R^+ > 90$, $B = 0$):

$$\frac{y_0 + y}{y_0} = \frac{E y^+}{C R^+} \tag{2694}$$

Substitute y^+ using Eqn.(446) and R^+ using Eqn.(467). Cancel common terms, giving:

$$\frac{y_0 + y}{y_0} = \frac{E y}{C r} \tag{2695}$$

Finally, assume that $y \gg y_0$, so that:

$$\frac{y}{y_0} = \frac{E y}{C r} \tag{2696}$$

Canceling common terms provides the relation to modify the STAR-CCM+ wall roughness equations to model an ABL:

$$r = \frac{Ey_0}{C} \tag{2697}$$

where:

- E and C are wall function coefficients.

To guarantee physical results, ($u^+ > 0$) make sure that $y \gg y_0$ is satisfied and deactivate the *Roughness Limiter* Wall Roughness Parameter.

Setting ABL Parameters in STAR-CCM+

To model an ABL, use a high Reynolds number turbulence model (see the recommended settings for K-Epsilon and K-Omega) and a high y^+ model.

Activate wall roughness and specify the **Roughness Height** using values from the **Standard Wall Function** and the **Wall Roughness Parameters** nodes using Eqn.(2612). For example, in a case with y_0 equal to 0.01 m, E equal to 8.75, and C equal to 0.25, and $y \gg y_0$, the calculated roughness height r is 0.35 m. To set the values that are mentioned:

1. For the wall boundary, select the **Physics Conditions > Wall Surface Specification** node and set the *Method* to **Rough**
2. Right-click the **Physics Values** node and select *Edit*

In the Edit panel, enter the values for the example which is specified above:

3. Expand the **Standard Wall Function** node and set the value of E to `8.75`
4. Expand the **Wall Roughness Parameters** node and set the value of C to `0.25`
5. Make sure that *RplusRough* (R^+) is set to `90` or greater so Eqn.(2609) is valid.
6. Make sure that the *Roughness Limiter* is deactivated because the assumption $y \gg y_0$ from Eqn.(2611) is valid.
7. Expand the **Roughness Height > Constant** node and set the resulting calculated *Value* to `0.35 m`

The panel appears as follows:


Edit
Expand/Contract Tree
Expand/Contract Values
Nodes
Values
Physics Values
Roughness Height
Method
Constant
Dimensions
Length
Constant
Value
0.35 m
Standard Wall Function
E
8.75
Kappa
0.42
Wall Roughness Parameters
C
0.25
B
0.0
RplusSmooth
2.25
RplusRough
90.0
Roughness Limiter
Multiple selections
Constant profile value ...
Close
Help

Recommended Values for ABL Modeling

Use the settings described in this section when creating an ABL model.

Turbulent Viscosity Ratio

- Set the Turbulent Viscosity Ratio to `1E10`

K-Epsilon

- Set *Cmu* (c_μ) according to the inlet velocity profile. Typical values are between 0.03 and 0.11
- Set *Ce1* ($c_{\epsilon 1}$) coefficient using the following equation:

$$c_{\epsilon 1} = c_{\epsilon 2} - \frac{\kappa^2}{\sqrt{c_\mu}\sigma_\epsilon} \tag{2698}$$

where the coefficients are defined in the turbulence modeling section.



Standard K-Epsilon - Properties	
Properties	
Constitutive Relation	Linear
Realizability Option	None
Convection	2nd-order
Expert	
Normal Stress Term	☐
Sigma_e	1.3
Sarkar	2.0
Sigma_k	1.0
Tke Minimum	1.0E-10
Secondary Gradients	On
Tdr Minimum	1.0E-10
Buoyancy Production of Dissip	Boundary Layer Orienta...
Cmu	0.09
C1e	1.44
C2e	1.92
Ct	1.0

K-Omega

- Set *BetaStar* (β^*) according to the inlet velocity profile, where $\beta^* \approx c_\mu$.
- Set *Beta* (β) using the following equation:

$$\beta = \frac{\beta^*}{1.25} \tag{2699}$$

- Set *Alpha* (a) using the following equation:

$$a = \frac{1}{1.25} - \frac{\sigma_\omega \kappa^2}{\sqrt{\beta^*}} \tag{2700}$$

- Set *Sigma_w* (σ_ω) to `0.5` (its default value)

where the coefficients are defined in the turbulence modeling section.



Mesh Considerations

From sensitivity studies [435], use the following values when setting up the mesh:

- Make sure the shortest distance from the wall to a cell centroid, y, so that $y > 0.2r$ and $y \geq 10y_0$.
- Make sure that $y^+ \geq 30$, due to the Richards Hoxley model [436] starting in the logarithmic region of the boundary layer.

Bibliography

Garratt, J.R. 1992. "The atmospheric boundary layer" *Cambridge University Press*, pp. 1–15.

Martinez, B. 2011. "Wind resource in complex terrain with OpenFOAM" *MSc Thesis*, **Riso DTU, Denmark**.

Richards, P. J., and Hoxey, R. P. 1993. "Appropriate boundary conditions for computational wind engineering models using the k-epsilon model" *Journal of Wind Engineering and Industrial Aerodynamics*, **46**, pp. 145–153.

Preparing CAD Models for Import

When you import a geometry file into STAR-CCM+, separate part surfaces are created for each named face. You can name faces in most CAD packages.

Instructions on how to name faces are provided for the following CAD packages:

Preparing CAD Models in CATIA

To name a face in CATIA you must extract the face you want to name and then sew it back onto the geometry. The name of the extracted feature is the name of the face.

To name a face in CATIA:

1. Launch CATIA and load the geometry.
2. Select **Start** > **Shape** > **Generative Shape Design**.
3. Click (**Extract**) and select a face to extract.
4. Click **OK**.
5. Select **Start** > **Mechanical Design** > **Part Design**.
6. Click (**Sew Surface**) and select the extracted surface.
7. Deactivate the *Simplify Geometry* option.
8. Click **OK**.
9. Hide the **Extract** feature.
10. Rename the **Sew Surface** feature.

    This feature is imported into STAR-CCM+ as a separate surface with the name you specified.
11. Save the geometry.

Preparing CAD Models in SolidWorks

In SolidWorks you can name a face by editing its face properties.

To edit the face properties:

1. Launch SolidWorks and load the geometry.
2. Right-click the face you want to name and click ⌄ to expand the menu.
3. Select **Face > Face Properties...**



4. In the *Entity Property* dialog, enter the name of the face.
5. Click **OK**.
6. Save the geometry.

Preparing CAD Models in NX

In NX, you can name faces by editing their properties.

To edit the face properties:

1. Launch NX and load the geometry.
2. Set the *Type Filter* to **Face**.
3. Right-click the face you want to name and select **Properties**.
4. In the *Face Properties* dialog, enter the name of the face.
5. Click **OK**.
6. Save the geometry.

Preparing CAD Models in Creo and Creo Elements/Pro

In Creo and Creo Elements/Pro, you can name faces through the Model Properties.

To name the faces:

1. Launch Creo and load the geometry.
2. Open the *Model Properties* dialog:
   - In Creo, select **File > Prepare > Model Properties**.
   - In Creo Elements/Pro, select **File > Properties**.
3. In the *Model Properties* dialog, locate the *Features and Geometry* group box and click **change** next to the *Names* property.



4. With the *Rename* dialog open, select the face that you want to rename on the geometry model.

   An entry is added to the *Rename* dialog. If the face corresponds to an existing entry, the name is highlighted.

5. Enter the new name for the face and press `<Enter>`.



6. Click **OK**.
7. Click **Close**.
8. Save the geometry.

Migrating from Other Codes

This section contains guidelines for migrating from the following codes:

- Fluent
- STAR-CD

Migrating from Fluent

This section lists details and limitations that guide you in migrating your Fluent files to STAR-CCM+.

To import Fluent `.cas` files into STAR-CCM+, remember:

- Only mesh data is imported. No boundary conditions, model selections, solver parameters or solution fields are imported. Therefore, any solution parameters or boundary conditions originally that were contained in Fluent must be redefined in STAR-CCM+.
- The Fluent importer in STAR-CCM+ reads the following sections of the file:

```
FLUENT_HEADER_SECTION = 1,
FLUENT_DIMENSION_SECTION = 2,
FLUENT_MACHINE_SECTION = 4,
FLUENT_NODE_SECTION = 10,
FLUENT_CELL_SECTION = 12,
FLUENT_FACE_SECTION = 13,
FLUENT_PERIODIC_SECTION = 18,
FLUENT_ZONE_SECTION = 39,
FLUENT_NAME_SECTION = 45,
```

- The following region, boundary, and interface types are recognized:

```
FLUENT_FLUID = 1,
FLUENT_INTERIOR = 2,
FLUENT_WALL = 3,
FLUENT_PRESSURE_INLET = 4,
FLUENT_PRESSURE_OUTLET = 5,
FLUENT_SYMMETRY = 7,
FLUENT_PERIODIC_SHADOW = 8,
FLUENT_PRESSURE_FAR_FIELD = 9,
FLUENT_VELOCITY_INLET = 10,
FLUENT_PERIODIC = 12,
FLUENT_POROUS_JUMP = 14,
FLUENT_SOLID = 17,
FLUENT_MASS_FLOW_INLET = 20,
FLUENT_INTERFACE = 24,
FLUENT_PARENT = 31,
FLUENT_HANGING_NODE_PARENT = 32,
FLUENT_OUTFLOW = 36,
FLUENT_AXIS = 37
```

  Boundaries with those types are assigned to equivalent boundary types in STAR-CCM+. If boundary types are not specified in Fluent, STAR-CCM+ tries to infer them from the boundary names. Any

name that cannot be recognized as a particular type is interpreted as a wall.

- Changing the internal boundaries to the porous jump type in Fluent helps retain them as interfaces in STAR-CCM+. Unlike Fluent, STAR-CCM+ only allows one set of interior faces per region. However, the STAR-CCM+ importer converts porous jump boundaries to interfaces.

Migrating from STAR-CD

This section lists details and limitations that guide you in migrating your STAR-CD files to STAR-CCM+.

To convert a STAR-CD simulation to run under STAR-CCM+, remember:

- STAR-CCM+ contains a *Model Selection* dialog where all physical and numerical models are chosen. The dialog contains an **Auto-select recommended models** checkbox which is ticked by default. Accept this setting unless you require special selections.
- The high-Reynolds number formulation of the "standard" K-Epsilon turbulence model is commonly used in STAR-CD simulations. On the other hand, the **Auto-select** option is likely to lead to a Realizable Two-Layer K-Epsilon model. In most cases, the Realizable Two-Layer K-Epsilon model has acceptable, or even better results.
- The default STAR-CD setting for initial k and ϵ values (as displayed on the STAR-CD GUI) is 0. When migrating your simulation, you are advised to use the STAR-CCM+ (non-zero) defaults. Better still, use values that are in keeping with representative velocity, turbulence intensity and turbulent viscosity values in your problem.
- There are important differences in the definition and required boundary conditions for inlet boundaries. Note in particular that, for mass flow inlets, the STAR-CD boundary conditions include velocity, density, static temperature and a "Flow Switch" setting. The corresponding STAR-CCM+ conditions are mass flow rate, flow angle, pressure and total temperature.
- For problems that involve porous media, STAR-CD uses settings for a and β. The corresponding settings in STAR-CCM+ are porous inertial resistance $\mathbf{P}_i$ and porous viscous resistance $\mathbf{P}_v$, respectively.
- Follow STAR-CCM+ guidelines on choosing suitable boundary conditions for your type of problem. For example, use pressure boundaries rather than flow-split outlets for compressible flows. As a result, change the definition of some of your existing STAR-CD boundary regions to more appropriate types.
- There is no direct equivalent for the STAR-CD Maximum Residual Error Tolerance for use as an overall solution convergence criterion in steady-state problems. Similar criteria exist in STAR-CCM+ and are available via the Monitors node, with each transport equation being solved. To form a composite stopping criterion:
    - Set up monitors for one or more of these equations.

  - Select the Create Stopping Criterion from Monitor option
  - Combine all monitor stopping criteria with the Maximum Steps criterion.

- The default setting for convective flux discretization is a likely cause for significant differences in convergence rate between the two codes. STAR-CD uses first-order upwind differencing. STAR-CCM+ uses a second-order differencing scheme which, in general, offers a better trade-off between accuracy and convergence rate.
- To produce the equivalent of a STAR-CD surface (or hidden-line) velocity vector plot near wall boundaries, use the **Cell Relative Velocity** field function.
- All of the STAR-CCM+ post-processing plot types are defined in a suitable **Scene**. At the end of the simulation run certain plots, such as the **Cell Relative Velocity**, require that you click the **Initialize** button before the plot can be produced.
- The STAR-CCM+, option **Thermodynamic Polynomial Data** for specific heat is equivalent to the Polynomial option in STAR-CD when specifying physical properties for chemical reaction components. The polynomial curve fits for Cp (plus enthalpy and entropy) are read in from the Chemkin database. With these curve fits, values for *Heat of Formation* and *Standard State Temperature* are implicit in the curve fit for enthalpy. STAR-CCM+ offers other options for specific heat: **Polynomial in T** (Cp is a piece-wise polynomial in temperature) or **Constant** (Cp is constant). If you use either option alone, specify the *Heat of Formation* and *Standard State Temperature* to determine the enthalpy in a simulation involving combustion.
- The default setting in STAR-CD for combustion problems using the Eddy Break-up model is to include reaction product mass fractions in the calculation of the chemical reaction rate (option Use Products). STAR-CCM+'s default setting for the rate calculation is 'Do not use products'. This means that the chemical reaction simulation starts as soon as fuel and oxidizer come into contact. If you tick the 'Use products' option, set up an ignition model explicitly under the **Ignitors** node .

Running

This part of the documentation describes preparation and procedures for running a STAR-CCM+ simulation.

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Modeling physics -- the concepts, theory and implementation details of the models
- Using tools in STAR-CCM+ to facilitate your work
- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format
- A guide to using STAR-View+, a viewer for three-dimensional STAR-CCM+ scenes

See Also

Using STAR-CCM+
Modeling Geometry
Meshing
Best Practices
Analyzing
Wizard Guide
Tutorial Guide

Using Initial Conditions

Initial conditions in a physics continuum specify the initial field data for the simulation.

The **Initial Conditions** manager node, which is subject to common menu items, contains the initial condition nodes for the continuum.



This section defines what initial conditions are in STAR-CCM+ and describes how to set initial conditions.

See Also

Troubleshooting the Solution
Getting Past a Poor Initial Solution

What Are Initial Conditions?

Initial conditions in the continua specify the initial field data for the simulation.

For steady-state simulations, the converged solution should be independent of the initial field. However, the path to convergence, and hence the computational effort that is required to reach convergence, is affected. Therefore, choose the initial conditions and values judiciously, particularly when the physics is complex.

Each model requires sufficient information for the primary solution data for the primary variables that are associated with the model to be set. In most cases, the initial conditions are set up directly. For some models, such as turbulence, there is the option of specifying the information in a more convenient form (turbulence intensity and turbulent viscosity ratio instead of turbulent kinetic energy and turbulent dissipation rate).

Examples of initial conditions are:

- Pressure
- Temperature
- Velocity components
- Turbulence quantities

Setting Initial Conditions

This section describes how to set initial conditions.

Initial conditions can be set in three ways:

- Using constant values
- Using field function values
- Using tabular data values

They can also be set for a particular region.

To change the method:

1. Open the **Initial Conditions** node and then open the node for the relevant condition.
2. In the *Properties* window, click the drop-down list corresponding to the *Method* property.



This action permits the selection of one of the following:

- **Constant**
- **Field Function**
- **Table (iteration)**
- **Table (r)**
- **Table (time)**
- **Table (x,y,z)**
- **User Code**

See Also

What Are Initial Conditions?

Using Constant Values

Generally, most initial conditions are specified as constants.

This value is entered as either a scalar or vector value, depending on the specific initial condition.



Using Field Function Values

Field functions are useful for specifying initial conditions that are not constant values.

For instance, if you initialized the velocity in the L-shaped duct that is shown schematically below:



An initial velocity of 0 m/s is acceptable. Although, a better initial condition would be 1 m/s in the x-direction in the first part of the duct and 1 m/s in the -y direction after the bend. To set up this scenario, use a vector field function which could be defined as follows:

```
[($$Centroid[0] < 5) ? 1 : 0, ($$Centroid[0] < 5) ? 0 :
-1, 0]
```

Suppose that this field function was named **Initial Velocity**. This field function would then be selected in the list corresponding to the *Vector function* property in the *Properties* window.



Although the above example was for a vector value, and thus required a vector field function, the process would be identical for scalar values using scalar field functions.

Take care to use a field function that is likely to be initialized before being used. For instance, problems are likely to occur if you initialize pressure in terms of velocity, since it is not clear that velocity is available at the time that pressure is initialized.

Using Tabular Data Values

Tabular data sets are useful for transferring the data from one solution to use as an initial condition for another solution.

Since the tabular data file format is not compact, this method is not a suitable method of data transfer between large simulations. However, this method is effective for transferring a coarse-mesh solution to a finer mesh.

If you were to do the following:

- Extract a tabular data set containing pressure from a previous simulation.
- Export the data set to a file called `initialData`.

- Read the tabular data file in the current simulation.
- Assign the name of this tabular data set to the *Table* property in the *Properties* window of the table node (**Table (iteration)** in the following screenshot) under **Pressure**.

Then the field in the tabular data corresponding to pressure would also be assigned to the *Data* property.



The tabular data set is not partitioned for a parallel case which means that the entire data set is loaded into each parallel process that is allocated to the simulation. This process could result in memory overflow, as multiple copies of the data set are created. For this reason, initialization using tables is not recommended for parallel runs.

Setting Initial Conditions for a Particular Region

This section describes how to customize initial conditions for an individual region.

By default, the initial conditions for every region in a simulation come from the settings in the continuum. However, you can customize initial conditions for an individual region.

To customize initial conditions:

1. Select the **Regions** > **[region]** > **Physics Conditions** > **Initial Condition Option** node.



2. In the Properties window, set *Option* to `Specify Region Values`.



A new **Initial Conditions** manager node appears within the region node.



Its sub-nodes, and their default values, appear identical to the STAR-CCM+ defaults of the **Initial Conditions** manager node in the continuum.



The initial condition nodes of the region, rather than the ones in the continuum, now control the initial condition values for this region.

The initial condition settings in the continuum still affect any other regions in the simulation that have not been changed in this way. However, any changes you make to the initial conditions of the continuum, either before or after you add initial condition objects to a region, will not affect the initial conditions of the region.

Using Solvers

This section outlines the use of solvers in a simulation.

Most solvers that are seen within the **Solvers** node of a simulation tree are closely related to the models selected in a continuum, and are documented in the relevant modeling section.



To access solver properties, select the node for that solver in the **Solvers** node as shown below.



When solver nodes are opened, they contain algebraic multigrid solvers. Selecting the **AMG Linear Solver** node allows you to adjust the multigrid parameters.

The solver parameters that most commonly need adjusted are as follows:

- The Coupled Implicit solver parameters and Coupled Explicit solver parameters for the Coupled Flow model and Coupled Energy model.
- The solver parameters for the pressure solver and the velocity solver of the Segregated Flow model.
- The solver parameters for the energy solver that the Segregated Fluid Energy model uses.

See Also

What Are Simulations?
Working with the Object Tree

Setting Up Stopping Criteria

You can use stopping criteria to control how the simulation runs.

The **Stopping Criteria** manager node, which has its own pop-up menu, contains all of the stopping criterion nodes in the simulation tree.



This section explains the following:

See Also

What Are Simulations?
Working with the Object Tree

What Are Stopping Criteria?

Stopping criteria allow you to specify how long the solution runs for and under what conditions it stops iterating and/or marching in time.

Each specified stopping criterion is evaluated at the completion of every simulation step and a logical rule is used to determine if the interaction of all of the criteria stops the solver.

Using Automatically Generated Stopping Criteria

Certain stopping criteria are generated automatically when a solver is chosen.

To view the available criteria, expand the **Stopping Criteria** node. For steady simulations, there are two stopping criteria:

- Maximum Steps
- Stop File



For unsteady simulations, there are four criteria:

- Maximum Inner Iterations
- Maximum Physical Time
- Maximum Steps
- Stop File



Automatically generated stopping criteria cannot be deleted, but the *Enabled* property can be activated or deactivated.

Setting the Maximum Number of Steps

The **Maximum Steps** stopping criterion allows you to specify the maximum number of iterations in a steady solver or the maximum number of time-steps in an unsteady solver.

The node of this criterion has its own properties, and is added automatically to the object tree by selecting any of the *Time* models (explicit unsteady, implicit unsteady, or steady).



This criterion cannot be manually deleted.

The stopping decision is based on the number of steps that the solver executes, including any steps that are executed in a previous session. If you Clear the solution, the counter resets to zero, if you initialize the solution it does not.

To change the number of steps, click the **Maximum Steps** criteria node and modify the value of the *Maximum Steps* property in the *Properties* window.

Properties

Enabled	Determines if this criterion is used for stopping the solver.	
	Activated	Use this criterion.

	Deactivated	Do not use this criterion.
Maximum Steps	The maximum number of steps to take before this criterion is satisfied. The default is 1000 steps.	
Criterion Satisfied	Indicates if the stopping criterion has been satisfied (read-only).	
Logical Rule	Defines how this criterion interacts with other stopping criteria that have been enabled.	
	AND	Requires that this criterion, along with one or more other criteria with the *Logical Rule* set to **AND**, all be satisfied before stopping the solver.
	OR	Indicates that satisfying only this criterion is sufficient to stop the solver.

See Also

Using Solvers
Clearing the Solution
Initializing the Solution
What Are Stopping Criteria?
Using Multiple Stopping Criteria
Common Pop-Up Menu Items

Setting a Stop File

The **Stop File** criterion allows you to specify the pathname of a file (named `ABORT` by default) that, once in place, causes the solver to stop.

This action can be useful for stopping a batch job, for example. The node of this criterion has its own properties, and is added automatically to the tree by selecting any of the *Time* models (explicit unsteady, implicit unsteady, or steady).



Note: Activating this criterion does not stop the run immediately because STAR-CCM+ checks once every 60 seconds for the stop file.

This criterion cannot be manually deleted.

To set the path name of the file, click the **Stop File** criterion node and modify the name and the path in the *Path* property. You can either modify the name directly with a text entry or click the ellipses button and enter the path in the *Stop File - Path* dialog.



To use this feature while a simulation is running, you can simply add a file as follows:

- On Windows, place a file with the correct name in the correct path; even an empty text file would suffice for this purpose.
- On Linux, use the line command:

```
% touch ABORT
```

Properties

Stop Inner Iterations	Controls if this criterion responds to the completion of an inner iteration or a time-step.	
	Activated	Stops after the current internal iteration.
	Deactivated	Stops when the current time-step is finished (when the appropriate number of internal iterations are complete).
Path	The fully qualified file name of the stop file, `ABORT` by default.	
Enabled	Determines if this criterion is used for stopping the solver.	
	Activated	Use this criterion.
	Deactivated	Do not use this criterion.
Criterion Satisfied	Indicates if the stopping criterion has been satisfied (read-only).	
Logical Rule	Defines how this criterion interacts with other stopping criteria that have been enabled.	
	AND	Requires that this criterion, along with one or more other criteria with the *Logical Rule* set to **AND**, all be satisfied before stopping the solver.

	OR	Indicates that satisfying only this criterion is sufficient to stop the solver.

See Also

Using Solvers
Working with the Object Tree
What Are Stopping Criteria?
Using Multiple Stopping Criteria
Common Pop-Up Menu Items

Setting the Maximum Physical Time

The **Maximum Physical Time** stopping criterion is based on the simulation time that has elapsed in a transient analysis.

The node of this criterion has its own properties, and appears when either the implicit unsteady model or the explicit unsteady model is chosen.



When used with the implicit unsteady model, this stopping criterion is linked to the *Time-Step* property of the **Implicit Unsteady** node in the **Solvers** node. If the *Time-Step* is set to 1 second, and the *Maximum*

Physical Time is set to 10 seconds, then the simulation runs for 10 time-steps.

To set the maximum time, click the **Maximum Physical Time** node and modify the *Max Physical Time* property in the *Properties* window.

Properties

Enabled	Determines if this criterion is used for stopping the solver.	
	Activated	Use this criterion.
	Deactivated	Do not use this criterion.
Max. Physical Time	The amount of physical time before this criterion is satisfied. The default is 1 second.	
Criterion Satisfied	Indicates whether the stopping criterion has been satisfied (read-only).	
Logical Rule	Defines how this criterion interacts with other stopping criteria that have been enabled.	
	AND	Requires that this criterion, along with one or more other criteria with the *Logical Rule* set to **AND**, all be satisfied before stopping the solver.
	OR	Indicates that satisfying only this criterion is sufficient to stop the solver.

See Also

What Are Stopping Criteria?
Using Multiple Stopping Criteria
Common Pop-Up Menu Items

Setting the Maximum Number of Inner Iterations

The **Maximum Inner Iterations** stopping criterion is based on the number of inner iterations that the solver executes for transient analyses.

The node of this criterion has its own properties, and appears when the implicit unsteady model is chosen.



If the implicit unsteady solver is used, this stopping criterion can control the number of inner iterations to be run at each physical time-step. The **Maximum Inner Iterations** stopping criterion differs from other stopping criteria. This stopping criterion does not control when to stop the solver, only when to stop the inner iterations of the implicit solver and march the solution to the next time-step.

You can change the number of inner iterations by clicking the **Maximum Inner Iterations** criterion node and then modifying the *Max Inner Iterations* property.

Properties

Enabled	Determines if this criterion is used for stopping the solver.	
	Activated	Use this criterion.

	Deactivated	Do not use this criterion.
Max. Inner Iterations	The number of inner iterations to take before this criterion is satisfied. The default is 5 iterations.	
Criterion Satisfied	A read-only property that indicates when the stopping criterion has been satisfied.	
Logical Rule	Defines how this criterion interacts with other stopping criteria that have been enabled.	
	AND	Requires that this criterion, along with one or more other criteria with the *Logical Rule* set to **AND**, all be satisfied before stopping the solver.
	OR	Indicates that satisfying only this criterion is sufficient to stop the solver.

See Also

Using Solvers
What Are Stopping Criteria?
Using Multiple Stopping Criteria
Common Pop-Up Menu Items

Disabling/Enabling Stopping Criteria

This section describes how to enable and disable stopping criteria.

To deactivate an automatically generated stopping criterion, click the criterion node and deactivate the *Enabled* property in the *Properties* window.



All criteria can be activated or deactivated at the same time. To activate or deactivate all criteria, right-click the **Stopping Criteria** node and select **Enable All Criteria** or **Disable All Criteria** respectively.



Using Manually Generated Stopping Criteria

Tis section describes the **Minimum Inner Iterations** stopping criterion.

Residuals do not always reflect changes after the first inner iteration in a transient analysis. The **Minimum Inner Iterations** stopping criterion prevents a time-step from stopping until the solver finishes the specified number of inner iterations. The node of this criterion has its own properties.



Minimum Inner Iterations Properties

Minimum Inner Iterations	The number of inner iterations to take before this criterion is satisfied. The default is 1 iteration.	
Enabled	Determines whether this criterion is used for stopping the solver.	
	Activated	Use this criterion.
	Deactivated	Do not use this criterion.
Criterion Satisfied	A read-only property that indicates whether the stopping criterion has been satisfied.	
Logical Rule	Defines how this criterion interacts with other stopping criteria that have been enabled.	
	AND	Requires that this criterion, along with one or more other

		criteria with the *Logical Rule* set to **AND**, all be satisfied before stopping the solver.
	OR	Indicates that satisfying only this criterion is sufficient to stop the solver.

Adding the Minimum Inner Iterations Stopping Criterion

To make this stopping criterion available, right-click on the **Stopping Criteria** node and select **Create New Criterion > Minimum Inner Iterations**. (This option is available for transient simulations only.)



See Also

Using Solvers
What Are Stopping Criteria?
Using Multiple Stopping Criteria

Using Monitors as Stopping Criteria

It is possible to create stopping criteria that are based on existing monitors. This function lets you use more meaningful criteria to judge convergence.

The node of a monitor-based stopping criterion has properties just like the automatically generated kind.

These monitor-based stopping criteria could include a reduction in a residual, or could be based upon some physical quantity that you are interested in obtaining from the solution. For example, if you are simulating the flow over an airfoil you can stop iterating when the drag and lift coefficients of the airfoil have reached steady values.

You can create stopping criteria that are based on any existing monitor. These criteria check the value of the associated monitor and return a satisfied condition when that value reaches some user-specified minimum, maximum, asymptotic limit or standard deviation. You can also copy properties of these criteria.

Using Monitor-Based Stopping Criteria in Transient Simulations

If you are using the implicit unsteady model, these criteria apply automatically to the inner iterations. To apply the criteria to the outer iterations, so that they stop the entire simulation, tick the *Stop Outer Iterations* property of the corresponding criterion node.

In a transient analysis these criteria, when satisfied, allow only one inner iteration by default. To require more than one inner iteration:

- Manually create the Minimum Inner Iterations criterion.
- Set the properties of the **Minimum Inner Iterations** criterion:
  - Set the *Minimum Inner Iterations* property to `3` or `4`.
  - Make sure that the *Logical Rule* property is set to `And`.
- Set the *Logical Rule* property of the other monitor-based criteria to `And`.

Properties

Monitor	The name of the monitor that is used to create this stopping criterion (read-only).
Stop Inner Iterations	Stops the inner iterations of the transient simulation when the criterion (as specified in *Criterion Option*) is satisfied.

	Activated	Use this criterion.
	Deactivated	Do not use this criterion.
Criterion Option	The method to use for evaluating this criterion.	
	Minimum	Specifies that this criterion is satisfied when the monitor reaches a user-specified minimum value. A Minimum Limit node is added as a sub-node.
	Maximum	Specifies that this criterion is satisfied when the monitor reaches a user-specified maximum value. A Maximum Limit node is added as a sub-node.
	Asymptotic	Specifies that this criterion is satisfied when the monitor has stabilized to a particular range during a number of iterations. The range and iteration window are specified in the Asymptotic Limit sub-node.
	Standard Deviation	Specifies that this criterion is satisfied when the monitor reaches a user-specified standard deviation. A Standard Deviation node is added as a sub-node.
Stop Outer Iterations	With a checkbox, stops the transient simulation when the criterion (as specified in *Criterion Option*) is satisfied.	

	Activated	Uses this feature.
	Deactivated	Does not use this feature.
Enabled	Determines if this criterion is used to stop the solver.	
	Activated	Use this criterion.
	Deactivated	Do not use this criterion.
Criterion Satisfied	Indicates if the stopping criterion has been satisfied (read-only).	
Logical Rule	Defines how this criterion interacts with other stopping criteria that have been enabled.	
	AND	Requires that this criterion, along with one or more other criteria with the *Logical Rule* set to **AND**, all be satisfied before stopping the solver.
	OR	Indicates that satisfying only this criterion is sufficient to stop the solver.

See Also

What Is a Monitor?
Using Solvers
Judging Convergence
Using Manually Generated Stopping Criteria
What Are Stopping Criteria?
Using Multiple Stopping Criteria

Creating a New Monitor-Based Criterion

This section describes how to create a stopping criterion from an existing monitor.

To create a stopping criterion from an existing monitor:

1. Right-click on the **Stopping Criteria** node and select **Create New Criterion > Create from Monitor**.
2. Select from the list of existing monitors.



A new stopping criterion node is added in the **Stopping Criteria** manager node. The new node is named after the selected monitor with `Criterion` appended to the name, and possibly a digit to make the new criterion unique among others with similar names. However, renaming the stopping criterion is possible. The new criterion is activated by default, with the *Criterion Option* set to `Minimum`, and the *Logical Rule* set to `OR`.

See Also

What Is a Monitor?
Renaming Monitor-Based Stopping Criteria

Setting Up a Min/Max Monitor

By default all criteria that are created from a monitor use a minimum limit to determine when to stop the simulation. The type of limit can be changed to a maximum, an asymptote or a standard deviation.

To change the type of limit, select the monitor stopping criterion node, then from the pull-down menu under the *Criterion Option* property in the *Properties* window, select the type. You can use a monitor to create more than one criterion. For example, a monitor could be used to stop the iteration when its value reaches a maximum or an asymptote.



If you select a *Minimum* or *Maximum* limit, you can set the limiting value by selecting the **Minimum Limit** or **Maximum Limit** node in the **Stopping Criteria** manager. In the properties window, you can then modify the *Minimum Value* or *Maximum Value* property.



Minimum Limit Properties

Minimum Value	The minimum monitor value that satisfies this criterion. The solver stops as soon as the monitor value is less than this specified value.

Maximum Limit Properties

Maximum Value	The maximum monitor value that satisfies this criterion. The solver stops as soon as the monitor value is greater than this specified value.

See Also

What Is a Monitor?
Using Solvers
Common Pop-Up Menu Items

Setting Up an Asymptotic Monitor

It is also possible to set the criterion to stop once it has reached an asymptotic limit. It allows for definition of the range over which the value can fluctuate in a given number of samples.

This feature is useful when the exact limiting value is unknown but you do know the maximum change in the monitored value that you consider sufficient for convergence.

To select the asymptotic limit option, select the monitor stopping criterion node and select the **Asymptotic** option in the drop-down list under the *Criterion Option* property in the *Properties* window. The limit is set by selecting a maximum permissible absolute change in the monitored value over a given number of iterations. Expand the monitor stopping criterion node, select the **Asymptotic Limit** child node and edit the absolute change in the monitored value and the number of samples in the *Properties* window.



Properties

|Max - Min|	The difference between the smallest and largest value that is collected during the previous *Number of Samples*. If this difference is not exceeded, then this criterion is satisfied.
Number of Samples	The number of samples that the monitor values check to see if the absolute change is exceeded.

See Also

What Is a Monitor?
Common Pop-Up Menu Items

Setting Up a Standard Deviation Monitor

You can use the standard deviation for a given number of samples.

To select this option, select the monitor stopping criterion node and select the **Standard Deviation** option in the drop-down list under the *Criterion Option* property in the *Properties* window. Expand the monitor stopping criterion node, select the **Standard Deviation** child node and edit the standard deviation and the number of samples in the *Properties* window.



Properties

Standard Deviation	The standard deviation of the samples, entered as text.
Number of Samples	The number of samples that the monitor values check.

See Also

What Is a Monitor?
Common Pop-Up Menu Items

Copying Properties of Monitor-Based Stopping Criteria

You can copy a property from one limiting value node to another.

If you have two (or more) monitor-based stopping criteria and you want them to have the same limiting value, you can copy that property from one limiting value node to another. You can only copy from the same type of limiting value node (minimum, maximum or asymptotic).

Renaming Monitor-Based Stopping Criteria

You can rename a monitor-based stopping criterion.

Similar to other objects, a monitor-based stopping criterion can be renamed within STAR-CCM+.

See Also

Renaming Objects

Using Multiple Stopping Criteria

If more than one stopping criterion is activated, it is necessary to construct a logical rule to determine when the criteria stops the solver.

For example, you can stop the solver when all criteria generated from monitors have been satisfied. However, you can also stop the solver when the number of iterations exceeds a limit.

Logical rules (AND, OR) are assigned to each individual criterion and are used to determine how they interact as a group. If a criterion is assigned an OR logical rule, the solver stops when it is satisfied. The solver also stops when all the criteria assigned the AND logical rule are satisfied. By default all criteria are assigned the OR logical rule. To change the logical rule that is associated with a criterion, select the criterion node then in the *Properties* window, set the option that you require for the *Logical Rule* property.



For example, if you have four criteria named A, B, C, and D, the Logical Rule is set to OR for A and B, then C and D have the Logical Rule set to AND. In that case, the boolean expression that is tested is:

```
A OR B OR (TRUE AND C AND D)
```

The `TRUE` at the beginning of the bracketed `AND`s caters for the (unusual) situation where there is only one criterion with an `AND`.

Pop-Up Menu of the Stopping Criteria Manager

This section describes the pop-up menu items that are available from the **Stopping Criteria** manager node.

The following are the pop-up menu items that are available when you right-click the **Stopping Criteria** manager node.

Create New Criterion	Creates a new stopping criterion of the type that is selected in the submenu.	
	Create From Monitor	Creates a stopping criterion that is based on any one of the monitors available in the monitor manager.
	Minimum Inner Iterations	Creates a Minimum Inner Iterations stopping criterion (available in transient simulations only).
Enable All Criteria	Activates *Enabled* for all stopping criteria in the stopping criteria manager. Activating this property is the equivalent of visiting each stopping criterion and activating the *Enabled* property.	
Disable All Criteria	Deactivates *Enabled* for all stopping criteria in the stopping criteria manager. Deactivating this property is the equivalent of visiting each stopping criterion and deactivating the *Enabled* property.	

See Also

What Are Stopping Criteria?
What Is a Monitor?
Using Manually Generated Stopping Criteria
Common Pop-Up Menu Items

Obtaining a Solution Interactively

This section explains how to use the interactive features of STAR-CCM+ to obtain a solution.

Other STAR-CCM+ features such as scenes and plots, which are normally associated with post-processing, can help you examine results in real time while the simulation is running. These functions are described in the chapter on analyzing.

The following basic operations are discussed:

See Also

Troubleshooting the Solution

Initializing the Solution

The solution must be initialized to apply the initial conditions.

Once initial conditions have been set for each model in the simulation, the solution must be initialized, either automatically or manually, to apply the initial conditions. In general, you do not need to initialize manually as the automatic initialization is always run at the appropriate times.

What Is Solution Initialization?

This section describes solution initialization.

During solution initialization, each model that is activated in the simulation is instructed to allocate storage for its solution data. The data is then set according to the specified initial values for the model.

If storage exists and is initialized, it is not changed. To reinitialize an existing solution completely, first clear the solution.

If the mesh is changed in any way, then mesh variables such as volume are cleared and need reinitialized. This phenomenon occurs in rotor stator problems, for example, each time the rotor is stepped and a new interface intersection is calculated.

See Also

Using Initial Conditions
Clearing the Solution

Initializing Automatically

In general, solution initialization requires no explicit action on your part.

Before running or stepping, the solution is initialized automatically as required.

See Also

Running the Simulation
Stepping the Solution

Initializing Manually

If you want to visualize the initialized solution, initialize manually.

To initialize, click the (**Initialize Solution**) button in the toolbar or use the **Solution** > **Initialize Solution** menu item.



If the solution has already been initialized (manually or automatically), then this action has no effect. Executing this action after one or more iterations does not reinitialize the solution data. Any values that have been calculated for the solution variables to that point are untouched.

If you then change or add models to the simulation, an initialization sets the solution variables for the changed models.

You can use this feature with a vector displayer to display the direction of quantities on boundaries, even before you run the simulation.

See Also

Viewing the Initialized Solution
Displaying Vector Quantities
Displaying the Direction of Quantities on Boundaries

Viewing the Initialized Solution

It can sometimes be informative to view the initialized solution to understand whether the initial data make sense.

Before scalars or vectors can be displayed, or reports generated, an initialized solution must exist. If it does not, no data is displayed. The

omission of a scale on the scalar or vector bar of the scalar or vector display indicate the lack of data (rather than just a constant initial value).

	
Before Initialization	After Initialization

Clearing the Solution

The action of clearing a solution causes all storage that is associated with the solution to be cleared from the memory of the computer.

This storage includes the contents of the Residuals plot of the *Graphics* window, as well as monitor information. If the simulation is saved with the solution cleared, no solution data is stored in the simulation file.

If you want to reinitialize a solution and start again, clear the solution. An alternative approach, however, is clearing the solution for a particular region or a particular boundary.

Storage that is associated with any model is not cleared when a model is deactivated. Clearing the solution is the only way of ensuring that unnecessary storage from deactivated models is not contained within the simulation. However, in some situations you can choose not to clear the solution after changing models. For example, if it is beneficial to use the solution data from one model as the initial conditions for another model.

To clear the solution:

1. Select **Solution** > **Clear Solution...**.



The *Clear Solution* dialog appears:



The options available on this dialog are described as follows:

- **History** – This option clears the current iteration number and time (for an unsteady case), as well as all monitors and residual plots. This option is always activated.
- **Fields** – Activating this option clears all data fields from the solution. To maintain the data fields so that they remain in place as initial conditions, deactivate this option.
- **Reset Mesh** – This option is only available when you have selected a motion model. To reset all mesh vertices to their original position at time zero, activate this option. This action will also clear the storage for initial position, so that you can manipulate regions before commencing the next simulation.
- **Lagrangian Dem Solution Fields** – This option is only available when you have selected the Discrete Element Model (DEM). To clear all DEM particles from the solution, activate this option.
- **Recorded Solution History** - This option is only available when you have created a simulation history file. To clear the solution data that is stored in the simulation history file, activate this option.

2. Deactivate any of the options if necessary, and click **OK**.

See Also

Troubleshooting the Solution

Clearing a Region's Solution Data

By clearing the solution for an individual region rather than for the simulation as a whole, you retain much of the material that is produced during the simulation run.

This information includes plots, monitor information and, for a transient analysis, the physical time (which is reset to zero upon the clearing of the entire solution). You can also step the solution extra iterations in addition to the iterations of the original run.

This option is also helpful if the solution results for one region are less satisfactory than for others. If you do not need to keep solution results for all regions, it can also help to reduce simulation file size.

To clear the solution for an individual region:

1. Right-click the node for that region and select **Solution** > **Clear**.



This step leaves in place the Residuals plot of the *Graphics* window, as well as monitor information.

See Also

Stepping the Solution

What Is a Monitor?

Clearing a Boundary's Solution Data

The procedure for clearing the solution for a single boundary is similar to the procedure for a particular region.

To clear the solution for a single boundary:

1. Right-click the node for that boundary and select **Solution** > **Clear**.



Resetting for Restart

This section describes how to reset the solution data.

Earlier versions of STAR-CCM+ had a **Reset for Restart** option on the **Solution** menu. This operation is now accomplished using the checkboxes that are provided on the *Clear Solution* dialog. To reset only the iteration number, current time, and residual and monitor plots, clear the **Fields** option on the dialog before clicking OK. Field data will be retained to act as the initial condition for the next simulation.

Running the Simulation

Once a simulation has been set up and there is reasonable confidence that the setup is correct, solution iteration can commence.

To start iterating:

1. Click the (**Run**) button in the toolbar, press `<Ctrl><R>` or use the **Solution** > **Run** menu item.



   If the solution is not already initialized before iterations commence, it is automatically initialized. The iterations are displayed in the residual plot in real time.

In addition to the residual plot, STAR-CCM+ provides other features that can help you examine results while the simulation is running. These features are described in the chapter on analyzing.

The simulation runs until the specified stopping criteria are met.

See Also

Watching Residuals
Analyzing
What Are Stopping Criteria?
Troubleshooting the Solution

Stepping the Solution

You can step the solution a fixed number of iterations, or time-steps if running an unsteady simulation.

To step the solution:

1. Click the (**Step**) button in the toolbar, press `<Ctrl><T>` on the keyboard or use the **Solution** > **Step** menu item.



If the solution is not already initialized before iterations commence, it is automatically initialized. The stopping criteria are ignored.

You can also adjust the size of a step with the step control feature.

The iteration or iterations are displayed in the residual plot.

See Also

Watching Residuals
Stepping the Solution
Running the Simulation
Setting a Stop File

Using Step Control

You can set the number of iterations or time-steps taken for each step.

The number of iterations or time-steps taken for each step has a default value of 1.

To change the step size:

1. Select the **Solution** > **Step Control** menu item.



The *Step Size* dialog appears.



2. Enter the step size and click **OK**.

   Each subsequent step uses this step size.

Watching Residuals

This section explains the concept of calculating residuals, as well as working with residual monitors, and residual plots.

For using residuals to monitor a solution, see the section on judging convergence.

What Is a Residual?

The residual in each cell represents the degree to which the discretized equation is satisfied. For example, consider the transport equation for a passive scalar and its component terms (the full description can be found at Modeling Physics > Modeling Flow and Energy > Fluid Flow Formulation > Passive Scalar Formulation > Transport Equation). When the solver is run, a discretized version of this equation is solved for each cell in the mesh. The residual of this equation is:

$$r = \left| \frac{\partial}{\partial t}\int_V a_i \rho \chi \phi_j \, dV + \oint a_i \rho \chi \phi_j \left(\mathbf{v} - \mathbf{v}_g\right) \cdot d\mathbf{a} - \left(\oint \left[a_i \chi \left(\frac{\mu}{\sigma} + \frac{\mu_t}{\sigma_t} \right) \nabla \phi_j \right] \cdot d\mathbf{a} + \int_V S_{\phi_j} \, dV \right) \right| \quad (2701)$$

where:

r is the absolute error in the solution of a particular variable. In a perfectly converged solution, the residual for each variable for each cell would be equal to machine roundoff.

In STAR-CCM+, the RMS (Root Mean Squared) value of a residual for all cells is determined from:

$$R_{rms} = \sqrt{\frac{1}{n}\sum_n r^2} \quad (2702)$$

where:

n is the number of cells.

Residual monitors keep a record of this global quantity for each of the transport equations solved in the continua within the simulation. As the residuals can vary widely in absolute value, by default a normalized residual, R_{pres}, is presented. This function allows you to compare multiple residuals as all residuals have a value from 0 through 1. R_{pres} is determined from:

$$R_{pres} = \frac{R_{rms}}{R_{norm}} \quad (2703)$$

The normalization value can be set manually to any non-zero value. However, if left at its default, R_{norm}, the normalization is automatic:

$$R_{norm} = \max\left\{\left|R_1\right|, \left|R_2\right|, \ldots, \left|R_m\right|\right\} \tag{2704}$$

where:

$R_1, R_2, \ldots, R_m$ are the residuals at each iteration.

m is the Normalization Number property that selects the number of iterations to compare before holding a single normalization value. By default, m is set to 5, as this setting removes the influence of the initialization values (reflected in R_1, R_2) but provides a sensible value for observing convergence that does not change after every iteration.

In general, use normalized residuals to monitor the behavior of the solvers at each iteration to make sure that the solution is converging toward the machine precision. If the simulation is steady state, the convergence is displayed over the total number of iterations. If the solution is time-dependent, make sure that convergence occurs within each time-step. For more details, see the section on judging convergence.

Adding Residual Monitors

Residual monitors are created automatically after model selection, and cannot be deleted or renamed.



When a solver runs, a residual plot is automatically created that displays the residual monitors on semi-log axes.

Residual Monitors Properties

Active	Read-only property to indicate if the models have been changed in this simulation after a simulation was run.	
	Activated	The monitor is active in the current simulation.
	Deactivated	Any monitors specific to the previous models become deactivated unless you delete them.
Trigger	Defines what action triggers the collection of a sample. The default setting depends on whether you are using the Steady model (iterations) or Implicit Unsteady model (time-steps).	
	Iteration	A sample is taken every `N` iterations.
	Time-Step	A sample is taken every `N` time-steps.
	Inactive	No sample is taken.
Maximum Plot Samples	Limits the number of sample points that used in the monitor plot.	
Frequency	Determines how often the sample is taken using the units that are specified in *Trigger*.	
Normalization Option	Defines how the monitor data is normalized.	
	Off	Do not normalize.
	Auto	Lets STAR-CCM+ do the normalization as derived in the Automatic Normalization node that is added as a sub-node when you choose this option.
	Manual	Lets you define a specific normalization factor in the Manual

		Normalization node that is added as a sub-node when you select this option.

What Is a Residual Plot?

The residual plot is a monitor plot that is automatically created from the active residual monitors on semi-log axes when iterating starts. A sample residual plot is shown below.



By default, all active residuals are displayed in the residual plot. The residual plot can be renamed but it cannot be deleted, though individual data series can be deleted from the plot.

The residual plot uses the same properties as the monitor plot.

Adding and Removing Residuals from the Residual Plot

All monitored residuals are displayed in the residual plot by default. Specific residual monitors can be removed from the plot by right-clicking on the node or nodes in the **Monitors** node under the **Residual Plot** node and selecting **Remove**. Alternatively, you can right-click on

the **Monitors** node and select **Edit...** to access the *Customizer* dialog to add and remove monitors from the plot.

The residual plot monitors customizer allows you to add and remove specific residual monitors from the plot.



Adding or removing residual monitors does not affect the currently selected models or the stored data; only the plotting of the residuals is affected. However, these changes are not saved to the simulation file. The residual plot is reset to include the residuals from all of the selected models when you change the currently selected models.

See Also

Normalizing Manually
Normalizing Automatically
Setting Maximum Plot Samples
Creating a Monitor Plot
Running the Simulation
Monitor Plot Properties

Stopping the Solution

This section describes how to stop the solution.

Once is solution is iterating, the solution can be interrupted by clicking the (**Stop**) button in the toolbar. The solution can be resumed again by running or stepping.

If the solution is stopped during an inner iteration, the time-step can be incomplete for transient simulations. In this case, a message similar to the one below appears in the *Output* window.



As indicated, stepping the solution once completes the time-step.

Although primarily intended for batch operation, the stop file stopping criterion also aborts the job while the solution is iterating interactively.

Stopping Lagrangian Iterations

This section describes how to stop Lagrangian simulations.

In some Lagrangian simulations, particularly those using DEM models, clicking the (**Stop**) button does not immediately stop the simulation. This phenomenon is due to the large number of sub steps that are left in the iteration.

You can force an immediate stop by using the (**Substepping Stop**) button. This button appears at the bottom of the simulation window, next to a substepping progress bar, during runs of Lagrangian simulations. It becomes active after 100 substeps. To stop a Lagrangian simulation without completing that iteration or the substepping process, click the **Substepping Stop** button.



After you click this button, particle time and iteration time fall out of synchronization. For this reason, do not use the **Substepping Stop** button as a routine way to control simulations.

Steering the Solution

STAR-CCM+ allows you to monitor the progress of a solution in several ways and to change simulation parameters while running.

STAR-CCM+ contains state-of-the-art numerical methods that are both accurate and robust. However, the equations of fluid flow are highly non-linear and can be difficult to solve for situations that involve complex physics, complicated geometries and/or poor-quality meshes. Human intervention is therefore often required to obtain a solution in the most efficient manner. For this reason, STAR-CCM+ allows you to monitor the progress of a solution in several ways and to change simulation parameters while running. These topics are discussed in the following sections, together with suggestions for how to judge convergence and what action to take after the stopping criteria are met.

Monitoring Solution Progress

In addition to the residual plot, the progress of the solution while iterating can be monitored either by using monitors, or by visualizing the solution with scalar displayers and/or vector displayers.

Both monitor update and displayer update policies can be set to regular intervals of iterations or time-steps if transient.

Other STAR-CCM+ features that can help you monitor solution progress are described in the chapter on analyzing.

Although the computational expense of visualization or monitors can be small compared to the overall solution iteration process, it is not always the case. For simulations that contain a small numbers of cells, the cost of monitoring is significant. Especially if the monitoring includes a number of scalar and vector displayers as well as monitors that work over a volume (for example, maximum, minimum or volume integrals). To reduce this monitoring overhead, you can reduce the update frequency of the displayers or monitors, or reduce the quantity of displayers and monitors. However, if monitoring is important, the overhead is likely to be worthwhile.

The ability to observe the inner convergence of the AMG solution of the linear systems that are associated with solvers can be useful for obtaining the most efficient solution or for debugging the simulation.

See Also

Debugging the Simulation Setup

Observing the AMG Inner Convergence

This section describes how to monitor the residuals of the AMG linear solver.

Solvers in STAR-CCM+ that solve systems of algebraic linear equations resulting from the discretization of transport equations, rely on the algebraic multigrid (AMG) linear solver. A residual measure, described by Eqn. (2514), also indicates the solution of these linear equations. In some situations, it is beneficial to monitor these residuals. To monitor the residuals, open the **AMG Linear Solver** node under the required solver. Then in the *Properties* window, set the *Verbosity* property to `low`.



When you continue iterating, extra information appears in the *Output* window. Nested within the customary residual output window are the AMG residuals. As a rule, the magnitude of the residuals is less important than how many iterations are being executed in a single iteration. Generally speaking, it is recommended that the number is between 2 and 5.

```
Output - intake
    201   1.816461e-06 3.877587e-04 1.063521e-04 1.046618e-04 5.551720e-01
    202   1.821734e-06 3.818299e-04 1.132056e-04 1.077223e-04 5.576124e-01
Iteration Residuals
--------- ------------ ------------ ------------ ------------ ------------
        0 1.644660e-02 3.014874e+00 4.736439e-01 5.118188e-01 5.001245e+03
        1 1.384656e-03 1.640371e-01 9.262962e-02 9.498136e-02 4.203382e+02
        2 1.803850e-04 2.036892e-02 1.553720e-02 1.757699e-02 5.479753e+01
        3 3.797281e-05 4.436430e-03 3.582350e-03 4.280312e-03 1.153817e+01
    203   1.912747e-06 3.942559e-04 1.156740e-04 1.204284e-04 5.838254e-01
Iteration Residuals
```

Changing the Simulation While Running

The client-server architecture of STAR-CCM+ allows you to make minor changes to the simulation while the solution is iterating.

The client-server architecture of STAR-CCM+ allows minor changes to be made to the simulation while the solution is iterating without stopping the solution:

- Opening plots and scenes
- Changing properties
- Adding monitors or reports

Therefore, it is acceptable to adjust solver properties such as Courant number (for coupled solvers) and under-relaxation factors (for segregated solvers) while continuing to iterate. However, if there is some risk of these changes causing solution convergence problems, then stop the solution, save the simulation, implement the changes, and continue with the solution.

However, major changes to the simulation are not advisable while the solution is iterating. As a general guideline, avoid any actions that involve the deletion of simulation objects or the creation of certain new ones (such as activating a new model) during this process.

Adding Monitors or Reports While Running

During a simulation run, it is possible to create new annotations, monitors, or reports, as well as new related annotations and plots, using the following methods:

- Creating a report
  - Creating an annotation from a report

  - Creating a monitor from a report
  - Creating a monitor and a plot at once from a report
- Creating a monitor
- Creating a plot from a monitor

The STAR-CCM+ server only handles such requests at the end of an iteration cycle. Pending requests appear in the object tree using the (**Pending**) icon. At the end of the iteration, the new object becomes available for use.

When you create a monitor during a simulation run, its *Trigger* property setting is `Inactive` by default. This default setting prevents error messages. To include the new monitor in the calculations:

- Set the *Field Function* property.
- Set the *Parts* property.
- Set any other appropriate properties.
- Select one of the two active *Trigger* options, `Time Step` or `Iteration`.

See Also

Plotting Results
Interacting with the Scene
Common Editing Techniques
Adding Reports
Creating a Report Annotation
Creating Report-Based Monitors
Creating Report-Based Monitors and Plots
Creating Field Function Monitors
Troubleshooting the Solution
Getting Past a Poor Initial Solution
Gradually Activating Models

Judging Convergence

Residual monitor plots are useful for judging the convergence (or divergence) of a solution.

Residual monitor plots are useful for judging the convergence (or divergence) of a solution, and they are created automatically within every simulation. However, it is important to understand both the

significance of residuals and their limitations. While it is true that the residual quantity tends toward a small number when the solution is converged, the residual monitors cannot be relied on as the only measure of convergence. The limitations of residuals are as follows:

- The amount that a residual decreases by depends on the particulars of the simulation. Therefore, a three-order-of-magnitude drop in residuals is possibly acceptable for one simulation, but not another. The initial guess also strongly influences the amount that residuals are reduced. If the initial solution satisfies the discretized equations perfectly, the residuals do not drop at all.
- There are two types of discretization errors: dissipative errors and dispersive errors. Dissipative errors are characteristic of first-order upwind schemes; they are inherently stabilizing and produce residual plots that tend to decrease monotonically. Dispersive errors are characteristic of second-order upwind schemes which tend to "smear" solutions less than first-order schemes. While dispersive errors tend to produce residual plots that are not monotonic. This outcome is generally an acceptable price to pay for the enhanced accuracy.

  In some cases, often because of poor mesh quality, dispersive errors result in oscillating solutions (that is, changing from one iteration to the next) within a few cells. The result is that the residual plots can indicate that the solution is not "converged". You have a choice to either accept the solution, or to try to stabilize it by choosing a lower-order numerical scheme. Frequently, it is better to accept the solution.
- Residuals do not necessarily relate to quantities of engineering interest in the simulation such as integrated forces, pressure losses, or mass flow rates.

With the issues above in mind, it is advisable to monitor quantities of engineering interest, such as integrated forces, pressure changes, or mass flow rates as well as the residuals. STAR-CCM+ features such as scenes and plots can help you examine these quantities while the solution progresses. The features are described in detail in the chapter on analyzing.

The choice of the engineering quantity, as well as the convergence criterion, is your judgment call. In the example below, taken from a large external aerodynamics solution, both lift and drag coefficients are monitored as well as the residuals. It is clear that not all quantities reach an asymptotic limit at the same time. Use your judgment and decide which coefficient is the most critical.



Monitor-based stopping criteria can be used to stop the simulation when an appropriate level of convergence has been reached, based on an appropriate combination of residual monitors or monitors of other physical quantities.

See Also

Plotting Results
Watching Residuals
Working with Scenes
Analyzing
Using Monitors as Stopping Criteria
Ramping Up Solution Accuracy

Taking Action After the Stopping Criteria Are Met

If you decide that the simulation is complete after the stopping criteria have been met, save the simulation and commence post-processing.

However, you can also change the stopping criteria and continue running. For instance, if a simple stopping criterion such as number of iterations was used, the iterations can stop before acceptable convergence has been reached. In this case, the stopping criteria can be changed and the solution continued by running again.

Saving Intermediate Solutions

This section deals with different approaches to saving the simulation while still iterating toward a final solution.

Three topics are discussed:

- Saving the simulation before making steering changes
- Saving multiple checkpoints
- Automatically saving at regular intervals

Saving Before Steering Changes

It is good practice to save the simulation before making changes to solver parameters.

The section on steering the solution describes how you can change the simulation while it is running. When solver parameters that could result in solution divergence are changed, such as under-relaxation factors or Courant numbers, it is good practice to first save the simulation. This step allows the saved solution to be recovered in case divergence occurs after the change. The same can be said for other major changes to the simulation, such as models changes or boundary condition changes.

See Also

Steering the Solution
Changing the Simulation While Running

Saving Multiple Checkpoints

It can be desirable to compare solutions from two different time-steps or two different times in a transient calculation.

You can save a copy of the simulation to allow this comparison to be done at a later stage.

See Also

Saving a Copy of a Simulation

Automatically Saving at Regular Intervals

When running a long simulation (either interactively or in batch mode), it is useful to activate the auto save feature.

This option causes STAR-CCM+ to save the simulation automatically at regular intervals during the run, when the **Save**, **Save As...** and **Save All** commands are deactivated. The auto save feature is also useful for recovering from computer system crashes, or saving transient analyses at predetermined times. This section includes a description of the properties of the *Auto Save* dialog.

To set up the auto save feature:

1. Select **File** > **Auto Save...** from the menu.



The *Auto Save* dialog appears.



2. Activate the auto save feature:
   a) To expand the *Update* property, use the node handle.
   b) Activate the *Enabled* checkbox.
3. Define the number of saves that you want STAR-CCM+ to create in the *Max Autosaved Files* property.

   This property controls the number of extra files that are created. The default setting of zero instructs STAR-CCM+ to keep overwriting the same file—the same as executing the **Save** command.
4. To save the simulation at the end of the run in a batch job, tick the *Autosave Batch Runs* checkbox.

   It is not necessary to specify this action in the macro. This feature works even if the simulation diverges.
5. If you want the simulation to save automatically after generating a volume mesh, activate the *Autosave After Volume Mesh* checkbox.
6. To save the simulation at the end of the next iteration or time-step, activate the *Autosave After Next Step* checkbox.

7. To trigger the auto save operation with a file, use the *Trigger File* property.
8. Enter the type of separator that you want to use in the *Separator* property.

   By default, when STAR-CCM+ saves copies of the simulation file, it places an `@` character to separate the base file name and the number of the trigger event. If the use of this character in this manner causes any problems with your system, you can enter a different one as the separator. If you attempt to leave this field blank, STAR-CCM+ puts the `@` character there by default.
9. Select the *Trigger Type* from the drop-down menu.

   This property specifies how the auto save operation is triggered. See Trigger Types. Unless you are running a transient analysis, leave the setting at `Iteration`.
10. To control the number of intervals that pass between auto saves, expand the **Trigger Frequency** node and enter the value that you want.

    See Trigger Properties.
11. Click **Close**.

STAR-CCM+ saves automatically during the next run at the intervals you specified. Depending on the Max Autosave Files that you specified, STAR-CCM+ runs and saves a separate copy of the simulation file at each instance of the auto save operation. These copies are saved in the same directory as the original simulation file. When the specified number of files is reached, the next auto save action deletes the oldest copy of the file. The name of each copy includes the existing simulation name, along with the number of the particular interval to ensure that each file is unique.

To deactivate this feature, reopen the *Auto Save* dialog and clear the checkbox of the *Enabled* property.

Note: You can use the auto save feature with the Thermal Comfort Wizard to save the thermal properties of the occupants during a TCM simulation.

See Also

Running the Client in Batch
Generating the Volume Mesh
Trigger Types
Triggers and Update Events
What is the Thermal Comfort Wizard?
Performing a TCM Simulation

Auto Save Properties

This section describes the properties of the Auto Save feature.

You can use the following properties to activate and customize this feature:

Max Autosaved Files	Sets the number of copies of the simulation file to be created.	
Autosave Batch Runs	Saves at the end of a batch run when the simulation has not been saved since the last iteration.	
	Activated	Automatically saves simulations at the end of a batch run.
	Deactivated	Does not automatically save simulations at the end of a batch run.
Autosave After Volume Mesh	Saves when a volume mesh is generated.	
	Activated	Automatically saves simulations when a volume mesh is created.
	Deactivated	Does not automatically save simulations when a volume mesh is created.
Autosave After Next Step	Saves when the next iteration or time-step is completed. This option triggers the same behavior as the *Trigger File* property. After the auto save occurs, the checkbox is deactivated automatically to avoid triggering the auto save again.	
	Activated	Automatically saves simulations after the next step.
	Deactivated	Does not automatically save

		simulations after the next step.
Trigger File	Uses a file to control the auto save operation.	
Separator	Specifies the character that appears between the base file name and the number of the trigger event (`@` by default).	
Format Width	Width of the iteration, time, or step number in the formatted file name.	
Update	When expanded, this property makes available the controls that are listed below.	
Enabled	Activates or deactivates the auto save feature.	
	Activated	Auto save is turned on.
	Deactivated	Auto save is turned off.
Trigger Type	Provides choices for triggering the automatic save. See Trigger Types.	
Update Frequency	This property controls the number of trigger events between saves. See Trigger Properties.	

The settings are saved in the simulation database.

See Also

Saving with a Checkpoint File
Trigger Types
Triggers and Update Events

Saving with a Checkpoint File

The *Trigger File* property of the *Auto Save* dialog allows you to specify the pathname of a file (named `CHECKPOINT` by default) that, once in place, triggers the auto save operation.

This feature can help you control the auto save process when accessing the *Auto Save* dialog is inconvenient, for example in a cluster environment with the compute nodes in a separate subnet.

If STAR-CCM+ finds the trigger file, the auto save occurs at the end of the current step (iteration in a steady case, time-step in an implicit unsteady case).

This feature is designed to use a particular file once. After the auto save, STAR-CCM+ renames the trigger file by appending a tilde character (~) to the filename.

To set the path name of the file:

1. Activate the *Auto Save* dialog.
2. Modify the name and the path in the *Trigger File* property.

   You can either modify the name directly with a text entry or click the ellipses button and enter the path in the *Trigger File* dialog.



   If you specify a relative path, or no path at all, STAR-CCM+ checks for the trigger file in the same directory as the simulation file.
3. To use this feature while a simulation is running, you can simply add a file as follows:
   - On Windows, place a file with the correct name in the correct path; even an empty text file would suffice for this purpose.
   - On Linux, use the line command:

```
% touch CHECKPOINT
```

Note: This feature does not execute the save action immediately because STAR-CCM+ checks once every 60 seconds for the trigger file.

Scripting the Application

Scripting allows you to work more efficiently by automating repetitive tasks.

In many situations, a java file is not needed. If the properties in your simulation file are correctly set, many batch simulations can be run using the default operations with no java file:

```
starccm+ -batch [SIMULATION]
```

For example, reports, auto-save, auto-export, and scene hardcopies can all be specified in the simulation properties. You would need a macro to customize extensively, or to automate the setting of the relevant properties so they are the same across simulations.

This chapter describes how to create and use macros.

What Is a Java Macro?

A STAR-CCM+ macro is a Java program that is compiled and executed within the STAR-CCM+ workspace.

The Java SDK (Software Development Kit) is required for this key reason - to compile the Java macros. Unlike other macro systems, STAR-CCM+ does not record every mouse movement, button click, and keystroke, but rather it translates your actions from the workspace into meaningful Java code.

The macros that you write are standard Java so you have access to all the programming constructs of that language, such as loops and conditional constructs. In addition, the STAR-CCM+ server exposes a number of objects that you instantiate and manipulate to carry out the sequence of tasks that you require. You can activate a macro with the `Play Macro...` command.

You can write a macro from scratch, but that would require up-front knowledge about all the objects, attributes, and methods that the server exposes. It is more effective to use the workspace to record the actions you want to perform. Then you can edit the Java file using your favorite text editor to get the exact effect you require. You can then run the macro as often as you need.

The macro is just a Java program, so it does not compile if there are syntax errors. In this instance, error messages are displayed so that you have specific file names, line numbers, and issues to fix.

See Also

Editing the Macro
Running the Macro in Batch Mode

Macro Files

The macros are stored in a `.java` file that contains a set of tasks for the program to perform.

Macro files are Java language files that can be recorded as a journal of your interaction with the application. While it is most convenient to create the macro files from within a STAR-CCM+ session, they can be manually edited in a text editor. For example, they can be specialized for specific applications using the full power of the Java programming language.

The two most significant macro commands are play and record. As their names suggest, play causes an existing macro file to be executed and record causes the file to be opened for recording. Recording occurs until the stop recording command is executed. A pause option is also available to preclude certain interactions form being recorded into a given macro file while it is still open for recording.

Recording the Macro

The easiest way to create your Java macro file is to record one in STAR-CCM+ while you perform all the steps that belong in the macro file.

To record a Java macro, select the **File > Macro > Start Recording...** menu item.



Alternatively, click the (**Start Recording**) toolbar button.

This action opens a variant of the standard, operating system-specific Save dialog that is set up for creating macro files.



Navigate to the directory where you want to place this macro and give it a name (the `.java` extension is automatically added to the file name). If the purpose of the macro does not involve specific adjustments to graphical elements such as scenes and displayers, deactivate the **Include Graphics Commands** option near the top of the dialog. This option excludes graphics code from the macro.

When the macro is recording, each action you perform in STAR-CCM+ adds an associated command to the macro file. You can see which commands are being added to the macro by selecting the macro tab in the *Output* window; the tab has the same name as the macro file. If you wish to perform some steps in the simulation, but do not want them to be included in the macro, it is possible to pause the recording. When you are ready to record more steps for the macro, you can resume recording. The macro is saved when you stop recording.

It is recommended that you plan which steps to include in the Java macro before you begin recording. This planning ensures that the macro is not cluttered with superfluous commands, and is easier to edit.

See Also

Saving the Macro

Pausing and Resuming the Recording

You can pause and resume recording the macro.

To interrupt recording of the macro:

1. Select **File** > **Macro** > **Pause Recording**.



2. Alternatively, click the (**Pause Recording**) toolbar button.
   Actions that you trigger in the workspace are not added to the macro file while recording is paused. The pause icon in the menus and the toolbars turns red ( ) when macro recording is paused. This action is only available if a macro is being recorded.

To resume recording:

3. Select the same menu item or toolbar button.



4. Alternatively, click the (**Pause Recording**) toolbar button.

   The pause icon in the menus and the toolbars reverts to blue ( ) once this action has been triggered.

Excluding Visualization Code from the Macro

This section describes how to exclude graphics commands from the macro.

By default, macros include code for all adjustments to displayers, scenes, and other graphical elements in the STAR-CCM+ workspace. The purpose of a macro can even be just to add an item like a derived part. This extra code can be inconvenient if you plan to run this macro in batch, or if you want to edit the macro.

To use this feature, simply deactivate the **Include Graphics Commands** checkbox in the *Save* dialog.



This action causes the macro recording feature to ignore changes to scenes, displayers, view, animation, and part groups for the objects affected.

See Also

Editing the Macro
Editing the Macro

Saving the Macro

This section describes how to save a macro.

To finish recording and save a macro:

1. Select **File** > **Macro** > **Stop Recording**.



2. Alternatively, click the (**Stop Recording**) toolbar button.

   This action is only available if a macro recording is in progress or is paused.

The macro file that you had named in the previous step contains the Java code that maps to the actions you carried out while the recorder was active.

See Also

Recording the Macro

Editing the Macro

The macro file is a regular text file containing Java code that can be edited and customized to suit your task.

For example, if you want to change the number of iterations that the batch job runs from 1 to 1000. Open the `.java` file in your favorite text editor and look for a line that looks something like:

```
simulation_0.getSimulationIterator().step(1);
```

and change the iteration count to whatever you need:

```
simulation_0.getSimulationIterator().step(1000);
```

or better still, use the `run` function:

```
simulation_0.getSimulationIterator().run();
```

Save the file and close it. Some examples of other typical edits are:

- Loading in a different case file
- Adding loops to change properties to simulate a rotor being spun up to speed
- Varying properties as a function of time

Path names in java macros can be made relative to the java file using the `resolvePath` function (see the Java API for details). It is also possible to edit the macro so that it pauses to prompt you for input.

You can also exclude graphics code from the macro.

See Also

Excluding Visualization Code from the Macro
Excluding Visualization Code from the Macro
What Is a Java Macro?
Using the Java API

Editing the Macro to Prompt for User Input

You can design a macro to prompt you for input at a given point in the simulation process.

This feature gives you the powerful capability of automating some parts of a simulation while having real-time discretion over settings in other parts.

Such settings include:

Prompting for an Integer Entry

You can edit the macro to prompt for an integer entry.

Integer entries are used in various STAR-CCM+ features, such as stopping criteria.



For this example, a simple macro was recorded to change the *Maximum Steps* property of the corresponding stopping criterion from `1000` to `500`.



The macro, named `integer.java`, contains the following code:

```
// STAR-CCM+ macro: integer.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
public class integer extends StarMacro {
  public void execute() {
    Simulation simulation_0 =
      getActiveSimulation();
    StepStoppingCriterion stepStoppingCriterion_0 =
      ((StepStoppingCriterion)
simulation_0.getSolverStoppingCriterionManager().
         getSolverStoppingCriterion("Maximum Steps"));
    stepStoppingCriterion_0.setMaximumNumberSteps(500);
  }
}
```

This typical macro code contains the value for the stopping criterion that is entered during recording—it remains constant. However, when the macro contains code for user input, it pauses to activate a dialog in which you can specify that value.



Once you click **OK**, the macro continues.

The code of `integer.java` is modified as follows to make this work:

```
// STAR-CCM+ macro: integer.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
public class integer extends StarMacro {
  public void execute() {
    Simulation simulation_0 =
      getActiveSimulation();
    StepStoppingCriterion stepStoppingCriterion_0 =
      ((StepStoppingCriterion)
simulation_0.getSolverStoppingCriterionManager().
        getSolverStoppingCriterion("Maximum Steps"));
    int ret = promptUserForInput("Maximum Steps",
1000); /* <---- NEW CODE */
```

```
stepStoppingCriterion_0.setMaximumNumberSteps(ret); /*
<---- "ret" replaces
         * the numeric entry that was recorded
automatically */
  }
}
```

Prompting for a Double-Precision Real Number Entry

You can edit a macro to prompt for a double-precision real number.

Double-precision real numbers are often used for specifying values in STAR-CCM+ such as boundary conditions.



For this example, a simple macro was recorded to change the pressure from `-3066.82` to `-4017.39`.



The macro, named `doublePrecision.java`, contains the following code:

```
// STAR-CCM+ macro: doublePrecision.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
import star.flow.*;
public class doublePrecision extends StarMacro {
  public void execute() {
    Simulation simulation_0 =
      getActiveSimulation();
    Region region_0 =

simulation_0.getRegionManager().getRegion("Fluid");
    Boundary boundary_0 =
region_0.getBoundaryManager().getBoundary("Free_Stream")
;
    StaticPressureProfile staticPressureProfile_0 =
boundary_0.getValues().get(StaticPressureProfile.class);
    ((ConstantScalarProfileMethod)
staticPressureProfile_0.getMethod()).getQuantity().
      setValue(-4017.39);
  }
}
```

This typical macro code contains the value for the static pressure that is entered during recording—it remains constant. However, when the macro contains code for user input, it pauses to activate a dialog in which you can specify that value.



Once you click **OK**, the macro continues.

The code of `doublePrecision.java` is modified as follows to make this work:

```
// STAR-CCM+ macro: doublePrecision.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
import star.flow.*;
```

```
public class doublePrecision extends StarMacro {
  public void execute() {
    Simulation simulation_0 =
      getActiveSimulation();
    Region region_0 =

simulation_0.getRegionManager().getRegion("Fluid");
    Boundary boundary_0 =
region_0.getBoundaryManager().getBoundary("Free_Stream")
;
    StaticPressureProfile staticPressureProfile_0 =
boundary_0.getValues().get(StaticPressureProfile.class);
    double ret = promptUserForInput("Static Pressure",
-3066.82); /* <---NEW CODE */
    ((ConstantScalarProfileMethod)
staticPressureProfile_0.getMethod()).getQuantity().setVa
lue(ret); /* <---- "ret"
         * replaces the numeric entry that was recorded
automatically */
  }
}
```

Prompting for a Text Entry

You can edit a macro to prompt for a text entry.

Text entries (strings) are used for various purposes in a STAR-CCM+ simulation, such as renaming objects.



For this example, a simple macro was recorded to change the name of a continuum from `Physics 1` to `Fluid`.



The macro, named `string.java`, contains the following code:

```
// STAR-CCM+ macro: string.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
public class string extends StarMacro {
  public void execute() {
    Simulation simulation_3 =
      getActiveSimulation();
    PhysicsContinuum physicsContinuum_0 =
      ((PhysicsContinuum)
simulation_3.getContinuumManager().
          getContinuum("Physics 1"));
    physicsContinuum_0.setPresentationName("Fluid");
  }
}
```

This typical macro code contains the value for the stopping criterion that is entered during recording—it remains constant. However, when the macro contains code for user input, it pauses to activate a dialog in which you can specify that value.



Once you click **OK**, the macro continues.

The code of `string.java` is modified as follows to make this work:

```
// STAR-CCM+ macro: string.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
public class string extends StarMacro {
  public void execute() {
    Simulation simulation_3 =
      getActiveSimulation();
    PhysicsContinuum physicsContinuum_0 =
      ((PhysicsContinuum)
simulation_3.getContinuumManager().
         getContinuum("Physics 1"));
    String ret = promptUserForInput("Rename", "Physics
1"); /* <---- NEW CODE */
    physicsContinuum_0.setPresentationName(ret); /*
```

```
<---- "ret" replaces
         * the text entry that was recorded
automatically */
  }
}
```

Providing an Interactive Dialog

You can edit a macro to provide an interactive dialog.

STAR-CCM+ offers a great many visualization controls that let you customize the appearance of the displays in your simulation. Settings for these controls are captured in the recording of macros. For example, a simple macro was created to add a geometry scene to a simulation and change the *Color Mode* property of the geometry displayer to `Constant`.



The next step in editing the macro involves changing the constant color to red.



The macro, named `partcolor.java`, contains the following code:

```
// STAR-CCM+ macro: partcolor.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
import star.vis.*;
public class partcolor extends StarMacro {
  public void execute() {
    Simulation simulation_1 =
      getActiveSimulation();

simulation_1.getSceneManager().createGeometrySceneAndWai
t("Geometry Scene", "Outline", "Geometry", 1);
    Scene scene_1 =
      simulation_1.getSceneManager().getScene("Geometry
Scene 1");
    scene_1.initializeAndWait();
    PartDisplayer partDisplayer_3 =
      ((PartDisplayer)
scene_1.getDisplayerManager().getDisplayer("Outline
1"));
    partDisplayer_3.initialize();
    PartDisplayer partDisplayer_4 =
      ((PartDisplayer)
scene_1.getDisplayerManager().getDisplayer("Geometry
1"));
    partDisplayer_4.initialize();
    PartDisplayer partDisplayer_5 =
      ((PartDisplayer) scene_1.getHighlightDisplayer());
    partDisplayer_5.initialize();
    CurrentView currentView_1 =
      scene_1.getCurrentView();
    currentView_1.setInput(new DoubleVector(new
double[] {-0.5333999395370483, 1.268258864962263,
0.0}), new DoubleVector(new double[]
{-0.5333999395370483, 1.268258864962263,
14.471860805953908}), new DoubleVector(new double[]
{0.0, 1.0, 0.0}), -1.0, 0);
    partDisplayer_4.setColorMode(1);
    partDisplayer_4.setDisplayerColor(new
DoubleVector(new double[] {1.0, 0.0, 0.0}));
  }
}
```

This typical macro code contains the one color selection that is made during recording—it remains the same. However, when the macro contains code for user input, it pauses to activate a dialog in which you can make your own selection.

Once you click **OK**, the macro continues.

To make this work, the code of `partcolor.java` has extensive new and modified code, as follows:

```
// STAR-CCM+ macro: partcolor.java
package macro;
import java.util.*;
import star.common.*;
import star.base.neo.*;
import star.vis.*;
import java.io.File;
import java.awt.*;
import javax.swing.*;
import star.base.neo.NeoMacro;
import star.base.neo.NeoScriptFile;
import star.base.neo.NeoDoubleField;
import star.base.neo.NeoIntegerField;
import org.openide.DialogDescriptor;
import org.openide.NotifyDescriptor;
import org.openide.DialogDisplayer;
public class partcolor extends StarMacro {
    public boolean promptUserForInput(String title,
JComponent panel) {
        if (System.getProperty("java.awt.headless") ==
null) {
            if (title == "") {
                title = "Input values";
            }
            DialogDescriptor dd = new
DialogDescriptor(panel, title, true,
NotifyDescriptor.OK_CANCEL_OPTION,
NotifyDescriptor.OK_OPTION, null);
            return DialogDisplayer.getDefault().
                notify(dd) ==
NotifyDescriptor.OK_OPTION;
```

```
        }
        return false;
    }
    public class ColorChoice extends JPanel {
        ButtonGroup group;
        String[] labels = { "Red", "Green", "Blue",
"Black", "White", "Gray", "Yellow" };
        Color[] colors = { Color.RED, Color.GREEN,
Color.BLUE, Color.BLACK, Color.WHITE, Color.GRAY,
Color.YELLOW };
        public ColorChoice() {
            super(new GridLayout(0, 1));
            setBackground(Color.lightGray);
            setBorder(BorderFactory.
                createTitledBorder("Colors"));
            group = new ButtonGroup();
            JRadioButton option;
            for (int i = 0; i < labels.length; i++) {
                option = new JRadioButton(labels[i]);
                if(i==0) option.setSelected(true);
                option.setForeground(colors[i]);
                group.add(option);
                add(option);
            }
        }
        public Color getColor() {
            Enumeration elements = group.getElements();
            while (elements.hasMoreElements()) {
                AbstractButton button =
(AbstractButton)elements.nextElement();
                if (button.isSelected()) {
                    String btext=button.getText();
                    for (int i = 0;i < labels.length; i+
+) {
                        if(btext==labels[i]) return
colors[i];
                    }
                }
            }
            return Color.RED;
        }
    }
    public void execute() {
        Simulation simulation_3 =
          getActiveSimulation();
simulation_3.getSceneManager().createGeometrySceneAndWai
t("Geometry Scene", "Outline", "Geometry", 1);
        Scene scene_3 =
simulation_3.getSceneManager().getScene("Geometry Scene
1");
        scene_3.initializeAndWait();
        PartDisplayer partDisplayer_9 =
          ((PartDisplayer)
scene_3.getDisplayerManager().getDisplayer("Outline
1"));
```

```
        partDisplayer_9.initialize();
        PartDisplayer partDisplayer_10 =
          ((PartDisplayer)
scene_3.getDisplayerManager().getDisplayer("Geometry
1"));
        partDisplayer_10.initialize();
        PartDisplayer partDisplayer_11 =
          ((PartDisplayer)
scene_3.getHighlightDisplayer());
        partDisplayer_11.initialize();
        CurrentView currentView_3 =
          scene_3.getCurrentView();
        currentView_3.setInput(new DoubleVector(new
double[] {-0.5333999395370483, 1.268258864962263,
0.0}), new DoubleVector(new double[]
{-0.5333999395370483, 1.268258864962263,
14.471860805953908}), new DoubleVector(new double[]
{0.0, 1.0, 0.0}), -1.0, 0);
        partDisplayer_10.setColorMode(1);
        ColorChoice p = new ColorChoice();
        if (promptUserForInput("Set Color", p)) {
        Color mycolor = p.getColor();

partDisplayer_10.setDisplayerColorColor(mycolor);
        }
    }
}
```

Playing the Macro

This section outlines the ways in which you can play a macro.

A macro can be played from the workspace and stopped or, maybe more typically, can be run in batch mode. It is also possible to replay recent macros.

Running a Macro from the Workspace

This section describes how to run a macro from the workspace.

To replay a macro from the workspace, select **File** > **Macro** > **Play Macro** from the menus.



Alternatively, click the (**Play Macro**) toolbar button.

This action opens the operating specific standard file open dialog, with the **Files of Type** drop-down list is automatically set to `.java` which

allows you to select a Java macro file. Once you click **Open** in the dialog, STAR-CCM+ attempts to compile the Java code, and then run it.

In addition, you can use the command line to open the STAR-CCM+ workspace, start a new simulation, and play a macro:

```
% starccm+ -new -m macro.java
```

Progress is reported in different ways:

- In the workspace status bar:



- In the *Output* window (including any output that manual inputs generate):



If the compilation does not work, you get errors reported in a dialog and in the *Output* window. These errors must be corrected before the script can run.

Resolving Errors

This section outlines how to resolve compilation errors in a macro.

If there are any compile errors, you receive an exception error dialog:



Click **OK**. You can see the error messages reported in the *Output* window under a new tab that is labeled `Compile`.



Use this information and edit the file until there are no further errors. In the example above, there is no such class called `Rotor`.

Stopping the Macro Run

This section describes how to stop the playing of a macro.

You can stop the playing of a macro by clicking the (**Cancel Macro**) button in the workspace status bar.



STAR-CCM+ stops the operations and shows acknowledgement by displaying an information dialog.



Running the Macro in Batch Mode

You can run a macro in batch mode.

Batch mode for macros is run from the command line using the `-batch` option:

```
% starccm+ -batch cavity.java file.sim
```

This option runs STAR-CCM+ with a script (`cavity.java` in the above example). The `-batch` option ensures that no GUI is displayed. If you do not want the server process to exit once the batch file has completed, then add the `-noexit` option to the command line. For example, maybe you want to connect to it to analyze results. If you wish to monitor the progress of a batch run, an additional client can access the server process interactively.

Alternatively, you can start a new simulation and run a macro in batch mode:

```
% starccm+ -new -batch macro.java
```

You can also add command-line arguments.

See Also

What Is a Java Macro?
Connecting from Multiple Clients

Adding Arguments to a Macro in Batch Mode

You can use command-line arguments in a batch run macro.

The following command-line arguments can be used in a batch run macro:

- Pass in property definitions:
  - On Linux: `-jvmargs -Dmyprop=myvalue`
  - On Windows: `-J-Dmyprop=myvalue`
- Include `System.getProperty("myprop")` in the macro.

Running Using the Recent Macros List

STAR-CCM+ keeps a record of macros that were run previously in any simulation and makes them available in a list.

To replay a recently run macro:

1. Click the **File** > **Macro** > **Play Recent** menu item to display the list of recently-run macros.
2. Select the macro that you want to play.



The selected macro begins running immediately.

Debugging a Macro Using an IDE

You can debug a STAR-CCM+ macro using a special command-line option and an Integrated Development Environment (IDE) program.

There are several IDEs available for free download including the following:

- BlueJ: www.bluej.org
- Eclipse: www.eclipse.org
- JDeveloper: www.oracle.com/technology/products/jdev
- NetBeans: www.netbeans.org

The instructions below are for debugging a macro using NetBeans 7.1. Similar techniques can be used for other IDEs.

The process for debugging a macro involves the following steps:

1. Loading the macro in the IDE
2. Running STAR-CCM+ in debugging mode
3. Stepping through the macro

Loading the Macro in the IDE

NetBeans requires the macro to be in a directory structure `<ProjectFolder>/<SourceFolder>/<Macro>.java`, so you must organize the macro and any input files to be in the appropriate locations before starting.

1. In NetBeans, select **File > New Project**.
2. In the *New Project* wizard, select **Java Project with Existing Sources** and click **Next**.



3. In the *Name and Location* panel specify a name for your project in Project Name, provide the location of your project directory in Project Folder, and click **Next**.



4. In the *Existing Sources* panel, click the **Add Directory** button for the **Source Package Folders**, add the source directory that you created earlier and click **Next**.



5. In the *Includes* panel, ensure that the required macro is highlighted.



6. Click **Finish** to create the project.

The next step is to add the STAR-CCM+ JAR files to the project so that the IDE can interpret the STAR-CCM+ Java classes. Add the JAR files that are held in the `ext/` directory that is created as part of the installation of STAR-CCM+.

7. In the *Projects* panel in NetBeans, right-click **Libraries** and select **Add JAR/Folder**.



8. In the STAR-CCM+ installation directory, select every JAR file in the `ext/` directory.

   On Linux and Windows, this directory is located in:

   `[INSTALLDIR]/star/lib/java/platform/modules/ext/`

   The error symbols disappear from the macro.



You are now ready to start STAR-CCM+ in debugging mode.

Running STAR-CCM+ in Debugging Mode

STAR-CCM+ can pass messages directly to the Java Virtual Machine on which the client runs. You can use this function to attach NetBeans to the instance of STAR-CCM+ through a specified port. In the examples

below the port number 8765 is used but you can use any unused port number.

1. Do one of the following:

   - On Windows, in a command prompt, enter the following command on a single line:

     ```
     "<InstallationDirectory>\star\bin\starccm+.exe" -
     jvmargs "-Xdebug -
     Xrunjdwp:transport=dt_socket,server=y,suspend=n,addr
     ess=8765"
     ```

   - On Linux, in a terminal window, enter the following command on a single line:

     ```
     <InstallationDirectory>/star/bin/starccm+ -jvmargs
     "-Xdebug -
     Xrunjdwp:transport=dt_socket,server=y,suspend=n,addr
     ess=8765"
     ```

   STAR-CCM+ opens.

Now attach the IDE to STAR-CCM+ through the port:

2. In NetBeans, select **Debug** > **Attach Debugger**, specify the correct port, and click **OK**.



   You are now ready to run the macro in debugging mode.

The macro executes to the specified breakpoint. You can use the IDEs stepping tools to work through each command from this point.

Stepping Through the Macro

Connecting STAR-CCM+ to an IDE lets you execute a macro up to any line and step through individual commands to understand each. Once the IDE is attached to STAR-CCM+, you can create one or more breakpoints in the macro before running.

To step through the macro, do the following:

1. In NetBeans, place a breakpoint at the start of your macro by clicking the line number.



2. In STAR-CCM+, load or create a case as necessary.
3. In STAR-CCM+, run the macro.

Running Multiple Simulations Consecutively

You can load and run multiple simulation files consecutively using a single macro. You do not need a separate shell or DOS script.

The following sample macro looks for all the `.sim` files in a specified directory, then for each simulation file it starts a server, iterates, and saves to a new filename. Starting new simulations requires the appropriate licenses.

```
// STAR-CCM+ macro: runMultiple.java
package macro;
import java.io.*;
import star.base.neo.*;
import star.common.*;
public class runMultiple extends StarMacro {
  public class SimFileFilter implements FilenameFilter {
    public boolean accept(File dir, String name) {
      return name.endsWith(".sim");
    }
  }
  public void execute() {
    File simDir = new File("c:\\users\\john\\documents\
\testMultiple");
    Simulation sim_0 = getActiveSimulation();
    sim_0.kill();
    for (File f : simDir.listFiles(new
SimFileFilter())) {
      startAndRun(f);
    }
  }
  public void startAndRun(File f) {
    System.out.println("\n Starting "+f);
    String fileName = f.getAbsolutePath();
    Simulation sim = new Simulation(fileName);
    // Uncomment next line to clear simulation before
running
    // sim.getSolution().clearSolution();
    sim.getSimulationIterator().run();
    String newFileName = fileName.replaceAll("\\.sim","-
new.sim");
    sim.saveState(newFileName);
    sim.kill();
  }
}
```

To use the sample macro:

1. Save the macro to a Java file and name it `runMultiple.java`.

2. In the macro, provide the correct path to the working directory. That is, replace "C:\\users\\john\\documents\\testMultiple" with the path to the actual working directory. The working directory contains the simulation files to run in sequence.
3. In the working directory, issue the following command:

   Windows:

   ```
   [INSTALL_DIR]\star\bin\starccm+.exe -batch
   runMultiple.java
   ```

   Linux:

   ```
   [INSTALL_DIR]/star/bin/starccm+ -batch runMultiple.java
   ```

Playing Macros Within Macros

For greater automation of your simulation work, you can create a macro that calls other macros.

To create this kind of macro, play a macro while you record one. For example, suppose that you have a group of simulations that involve similar subtasks, such as working with the same types of field functions and reports.

Do the following:

1. Record a macro for each subtask, such as `createFieldFunctions.java` or `createReports.java`.
2. As you record the simulation macro, play the subtask macro.



Alternatively, for your subsequent simulations, you can manually edit the macro files to copy and paste the references to the macros of the subtasks.

Using the Java API

The Java API, or Application Programmer Interface, contains a comprehensive range of Java code that functions in STAR-CCM+ macros.

To access this material, select the **Help** > **Java API** menu item.



This action makes the STAR-CCM+ API resources available through your web browser.



This set of Java resources only includes Java code that applies to STAR-CCM+ macros. For a more extensive Java API, visit the Oracle Java page: http://www.oracle.com/technetwork/java/index.html.

Do not reuse any of the class names belonging to the STAR-CCM+ Java API, as it can cause conflicts at compile time.

If your browser fails to respond, change the browser path or change the default browser that STAR-CCM+ uses.

See Also

Editing the Macro

Programming in Java

This section contains links to several tutorials on Java programming.

These tutorials are especially useful if you already have some experience with another programming language, or at least some familiarity with basic programming concepts.

http://www.mindview.net/Books/TIJ/ Free Electronic Book: Thinking in Java, 3rd Edition, by Bruce Eckel. An electronic book that you can download to your local drive. It is one of the better learning tools for Java programming, though it presumes prior programming knowledge.

http://www.dickbaldwin.com/tocint.htm Java and JavaScript Programming, by Richard G. Baldwin. An excellent, comprehensive tutorial that makes the subject easy to learn.

http://www.ibiblio.org/javafaq/books/jdr/chapters/ The Java Developer's Resource by Elliotte Rusty Harold. An online version of the book by the same name. The book was developed from the same author's Brewing Java: A Tutorial, which can be found at http://www.ibiblio.org/javafaq/javatutorial.html.

http://docs.oracle.com/javase/tutorial/ The Java Tutorial by Oracle. A well-structured introduction to the Java language.

http://www.oracle.com/technetwork/java/langenv-140151.html White Paper - The Java Language Environment by Oracle. A thorough description of the Java language with a useful feature: comparison of Java with C and C++.

Working with Default Macros

The default macro is a Java macro that is automatically created at the start of each interactive session.

During the interactive session, STAR-CCM+ automatically records the Java code for every action in the default macro.



You can use the default macro to repeat the sequence of actions that you performed during the session when you play it using the Play Macro command. A copy of the default macro can be saved through the STAR-CCM+ GUI at any point in the simulation.

Default macros make it easier for you to submit cases to CD-adapco support engineers for debugging. However, when sending a macro to CD-adapco support, also send any initial simulation file or input files that the macro requires.

The path to the directory in which default macros are stored on your system is displayed in the *Output* window whenever a new session is started, as shown in the preceding screenshot.

The default macros feature is activated as standard in STAR-CCM+, but can be deactivated if desired. As a default macro is generated for each interactive session, it is good practice to empty the default macros directory occasionally to make sure that it does not become cluttered.

See Also

What Is a Java Macro?
Running a Macro from the Workspace

Disabling Default Macros

This section describes how to disable the default macros.

To stop the creation of default macros:

1. Open the *Options* dialog by selecting **Tools** > **Options** in the menu bar.

2. On the left-hand side of the dialog, select STAR-CCM+ > **Environment**, then on the right-hand side of the dialog, deactivate the *Default macros* property.
3. Finally, click **Close**.



Changes that are made to this setting only come into effect in subsequent interactive sessions. For example, when a new simulation is started, or an existing `.sim` file is opened. The setting is maintained if STAR-CCM+ is closed and reopened.

See Also

What Is a Java Macro?
Running a Macro from the Workspace

Enabling Default Macros

This section describes how to enable the creation of default macros.

To activate the creation of default macros:

1. Open the *Options* dialog by selecting **Tools** > **Options** in the menu bar.
2. On the left-hand side of the dialog, select STAR-CCM+**Environment**, then on the right-hand side of the dialog, activate the *Default macros* property.
3. Finally, click **Close**.



Default macros are activated as standard in STAR-CCM+. Changes that are made to this setting only come into effect in subsequent interactive sessions, that is when a new simulation is started or an existing `.sim` file is opened. The setting is maintained if STAR-CCM+ is closed and reopened.

See Also

What Is a Java Macro?
Running a Macro from the Workspace

Saving a Default Macro

A copy of the default macro can be saved at any point during the interactive session.

The macro contains every action that is performed in the session up to the point of saving.

To save the default macro:

1. Select **File > Macro > Save Default...**



   This action opens the standard, operating system-specific *Save* dialog.

2. Navigate to the directory where you want to save a copy of the default macro and give it a name.

   The `.java` extension is automatically added to the file name.

See Also

What Is a Java Macro?
Running a Macro from the Workspace

Emptying the Default Macros Folder

It is good practice to empty the default macros folder occasionally to ensure that it does not become large and cluttered.

All default macros are created and stored in a dedicated folder that is generated during the installation of STAR-CCM+. The location of this folder is given in the *Output* window at the start of each interactive session as shown below.



It is good practice to empty the default macros folder occasionally to ensure that it does not become large and cluttered. You can empty the default macros folder in the STAR-CCM+ GUI through the *Reset User Settings* dialog. This dialog is opened by selecting **Tools > Reset User Settings...** in the menu bar.



In the *Reset User Settings* dialog, tick the *Default Macro Files* checkbox and click **OK** to schedule the deletion of all default macros in the folder.



A message is displayed that prompts you to close and restart the program for the action to be completed.



See Also

What Is a Java Macro?
Running a Macro from the Workspace

Working with User Code

This section outlines the main concepts of working with user code. Instructions for working with user code are provided.

The **User Code** node, which has a pop-up menu, is the manager object for all user code libraries that have been saved in the simulation.



It exists even when empty to allow creation of the first user code library. The user code manager can hold any number of user code libraries up to the memory limits of the computer.

User functions and library registration functions can be coded in any language, as long as that language is able to bind like C functions or Fortran subroutines. For convenience, the C language interface and the Fortran language interface are documented separately, although conceptually they are identical. Functions in other languages should follow one of these interfaces. A list of the variables available from STAR-CCM+ and some examples of user code are also provided.

FORTRAN is not supported on Windows.

What Is User Code?

User code allows STAR-CCM+ to be customized with functions written in a compiled language such as C, C++ or Fortran.

User code takes the form of one or more user libraries that are attached to STAR-CCM+, each of which contains one or more user functions and a library registration function. When a user library has been attached, its user functions are available for use and are provided in drop-down lists for any operation that needs them. You can use these functions to specify values on a region or a boundary, initialization, and visualization.

User functions can only be accessed on a part if the data used to define the user function exists on that part. For the data to exist, it must be allocated on the specific part and must have been initialized.

You can view the status of the user code where the currently attached user libraries are shown, under the **Tools** > **User Code** node.



See Also

What Are User Libraries?
What Are User Functions?
What is a Library Registration Function?
What Are Values?
What Are Regions?
What Are Boundaries?
What Are Initial Conditions?
Visualizing the Solution

Accessing Solution Data through Derived Parts

What Are User Libraries?

User libraries are shared-object libraries (files with the platform-dependent extension `.so`, `.sl` or `.dll`) containing one or more user functions and a library registration function.

They are created by compiling and linking the user functions. Each of the user libraries that have been attached to STAR-CCM+ have a node in the simulation tree. Each of these nodes has its own properties and a pop-up menu.



The path to the user library is relative to the location of the simulation file.

A user library that is compiled on one platform is not binary compatible with another platform. For example, you are not able to compile a user library on Windows and then run the simulation on Linux. Compile user libraries for each platform on which you run the simulation.

The *Language* of the library is deduced from the library registration function. It represents the style of binding which the library functions require, not necessarily their language. For example, make C++ functions bind like C.

User Library Properties

These properties determine the link between the user library object in STAR-CCM+ and the file with your user code.

Language	The programming language of the library. STAR-CCM+ determines this language when it reads the library (read-only).
Library Name	The path to the file containing the user coding. This path is relative to the location of the simulation file.

User Library Pop-Up Menu

Reload Library	Executes the reload library action.

See Also

What Are User Functions?
What is a Library Registration Function?
Reloading A User Library
Common Pop-Up Menu Items

What Are User Functions?

User functions are functions or subroutines that are written in a compiled language such as C, C++ or Fortran.

Each function can request arguments from STAR-CCM+ and must return a result consistent with its type. The two principal categories of user functions are:

- User-coded profile
- User-coded field function

The most common use of user functions is the specification of profiles of values on regions or boundaries. Examples are initial conditions, boundary conditions, and source terms. User functions which are registered as scalar field functions or vector field functions can also be used in the same applications as field functions, such as visualization. The principal reason for choosing a user function instead of a field function is power: user functions can implement more complex functionality than field functions.

In order for STAR-CCM+ to know which arguments to pass to a user function and what type of result to expect, register the user function using a library registration function. Available user functions can then be viewed in the user library nodes in the **User Code** node under **Tools** in the object tree.



The display name for the user function, together with its *Type*, its *Args*, and their *Sizes*, is obtained from the library registration function.

What Are User-Coded Profiles?

A user-coded profile is a profile that is implemented through a user function. Either of the following are possibilities:

- A boundary profile
- A region profile

Once a user-coded profile has been compiled and its library is loaded, it is available for use in boundary and region definitions, which are shown in drop-down lists of user functions.

What Are User-Coded Field Functions?

A user-coded field function is a field function that is implemented through a user function. As with all field functions, a user-coded field function can be either of the following:

- A scalar field function
- A vector field function

Once a user-coded field function has been compiled and its library is loaded, it is available for use. This availability is the same as with any other field function that is shown in scalar or vector drop-down lists of field functions.

See Also

Variables Available to User Functions
What is a Library Registration Function?
Creating a New User Library
Accessing a User Library
Common Pop-Up Menu Items

Using User Functions in Parallel Simulations

This section provides information on using user functions in parallel simulations.

Every array which is passed into a user function is given a size which depends on the number of relevant entities of the region or boundary on that partition. Depending on the partitioning, some partitions can end up with 0 relevant entities.

Never assume that there is at least one relevant entity present on all partitions. For example, the following code causes a segmentation fault when used in a partition that does not contain any relevant entities:

```
// This will seg fault if the iter array is sized 0

int currentIteration = iter[0];
```

What is a Library Registration Function?

This section describes the library registration function.

Each user library must contain a library registration function, `uclib` (for C), or `uflib` (for Fortran). STAR-CCM+ calls this function when the user library is first attached. The purpose of calling this function is to:

- Discover which user functions are contained in the library
- Determine which arguments they need
- What type of result they return

Ensure that the library registration functions follow the C template or Fortran template, either by coding in those languages or by replicating their bindings.

Creating a New User Function

This section provides information on creating a new user function.

Ensure that the user functions follow the C template or Fortran template, either by coding in those languages or by replicating their bindings. For each user function in a library, add code to the library registration function to register it. STAR-CCM+ does not impose restrictions on how many user functions can be contained in a source file, nor on the names of the source files. Once a set of user function source files has been completed, together with a library registration function, the next stage is to create a user library from them.

Creating a New User Library

This section provides information about creating a user library on Linux or Windows.

A user library is created from its constituent user functions and a library registration function.

Creating a User Library on Linux

This section provides information on creating a user library on Linux.

On Linux, STAR-CCM+ does not require a particular compiler to be used for this process. For a selection of generic compilers, the table below gives example commands to create a library with base-name `user` from C or Fortran source files. Optional compilation flags, such as for optimization, can also be used; consult the compiler manual for more information.

Compiler	Precision	Compilation Command
GCC	mixed	`gcc -fPIC -shared *.c -o libuser.so`
GCC	double	`gcc -DDOUBLE_PRECISION -fPIC -shared *.c -o libuser.so`
GNU Fortran 77	mixed	`g77 -fPIC -shared *.f -o libuser.so`
G95	mixed	`g95 -fPIC -shared *.f -o libuser.so`

Creating a User Library on Windows

This section describes how to create a user library on Windows.

FORTRAN is not supported on Windows.

Microsoft Visual C++ 2008 and 2010 are reported to work with STAR-CCM+ user code. Link against `UserFunctions.lib` found in the

compiler subdirectory of your STAR-CCM+ installation directory, for example:

```
C:\Program Files\CD-adapco\STAR-CCM
+7.02.008\win64\intel11.1
```

It is possible to use the Microsoft Visual C++ Express editions for compiling STAR-CCM+ user code, although these editions only come with support for 32-bit libraries. Link against the `UserFunctions.lib` supplied with a 32-bit installation of STAR-CCM+ for this option to work.

To create a user library on Windows, follow the steps below:

1. Open a **Visual Studio x64 Cross Tools Command Prompt** to compile a 64-bit library (or the **Visual Studio Command Prompt** for 32 bit).
2. Compile the source code into object files using the command:

   ```
   cl /MD /D_WINDOWS -c *.c
   ```

   For `*.cpp` files, substitute `*.c` with `*.cpp`.

   For double precision code, include the compiler option `/DDOUBLE_PRECISION`.
3. Link the objects into a Dynamic Link Library (DLL) using the `UserFunctions.lib` found in the compiler subdirectory of your STAR-CCM+ installation directory:

   ```
   link -dll /out:libuser.dll *.obj "C:\Program Files
   \CD-adapco\STAR-CCM+
   7.02.008\win64\intel11.1\UserFunctions.lib"
   ```

   For a 32-bit compilation substitute `win64` with `win32`.
4. Check that you have a function that is called `uclib` exported from the DLL. STAR-CCM+ looks for this function in the DLL:

   ```
   dumpbin/exports libuser.dll
   ```

   An example of a dump output is given below:

   ```
   Microsoft (R) COFF/PE Dumper Version 10.00.30319.01
   Copyright (C) Microsoft Corporation.  All rights
   reserved.
   Dump of file libuser.dll
   File Type: DLL
     Section contains the following exports for
   libuser.dll
       00000000 characteristics
       4F38F59E time date stamp Mon Feb 13 11:35:58 2012
   ```

```
        0.00 version
           1 ordinal base
           1 number of functions
           1 number of names
    ordinal hint RVA      name
          1    0 00001240 uclib
  Summary
        1000 .data
        1000 .pdata
        1000 .rdata
        1000 .reloc
        1000 .text
```

5. Embed the manifest so that it moves with the DLL file when the library file is moved to another directory:

```
mt -manifest libuser.dll.manifest -
outputresource:libuser.dll;2
```

The manifest file can be deleted after embedding.

Accessing a User Library

This section describes how to load or reload a user library in a STAR-CCM+ simulation.

To load a new user library:

1. Right-click on the **User Code** node under the **Tools** object tree, then select **New User Library**.



   A standard *Open* dialog appears, with the filter **Files of Type:** set to compatible formats.

2. Navigate to the file and open it.



The path to the user library is relative to the location of the simulation file.

The requested library is loaded and its registration function is invoked to set the *Language* of the user functions. Each user function has a node within its parent user library folder.



If you are working in parallel, it is also possible to reset the library path with a command line.

Note: User libraries can be reloaded, for example if they have been recompiled since they were last loaded.

To reload a user library:

3. Right-click the node of the user library and select **Reload Library**.



User libraries can be deleted from an existing simulation as long as they are not in use.

To prevent a user library that is in use from being deleted, the **Delete** menu item for that user library is grayed out.

User Function Interface Reference (C)

This section introduces the interface between STAR-CCM+ and C user functions.

The following topics cover the interface between STAR-CCM+ and C user functions:

- Type definitions
- The template user function
- The template library registration function (uclib)
- Function calls in `uclib` to register the user functions (ucfunc and ucarg)
- The variables available from STAR-CCM+

The usage of this interface is illustrated in the examples.

Type Definitions (C)

The types of variables and functions that are used in C user functions are defined in the file `uclib.h`.

This file is not part of the standard STAR-CCM+ installation. Create it in the same directory where you are creating your user functions. The contents of `uclib.h` is displayed as shown here:

```
#ifndef UCLIB_H
#define UCLIB_H
#ifdef DOUBLE_PRECISION
typedef double Real;
#else
typedef float Real;
#endif
typedef double CoordReal;

#ifdef __cplusplus
extern "C" {
#endif
#if defined(WIN32) || defined(_WINDOWS) ||
defined(_WINNT)
# define USERFUNCTION_EXPORT __declspec(dllexport)
# define USERFUNCTION_IMPORT __declspec(dllimport)
#else
# define USERFUNCTION_EXPORT
# define USERFUNCTION_IMPORT
#endif
```

```
extern void USERFUNCTION_IMPORT ucarg(void *, char *,
char *, int);
extern void USERFUNCTION_IMPORT ucfunc(void *, char *,
char *);
extern void USERFUNCTION_IMPORT ucfunction(void *, char
*, char *, int, ...);

void USERFUNCTION_EXPORT uclib();
#ifdef __cplusplus
}
#endif
#endif
```

Template User Function (C)

This section provides a template that you can use when creating a user function.

The generic purpose of a user function is to populate a `result` array with values obtained by operating on `args`, as shown in the examples, using this template:

```
#include "uclib.h"
void name(result, int size, args ...);
```

`name`: the function name. This name is arbitrary: STAR-CCM+ does not impose particular function names for particular functionality. There are, however, some protected names which must be avoided: `ucarg, ucfunc, ucfunction, uclib, ufarg, uffunc, uflib`.

`result`: the array of values that the user function returns. It is declared as:

```
Real *result
```

for a scalar function, or:

```
Real (*result)[3]
```

for a vector function. Each element of the `result` array represents a cell or face, depending on the type of the user function.

`size`: the number of elements in the `result` array.

`args`: the arguments which STAR-CCM+ passes to the function, which is chosen from the list of available variables and requested through ucarg. The form of declaration for an array `arg` with elements of various types is given in the table below.

Elemental Type	Declaration
`int`	`int *arg`
`Real`	`Real *arg`
`unsigned int`	`int *arg`
`Vector<2, unsigned int>`	`int (*arg)[2]`
`Vector<3, CoordReal>`	`CoordReal (*arg)[3]`
`Vector<3, Real>`	`Real (*arg)[3]`

For example,

```
CoordReal (*centroid)[3], Real *temperature
```

would declare centroid position and temperature as arguments.

The `Real` and `CoordReal` types define the precision of floating-point variables, and are defined in the uclib.h file.

Template Library Registration Function uclib (C)

This section provides a template that you can use when creating a library registration function.

Register each user function in the library with calls to `ucfunc` and `ucarg` in `uclib`, as shown in the examples. The template `uclib` is:

```
#include "uclib.h"
void uclib()
{
/* Register user functions here */
}
```

Create the header file `uclib.h` in the same directory as `uclib.c` to declare the functions that are used in `uclib`.

User Function Name Registration ucfunc (C)

This section describes the initial registration of a C user function.

The initial registration of a C user function is made with a call in uclib to `ucfunc`:

```
ucfunc(void *func, char *type, char *name);
```

`func`: a pointer to the user function.

`type`: a character string giving the type of the user function, one of the following:

- "BoundaryProfile"
- "RegionProfile"
- "ScalarFieldFunction"
- "VectorFieldFunction"

`name`: a character string giving the name with which the function is displayed as in drop-down lists in STAR-CCM+.

The arguments that the user function requires are registered by calling ucarg.

User Function Argument Registration ucarg (C)

This section describes the user function argument registration.

Each argument of a C user function which has been initially registered with ucfunc is requested with a call in uclib to `ucarg`:

```
ucarg(void *func, char *type, char *variable, int size);
```

`func`: a pointer to the user function.

`type`: a character string giving the type of the argument, either of the following:

- `"Cell"` for a cell field
- `"Face"` for a face field

`variable`: a character string naming the requested argument, which must be a valid STAR-CCM+ variable.

`size`: the storage in bytes required for one element of the `variable` array. This size is used to ensure that the precision of the user function matches the precision of STAR-CCM+. Examples are:

- `sizeof(Real)` if `variable` is a `Real` scalar per element
- `sizeof(CoordReal[3])` if `variable` is a `CoordReal` vector per element
- `sizeof(int[2])` if `variable` is a pair of integers per element

The `Real` and `CoordReal` types define the precision of floating-point variables, and are defined in the uclib.h file.

Calls to `ucarg` must be made in the order the user function needs the arguments in.

Fortran User Interface Reference

This section introduces the interface between STAR-CCM+ and Fortran user functions.

The interface between STAR-CCM+ and Fortran user functions is covered in the topics below:

- Type definitions
- The template user function
- The template library registration function (uflib)
- Function calls in `uflib` to register the user functions (uffunc and ufarg)
- The variables available from STAR-CCM+

The usage of this interface is illustrated in the examples. Although the interface is compatible with Fortran 77, code that is presented in this documentation assumes Fortran 90 as a minimum.

Type Definitions (Fortran)

The types of variables and functions that are used in Fortran user functions are defined in the file `StarReal.f`.

This file is not part of the standard STAR-CCM+ installation. Create it in the same directory where you are creating your user functions. The contents of `StarReal.f` appear as shown here:

```
module StarRealMod
        integer, parameter :: StarInt = kind(1)
        integer, parameter :: StarReal = kind(1.0)
        integer, parameter :: CoordReal = kind(1D0)
        integer, parameter :: StarIntSize = StarInt
        integer, parameter :: StarRealSize = StarReal
        integer, parameter :: CoordRealSize = CoordReal
      end module StarRealMod
```

When working with double precision, declare `StarReal` with:

```
integer, parameter :: StarReal = kind(1D0)
```

Template User Function (Fortran)

This section provides a template to help you create a user function.

The generic purpose of a user function is to populate a `result` array with values obtained by operating on `args`, as shown in the examples, using this template:

```
subroutine name(result, size, args ...)
use StarRealMod
integer, intent(in) :: size
```

`name`: the function name. This name is arbitrary: STAR-CCM+ does not impose particular function names for particular functionality. There are some protected names however which must be avoided: `ucarg, ucfunc, ucfunction, uclib, ufarg, uffunc, uflib`.

`result`: the array of values that the user function returns. It is declared as:

```
real(StarReal), intent(out) :: result(size)
```

for a scalar function, or:

```
real(StarReal), intent(out) :: result(3,size)
```

for a vector function. Each element of the `result` array represents a cell or face, depending on the type of the user function.

`size`: the number of elements in the `result` array.

`args`: the arguments which STAR-CCM+ passes to the function are chosen from the list of available variables and requested through `ufarg`. The form of declaration for an array `arg` with elements of various types is given in the table below.

Elemental Type	**Declaration (all `intent(in)`)**
`int`	`integer arg(*)`
`Real`	`real(StarReal) arg(*)`
`unsigned int`	`integer arg(*)`
`Vector<2, unsigned int>`	`integer arg(2,*)`
`Vector<3, CoordReal>`	`real(CoordReal) arg(3,*)`
`Vector<3, Real>`	`real(StarReal) arg(3,*)`

For example,

```
real(CoordReal), intent(in) :: centroid(3,*)
real(StarReal), intent(in) :: temperature(*)
```

would declare centroid position and temperature as arguments.

The `StarReal` and `CoordReal` types define the precision of floating-point variables, and are defined in the `StarReal.f` file.

Template Library Registration Function uflib (Fortran)

This section describes how to register each user function in the library.

Register each user function in the library with calls to `uffunc` and `ufarg` in `uflib`, as shown in the examples. The template `uflib` is:

```
subroutine uflib()
      use StarRealMod
C Register user functions here
      return
      end
```

User Function Name Registration uffunc (Fortran)

This section describes the initial registration of a Fortran user function.

The initial registration of a Fortran user function is made with a call in `uflib` to `uffunc`:

```
external func
call uffunc(func, type, name)
```

`func`: a reference to the user function (called `func` in this case).

`type`: a character string giving the type of the user function, one of the following:

- 'BoundaryProfile'
- 'RegionProfile'
- 'ScalarFieldFunction'
- 'VectorFieldFunction'

`name`: a character string giving the name with which the function is shown as in drop-down lists in STAR-CCM+.

The arguments that the user function requires are registered by calling `ufarg`.

User Function Argument Registration ufarg (Fortran)

This section describes user function argument registration.

Each argument of a Fortran user function which has been initially registered with `uffunc` is requested with a call in `uflib` to `ufarg`:

```
external func
integer size
call ufarg(func, type, variable, size)
```

`func`: a reference to the user function (called `func` in this case).

`type`: a character string giving the type of the argument, either of the following:

- `'Cell'` for a cell field
- `'Face'` for a face field

`variable`: a character string naming the requested argument, which must be a valid STAR-CCM+ variable.

`size`: the storage in bytes required for one element of the `variable` array. This value is used to ensure that the precision of the user function matches the precision of STAR-CCM+. Examples are:

- `StarRealSize` if `variable` is a `real(StarReal)` scalar per element
- `3*CoordRealSize` if `variable` is a `real(CoordReal)` vector per element
- `2*StarIntSize` if `variable` is a pair of integers per element

The `StarReal` and `CoordReal` types define the precision of floating-point variables, together with StarRealSize, CoordRealSize, and StarIntSize, and are defined in the `StarReal.f` file.

Calls to `ufarg` must be made in the order the user function requires the arguments.

Variables Available to User Functions

This section introduces the variables that are available to user functions.

User functions can access variables of two principal types:

Accessing Field Functions in User Functions

This section describes how to access field functions in user functions.

Field functions can be accessed in user functions by requesting them with their names enclosed in double-quotes; the syntax depends on whether they are scalar or vector field functions. Temperature would be referred to as `"$Temperature"`, for example, and velocity as `"$$Velocity"`. Additionally:

- Field function arguments are always arrays of cell or face values, even when the fundamental data that they represent is a single value.

  For example `"$Iteration"` is a cell or face array, each element of which is the iteration number.
- Scalar field functions have `CoordReal` elemental type, even when the fundamental data that they represent has a different type.

  For example, whereas the iteration number is an integer, each element of `"$Iteration"` is the iteration number in `CoordReal` precision.
- Similarly, vector field functions have `Vector<3, CoordReal>` type, even when the fundamental data that they represent has a different type.
- The report can be accessed as a field function.

  For example, the **Cell Count** report spawns a **CellCountReport** field function, which can be requested in a user function as `"$CellCountReport"`.

Accessing Stored Variables in User Functions

This section describes how to access stored variables in user functions.

Some fundamental variables in STAR-CCM+ are stored, and can be accessed by bypassing the field function mechanism. You can access stored variables for efficiency, or to access variables which are not

available as field functions. The table below is a partial list of stored variables. Additionally:

- Each variable has a distinct elemental type.
- Stored variables can be distinguished from field functions by their names.

  For example, "$Temperature" (CoordReal) is the temperature field function, whereas "Temperature" (Real) is the stored temperature.
- The string that is used to access stored variables for individual phases in a Eulerian multiphase simulation takes the form Phase<Index - 1>::<Variable>.

  For example, the temperature of phase 1 is accessed as Phase0::Temperature. The phase index is listed as a property of each phase node.
- The availability of a variable can depend on the activation of a relevant model (for example, UU_Stress needs a Reynolds stress turbulence model).
- Some variables are only available as Face or Cell fields (for example, Volume is a Cell field, Area is a Face field).

Variable	Elemental Type
"Area"	Vector<3, Real>
"BoundaryHeatFlux"	Real
"Centroid"	Vector<3, CoordReal>
"Density"	Real
"FaceCellIndex"	Vector<2, unsigned int>
"Pressure"	CoordReal
"ProstarCellIndex"	int
"ProstarCellType"	int
"Temperature"	Real
"Phase0::Temperature"	Real
"TotalEnthalpy"	Real
"TurbulentDissipationRate"	Real
"TurbulentKineticEnergy"	Real
"TurbulentViscosity"	Real
"UU_Stress"	Real
"UV_Stress"	Real

Variable	Elemental Type
“UW_Stress“	Real
“U_Velocity“	Real
“VV_Stress“	Real
“VW_Stress“	Real
“V_Velocity“	Real
“Volume“	Real
“WW_Stress“	Real
“W_Velocity“	Real
“WallShear “	Vector<3, Real>
“WallVector“	Vector<3, Real>

User Code Examples

This section provides some examples of user functions.

The examples include:

- Defining a user-coded boundary profile
- Defining a user-coded region profile
- Defining a user-coded field function
- Registering user functions with STAR-CCM+
- Requesting and using a Cell field argument
- Requesting and using a Face field argument

The examples form a single library and are presented as both C and Fortran 90 coding.

A User-Coded Boundary Profile

This section contains an example of a user-coded boundary profile.

Suppose that you wish to specify the temperature on a boundary as the same temperature that is in the neighboring cell. You could create a user function that is called `zeroGradT` which takes the temperature and face-cell index fields as arguments and returns the boundary temperature.

In C, the following could be coded in a file `zeroGradT.c`, using the uclib.h file:

```
  #include "uclib.h"

  /* Set boundary temperature equal to cell temperature
*/
  void
  USERFUNCTION_EXPORT zeroGradT(Real *result, int size,
int (*fc)[2], Real *T)
  {
    int i;

  /* Loop through all entities applying T_boundary =
T_cell *
   * fc[i][0] is the cell next to
i                          */
    for (i = 0; i != size; ++i)
    {
      result[i] = T[fc[i][0]];
```

```
    }
  }
```

The equivalent in Fortran 90 could be a file `zeroGradT.f`, using the StarReal.f file:

```
C Set boundary temperature equal to cell temperature
      subroutine zeroGradT(result,size,fc,T)
      use StarRealMod
      implicit none
      integer, intent(in) :: size
      real(StarReal), intent(out) :: result(size)
      integer, intent(in) :: fc(2,*)
      real(StarReal), intent(in) :: T(*)
      integer i
C Loop through all entities applying T_boundary = T_cell
C fc(1,i) is the cell next to i
      do i = 1,size
        result(i) = T(fc(1,i))
      end do

      return
      end
```

This user function can then be registered with STAR-CCM+.

A User-Coded Region Profile

This section contains an example of a user-coded region profile.

Suppose that you wish to initialize the velocity to a constant swirl condition. You could create a user function that is called `initVelocity` which takes the centroids as an argument and returns the velocity vector at those centroid locations.

In C, the following could be coded in a file `initVelocity.c`, using the uclib.h file:

```
  #include "uclib.h"

  /* Initial velocity based on uniform swirl */
  void USERFUNCTION_EXPORT
  initVelocity(Real (*result)[3], int size, CoordReal
(*centroid)[3])
  {
  /* Angular velocity and origin of rotation */
    CoordReal omega[3] = {0.0, 0.0, 100.0};
    CoordReal origin[3] = {0.0, 0.0, 0.0};
    CoordReal dr[3];
```

```
    int i;

  /* Loop through all entities applying u = omega x
(centroid - origin) */
    for (i = 0; i != size; ++i)
    {
      dr[0] = centroid[i][0] - origin[0];
      dr[1] = centroid[i][1] - origin[1];
      dr[2] = centroid[i][2] - origin[2];
      result[i][0] = (Real)(omega[1]*dr[2] -
omega[2]*dr[1]);
      result[i][1] = (Real)(omega[2]*dr[0] -
omega[0]*dr[2]);
      result[i][2] = (Real)(omega[0]*dr[1] -
omega[1]*dr[0]);
    }
  }
```

The equivalent in Fortran 90 could be a file `initVelocity.f`, using the StarReal.f file:

```
C Initial velocity based on uniform swirl
      subroutine initVelocity(result,size,centroid)
      use StarRealMod
      implicit none
      integer, intent(in) :: size
      real(StarReal), intent(out) :: result(3,size)
      real(CoordReal), intent(in) :: centroid(3,*)
      integer i
      real(CoordReal) dr(3)
C Angular velocity and origin of rotation
      real(CoordReal), parameter :: omega(3) = (/
0.0,0.0,100.0/)
      real(CoordReal), parameter :: origin(3) = (/
0.0,0.0,0.0/)
C Loop through all entities applying u = omega x
(centroid - origin)
      do i = 1,size
        dr(1) = centroid(1,i) - origin(1)
        dr(2) = centroid(2,i) - origin(2)
        dr(3) = centroid(3,i) - origin(3)
        result(1,i) = omega(2)*dr(3) - omega(3)*dr(2)
        result(2,i) = omega(3)*dr(1) - omega(1)*dr(3)
        result(3,i) = omega(1)*dr(2) - omega(2)*dr(1)
      end do

      return
      end
```

This user function can then be registered with STAR-CCM+.

A User-Coded Field Function

This section provides an example of a user-coded field function.

Suppose that you would like to define the Sutherland viscosity law as a field function. Defining this field function means that you can visualize the Sutherland viscosity for a given temperature field without having to apply it as the dynamic viscosity in your region of interest. The law is simple enough to be implemented as a user field function, but here you can see how it can be implemented as a user function. The function could be called `sutherlandViscosity`, taking the temperature as an argument and returning the viscosity.

In C, the following could be coded in a file `sutherlandViscosity.c`, using the uclib.h file:

```
  #include <math.h>
  #include "uclib.h"
  /* Dynamic viscosity based on Sutherland's law */
  void USERFUNCTION_EXPORT
  sutherlandViscosity(Real *result, int size, Real *T)
  {
  /* Reference viscosity, Sutherland constant and
reference temperature */
    Real v0 = 1.716E-5;
    Real Cs = 110.0;
    Real T0 = 273.15;
    int i;
  /* Loop through all entities applying Sutherland's
law */
    for (i = 0; i != size; ++i)
    {
      result[i] = v0 * pow(T[i]/T0, 1.5) * (T0 + Cs)/
(T[i] + Cs);
    }
  }
```

The equivalent in Fortran 90 could be a file `sutherlandViscosity.f`, using the StarReal.f file:

```
C Dynamic viscosity based on Sutherland's law
      subroutine sutherlandViscosity(result,size,T)
      use StarRealMod
      implicit none
      integer, intent(in) :: size
      real(StarReal), intent(out) :: result(size)
      real(StarReal), intent(in) :: T(*)
      integer i
C Reference viscosity, Sutherland constant and
reference temperature
```

```
      real(StarReal), parameter :: v0 = 1.716E-5
      real(StarReal), parameter :: Cs = 110.0
      real(StarReal), parameter :: T0 = 273.15

C Loop through all entities applying Sutherland's law
      do i = 1,size
        result(i) = v0 * (T(i)/T0)**1.5 * (T0 + Cs)/
(T(i) + Cs)
      end do

      return
      end
```

This user function can then be registered with STAR-CCM+.

Defining a Library Registration Function

This section provides an example of a library registration function.

The example functions zeroGradT, initVelocity, and sutherlandViscosity must be registered with STAR-CCM+ through the library registration function.

In C, this function could be achieved by coding the following in `uclib.c`, using the uclib.h file:

```
  #include "uclib.h"

  void zeroGradT(Real*, int, int*, Real*);
  void initVelocity(Real*, int, CoordReal*);
  void sutherlandViscosity(Real*, int, Real*);

  void
  USERFUNCTION_EXPORT uclib()
  {
    /* Register user functions here */
    ucfunc(zeroGradT, "BoundaryProfile", "Zero Gradient
Temperature");
    ucarg (zeroGradT, "Face", "FaceCellIndex",
sizeof(int[2]));
    ucarg (zeroGradT, "Cell", "Temperature",
sizeof(Real));

    ucfunc(initVelocity, "RegionProfile", "Initial
Velocity");
    ucarg(initVelocity, "Cell", "Centroid",
sizeof(CoordReal[3]));

    ucfunc(sutherlandViscosity, "ScalarFieldFunction",
"Sutherland Viscosity");
    ucarg(sutherlandViscosity, "Cell", "Temperature",
```

```
sizeof(Real));
  }
```

The equivalent library registration function in Fortran 90 could be coded in `uflib.f`, using the StarReal.f file:

```
      subroutine uflib()
      use StarRealMod
      implicit none
C Register user functions here
      external
zeroGradT,initVelocity,sutherlandViscosity
      call uffunc(zeroGradT, "BoundaryProfile",
     &            "Zero Gradient Temperature")
      call ufarg (zeroGradT, "Face",
     &            "FaceCellIndex", 2*StarIntSize)
      call ufarg (zeroGradT, "Cell",
     &            "Temperature", StarRealSize)

      call uffunc(initVelocity, "RegionProfile",
     &            "Initial Velocity")
      call ufarg(initVelocity, "Cell",
     &           "Centroid", 3*CoordRealSize)

      call uffunc(sutherlandViscosity,
"ScalarFieldFunction",
     &            "Sutherland Viscosity")
      call ufarg(sutherlandViscosity, "Cell",
     &           "Temperature", StarRealSize)
      return
      end
```

Using User Code on Particles and Parcels

This section gives examples of cases where user coding applies to particles and parcels in multiphase modeling.

The first example covers the topic in detail, while the subsequent examples build on the first. The examples are listed below:

Changing the Distribution of Diameters in a Transient Injection

This example examines the case of particles such as droplets being injected into the simulation.

STAR-CCM+ uses the Rosin-Rammler distribution to describe droplet diameters. The probability of finding a particle with diameter D in a given set of particles (the probability density function, or PDF) is:

$$PDF\left(D\right) = \exp\left[\left(\frac{-D}{D_{ref}}\right)^{q}\right] \tag{2705}$$

where D_{ref} is the Rosin-Rammler diameter and q is the Rosin-Rammler exponent.

On a log-log graph, this distribution is a straight line; D_{ref} is the intercept and q is the slope. The cumulative probability density function (or CDF) is the integrated probability:

$$Y\left(D\right) = 1 - \exp\left[\left(\frac{-D}{D_{ref}}\right)^{q}\right] \tag{2706}$$

where $Y(D)$ is the mass fraction of droplets with diameter less than D and is bounded between zero and one. This relation can be solved for D to obtain a sample of diameters at different mass fraction points:

$$D = D_{ref}\left[-\ln\left(1-Y\right)\right]^{\frac{1}{q}} \tag{2707}$$

This relation can be used to return diameters for a given random sample of Y in the range [0,1]. Writing a routine to return the diameter is straightforward in Fortran or C, but you must provide some additional coding to link the routine into STAR-CCM+.

This example requires four files:

- A file that defines the size of variables in a module for a shared Fortran object library, StarReal.f
- A file containing definitions that link the variables inside STAR-CCM+ to the external library created for the case, uflib.f
- Source code, returning the appropriate diameter, rosinRammlerDiameter.f
- A makefile, for convenience

Module File StarReal.f

The first file, StarReal.f, defines a module which denotes the standard variable sizes that are used in STAR-CCM+. Compiling StarReal.f creates the Fortran shared object library. The file is:

```
module StarRealMod
     integer, parameter :: StarInt = kind(1)
     integer, parameter :: StarReal = kind(1.0)
     integer, parameter :: CoordReal = kind(1D0)
     integer, parameter :: StarIntSize = StarInt
     integer, parameter :: StarRealSize = StarReal
     integer, parameter :: CoordRealSize = CoordReal
 end module StarRealMod
```

Note: When using double precision, define `StarReal` as `kind(1.0d0)`.

Library File uflib.f

The library file, uflib.f, registers the user routine and defines the internal variables to be passed into your code. The file is:

```
subroutine uflib()
     use StarRealMod
     implicit none

  c Register user functions here

     external rosinRammlerDiameter

     call uffunc(rosinRammlerDiameter,
     &    "ParcelProfile",
     &    "Rosin-Rammler Diameter")

     return

  end
```

In this case, the internal routine uffunc looks for a module named rosinRammerDiameter, which is being applied internally to a ParcelProfile.

The third argument in the list, `Rosin-Rammler Diameter`, is the title that appears in the interface when the library is loaded.

Subroutine rosinRammlerDiameter.f

This subroutine calculates the diameters by sampling the CDF uniformly. The example here uses the built-in Fortran random number generator `rand()`, which means that this implementation depends on the machine and compiler combination. Other suitable pseudorandom number generators are available from several sources.

```
subroutine rosinRammlerDiameter(result,size)

  c particle diameter example

     use StarRealMod

     implicit none
     integer, intent(in) :: size
     real(StarReal), intent(out) :: result(size)

     real(StarReal), parameter :: meandiam = 1e-5
     real(StarReal), parameter :: rrexpon = 1.0/3.5

     integer i

     do i = 1,size
      result(i) = meandiam*(-log(1.0-rand()))**rrexpon
     end do

     return
     end
```

The makefile

Below is an example of a makefile that contains the compilation flags for gnu Fortran95:

```
default: libuser.so

  libuser.so: StarReal.f uflib.f rosinRammlerDiameter.f
          f95 StarReal.f uflib.f rosinRammlerDiameter.f
-fPIC -shared -o $@
```

Running the makefile produces a shared object library libuser.so. After building the library, you connect the function to STAR-CCM+.

Connecting User Code to STAR-CCM+

To connect the user code:

1. Right-click the **Tools** > **User Code** node and select **New User Library...**.
2. Browse for libuser.so in the file system and expand the **User Code** node:



Assuming an appropriate injection has been defined for the case, the **User Code** option appears as an option for the **Particle Diameter** definition.



3. To select the Rosin-Rammler distribution, edit the *Function* property of the **User Code** node.

The case that is shown here defines two injectors. In one, the particle diameter is given by the Rosin-Rammler particle size distribution as defined internally by STAR-CCM+. In the other, the user coding defines the Rosin-Rammler distribution. The injectors are otherwise the same.

The following graph is the histogram of particle diameter for the internally defined Rosin-Rammler diameter distribution:



The following graph is the histogram of particle diameter for the user-coded Rosin-Rammler diameter distribution:



There are slight differences in mean diameter and actual number distribution for the two definitions. These differences are due mostly to the choice of sampling for the random numbers and clipping of maximum and minimum values between the methods. In both cases, the Rosin-Rammler diameter D_{ref} was 1.0E-5 meters and the Rosin-Rammler exponent q was set to 3.5. The differences become smaller for a larger number of samples.

Changing the Mass Flow Rate in a Transient Injection

This section provides an example of changing the mass flow rate in a transient injection.

By default, each parcel in STAR-CCM+ represents the same amount of mass. The mass of a parcel is the product of the mass of an individual particle $\left(\rho\left(\frac{4}{3}\right)\pi r^3\right)$ and the parcel count. The parcel count is the number of particles contained in the parcel. The parcel count is a floating point number and can also be considered the probability of finding a particle with that mass and diameter at that spatial location.

It is straightforward to change the amount of mass in each parcel as a function of diameter with user code. In this example, the mass flow rate of each parcel increases linearly with diameter.

To pass an argument to a subroutine, the argument must be registered with a call to the library routine ufarg():

```
subroutine uflib()
use StarRealMod
```

```
      implicit none

  c Register user functions

      external linearDiameter
      external parcelMdot

      call uffunc(linearDiameter,
      &     "ParcelProfile",
      &     "Linear Diameter")

      call uffunc(parcelMdot,
      &     "ParcelProfile",
      &     "Parcel Mass Flow")

      call ufarg(parcelMdot,
      &     "Parcel",
      &     "$ParticleDiameter",
      & CoordRealSize)

      return
      end
```

To set a linear diameter distribution, modify the rosinRammlerDiameter.f routine by replacing the line that sets the result. In this example, the minimum diameter is set to 1 micron and the maximum is set to 100 microns. These diameters can be set as parameters. The linear distribution can be coded as:

```
      delta = (maxdiam-mindiam)/(size-1)
      do i = 1, size
        result(i) = mindiam+(i-1)*delta
      enddo
```

As the example routine below sets the mass flow rate for each parcel in the injection, you must set the mass flow rate condition to Per Parcel Stream. This setting is valid only if the injector is on a single partition—it does not work in parallel if the partition cuts through the injector.

```
      subroutine parcelMdot(result,size,D)

      use StarRealMod

      implicit none
      integer, intent(in) :: size
      real(StarReal), intent(out) :: result(size)
      real(CoordReal), intent(in) :: D(*)

  c parcel mass is count*rho*volume, so for an exponent
of
  c -3 parcel mass is constant (default behavior)
  c -2 parcel mass increases with diameter linearly
```

```
c -1 parcel mass increases with diameter quadratically
c etc.

    real(StarReal), parameter :: expon = -2.0

    real(StarReal) sum, ratio, mdot
    integer i
c a constant mass flow rate for all particles
    mdot = 1.0e-3*size
c set the Parcel Count
    sum = 0.0
    do i = 1, size
     result(i) = D(i)**expon
     sum = sum + result(i)*D(i)**3
    end do

c scale for flow rate per parcel
    if (sum > 0.0) then
     ratio = mdot/sum
     do i = 1, size
       result(i) = ratio*result(i)*D(i)**3
     enddo
   endif

   return
   end
```

An XY plot of the parcel mass as a function of particle diameter shows the default behavior of STAR-CCM+ (Phase 1) and the results of the user code (Phase 2) with the parcel mass increasing linearly with diameter.



Defining a User-Coded Energy Source in a Parcel

In simple cases, internal heating of particles can be modeled using field functions, but more complex cases can be modeled with user coding. This example compares the two methods for a transient simulation injecting aluminum particles into a square box.

A constant energy source of 1.0E8 W/m^3 heats the particles during the 0.05 seconds of simulation time, raising their temperature from 300 to 301.12 degrees. Since the heat source is constant, the temperature increases linearly with residence time.

Parts of the parcel profile can become invalid due to other physical processes occurring during a substep of the Lagrangian solution procedure. The user code must check the `parcelId` for validity before operating on the profile. The profile must have the `parcelId` passed in as an argument in addition to the arguments necessary to calculate the heat source.

The subroutine uflib.f registers the `parcelId` and the temperature:

```
subroutine uflib()
     use StarRealMod
     implicit none

 c Register user functions

    external parcelQdot

    call uffunc(parcelQdot,
     &    "ParcelProfile",
     &    "Parcel Heat Source")

    call ufarg(parcelQdot,
     & "Parcel",
     & "ParcelId",
     & StarInt)

    return
     end
```

The user code that applies the heat source to the parcels in the domain is:

```
subroutine parcelQdot(result,size,Id)

 c parcel heat source example

    use StarRealMod

    integer, intent(in) :: size
     real(StarReal), intent(out) :: result(size)
```

```
     real(StarInt), intent(in) :: Id(*)

    real(StarInt), parameter :: NULL_INDEX = -1
     integer i

    do i = 1, size
       if (Id(i) .gt. NULL_INDEX) then
         result(i) = 1e8
       endif
     end do

    return
     end
```

When compared to the internal user-defined heat source, the field functions and the user code above have identical results:



In the plot above, Phase 1 corresponds to the internally calculated parcel heat flux and Phase 2 corresponds to the user-coded parcel heat flux. Since the particle times are staggered slightly during the injection process for this case, the residence times differ slightly. In both cases, however the particle temperature is linearly proportional to residence time.

Defining a User-Coded Body Force on a Parcel Cloud

Defining a vector body force on a parcel cloud provides considerable flexibility in modeling phenomena acting on a parcel cloud. This example demonstrates such a body force, simulated first with field functions and then with user code.

In this simple case, the user code is longer and more difficult than the field functions, but in more elaborate situations, user code can be easier to write and debug.

In both methods:

- The body force on the parcel cloud is inversely proportional to the cube of the distance from the origin to the parcel position.
- The minimum radial distance is clipped to 1 cm to eliminate the singularity.
- Particles are injected for 10 ms from the top surface with a velocity of -10 m/s.
- On the time-step after injection, the body force begins to act, attracting the particles to the origin at the center of the domain.
- After 250 ms, the parcels are (unphysically) gathered at the origin.

Both the field function method and the user code method produce the same results.



For both field functions and user code, the injection mass flow rate is set to a field function that turns the injection off after 10 ms (that is, using a field function of: `($Time < 0.01) ? 0.001 : 0.0)` in the injection definition.

Field Function Method

Two field functions are defined to describe the body force:

- **rCubed**, defined as:

```
max(0.01,pow($$ParcelCentroid.mag(),3))
```

- **UserParticleBodyForce**, defined as:

```
($Time > 0.01) ? -9.81e5*($$ParcelCentroid)/
$rCubed : [0,0,0]
```

User Code Method

The user code defining the same body force is:

```
subroutine parcelBodyForce(result,size,centroid,Id,Time)
 c parcel body force example
   use StarRealMod
   integer, intent(in) :: size
   real(StarReal), intent(out), dimension(3,size) ::
result
   real(CoordReal), intent(in), dimension(3,*) ::
centroid
   real(StarInt), intent(in), dimension(*) :: Id
   real(CoordReal), intent(in) :: Time
 c body force is proportional to 1/r^2 after 0.01
seconds
   real(StarReal), parameter:: aScale = -9.81e5
   real(StarReal), parameter:: center(3) = (/0.0, 0.0,
0.0/)
   real(StarInt), parameter :: NULL_INDEX = -1
   real(StarReal) radius(3), rmag
   integer i
 c set the magnitude and direction for valid parcels
   do i = 1, size
     if ((Id(i) .gt. NULL_INDEX) .and. (Time > 0.01))
then
       radius(:) = centroid(:,i)-center(:)
       rmag = max(0.01,sqrt(sum(radius**2))**3)
       result(:,i) = aScale*(radius)/rmag
     endif
   end do
   return
   end
```

The user code specifies a vector force on the parcel cloud, returned in the result vector. The parcel centroid is passed into the user code as a vector field function with a size of `3*CoordReal`. The uflib.f function that registers this function and argument list is:

```
subroutine uflib()
    use StarRealMod
 c Register user functions
    external parcelBodyForce
    call uffunc(parcelBodyForce,
    &    "ParcelProfile",
    &    "Parcel Body Force")
    call ufarg(parcelBodyForce,
    &    "Parcel",
    &    "$$ParcelCentroid",
    &    3*CoordRealSize)
    call ufarg(parcelBodyForce,
    &    "Parcel",
    &    "ParcelId",
    &    StarInt)
    call ufarg(parcelBodyForce,
```

```
    &    "Parcel",
    &    "$Time",
    &    CoordRealSize)
    return
    end
```

In this case, both the ParcelCentroid and Time field functions are required in the routine.

Note: The registration calls that register the arguments in STAR-CCM+ must be performed in the same order as the arguments appear in the parcelBodyForce function.

Note: These examples are not coupled with the flow field.

Pop-Up Menu of the User Code Manager

This pop-up menu item is available when you right-click the **User Code** manager node.

New User Library	Executes the create user library action and adds a user code library to the **User Code** node.

Working with the Plugins Manager

The Plugins Manager makes it possible to install and manage NetBeans modules for STAR-CCM+.

The Plugins Manager is opened when you select **Tools** > **Plugins** from the top menus.



The NetBeans modules (`.nbm` files) can be thought of as plugins that permit you to extend the STAR-CCM+ GUI dynamically. Modules range in complexity from a single Java class that executes a single

operation to a group of classes to control the generation of application-specific models.

The Macro Hut site contains examples of NetBeans modules. To learn how to create your own modules, contact CD-adapco support.

This section provides information on the following:

- Installing a NetBeans Module file
- Deactivating an Installed Plugin

Installing a NetBeans Module File

To install a plugin, it is necessary to modify the STAR-CCM+ installation as follows:

1. Choose **Tools > Plugins** from the main menu and select the *Settings* tab.
2. If you want to make the plugin available for all users, activate *Force install into shared directories*, otherwise deactivate this function.



3. Select the *Downloaded* tab.
4. Click **Add Plugins** and browse to the location of your created or downloaded .nbm file.



5. Click **Install** and complete the installation.



6. Restart STAR-CCM+ if prompted to do so.

   When the installation is complete, the plugin appears in the *Installed* tab as activated.


Plugins
Updates
Available Plugins
Downloaded
Installed (3)
Settings
Search:
Select
Name
Category
Active
RCP Platform
RCP Platform
STAR-CCM+
STAR-CCM+
2D Profile Wizard
STAR-CCM+
2D Profile Wizard
Version: 1.0.0.3
Source: twoDWizard.nbm
Plugin Description
Adds the 2D Profile Simulation
Activate
Deactivate
Uninstall
Close
Help

Deactivating an Installed Plugin

STAR-CCM+ allows you to deactivate installed plugins that you do not need to minimize startup time and save memory.

Note: Do not deactivate the STAR-CCM+ plugin.

To deactivate a plugin:

1. Choose **Tools** > **Plugins** from the main menu and select the *Installed* tab.
2. Select the plugin.
3. Click **Deactivate**.



4. Restart STAR-CCM+ if prompted to do so. If you require it later, the plugin can be reactivated using the **Activate** button.

Uninstalling a plugin follows a similar workflow to deactivating with the exceptions that you click the **Uninstall** button and the plugin cannot be reactivated. The STAR-CCM+ plugin cannot be uninstalled.

Troubleshooting the Solution

This section provides information on troubleshooting the solution

Divergence of the solution is the situation where cell residuals increase from one iteration to another. If left, it can eventually result in a floating point exception being reported. Generally, the behavior of the residual monitor plots indicates when divergence happens.

An increase in one or more of the residual monitors during the first few iterations is normal and acceptable. For example, consider that a (constant) initial solution often satisfies the discretized equations exactly, everywhere except at the boundaries. The boundary conditions propagate into the interior as the solution iterates, causing the global residuals to increase for several iterations before decreasing. However, if this situation does not change after 5-10 iterations, it is likely to be an indication of divergence.

Solution divergence could be due to one of two reasons:

- In a case with difficult physics, the initial conditions specify that the initial solution is inappropriate and special steps must be taken to get past a poor initial solution.
- There is some other fundamental issue with the simulation, such as the mesh quality, the boundary conditions or the solution parameters, and special steps must be taken to debug the simulation setup.

If divergence occurs, you do not need to exit from the simulation. Instead, you can clear the solution and run it again after investigating the cause of divergence and making the required changes to the simulation.

See Also

Obtaining a Solution Interactively

Using Initial Conditions
Clearing the Solution
Running the Simulation
Changing the Simulation While Running

Getting Past a Poor Initial Solution

This section provides information on dealing with divergence that is related to a poor initial solution.

A possible cause of divergence can be related to the initial solution being radically different from a converged solution. This situation occurs because implicit numerical methods rely on linearization techniques to advance the solution of equations toward convergence, and it is possible that the linearization is inadequate to advance the solution. There are two techniques to deal with this issue:

- Improve the initial conditions. Alternatively, it can be achieved using techniques such as gradually activating models, changing boundary conditions, activating unsteady models from a steady solution, or ramping up solution accuracy.
- Advance the solution less by ramping up Courant numbers and relaxation factors.

When using these techniques, be sure to follow the guidelines which tell you if you are required to stop the solution or not before you change the conditions. These guidelines are described in the section that deals with changing the simulation while running.

See Also

Using Initial Conditions
Changing the Simulation While Running

Gradually Activating Models

There is no need to activate all the models at the beginning of the solution; models can be activated during the solution.

Examples of where it can be beneficial to activate models during the solution process are as follows:

- Compressible flows can be started inviscid (or laminar) before turning on turbulence. However, take care to ensure that the meshes are appropriate for the substituted models. It is possible that the Courant numbers and relaxation factors need ramped up after the models are changed.
- Unsteady models can be activated from a steady solution.

See Also

Changing the Simulation While Running
Ramping Up Courant Numbers and Relaxation Factors
Activating Unsteady Models from a Steady Solution

Changing Boundary Conditions or Ramping Boundary Values

This section provides information on changing boundary conditions or boundary values.

In some simulations, it can be difficult to start the solution with the final boundary conditions and boundary values. Possible examples of where it is advantageous to start with one set of boundary conditions or boundary values before changing them are as follows:

- When solving high-Mach number flows over bluff bodies, it can be beneficial to ramp-up the free-stream Mach number gradually.
- In some complex turbomachinery cases that use a combination of stagnation and pressure boundaries, establishing the correct flow direction at the inlet boundary can be difficult. It is therefore prudent to obtain an initial solution with another sort of inlet boundary (mass flow or velocity inlet) before switching to a stagnation inlet.

When changing the boundary type, take care to ensure that the types are compatible. For example, exercise caution in the following scenarios:

- For a highly compressible flow, it is possible that a velocity inlet is not a well-posed type.
- In incompressible flows, exercise caution when changing a flow split outlet to a pressure outlet. The pressure at a flow split outlet is not specified and is potentially considerably different from the pressure that is specified at a pressure outlet.

Activating Unsteady Models from a Steady Solution

It is common practice to start an unsteady simulation with a steady solution.

Potentially advantageous examples of this practice are:

- When doing a simulation of vortex shedding from a bluff body, you can often first obtain a steady-state solution, though there is no

guarantee that the solution you obtain is converged. In some cases, if a converged steady-state solution exists, it is possible that a change in inlet boundary conditions for a few time-steps is necessary. This action can cause a perturbation and prompt the shedding to occur.

- In some turbomachinery problems where a transient, sliding mesh calculation is required, you can often start with a steady solution using the multiple-reference frame (MRF) approach. Then you can initiate the unsteady and mesh motion models.

See Also

Gradually Activating Models

Ramping Up Solution Accuracy

It can be beneficial to start the solution with a first-order numerical scheme.

As discussed in the section on judging convergence, the dissipative errors resulting from first-order solutions are stabilizing and can enhance convergence. Therefore, it can be beneficial to start the solution with a first-order numerical scheme. For example:

- Models involving convective transport equations have either a *Convection* property (for segregated models), or a *Discretization* property (for the Coupled Flow model), that can be set to first order. The *Convection* property for segregated models affects only the upwind scheme for convection, leaving the diffusion and other gradients still fully second-order. The *Discretization* property for the Coupled Flow model neglects higher-order reconstruction for convection and diffusion.
- Solvers that are related to models that involve convective transport equations have a *Reconstruction Zeroed* expert property. Irrespective of the model type, this property causes reconstruction gradients to be neglected, effectively resulting in a first-order scheme.

See Also

Judging Convergence

Ramping Up Courant Numbers and Relaxation Factors

This section provides information on modifying values of Courant number and under-relaxation factors.

The default values of Courant number and the under-relaxation factors in STAR-CCM+ work in many situations, but are unlikely to be optimal in all situations, particularly when dealing with poor initial solutions. Suggestions for suitable alternative values are contained in the troubleshooting sections of the following relevant modeling guides:

- Coupled Flow and Coupled Energy
- Segregated Flow
- Segregated Fluid Energy
- Turbulence

See Also

Gradually Activating Models

Debugging the Simulation Setup

This section provides information on debugging the simulation setup.

When a simulation is diverging or a floating-point error appears, try to investigate where the problem is occurring; either in the solution process or its physical location in the mesh. STAR-CCM+ provides several tools to aid debugging that are not necessarily typically employed. Examples include:

- Using maximum and minimum reports
- Using threshold parts and cell surfaces
- Turning on temporary storage
- Zeroing reconstruction gradients
- Omitting secondary gradients

Using Maximum and Minimum Reports

In cases where the solution is “blowing” up, it is a good idea to create report-based monitors and to plot them to see when the solution starts diverging.

Typically a maximum report of velocity magnitude is a good indicator. In addition, minimum values of positive-definite quantities such as density, temperature, or turbulence quantities can be equally informative.

See Also

Using Report-Based Monitors
Creating a Monitor Plot

Using Threshold Parts and Cell Surfaces

This section provides information to help you to identify the physical location of the problem.

Once it has been determined that the solution is growing unbounded, it is useful to identify the physical location of the problem. Locating the problem can help identify if it is occurring near a boundary, or near a poor-quality cell.

Threshold parts are often the best way to view specific cells. Set the threshold to the required value (of, say, velocity magnitude), and once the threshold part is created, drag it into the scene. You can also create a cell surface from any other surface which lets you visualize a cell with a bad solution on a cutting plane.

See Also

Defining a Threshold
Adding Parts to Displayers
Defining a Cell Surface

Turning on Temporary Storage

You can turn on temporary storage for a particular solver in the Expert Properties category of that solver.

That action adds the field functions that let you view data such as Ap coefficients, residuals, corrections, reconstruction gradients, and cell gradients.

The Ap coefficients are the diagonal coefficients in the matrix of the linear system. Ensure that they are always positive-definite (that is, greater than 0).

The residuals are the imbalance in the discretized equations. The corrections are what is added to the solution of the previous iteration to get the updated solution. Generally, if the solution is not converging you can look where the residuals or the corrections are comparatively large. (Large corrections typically correspond to large residuals.)

Reconstruction gradients are used to interpolate cell values to faces, while gradients are used to compute values such as diffusion and strain rates. Visualizing these quantities can help to diagnose if something is going wrong at a flow boundary, say, where large gradients are not expected.

Field functions added when temporary storage is turned on are not saved to the simulation file. It is necessary that at least one step is iterated before the additional field functions become available.

Zeroing Reconstruction

When you zero reconstruction, the interpolation to the faces is a simple arithmetic averaging.

This averaging causes a less accurate solution, but improves the robustness. It is a useful thing to do if your solution is diverging, since it can reduce the magnitudes of variables like pressure and velocity, and keep the parts that are diverging more “local”. However, a converged solution with reconstruction zeroed does not have high accuracy.

See Also

Face

Omitting Secondary Gradients

Secondary gradients (also called cross-diffusion) are used in the diffusion terms and also in the pressure-velocity coupling for SIMPLE.

When the skewness angle gets large (89 degrees or more), these secondary gradients can become destabilizing. You can turn them off in the model, such as the segregated flow model. Do not accept a solution with secondary gradients off, since it is inaccurate unless the mesh is orthogonal. However, it can provide a clue and show if the diffusion terms are causing problems with the solution. If that is the case, improve the mesh and reduce the skewness angles. There are two separate switches: one for the boundary faces, and one for the interior faces.

See Also

Checking The Skewness Angle

Analyzing

STAR-CCM+ allows you to analyze the solution while the simulation is running, as well as when it completes.

Field functions provide access to raw simulation data, and to data computed from the raw data. These field functions can be analyzed using three general strategies:

- Creating reports of specific quantities. Four different types of report can be created, including one that accepts a user-defined expression in the same format as a user-defined field function. In order to obtain a history of the reported value, reports can be monitored while a simulation is running.
- Visualizing the solution data in scalar or vector scenes. If these scenes are created before you run the simulation, they are updated continuously as the solution progresses.
- Plotting sets of data that are extracted from the solution. You can create monitor and residual plots, XY plots, and histogram plots.

For each of these strategies, you generally choose the parts from which the solution data is extracted. These parts can be regions, boundary surfaces, or derived parts, depending on the type of analysis. Derived parts are additional lines and surfaces that you create in addition to the surfaces already present on geometry parts and regions.

Monitors are used to sample and save solution data while the simulation is running. Some monitor types are made available in response to the solvers you have active. Additional monitors can be created from a report.

Several features of STAR-CCM+ require that you define a trigger for when an update event takes place. See Triggers and Update Events.

For data that have been obtained over time in a transient analysis, such as pressure data at a point, data set functions are provided. These data set functions permit spectral analysis through Fast Fourier Transforms and related tools.

Accessing Solution Data through Derived Parts

Although STAR-CCM+ can display, analyze, and report simulation data on regions and boundaries, there are many cases when you want to access data in other parts of the solution space. Derived parts provide a mechanism for defining additional entities on which you can access solution data.

This section includes general information on working with derived parts, information for specific derived parts, and a method for calculating the surface area of a part.

What Is a Derived Part?

A part that is based on another part is known as a derived part.

A derived part is created using one or more other parts as parent parts, or inputs, and it is dependent on these parent parts. If any of the parent parts change, all parts depending on those parent parts are automatically recalculated and redisplayed to reflect the change.

For example, consider the following: An isosurface of pressure is created in a three-dimensional region and a plane-section is created from the isosurface resulting in a curve in space. The parent of a plane is the isosurface, and the parent of an isosurface is the three-dimensional region. If the region is modified somehow, for instance by either the mesh or the solution changing, first the isosurface is recalculated, and the plane follows. In this way, consistency is maintained for all parts.

Working with Derived Parts

Derived parts are used to analyze data within a STAR-CCM+ simulation.

The **Derived Parts** manager node, which has its own pop-up menu, contains all of the derived part nodes in the object tree.



You can add and edit any of various derived parts and use them to examine solution data. Creating and editing can be done visually using the powerful in-place dialog. A derived part gets its data when you select input (parent) parts for it. Certain derived parts can be defined using local coordinate systems.

You can control the naming of derived parts. If a derived part is no longer needed, it can be deleted. You can also copy a derived part to use it as a starting point for a new part, or to create a deep copy for use in a different simulation. It is possible to manipulate several of these derived parts visually using tools. Some derived parts use data from an external file. You can also export data that the derived parts extract.

Adding a New Derived Part

You can add derived parts with the context menu.

To add a new derived part, begin by right-clicking the **Derived Parts** node and choosing a part to create from the **New Part** submenu.



Three items in the submenu are derived part categories with their own submenus.



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

If the conditions for activating this dialog are not met, then this action causes the new derived part to appear in the simulation tree immediately. You can then define the part using its properties.

Creating a New Derived Part with the In-Place Dialog

When you execute the action for adding a derived part, typically an in-place dialog appears in the *edit* tab of the simulation window where you can define the part in detail.

The following screenshot is an example of the in-place dialog; if you are working with another type of part, the dialog that you use looks different from the one in this illustration.



This dialog is connected to the particular scene that is active when it opens. It disappears when the scene becomes inactive (for example,

when another display obscures it), and reappears when the scene becomes active again.

Click the **Input Parts** button to select the input parts that you want to include in the derived part. The other features in the dialog you are using depend on the type of derived part.

The **Create** button implements the settings that you entered in the rest of the derived part dialog. When you finish setting the parameters, click `Create` to put your specifications into effect, and to add the part to the simulation tree, then click `Close`. Alternatively, click `Close` to cancel the creation process and discard all your settings.

If you click **Create** more than once, you get an additional identical derived part each time. Also, if the Display group box is set for a new displayer, each extra click of the **Create** button produces an additional displayer.

Creating a New Derived Part with the Properties Window

You can specify settings for a new derived part in the Properties window as well as with the in-place dialog.

If you do not want the in-place dialog visible, make sure that none of the scene displays are active. Then right-click the **Derived Parts** node and choose a part to create.



A new node appears under the **Derived Parts** node, which you can visit and adjust the properties that define the part.



Editing Derived Parts

You can edit existing derived parts as needed.

Two methods are available for editing a derived part.

Editing a Derived Part with the In-Place Dialog

The in-place dialog lets you change the part with properties that are easier to use plus, in some situations, graphical tools.

To use the dialog, make sure that the conditions for its activation are met. Right-click the derived part (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**. This menu item is disabled (grayed out) if the dialog cannot be activated.



If you are editing a part that is not in the active scene, a list of displayers appears as a separate dialog ahead of the in-place dialog.



This dialog allows you to select a displayer for the part in advance. If you click **OK**, the in-place dialog appears near the scene.

If you click **Cancel**, the in-place dialog appears with the same list. The following screenshot is an example; if you are working with another type of part, the dialog that you use looks different from the one in this illustration.



The (**Up-Arrow**) button, at the top of the *Display* group box, lets you minimize the group box temporarily. When the group box is minimized,

its only visible parts are the title bar and a (**Down-Arrow**) button. Click anywhere within the title bar to restore the group box.

If the active scene already contains the part in a displayer, or if you made a selection in the *Select Displayer* dialog, then the Display group box does not appear.



This dialog is connected to the particular scene that is active when it opens. It disappears when the scene becomes inactive (for example, when another display obscures it), and reappears when the scene becomes active again.

Click the **Input Parts** button to select the input parts that you want to include in the derived part. The other features in the dialog you are using depend on the type of derived part.

The **Apply** button implements the settings that you entered in the rest of the derived part dialog. When you finish editing, click **Apply** to put your changes into effect, then click `Close`. Alternatively, click `Close` to cancel the editing process and discard all your changes.

The part that you are editing is modified each time that you click **Apply**. Also, if the Display group box is set for a new displayer, clicking the **Apply** button produces the additional displayer.

Editing a Derived Part with the Properties Window

The *Properties* window offers a more familiar approach to specifying the characteristics of your derived part, like many other objects in STAR-CCM+.

If you prefer to use the *Properties* window, select the node of the derived part and modify its properties.



Using the In-Place Dialog

This section provides details about the common features of the in-place dialog that is used for creating or editing derived parts.

The common features of this dialog are:

- Using the input part selector
- Selecting the display option

However, each dialog has other features that are specific to that type of derived part. For certain derived parts, this dialog also makes tools available.

When activated, this dialog causes the *Properties* window to disappear from the workspace. It reappears when the dialog is closed.

See Also

Using the Properties Window

Under Which Conditions Does the In-Place Dialog Appear?

If you prefer to use the in-place dialog, see this section about what is required for the dialog to be available.

For creating a derived part with the in-place dialog, the following conditions must be met in order for the in-place dialog to appear in a scene display:

- A scene display is active.
- Another derived part is not being edited in the current scene.

For editing a derived part with the in-place dialog, the following conditions must be met in order for the **Edit Part in Current Scene** pop-up menu item to be available:

- A scene display is active.
- Another derived part is not being edited in the current scene.
- The derived part that you are selecting is not already being edited in a scene.

Using the Input Part Selector

The **Input Parts** button lets you select the input parts that you want to include in the derived part.

For some derived parts, a region is selected by default since that type of part is a logical choice of parent part.



In other cases, the button is marked with empty brackets to indicate that no parts have been selected.



After the selection process, the input parts are named within the brackets. Since space in the button does not suffice for a complete list, you can hover with the mouse to display the information completely.



Selecting the Display Option

The *Display* group box allows you to select whether and how to add the part to the currently active scene.

This group box does not appear anywhere if you are editing a part that is already displayed in the active scene.



When you select one of the displayers that are listed in the group box and click **Create** or **Apply**, STAR-CCM+ automatically adds the part to that displayer. If you select one of the options for a new displayer, STAR-CCM+ creates a displayer of that type each time you press the **Create** button.

The (**Up-Arrow**) button, at the top of the group box, lets you minimize the group box temporarily. When the group box is minimized, its only visible parts are the title bar and a (**Down-Arrow**) button. Click anywhere within the title bar to restore the group box.

If you are editing a part that is not in the active scene, the list appears as a separate dialog ahead of the in-place dialog.



Using Type-Specific Dialogs

Some portions of in-place derived part dialogs are unique to that specific type of part.

Features of in-place dialogs that are specific to a particular derived part are explained in the section relevant to that part.

Selecting Input Parts

An input part can be envisioned as a parent or source part from which the derived part obtains its data.

To select the input boundaries or regions that you want to include in a derived part, click the right half of the *Parts* property of the node of the derived part. Alternatively, click the **Input Parts** button in the in-place dialog. This action activates the object selector.

Using Coordinate Systems with Derived Parts

You can link a derived part to a local coordinate system, and thus simplify your direction calculation by changing the frame of reference.

Linking to a coordinate system can be useful if, for example, you are modeling a certain part that is oriented differently from standard (laboratory) coordinates. Additionally, you can change the coordinates for more than one derived part by editing a single object, rather than modifying each derived part individually.

The following types of derived parts use local coordinate systems:

- Plane Section
- Constrained Plane Section
- Cylinder Section
- Sphere Section
- Point
- Line-Probe
- Presentation Grid

To use a local coordinate system with a derived part, first add it to the simulation. Then select the node of the derived part and choose that coordinate system from the drop-down list of the *Coordinate System* property.



Naming Derived Parts

When created, derived parts have a default name that is based on their type.

When an additional derived part of the same type is created in a simulation, a number is added to its default name, such as `plane section N`. `N` represents one plus the number of derived parts of that type existing in the simulation.

Derived parts can be renamed in STAR-CCM+.

Using the Tools

Included in STAR-CCM+ are tools that let you manipulate derived parts visually.

If a part becomes highlighted while you are using a tool, you can hold down the `<Ctrl>` key and click anywhere in the display to deactivate the highlighting.

Using the Point Tool

The point tool allows you to position a point in three-dimensional space using a three-dimensional cursor.

The cursor has an outline bounding box, axes-aligned cross-hairs, and axes shadows.



This widget operates in the following ways:

- To move the point, click any widget line and “slide” the point into position.
- To scale the widget, click and drag with the right mouse button up or down. The upward direction makes the widget bigger, while the opposite direction makes it smaller.
- To translate the widget, drag it while holding down the middle mouse button. All of the translation interactions can be constrained to one of the x, y, or z axes by using the `<Shift>` key.

  The constrained translation/sliding action (when the `<Shift>` key is depressed) along the axes is based on a combination of a “hot” spot around the cursor focus plus the initial mouse motion after selection. That is, if you select an axis outside of the hot spot, then the motion is constrained along that axis. If you select the point widget near the focus (within the hot spot), the initial motion defines a vector which is compared to the x-y-z axes. The motion is constrained to the axis that is most parallel to the initial motion vector.
- The widget can be picked even when it is “behind” other actors. It is not a bug: this feature is intended.

Using the Line Tool

The three-dimensional line tool defines a line that can be interactively placed in a scene.

The line has two handles (at its endpoints), and the line can be picked to translate it in the scene.



- To orient and stretch the line, click one of the two end-point handles and drag it. (The other end point remains fixed.)
- To scale the line from its center, click and drag with the right mouse button up or down. The upward direction makes the line bigger, while the opposite direction makes it smaller.
- To translate the line, drag it while holding down the middle mouse button. All of the translation interactions can be constrained to one of the x, y, or z axes by using the `<Shift>` key.
- The handles and line can be picked even when it is “behind” other actors. It is not a bug: this feature is intended.

Using the Implicit Plane Tool

This three-dimensional tool defines an infinite plane that can be interactively placed in a scene.

The widget appears as a plane with a normal vector; the plane is contained by a bounding box. The intersection edges of the plane and the bounding box are shown tubed.



The normal vector handle can be selected and moved to rotate the plane; the plane itself can be selected and translated in various directions.

- The scale of the bounding box is set automatically to the combined size of the input parts that you select for the plane.
- The widget can be picked even when it is “behind” other actors. It is not a bug: this feature is intended
- If you select the normal vector, you can rotate the plane arbitrarily. You can also translate the plane along the normal by clicking and dragging.
- You can arbitrarily move the plane origin by clicking and dragging the plane with the middle mouse button.
- You can constrain the translation of the plane by ticking the **Snap To Part** option in the *Plane Parameters* group box, and selecting the part in the drop-down menu. This option makes it easy to align a plane with a certain part. (This action is essentially snapping to the bounding box of the part.)
- To refine the position of the plane, you can adjust the coordinates with text entries.

  After each text entry, press `<Enter>` to apply your adjustment.



Alternatively, you can define the plane with two mouse clicks in the display. First, position the object in the scene so that your view is in the direction of the plane you want to define. Click the button to the left of the coordinate text boxes.



Click once in the display for the first point, and then click the second point to draw the line defining the plane.



This action repositions the plane tool.



Using the Bounded Implicit Plane Tool

This tool lets you create a loop that would limit the extent of the implicit plane.

This device gives you greater control of the scope of data that you sample from the plane section. Before using this tool, finish using the implicit plane tool.

Select the **Modify Contour** button to activate the loop feature.



To begin creating the loop, click once in the plane tool to define the first point. Each subsequent click in the plane draws a line connecting your last point with the new one, drawing the bounding loop.



To close the loop, click the first point. You can reverse your steps by using the **Delete Last Point** button, or you can start over by clicking the **Clear Loop** button.

Using the Explicit Plane Tool

This three-dimensional tool defines a finite (bounded) plane that can be interactively placed in a scene.

The plane has four handles (at its corner vertices), a normal vector, and the plane itself. The handles are used to resize the plane; the normal vector to rotate it, and the plane can be picked and translated.



- If you click the plane and drag it while holding down the `<Ctrl>` key, it rotates around the normal.
- To resize the plane, click and drag one of the four handles on the corners.
- If you select the normal vector, you can rotate the plane arbitrarily.
- You can translate the plane along its normal by clicking and dragging it with the middle mouse button.
- To scale the plane from its center, click and drag with the right mouse button up or down. The upward direction makes the plane bigger, while the opposite direction makes it smaller.
- You can constrain the translation of the plane by ticking the **Snap To Part** option in the *Plane Parameters* group box, and selecting the part in the drop-down menu. This option makes it easy to align the plane with a certain part. (This action is essentially snapping to the bounding box of the part.)
- To refine the position of the plane, you can adjust the coordinates with text entries. After each text entry, press `<Enter>` and click `Update Tool` to apply your adjustment.
- The plane can be picked even when it is “behind” other actors. It is not a bug: this feature is intended.

Using the Sphere Tool

This three-dimensional tool defines a sphere that can be interactively placed in a scene.

You can click and drag the sphere to change its position in the scene.



- To resize the sphere, you can change the radius with a text entry. After the entry, press `<Enter>` to apply your adjustment.
- To refine the position of the sphere, you can adjust the coordinates of the origin with text entries. After each text entry, press `<Enter>` to apply your adjustment.
- You can constrain the translation of the sphere by ticking the **Snap to Part** option in the *Sphere Parameters* group box, and selecting the part in the drop-down menu. This option makes it easy to align the sphere with a certain part. (The sphere is snapped to the bounding box of the part.)

Using the Resampled Volume Tool

This three-dimensional tool defines a resampled volume that can be interactively placed in a scene.

To control the size and position of the resampled volume, use the handles and the bounding box.



- To resize the resampled volume, click and drag the handle at any of the sides of the bounding box.
- To rotate the resampled volume, click and drag the outline of the bounding box.
- To translate the resampled volume, click and drag the handle at the center of the bounding box.
- To constrain the translation of the resampled volume, activate the **Snap to Part** option and select a part. This option makes it easy to align the resampled volume with a certain part. (The resampled volume is snapped to the bounding box of the part.)

Reading an External File

An arbitrary section, arbitrary probe, or constrained plane section can read data from an external file (STL or VTK).

Note: Before importing the data for a constrained plane section, create an arbitrary constrained planar section.

1. To access this file, right-click the node of the derived part and select **Re-Import Data...**



2. Select an available file from the *Open* dialog, then click **Open**. This action adds the data from the specified file.



The **Decimate input** checkbox, which is activated by default, facilitates creation of an arbitrary section.

Exporting to an External File

You can export the data of derived parts to an external file in STL or VTK format.

Both file formats can be binary or ASCII, and STL files are a universal method of transferring surface data using the least possible number of triangles.

1. To export the data of a derived part, right-click the node of the part in the object tree and select **Export...**



2. In the *Save* dialog, enter a name for the `.stl` or `.vtk` file and ensure that the correct representation is selected in the **Representations** menu. By default, the file is saved in binary format. To save the file in ASCII format, tick the **Export as Ascii** checkbox.



For a constrained plane section, arbitrary section or arbitrary probe, the file that this action generates is different from the imported file, as only the extracted triangle data is exported. When streamlines are exported as an STL or VTK file, they are defined as thin ribbons.

The derived part must have one or more input parts, otherwise STAR-CCM+ considers it empty. An attempt to export it fails with a message like this one:



Pop-Up Menu of the Derived Parts Manager

This section lists the options in the pop-up menu of the **Derived Parts** manager node.

These pop-up menu items are available when you right-click the **Derived Parts** manager node.

New Part	Creates a new derived part of the selected type, using a graphical editor when a visualization display is open.		
	Cell Surface...	Creates a cell surface.	
	Isosurface...	Creates an isosurface.	
	Probe	Creates a probe derived part of the type you select in the submenu.	
		Point...	Creates a point.
		Line...	Creates a line probe.
		Presentation Grid...	Creates a presentation grid.
		Arbitrary...	Creates an arbitrary probe.
	Section	Creates a section derived part of the type you select in the submenu.	
		Plane...	Creates a plane section.
		Cylinder...	Creates a cylinder section.
		Sphere...	Creates a sphere section.
		Constrained Plane...	Creates a constrained plane section.
		Arbitrary...	Creates an arbitrary section.
	Streamline...	Creates a streamline.	
	Constrained Streamline...	Creates a constrained streamline.	
	Threshold...	Creates a threshold.	
	(m′, theta) warp	Creates an (m′, theta) warp.	

	Resampled Volume...	Creates a resampled volume.
	Warp...	Creates a warp.
New Group	Facilitates the grouping of objects by creating a group node within the manager node.	
Group By	Allows you to group the objects according to any of the properties of the objects in the manager node, which are listed in a submenu of this menu item.	
Ungroup	Reverses the grouping process.	

Working with Implicit Sections

Implicit sections are derived parts that are based on planes or geometrical shapes, which can be cut through one or more volumes or surfaces.

These parts are implicit because they are functions of coordinates, rather than combinations of mesh cells.

Within an implicit section part, it is possible to create multiple instances of the section by defining offset copies of the original. This technique allows you to expand the potential range of sample data for the part.

Defining a Plane Section

The plane section is an implicit part that is derived from one or more input parts.

The node of a plane section has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



A plane section uses a plane as the implicit function to cut. The resulting output is dependent on the topology and discretization of the input parts. The exact form of output depends on the dimensionality of the input parts. Input data consisting of three-dimensional cells result in output surfaces, input data consisting of two-dimensional cells result in output lines, and input data consisting of one-dimensional data result in points. STAR-CCM+ allows the creation of a planar section from parts comprising regions, boundaries, and geometry parts, however input parts are enforced to be of the same dimensionality.



When a plane section is created, you can specify the plane properties manually with one of the following methods:

- Entering the coordinates of the plane origin and normal vector
- Using the plane tool
- Using the toolbar

Example of Creating a Plane Section with a Cylindrical Coordinate System

If you create a plane section with a cylindrical coordinate system, specify r, θ, and z coordinates. The resulting normal vector consists of the r-direction (the cylinder radius), the z-direction (the cylinder length), and θ, the angle of rotation about the z-axis. The plane section is perpendicular to the normal.

In the example illustrated below, the normal is 1 m long, with values of 45 degrees for θ and 0 for z, so the normal lies in the r-θ plane with the plane section being perpendicular to it.



Plane Section Properties

You can use the properties that are listed below to adjust the specifications of the plane section.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part. See also Example of Creating a Plane Section with a Cylindrical Coordinate System.
Origin	A Cartesian position vector that defines the origin of the plane. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).
Normal	A Cartesian position vector that defines the second point through which the plane normal passes. The first point is the origin of the selected *Coordinate System*. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).

Section Mode	Defines whether the part includes a single or multiple sections.	
	Single Section	Specifies a single section (default). A **Single section** node appears under the node of the derived part.
	Multiple Sections (range)	Uses a range technique for creating multiple sections. A **Multiple Sections (range)** node appears under the node of the derived part.
	Multiple Sections (delta)	Uses a delta technique for creating multiple sections. A **Multiple Sections (delta)** node appears under the node of the derived part.
	Multiple Sections (arbitrary)	Uses an arbitrary technique for creating multiple sections. A **Multiple Sections (arbitrary)** node appears under the node of the derived part.
Displayed Index	Controls the display of sections when a multiple mode is used. Enter either a positive integer to display the corresponding section, or a value of -1 to display all sections.	
Parts	Selects the input parts as the source for deriving this plane section.	

Plane Section Pop-Up Menu

Edit Part in Current Scene	Executes the define plane section action allowing you to set the properties of this part interactively. This option is only available if a scene is open.

Export...	Exports the extracted data of the part to an external STL or VTK file.

Creating a Plane Section

You can create a plane section with the pop-up menu or the toolbar button.

To create a plane section, right-click the **Derived Parts** node and select **New Part > Section > Plane...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

If the conditions for activating this dialog are not met, then this action causes the new plane section to appear in the simulation tree immediately. You can then define the part using the properties.

Creating a Plane Section with the Toolbar Button

You can use the (**Create Plane Section**) toolbar button to create a plane section in an open scene display quickly.

The plane is defined with two points that you specify. First position the image in the scene so that your view is exactly perpendicular to the plane. Click the toolbar button, then click at two points in the display to define the plane.



Select a displayer in the dialog. You can choose one of the existing displayers, create a new one, or create the derived part without assigning it to any displayer.



The plane section appears in the scene. By default, it has the appropriate input parts from among the ones that are visible in the scene. You can edit this derived part to change the input parts or other settings. (This screenshot shows a transparent scene.)



Editing a Plane Section

Multiple techniques are available for editing a plane section.

You can edit a plane section using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use its properties.

Using the Plane Section In-Place Dialog

This dialog facilitates the definition of the plane section with interactive features.

If certain basic conditions are met when you start the process to create a plane section or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the plane section.

You can use the plane tool in the scene display, which is activated automatically with the dialog, to position the plane visually. By clicking the button to the left of the coordinate text boxes, you can position the plane tool with two mouse clicks in the display. Also, you can define the plane section with the following specifications in the *Plane Parameters* group box:

- origin—sets x, y, and z coordinates for the origin. You can include units in the entry.
- normal—sets x, y, and z coordinates for a point that is used to define the surface normal. The normal is defined with the vector between the origin of the selected coordinate system and this point. You can include units in the entry.
- A drop-down list that selects from among the coordinate systems existing in the simulation (Laboratory by default).
- **Snap To Part**—selects a part along which you can position the plane; available when the plane tool is active. (This action is essentially snapping to the bounding box of the part.)
- A drop-down list and text boxes for specifying single or multiple sections

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your plane section, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the plane section, that is shown in brown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining a Constrained Plane Section

A constrained plane section is similar to a plane section, except that it is bounded in a loop (a set of coordinates that you specify).



A volume mesh must be present for the plane to appear. The edges of the constrained plane snap to the nearest whole cell, as shown above.

A constrained plane section node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



The constrained plane section is also different to a standard plane section in that it can read shape data from an external file. An arbitrary constrained planar section must be created first, before an external file can be read. The external file must be in .stl or .vtk file format, and contain a surface that is used as the constrained planar section. The loop of the constrained plane consists of the perimeter edges of the imported surface. If the surface you are importing was created in a commercial CAD package, adjust its position and orientation as needed using the in-place dialog. This task is done by editing the constrained plane. Ensure that the units of the exported `.stl` file and the STAR-CCM+ file match.

The loops that you create in STAR-CCM+ can be exported by right-clicking on the constrained plane section node in the object tree and selecting **Export Loop...** The loop is saved in `.stl` format.



Constrained Plane Section Properties

You can use the properties that are listed below to adjust the specifications of the constrained plane section.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part.	
Normal	A Cartesian position vector that defines the second point through which the plane normal passes. The first point is the origin of the selected *Coordinate System*. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button ( ).	
Origin	A Cartesian position vector that defines the location of the plane. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button ( ).	
Section Mode	Defines single or multiple sections.	
	Single Section	Specifies a single section (default). A

		Single section node appears under the node of the derived part.
	Multiple Sections (range)	Uses a range technique for creating multiple sections. A **Multiple Sections (range)** node appears under the node of the derived part.
	Multiple Sections (delta)	Uses a delta technique for creating multiple sections. A **Multiple Sections (delta)** node appears under the node of the derived part.
	Multiple Sections (arbitrary)	Uses an arbitrary technique for creating multiple sections. A **Multiple Sections (arbitrary)** node appears under the node of the derived part.
Displayed Index	Controls the display of sections when a multiple mode is used. Enter either a positive integer to display the corresponding section, or a value of -1 to display all sections.	
Parts	Selects the input parts as the source for deriving this constrained plane section. See Selecting Input Parts.	

Constrained Plane Section Pop-Up Menu

Edit Part in Current Scene	Allows you to set the properties of this part interactively. This option is only available if a scene is open. See Using the Constrained Plane Section In-Place Dialog.
Export...	Exports the extracted data of the part to an external STL or VTK file. See Exporting to an External File.

Re-Import Data...	Imports data from either an STL or VTK file. This option can be used to replace data that are imported during the creation of this derived part.
Export Loop...	Exports the polygon bounding the constrained plane section to an external an STL file.

Creating a Constrained Plane Section

You can create a constrained plane section with the pop-up menu.

To create a constrained plane section, right-click the **Derived Parts** node and select **New Part > Section > Constrained Plane...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new constrained plane section to appear in the simulation tree immediately. You can define the part using the properties.

Editing a Constrained Plane Section

Multiple techniques are available for editing a constrained plane section.

You can edit a constrained plane section using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Constrained Plane Section In-Place Dialog

This dialog facilitates the definition of the constrained plane section with interactive features.

If certain basic conditions are met when you start the process to create a new constrained plane section or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the constrained plane section.

To define the limits of the constrained plane section, you would use the bounded implicit plane tool, which can be activated with the **Modify Contour** button in the *Widget Mode* group box. Before doing so, however, finalize the position of the plane with the plane tool, which is selected by default with the **Modify Plane** button.

The plane tool lets you position the plane visually in the scene display.

By clicking the button to the left of the coordinate text boxes, you can position the plane tool with two mouse clicks in the display. Also, you can define the plane section with the following specifications in the *Plane Parameters* group box:

- origin—sets x, y, and z coordinates for the origin. You can include units in the entry.
- normal—sets x, y, and z coordinates for a point that is used to define the surface normal. The normal is defined with the vector between the origin of the selected coordinate system and this point. You can include units in the entry.
- A drop-down list that selects from among the coordinate systems existing in the simulation (Laboratory by default).
- **Snap To Part**—selects a part along which you can position the plane; available when the plane tool is active. (This action is essentially snapping to the bounding box of the part.)

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your constrained plane section, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the constrained plane section, which is shown in brown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining a Cylinder Section

The cylinder section is an implicit part that is created from one or more input parts.

A cylinder section node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



A cylinder section uses a cylinder as the implicit function to cut. The exact form of output depends on the dimensionality of the input parts. Input data consisting of three-dimensional cells result in output surfaces, input data consisting of two-dimensional cells result in output lines, and input data consisting of one-dimensional data result in points. STAR-CCM+ allows the creation of a cylinder section from parts comprising regions, boundaries, and other parts, provided all Input parts are of the same dimensionality. When a cylinder section is created, the location and orientation of the cylinder is specified by entering the coordinates of the cylinder origin and axis.



Cylinder Section Properties

You can use the properties that are listed below to adjust the specifications of the cylinder section.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part.
Origin	A Cartesian position vector that defines the origin of the cylinder axis. You can include units in the entry, or select the units in the

	dialog that is activated by clicking the property customizer button ( ).	
Orientation	A Cartesian position vector that defines the second point through which the cylinder axis passes. The first point is the origin of the selected *Coordinate System*. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button ( ).	
Radius	The radius of the cylinder. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button ( ).	
Section Mode	Defines a single section or multiple sections.	
	Single Section	Specifies a single section (default). A **Single section** node appears under the node of the derived part.
	Multiple Sections (range)	Uses a range technique for creating multiple sections. A **Multiple Sections (range)** node appears under the node of the derived part.
	Multiple Sections (delta)	Uses a delta technique for creating multiple sections. A **Multiple Sections (delta)** node appears under the node of the derived part.
	Multiple Sections (arbitrary)	Uses an arbitrary technique for creating multiple sections. A **Multiple Sections (arbitrary)** node appears under the

		node of the derived part.
Displayed Index	Controls the display of sections when a multiple mode is used. Enter either a positive integer to display the corresponding section, or a value of -1 to display all sections.	
Parts	Selects the input parts as the source for deriving this cylinder section. See Selecting Input Parts.	

Cylinder Section Pop-Up Menu

Edit Part in Current Scene	Executes the define cylinder action allowing you to set the properties of this part interactively. This option is only available if a scene is open. See Using the Cylinder Section In-Place Dialog.
Export...	Exports the extracted data of the part to an external STL or VTK file. See Exporting to an External File.

See also: common menu items.

Creating a Cylinder Section

You can create a cylinder section with the pop-up menu.

To create a cylinder section, right-click the **Derived Parts** node and select **New Part > Section > Cylinder...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new cylinder section to appear in the simulation tree immediately. Define the part using the properties.

Editing a Cylinder Section

Multiple techniques are available for editing a cylinder section.

You can edit a cylinder section using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Cylinder Section In-Place Dialog

This dialog facilitates the definition of the cylinder section with interactive features.

If certain basic conditions are met when you start the process to create a cylinder section or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the cylinder section. See Selecting Input Parts.

In the *Cylinder Parameters* group box, you can define the cylinder section with the following specifications:

- origin—sets x, y, and z coordinates for the origin of a rotation axis through the center of the cylinder. You can include units in the entry.
- orientation—sets x, y, and z coordinates for a point that is used to define the central axis of the cylinder. The axis is defined by the vector between the origin of the selected coordinate system and this point. You can include units in the entry.
- A drop-down list that selects from among the coordinate systems existing in the simulation (Laboratory by default).
- **Radius**—specifies distance in a perpendicular direction from the center axis, defining the outer wall of the cylinder. You can include units in the entry.
- A drop-down list and text boxes for specifying single or multiple sections

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your cylinder section, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the cylinder, which is shown in brown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.

You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining a Sphere Section

STAR-CCM+ allows the creation of a sphere section from parts comprising regions, boundaries, and other parts.

The sphere section is an implicit part that is created from one or more input parts. A sphere section node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



A sphere section uses a sphere as the implicit function to cut. The exact form of output depends on the dimensionality of the input parts. Input data consisting of three-dimensional cells result in output surfaces, input data consisting of two-dimensional cells result in output lines, and input data consisting of one-dimensional data result in points. Input parts are enforced to be of the same dimensionality. The location of the sphere section is specified by entering the coordinates of the sphere origin and its radius.



Sphere Section Properties

You can use the properties that are listed below to adjust the specifications of the sphere.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part.	
Origin	A Cartesian position vector that defines the location of the sphere. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).	
Radius	The radius of the sphere. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).	
Section Mode	Defines a single section or multiple sections.	
	Single Section	Specifies a single section (default). A **Single section** node

		appears under the node of the derived part.
	Multiple Sections (range)	Uses a range technique for creating multiple sections. A **Multiple Sections (range)** node appears under the node of the derived part.
	Multiple Sections (delta)	Uses a delta technique for creating multiple sections. A **Multiple Sections (delta)** node appears under the node of the derived part.
	Multiple Sections (arbitrary)	Uses an arbitrary technique for creating multiple sections. A **Multiple Sections (arbitrary)** node appears under the node of the derived part.
Displayed Index	Controls the display of sections when a multiple mode is used. Enter either a positive integer to display the corresponding section, or a value of -1 to display all sections.	
Parts	Selects the input parts as the source for deriving this sphere section.	

Sphere Section Pop-Up Menu

Edit Part in Current Scene	Executes the define sphere action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL or VTK file.

See also: common menu items.

Creating a Sphere Section

You can create a sphere section with the pop-up menu.

To create a sphere section, right-click the **Derived Parts** node and select **New Part > Section > Sphere...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new sphere section to appear in the simulation tree immediately. Define the part using the properties.

Editing a Sphere Section

Multiple techniques are available for editing a sphere section.

You can edit a sphere section using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Sphere Section In-Place Dialog

This dialog facilitates the definition of the sphere section with interactive features.

If certain basic conditions are met when you start the process to create a sphere section or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the sphere section.

You can use the sphere tool in the scene display, which is activated automatically with the dialog, to position the sphere visually. Also, you can define the sphere section with the following specifications in the *Sphere Parameters* group box:

- sphere origin—sets x, y, and z coordinates for the sphere origin. You can include units in the entry.
- **Radius**—sets distance from the sphere origin, defining the outer wall of the sphere. You can include units in the entry.
- A drop-down list that selects from among the coordinate systems existing in the simulation (Laboratory by default).
- **Snap To Part**—selects a part along which you can position the sphere. (This action is essentially snapping to the bounding box of the part.)
- A drop-down list and text boxes for specifying single or multiple sections

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your sphere section, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the sphere, which is shown in brown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.

You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining an Arbitrary Section

The arbitrary section is an implicit part that samples one or more input parts or regions.

An arbitrary section node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree. Arbitrary sections should be used when a nonplanar sampling surface is desired.



The geometry of an arbitrary section can be created from either a CAD part surface or directly from a STL file. This representation of the input geometry is mesh-resolution-dependent.



A lower polygon count in an arbitrary section results in a faster intersection computation between the section and the sampled region.

Arbitrary Section Properties

Parts	Selects the input parts as the source for deriving this arbitrary section. See Selecting Input Parts.	
Surface Input Option	Specifies the source of the input geometry.	
	File	Reads in arbitrary STL data.
	Part Surfaces	Selects one or more of the existing CAD part surfaces in the simulation.
Extract Approximate Faces	Takes data from the closest faces of the mesh cells that intersect the arbitrary section, and creates an approximation of the surface.	
	Activated	Creates an approximation of the surface that maps to the closest mesh cell faces. This option is faster because it does not have to compute the intersection between the section's surface geometry and the part or region.
	Deactivated	The arbitrary section retains its original surface geometry and slices through the sample region.
Flip Orientation	Reverses the directions of the normal vectors of the sample surface faces.	
	Activated	Changes the orientation. The results become evident in reporting for

		relevant quantities such as mass flux.
	Deactivated	The face normals are right-hand oriented with respect to the counter-clockwise winding order of the original triangles in the section geometry.

Arbitrary Section Pop-Up Menu

Edit Part in Current Scene	Allows you to set the properties of this part interactively. This option is only available if a scene is open. See Using the Arbitrary Section In-Place Dialog.
Export...	Exports the geometry of the resulting intersection to an external STL or VTK file. See Exporting to an External File.
Re-Import Data...	Imports the surface geometry from an external STL or VTK file.

See also: common menu items.

Creating an Arbitrary Section

You can create a arbitrary section with the pop-up menu.

To create an arbitrary section, right-click the **Derived Parts** node and select **New Part** > **Section** > **Arbitrary...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action. However, if the conditions for activating this dialog are not met, then this action causes the *Select Input Surfaces* dialog to appear instead. Each of these dialogs allows you to create the section geometry either by using part surfaces or by importing an STL file.



To work with part surfaces:

- Select the **Part Surfaces** option.

  The dialog will change to allow you to activate the object selector.
- Make your selections and close the object selector.
- In the *Select Input Surfaces* dialog, click **OK**.

  This action adds the data from the specified part surfaces.

To import an STL file:

- Select the **File** option. When you click **OK**, the *Open* dialog appears.



- Select an available STL file from this dialog.

  During the import, STAR-CCM+ retriangulates the surface to reduce the number of triangles. The resulting coarser resolution can reduce processing time significantly. However, this retriangulation can also cause problems for surfaces that have curves or other complex characteristics. To disable this feature and retain the resolution of the input file, clear the **Decimate input** checkbox in the dialog (which is activated by default). Since the surface decimation algorithm is based on the local curvature, planar sections of the surface will not be retriangulated.

- Click **Open**.

  This action adds the data from the specified file.

The new arbitrary section appears in the simulation tree immediately after you select the part surface or import the external file. You can define the part using the properties.

Editing an Arbitrary Section

Multiple techniques are available for editing a arbitrary section.

You can edit an arbitrary section using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Arbitrary Section In-Place Dialog

This dialog facilitates the definition of the arbitrary section with interactive features.

If you have an appropriate scene open when creating or editing an arbitrary section, STAR-CCM+ displays an in-place dialog. Use this dialog to define the section.



Click the **Input Parts** button to select the input parts that you want to include in the arbitrary section. See Selecting Input Parts.

To define the limits of the arbitrary section, get shape data from part surfaces or an external file. Choose a source by selecting one of the options in the *Surface Input Options* group box. (See Creating an Arbitrary Section for details.)

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your arbitrary section, you can view it in the assigned scene.

Once you have created the arbitrary section, rotate and zoom in the display to get a clear view of it. An example is shown highlighted in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Using Multiple Implicit Sections

Within an implicit section part, it is possible to create multiple instances of the section by defining offset copies of the original.

This feature gives you the benefit of sampling data from multiple sections while you still work with just one part.

The following screenshot shows how you can use multiple sections in the different implicit parts.



By default an implicit section is a single section. To define a pattern of sections, three additional options are available:

Working with a Single Section

By default, implicit sections are single sections, meaning that there is only one instance of the section within the derived part.

The setting for this section is defined in the *Section Mode* property of the derived part, and adds a **Single Section** node to the part node. Within the properties of this node, you can specify an *Offset* value for the implicit section, which allows you to offset the section from its original location. For planes, this value is the offset distance along the normal from the origin, and for cylinders or spheres, it the offset value for the radius. The node is also subject to common menu actions. See Common Pop-Up Menu Items.



In the in-place dialog for plane, cylinder or sphere sections this setting is available in the Mode drop-down list and you can enter an offset value in the **Offset** text box



Single Section Properties

Offset	The offset distance along the normal of the plane from the origin. For spheres or cylinders, the offset value for the radius. This property can be set with a slider. See Using a Slider.

Working with Multiple Sections (Range)

Defining multiple sections using the Range option allows you to specify a range of sections that are evenly distributed between the start and end offset section positions.

The range setting is applied by selecting the **Multiple Sections (range)** option for the *Section Mode* property of the derived part, and adds a **Multiple Sections (range)** node to the part node. Within the properties of this node, you can specify the total number of sections to create, and the offset position of the first and last sections. For planes, the offset

position is the distance along the normal from the origin, and for cylinders or spheres, it is the offset value for the radius. The node is also subject to common menu items. See Common Pop-Up Menu Items.



In the in-place dialog for plane, cylinder or sphere sections this setting is available in the Mode drop-down list and you can enter the properties in the corresponding textboxes.



Multiple Sections (Range) Properties

Number of Sections	The number of sections to distribute over the range.
Starting Offset	The distance of the first section along the axis normal to the section, or the first offset radius for spheres or cylinders. This property can be set with a slider. See Using a Slider.
Ending Offset	The distance of the last section along the axis normal to the section, or the last offset radius for spheres or cylinders. This property can be set with a slider. See Using a Slider.

Working with Multiple Sections (Delta)

Defining multiple sections using the Delta option allows you to specify a pattern of sections. This pattern uses the total number of sections and the incremental distance between each section.

The setting is applied by selecting the **Multiple Sections (delta)** option for the *Section Mode* property of the derived part, and adds a **Multiple Sections (delta)** node to the part node. Within the properties of the node, you can specify:

- The total number of sections to create
- The initial offset for the first section
- The offset increment to apply from one section to the next

For planes, the offset position refers to the distance along the normal from the origin, and for cylinders or spheres, it is the offset value for the radius. The node is also subject to common menu items. See Common Pop-Up Menu Items.



In the in-place dialog for plane, cylinder or sphere sections this setting is available in the Mode drop-down list and you can enter the properties in the corresponding textboxes.



Multiple Sections (Delta) Properties

Delta Value	The incremental offset between successive sections.
Number of Sections	The total number of implicit sections in the derived part.
Starting Offset	The distance of the first section along the axis normal to the section, or the first offset radius for spheres or cylinders. This property can be set with a slider. See Using a Slider.

Working with Multiple Sections (Arbitrary)

Defining multiple sections using the Arbitrary option allows you to specify a series of sections that are based on individual offset values.

The setting is applied by selecting the **Multiple Sections (arbitrary)** option for the *Section Mode* property of the derived part, which adds a **Multiple Sections (arbitrary)** node under the part node. Within the properties of this node, you can specify a series of sections by entering offset values for the sections in a comma-separated list. For example entering `[0.0, 0.01, 0.05, 0.09, 0.1]` would create 5 offset sections at the specified offset distances. For planes, the offset value is the distance along the normal from the origin, and for cylinders or spheres, it is the offset value for the radius. The node is also subject to common menu items. See Common Pop-Up Menu Items.



To specify arbitrary multiple sections in the in-place dialog, select the **Multiple (arbitrary)** setting in the Mode drop-down list and enter the offset values in the **Offsets** textbox. This feature is available for plane, cylinder, or sphere sections.



Each value that is entered in the **Offsets** box is another subdivision.

Multiple Sections (Arbitrary) Properties

Offsets	An array of offset values with comma-separated text entries.

Working with Scalar Value-Based Parts

This section explains the derived parts that are based on scalar values that you select.

Defining an Isosurface

The isosurface is a surface representing a constant-value scalar function.

An isosurface node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



The exact form of the output of the isosurface depends on the dimensionality of the input parts. Input data consisting of three-dimensional cells result in output isosurfaces, input data consisting of two-dimensional cells result in iso lines, and input data consisting of one-dimensional data result in points. STAR-CCM+ allows the creation of an isosurface from parts comprising regions, boundaries, and other parts, however input parts are enforced to be of the same dimensionality. The properties of an isosurface part are the scalar quantity to isosurface, the mode (either single or multiple) and the parent parts from which to extract the isosurface.



Isosurface Properties

You can use the properties that are listed below to adjust the specifications of the isosurface.

Scalar Field	Scalar field function upon which the isosurface is based.	
Parts	Selects the input parts as the source for deriving this isosurface. See Selecting Input Parts.	
Mode	The mode for isosurface creation, specifying whether a single surface or multiple surfaces are created.	
	Single	Indicates that a single isosurface based on a single scalar value is created. See Using the Single Isovalue
	Multiple (range)	Selects a range technique for creating multiple isosurfaces. A

		Values (range) node appears in the node of the derived part. See Using Multiple Isovalues (Range).
	Multiple (arbitrary)	Selects an arbitrary technique for creating multiple isosurfaces. A Values (arbitrary) node appears in the node of the derived part. See Using Multiple Isovalues (Arbitrary).

Isosurface Pop-Up Menu

Edit Part in Current Scene	Executes the define isosurface action allowing you to set the properties of this part interactively. This option is only available if a scene is open. See Using the Isosurface In-Place Dialog.
Export...	Exports the extracted data of the part to an external STL file. See Exporting to an External File.

See also Common Pop-Up Menu Items.

Creating an Isosurface

You can create a isosurface with the pop-up menu.

To create an isosurface, right-click the **Derived Parts** node and select **New Part > Isosurface...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new isosurface to appear in the simulation tree immediately. Define the part using the properties.

Editing an Isosurface

Multiple techniques are available for editing a isosurface.

You can edit an isosurface using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Isosurface In-Place Dialog

This dialog facilitates the definition of the isosurface with interactive features.

If certain basic conditions are met when you start the process to create an isosurface or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the isosurface. See Selecting Input Parts.

Specify the characteristics of the isosurface with the following features:

- Scalar drop-down list—this list selects the scalar field function as the basis for the isosurface. The contents of the scalar list depends on the models that you selected.

- *Scalar Range* group box—for automatic selection of the scalar range:
  - **min**—sets lowest scalar value (read-only)
  - **max**—sets highest scalar value (read-only)
  - `Query` button—instructs STAR-CCM+ to probe the selected scalar and fill in the values
- Extraction Mode drop-down list:
  - **Single Value**—extracts selected scalar data from a single isosurface
    - **isovalue**—available when **Single Value** is selected. You can enter the single isovalue here, or use the value that is generated by STAR-CCM+ when you click the `Query` button. STAR-CCM+ calculates this value as the average of the minimum and maximum values.
  - **Multiple (range)**—extracts selected scalar data from multiple isosurfaces that are spaced evenly within a specified range, and makes the following text boxes available for entering values:
    - **Number of isosurfaces**—sets the number of equally spaced isosurfaces.
    - **Minimum isovalue**—enter the lowest scalar value here, or use the value that STAR-CCM+ generates when you click the `Query` button.
    - **Maximum isovalue**—enter the highest scalar value here, or use the value that STAR-CCM+ generates when you click the `Query` button.
  - **Multiple (arbitrary)**—extracts scalar data from multiple isosurfaces corresponding to a set of scalar values that you specify with numeric entries. This option makes an **Isovalues** text box available.

The units can be changed by typing them in after the values.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your isosurface, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the isosurface, which is shown in blue in

the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display. Finally, you can show a smooth surface by activating the *Smooth Shade* expert property of the geometry displayer. See Smooth Shading Displayer Surfaces.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Using the Single Isovalue

Single Value mode for the isosurface involves the creation of one isosurface corresponding to a single scalar value.

The **Value** node, which has properties, is automatically added as a subnode of a specific isosurface node when `Single` is specified for the mode property of the isosurface.



Value Properties

Isovalue	The scalar value for the isosurface.

Clicking the right half of the *Isovalue* property activates a slider with which you can adjust the isovalue.



If you want a specific value, you can enter it into the **value** text box.



Using Multiple Isovalues (Range)

The range multiple-value mode for the isosurface involves the creation of multiple isosurfaces corresponding to a range of equally spaced scalar values.

The **Values (range)** node, which has properties, is automatically added as a subnode of a specific isosurface node when `Multiple (range)` is specified for the mode property of the isosurface.



Values (Range) Properties

Number of isosurfaces	The total number of equally spaced isosurfaces in the range.
Minimum isovalue	The lowest value in the range of isosurfaces.
Maximum isovalue	The highest value in the range of isosurfaces.

Using Multiple Isovalues (Arbitrary)

The arbitrary multiple-value mode for the isosurface involves the creation of multiple isosurfaces corresponding to a set of scalar values that you specify with numeric entries.

The **Values (arbitrary)** node, which has properties, is automatically added as a subnode of a specific isosurface node when `Multiple (arbitrary)` is specified for the mode property of the isosurface.



Values (Arbitrary) Properties

Isovalues	A list of scalar values for the multiple isosurfaces.

Clicking the property customizer (...) of the *Isovalues* property activates a *Values (arbitrary) - Isovalues* dialog.



Use the **Query** button to select a scalar range automatically. You can click **Auto Generate Values** to select the isovalues automatically as well.

Defining a Threshold

The threshold part isolates a group of cells based on their relationship to specified scalar values.

A threshold node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



A threshold part can also use a cell set as a scalar value. One application for this feature is a coolant drain problem, in which several pockets of coolant are left behind in the water jacket. A cell set threshold can help you determine the amount of fluid that is left in each pocket.

When created from a region, a threshold represents the boundaries of the cells containing data corresponding to the threshold criterion. When created from a boundary, it shows that boundary clipped to the specified threshold criterion. STAR-CCM+ allows the creation of a threshold part from one or more parent parts, provided the set of parent parts are of the same dimensionality.



Threshold Properties

You can use the properties that are listed below to adjust the specifications of the threshold.

Mode	The mode for threshold creation. These options relate to the *Range* property below.	
	Between	The threshold encompasses cells with all values

		between the minimum and the maximum of the *Range*.
	Above Max	The threshold encompasses cells with values greater than the maximum of the *Range* property.
	Below Min	The threshold encompasses cells with values less than the minimum of the *Range* property.
	Outside	The threshold encompasses cells with values less than the minimum and greater than the maximum of the *Range* property.
Scalar Field	The field function that is used to determine the threshold range.	
Range	A vector of two values, corresponding to the minimum and maximum of the threshold range.	
Parts	Selects the input parts as the source for deriving this threshold. See Selecting Input Parts.	

Threshold Pop-Up Menu

Edit Part in Current Scene	Executes the define threshold action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file. See Exporting to an External File.

See also: common menu items.

See Also

Using Threshold Parts and Cell Surfaces

Creating a Threshold

You can create a threshold with the pop-up menu.

To create a threshold, right-click the **Derived Parts** node and select **New Part > Threshold...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new threshold to appear in the simulation tree immediately. Define the part using the properties.

Editing a Threshold

Multiple techniques are available for editing a threshold.

You can edit a threshold using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Threshold In-Place Dialog

This dialog facilitates the definition of the threshold with interactive features.

If certain basic conditions are met when you start the process to create a threshold part or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the threshold.

Use the Scalar drop-down list to select the scalar field function that you want to use as the basis for the threshold. (The contents of the scalar list depend on the models that you selected.)

Use the Extraction Mode drop-down list to select the extraction mode from the choice of cells where the scalar values are:

- **All Between**—between a maximum and a minimum

- **All Above**—greater than a particular value
- **All Below**—less than a particular value
- **All Below Min & Above Max**—less than a minimum and greater than a maximum

In the *Scalar Range* group box, specify the scalar range:

- **min**—sets the lowest scalar value in the range
- **max**—sets the highest scalar value in the range
- **threshold value**—available only when **All Above** or **All Below** is selected as the extraction mode

  You can enter a number, or use the default value that is generated by STAR-CCM+ if you click the `Query` button.
- `Query` button—this instructs STAR-CCM+ to probe the selected scalar and fill in the values. STAR-CCM+ calculates this value as the average of the minimum and maximum values.

The units can be changed by typing them in after the values.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your threshold, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the threshold, which is shown in pink in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining a Scalar Warp Surface

The scalar warp surface moves the point coordinates on the surface of its input parts along the point normal vectors by an amount equal to the scalar value times the scale factor.

A scalar warp node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.

This feature is useful for showing flow profiles, mechanical deformation, or other desired user-defined functions. An example of the

latter is shown below, which corresponds to visualizing an approximation of material loss due to particle erosion in a bend.



Scalar Warp Surface Properties

You can use the properties that are listed below to adjust the specifications of the scalar warp surface.

Scalar Field	Scalar field function upon which the warp surface is based.
Scale Factor	Specifies the scale factor for the warp surface with a text entry.
Parts	Selects the input parts as the source for deriving this warp surface. See Selecting Input Parts.

Scalar Warp Surface Pop-Up Menu

You can use the options that are listed below to edit the properties of the scalar warp surface, or export its extracted data.

Edit Part in Current Scene	Executes the define warp surface action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file. See Exporting to an External File.

Creating a Scalar Warp Surface

You can create a scalar warp surface with the pop-up menu.

To create a scalar warp surface, right-click the **Derived Parts** node and select **New Part > Warp > Scalar Warp...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new warp surface to appear in the simulation tree immediately. Define the part using the properties.

Editing a Scalar Warp Surface

Multiple techniques are available for editing a scalar warp surface.

You can edit a scalar warp surface using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Scalar Warp Surface In-Place Dialog

This dialog facilitates the definition of the scalar warp surface with interactive features.

If certain basic conditions are met when you start the process to create a warp surface or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the warp surface. See Selecting the Input Parts.

Specify the characteristics of the warp surface with the following features:

- Scalar drop-down list—selects the scalar field function as the basis for the warp surface.
- *Scale Factor* text box—allows text entry of the scale factor.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Scalar Warp Surface Result

When you finish creating your scalar warp surface, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the warp surface, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display, and/or display its vectors with a new vector displayer.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Working with Vector Value-Based Parts

This section explains the derived parts that are based on vector values that you select.

Defining a Warp Surface

The warp surface is a part that modifies input surface point coordinates by moving points along the selected vector multiplied by the scale factor.

A vector warp node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



This feature is useful for showing flow profiles or mechanical deformation.



Warp Surface Properties

You can use the properties that are listed below to adjust the specifications of the warp surface.

Vector Field	Vector field function upon which the warp surface is based.
Scale Factor	Specifies the scale factor for the warp surface with a text entry.

Parts	Selects the input parts as the source for deriving this warp surface. See Selecting Input Parts.

Warp Surface Pop-Up Menu

Edit Part in Current Scene	Executes the define warp surface action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file. See Exporting to an External File.

Creating a Warp Surface

You can create a warp surface with the pop-up menu.

To create a warp surface, right-click the **Derived Parts** node and select **New Part > Warp > Vector Warp...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new warp surface to appear in the simulation tree immediately. Define the part using the properties.

Editing a Warp Surface

Multiple techniques are available for editing a warp surface.

You can edit a warp surface using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Warp Surface In-Place Dialog

This dialog facilitates the definition of the warp surface with interactive features.

If certain basic conditions are met when you start the process to create a warp surface or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the warp surface. See Selecting Input Parts.

Specify the characteristics of the warp surface with the following features:

- Vector drop-down list—selects the vector field function as the basis for the warp surface.
- *Scale Factor* text box—allows text entry of the scale factor.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your warp surface, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the warp surface, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display, and/or display its vectors with a new vector displayer.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Working with Streamlines

This section explains the derived parts that represent particle paths.

These parts are:

Defining a Streamline Part

The streamline part represents particle paths that you select within a region, computed from a vector field.

A streamline part node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



The streamline part may be computed from any vector field, although it generally makes the most physical sense to construct streamlines from a velocity field. STAR-CCM+ allows the creation of a streamline part from parent parts consisting of regions, or other parts which define a volume such as a threshold part. A streamline part has a seed object which defines the particle seed points. There are three types of seeds:

- A part seed, which takes its points from an existing part
- A line seed, which takes its points from a line
- A point seed, which allows the definition of either a single point or a cloud of points with a given radius

Streamlines are calculated using a second-order Runge-Kutta integration, accessible through RK nodes in the object tree, and can integrate both forward and backward. The integration is controlled by the following properties:

- *Maximum Propagation*—specifies the maximum arc length of a streamline defined in STAR-CCM+ internal SI units. The streamline computation is started at an initial position p0. The next position, p1,

is a small increment along the direction vector located at p0. Subsequent positions are similarly evaluated. While stepping tangential to the vector field, the distances between the positions are accumulated. This accumulated distance is the arc length of the streamline, and is the quantity that STAR-CCM+ compares to the Maximum Propagation value.

- *Max Steps*—specifies the maximum number of steps to execute in the integration.

A streamline terminates when either of these criteria are met. A streamline may also terminate because it exited the domain, or entered a stagnation region. Integration quality is controlled by the *Initial Integration Step* property which is the step-size that is used by the integration. If chosen to be much smaller than the average cell-length, the resulting streamline likely contains many more points than desired. Conversely if the step size is too large, cells may be skipped during the integration, resulting in a streamline of degraded quality.

In the following screenshot, particles are released from a line seed, and a line tool is used to assist with positioning the line endpoints.



Streamline Properties

You can use the properties that are listed below to adjust the specifications of the streamline.

Vector Field	Selects the type of vector field function to show, such as velocity or flow direction.	
Parts	Selects the input parts as the source for deriving this streamline part. See Selecting Input Parts.	
Seed Type	Specifies how the seed point is clustered and defined.	
	Point	Uses a single point or a cloud of points as the seed. This causes a point seed node to be added as a subnode.
	Line	Uses a line of points as the seed. This causes a line seed node to be added as a subnode.
	Part	Uses an existing surface, line, or point part as the seed. Regions cannot be selected. Adds a part seed node as a subnode.
Integration Solver	Selects from a choice of integration solvers.	
	2nd-Order RK	Selects the second-order integration solver. See Using the Second-Order Integrator.
	4th-Order RK	Selects the fourth-order integration solver. See Using the Fourth-Order Integrator.
	Adaptive-Step RK	Selects the adaptive streamline integrator. See Using the Adaptive Integrator.

Streamline Expert Properties

Rotation Scale	The scale factor for twist.

Streamline Pop-Up Menu

Edit Part in Current Scene	Executes the define streamline action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file. See Exporting to an External File.

Creating a Streamline

You can create a streamline with the pop-up menu.

To create a streamline, right-click the **Derived Parts** node and select **New Part > Streamline...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new streamline part to appear in the simulation tree immediately. Define the part using the properties.

Creating a Streamline on an Arc

It is possible to create a streamline on an arc or a circle through the following steps:

- Create a cylinder section. See Creating a Cylinder Section.
- Create a plane section to intersect with the cylinder section (by selecting the cylinder section as the input part). Depending on the angle at which the two sections intersect, the resulting plane section will be a circle, an ellipse or an arc. See Creating a Plane Section.
- Create streamlines, setting the seed type as `Part` and selecting the plane section as the seed part.

Editing a Streamline

Multiple techniques are available for editing a streamline.

You can edit a streamline using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Streamline In-Place Dialog

This dialog facilitates the definition of the streamline with interactive features.

If certain basic conditions are met when you start the process to create a streamline part or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the streamlines.

From the Vector Field drop-down list, select the vector type.

From the Seed Mode drop-down list, select one of three available types of seeds:

- Part seed (the default which is shown above)
- Point seed
- Line seed

The **Snap To Part** option, disabled in the preceding screenshot, becomes available when a tool is in use. When you activate this option by ticking its checkbox, the adjacent drop-down list allows you to select a part along which you can fix the location of the derived part. (This action is essentially snapping to the bounding box of the part.)

The units can be changed by typing them in after the values. Coordinates are relative to the laboratory (Cartesian) coordinate system, which cannot be changed.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your streamline, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the streamline part, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display. Finally, to display scalar data on this derived part effectively, use the techniques for displaying streamlines.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining a Constrained Streamline Part

Constrained streamlines are particle paths that are constrained to lie in a surface. Use them to simulate effects such as oil streaks.

This type of derived part is calculated using a fourth-order adaptive Runge-Kutta integration, accessible through an RK node in the object tree. Specifically, the constrained streamline results from integration of the velocity field near the surface of interest. This velocity field is modified by subtracting the dot product of the local velocity with the surface normal, and the result is projected onto the surface. Constrained streamlines can integrate both forward and backward in STAR-CCM+.

You can create a constrained streamline part from parent parts such as boundaries, or from surface-type parts such as sections. This type of derived part includes a seed object that defines the particle seed points.

Select the **constrained streamline** node to access the properties and pop-up menu of this derived part.



To control the integration, use the following properties:

- To specify the maximum length of the constrained streamline (in units of cell-length), use the *Maximum Propagation* property.
- To set the maximum number of steps to execute in the integration, use the *Max Steps* property.

  Since the integrator is adaptive, it adjusts the step-size during the integration to keep the error below a certain threshold.

A constrained streamline terminates when either of these criteria are met. This derived part also terminates when it exits the domain or enters a stagnation region.

Constrained Streamline Properties

You can use the properties that are listed below to adjust the specifications of the constrained streamline.

Parts	Selects the input parts as the source for deriving this constrained streamline.	
Seed Type	Specifies how the seed point is clustered and located.	
	Point	Uses a single point or a cloud of points as the seed. This option causes a point seed node to be added as a sub-node.
	Line	Uses a line of points as the seed. This option causes a line seed

		node to be added as a sub-node.
	Part	The surface, line, or point part that is used to seed. The number of lines that are generated depend on setting in the part seed node that gets added as a sub-node.
Vector Field	Selects the type of vector field function for analysis, such as velocity or flow direction.	
Integration Solver	Selects from a choice of integration solvers.	
	2nd-Order RK	Selects the second-order integration solver.
	4th-Order RK	Selects the fourth-order integration solver.
	Adaptive-Step RK	Selects the adaptive streamline integrator.

Constrained Streamline Expert Properties

Rotation Scale	A scale factor for twisting streamline ribbons.

Constrained Streamline Pop-Up Menu

Edit Part in Current Scene	Executes the creation of a constrained streamline, allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file.

See also: common menu items.

Creating a Constrained Streamline Part

You can create a constrained streamline with the pop-up menu.

Before you begin, initialize the solution to enable creation of this streamline.

To create a constrained streamline, right-click the **Derived Parts** node and select **New Part > Constrained Streamline...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new constrained streamline part to appear in the simulation tree immediately. Define the part using the properties.

Editing a Constrained Streamline Part

Multiple techniques are available for editing a constrained streamline.

You can edit a constrained streamline using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Constrained Streamline In-Place Dialog

This dialog facilitates the definition of the constrained streamline with interactive features.

If certain basic conditions are met when you start the process to create a constrained streamline part or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the streamlines.

From the Seed Mode drop-down list, select one of three available types of seeds:

- Part seed (the default which is shown above)
- Point seed

- Line seed

The **Snap To Part** option, disabled in the preceding screenshot, becomes available when a tool is in use. When you activate this option by ticking its checkbox, the adjacent drop-down list allows you to select a part along which you can fix the location of the derived part. (This action is essentially snapping to the bounding box of the part.)

The units can be changed by typing them in after the values. Coordinates are relative to the laboratory (Cartesian) coordinate system, which also cannot be changed.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your constrained streamline, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the constrained streamline, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display. Finally, to display scalar data on this derived part effectively, use the techniques for displaying streamlines.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Working with Seed Objects

This section describes the seed objects that help define a streamline or constrained streamline part.

Using the Part (Source) Seed

The part (source) seed takes its points from one or more existing parts, therefore it is certain to create a streamline or constrained streamline that lies directly on the part surfaces. This type of seed is the recommended option.

The **Source Seed** node, with its own properties, is automatically added as a subnode of a specific streamline or constrained streamline derived part node when `Part` is specified for the seed type property of the streamline.



The Seed Mode drop-down list contains this option in the in-place dialog of the streamline or constrained streamline.



Click the **Seed Parts** button to select the parts for the seed. Typically the best results are obtained when the part seed is the same as the input part.

For streamlines, the Direction drop-down list lets you specify the direction of integration from the seed location: forward, backward or both.

Specifying Part Resolution in the U- and V- Directions

The **Part U-Resolution** and **Part V-Resolution** text boxes let you specify the maximum number of points extracted from the part in the u- and v- directions, respectively. The u- and v- directions are part of the computation of an average plane, based on the specified parts, for the seed source to use.

The u-direction is the vector in this plane that is most closely aligned with the global, Cartesian x-axis. If the average plane is an x-y plane (perpendicular to the z-axis), u corresponds exactly to global x. If the average plane is skewed from x-y, the u-axis points toward the global x-axis. The y-axis is then the cross-product of the computed x-axis and the normal. For the x-y plane, this y-axis is exactly in the y-direction.

Once the average plane for the seed source is computed, the plane size is determined from the projection of the seed source onto the average plane. Suitable u- and v- directions are computed, then the plane is divided into the specified number of points along the u- and v- directions. The points are then projected back onto the boundary and used as the seed points for the part.

Note: This feature is only designed to work with planar or reasonably planar boundaries. Therefore, if you are working with a curved surface, STAR-CCM+ attempts to project it to the average plane, but accuracy is

limited. Accuracy decreases with the amount of curvature in the surface that you project.

Source Seed Properties

Seed Parts	Selects the parts to be used as the seed. Regions are not included among the options.	
On Ratio	The ratio of available vertices to the maximum numbers of seed points specified for the u- and v-directions. This value is automatically calculated by STAR-CCM+. It is recalculated each time you change the u or v points or the *Seed Parts* property.	
Randomize	Use existing vertices or a random pattern for the seed points. It is useful to randomize if the *Part* uses a highly unstructured grid or mixture of fine and coarse grids.	
	Activated	Use a random pattern of points.
	Deactivated	Use the vertices of the part.
N Grid Points	The maximum numbers of seed points to use in the u- and v- directions. If the number of vertices on the part is less than either value, the number of vertices is used.	

Using the Point Seed

The point seed allows the definition of either a single point or a cloud of points with a given radius.

The **Point Seed** node, with its own properties, is automatically added as a subnode of a streamline or constrained streamline derived part when `Point` is specified for the seed type property of the streamline.



Use this option only when you have a particular set of points, which you could obtain by cutting a surface with a plane or drawing an isosurface.

The Seed Mode drop-down list contains this option in the in-place dialog of the streamline or constrained streamline.



The Seed Position column provides text boxes for you to specify the x, y, and z coordinates for the point.

A drop-down list selects from among the coordinate systems existing in the simulation (Laboratory by default).

The **Seed Radius** text box lets you specify the size of the cloud of points. It should remain at the default value of zero if you use only one point. The units can be changed by typing them in after the values.

Use the **Number of Points** text box to determine whether you will use a single point or a cloud of them.

The **Display Tool** checkbox, when activated, lets you use the point tool to position the seed visually.

The Direction drop-down list lets you specify the direction of integration from the seed location: forward, backward or both.

Point Seed Properties

Center	A Cartesian vector specifying the position of the point cloud.
Radius	The radius of the point cloud.
Resolution	The number of points in the point cloud.

Using the Line Seed

The line seed takes its points from a line. As with the point seed, you should have a specific set of points to use this option. Moreover, the line seed would be effective only on a perfectly flat surface.

The **Line Seed** node, with its own properties, is automatically added as a subnode of a specific streamline or constrained streamline derived part when `Line` is specified for the seed type property of the streamline.



The Seed Mode drop-down list contains this option in the in-place dialog of the streamline or constrained streamline.



The point 1 column provides text boxes for you to specify the x, y, and z coordinates for the first point of the line.

The point 2 column provides text boxes for you to specify the x, y, and z coordinates for the second point of that line, defining its length and direction.

A drop-down list selects from among the coordinate systems existing in the simulation (Laboratory by default).

The **Resolution** text box lets you specify the number of evenly spaced points along the line.

The **Display Tool** checkbox, when activated, lets you use the line tool to position the seed visually.

The Direction drop-down list lets you specify the direction of integration from the seed location: forward, backward or both.

Line Seed Properties

Line Point 1	A Cartesian vector specifying the position of one end point of the line.
Line Point 2	A Cartesian vector specifying the position of the other end point of the line.
Resolution	The number of evenly spaced points along the line.

Working with Runge-Kutta Objects

This section describes the Runge-Kutta (RK) integration solvers for either a streamline or constrained streamline part.

Using the Second-Order Integrator

This is the second-order Runge-Kutta integration solver.

The **2nd-Order Integrator** node, with its own properties, is automatically added as a subnode of a specific streamline or constrained streamline derived part node when `2nd-Order RK` is specified for the *Integration Solver* property of the derived part.



Using the Fourth-Order Integrator

This is the fourth-order Runge-Kutta integration solver.

The **4th-Order Integrator** node, with its own properties, is automatically added as a subnode of a specific streamline or constrained streamline derived part node when `4th-Order RK` is specified for the *Integration Solver* property of the derived part.



Using the Adaptive Integrator

This is the adaptive-step Runge-Kutta integration solver.

The **Adaptive Integrator** node, with its own properties, is automatically added as a subnode of a specific streamline or constrained streamline derived part node when `Adaptive-Step RK` is specified for the *Integration Solver* property of the derived part.



Adaptive Integrator Properties

Integration Direction	Specifies the direction of integration from the seed location.	
	Forward	The integration proceeds in the forward direction (for instance with the flow if velocity is the vector field).
	Backward	The integration proceeds in the backward direction (for instance against the flow if velocity is the vector field).
	Both	The integration proceeds both forward and backward from the seed location.

Adaptive Integrator Expert Properties

Initial Integration Step	The initial integration step size in units of cell length.
Maximum Error	The maximum allowable error in integration before step size is adjusted.
Maximum Integration Step	The maximum integration step size in units of cell length.
Minimum Integration Step	The minimum integration step in units of cell length.
Maximum Propagation	The maximum length of the streamlines in units of cell length.
Max Steps	The maximum number of integration steps.

Integrator Properties

You can adjust various features of the integration solvers with the object properties.

Integration Direction	Specifies the direction of integration from the seed location.	
	Forward	The integration proceeds in the forward direction (for instance with the flow if velocity is the vector field).
	Backward	The integration proceeds in the backward direction (for instance against the flow if velocity is the vector field).
	Both	The integration proceeds both forward and backward from the seed location.

Integrator Expert Properties

Initial Integration Step	The initial integration step size in units of cell length. Reducing the value creates a higher resolution for the streamlines.
Maximum Propagation	The maximum length of the streamlines in units of cell length.
Max Steps	The maximum number of integration steps.

Working with (m', theta) Warps

The derived parts that are described in this section unwrap results to a flat plane.

Defining an (m', theta) Warp

In STAR-CCM+, the (m', theta) warp is used to present a flattened surface of revolution.

The **(m', theta) warp** node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



For turbomachinery, such surfaces are blade-to-blade. They usually originate from a cylindrical plane at constant radius, or an isosurface of the block-mapped coordinate that varies from hub to case.

This derived part requires three pieces of information in its properties:

- The *Origin* of the axis of rotation
- The vector specifying the *Rotation Axis*
- A vector specifying the *Tangential Axis*—this axis is used to compute the angular coordinate.

The m′ coordinate is the arc length along the surface projected into the (z,r) plane and normalized by the local radius (Drela and Youngren):

$$m' = \int \left(\frac{\sqrt{dr^2 + dz^2}}{r} \right) \tag{2708}$$

- z is the axial coordinate along the axis of rotation.
- r is the radial coordinate perpendicular to the axis of rotation.

Theta is calculated as the angle from the tangential axis:

$$\theta = 2\pi + \arctan\left(\frac{-y}{-x}\right) \tag{2709}$$

where y is the distance in any direction from the plane that is formed by the tangential and rotation axes, x is the tangential axis, and z is the rotation axis.



For example, if you take the points of the surface and project them into the (z,r) plane, you can expect them to fall along the defining line of the surface (the generating curve). This result is depicted in the following diagram.



The meridional coordinate is defined as:

$$m = \sqrt{\left(r^2 + z^2\right)} \tag{2710}$$

The *Tangential Axis* property does not have to be orthogonal to the *Rotation Axis* since only the component orthogonal to the rotation axis is used.

The following example uses the (m′, theta) derived part:



- The rotation axis is [0, 0, 1] (aligned with the z-axis).
- The origin is upstream of the surface at [0, 0, -10].
- The tangential axis, [0, 1, 0], is orthogonal to the rotation axis.
- The following screenshot shows the view down the z-axis. The derived part maps the m′ coordinate to x and the theta coordinate to y.



Limitations

- The derived part is computed by determining the generating curve from the part. So if there are multiple parts in a displayer, they must form a continuous surface for the derived part to be valid.

(m', theta) Bibliography

Drela, M. and Youngren, H. 2008. "A User's Guide to MISES 2.63", MIT Aerospace Computational Design Laboratory.

(m', theta) Warp Properties

You can use the properties that are listed below to adjust the specifications of the (m', theta) warp.

Coordinate System	A list of coordinate systems in the simulation for use by the (m', theta) warp.
Rotation Origin	Sets the rotation origin for an (m', theta) warp.
Rotation Axis	Sets the rotation axis for an (m', theta) warp.
Tangential Axis	Sets the tangential axis for an (m', theta) warp.
Parts	Selects the input parts as the source for deriving this (m', theta) warp.

(m', theta) Warp Pop-Up Menu

Export...	Exports the extracted data of the part to an external STL file.

Creating an (m', theta) Warp

You can create an (m', theta) warp with the pop-up menu.

To create an (m′, theta) warp, right-click the **Derived Parts** node and select **New Part > (m′, theta) warp**.



This action causes the new (m′, theta) warp part to appear in the simulation tree immediately. Define or edit the part using the properties.

Viewing the Final Result

When you finish creating your (m', theta) warp, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the (m', theta) warp part, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Working with Separation/Attachment Parts

These derived parts detect where flow separates from or attaches to a surface.

The available parts are:

- Separation/attachment line

Defining a Separation/Attachment Line

Separation/attachment lines are lines along which flow separates from or attaches to a surface.

This phenomenon is induced by viscous effects that take place in direct proximity to the point of separation or attachment. When a fluid flows over a no-slip wall boundary, the velocity magnitude goes to zero as one approaches the surface along a normal direction. Therefore, the extraction of separation and attachment lines focuses on the non-zero, tangential shear-stress vector field defined over the surface that exhibits the same flow patterns as nearby streamlines.

The following image shows how STAR-CCM+ uses the wall shear stress vector field to identify separation lines (the line is shown in magenta for clarity). However, several other parameters are used to identify a separation line, and you can also use cell relative velocity as the vector field. The cell relative velocity is the fluid flow velocity at the cell centroid relative to the cell centroid due to motion.



The **separation/attachment line** node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



Once you have extracted a separation/attachment line, you can display it in a graphics scene with a scalar color, if desired. You can also use it as the input part for a streamline displayer, so that the streamlines begin along the separation line.

Separation/Attachment Lines Properties

Display Attachment Lines	Toggles whether lines displaying flow attachment are displayed or not.
Display Separation Lines	Toggles whether lines displaying flow separation are displayed or not.
Vector Field	Specifies the vector field that is used by the algorithm to extract the separation/attachment line. This must be set to wall shear stress or cell relative velocity. If you use a different three-dimensional vector field, it does not produce valid results.
Line Quality	Filters lines where the angle between the extracted separation/attachment line and the flow vector are below the user-defined value. The quality is entered on a scale of 0.0 to 90.0 degrees.

Line Strength	Sets the strength of the separation/attachment. This setting is used to filter out weak results. Line strength is not a physical quantity, but a filtering property. The strength is entered on a scale of 0–1, where 0 includes all separation/attachments, and 1 includes only the strongest results.
Minimum Line Length	Sets the minimum separation/attachment line length. Separation/attachment lines which are shorter than this value are eliminated.
Parts	Selects the input parts as the source for deriving this separation/attachment line. Typically, the input part is a no-slip wall boundary.

Separation/Attachment Lines Pop-Up Menu

Edit Part in Current Scene	Executes the define separation/attachment action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file.

Creating a Separation/Attachment Line

You can create a separation/attachment line with the pop-up menu.

To create a separation/attachment line:

- Right-click the **Derived Parts** node and select **New Part > Separation/Attachment Line...**



Typically, an in-place dialog is used to edit the part, which appears as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, a new separation/attachment line node is added to the simulation tree immediately. You can then define the part using its properties, or open a graphics scene and edit the part

Editing a Separation/Attachment Line

Multiple techniques are available for editing a separation/attachment line.

To edit a separation/attachment line:

- Right-click a separation/attachment line node in the simulation tree and select **Edit Part in Current Scene**.



Typically, an in-place dialog is used to create the part, which appears as soon as you execute this pop-up menu action.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can click the separation/attachment line node and use its properties.

Using the Separation/Attachment Line In-Place Dialog

This dialog facilitates the definition of the separation/attachment line with interactive features.

If certain conditions are met when you create a separation/attachment line or edit an existing one, an in-place dialog appears in the *edit* tab of the *Simulation* window.



Vector Field selects the vector field used by the algorithm to calculate the separation/attachment line. This must be either wall shear stress or cell relative velocity. If you select another vector field, a separation/attachment line is still created, but the result will be invalid.

Input Parts selects the boundaries that are used to generate the separation/attachment line. Typically, this part is a no-slip wall, over which fluid is flowing.

The *Display Attachment Lines* and *Display Separation Lines* checkboxes specify whether attachment or separation lines are included in the derived part. Attachment occurs when flow meets a wall boundary and flows along the surface. Separation occurs when flow leaves the surface.

Line Strength does not represent a physical quantity, but is used by the code to filter out weak results (that is, lines that do not represent a strong flow separation).

Minimum Line Length excludes lines from the derived part that are shorter than the user-defined length.

Line Quality is the angle between the separation/attachment line and the flow vector, which filters lines from the derived part that are at a large angle to the flow.

The Display group box and the Create button are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your separation/attachment line, you can view it in the assigned scene.

Once you have created a separation/attachment line, you can rotate and zoom in the display to get a clear view of the part. If you did not add the separation/attachment line to the display using the Display group box, you can add the part manually.

You can display separation/attachment lines in a scene using geometry or streamline displayers. With streamline displayers, you can display the part as a tube or ribbon, or color it using a scalar.

The following image shows a separation line that is displayed using a streamline displayer. The thickness of the line has been increased to make it more prominent on the geometry surface. The velocity vectors illustrate where flow separation occurs, which matches the separation line location.



You can also use separation/attachment lines as a seed part for a streamline derived part. This can be used to analyze the flow path after the point of flow separation, which is illustrated in the example below.



You can remove this part from the display by using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Calculating the Surface Area of a Part

When a volume representation is active, the surface area of a part such as a plane section, a cylinder surface, a sphere, or even an isosurface can be calculated using a surface integral report with a value of 1.

To calculate the surface area:

1. Create a user-defined scalar field function with the constant value of 1. Such a function is shown below.



2. Create a surface integral report using the right-click menu of the top-level Report node. Set the *Parts* property of the new report node to the part whose surface area you require, and set the *Scalar Field Function* property to the user-defined field function created in the previous step. An example of such a report is shown below.



3. Right-click the newly created report node and choose **Run Report**. The surface area of the part is displayed in the *Output* window.

Working with Vortex Cores

This section describes derived parts that detect vortex cores in a flow field.

Defining a Vortex Core Part

In STAR-CCM+, the vortex core part automatically extracts vortex cores in a flow field.

Vortices are features in the flow data which are of great importance in fluid dynamics, both in theoretical and practical research. In some cases, vortices (turbulence) have to be impelled, for example to stimulate mixing of fluids, or to reduce drag. In other cases, vortices have to be prevented, for example around aircraft, where they can reduce lift.

The vortex core is the axis line around which the flow field rotates. The core provides a precise position of the vortex and also allows you to distinguish between different vortical structures close to each other.

Vortex core visualization belongs to the category of feature-based flow visualization techniques. The methods and algorithms in this category present the flow data in a higher level of abstraction. The flow data is described by a set of interesting and important objects or structures in the data called features. In fluid dynamics vortices, shock waves, separation lines, and attachment lines are examples of such features. Since only a small percentage of the data is of interest, and the features can be described compactly, it is possible to achieve considerable data reduction. This enables an interactive visualization of large amounts of fluid simulation data.

The **vortex core** node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



Once you have extracted a vortex core, you can display it on a geometry scene or as parts in streamline displayers. You can also color-code vortex cores in a scalar scene.

Understanding How Vortex Cores Work

The vortex core extraction algorithm needs the velocity field and a three-dimensional region to generate poly lines representing vortex cores. It has three filtering parameters to filter out unwanted vortex cores.

There are many different definitions of vortices and likewise many different vortex detection algorithms. These definitions are based on certain field values being high or low, like high vorticity magnitude, high helicity magnitude, or low-pressure value. However, there is no accepted threshold value for what is considered high or low.

Definitions of features are usually implicit because they are procedural. However, precise mathematical definitions are attainable with the parallel vectors operator proposed by Peikert and Roth (1999, 2000). This operator is an elementary operation on two vector fields, which in turn make it possible to define and calculate global line-type features of vector or scalar fields. The parallel vectors operator can be extended to unsteady fields. Advantages include:

- Comparing feature definitions
- Reusing algorithms and implementations

The parallel vectors operator causes the feature definition to work separately from the actual implementation. According to Roth and Peikert, along with various line-type features, the operator can implement algorithms to extract the following:

- Vortex core lines
- Lines of separation or attachment

This version of STAR-CCM+ uses the parallel vectors operator to implement the Eigenvector method for vortex core extraction proposed by Sujudi and Haimes (1995, 1999). According to these authors, a rotational motion is present when the Jacobian (velocity gradient matrix) has a conjugate complex pair of Eigenvalues. The description of how the parallel vectors operator can be applied to the Eigenvector method is in the papers by Peikert and Roth (1998, 1999) as well as in the dissertation by Roth.

Vortex Core Properties

Vortex Quality	Computes the angle between the extracted core line and the velocity vector at this position. Small angles mean high quality; this may not apply to all simulation data.

Vortex Strength	Sets the strength of the rotational motion. This setting is used to filter out weak vortices.
Minimum Core Length	Sets the minimum length of a line to be used by a vortex core. Below this setting, the core lines are eliminated.
Vector Field	The vector field that is used by the algorithm to extract the vortex core. This must be set to velocity. If you use a different three-dimensional vector field, this would not lead to a valid result.
Parts	Selects the input parts as the source for deriving this vortex core.

Vortex Core Pop-Up Menu

Edit Part in Current Scene	Executes the define vortex core action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file.

Creating a Vortex Core Part

You can create a vortex core with the pop-up menu.

To create a vortex core, right-click the **Derived Parts** node and select **New Part > Vortex Core...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new vortex core part to appear in the simulation tree immediately. Define the part using the properties.

Editing a Vortex Core Part

Multiple techniques are available for editing a vortex core.

You can edit a vortex core using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Vortex Core In-Place Dialog

This dialog facilitates the definition of the vortex core with interactive features.

If certain basic conditions are met when you start the process to create a vortex core part or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the Vector Field button to select the vector field to be used by the algorithm to extract the vortex core.

Click the **Input Parts** button to select the input parts that you want to include in the vortex cores.

Specify the **Vortex Core Options**. These include the following:

- Vortex strength
- Vortex quality
- Minimum core length

See Defining a Vortex Core Part.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your vortex core, you can view it in the assigned scene.

Once you have finished creating a vortex core part, you can rotate and zoom in the display to get a clear view of the vortex core. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.

You can display vortex cores on a geometry scene or as Parts in streamline displayers. Within streamline displayers, you can display them as lines, ribbons, or tubes and color code them by a scalar as shown below.



The vortex core can be used to seed streamlines in order to visualize the vortical structure as shown in the figure below.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Vortex Core Bibliography

This section contains bibliographic references related to vortex cores.

Haimes, R. and Kenwright, D. N. 1999. "On the velocity gradient tensor and fluid feature extraction", AIAA 14th Computational Fluid Dynamics Conference, Paper 99-3288.

Peikert, R. and Roth, M. 1999. "The 'parallel vectors' operator -- a vector field visualization primitive", Proceedings of the Conference on Visualization '99, San Francisco, California, USA, pp. 263–270.

Roth, M. 2000. "Automatic Extraction of Vortex Core Lines and Other Line-Type Features for Scientific Visualization", PhD thesis, Swiss Federal Institute of Technology, Zurich, Switzerland.

Roth, M. and Peikert, R. 1998. "A higher-order method for finding vortex core lines", Proceedings of the Conference on Visualization '98, IEEE Computer Society Press, pp. 143–150.

Sujudi, D. and Haimes, R. 1995. “Identification of swirling flow in 3-D vector fields”, AIAA 12th Computational Fluid Dynamics Conference, Paper 95-1715.

Working with Probes

This section explains the derived parts that take samples of data in a region based on coordinates or intervals that you specify.

These parts are:

- Point
- Line-probe
- Presentation grid
- Arbitrary probe

How data for a particular point is extracted depends on your selection of input parts.

Input Parts Type	Reported Data Value
Region	Cell value
Region where Probe Point is outside of Region	Meaningless value (Probe Point is discarded)
Boundary	Surface cell value of the face that is closest in Eucledian distance to the Probe Point
Region and Boundary	Cell value or surface cell value whichever is closest in Eucledian distance to the Probe Point
Region and Boundary where Point is outside of Region	Surface cell value of the face that is closest in Eucledian distance to the Probe Point.

Defining a Point

The point part samples its input part at a set of coordinates in a region.

A point node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



STAR-CCM+ allows the creation of a point part from parent parts consisting of regions and boundaries only. The point location is specified by entering the coordinates of the point. A point tool is available to assist with positioning the point graphically. The point is represented in the scene by a sphere of constant radius in screen coordinates.



Point Properties

You can use the properties that are listed below to adjust the specifications of the point.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part.	

Parts	Selects the input parts as the source for deriving this point.	
Follow Motion	Controls whether the point probe moves together with a moving mesh or not.	
	Activated	The point probe attaches itself to the moving mesh (if present) and follows its motion.
	Deactivated	The point probe remains stationary (assuming that its assigned coordinate system is not moving).
Point	A Cartesian position vector that defines the location of the point.	

Point Pop-Up Menu

Edit Part in Current Scene	Executes the define point action allowing you to set the properties of this part interactively. This option is only available if a scene is open.	
Export...	Exports the extracted data of the part to an external STL file.	

See also: common menu items.

Creating a Point Part

You can create a point part with the pop-up menu.

To create a point part, right-click the **Derived Parts** node and select **New Part > Probe > Point...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new point part to appear in the simulation tree immediately. Define the part using the properties.

Using a Point Part on a Moving Mesh

Point parts can be used to monitor data on a moving mesh.

The point attaches itself to the moving mesh and keeps track of the cell it is in and its relative position within that cell. Every time the mesh moves, the point checks its position to ensure that it is inside the correct cell and in the correct position within that cell.

To enable this feature, select the point probe node on the STAR-CCM+ object tree, and tick the *Follow Motion* checkbox in the *Properties* window.



The original position of the probe can be seen below.



As the mesh rotates, the probe moves with it.



Note: The principles described above apply to all types of motion. For example, a point part could also be used to monitor data on a morphing mesh.

Editing a Point Part

Multiple techniques are available for editing a point part.

You can edit a point part using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Point In-Place Dialog

This dialog facilitates the definition of the point part with interactive features.

If certain basic conditions are met when you start the process to create a point or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the point.

The point tool positions the point visually. The **Snap To Part** option lets you select a part along which you can position the point. (This action is essentially snapping to the bounding box of the part.)

The Point column provides text boxes for you to specify the x-, y-, and z-coordinates of the point.

The Lock column contains checkboxes that you can tick to fix the position of any coordinate or combination of them.

A drop-down list selects from among the coordinate systems existing in the simulation (Laboratory by default).

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your point part, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the point, which is shown in red in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Extracting Data at a Point Probe

You can use reports to extract data from point probes.

To extract data at Point Probes previously specified, create a Report of the type Maximum or Minimum (see Reporting Results). When the

Smooth Values expert option in the *Properties* window of the report is selected, vertex values for the cell are interpolated to the Point Probe.

Defining a Line-Probe

The line-probe part samples its data from evenly spaced points along a line in a region.

A line-probe node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



STAR-CCM+ allows the creation of a line-probe part from parent parts consisting of regions and boundaries only. The line location is specified by entering the coordinates of the line. A line tool is available to assist with positioning the line endpoints graphically. A line-probe can be incorporated into an XY plot.



The line-probe can be useful for extracting data along a line in a region, and can be used to display scalars or vectors. The most common use of data extracted with a line-probe is selecting it for an XY plot.



Line Probe Properties

You can use the properties that are listed below to adjust the specifications of the line-probe.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part.
Parts	Selects the input parts as the source for deriving this line-probe.
Point 1	A Cartesian position vector defining one end of the line.
Point 2	A Cartesian position vector defining the other end of the line.
Resolution	The number of evenly spaced segments along the line. The actual number of points in the probe is one greater than the number of segments; fewer points may be generated if part of the probe lies outside the model geometry.

Line Probe Pop-Up Menu

Edit Part in Current Scene	Executes the define line probe action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file.

Creating a Line-Probe

You can create a line-probe with the pop-up menu.

To create a line-probe, right-click the **Derived Parts** node and select **New Part > Probe > Line...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new line-probe to appear in the simulation tree immediately. Define the part using the properties.

Editing a Line-Probe

Multiple techniques are available for editing a line-probe.

You can edit a line-probe using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Line-Probe In-Place Dialog

This dialog facilitates the definition of the line-probe with interactive features.

If certain basic conditions are met when you start the process to create a line-probe or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the point.

The point 1 and point 2 columns provide text boxes for you to specify the x-, y-, and z-coordinates of the first and second points of the line, respectively.

A drop-down list selects from among the coordinate systems existing in the simulation (Laboratory by default).

The **Resolution** text box lets you set the number of evenly spaced segments along the line. The overall number of points is one greater than this; fewer points may be generated if part of the probe lies outside the model geometry.

The line tool positions the line visually. The **Snap To Part** option lets you select a part along which you can position the line. (This action is essentially snapping to the bounding box of the part.)

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your line-probe, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the line-probe, which is shown in brown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining a Presentation Grid

The presentation grid part samples data from regularly spaced intervals on a finite plane in a region.

This part is similar to the plane section, the difference being that the output points of a presentation grid are sampled on a regular grid, whereas points from the plane section are based on the topology and discretization of the underlying mesh. A presentation grid node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



STAR-CCM+ allows the creation of a presentation grid part from parent parts consisting of regions and boundaries only. The finite plane can be defined by specifying a point for the origin, and two other points which, together with the origin define two axes for the plane. An explicit plane tool is available to assist with positioning the plane graphically. The presentation grid can be useful for extracting data at structured intervals in a region.



Presentation Grid Properties

The following is a reference for presentation grid properties. A detailed explanation of the properties and their use is provided in another section of the User Guide.

Coordinate System	A list of coordinate systems in the simulation for use by this derived part.
Parts	Selects the input parts through which the grid is placed.
Normal	A Cartesian unit vector defining the normal to the plane of the grid. Use this property as an alternative to *Point 1* and *Point 2* to define the plane. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).
Origin	The origin of the grid. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).
Point 1	A Cartesian position vector specifying the extent of the first axis. Use this property along with *Point 2* as an alternative to *Normal* to define the plane. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).
Point 2	A Cartesian position vector specifying the extent of the second axis. Use this property along with *Point 1* as an alternative to *Normal* to define the plane. You can include units in the entry, or select the units in the dialog that is activated by clicking the property customizer button (...).
X Resolution	The number of evenly spaced points along the axis direction from *Origin* to *Point 1*.
Y Resolution	The number of evenly spaced points along the axis direction from *Origin* to *Point 2*.

Presentation Grid Pop-Up Menu

Edit Part in Current Scene	Executes the define presentation grid action allowing you to set the properties of this part interactively. This option is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file.

Creating a Presentation Grid

You can create a presentation grid with the pop-up menu.

To create a presentation grid, right-click the **Derived Parts** node and select **New Part > Probe > Presentation Grid...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new presentation grid to appear in the simulation tree immediately. Define the part using the properties.

Editing a Presentation Grid

Multiple techniques are available for editing a presentation grid.

You can edit a presentation grid using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Presentation Grid In-Place Dialog

This dialog facilitates the definition of the presentation grid with interactive features.

If certain basic conditions are met when you start the process to create a presentation grid or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.


Servers
intake
simulation
scene/plot
edit
Create Presentation Grid
Input Parts
[Fluid]
Plane Parameters
center
normal
X -0.53339994 m 1 m
Y 1.26825886 m 0 m
Z 0 m 0 m
Coordinate System
Laboratory
origin
point 1
point 2
X -0.53339994 m -0.53339994 m -0.53339994 m
Y -0.00063551 m 2.53715324 m -0.00063551 m
Z -2.54000258 m -2.54000258 m 2.54000258 m
X Resolution 20
Y Resolution 20
Snap to Part
plane section
Display
No Displayer
New Geometry Displayer
New Scalar Displayer
New Vector Displayer
Existing Displayer
Outline 1
Create
Close
Help

Click the **Input Parts** button to select the input parts that you want to include in the presentation grid.

You can use the plane tool in the scene display, which is activated automatically with the dialog, to position the plane visually. Also, you can set up the presentation grid with the following specifications in the *Plane Parameters* group box:

- center—sets x, y, and z coordinates for the center. You can include units in the entry.
- normal—sets x, y, and z coordinates for the normal. You can include units in the entry.
- A drop-down list that selects from among the coordinate systems existing in the simulation (Laboratory by default).
- origin—the x, y, and z coordinates for this are generated by STAR-CCM+ automatically based on your entries for the center and normal. This ensures that a proper plane is created.
- point 1 and point 2—these are generated by STAR-CCM+ automatically based on your entries for the center and normal. This ensures that a proper plane is created.

  The automatic calculations of point 1, point 2 and the origin, as well as the visual use of the plane tool, are designed for convenience. If you prefer to choose the proper combination of values yourself, finish creating an approximate presentation grid in the dialog to add its node to the object tree. Then you can enter these values directly in the properties of the node.
- **X Resolution** and **Y Resolution**—these set the resolution of the grid in the x-direction and y-direction, respectively. An increased resolution can, for example, show more vectors within a given area.
- **Snap To Part**—selects a part along which you can position the plane. (This action is essentially snapping to the bounding box of the part.)

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Using the Presentation Grid Properties

You can use the presentation grid properties to edit the size and position of a presentation grid, as well as the X-Y resolution and associated parts.



Editing the Grid Resolution

X Resolution is the number of intervals between data points along the x-axis (that is, Number of Data Points = X Resolution + 1). The x-axis being the axis between the *Origin* and *Point 1*.

Y Resolution is the number of intervals between data points along the y-axis. The y-axis being the axis between the *Origin* and *Point 2*.

For example, the following grid has an *X Resolution* of `5` and a *Y Resolution* of `10`.



Editing the Grid Size and Position

Coordinate System defines the coordinate system to which *Point 1*, *Point 2*, *Normal*, and *Origin* are relative. If using a cylindrical or spherical coordinate system, the units for position vectors update accordingly.

Note: Using a cylindrical or spherical coordinate system does not create a cylindrical or spherical presentation grid.

The diagram below illustrates the *Point 1*, *Point 2*, *Normal* and *Origin* properties on the presentation grid.



Point 1 defines the lower-right corner on the presentation grid, relative to the selected coordinate system. Entering a new value for *Point 1* automatically adjusts the *Normal*, in order to define a grid that matches the current properties.

Point 2 defines the upper-left corner on the presentation grid, relative to the selected coordinate system. Entering a new value for *Point 2* automatically adjusts the *Normal*, in order to define a grid that matches the current properties.

Origin defines the lower-left corner on the presentation grid, relative to the selected coordinate system. Entering a new value for *Origin* automatically adjusts the *Normal*, in order to define a grid that matches the current properties.

Normal defines the vector passing through the center of the presentation grid at a right-angle. Entering a new value for *Normal* automatically adjusts the values for *Point 1*, *Point 2* and *Origin*, in order to define a grid that matches the current properties.

Note: The *Normal*, *Point 1*, *Point 2* and *Origin* properties override one another according to the most recently entered value.

Editing the Associated Parts

Parts define the association of a presentation grid to selected geometry parts or regions. For example, if a presentation grid spans two parts, selecting only one part will ignore all data points in the second part.



Editing the *Parts* property opens the object selection dialog.



Viewing the Final Result

When you finish creating your presentation grid, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the presentation grid, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Defining an Arbitrary Probe

The arbitrary probe is an implicit part that samples its input part or region from a set of specific points.

An arbitrary probe node has its own properties and pop-up menu, and it is contained in the **Derived Parts** node of the simulation object tree.



An arbitrary probe works in a manner similar to the other probe derived parts, except that it extracts its sample points from either a CAD part surface or from an STL or VTK file.



Arbitrary Probe Properties

Parts	Selects the input parts or regions as the source for deriving the field values at the points of the probe.	
Surface Input Option	Specifies the source of the input geometry.	
	File	Reads in points from a STL or VTK file.
	Part Surfaces	Select one or more CAD part surfaces from which to extract the sample points.

Arbitrary Probe Pop-Up Menu

Edit Part in Current Scene	Executes the define arbitrary probe action allowing you to set the properties of this part interactively. This is only available if a scene is open.
Export...	Exports the sample points of the probe to an STL or VTK file. An exported VTK file will also contain all the field values visualized at the probe points up to the moment the probe was exported.
Re-Import Data...	Loads new sample points from an STL or VTK file.

Creating an Arbitrary Probe

You can create an arbitrary probe with the pop-up menu.

To create an arbitrary probe, right-click the **Derived Parts** node and select **New Part > Probe > Arbitrary...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action. However, if the conditions for activating this dialog are not met, then this action causes the *Select Input Surfaces* dialog to appear instead. This dialog allows you to extract sample points from one or more CAD part surfaces or from an STL or VTK file.



The in-place dialog also includes this choice of inputs, which is the first step in creating this derived part.

To work with part surfaces:

- Select the **Part Surfaces** option.

  The dialog will change to allow you to activate the object selector.
- Make your selections and close the object selector.
- In the *Select Input Surfaces* dialog, click **OK**.

  This action adds the data from the specified part surfaces.

To import an external file:

- Select the **File** option. When you click **OK**, the *Open* dialog appears.



- Select an available STL or VTK file from this dialog.
- Click **Open**.

  This action adds the data from the specified file.

The new arbitrary probe will be added to the simulation tree immediately after you select the part surface or import the external file. You can define the part using the properties.

Editing an Arbitrary Probe

Multiple techniques are available for editing an arbitrary probe.

You can edit an arbitrary probe using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Arbitrary Probe In-Place Dialog

This dialog facilitates the definition of the arbitrary probe with interactive features.

If you have an appropriate scene open when creating or editing an arbitrary probe, STAR-CCM+ displays an in-place dialog. Use this dialog to define the probe.



Click the **Input Parts** button to select the input parts that you want to include in the arbitrary probe.

To define the limits of the arbitrary probe, get shape data from part surfaces or an external file. Choose a source by selecting one of the options in the *Surface Input Options* group box. (See Creating an Arbitrary Probe for details.)

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your arbitrary probe, you can view it in the assigned scene.

Once you have created the arbitrary probe, rotate and zoom in the display to get a clear view of it. An example is shown with scalar data in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Working with Topological Parts

This section explains the derived parts that take data from one or more mesh cells.

The topological parts are:

Defining a Cell Surface

The cell surface part represents all the cells adjacent to or intersected by a given part.

A cell surface node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



A cell surface part can be useful for extracting boundary or interior cells. For example, a cell surface created from a boundary shows the cells adjacent to the boundary, whereas a cell surface created from a plane section shows all cells intersected by the plane. A cell surface part can be created using any part of dimensionality two or lower.



Cell Surface Properties

You can use the property that is listed below to adjust the specifications of the cell surface.

Parts	Selects the input parts as the source for deriving this cell surface.

Cell Surface Pop-Up Menu

Edit Part in Current Scene	Allows you to set the properties of this part interactively. This is only available if a scene is open.
Export...	Exports the extracted data of the part to an external STL file.

See Also

Using Threshold Parts and Cell Surfaces

Creating a Cell Surface

You can create a cell surface with the pop-up menu.

To create a cell surface, right-click the **Derived Parts** node and select **New Part > Cell Surface...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new cell surface to appear in the simulation tree immediately. Define the part using the properties.

Editing a Cell Surface

Multiple techniques are available for editing a cell surface.

You can edit a cell surface using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Cell Surface In-Place Dialog

This dialog facilitates the definition of the cell surface with interactive features.

If certain basic conditions are met when you start the process to create a cell surface or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.



Click the **Input Parts** button to select the input parts that you want to include in the cell surface.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Viewing the Final Result

When you finish creating your cell surface, you can view it in the assigned scene.

Once you have finished the creation process, rotate and zoom in the display to get a clear view of the cell surface, which is shown in green in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Visualizing the Solution

The visualization tools in STAR-CCM+ are integrated with the analysis, allowing for the interactive extraction and viewing of solution data from either a running or converged simulation.

It is possible to watch a flow field evolve as the simulation iterates, change parameters and immediately see the effects of those changes. This level of feedback allows for the dynamic supervision and control of the simulation, as well as allows you to gain insight on a particular aspect of the simulation. STAR-CCM+ provides identical visualization capabilities in both serial and parallel modes, with no additional user intervention. Since visualization in STAR-CCM+ operates in a client-server environment, the bulk of data processing takes place on the server process. Only the lightweight graphics data is sent to the client, allowing for a massively parallel simulation to be visualized from a client workstation.

Working with Scenes

This section defines what is meant by a scene, and describes how to use scenes to visualize the results of a simulation.

The special scenes available in STAR-CCM+ for visualizing models that are created using 3D-CAD are described in Visualizing A 3D CAD Model.

What Is a Scene?

A scene allows you to visualize the model.

The scene node, which has its own properties and pop-up menu, defines annotations and coordinate systems to be displayed, and serves as a manager node for additional objects that control various aspects of that scene.



In STAR-CCM+, a scene encapsulates the complete representation of the components that are required to generate an image or animation including: lights, cameras, parts, properties, transformations, geometry, textures, and other pertinent information.

To assist in managing the collection of objects in a scene, STAR-CCM+ provides another high-level object that is called a displayer. A displayer brings together parts and other attributes (such as scalars, vectors, and surface properties) to produce graphics primitives which are displayed in the scene. In order to construct visualizations it is important to understand the concepts of both scenes and displayers and how they interact. A scene can be comprised of one or more displayers, allowing for a flexible way to view parts with various different attributes. For example, a scene could have two displayers, one showing a part in outline mode, and another showing the same part with a scalar field.

A scene also manages the view. Using the mouse and keyboard, you can manipulate the scene view, select parts and other objects and interact with them.

If the scene is open, all scenes are continuously updated as the simulation changes for the duration of the session.

See Also

Refreshing the Display

Scene Properties

Use the scene properties to control characteristics such as transparency and mesh display.

Transparency Override	Overrides displayer settings for opacity in the scene.	
	Make Scene Transparent	Makes all parts transparent so that you can see all of them from any angle, and overrides the opacity settings of all the displayers in the scene.
	Make Scene Opaque	Makes all parts in the scene completely opaque, and overrides the opacity settings of all the displayers in the scene.
	Use Displayer Property	Reverts to the individual displayer settings for opacity.
Mesh Override	Overrides displayer settings for displaying the mesh in the scene.	
	Show All Meshes	Adds mesh lines, and overrides the mesh display settings of all the displayers in the scene.

	Hide All Meshes	Removes mesh lines, and overrides the mesh display settings of all the displayers in the scene.
	Use Displayer Property	Reverts to the individual displayer settings for mesh display.
Coordinate Systems	A list of the coordinate systems added to the scene. Editing the property accesses the object selection dialog to add or remove coordinate systems.	

Scene Expert Properties

Width	Width of the scene display (read-only).	
Height	Height of the scene display (read-only).	
Transparency Mode	Provides a choice between visual accuracy and speed of rendering of transparent images.	
	Alpha Blending	Renders the image faster.
	Depth Peeling	Provides a more accurate transparency.

Scene Pop-Up Menu

The pop-up menu of the scene includes refreshing the display and exporting the image to a graphics file.

Open	Executes the open scene action to make the associated visualization display visible. See Interacting with the Scene.
Refresh	Executes the refresh display action.
Restore Hidden Part	Executes the restore hidden part action to redisplay a part that is hidden by the hide part action. You choose which part to restore from a submenu of this menu item. See Restoring a Hidden Part in a Scene.
Restore All Hidden Parts	Executes the restore hidden parts action to redisplay all parts that is hidden by previous

	hide part actions. See Restoring All Hidden Parts in a Scene.
Hardcopy...	Executes the save scene to file action to create an image file of the current scene. See Saving Scene Image Files.
Report	Executes the test graphics action to display graphics hardware-related information in the *Output* window.

Creating a Scene

This section explains the various options for creating a scene.

STAR-CCM+ offers five default scene templates:

- Geometry—for viewing surface geometry
- Mesh—for viewing mesh grid lines and surfaces
- Scalar—for viewing scalar fields
- Vector—for viewing vector fields
- Empty—to start from scratch

Each default scene template contains a specific combination of displayers with certain properties set, and with a default set of parts, all of which you can change. For example, you can create a scalar scene and later add or remove displayers to suit your needs.

By default a scene is opened and displayed upon creation. A scene is continuously updated and kept current with the running simulation while the scene remains open.

Scenes are given default names when they are created. However they can be renamed at any time using any of the standard methods, as well as from within the visualization display. If multiple simulations are open at once, then the tab includes the name of the corresponding simulation. That part of the tab name cannot be changed unless you save a copy of the simulation with a new name.

Scenes can also be copied and deleted.

Creating a Geometry Scene

You can create a geometry scene to view outlines and surfaces of parts in the simulation.

When you create a geometry scene, you also create:

- A geometry displayer, which contains all boundaries in the simulation and is preset to display the outlines of the parts
- A second geometry displayer, which contains all boundary parts in the simulation and is preset to display the shaded surface

To create a geometry scene:

1. Right-click the **Scenes** node and select **New Scene** > **Geometry** from the pop-up menu.



An example of a geometry scene is shown below.



Creating a Mesh Scene

You can create a mesh scene to view the mesh grid along with surfaces of parts in the simulation.

When you create a mesh scene, you also create:

- A geometry displayer, which contains all boundaries in the simulation and is preset to display the outlines of the parts
- A second geometry displayer, which contains all boundary parts in the simulation and is preset to display the mesh cells on the shaded surface

To create a mesh scene:

1. Right-click the **Scenes** node and select **New Scene** > **Mesh** from the pop-up menu.



An example of a mesh scene is shown below.



Creating a Scalar Scene

You can create a scalar scene to view scalar contours on parts. You can choose to view contours of any scalar field function.

When you create a scalar scene, you also create:

- A geometry displayer, which contains all boundaries in the simulation and is preset to display the outlines of the parts
- A scalar displayer, which contains an undefined scalar field and no parts

You can add any valid part on which you wish to see scalar data. For a two-dimensional simulation, all region parts are added to the displayer by default.

To create a scalar scene:

1. Right-click the **Scenes** node and select **New Scene** > **Scalar** from the pop-up menu.



An example of a scalar scene is shown below.



Creating a Vector Scene

You can create a vector scene to visualize the velocity field on parts. You can choose to view glyphs of any vector field function.

When you create a vector scene, you also create:

- A geometry displayer, which contains all boundaries in the simulation and is preset to display the outlines of the parts
- A vector displayer, which contains a vector field (set to velocity by default) and with no parts

You can add any valid part on which you wish to see vector data. For a two-dimensional simulation, all region parts are added to the displayer by default.

To create a vector scene:

1. Right-click the **Scenes** node and select **New Scene** > **Vector** from the pop-up menu.



An example of a vector scene is shown below.



Creating an Empty Scene

An empty scene is a blank canvas to which you can add any type of displayer.

Each new displayer contains no parts, and is set to the default settings. Alternatively, you can add parts directly using the drag-and-drop technique.

To create a new empty scene:

1. Right-click the**Scenes** node and select **New Scene** > **Empty** from the pop-up menu.



Interacting with the Scene

A scene also provides simple but powerful ways to manipulate the image.

The main features are discussed in the Using STAR-CCM+ section; however, they include the ability to change the view dynamically by panning, rotating, and zooming the model.

In addition to simple orientation manipulations you can also change the properties of the scene, as well as manipulate the components of the model within the scene using the right-click functionality. For example, you can manipulate the regions and boundaries of the model from within the scene without having to revert to the tree. This provides a powerful visual aid to model manipulations.

The visualization toolbar at the top of the workspace also provides easy access to common functionality, including a back and forward button between recent views, and easy access to a transparency mode.

Any type of scene can be opened by double-clicking its node in the simulation tree, or by right-clicking that node and selecting **Open**. Either operation shows the selected scene in a new display in the *Graphics* window, or brings that display to the foreground if the scene is already open. A scene is opened with the view set to the last setting used, or with all parts fitted in the display looking down the Z axis with the Y axis up if it has never been opened before.

When more than one simulation is open, the tab at the top of the display includes the name of the corresponding simulation. This lets you know at a glance which simulation a display represents.

See Also

Refreshing the Display

Managing Parts Through a Scene

You can use the pop-up menu to work with parts in a scene.

Most of the work you do with parts in a scene, including adding and removing them, is through individual displayers. However, the following actions are possible with pop-up menu items of scenes:

Hiding a Part in a Scene

Hiding a part in a scene can be achieved by right-clicking on the part in the visualization display and selecting **Hide** from the pop-up menu.



This pop-up menu item is also available if you right-click a displayer part node in the object tree.

See Also

Accessing Objects Obscured from View

Hiding All Parts in a Scene

You can hide all the parts of a scene at once by right-clicking in the visualization display background and selecting **Hide All Parts** from the pop-up menu.



This pop-up menu item is also available if you right-click the scene node or the displayer Parts manager node in the object tree.

Restoring a Hidden Part in a Scene

To restore a hidden part, right-click in the visualization display background, select **Restore Hidden Part** from the pop-up menu, and then select the part from the submenu that appears.



You can also find this pop-up menu item by right-clicking the scene node or the displayer Parts manager node in the object tree.

This menu item is disabled (grayed out) if no parts in the scene are hidden.

Restoring All Hidden Parts in a Scene

To restore all hidden parts at once, right-click in the visualization display background and select **Restore All Hidden Parts** from the pop-up menu.



You can also find this pop-up menu item by right-clicking the scene node or the displayer Parts manager node in the object tree.

This menu item is disabled (grayed out) if no parts in the scene are hidden.

Adding Extra Displayers

In addition to the default displayers, it is also possible to add displayers to scenes.

Except for empty scenes, by default scenes are created with two displayers. The first displayer adds the mesh outline of all the boundaries. The second displayer depends upon the type of scene. For example, a scalar scene has a scalar displayer with no default parts, on which is displayed an undefined scalar.

To add a displayer, expand the **Scenes** node and the node of the chosen scene. Right-click on the **Displayers** node, select **New Displayer** from the pop-up menu and choose the type of displayer from the submenu.



Any new displayers added to the scene has default properties. The new displayer is named `Displayer N` where `Displayer` maps to the type of displayer added, and `N` is a numeric suffix to ensure that all displayers are uniquely named. All new displayers added with this action have an empty *Parts* manager, a new scalar displayer and streamline displayer have an undefined scalar field node, and a vector displayer has a vector field node set to velocity by default.

Removing Displayers

To remove a displayer from a scene, simply right-click on the displayer node you wish to remove and select the **Delete** action.



As a result, the displayer object and all associated graphics entities are removed from the scene.

Copying Scenes

When a scene is copied, all displayers and properties are copied, resulting in an exact duplicate of the original.

In addition to duplicating content, this copy action captures the open/closed state of the original: copy an open scene and the duplicate opens, copy a closed scene and the duplicate stays closed.

This process differs from copying properties of a scene.

Copying Properties of a Scene

You can copy the properties of a scene, though certain limitations exist.

When properties are copied from one scene to another, a new set of displayers appears in the destination scene alongside the originals. Those new displayers match the ones in the source scene. Applying the displayers' properties requires copying those properties separately. See Copying Properties of a Displayer.

Managing Scenes

STAR-CCM+ provides for easy management of scenes.

The **Scenes** node, which has a pop-up menu, is the manager object for working with scenes in the simulation tree.



It exists even when empty to allow creation of the first scene. The scenes manager can hold any number of scenes up to the memory limits of the computer.

When a scene is created or opened, it creates a visualization display. The scene is docked in the workspace in the same way as any other display, using a tabbed layout. When there are multiple scenes present, you can select which scene is active by selecting the appropriate tab. If multiple

simulations are open at once, then the tab includes the name of the corresponding simulation.

A scene can be closed by clicking the **X** button on the visualization display tab, or by right-clicking the tab and selecting **Close Window** from the pop-up menu. When a scene is closed, it does not continuously update during the iteration.



To open a scene expand the **Scenes** node, right-click on the scene and select **Open** from the pop-up menu.



To open all scenes at once, right-click on the **Scenes** node and select **Open All Scenes**.



A scene can be maximized to occupy the entire workspace by double-clicking its tab, and minimized by double-clicking the tab when in the maximized state. These operations are also available as a right-click action on the tab.

Using drag-and-drop, scenes can be moved anywhere in the workspace. To move a scene, click the scene's header or tab and drag it to the desired location. A red box shows where the scene is placed when the scene is dropped. You can do any of the following:

- Split any open scene. Drag the scene tab to any of the edges of the open scene. When the red box shows a split in the open scene, drop the scene.
- Move a scene from one open window to another. Drag the scene tab over the open tabs of the destination window. When the red box shows a tab in the new window, drop the scene.
- Arrange the order of the scenes within a window. Drag the scene tab to any location in the window's tabs and drop it.

Pop-Up Menu of the Scenes Manager

These pop-up menu items are available when you right-click the **Scenes** manager node.

New Scene	Executes the create scene action to create a scene of the selected type.	
	Geometry	Creates a geometry scene. See Creating a Geometry Scene.
	Scalar	Creates a scalar scene. See Creating a Scalar Scene.

	Vector	Creates a vector scene. See Creating a Vector Scene.
	Empty	Creates a new empty scene. See Creating an Empty Scene.
Open All Scenes	Executes the open all scenes action to display all scenes in the scenes manager.	
Test Graphics	Executes the test graphics action to report on the state of your graphics hardware. See A Scene Will Not Open.	
New Group	Facilitates the grouping of objects by creating a group node within the manager node. See Grouping Objects.	
Ungroup	Reverses the grouping process.	

See also Common Pop-Up Menu Items.

Changing the Attributes of a Scene

The **Attributes** node includes objects that let you control various characteristics of the scene, such as background colors, lighting, and animations.

This object exists within every scene node.



Working with the Model Window

The model window feature lets you crop part of the scene to show the model, while you use the rest of the space for elements such as annotations or scalar bars.

The **Model Window** node, which has properties and is subject to common menu items, represents this feature.



Use the properties to set the width, height, and position. In the example below, the width and height have been set to 0.7 of the display dimensions, allowing space for a simple text annotation, and a scalar color bar that is set to a vertical orientation. (The border of the model window appears only when the node of this feature is selected.)



Model Window Properties

Use these properties to control the width, height, and position of the model window.

Width	The width of the model window as a fraction of the visualization display width, where 1.0 equals the full width.
Height	The height of the model window as a fraction of the visualization display height, where 1.0 equals the full height.
Position	The position of the lower left corner of the model window in the visualization display.

Setting the Background Color

You can control the mode of the background color of the scene: solid or gradient.

The **Background Color** node has properties and is subject to common menu items.



The default mode is `Solid`. To change it, expand the **Scenes** node, the individual scene node, and the **Attributes** node. Select the **Background Color** node.

In the *Properties* window, select `Gradient`.



If the scene is open, the background shows a gradient.



Changing the Solid Background Color

The **Solid Background** node lets you change the solid background color of the scene.



To access the color controls, expand the **Scenes** node, the individual scene node, the **Attributes** node, and the **Background Color** node. Select the **Solid Background** node.

In the *Properties* window, specify the new color one of two ways. You can select it from the drop-down list next to the *Color* property.



Alternatively, to choose a color visually, you can use the *Color Selector* by clicking ... (**Custom Editor**) in that same property.

The results appear in the scene display if you have it open.



Solid Background Properties

Color	The background color of the visualization display. The color can be input directly in RGB format (for example, white would be entered as `255,255,255`). A comma must separate the numeric values. The property customizer executes the define color action. See Defining Colors.

Changing the Gradient Background Color

The **Gradient Background** node lets you change the colors and style of the gradient background color of the scene.



To access the gradient controls, expand the **Scenes** node, the individual scene node, the **Attributes** node, and the **Background Color** node. Select the **Gradient Background** node.

In the *Properties* window, specify each new color one of two ways. You can select it from the drop-down list next to the *Color 1* or *Color 2* property.



Alternatively, to choose a color visually, you can use the *Color Selector* by clicking ... (**Custom Editor**) in that same property.

The new color shows up immediately in the scene if the scene is open.



Repeat this procedure for the other color property if desired.

To change the style of the gradient, use the *Gradient Mode* property.



The new gradient mode appears, using the selected colors.



Gradient Background Properties

Color 1	The first color that is used in the gradient background color of the visualization display. The color can be input directly in RGB format (for example, white would be entered as `255,255,255`). A comma must separate the numeric values. The property customizer executes the define color action. See Defining Colors.	
Color 2	The second color that is used in the gradient background color of the visualization display. The techniques for editing this property are the same as for *Color 1*.	
Gradient Mode	Controls the direction from Color 1 to Color 2.	
	Top-Bottom	Color 1 at the top to Color 2 at the bottom.
	Bottom-Top	Color 1 at the bottom to Color 2 at the top.
	Left-Right	Color 1 on the left to Color 2 on the right.
	Right-Left	Color 1 on the right to Color 2 on the left.

	Radial	Color 1 at the center to Color 2 at the outer edges.

Background Color Properties

The Background Color properties let you choose between a solid and gradient color background.

Background Color Mode	Selects the color mode of the scene display background.	
	Solid	Uses a solid color mode.
	Gradient	Uses a gradient color mode (a blend of two colors or shades).

Working with the Current View of the Scene

Information regarding the layout of the display contents of a scene, such as the projection and view angle, is stored in STAR-CCM+ as a view.

Specific views can be created and saved, and are stored in the simulation tree under the **Tools > Views** manager node. Each scene also has its own **View** node, under the **Attributes** node for the scene, which provides details of the current view in the selected scene. The property values of this node change dynamically if the view is altered by using the mouse, and can be controlled interactively by editing the values. This node has its own properties and pop-up menu.



Controlling the Current View

The current view of the scene can be controlled by right-clicking on the **Attributes** > **View** node of a scene and selecting **Edit...**

Use the *Edit View* dialog to define each of the view properties. This technique can also be used to edit the properties of an existing view under the **Tools** > **Views** manager node.



It is also possible to change the value of each property directly in the *Properties* window. When this method is used, STAR-CCM+ attempts to reconcile the new value with the existing view properties in order to

update the scene. As a result, STAR-CCM+ may alter some of the properties values, and it is therefore recommended that the view properties are controlled using the *Edit View* dialog, as described above.

The *Parallel Scale* property in the dialog is used in parallel projection mode to scale the image. The parameter works as an inverse scale, in that larger numbers produce smaller images. Any user-defined value must be greater than zero. When you activate parallel projection mode, a parallel scale subnode is added to the **View** node.

Setting the Parallel Scale in a Scene

The **Parallel Scale** subnode of the **View** node is used to scale the image in parallel projection mode.

This subnode is made available when you set the *Projection mode* property of **View** node to `Parallel`. You can specify the scale in this object's properties.



Parallel Scale Properties

Parallel Scale	The user-defined value of the scale.

The parameter works as an inverse scale, in that larger numbers produce smaller images. Any user-defined value must be greater than zero.

	
Setting of 0.2	Setting of 1.2

View Properties

Use the view properties to control the position, coordinate system, and projection mode.

Focal Point	Cartesian vector specifying the position of the focal point. This property also changes when you zoom, pan, or rotate the visualization display.	
Coordinate System	A list of coordinate systems in the simulation for use by this view.	
Position	Cartesian vector specifying the position of the camera. This property also changes when you zoom, pan, or rotate the visualization display.	
Projection Mode	Defines the projection mode of the objects in the visualization display.	
	Perspective	Displays the objects in perspective mode with a vanishing point.
	Parallel	Displays the object with edges parallel to each other.

View Up	Cartesian vector specifying the direction where the camera is pointing. This property also changes when you zoom, pan, or rotate the visualization display.

View Pop-Up Menu

The view pop-up menu lets you edit views and store new ones.

Edit...	Opens the *Edit View* dialog, which can be used to control the current view of the scene. See Controlling the Current View.
Store View	Executes the store view action to save the view of the current scene. See Saving a View.

Updating Scenes

You can control how open scenes update by changing the trigger type that is located in the **Update** node.

This node, which has its own properties and is subject to common menu items, customizes how objects in the display of the *Graphics* window are updated as the simulation iterates. See Trigger Types.



The **Update** node has a sub-node with which to control the frequency of the update. The properties of the **Update Frequency** child node are used to change the frequency of updates. See Trigger Properties.



For more information on selecting the trigger type and defining the update frequency, refer to Setting the Trigger for an Update.

Updating a Scene Manually

The scene can also be refreshed manually. Expand the **Scenes** node and right-click on the chosen scene. Choose **Refresh** from the pop-up menu.



Generating Images from a Scene

You can also generate an image of the scene after each iteration or time-step to compile an animation of the scene. This is different from the single image that is generated when you save a hardcopy of the scene.

If all clients are disconnected from a running server, no graphics scenes are saved to files.

Note:

- Multiple displayers within a scene that is displaying different scalar quantities update simultaneously whenever an **Update Frequency** setting triggers an update. For example, the scalar quantities that are displayed in a scalar displayer are updated to represent the new solution. Closed scenes are not updated until they are opened.
- Multiple scenes displaying the same scalar quantity update simultaneously whenever an **Update Frequency** setting triggers an update. For example, if you have a simulation with a 1 s time-step, and have one scene set to update every 5 time-steps (5 s), and a second scene set to update with a delta time of 2 s, you ultimately get updates to both scenes at: 2 s, 4 s, 5 s, 6 s, 8 s, 10 s, and so on.
- Multiple scenes using the same parts with the same or different scalar quantities also update simultaneously whenever an **Update Frequency** setting triggers an update. For example, if you have two scenes displaying the same part, and their **Update Frequency** settings are 10 and 15, both of them will update after every multiple of 10 or 15 iterations, from the start to the completion of the simulation run. This reflects the philosophy that all client scenes reflect an up-to-date snapshot of the server data.
- Hardcopies are only written to disk at the intervals that are dictated by the **Update Frequency** setting for a particular scene.

Update Properties

The scene update properties allow you to control how the display is updated, as well as whether and how to generate image files of the display during a simulation run.

Trigger	Determines how the display is updated to reflect changes in the simulation. Change the trigger type to control the method that updates the scene. See Trigger Types.
Save to File	Determines whether multiple images of the scene display are saved as the solution progresses. This property can be used to create

	an animation showing how the object changes while the simulation runs. STAR-CCM+ brings the selected scene to the front when it produces a screenshot of that frame.
Output Directory	Path to where the files specified in the *Base Filename* property are saved.
Base Filename	The prefix of the file name to use when saving scenes to file. The full file name is `prefixNNNNN.ext` where `prefix` is this property, `NNNNN` is a number to indicate the sequence of the files, and `ext` is an extension based on the graphics format. If you leave the text box blank, STAR-CCM+ restores the name that was there previously once that text box no longer has the focus.
File Format	The graphics format that is used to save images of the visualization display. You have the same options as provided in the save scene to file action.

Expert Properties

These properties let you control the specifications of the images generated by the scene updates. However, if all clients are disconnected from a running server, no graphics scenes are saved to files.

Output Magnification	Produces a larger image than the image in the original display, based on the factor you enter.
Offscreen Width	Sets the width of the image of the scene display when the save scene to file action is executed. To make this specification effective, change it while recording a macro for saving scene images in batch.
Offscreen Height	Sets the height of the image of the scene display when the save scene to file action is executed. To make this specification effective, change it while recording a macro for saving scene images in batch.

Creating Images for an Animation

You can create images for an animation of the simulation, since STAR-CCM+ lets you output the image of the scene to a unique graphic file after each update of a scene during a simulation run.

Note: If all clients are disconnected from a running server, no graphics scenes are saved to files.

In the tree, expand the **Scenes** node, the node of the chosen scene and the **Attributes** node. Select the **Update** node.



Decide how frequently you want the scene to be updated by changing the properties of the **Update Frequency** child node. The default setting of 1 for the **Iteration** and **Time-Step** options means that the scene will update after each iteration or time-step. The default setting of 1 s for the **Delta Time** option means that the scene will update after 1 s in actual time. If you are working with a transient analysis, change the *Update Policy* from **Iteration** to **Time-Step** or **Delta Time**.



The format of the file can be selected from the options in the drop-down menu under the *File Format* property.



You have the same format options as those provided when you save a scene to a file. See Saving Scene Image Files.

In the *Output Directory* property, you can control where the files are created one of two ways. If you know the full path, simply enter it into the text box of the property.



Otherwise, click the property customizer (...) to activate the browser dialog.



Choose the location carefully; if, for example, you generate one image per iteration and you run 1000 iterations, then 1000 image files are created.

Change the *Base Filename* to something other than `image` if you prefer.



Tick the checkbox of the *Save to File* property.



The update feature is ready for you to begin the run. Before you do, however, make sure that your visualization display is not obscured by anything else on your screen. If anything is covering the display at the time an image file is saved, it appears in the image. For example, be sure to deactivate your screen saver; otherwise all you have at the end of the run is an animation of your screen saver.

Also, have only one scene active at a time. Depending on your system and your version of STAR-CCM+, the scenes may conflict with each other and result in an incomplete set of images.

Run the simulation for the desired number of steps. As it runs STAR-CCM+ will save an image file after each update. The name of each image file (the base file name) is concatenated with the number of the update to ensure that each file is unique.

These files can then be animated with your media program of choice.



Controlling the Frame Rate of the Scene Animation

You can adjust the number of frames per second of all scene animations with the **Animation** node.

This object is in the **Attributes** subnode of a scene node.



Animation Properties

In the Animation properties, set the frame rate of the animation.

Target frame rate (fps)	Maximum number of animation frames that are displayed in one second (less for complex scenes, or when less computing power is available).

Working with the Axes

It is possible to control the orientation axes in the visualization display. The **Axes** node, which has its own properties, lets you do this.



To hide the axes, expand the **Scenes** node, the node of the chosen scene and the **Attributes** node. Select the **Axes** node. In the *Properties* window clear the checkbox of the *Visible* property.

To show the axes again, tick the checkbox of that property.

Axes Properties

The properties of the **Axes** node let you specify various characteristics of the axes in the scene display.

Text Color	Controls the color of the X, Y, and Z labels of the axes. The property customizer executes the define color action. See Defining Colors.

Visible	Sets visibility of the axis indicator. You may want to hide the axis indicator when it overlaps with an object in the visualization display, or when you want to produce a printout or a screenshot.	
	Activated	Axis indicator is visible.
	Deactivated	Axis indicator is hidden.
Viewport	Gives the coordinates, in the form `[x0, y0, x1, y1]`, of the lower left and upper right corners of a rectangle containing the axes. By changing the coordinates, you can control the position and size of the axes. Units are fractions of the viewport height and width. The axes do not fill the rectangle.	

Lighting

Each scene includes a sophisticated set of lights that function like those in a real display setting. The **Lights** manager node contains all of the light nodes in the scene.

These lights can help you study an object's design, or generate images for presentation documents. Unlike the lighting of displayer surfaces, which determines how parts within a particular displayer absorb the light, this lighting affects the scene as a whole.



In particular, the properties of the manager object offer a choice of lighting models. The pop-up menu of the manager allows you to restore default settings, use the head light feature, and refresh the lighting.

Restoring Default Lighting Settings

To reverse all of your changes to the light settings, simply right-click on the **Lights** manager node and select **Restore Default Lighting**.



If no adjustments have been made to properties of any lights during this session, then this action has no effect.

Using the Head Light Feature

To focus Light 1 on the object in the visualization display, right-click on the **Lights** manager node and select **Head Light**.



This action does the following:

- Turns off the other three lights
- Sets the intensity of Light 1 at the maximum
- Reverts the color of Light 1 to white
- Sets the azimuth and elevation of Light 1 to zero

Using the Refresh Action

After you finish changing the light settings, you may need to refresh the lighting to see the effects of your changes.

To refresh the lights, right-click the **Lights** manager node and select **Refresh**.



Lights Manager Properties

To specify the lighting model, use the properties of the **Lights** manager node.

Lighting Model	Chooses a mode for lighting and shading.	
	Per Vertex Lighting	This option calculates lighting by interpolating among vertices. It is relatively economical with graphics processing requirements. 
	Per Pixel Lighting	This option calculates lighting based on pixels. While it requires more graphics processing power than the **Per Vertex Lighting** option, it also provides a more precise image. Also,

		this option makes it possible to light streamlines. 

Pop-Up Menu of the Lights Manager

These pop-up menu items are available when you right-click the **Lights** manager node.

Restore Default Lighting	Executes the restore default lighting action. See Restoring Default Lighting Settings.
Head Light	Executes the head light action. See Using the Head Light Feature.
Refresh	Executes the refresh action. See Using the Refresh Action.
Edit	Executes the edit action. See Using the Table Editor.

Working with Lighting Objects

A **Light** node, which has properties and is subject to common menu items, contributes to lighting in the scene.



It is not possible to create, delete, or rename these objects. There are four of them in a scene, all identical except that Light 1 cannot be disabled with the *Enabled* property.

Restoring default settings and using the head light feature are part of working with the **Lights** manager node.

How Do Lighting Objects Work?

This section explains the principles by which lights in STAR-CCM+ scene displays work.

All lights are directional lights (infinitely far away with no falloff). Lights move with the camera, which is controlled by the view.

For simplicity, the position of lights in STAR-CCM+ can only be specified using angles. The *Elevation* (latitude) and *Azimuth* (longitude) properties of each light are expressed in degrees with respect to the camera. (Lights always shine on the camera's focal point.) For example:

- A light at (*Elevation* = `0`, *Azimuth* = `0`) is at the camera (a headlight).



- A light at (*Elevation* = `90`, *Azimuth* = `0`) is above the focal point, shining down.



- Negative azimuth values move the lights counter-clockwise as seen from your point of view, positive values clockwise. So, a light at (*Elevation* = `45`, *Azimuth* = `-20`) is above and in front of the object, shining slightly from the left side.



Deactivating and Reactivating Lighting Objects

Individual lights can be deactivated or reactivated, except for **Light 1**.

To turn off a particular light, select its node in the object tree.



Clear the checkbox of its *Enabled* property.



(Although this cannot be done for **Light 1**, you can turn off that light by setting its *Intensity* property to 0.) To turn the light back on, tick the checkbox.

Controlling the Intensity of Lighting Objects

The *Intensity* property controls the brightness of the light.



Zero is dim.



One is maximum brightness.



Controlling the Color of Lighting Objects

The *Color* property of lights lets you add another visual effect to the display.

You can use colors for highlighting parts of the model, or simply for a preferred aesthetic effect.



Light Properties

The properties of a typical light object lets you activate or deactivate it, as well as control its lighting characteristics.

Azimuth	Azimuthal location of light in degrees.	
Elevation	Elevation of light in degrees.	
Enabled	Uses a checkbox to turn on or off the light.	
	Activated	Turns on the light (default).
	Deactivated	Turns off the light.

Intensity	Controls the strength of the light with text entries as fractions of 1. An entry of 0 effectively turns the light off, while an entry of 1 sets the light to its maximum intensity. Any entry greater than 1 yields an error message.
Color	Controls the color of the light. The property customizer executes the define color action. See Defining Colors.

Using Clip Planes

In every scene, you can render portions of the model invisible with clip planes.

The **Clip Planes** manager node, which has its own pop-up menu, contains all of the clip plane object nodes in the scene. By default it is disabled, as shown by the icon's gray color.



Clipping parts in a scene differs from removing them in that they are still in that scene. It also differs from hiding a part in that you are cropping the view by a set of coordinates that you control, rather than excluding an entire part from view.

This feature allows you to display information, such as scalars or vectors, within a space of your choice. In addition, it allows you to display a smaller portion of the model.

Enabling Clip Planes

To enable the clip planes in a scene, right-click the **Clip Planes** manager node and select **Toggle Clip Enabled**.



The scene shows the objects clipped along the planes, based on their default coordinates.



To reverse this step, you can disable the clip planes with the same pop-up menu item. You can also disable and enable individual clip planes. See Disabling and Enabling Clip Plane Objects.

Modifying Clip Planes

Multiple techniques exist for modifying clip planes.

To edit a clip plane, select its node in the object tree and edit the properties.



Alternatively, to modify the planes visually, right-click the **Clip Planes** manager node and select **Edit Clip Planes**.



An interactive in-place dialog appears in the scene. (If the scene display is not in view, the **Edit Clip Planes** action causes the display to be opened or become active.)



Use the (**Up-Arrow**) button to minimize the in-place dialog to see the entire scene.

You can adjust the tool by clicking and dragging any of the end point handles of the lines, or any of the sides of the bounding box.



When done with the tool, restore the dialog with the (**Down-Arrow**) button. Click **Apply** to implement your changes. If you have no further changes to make with the tool, click **Close** to remove the dialog.

Inverting the Coordinates of Clip Planes

This action changes all normal values from positive to negative, and vice versa.

To invert the normal values of all clip planes in a scene at once, right-click the **Clip Planes** manager node and select **Invert Clip**.



Pop-Up Menu of the Clip Planes Manager

You can carry out a variety of operations with the pop-up menu of the **Clip Planes** manager node.

Toggle Clip Enabled	Enables and disables the clip planes.
Edit Clip Planes	Opens an in-place dialog to modify the clip planes visually.
Invert Clip	Switches the normals of the clip planes to their negative equivalents.
Paste	Pastes the copied clip plane to the clip planes object tree.
Edit	Opens the edit dialog to modify the clip planes properties.

Working with Clip Plane Objects

Clip plane objects represent individual sets of adjustable coordinates for clipping an object along a particular plane in a scene.

There are six of these clip plane nodes in any scene, all with identical properties.



It is not possible to create or delete these objects, though they can be renamed with the inline editor, copied using the clipboard, as well as individually disabled or enabled.

Disabling and Enabling Clip Plane Objects

You can deactivate an individual clip plane by selecting its node and clearing the *Enabled* property checkbox.



Tick the same checkbox again to reactivate it.

Clip Plane Object Properties

Use the properties of a clip plane object to activate it, deactivate it, and specify its coordinates.

Enabled	Uses a checkbox to turn on or off the clip plane.	
	Ticked	Turns on the clip plane (default).
	Cleared	Turns off the clip plane.
Normal	Sets x, y, and z coordinates for the second point of the plane.	
Origin	Sets x, y, and z coordinates for the origin of the plane.	

Customizing an Annotation in a Scene

Scenes can have their own annotation objects that you can customize separately from the original annotations.

Customized annotations are also called applied annotations, and they are stored in the **Attributes** > **Annotations** node within the scene node. The **Annotations** node has properties and is subject to common menu items.



What Are Applied Annotations?

An applied annotation node represents the addition of an annotation to a scene. Its name matches that of the original annotation.

Applied annotations have properties (except for the **Background** type) and share a pop-up menu.



To create an applied annotation, simply follow the procedures for adding an annotation to a scene. Once an applied annotation is in a scene, it can only be changed by modifications to its own properties. If you edit the original annotation subsequently, it does not affect the applied annotation. The exception is the text content of a simple text annotation, which can only be edited from the original.

If you have an annotation that you copied to multiple scenes, it is possible to apply a change globally with the following steps:

- Modify the original annotation.
- Delete the applied annotations that you wish to update.
- Add the modified annotation to all the scenes, creating replacement applied annotations.

To rename an applied annotation, rename the original annotation. To hide or show the applied annotation in the scene, clear or tick the checkbox of its *Visible* property.

Applied Annotation Properties

These properties are shared by all applied annotations other than the **Background** type, except where otherwise noted.

Width	The width of the annotation as a fraction of the visualization display width, where 1.0 equals

	the full width (available for **3D Image** and **Plot Image** types only).	
Height	The height of the annotation as a fraction of the visualization display height, where 1.0 equals the full height (not available for **Scene Grid** type).	
Position	The position of the lower left corner of the text or graphic in the visualization display based on the coordinate system specified in the *Coordinate System* expert property (not available for **Scene Grid** type).	
Rotation	Rotates the image in a counterclockwise direction (**3D Image** type only).	
Normal	Sets x, y, and z coordinates for the second point of the image plane (**3D Image** type only).	
Visible	Shows or hides the annotation in the scene (available for all types, including **Background**).	
	Ticked	The annotation is visible.
	Cleared	The annotation is hidden.
Width Stretch Factor	Scales the width of the annotation, with the default of 1.0 equaling 100% of the original width (not available for **3D Image**, **Plot Image** and **Scene Grid** types).	

Applied Annotation Expert Properties

These expert properties are not available for the **3D Image** and **Background** types.

Location	Controls the visibility of the annotation relative to the objects in the display.	
	Background	Shows the objects in front of the annotation.
	Foreground	Shows the annotation in front of the objects.
Representation	Controls which mesh representation is used in the scene. Select the desired representation from the drop-down list. The **Solution Time**, **Iteration**, and **Time Step** annotations can display data from a recorded solution. The	

	annotations update automatically as you retrieve results from the solution history file. See What Are Representations?.	
Coordinate System	Selects the type of coordinates to use for position (not available for **Scene Grid** type).	
	Normalized Viewport	Shows x-y values as fractions of the length and height of the *Graphics* window display, respectively. The origin is the lower left-hand corner of that display.
	World	Shows x-y-z values as global coordinates. As you rotate the object, the annotation keeps its position relative to that object.
Interactive	Controls whether the annotation can be moved in the scene display with the mouse (not available for **Scene Grid** type).	
	Activated	Allows you to move the annotation in the display.
	Deactivated	Locks the position of the annotation in the display.

Applied Annotation Pop-Up Menu

Edit Annotation...	Opens the table editor for the parent annotation. See Using the Table Editor.
Remove from Scene	Removes the applied annotation from the active scene.

Annotations Manager Properties

To add or remove annotations in a particular scene, access the properties of the **Annotations** manager node.

Annotations	A list of the annotations added to the scene. Editing the property accesses the in-place object selector to add or remove annotations. See Using the Part Selector.

Displaying Geometry and Mesh

Geometry displayers are used to display the outline and surfaces of the parts.

The **Geometry** displayer node controls the appearance of the surface and mesh aspects of the parts in the scene. This node has its own properties and shares a pop-up menu with other types of displayers.



Use the displayer to specify whether the model appears with outlines, a complete grid or shaded surfaces or a mixture of these aspects.

Creating Geometry Displayers

Two geometry displayers are created automatically when a geometry scene is created. Alternatively, you may wish to add a geometry displayer to any scene.

The first displayer is set to display the outline of all the boundary parts whereas the second is set to show the surfaces of all the boundary parts. A similar setup exists when you create a mesh scene, with the surface geometry displayer also set to display the mesh grid. In addition, a geometry displayer is created automatically when you create either a scalar scene or a vector scene. In these cases, the geometry displayer is set to display the outline of the boundary parts.

To add a geometry displayer manually, open the **Scenes** node and the node of the chosen scene. Right-click on the **Displayers** node and select **New Displayer > Geometry**.



By default the geometry displayer will contain no parts after creation and you must add parts to the displayer. This feature allows you to choose only those parts of the model that you wish to view.

The default name of the displayer is based on its type. However, renaming the displayer is possible.

Using Surface and Outline Displays

Geometry displayers control the visibility of surfaces and outlines in the scene display, and allow you to show or hide them.

Often, outlines are visible through a geometry displayer named **Outline**, which accompanies another displayer in most types of scenes (any default geometry, scalar or vector scene).

To make the outlines invisible, expand the **Scenes** node and expand the scene that includes a geometry displayer. Expand the **Displayers** node and select the geometry displayer node, commonly named **Outline**. (In an empty scene in which you create a geometry displayer, the default name of the displayer might be different but this procedure works the same way.)



In the *Properties* window, deactivate the checkbox of the *Outline* property.



In the scene display, the outlines are gone. In this example with a geometry scene, only the surfaces are displayed by the other geometry displayer. (The boundary type of the part dictates the color of the surface unless this has been changed.)



To restore the outlines, simply activate the same checkbox.

To make the surfaces invisible, expand the **Scenes** node and expand the scene that includes the geometry displayer. Expand the **Displayers** node and select the geometry displayer node.



In the *Properties* window, deactivate the checkbox of the *Surface* property.



In the scene display, the surfaces are gone. To restore them, simply activate the same checkbox.



Displaying Mesh Outlines and Surfaces

You may display the mesh on the surface of any part in a geometry displayer.

To display a mesh in a geometry displayer, expand the **Scenes** node and expand any scene that includes a geometry displayer (any default geometry, scalar or vector scene, or an empty scene in which you create a geometry displayer). Expand the **Displayers** node and select the geometry displayer node.



In the *Properties* window, activate the checkbox of the *Mesh* property.



The mesh is visible in the scene display.



The *Display Mesh* property of the scalar displayer works the same way.

In addition to displaying the mesh outlines, the surfaces of each cell may be illustrated by shrinking each surface element.

Begin by making sure that the surface is visible.

To shrink the elements, click the right column of the *Element shrink* property in the *Properties* window. In the slider box that appears, select a value between `0` and `1`, where `0` results in no surface and `1` results in whole surface elements. You can also enter your selection directly in the **value** text box above the slider. Click anywhere in the workspace around the slider box to close it.



The separate cell elements are visible in the scene display.



Displaying Feature Lines

You can capture and display feature lines. This enables you to display mesh edges that are extracted by feature angle (the angle between normal vectors of adjacent polygons).

The **Feature Angle** node controls the feature lines. This node has properties and is subject to common menu items. The node appears only when the *Feature lines* property of the displayer is activated.



To make this object available, expand the **Scenes** node and expand any scene that includes a geometry displayer (any default geometry, scalar or vector scene, or an empty scene in which you create a geometry displayer). Expand the **Displayers** node and select the geometry displayer node.



Activate the checkbox of the *Feature lines* property in the *Properties* window.



Once you activate the node, you can use its *Feature Angle* property to change the angle between surface elements that defines a feature. The smaller the angle, the more edges are extracted.



Feature Angle Properties

Feature Angle	Sets the minimum angle between polygonal facets when displaying feature lines.

Geometry Displayer Properties

Use the properties of the geometry displayer to control the appearance of the surface, outline, and mesh lines.

Outline	Controls the appearance of the outlines.	
	Activated	Shows the outlines of the parts.
	Deactivated	Hides the outlines of the parts. When you choose this option and activate the *Surface* property, you generate a more realistic image of the model.
Surface	Controls the appearance of parts of the model.	
	Activated	Activates a solid coloring of the parts. When you choose this option, you generate a more realistic image of the model.
	Deactivated	Leaves no shading; the surfaces are completely transparent except for outlines, if outlines are activated.
Mesh	Controls the appearance of the mesh grid.	
	Activated	Shows the complete grid and adds the **Mesh Color** subnode to the displayer node to control the color of the mesh lines. This option lets you examine the density

		and cell structure of various parts of the model, and make judgments about accuracy.
	Deactivated	Hides the grid except for the outlines. When you choose this option and activate the *Surface* property, you generate a more realistic image of the model.
Feature Lines	Controls the appearance of the feature lines.	
	Activated	Adds the **Feature Angle** subnode to the displayer node to show feature lines.
	Deactivated	Hides the feature lines and removes the **Feature Angle** subnode. When you choose this option and activate the *Surface* property, you generate a more realistic image of the model.
Element Shrink	Contracts the vertices of each polygon toward the polygon centroid when set below 1.0. This lets you see the surface cells of the model more clearly. If you enter `0`, the shading disappears; only the edges are visible if they are activated. The *Surface* property must be activated in order for the effects of this function to be visible. This property can be set with a slider.	
Color Mode	Controls the arrangement of colors.	
	Type (Default)	Colors the part surfaces according to their boundary type.
	Constant	Adds the **Color** subnode to the displayer node to

		display surfaces using a custom color.
	Distinguish Regions	Colors the surfaces by region, so that different regions can be identified.
	Region Type	Colors the surfaces by region type, that is, fluid, porous or solid.
	Distinguish Inputs	Colors regions by boundaries, or parts by part surfaces - depending if regions or parts are displayed in the scene.
Opacity	Adjusts the transparency of the part. The settings range from 0.0 (invisible) to 1.0 (opaque), and can be controlled with a slider. This property allows you to show an overlay of different parts of the model.	

Geometry Displayer Expert Properties

Representation	Selects a representation. Additional selections would appear as representations are added to the simulation.	
Point Size	Sets the point size in pixels. This property can be set with a slider.	
Line Width	Adjusts the width in pixels of lines that are used in the display, depending on the type of displayer. For example, in a scalar displayer, this setting would be effective only if the *Contour Style* property were set to `Line`. This property can be set with a slider.	
Lighting	Controls the lighting of the part. Accessing the property customizer opens the *Lighting* dialog.	
Smooth Shade	Controls the transition between segments of the image.	
	Activated	Makes the transitions smooth for a more realistic effect.

	Deactivated	Leaves this function off.
Ignore Clip	Controls whether the parts in the displayer are affected by any enabled clip plane.	
	Activated	Shows the parts unaffected by the clip plane.
	Deactivated	Shows the parts that are trimmed by the clip plane.
Face Culling	Controls whether and how to cull polygons.	
	None	Does not cull any polygons.
	Front	Culls (omits from the display) all polygons that are oriented toward the front of the view.
	Back	Culls (omits from the display) all polygons that are oriented toward the back of the view.
Immediate Mode	Controls whether parts in the displayer are rendered from CPU memory or from OpenGL display lists held in the GPU of the graphics hardware. Using display lists may provide faster performance for static data.	
	Activated	Parts are rendered from CPU memory.
	Deactivated	Parts are rendered from display lists on the graphics hardware.
Transform	Selects a transform to apply to all parts in the displayer. **None** is the default selection. Additional selections would appear when you create transforms. This feature allows for the creation of symmetric and periodic repeats, as well as arbitrary transforms to be used.	

Displaying Scalar Quantities

Scalar displayers are used to display scalar field functions or field function monitors on parts using contours.

The **Scalar** displayer node, which has its own properties and shares a pop-up menu with other types of displayers, controls the display of data along the surface of a part.



In the icon, the numeral next to the paint brush signifies the depiction of quantities.

See Also

Using the Scalar and Vector Bars

Creating Scalar Displayers

A scalar displayer is created automatically when a scalar scene is created. Alternatively, you may wish to add a scalar displayer to any scene.

To add a scalar displayer manually, open the **Scenes** node and the node of the chosen scene, right-click the **Displayers** node and select **New Displayer > Scalar**.



By default the scalar displayer will contain no parts after creation and you must add parts to the displayer. This feature allows you to choose only those parts of the model on which you wish to display scalar quantities.

The default name of the displayer is based on its type. However, renaming the displayer is possible.

Changing the Scalar Quantity

The scalar displayer and streamline displayer have no default scalar, so you must select one from a list of scalar field functions to display data in a scene.

The list of field functions is available through the **Scalar Field** node, which has properties and is subject to common menu items.



To access the **Scalar Field** node, expand the following nodes: the **Scenes** node, the node of a scene with a scalar displayer, the **Displayers** node within that scene, and the scalar displayer node.

Once you select the node, the scalar field function can be chosen from the list in the object selection dialog. To activate the dialog, edit the *Function* property.



This list also allows you to select the individual components of certain functions, such as positions and centroids, in local coordinate systems.



You may also choose the scalar quantity using the same drop-down list by right-clicking the scalar bar in the visualization display. This may be more convenient when you manipulate the scene.



Scalar Field Properties

Function	Selects the type of scalar field function or field function monitor to show, such as temperature or velocity magnitude. See What Is a Field Function Monitor?	
Min	Sets the minimum value in the displayed range.	
Max	Sets the maximum value in the displayed range.	
Clip	Determines how to display values outside of the range that is specified by *Min* and *Max*. See Changing the Scalar Range.	
	Off	Values above *Max* are displayed in the maximum color, and values below *Min* are displayed in the minimum color.
	Below Min, Above Max	Does not color values outside the range--

		uses the normal surface color.
	Below Min	Does not color values below *Min*--uses the normal surface color.
	Above Max	Does not color values above *Max*--uses the normal surface color.
Units	Selects the type of units to display. See Changing the Scalar Units.	
Auto Range	Determines whether to set the range of values to display automatically. See Changing the Scalar Range.	
	Off	Does not use auto ranging, rather use the values that are set in *Min* and *Max*.
	Min and Max Values	Automatically determines the range from the field function data.
	Min Value	Automatically determines the minimum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Max Value*.
	Max Value	Automatically determines the maximum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Min Value*.

Auto Min	Indicates the minimum value among all parts in the displayer (read-only).	
Auto Max	Indicates the maximum value among all parts in the displayer (read-only).	
Auto Range Mode	Chooses the domain for auto-range minimum and maximum value computation that is used for color coding.	
	Displayer Parts	Parts added to the particular displayer.
	Displayer Regions	Regions from which the displayer extracts data.
	All Regions	All regions in the simulation.

Changing the Contour Style

The *Contour Style* property has several options for showing scalar contours. By default it is set to show them as filled with solid colors.

To change the contour style of the displayer, open the **Scenes** node, the node of the particular scene which includes a scalar displayer, and the **Displayers** node within that scene. Select the node of the scalar displayer.



Select the new style from the drop-down list of the *Contour Style* property in the *Properties* window.



The property can be set as follows:

- `Filled`



The filled contour style fills each cell face of the surfaces displayed with the scalar value stored at the center of its corresponding cell.

- `Smooth Filled`



The smooth filled contour is generated using vertex values rather than cell centroid-based values and appears smoother than the default filled contour style. However, for most scalar fields values are interpolated from the cell centers to the vertices (or nodes) of each cell. Thus, data extraction with the smooth-filled option is more expensive in terms of computing power.

- `Smooth Filled + Lines`



- `Line` (The illustration below shows a scene with a geometry displayer that is added to provide a background for the lines.)



The line contour style displays isolines of the vertex-based scalar quantity that is colored by the value of the scalar.

- `Smooth Blended`



The smooth blended option is based on vertex values, similar to the smooth filled option. However, it is less computationally expensive than the smooth filled option due to the technique used to blend colors across the cell surfaces.

Understanding Smoothed Values

To avoid the default stepped appearance of a scalar scene, STAR-CCM+ also provides options that apply smoothing to scalar scenes.

When STAR-CCM+ runs a simulation, it stores many computed quantities at the location of the cell centroid. When you create scalar scenes, the default is to display results based on these cell center values only. While this option is essential for accurate reporting of computed quantities, it results in scalar scenes that have a stepped appearance. An example of this is shown below.



When the smooth filled option is chosen, for example, the following image results.



To obtain the nodal values that are used in smoothed scalar scenes, STAR-CCM+ first maps the cell center values (where the raw simulation data is stored), to the vertices of the mesh. It then applies a weighted interpolation between these vertex values and the points at which a chosen section (for example, a plane section derived part) cuts through the edges of the mesh between the mesh vertices. This is illustrated in the following diagram.



When you choose one of the smooth or blended options for the scalar displayer, the results you see is based on this mapping and interpolation.

Note: If you apply the smooth filled option to a section to visualize field variables that are only defined on boundaries, the face data is interpolated onto the boundary vertices to create vertex data. The section then takes up these values from the boundary vertices.

Labeling Line Contour Values

You can plot the actual values of the scalar contour lines on the lines themselves.

When you set the Contour Style property of the scalar displayer to either `Smooth Filled + Lines` or `Line`, the **Isoline Labels** node is made available in the scalar displayer node. This node has properties and is subject to common menu items.



To make the labels appear in the scene, activate the checkbox in the *Visible* property of the node.



Isoline Labels Properties

Point Size	Sets the point size in pixels.	
Visible	Shows or hides the labels and points in the scene.	
	Activated	The labels and points are visible.

	Deactivated	The labels and points are hidden.
Label Font	Sets the font of the labels.	
Label Style	Uses a drop-down list to display the labels as normal, bold, italic or bold and italic.	
Label Color	Sets the color to display the labels. The property customizer executes the define color action. See Defining Colors.	
Label Precision	Specifies the total number of digits in the label. Raising this number adds decimal places.	
Shadow	Activates a shadow effect for the labels.	
	Activated	Shadow effect is on.
	Deactivated	Shadow effect is off.
Point Color	Sets the color to display the points. The property customizer executes the define color action. See Defining Colors.	
Label Size	Sets label size as a percentage of window size.	

Changing the Scalar Range

By default the range of the scalar quantity is automatically based upon the range of the field function on the parts in the displayer. However, the range can also be set manually.

Expand the following nodes: the **Scenes** node, the node of a scene with a scalar displayer, the **Displayers** node within that scene, and the scalar displayer node. Select the **Scalar Field** node.



In the *Properties* window, enter the minimum value under the *Min* property.



This causes the *Auto Range* property to switch to `Max Value` automatically, and the result appears in the scalar scene display. If you set the *Max* property first, the *Auto Range* property setting switches to `Min Value`.

Now enter a value for the *Max* property, and the *Auto Range* property sets itself to `Off`, since you have manually set both values.

If the manually entered range is smaller than the bounds of the scalar quantity, the contours are clipped by default.



If you select `Off` in the drop-down list of the *Clip* property, scalar quantities out of range appear in the color contours at the extremes of the color gradient.



Alternatively, you can set the *Clip* property to `Below Min` to exclude only the scalar quantities below the range, or `Above Max` to do this with the quantities above the range.

You can revert to the automatic range setting anytime by selecting `Min and Max Values` in the drop-down list of the *Auto Range* property.



Changing the Scalar Units

If alternative units are available, the units displayed with the scalar legend in the scene can be changed.

Expand the following nodes: the **Scenes** node, the node of a scene with a scalar displayer, the **Displayers** node within that scene, and the scalar displayer node. Select the **Scalar Field** node.



Select a unit from the drop-down list of the *Units* property in the *Properties* window.



The units along the scalar bar change accordingly.

	
Previous units	New units

Working with Point Scale Modes

The screen display size of parts that are represented by points (such as Lagrangian parcels or points on a presentation grid) can either be set to a constant pixel size or be scaled according to a scalar function.

The display size options are available in the *Point Scale Mode* property of the scalar displayer. The option you select depends largely on what you are displaying.

What Is the Point Size?

This parameter, represented by the **PointSize** node, is the default setting. It sets the screen display size of a point in pixels in its properties, independent of the physical size of the point. This option is likely to be useful when you want to display the points on a presentation grid. Lagrangian particles or parcels are not sized according to their physical size if this option is chosen.



PointSize Properties

Size	Sets the point size in pixels.

What Is the Point Size Scale Field?

The Point Size Scale Field lets you scale a point's screen dimension based on scalar quantity, such as particle size or fluid temperature. It is represented by the **Point Size Scale Field** node, which is activated when you set the *Point Scale Mode* property of the scalar displayer to `Scalar`.



In the properties, select a scalar field function on which to base the screen point size. The Point Size Scale Field can be useful in cases involving particles in Discrete Element modeling or parcels in Lagrangian Multiphase modeling. It is recommended that you use the scalar Particle Size when you wish to plot particles according to their actual size. The size of the particle on the screen is calculated relative to the overall dimensions of the mesh. An exaggerated example of this is shown below.



If you want to scale by something else (such as velocity), map the velocity range to a physical size range, where the physical size range sets the minimum and maximum size of particles relative to the mesh dimensions. Activate the *Map values* property to activate this mapping. The formula that is used for the mapping is:

$$\text{Point Size} = (\text{Scalar Value} - \text{Min})\left(\frac{\text{Map size maximum} - \text{Map size minimum}}{\text{Max} - \text{Min}}\right) + \text{Map size minimum}$$

which means that is mapping values in interval [Min, Max] to interval [Map size minimum, Map size maximum].

Point Size Scale Field Properties

Function	Selects the type of scalar field function or field function monitor to show, such as temperature or velocity magnitude. See What Is a Field Function Monitor?	
Auto Min	Indicates the minimum value among all parts in the displayer (read-only).	
Auto Max	Indicates the maximum value among all parts in the displayer (read-only).	
Auto Range	Determines whether to set automatically the range of values to display.	
	Off	Does not use auto ranging, rather uses the values that are set in *Min* and *Max*.
	Min and Max Values	Automatically determines the range from the field function data.
	Min Value	Automatically determines the minimum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Max Value*.
	Max Value	Automatically determines the maximum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Min Value*.
Auto Range Mode	Chooses the domain for auto-range minimum and maximum value computation that is used for color coding.	

	Displayer Parts	Parts added to the particular displayer.
	Displayer Regions	Regions from which the displayer extracts data.
	All Regions	All regions in the simulation.
Clip	Determines how to display values outside of the range that is specified by *Min* and *Max*. See Changing the Scalar Range.	
	Off	Values above *Max* are displayed in the maximum color, and values below *Min* are displayed in the minimum color.
	Below Min, Above Max	Does not color values outside the range--uses the normal surface color.
	Below Min	Does not color values below *Min*--uses the normal surface color.
	Above Max	Does not color values above *Max*--uses the normal surface color.
Units	Selects the type of units to display. See Changing the Scalar Units.	
Min	Sets the minimum scalar value to include in the displayed range.	
Max	Sets the maximum scalar value to include in the displayed range.	
Map values	Activates or deactivates the setting of a displayed range with the *Min* and *Max* properties.	
	Activated	Mapping of values is active.
	Deactivated	Mapping of values is not active.

Map size minimum	The physical size of the particle that represents the Min scalar value. The size is relative to the overall dimension of the mesh.
Map size maximum	The physical size of the particle that represents the Max scalar value. The size is relative to the overall dimension of the mesh. Some graphics hardware may impose a limit on how large this size can be.

Copying Properties of a Scalar Field

It is possible to copy certain properties of one scalar field object onto another.

When properties are copied from one scalar field to another, the following properties are changed:

- *Auto Range*
- *Clip*
- *Function*

However, since the other properties are affected by which parts are displayed in the scene, you must copy the displayer properties to make the destination scalar field match the source one completely.

Defining a Custom Scalar

In addition to the standard system field functions, you can also display custom user-defined field functions.

Once you create the desired field function, simply select it from the list of available field functions as before.

Scalar Displayer Properties

Use the properties of the scalar displayer to control the contour style, point scale, and smoothness of the contours.

Contour Style	Controls how the values, represented by colors, are separated from each other. See Changing the Contour Style.

	Filled	Assigns values based on cell values. The values that are shown are the values stored at the cell centroids on the server. Scenes have a stepped appearance.
	Smooth Filled	Assigns values by node. See Understanding Smoothed Values.
	Smooth Filled + Lines	Adds lines to separate the colored values in the `Smooth Filled` style. Smoothing is applied to get the node values. See Understanding Smoothed Values.
	Line	Draws lines between the different values.
	Smooth Blended	Displays a smooth transition from one value to another. Smoothing is applied to get the node values.
Display Mesh	Controls the appearance of the mesh grid. See Displaying Mesh Outlines and Surfaces.	
	Activated	Shows the complete grid and adds the **Mesh Color** subnode to the displayer node to control the color of the mesh lines. This lets you examine the density and cell structure of various parts of the model, and thereby make judgments about accuracy. See Changing the Color of

		the Mesh Lines Using the Displayer .
	Deactivated	Hides the grid except for the outlines. When you choose this option and activate the *Surface* property, you generate a more realistic image of the model.
Opacity	Adjusts the transparency of the part. The settings range from 0.0 (invisible) to 1.0 (opaque), and can be controlled with a slider. This property allows you to show an overlay of different parts of the model. See Making Displayer Surfaces Transparent.	

Scalar Displayer Expert Properties

Representation	Selects a representation. Additional selections would appear as representations are added to the simulation.	
Point Scale Mode	Selects a mode for scaling points. See Working with Point Scale Modes.	
	Pixels	Activates the **PointSize** node to let you set the display size of points on the screen in pixels.
	Scalar	Activates the **Point Size Scalar Field** node to let you scale points based on a scalar function.
Line Width	Adjusts the width in pixels of lines that are used in the display, depending on the type of displayer. For example, in a scalar displayer, this setting would be effective only if the *Contour Style* property were set to `Line`. This property can be set with a slider.	

Lighting	Controls the lighting of the part. Accessing the property customizer opens the *Lighting* dialog. See Lighting Displayer Surfaces.	
Smooth Shade	Controls the transition between segments of the image.	
	Activated	Makes the transitions smooth for a more realistic effect.
	Deactivated	Leaves this function off.
Ignore Clip	Controls whether the parts in the displayer are affected by any enabled clip plane. See Using Clip Planes.	
	Activated	Shows the parts unaffected by the clip plane.
	Deactivated	Shows the parts that are trimmed by the clip plane.
Face Culling	Controls whether and how to cull polygons. See Culling Polygons.	
	None	Does not cull any polygons.
	Front	Culls (omits from the display) all polygons that are oriented toward the front of the view.
	Back	Culls (omits from the display) all polygons that are oriented toward the back of the view.
Immediate Mode	Controls whether parts in the displayer are rendered from CPU memory or from OpenGL display lists held in the GPU of the graphics hardware. Using display lists may provide faster performance for static data.	

	Activated	Parts are rendered from CPU memory.
	Deactivated	Parts are rendered from display lists on the graphics hardware.
Transform	Selects a transform to apply to all parts in the displayer. **None** is the default selection. Additional selections would appear when you create transforms. See Creating a New Transform. This feature allows for the creation of symmetric and periodic repeats, as well as arbitrary transforms to be used.	

Displaying Vector Quantities

Vector displayers are used to display glyphs that represent vector field functions.

The **Vector** displayer node, which has its own properties and shares a pop-up menu with other types of displayers, controls the display of vectors along the surface of a part.



In the icon, the arrow next to the paint brush signifies the depiction of vectors. The vector displayer also includes an undefined vector field node.

You can use this displayer to display the direction of quantities on boundaries without running the simulation.

As an alternative to glyphs, this displayer can visualize vector fields with a texture-based technique called line integral convolution.

See Also

Using the Scalar and Vector Bars
Initializing Manually

Creating Vector Displayers

Vector displayers are created automatically when a vector scene is created. Alternatively you may wish to add a vector displayer to any scene.

To add a vector displayer manually, open the **Scenes** node and the node of the chosen scene, right-click on the **Displayers** node and select **New Displayer > Vector**.



By default the vector displayer will contain no parts after creation and you must add parts to the displayer. This feature allows you to choose only those parts of the model on which you wish to view vectors.

Also by default a vector displayer uses velocity as the vector quantity. Glyphs are plotted on the parts' surfaces to represent the direction and the magnitude of the vector quantity at each boundary face centroid. The glyphs are also colored according to their magnitude.



The default name of the displayer is based on its type. However, renaming the displayer is possible.

Changing the Vector Quantity

By default the vector displayer uses the Velocity field function; however any vector field function can be used in a vector displayer.

Expand the **Scenes** node, and select any scene which includes a vector displayer (any default vector scene). Expand the **Displayers** and vector displayer nodes, and select the **Vector Field** node.



Choose the vector field function from the list in the object selection dialog. To activate the dialog, edit the *Function* property.



Alternatively you can change the vector field function from the visualization display. Right-click on the vector bar and select a vector field function from the same drop-down menu.



Displaying the Direction of Quantities on Boundaries

There is a simple way to show the direction of a quantity on a boundary even before you get a solution.

For example, the following steps show how to display the direction of velocity on an inlet in a geometry scene:

- Initialize the solution by clicking the (**Initialize Solution**) button in the toolbar or by using the **Solution** > **Initialize Solution** menu item.
- Create a vector displayer in the scene.



- Add the inlet through the drag-and-drop technique or using the Parts Selector from the properties of the new vector displayer. The velocity vectors appear in the scene. (A convenient way to view them completely is using the (**Make Scene Transparent**) toolbar button, as shown in the following screenshot.)



Working with Line Integral Convolution

Line integral convolution (LIC) is a texture-based technique for visualizing vector fields.

The **Line Integral Convolution** node represents the activation of this option in the STAR-CCM+ vector displayer. Select this node to access its properties.



Note: Line integral convolution is not supported on Intel graphics chipsets.

To activate LIC:

- Select `Line Integral Convolution` for the Display Mode property of a vector displayer node.

Cabral and Leedom first introduced LIC for imaging two-dimensional vector fields. Generally this technique involves convoluting a white noise image along streamlines computed from the vector field. In the resulting image, these streamlines cover the entire domain of the vector field. This technique has the advantage of being able to visualize large and detailed vector fields in a reasonable display area. Compared with simpler integration-type approaches, which entail following the flow vector at each point to produce a line, LIC produces a whole image at every step.



Limitations

- This feature is designed to use the Graphics Processing Unit (GPU) of your graphics hardware. It can work slowly in conjunction with software (mesa) rendering.
- For hardcopy with a magnification factor greater than 1, seams appear. These seams come from an attempt by STAR-CCM+ to perform tiled rendering, with results differing at the boundaries of adjacent tiles.



Line Integral Convolution Properties

Enhanced	Controls whether to use enhanced LIC computation.	
	Activated	Activates enhanced LIC computation. This option produces better-looking results, but is computationally more expensive.
	Deactivated	Leaves this feature deactivated.
Number Steps	Number of integration steps to perform during the LIC image generation. This property can be set with a slider. See Using a Slider.	
Blending Factor	The contribution of the LIC image when blending with the surface color, using linear interpolation between the LIC image color and the surface color. Ranges are from 0 to 1. 0 for surface color only, and 1 for LIC only	

	(grayscale image). This property can be set with a slider.
Step Size	The length of the integration step (in pixels). This property can be set with a slider.

Line Integral Convolution Bibliography

Cabral, B., Leedom, L.C., 1993. "Imaging Vector Fields Using Line Integral Convolution", Proceedings of ACM SigGraph 93, August 2-6, Anaheim, California, pp. 263-270.

Customizing Glyphs

The **Glyph** manager node allows you to customize the size and spacing of the glyphs.

Open the node to access its sub-nodes. Select the node to access its properties.



Properties of the Glyph Manager

Vector Spacing	Defines how the vectors are distributed. See Controlling the Spacing of Glyphs.	
	Geometric	Vectors are placed on cell faces, and the

		density of the vectors can be controlled by changing the *On Ratio* and *Randomize* properties.
	Uniform Screen	Vectors are placed at each grid point on a uniform grid that lies in the plane of the screen. The density of the vectors can be controlled by changing the grid dimensions with the *N Grid Points* property.
Vector Scale	Defines how to scale the vectors. See Changing the Size of Glyphs.	
	Absolute	Scales the vectors based on a specified length in the *Glyph Length* property of the **Absolute Length** node.
	Parts Size	Scales the vectors relative to the size of the parts in the *Glyph Length* property of the **Relative Length** node.
	Screen Size	Scales the vectors relative to the display size in the *Glyph Length* property of the **Relative Length** node.
	Direct	Scales the vectors based on a proportion specified in the *Glyph Length* property of the **Direct Length** node.
Vector Length	Quantity that determines vector length.	
	Vector Magnitude	Uses relative vector magnitude, that is, the scales of all vectors are

		fractions of the largest one.
	Constant Length	Lets you scale vectors on the order of the model rather than on the data.
Vector Style	Defines how the vector are displayed. See Changing the Style of the Glyphs.	
	2D Wireframe	A simple line.
	2D Filled	A line with finite diameter.
	3D Head, 2D Tail	Conical head.
	3D Head, 3D Tail	Conical head and tail.
	Cone	Cone only.
	Sphere	Sphere only.

Expert Properties of the Glyph Manager

Glyph Tip Scale	Sets the scale of the glyph heads. See Controlling the Scale of the Glyph Heads.	
Constant Tip Size	Controls whether to keep the size of the glyph heads constant.	
	Activated	Causes the head of the arrow in the vector displayer to be scaled equally for all arrows independent of the underlying velocity function.
	Deactivated	Leaves this function off.

Changing the Size of Glyphs

The glyph manager provides several features with which you can control the size of the glyphs.

These features include the **Relative Length**, **Absolute Length** and **Direct Length** objects, whose nodes have properties.



The **Relative Length** node appears when the *Vector Scale* property of the glyph manager is set to `Parts Size` or `Screen Size`, the **Absolute Length** node appears when this property is set to `Absolute`, and the **Direct Length** node appears with a setting of `Direct`.

In addition to the *Vector Scale* property of the glyph manager, its *Vector Length* property provides more flexibility. The relationships among these properties and the relative or absolute length of the vector glyphs are outlined in the table below.

		Vector Length	
		`Vector Magnitude`	`Constant Length`
Vector Scale	`Absolute`	• Glyph size remains the same independent of the parts or the display. • *Glyph Length* is a value; applies to size of largest glyph.	• Glyph size remains the same independent of the parts or the display. • *Glyph Length* is a value; applies to size of all glyphs.
	`Parts Size`	• Glyph size changes with parts size. • *Glyph Length* is a percentage; applies to size of largest glyph.	• Glyph size changes with parts size. • *Glyph Length* is a percentage; applies to size of all glyphs.
	`Screen Size`	• Glyph size remains the same while parts size changes. • *Glyph Length* is a percentage; applies to size of largest glyph.	• Glyph size remains the same while parts size changes. • *Glyph Length* is a percentage; applies to size of all glyphs.

		Vector Length	
	`Direct`	• Glyph size represents the vector quantity, based on the scale factor entered in the *Glyph Length* property. • *Glyph Length* is the scale factor of the unit chosen.	• Glyph size remains the same. • *Glyph Length* is a value; applies to size of all glyphs.

By default the glyphs are scaled using the magnitude of the vector. To change the scaling mode, expand the **Scenes** node and expand a scene which contains a vector displayer (all default vector scenes). Expand the **Displayers** node and the vector displayer node, select the glyph manager node and choose settings for the *Vector Length* and *Vector Scale* properties.

Although the **Relative Length**, **Absolute Length** and **Direct Length** nodes seem to have the same *Glyph Length* property, entries are retained separately for each combination of settings that are listed in the table above. For example, if your entry for *Glyph Length* with an `Absolute` vector scale causes the glyph to appear visibly longer than it did with a `Parts Size` (relative) vector scale, and then you switch back to `Parts Size`, you see the glyphs revert to their earlier size.

The Direct Length feature allows you to define a vector scale as a function of velocity. In this manner you can create animations where the vector length is consistent between frames, and also compare similar simulations with different maximum velocities.

Relative Length Properties

Glyph Length (%)	Specifies the length of the largest glyph or all of them as a percentage of parts size or screen size. (See preceding table.)

Properties of Absolute Length and Direct Length

Glyph Length	Specifies the length of the largest glyph or all of them as a specific value or a proportion of a unit. (See preceding table.)
Units	Selects the type of units to use.

Controlling the Spacing of Glyphs

The spacing and density of glyphs can be specified in two ways: geometric spacing and uniform spacing.

Geometric spacing places a vector glyph at the vertex or at the centroid of each cell face in the mesh. This type of spacing is well suited to cases where there is little variation in mesh cell size. Uniform spacing places a vector glyph at each intersection in a user-specified uniform grid that lies in the plane of the screen, and is particularly useful when the mesh has a large variation in cell size.

To access these options select a scene containing vectors, expand the **Displayers** node and the vector displayer node, and select the **Glyph** node. In the *Properties* window, select the *Vector Spacing* property and use the drop-down list to select either **Geometric** for geometric spacing, or **Uniform Screen** for uniform spacing. Depending on the selection either a **Geometric Spacing** node or a **Uniform Spacing** node is added to the tree, and these nodes have properties and are subject to common menu actions.



Geometric Spacing

When the **Geometric Spacing** node is selected, the density of glyphs can be controlled by increasing the *On Ratio* property (using integers only).



The number of glyphs that are shown are decreased as the integer value under the *On Ratio* property is increased. (In the following screenshots the vector style has been changed to **Cone** to show the effects more clearly.)



On Ratio = 1 (default number)



On Ratio = 11

The choice of glyphs to display can be randomized by checking the box for the *Randomize* property.

By default, the location of glyphs is the same as their native data location, but you can reposition the glyphs with the *Glyph Location* property. The native data location can be the cell face centroid or the vertex, depending on the type of field function and the part that is assigned to the displayer. For example, Position is a natively vertex-based field function, while Vorticity is cell-centroid based.

Vector glyphs are cell-centroid based by default under the following conditions:

- They appear in simulation files that were saved in versions of STAR-CCM+ prior to 9.06.
- Their native data location cannot be determined, due to either of the following causes:
  - The input parts of a derived part exhibit conflicting native data locations.
  - The internal field function definition is ambiguous or undefined.

Note: Certain parts only allow vertex-based vector glyphs: feature lines, streamlines, and isosurfaces.

Table 1: Geometric Spacing Properties

Randomize	Determines whether the pattern of vector glyphs is defined by the mesh grid.

	Activated	Ignores the structure of the model, including differences in density of the grid. This can result in a better distribution of vectors if the mesh is highly unstructured or has great variation in cell surface area.
	Deactivated	Takes the structure of the model into account.
On Ratio	The number of cells that display a single vector quantity. You can use this function to control the density of vectors, and thereby study broader or narrower patterns. At the default setting of 1, every cell displays a vector. Higher settings cause fewer vectors to be displayed. The value must be an integer.	
Glyph Location	Sets the location of the vector glyphs.	
	Cell	Sets the glyph location at the cell face centroid.
	Vertex	Sets the glyph location at the cell vertex.
	Native	Automatically uses the native data location, which depends on the assigned field function and the displayer part.

Uniform Spacing

When the **Uniform Spacing** node is selected, the density of glyphs can be controlled by setting the dimensions of the grid. The grid is bound by the extents of the scene display, so the values entered for the *N Grid Points* property control how many grid points are placed horizontally and vertically in the scene, or `u` and `v`, respectively.



A vector glyph is placed at each grid point on the parts in the display.

	
N Grid Points = [20, 20] (default)	N Grid Points = [5, 25]

Zooming in does not affect the grid size; in the following screenshot the *N Grid Points* property is still set to `[5, 25]` and the zoom level has been increased.



The uniform grid is linked to the size of the scene display, so shrinking the window increases the glyph density, and enlarging the window decreases the glyph density.

Table 2: Uniform Spacing Properties

N Grid Points	Specifies the u and v values of the grid, where u is the number of grid points in the horizontal direction and v is the number of grid points in the vertical direction.

Changing the Style of the Glyphs

By default the glyph style is a two-dimensional wire-frame arrow. Five alternative styles are available.

The styles are:

- Two-dimensional with a filled arrow head
- Three-dimensional arrow head with two-dimensional tail
- Three-dimensional arrow head with three-dimensional tail
- Cone
- Sphere

Expand the **Scenes** node and expand a scene node that includes a vector displayer. Expand the **Displayers** node and the vector displayer node, and select the **Glyph** node.



Select the glyph style from the drop-down menu under the *Vector Style* property in the *Properties* window.



This changes the appearance of the glyphs in any open vector scene display.



Controlling the Scale of the Glyph Heads

It is possible to control the scale of the head of the vector glyph independent of the tail.

Expand the **Scenes** node and expand a scene node that includes a vector displayer. Expand the **Displayers** node and the vector displayer node, and select the **Glyph** node.



In the *Glyph Tip Scale* expert property, enter a value between `0.0` and `1.0`.



This changes the scale of the glyph heads with no change to the appearance of the tail.



0.2 (default number)	0.6

You can also keep the size of the glyph heads constant.

Keeping the Glyph Head Size Constant

The head of the arrow in the vector displayer can be scaled equally for all arrows independent of the underlying velocity function. Expand the **Scenes** node and expand a scene node that includes a vector displayer. Expand the **Displayers** node and the vector displayer node, and select the **Glyph** node.



In the *Properties* window, activate the checkbox of the *Constant Tip Size* expert property.



This feature works as follows:

- The head scaling is controlled by the *Glyph Tip Scale* expert property only.
- The head size is never bigger than the actual arrow size; therefore the arrows still disappear where there is no flow.
- If the *Vector Style* property is set to `3D Head and Tail`, the tail radius can be scaled by changing the *Line Width* expert property. The *Line Width* expert property can be set with a slider.
- If the *Vector Style* property is set to `Cone` or `Sphere`, then the glyphs are unaffected by the *Glyph Tip Scale* or *Constant Tip Size* options, since the velocity function determines the scaling.

- If the *Vector Length* property is set to `Constant Length`, then the *Constant Tip Size* option has no effect on the display.

Controlling the Projection of the Glyphs

You can project the vector glyphs using only the normal component or the tangential component.

Expand the **Scenes** node and expand a scene node that includes a vector displayer. Expand the **Displayers** node and select the vector displayer node.



Select an option in the drop-down list of the *Projection Mode* property in the *Properties* window.



The options work as follows:

- `None`, the default, leaves the vectors with both normal and tangential components.
- `Normal` removes the tangential component of the vectors, projecting them with the normal component.

- `Tangential` removes the normal component of the vectors, projecting them with the tangential component.

Working with the Vector Field Object

The **Vector Field** node controls data in a scene with a vector displayer; by default it is set to velocity.

Use the properties of this node to control what is displayed and the range of values.



Vector Field Properties

Function	Selects the type of vector field function to show, such as velocity or flow direction.	
Min Magnitude	Sets the minimum value in the displayed range.	
Max Magnitude	Sets the maximum value in the displayed range.	
Clip	Determines how to display values outside of the range that is specified by *Min* and *Max*. See Changing the Scalar Range.	
	Off	Values above *Max* are displayed in the maximum color, and values below *Min* are

		displayed in the minimum color.
	Below Min, Above Max	Does not color values outside the range--uses the normal surface color.
	Below Min	Does not color values below *Min*--uses the normal surface color.
	Above Max	Does not color values above *Max*--uses the normal surface color.
Units	Selects the type of units to display. See Changing the Scalar Units.	
Auto Range	Determines whether to set the range of values to display automatically. See Changing the Scalar Range.	
	Off	Does not use auto ranging, rather uses the values that are set in *Min* and *Max*.
	Min and Max Values	Automatically determines the range from the field function data.
	Min Value	Automatically determines the minimum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Max Value*.
	Max Value	Automatically determines the minimum value from the field function data. The property switches automatically from `Min and Max Values`

		to this setting if you make an entry for *Min Value*.
Auto Min	Indicates the minimum value among all parts in the displayer (read-only).	
Auto Max	Indicates the maximum value among all parts in the displayer (read-only).	
Auto Range Mode	Chooses the domain for auto-range minimum and maximum value computation that are used for color coding.	
	Displayer Parts	Parts added to the particular displayer.
	Displayer Regions	Regions from which the displayer extracts data.
	All Regions	All regions in the simulation.

Copying Properties of a Vector Field

It is possible to copy certain properties of one vector field object onto another.

When properties are copied from one vector field to another, the following properties are changed:

- *Auto Scale*
- *Function*
- *Units*

However, since the other properties are affected by which parts are displayed in the scene, you must copy the displayer properties to make the destination vector field match the source one completely.

Defining a Custom Vector

In addition to the standard system field-functions, you can also display custom user-defined field functions.

Once you create the desired field function, simply select it from the list of available field functions as before.

Using the Color Bar to Modify the Vector Display

The vector bar of the scene display shows the scale of the magnitude of the chosen vector field function.

The appearance of the vector bar can be manipulated using the color bar in the object tree.



The range that is shown on the vector bar corresponds to the full range of the vector magnitude on all parts in the displayer. The units may be changed in the same way as the units of a scalar displayer.

Coloring Glyphs Based on Scalar Values

You can display scalar values on the vector glyphs.

The **Color Field** node, which has properties and is subject to common menu items, appears under the vector displayer node when the displayer's *Color Mode* property is set to `Scalar Field`.



It colors the glyphs based on the values of the scalar field function that you select.

Color Field Properties

Function	Selects the type of scalar field function to show, such as temperature or velocity magnitude.	
Auto Min	Indicates the minimum value among all parts in the displayer (read-only).	
Auto Max	Indicates the maximum value among all parts in the displayer (read-only).	
Auto Range	Determines whether to set the range of values to display automatically. See Changing the Scalar Range.	
	Off	Does not use auto ranging, rather uses the values that are set in *Min* and *Max*.
	Min and Max Values	Automatically determines the range from the field function data.
	Min Value	Automatically determines the minimum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Max Value*.
	Max Value	Automatically determines the minimum value from the field function data. The property switches automatically from `Min and Max Values` to this setting if you make an entry for *Min Value*.
Auto Range Mode	Chooses the domain for auto-range minimum and maximum value computation that is used for color coding.	

	Displayer Parts	Parts added to the particular displayer.
	Displayer Regions	Regions from which the displayer extracts data.
	All Regions	All regions in the simulation.
Clip	Determines how to display values outside of the range that is specified by *Min* and *Max*. See Changing the Scalar Range.	
	Off	Values above *Max* are displayed in the maximum color, and values below *Min* are displayed in the minimum color.
	Below Min, Above Max	Does not color values outside the range--uses the normal surface color.
	Below Min	Does not color values below *Min*--uses the normal surface color.
	Above Max	Does not color values above *Max*--uses the normal surface color.
Units	Selects the type of units to display. See Changing the Scalar Units.	
Min	Sets the minimum value in the displayed range.	
Max	Sets the maximum value in the displayed range.	

Vector Displayer Properties

Use the properties of the vector displayer to control characteristics such as display mode, colors, and projection.

Display Mode	Controls how the vector field is displayed.

	Glyph	Uses glyphs to display vectors. Activates the Glyph manager node. See Customizing Glyphs.
	Line Integral Convolution	Uses a texture-based technique to display vectors. Activates the Line Integral Convolution node. See Working with Line Integral Convolution.
Color Mode	Controls the arrangement of colors.	
	Vector Magnitude (Default)	Colors the glyphs based on values of a selected vector field.
	Scalar Field	Adds the **Color Field** subnode to the displayer node to color glyphs based on values of a selected scalar field. See Coloring Glyphs Based on Scalar Values.
	Constant	Adds the **Color** subnode to the displayer node to display using a custom color. See Changing the Color of the Surface Using the Displayer.
	Region	Colors the glyphs by region.
Projection Mode	Controls the projection of vectors relative to the surface. See Controlling the Projection of the Glyphs.	
	None (Default)	Leaves the vectors with both normal and tangential components.

	Normal	Removes the tangential component of the vectors, projecting them with the normal component.
	Tangential	Removes the normal component of the vectors, projecting them with the tangential component.
Opacity	Adjusts the transparency of the part. The settings range from 0.0 (invisible) to 1.0 (opaque), and can be controlled with a slider. This property allows you to show an overlay of different parts of the model. See Making Displayer Surfaces Transparent.	

Vector Displayer Expert Properties

Representation	Selects a representation. Additional selections would appear as representations are added to the simulation.	
Line Width	Adjusts the width in pixels of lines that are used in the display, depending on the type of displayer. For example, in a scalar displayer, this setting would be effective only if the *Contour Style* property were set to `Line`. This property can be set with a slider.	
Lighting	Controls the lighting of the part. Accessing the property customizer opens the *Lighting* dialog. See Lighting Displayer Surfaces.	
Smooth Shade	Controls the transition between segments of the image.	
	Activated	Makes the transitions smooth for a more realistic effect.
	Deactivated	Leaves this function off.
Ignore Clip	Controls whether the parts in the displayer are affected by any enabled clip plane. See Using Clip Planes.	

	Activated	Shows the parts unaffected by the clip plane.
	Deactivated	Shows the parts that are trimmed by the clip plane.
Immediate Mode	Controls whether parts in the displayer are rendered from CPU memory or from OpenGL display lists held in the GPU of the graphics hardware. Using display lists may provide faster performance for static data.	
	Activated	Parts are rendered from CPU memory.
	Deactivated	Parts are rendered from display lists on the graphics hardware.
Face Culling	Controls whether and how to cull polygons. See Culling Polygons.	
	None	Does not cull any polygons.
	Front	Culls (omits from the display) all polygons that are oriented toward the front of the view.
	Back	Culls (omits from the display) all polygons that are oriented toward the back of the view.
Transform	Selects a transform to apply to all parts in the displayer. **None** is the default selection. Additional selections would appear when you create transforms. See Creating a New Transform. This feature allows for the creation of symmetric and periodic repeats, as well as arbitrary transforms to be used.	

Displaying Streamlines

Streamline displayers are used to display derived streamline parts in a scene (provided you have created such parts). They cannot be used to create streamlines.

The **Streamlines** displayer node, which has its own properties and shares a pop-up menu with other types of displayers, controls the appearance of those streamlines.



In the icon, the lines next to the paint brush signify the depiction of streamlines.

The scene that displays streamlines includes a streamline bar which functions like a scalar bar. This bar shows the scale of the magnitude of the field function that is used to generate the streamlines (by default velocity magnitude). Its appearance can be manipulated using the color bar in the object tree.

Creating Streamline Displayers

Streamline displayers are not created automatically through the creation of any scene and have to be added.

To add a streamline displayer, open the **Scenes** node and the node of the chosen scene, right-click on the **Displayers** node and select **New Displayer** > **Streamline**.



By default streamline displayers contain no parts and only derived streamline parts may be added. Thus before you can display streamlines you must create at least one derived streamline part. Any number of derived streamline parts can be added to a streamline displayer.



The default name of the displayer is based on its type. However, renaming the displayer is possible.

Changing the Streamline Style

The streamlines can be displayed using various styles: lines, ribbons, or tubes.

By default, the streamlines are displayed as lines. To change the style of the streamlines, expand the **Scenes** node and expand a scene which contains a streamline displayer. Expand the **Displayers** node and select the streamline displayer node.



Set the style in the drop-down menu under the *Mode* property in the *Properties* window.



The `Tubes` option displays the streamlines as cylinders with the diameter of the cylinder being governed by the *Width* property.



The `Ribbons` mode displays the streamlines as flat surfaces with a width governed by the *Width* property.



In addition, the ribbons can be twisted based on the setting of the *Orientation* property.



Changing the Size of the Streamlines

If either the ribbon or tube mode is used to display the streamlines, the size of the ribbons or tubes can be changed.

Expand the **Scenes** node and expand a scene which contains a streamline displayer. Expand the **Displayers** node and select the streamline displayer node.

Adjust the width of the ribbons or tubes using the *Width* property in the *Properties* window.



Displaying a Scalar on Streamlines

Streamlines can be colored using any scalar field function.

Expand the **Scenes** node and expand a scene which contains a streamline displayer. Expand the **Displayers** node and the streamline displayer node, and select the **Scalar Field** node.



By default the function is not set. To set the scalar field function, choose from the list in the object selector dialog. To activate the dialog, edit the *Function* property.



The scalar values appear on the streamlines in the display.



The units and range of the scalar field function can be changed in the same way as in the Scalar Displayer.

Adding Arrows to Streamlines

Arrows can be added to streamlines in order to give them a direction.

Expand the **Scenes** node and expand a scene which contains a streamline displayer. Expand the **Displayers** node and the streamline displayer node, and select the **Markers** node.



By default, there are no arrows that are displayed on streamlines. To display arrows on streamlines, select the desired function in the *Properties* window.



The following properties can be set in the *Markers - Properties* window:

- *Display start point markers*: activating this checkbox displays spheres at the start point of the streamlines.



- *Display end point markers*: activating this checkbox displays conical arrow heads at the end points, pointing in the direction of the flow.



- *Display intermediate markers*:
  - **None**: no arrow appears between the start points and the end points on the streamlines.
  - **Geometric Spacing**: selecting this adds arrows, equally spaced out, between the start points and end points on the streamlines.



Selecting this option creates a node in the **Markers** folder where the distance between the arrows can be controlled using the *Step size* in the *Geometric Distance - Properties* window.



- **Integration time based**: this adds arrows between the start points and end points on the streamlines that are displayed with relation to the time step.



Selecting this option creates a node in the **Marker** folder where the gap between the arrows can be controlled by adjusting the *Time interval* in the *Integration time interval - Properties* window. This option is useful for visualizing the speed of the flow.



The color of the arrows is controlled by the scalar field chosen for the streamlines or by the color selected in the *Properties* window when the Color Mode is set to **Constant**. The size of the spheres and cones is scaled according to the size of the streamlines and can be controlled using the *Width* property in the *Properties* window.

Creating Streamlets

Reducing the streamline length and increasing the number of seed points can produce streamlets (short streamlines) as shown below.



One way of achieving this effect is to create a constrained streamline on a plane or part, reduce the integration time of the streamlines (*Initial Integration Step* and *Maximum Propagation Step*), and increase the number of Seed Points. The integrator properties that are used for the above example are shown below.



To achieve the streamlets effect shown above, it is best to only have start_point and end_point markers displayed, although it is possible to have streamlets with no markers displayed. The streamlets effect can be achieved with any type of streamlines.

Streamline Displayer Properties

Use the properties of the streamline displayer to control characteristics such as streamline style, width, and colors.

Width	Controls the size of the ribbons or tubes. The units are of the mesh that is imported into the simulation. See Changing the Size of the Streamlines.	
Orientation	Defines how ribbon streamlines are twisted. Only valid if the *Mode* property is set to `Ribbons`.	
	Normals	Normal twisting.
	Angular Velocity	Use the angular velocity of the vector to twist then ribbon.
	Vorticity	Use the local vorticity to twist the ribbon.
	Rotation	Use rotation to twist the ribbon.
Mode	Defines the style of streamline. See Changing the Streamline Style.	
	Lines	Shows streamlines in a linear format. This is the default setting.
	Ribbons	Shows streamlines as flat bands. This format depicts the local vorticity of particles moving through flow. Unlike lines, ribbons provide surfaces that can be viewed clearly with the Surface Lighting feature. See Lighting Displayer Surfaces.
	Tubes	Shows streamlines as tubes. Tubes are more suitable for presentation purposes,

		since, unlike ribbons, they do not indicate particle orientation.
Color Mode	Controls the arrangement of colors.	
	Scalar (Default)	Colors the streamlines according to the selected scalar field function.
	Constant	Adds the **Color** subnode to the displayer node to display surfaces using a custom color.
	Region	Colors the surfaces by region.
Opacity	Adjusts the transparency of the part. The settings range from 0.0 (invisible) to 1.0 (opaque), and can be controlled with a slider. This property allows you to show an overlay of different parts of the model. See Making Displayer Surfaces Transparent.	

Streamline Displayer Expert Properties

Representation	Selects a representation. Additional selections would appear as representations are added to the simulation.	
Line Width	Adjusts the width in pixels of lines that are used in the display, depending on the type of displayer. For example, in a scalar displayer, this setting would be effective only if the *Contour Style* property were set to `Line`. This property can be set with a slider.	
Lighting	Controls the lighting of the part. Accessing the property customizer opens the *Lighting* dialog. See Lighting Displayer Surfaces.	
Smooth Shade	Controls the transition between segments of the image.	
	Activated	Makes the transitions smooth for a more realistic effect.

	Deactivated	Leaves this function off.
Ignore Clip	Controls whether the parts in the displayer are affected by any enabled clip plane. See Using Clip Planes.	
	Activated	Shows the parts unaffected by the clip plane.
	Deactivated	Shows the parts that are trimmed by the clip plane.
Face Culling	Controls whether and how to cull polygons. See Culling Polygons.	
	None	Does not cull any polygons.
	Front	Culls (omits from the display) all polygons that are oriented toward the front of the view.
	Back	Culls (omits from the display) all polygons that are oriented toward the back of the view.
Immediate Mode	Controls whether parts in the displayer are rendered from CPU memory or from OpenGL display lists held in the GPU of the graphics hardware. Using display lists may provide faster performance for static data.	
	Activated	Parts are rendered from CPU memory.
	Deactivated	Parts are rendered from display lists on the graphics hardware.
Transform	Selects a transform to apply to all parts in the displayer. **None** is the default selection. Additional selections would appear when you create transforms. See Creating a New Transform.	

	This feature allows for the creation of symmetric and periodic repeats, as well as arbitrary transforms to be used.

Displaying Dynamic Fluid Body Interaction (DFBI)

It is possible to display DFBI entities such as bodies (including the center of mass), external forces and moments, and body couplings (springs and catenaries) in a scene.



To make the DFBI displayer available, activate DFBI in the simulation.

A DFBI displayer can be created using the pop-up menu of the displayer manager. The **DFBI** displayer node, which has properties and shares a pop-up menu with other displayer types, controls the display of this content. Specifically, it contains subobjects that let you change the size of glyphs.



Creating DFBI Displayers

DFBI displayers are not created automatically through the creation of any scene and have to be added.

To add a DFBI displayer, open the **Scenes** node and the node of the chosen scene, right-click on the **Displayers** node and select **New Displayer > DFBI**.



By default DFBI displayers contain no parts and only 6-DOF bodies, external forces and moments, or body couplings may be added. Thus before you can display DFBI you must create at least one of these related objects. Any number of 6-DOF bodies, external forces and moments, or body couplings can be added to a DFBI displayer.

The default name of the displayer is based on its type. However, renaming the displayer is possible.

Changing the Size of Glyphs

The DFBI displayer provides several features with which you can control the size of the glyphs.

These features include the **Relative Size**, **Absolute Size** and **Direct Size** objects, whose nodes have properties and are subject to common menu items.



The **Relative Size** node appears when the *Glyph Scale* property of the DFBI displayer is set to `Model Size` or `Screen Size`, the **Absolute Size** node appears when this property is set to `Absolute`, and the **Direct Size** node appears with a setting of `Direct`.

Vectors are scaled according to the *Glyph Scale* property of the DFBI displayer and its *Vector Length* property. The relationships among these properties and the relative or absolute length of the vector glyphs are outlined in the table below.

		Vector Length	
		`Vector Magnitude`	`Constant Length`

		Vector Length	
Glyph Scale	`Absolute`	• Glyph size remains the same independent of the model or the display. • *Vector Glyph Length* is a value; applies to size of largest vector glyph.	• Glyph size remains the same independent of the model or the display. • *Vector Glyph Length* is a value; applies to size of all vector glyphs.
	`Model Size`	• Glyph size changes with model size. • *Vector Glyph Length* is a percentage; applies to size of largest vector glyph.	• Glyph size changes with model size. • *Vector Glyph Length* is a percentage; applies to size of all vector glyphs.
	`Screen Size`	• Glyph size changes with screen size. • *Vector Glyph Length* is a percentage; applies to size of largest vector glyph.	• Glyph size changes with screen size. • *Vector Glyph Length* is a percentage; applies to size of all vector glyphs.
	`Direct`	• Glyph size represents the vector quantity,	• Glyph size remains the same independent

		Vector Length	
		based on the scale factor entered in the *Vector Glyph Length* property. • *Vector Glyph Length* is the scale factor of the unit chosen.	of the model or the display. • *Vector Glyph Length* is a value; applies to size of all vector glyphs.

By default the vector glyphs are scaled to constant length. To change the scaling mode, expand the **Scenes** node and expand a scene which contains a DFBI displayer. Expand the **Displayers** node, select the DFBI displayer node and choose settings for the *Vector Length* and *Glyph Scale* properties.

Although the **Relative Size**, **Absolute Size** and **Direct Size** nodes seem to have the same properties, entries are retained separately for each combination of settings that are listed in the table above. For example, if your entry for *Vector Glyph Length* with an `Absolute` glyph scale causes the glyph to appear visibly longer than it did with a `Model Size` (relative) vector scale, and then you switch back to `Model Size`, you see the glyphs revert to their earlier size.

The Direct Length feature allows you to define a vector scale as a function of the absolute value of the vector. In this manner you can create animations where the vector length is consistent between frames, and also compare similar simulations with different maximum vector lengths.

Relative Size Properties

Vector Glyph Length (%)	The user-defined length of a vector glyph, in percent of the model size or the screen size, depending on the setting of the *Glyph Scale* property.
Coupling Glyph Diameter (%)	The user-defined diameter of a body coupling glyph, in percent of the model size or the screen size, depending on the setting of the *Glyph Scale* property.

Center Of Mass Glyph Diameter (%)	The user-defined diameter of a center of mass glyph, in percent of the model size or the screen size, respectively, depending on the setting of the *Glyph Scale* property.

Properties of Absolute Size and Direct Size

Units	The unit in which the size of a glyph is specified.	
Vector Glyph Length	The user-defined length of a vector glyph, given in the unit specified in the *Units* property.	
Coupling Glyph Diameter	The user-defined diameter of a body coupling glyph, given in the unit specified in the *Units* property.	
Center Of Mass Glyph Diameter	The user-defined diameter of a center of mass glyph, given in the unit specified in the *Units* property.	

DFBI Displayer Properties

Use the properties of the DFBI displayer to control characteristics such as DFBI entities, glyph scale, and colors.

Color Mode	Controls the arrangement of colors.	
	Type (Default)	Colors the glyphs according to their type (body hull, body center of mass, force vector, moment vector, linear spring coupling, catenary coupling). The default colors are as follows: • Body hull: slate gray • Body center of mass: dark orange • Force vector: red

		• Moment vector: blue • Linear spring coupling: steel blue • Catenary coupling: beige
	Distinguish Forces and Moments	Colors forces and moments according to their specific type (External Fluid Force and Moment, External Gravity Force, Force CM, Moment CM, Generalized Force, Propulsion Force, Damping Force, Damping Moment, Linear Spring Force, Catenary Force). All other parts are colored as in the setting `Type (Default)`.
	Constant	Adds the **Color** subnode to the displayer node to display parts using a custom color. See Changing the Color of the Surface Using the Displayer.
Glyph Scale	Defines how to scale the glyphs. See Changing the Size of Glyphs.	
	Absolute	Scales the glyphs based on the specified sizes in the **Absolute Size** node.
	Model Size	Scales the glyphs relative to the maximum body length based on the specified sizes in the **Relative**

		Size node. The maximum body length is the length of the largest body in the simulation. The length of a body is its largest extent measured in the three directions given by the axes of the body coordinate system.
	Screen Size	Scales the glyphs relative to the display height based on the specified sizes in the **Relative Size** node.
	Direct	Scales the glyphs based on the proportions specified in the **Direct Size** node.
Bodies	Controls the appearance of bodies.	
	Activated	Bodies are displayed in the scene.
	Deactivated	Bodies are not displayed.
Forces	Controls the appearance of forces.	
	Activated	Forces are displayed in the scene.
	Deactivated	Forces are not displayed.
Moments	Controls the appearance of moments.	
	Activated	Moments are displayed in the scene.
	Deactivated	Moments are not displayed.
Couplings	Controls the appearance of body couplings.	
	Activated	Body couplings are displayed in the scene.

	Deactivated	Body couplings are not displayed.
Vector Length	Mode that determines vector length.	
	Vector Magnitude	Uses relative vector magnitude, that is, the scales of all vectors are fractions of the largest one.
	Constant Length	Lets you scale vectors on the order of the model rather than on the data.
Opacity	Adjusts the transparency of the parts. The settings range from 0.0 (invisible) to 1.0 (completely opaque), and can be controlled with a slider. This property allows you to show an overlay of different parts of the model. See Making Displayer Surfaces Transparent.	

DFBI Displayer Expert Properties

Vector Tip Scale	Sets the scale of the vector heads. See Controlling the Scale of the Glyph Heads.	
Lighting	Controls the lighting of the parts. Accessing the property customizer opens the *Lighting* dialog. See Lighting Displayer Surfaces.	
Ignore Clip	Controls whether the parts in the displayer are affected by any enabled clip plane. See Using Clip Planes.	
	Activated	Shows the parts unaffected by the clip plane.
	Deactivated	Shows the parts that are trimmed by the clip plane.
Immediate Mode	Controls whether parts in the displayer are rendered from CPU memory or from OpenGL display lists held in the GPU of the graphics hardware. Using display lists may provide faster performance for static data.	

	Activated	Parts are rendered from CPU memory.
	Deactivated	Parts are rendered from display lists on the graphics hardware.
Face Culling	Controls whether and how to cull polygons. See Culling Polygons.	
	None	Does not cull any polygons.
	Front	Culls (omits from the display) all polygons that are oriented toward the front of the view.
	Back	Culls (omits from the display) all polygons that are oriented toward the back of the view.
Transform	Selects a transform to apply to all parts in the displayer. **None** is the default selection. Additional selections would appear when you create transforms. See Creating a New Transform. This feature allows for the creation of symmetric and periodic repeats, as well as arbitrary transforms to be used.	

Animating Scenes

You can animate information in a scene by activating the animation object within the displayer of the scene.

This object is represented by the **Animations** node, which has properties and is subject to common menu items.



Which Kinds of Animation Are Available?

Scalars can be animated with the sweeping of a plane section or an iso surface. However, their other characteristics are controlled through the scalar displayer.



The scalars that are shown here represent wall distance.

Vector glyphs are animated with the **Animations** node. They can also move while a sweeping derived part is moving. However, their other characteristics are controlled through the vector displayer.



For example, the preceding screenshot shows pulsing vectors with a glyph length setting of 0.5% and a field function of **Velocity**.

Streamline animation allows you to depict flow in motion with a series of objects that are called tracers, which move along the streamline paths. Each tracer consists of a particle, or head, and a tail.



Many of the characteristics of the tracers can be adjusted through the properties of the **Streamline Settings** node. However, the shape of the tail is set by the streamline style, and the color indicates the streamlines' scalar selection. For example, the preceding screenshot shows tracers with a streamline setting of **Ribbons** and a scalar setting of **Velocity: Magnitude**.

Controlling the Animation

The various types of animation of a scene share several common controls.

These controls are:

- Setting up the Animation
- Controlling the Movement of the Animation
- Controlling the Speed of the Animation

Setting up the Animation

To set up the animation, select the **Animations** node of the displayer.



In the *Properties* window, change the *Animation Mode* property from `Off` to `Pulse Glyphs` for vectors (as shown in the following screenshot), to `Tracers` for streamlines, or to `Sweep` for a derived part.



This step adds a subnode to the **Animations** node.

The **Animation** toolbar is available by default. If not, display it.



Controlling the Movement of the Animation

After you set up the animation, click the (**Play/Pause Animation**) toolbar button to start it. This causes the animation to move at its default settings.



To pause the motion, click the (**Play/Pause Animation**) button again. This button toggles between motion and freeze-frame.

When the animation is in freeze-frame, you can move the animation ahead or back one frame at a time. Click the (**Single Step Forward**) to move ahead and the (**Single Step Backward**) button to move back.

To stop the motion, click the (**Stop Animation**) button. In streamline animations, this restores the standard appearance of the streamlines.

Controlling the Speed of the Animation

To make the tracer, vector glyph, or sweeping derived part move faster or slower, select its node and enter a number of seconds in the *Cycle time* property.



Additionally, the number of frames per second can be controlled in the scene attributes.

Recording the Animation onto Separate Files

It is possible to record a segment of the animation onto separate files: either a commonly used type of video file, `.avi`, or a set of frames in `.png` graphics format.

For best results, follow these steps:

1. Make sure that the particles appear and move exactly as you want them. Adjust all the properties and view the particles in motion.
2. Press the (**Write Movie**) toolbar button in the **Animation** toolbar.
3. In the *Write animation* dialog, prepare the specifications for the recording:



1. **Frame Rate** -- Decide on the number of frames per second (fps). More frames show smoother movement but take more time and computing power for recording and playback. Fewer frames allow the file to run more efficiently, but may show breaks in the movement. The default rate for animation in STAR-CCM+ is 25 fps. To try an alternate rate, change it in the properties of the **Animation** object in the scene attributes, and then preview the animation again as in Step 1.
2. **Start Time** -- Start the animation at a time other than t = 0.
3. **Animation Length** -- Decide on the time length of the animation.

4. **Size** -- Specify the movie size for viewing.
5. *Movie Type* -- Select the method for recording the animation:
   - **AVI movie** -- This is the default setting.
   - **Directory of PNG Frames** -- Click this radio button if you prefer to record each individual frame as a graphic file. Type the directory path name into the **File Name** text box of the dialog. The frames are saved in that directory as the name `frameXXXXX.png` where `XXXXX` is the five-digit frame number (padded with zeros if necessary).



Note: If you choose to create `.png` files, consider the previous specifications carefully, since you may create many graphics files. For example, the default settings would record 25 frames per second for 15 seconds, generating 375 `.png` files.

4. Navigate to the file location of your choice. For an `.avi` movie, enter a file name as you would for saving any other file.
5. Click **Save**. STAR-CCM+ generates the file(s).

   The `.png` files can then be animated with your media program of choice. The `.avi` file can be played in a typical Windows Media Player, or its equivalent in Linux.



Using Vector Glyphs

You can animate vectors by activating the **Animations** > **Settings** node.



Animations Settings Properties

Cycle Time (sec)	Sets the total time for a glyph's travel from the inlet to the outlet.

Using Streamline Tracers

You can animate streamlines by activating the **Streamline Settings** node.



Streamline Settings Properties

Cycle time (sec)	Sets the total time for a tracer's travel from the inlet to the outlet.
Delay between tracers (sec)	Sets the time elapsed between repetitions of the appearance of a particular tracer.
Head size	Specifies the radius of the particle in the units that are used in the simulation.
Tail length (sec)	Defines the length of the tail in terms of how far the particle travels in the specified length of time.

Changing the Tracer Frequency

You can make the particles appear in a sparse or dense pattern.



The *Delay between tracers* property specifies how soon the next set of tracers is injected into the flow.

Changing the Tracer Dimensions

You can change the dimensions of the tracer.

- The *Tail length* property controls the length of the tail. Select the **Streamline Settings** node and enter a number of seconds.



  These seconds are the time taken for the particle head to travel a certain distance: the longer the time, the longer the distance, and hence the longer the tail length.

- The *Head size* property specifies the radius of the particle in the units that are used in the simulation.



Using Sweeping Derived Parts

You can use two kinds of derived parts to animate scalars or vectors in a scene: an isosurface or a plane section.

The settings nodes for these parts are named according to the type of part, but they have identical properties.



Each controls the animation for all the parts of that type in the displayer. If the displayer contains an isosurface and a plane section, the nodes that represent them can both exist at the same time in an **Animations** node.

The following screenshot shows an example of a plane section with vectors.



Additionally, you can animate the vector glyphs themselves (in a separate displayer) while they are being swept on the derived part.

This example shows an isosurface being used to sweep scalars.



You can choose between two types of animation cycles in the properties:

- Repeat -- this default setting causes the sweeping motion to go from the beginning to the end of the range, and then start again from the beginning.
- Bounce -- this setting specifies a sweep from beginning to end, then back to the beginning.

If you have multiple plane sections in a displayer, they must all have the same cycle time, but an isosurface in that displayer can have a different cycle time. To have different cycle times among plane sections, put them in different displayers.

To obtain consistent color mapping over the whole animation time, use a fixed scalar range.

Note: The animated derived parts require default units of STAR-CCM+ for the start and end values, not the preferred units of your simulation. For example, with a plane section, the units are always in meters. To animate with an isosurface, work with the default STAR-CCM+ units for the particular scalar that you have in that isosurface.

Using Precomputed Sweep Animation

You can animate sweeping derived parts with precomputed steps. By default, STAR-CCM+ must access the server during each animation step to calculate values and geometry for the derived part. When you activate precomputation, STAR-CCM+ prepares these calculations before you start the animation, allowing you to animate the derived part without constant server access. Precomputed steps improve the performance of the animation with a higher average frame rate, but the precomputation requires more time, and the animation requires more memory. The requirements for time and memory increase proportionally to the specified number of steps.



To prepare a sweep animation with precomputed steps:

1. Select the **Animations** node in the displayer node.
2. Make sure that *Animation Mode* is set to **Sweep**.
3. Expand the **Animations** node and select the sub-node that corresponds to the derived part (plane section or isosurface).
4. Activate the *Precompute Steps* property.

   In the **Precompute Options** sub-node that appears, specify the number of precomputed steps in the *Number of Steps* property.

Note: While increasing *Number of Steps* generates a smoother motion in the animation, it also increases memory consumption.

Sweep Settings Properties

Cycle time (sec)	Defines how many seconds it should take to sweep over the whole configured range.	
Auto Range	Determines whether to set the range of values to display automatically.	
	Activated	Automatically determines the range of values. This range begins with the starting offset for sections and the lowest iso value for

		isosurfaces. It ends with the ending offset for sections and the highest iso value for isosurfaces.
	Deactivated	Does not use auto ranging, rather uses the values that are set in *Min* and *Max*.
Precompute Steps	Controls whether to use precomputed steps for the sweep animation.	
	Activated	A user-defined number of steps is precomputed over the configured sweep range. A suitable precomputed step is chosen and rendered while the sweep animation is running in the scene display, resulting in an increased average frame rate.
	Deactivated	All sweeping parts are continuously updated while the sweep animation is running. The playback performance within a scene display is limited by the computation performance of sweeping derived parts.
Start Value	The starting offset for sections and the lowest iso value for isosurfaces. See Working with a Single Section.	
End Value	The ending offset for sections and the highest iso value for isosurfaces. See Working with a Single Section.	

Repeat Mode	Controls how the derived part moves once a sweep is finished.	
	Repeat	Sweeps from the beginning to the end of the range, and then starts over from the beginning (default). One cycle is from beginning to end.
	Bounce	Sweeps from beginning to end, then back to the beginning. One cycle is from beginning to end to beginning, so a bouncing derived part moves twice as fast as a repeating one.

Animations Properties

These properties vary depending on the type of displayer.

Animation Mode	Controls the animation feature for vectors and streamlines.	
	Pulse	Makes glyph animation available. Adds a Settings subnode to the **Animations** node (vector displayer only). See Using Vector Glyphs.
	Tracers	Makes tracer animation available. Adds a Streamline Settings subnode to the **Animations** node (streamline displayer only). See Using Streamline Tracers.

	Sweep	Makes sweep animation available. Adds an **iso** and/or **plane section** subnode to the **Animations** node, depending on whether the displayer contains a derived part of the respective type. See Using Sweeping Derived Parts.
	Off	Leaves animation feature off.

Using Displayers

This section defines what a displayer is and how to use displayers to build scenes in STAR-CCM+.

The **Displayers** node, which has a pop-up menu, is the manager object for all displayers in a scene.



It is included automatically in every scene. The displayer manager can hold any number of displayers up to the memory limits of the computer.

What Is a Displayer?

A displayer is the basic building block used to build a scene. It provides a flexible way of defining and controlling the graphics entities appearing in a scene.

There are four types of displayers available in STAR-CCM+, and they share a common pop-up menu:

- A part displayer to control the appearance of geometry
- A scalar displayer to control the appearance of scalar contours
- A vector displayer to control the appearance of vector glyphs
- A streamline displayer to control the appearance of streamlines

A displayer includes one or more parts, an attribute such as a scalar or vector field, and other properties such as surface lighting, color, opacity, normals for smooth-shading, and transforms. A displayer takes these parts and attributes and produces the appropriate graphics entities to represent the parts in the scene.

A displayer provides a convenient way to control the appearance of many parts with the same attributes. A displayer also provides a way to

view the same part in many different ways in the same scene. For example, consider viewing the scalar field Mach number on a wing and fuselage. The displayer manages the wing and fuselage parts, and produces an image of the Mach number contours on those parts. If you want now to see pressure on those parts, it is a simple matter of changing the scalar field of the displayer to pressure. Next, consider viewing the wing surface with velocity vectors. In this case, you can create two displayers, one a geometry displayer showing the wing surface, the other a vector displayer showing the wing velocity vectors.

Displayers can be transformed and repeated, allowing for any user-defined or solver-defined transform to be applied to the parts within the displayer. This allows for the parts in the scene to be repeated or moved from their default location within a given scene.

A displayer can be renamed or deleted using the standard techniques. When a displayer is deleted, all part representations controlled by that displayer are removed from the scene as well.

Displayer Pop-Up Menu

Use the pop-up menu of the displayer to toggle its visibility.

Toggle Visibility	Renders all the parts in the displayer visible or invisible with a single command.

See also Common Pop-Up Menu Items.

Displaying Cell Values in Scalar Scenes

To view the cell values for a region on one side of the interface only, you must create a threshold derived part and use this in the scalar displayer.

When a case contains multiple regions, such as in a Conjugate Heat Transfer (CHT) analysis, a scalar scene will show averaged values for boundaries that belong to interfaces between regions. STAR-CCM+ calculates the averaged values based on solution data from both sides of the interface.

The recommended steps to follow are:

- Create a threshold derived part. The Create Threshold in-place dialog appears.
- For the *Input Parts* property, select the parts that you want to view the cell values on.

- For the *Scalar* property, select the field function that you want to view on the cells.
- Ensure that the *Extraction Mode* is set to **All Between**.
- Click the **Query** button to populate the *Min* and *Max* values.

A completed dialog is shown below.



- Create the threshold.
- Select the threshold part for the scalar displayer in the active scene.

- In the scalar scene, apply the field function that you want to view.

An example scene that displays the cell values is shown below.



In this case, the eight memory chips on the graphics card have the same volumetric heat source, while the GPU chip has a different value. The remaining parts have no volumetric heat source.

Managing Parts in Displayers

The **Parts** manager object contains all the parts in a scene displayer.

The parts are represented as a distinct group of part nodes in the object tree. The **Parts** node has its own properties and pop-up menu.



This section explains other techniques for adding parts to a scene and removing them.

Hiding and Showing Parts with the Displayer Parts Manager

The pop-up menu of the **Parts** manager node includes actions for hiding and showing parts in a scene.



Adding Parts to Displayers

This section explains two techniques for adding parts: drag-and-drop and the object selector.

Adding Parts Using Drag-and-Drop

You can drag-and-drop a region, boundary, or derived part into a scene, even if that scene is empty (no displayers). To do this, select a part from

the **Regions** and **Boundaries** nodes and/or the **Derived Parts** node. Drag the selected part into the visualization display.



Once the part has been dragged to the visualization display, STAR-CCM+ responds depending on the number of displayers in the scene:

- If there are no displayers, a geometry displayer showing surfaces is created automatically by default, and the part is added to it.
- If there is only one displayer, the part is added automatically to that displayer.
- If there are more than one displayer, a pop-up menu lets you choose which displayer should receive the part.



After the part has been dropped into the visualization display, it has the properties that are prescribed by the displayer. In the example below,

the displayer has the mesh display enabled so the mesh is displayed on the new part.



If a region part is added to a displayer, all of the boundaries of the region are also added to the displayer. Multiple parts can be added at the same time by multi-selecting nodes.

Alternatively parts can be dragged into the displayer node.

Adding Parts Using the Object Selector

Parts may also be added to the displayer using the object selector. Expand the **Scenes** node, the node of the chosen scene, the **Displayers** node and the displayer node to which you wish to add the part. Select the **Parts** node.



Open the object selector dialog by clicking the right column of the *Parts* property in the *Properties* window.

Click **Close** when you have finished selecting parts.

See Also

Using Threshold Parts and Cell Surfaces

Removing Parts from Displayers

Parts can be removed from a displayer with one of multiple techniques.

The available techniques are:

- Using the same part selector that is used for adding parts to the displayer
- Right-clicking a displayer part node
- Directly from the visualization display - right-click the part in the display, select **Displayers** from the pop-up menu to activate a set of checkboxes, and deactivate the **in** checkbox that corresponds to the displayer from which you want to remove the part.



Working with Displayer Part Nodes

The Parts manager node of a displayer contains subnodes resembling those of regions, boundaries, derived parts and/or geometry parts.

These subnodes, which have their own properties and pop-up menu, are a visual representation of the parts in the displayer.



Hiding and Showing Parts

You can use these nodes to control the visibility in the scene of the parts they represent.

To hide a part from view, either deactivate the checkbox in its *Visible* property, or right-click its node and select **Hide** from the pop-up menu.



To reverse this step, activate the checkbox in the node's *Visible* property or select **Show** from the pop-up menu.

To make only one or more parts visible, right-click the corresponding node(s) and select **Show Only This Part** from the pop-up menu. This has the same effect as selecting all the other parts and hiding them.



When you hide or show an entire displayer, STAR-CCM+ remembers the settings you applied with these actions to the individual parts.

Removing Parts

To remove a part from the displayer completely, right-click its node and select **Remove** from the pop-up menu.



Displayer Part Properties

Visible	Shows or hides the part in the scene.	
	Activated	The part is visible.
	Deactivated	The part is hidden.

Displayer Part Pop-Up Menu

Highlight	Highlights the part in the display by changing its color and labeling it.
Hide	Executes the hide part action.
Show	Executes the restore hidden part action to redisplay a part that was hidden by the hide part action.
Show Only This Part	Isolates one or more selected parts by executing the hide part action for all the other parts in the displayer.
Align View to Part Normal	Aligns the view to the surface normal of a part. See Aligning to a Normal.

Remove	Executes the remove part from displayer action to remove the selected part from the relevant displayer's **Parts** node.

Properties of the Parts Manager

Use the properties of the Parts manager to view and control which parts appear in the displayer.

Parts	List of parts in this displayer. Editing the property accesses in-place object selector to add or remove parts. See Using the Part Selector.

Pop-Up Menu of the Parts Manager

Use the pop-up menu of the Parts manager to hide and restore parts in the scene.

Hide All Parts	Executes the hide all parts action to make all parts that are controlled by that displayer invisible in the scene. See Hiding All Parts in a Scene.
Restore Hidden Part	Executes the restore hidden part action to redisplay a part that is hidden by the hide part action. You choose which part to restore from a submenu of this menu item. See Working with Displayer Part Nodes.
Restore All Hidden Parts	Executes the restore hidden parts action to redisplay all parts that are hidden by previous hide part actions. See Working with Displayer Part Nodes.

See also: Common Pop-Up Menu Items.

Changing the Color of the Surface

By default the part surfaces are colored according to their boundary type. However it is possible to change the color of the parts included within a displayer.

Color arrangements can be as follows:

- A different color for each region

- The same color for all parts in the displayer
- A different color for each part

Expand the **Scenes** node and expand any scene that includes a geometry displayer (any default geometry, scalar or vector scene, or an empty scene in which you create a geometry displayer). Expand the **Displayers** node and select the geometry displayer node.



For coloring by region, set the *Color Mode* property to `Distinguish Regions`.



To apply the displayer color, set it to `Constant`. This makes the **Color** subnode available.



To change the displayer color, either enter the RGB value under the *Color* property of this node or click ... (**Custom Editor**) to open the color selector dialog.



If you wish to color different parts with different colors, set the *Color Mode* property to `Distinguish Inputs`.



Changing the Color of Parts in a Displayer

You can change the colors of surfaces or mesh lines of parts within an individual displayer.

The **Color** and **Mesh Color** nodes are subnodes of the displayer node when activated, and let you control the surface color and mesh color, respectively.



Changing the Color of the Surface Using the Displayer

The **Color** node is activated when the displayer's *Color Mode* property is set to `Constant`.

This node controls the color of surfaces on parts that are included in the displayer.

Color Properties

Color	Sets the color to display a surface. The property customizer executes the define color action. See Defining Colors.

Changing the Color of the Mesh Lines Using the Displayer

The **Mesh Color** node is activated when the checkbox of the displayer's *Mesh* property is activated.

This node controls the color of mesh lines on parts included in the displayer.

Mesh Color Properties

Mesh Color	Sets the color of the mesh lines. The property customizer executes the define color action. See Defining Colors.

Using the Color Bar

Scene displays include multi-colored bars to map values to colors for scalars, vectors or streamlines.

You can control the appearance of the bar in the scene with the displayer's color bar object.

See Also

Using the Scalar and Vector Bars

What Is the Color Bar?

The color bar controls the appearance of a scalar bar or vector bar in a scene display.

The **Color Bar** node has properties and is subject to common menu items.



Color Bar Properties

Color Map	Selects colors to indicate values from lowest to highest. This property lets you customize the use of colors in the scalar scene.	
	blue-red	From blue to red.
	blue-red-band	Rainbow color map with the opacity decreasing from center towards minimum and maximum (band pass filter).
	blue-red-highpass	Rainbow color map with the opacity increasing from minimum to maximum (high pass filter).
	blue-red-lowpass	Rainbow color map with the opacity

		increasing from maximum to minimum (low pass filter).
	cool-warm	From blue to red, in softer shades.
	cool-warm-band	Cool-warm color map with the opacity increasing from center towards minimum and maximum (band elimination filter).
	grayscale	From black to white. This option is useful for illustrating a printed report.
	high-contrast	Shades of colors that contrast sharply for easy viewing.
	spectrum	Colors in order of the visible light spectrum.
Levels	Sets the number of levels into which the range of data should be divided. Higher numbers produce finer blends between values. The default setting is 32.	
Width	Controls the width of the color bar. The number is a fraction of the display width; 1.0 equals the full width of the display.	
Font	Sets the font of the color bar text.	
Background	Activates or deactivates the color bar background.	
	Activated	Background color is active.
	Deactivated	Background color is not active.
Background Color	Controls the color of the background. The property customizer executes the define color action. See Defining Colors.	
Shadow	Activates a shadow effect for the text of the color bar.	

	Activated	Shadow effect is on.
	Deactivated	Shadow effect is off.
Reverse	Reverses the sequence of the selected colormap. See Using Colormaps.	
	Activated	Color sequence is reversed (for example, red to blue rather than blue to red).
	Deactivated	Uses the default color sequence.
Text Position	Sets the position of the text relative to the color bar.	
	Preceding	The text appears before (above) the color bar.
	Succeeding	The text appears after (below) the color bar.
Height	Sets the height of the color bar and its label.	
Title Height	Specifies a height for the title (field function label) of the color bar relative to the display height.	
Label Height	Specifies a height for the numeric labels of the color bar relative to the display height.	
Orientation	Defines orientation of the color bar. See Changing the Position and Orientation of the Bar.	
	Horizontal	Color bar is horizontal.
	Vertical	Color bar is vertical.
Position	Specifies the position in the display of the visualization display as X and Y coordinates from 0 to 1.	
Visible	Shows or hides the color bar.	
	Activated	Color bar is visible.
	Deactivated	Color bar is hidden.
Scale Mode	Activates a logarithmic scale mode for the color bar. See Switching the Scale Mode.	

	Linear	Uses linear scaling.
	Log10	Uses logarithmic scaling.
Text Color	Controls the color of the text. The property customizer executes the define color action. See Defining Colors.	

Expert Properties

Opacity	Adjusts the transparency of the color bar. The settings range from 0.0 (invisible) to 1.0 (opaque). This property allows you to show the bar as an overlay in front of the model.	
Label Format	Defines the numeric formatting of the numbers that are displayed along the color bar. The property customizer activates a dialog with interactive controls. See Controlling the Labels. Alternatively, text elements that can be entered here are identical to those described in the section What Is the Number Format?	
Ramp Type	Sets the shape of the table ramp to either linear, S-curve, or square-root.	
	Linear	Specifies a linear ramp, .
	S-curve	Uses the S-curve, which tails off gradually at each end: $y = \frac{\sin((x-1/2)\pi)+1}{2}$
	Sqrt	Uses the square root, $y = \sqrt{x}$.
Number of Labels	Controls the number of value labels along the color bar.	

Changing the Position and Orientation of the Bar

The scalar bar or vector bar can be repositioned in the scene.

Click in the interior of the bar and drag it to the desired location.



As the bar is moved close to the edges of the visualization display, the orientation of the bar automatically changes between horizontal and vertical.



The bar may also be repositioned by entering the normalized coordinates of the location within the visualization display.

Expand the following nodes: the **Scenes** node, the node of a scene with a scalar or vector displayer, the **Displayers** node within that scene, and the displayer node. Select the **Color Bar** node.



Either directly enter the normalized coordinates under the *Position* property in the *Properties* window or click the property customizer (...) to enter the coordinates through a text input dialog.



The orientation (horizontal or vertical) of the bar does not change automatically but you may change the orientation by selecting from the drop-down menu under the *Orientation* property.



The bar may be removed completely from the display by deactivating the checkbox of the *Visible* property, or by dragging the bar to locations outside the scene. If you drag the bar out of view by mistake, simply activate the *Visible* property's checkbox to restore it.

Resizing the Bar

The scalar bar or vector bar can be resized by clicking the perimeter of the bar and dragging to enlarge or shrink it.



As the bar is enlarged, the text automatically changes size to maintain a good balance between the text size and bar size.

The size of the bar may also be changed by entering the width and height.

Expand the following nodes: the **Scenes** node, the node of a scene with a scalar or vector displayer, the **Displayers** node within that scene, and the displayer node. Select the **Color Bar** node.



In the *Properties* window, enter the height.



Then enter the width.



Switching the Scale Mode

It is possible to switch the scale of the scalar bar or vector bar from linear to logarithmic, or vice versa. The logarithmic scale is helpful if you have a wide range of values to display.



To use this feature, expand the following nodes: the **Scenes** node, the node of a scene with a scalar or vector displayer, the **Displayers** node within that scene, and the displayer node. Select the **Color Bar** node.



The scale can be changed using the drop-down menu under the *Scale Mode* property in the *Properties* window.



This feature works only if the function field ranges are positive or negative. If the ranges cross zero, then the display and the bar are disabled, showing the geometry model colored by default colors.

Changing the Text Properties of the Color Bar

The color bar text consists of the title (such as Velocity Magnitude) and the labels (numerical values). Controls are available for the entire text, the title, and the labels.

- To gain access to the color bar controls, select the **Color Bar** node.



Controlling the Text

Controls for the text (title and labels combined) are available in the properties:

- To change the font, or to make the text bold and/or italic, click ... (**Custom Editor**) of the *Font* property.
- To change the color of the text, select a color from the list in the *Text Color* property. Alternatively, click ... (**Custom Editor**) to activate the color selector.
- To control whether the text is above or below the color bar, use the *Text Position* property:
  - The `Preceding` option (default) puts the text above the bar.
  - The `Succeeding` option puts the text below the bar.
- To adjust the size of the entire text, resize the bar in the display.

Controlling the Title

The title of the color bar comes from the name of the selected scalar or vector.

- To specify a height for the title, enter a value for the *Title Height* in the properties. The value is a fraction of the height of the display window, where 1.0 is the full height.

Controlling the Labels

- To change the number of labels in the color bar, enter a different value for the *Number of Labels* expert property.
- To make other changes to the labels:
  - To specify a height for the labels, enter a value for the *Label Height* property. The value is a fraction of the height of the display window, where 1.0 is the full height.
  - In the *Label Format* expert property, enter text elements that are identical to those described in the section What Is the Number Format?.
  - Alternatively, click ... (**Custom Editor**).

    This action activates the *Color Bar - Label Format* dialog.



- To adjust the appearance of the labels, use the options in the *Overall Formatting* group box:
  - To justify the labels toward the left end of the color bar, activate the **Left-Justified Labels** checkbox.

  - To indicate whether the values are positive or negative, activate the **Print Sign** checkbox.
  - To include zeros at the left of the labels (for a uniform number of digits), activate the **Zero Padding** checkbox.
  - To display decimal points for all numbers, including integers, activate the **Print Decimal Point** checkbox.
- To specify the minimum number of digits overall, and the number of digits after the decimal point, use the text boxes in the *Precision Options* group box. Press `<Enter>` after making an entry in each of these text boxes.
- In the *Style Options* group box, select a style from the drop-down list:
  - Best Fit - this default option presents the data in a visually optimal manner.
  - Decimal - this option displays the full decimal values based on the entry in the **Number Decimal Digits** text box.
  - Scientific - this option displays the values in scientific notation. To use an uppercase "e", activate the **Use Upper Case** checkbox.

Changing the Background Color of the Bar

To change the background color of the scalar bar or vector bar, select the **Color Bar** node.

Within the *Properties* window, activate the *Background* property checkbox.



Then, click the ellipsis button to the right of the *Background Color* property.



This opens the *Background Color* dialog, where colors can be selected.



Changing the Colormap of the Bar

Various colormaps (combinations of colors) are available for you to use in the display.

By default the colormap of the scalar bar or vector bar uses the blue-red gradient.



To switch from the default, expand the following nodes: the **Scenes** node, the node of a scene with a scalar or vector displayer, the **Displayers** node within that scene, and the displayer node. Select the **Color Bar** node.



Select a colormap from the drop-down menu under the *Color Map* property in the *Properties* window.



The options besides the default include:

- `Cool-warm`



- `Grayscale` (useful for illustrating a printed report)



- `High-contrast` (shades that contrast sharply for easy viewing)



- `Spectrum`



In addition, the *Reverse* property lets you display the color in the opposite sequence when you activate the checkbox.



For example, the default blue-red setting appears as red-blue when the *Reverse* property is activated.



The number of color bands can be changed using the *Levels* property.

In the *Properties* window, enter the new number of color levels.



The change is visible in the display for both the bar and the visible parts.



32 levels (default number)	64 levels

Copying Properties of a Color Bar

When properties are copied from one color bar node to another, all of the properties are applied so that the destination color bar matches the source one.

Because color bars in scalar and vector scenes have the same set of properties, you can copy color bar properties between the two types of scenes.

Using a Combination of Displayers

Combinations of displayers can be used to construct visualizations with a combination of surfaces, scalars, vectors and streamlines.

By default: geometry, scalar, and vector scenes use more than one displayer to show both the outline of the boundary parts and either surfaces, scalars and vectors.

This section describes a simple example in which a combination of displayers are used to construct a scene.

Initial Geometry Scene

Initially a geometry scene is opened which includes two geometry displayers. The first displayer displays the outline of all the boundary parts. In the example, some parts have been removed from both of the displayers. The scene is viewed through the symmetry plane to display the inside of the cowl and manifold.



Displayer Transformation

Both displayers are then transformed using a transform automatically generated from the symmetry plane part. The opacity of the surface geometry displayer is decreased to provide transparency. This is useful to maintain surface definition but allow inspection of parts beyond the surface.

Addition of a Scalar Displayer

A scalar displayer is then added to the scene and transformed using the same transform. Only the cowl and splitter parts are added to the displayer and the Pressure field function is displayed on the surfaces.



Addition of a Second Scalar Displayer

Another scalar displayer is then created with only one side of the manifold added. Again the Pressure field function is displayed on the surface. The only difference with the previously added scalar displayer is that it is not transformed, maintaining the ability to view inside the manifold. Since both scalar displayers use the same field function the second scalar bar is rendered invisible.



Addition of a Streamline Displayer

Finally a streamline displayer is added to the scene with the same transform applied and the Velocity Magnitude field function used to color the streamlines. The streamline bar is moved to the left-hand side of the visualization display which automatically changes its orientation.



The result is a complex visualization constructed using a few simple steps.

Lighting Displayer Surfaces

The surface lighting properties of all parts in a displayer can be adjusted using a combination of parameters.

The parameters are:

- Ambient lighting - accounts for light reflected or scattered from other surfaces in the scene
- Diffuse lighting - accounts for the angle of incidence of light onto a surface
- Specular lighting - accounts for the direct reflection of the light source off a shiny object
- Specular power - adjusts the intensity of specular lighting

The combination of these four parameters represents how the surface reflects and absorbs the scene lighting. By default, the ambient parameter is set to 0.0, and the diffuse parameter is set to 1.0, resulting

in a relatively good setting to distinguish easily the depth and proximal relationships between three-dimensional objects in the scene. The default setting may not always be suitable for viewing scalar contours, as surfaces which are perpendicular to the view normal may become dark. In this case, it is recommended to set the ambient parameter to 1.0 and the diffuse parameter to 0.0 to see the scalar colors clearly.

Your settings of the specular and specular-power parameters depend on how you want your simulated surface to appear. The specular parameter controls how the light is reflected from the surface. Specular light comes from a particular direction, and it tends to bounce off the surface in a preferred direction. The level of specularity is essentially how “shiny” a surface is. Shiny metal or plastic has a high specular component, and chalk or carpet has almost none. The specular-power parameter controls the tightness of the light source: higher values of specular power produce a tighter light source, while lower values broaden the light source.

To change any of the parameters expand the **Scenes** node, any scene node and the **Displayers** node. Select any displayer which includes surfaces (geometry or scalar surfaces).



The value of each parameter may be entered directly under the *Lighting* expert property in the *Properties* window.



Alternatively, you can click ... (**Custom Editor**) to open the *Lighting* dialog through which the parameters can be set using sliders.



Hiding and Showing All Parts in a Displayer

You can render all the parts in the scene displayer visible or invisible with a single command.

To do this, right-click the displayer node and select **Toggle Visibility**.



When you use this feature, STAR-CCM+ remembers the visibility settings of the individual parts in the scene. Hence, when you render the displayer visible again, the individually hidden parts remain hidden.

You can also control the visibility of parts with the *Parts Visibility* expert property of the displayer. It lets you show all parts, hide all parts, or revert to the visibility settings of the individual parts.



Making Displayer Surfaces Transparent

Displayer surfaces can be made transparent. This is useful when surface definition is desirable but you wish to see what is beyond the surface.

To make displayer surfaces transparent expand the **Scenes** node, any scene node and the **Displayers** node. Expand any displayer which displays surfaces (geometry or scalar displayers).



Expand the **Expert** category in the *Properties* window and click the right half of the *Opacity* property.



Use the slider or the text box to set a value between 0 and 1.0, where 0 means the parts are fully transparent and 1.0 means the parts are fully opaque.



In the example below two geometry displayers are used to display different parts with different degrees of transparency.



Smooth Shading Displayer Surfaces

The normals of the displayer surfaces determine how light is reflected from the surface and thus the appearance of the surface.

By default the normals are calculated at the center of the boundary face polygons which result in a faceted surface appearance. When the normals are calculated at boundary face vertices, the resulting appearance of the surface is smooth.



Faceted surface	Smooth surface

To obtain smooth part surfaces expand the **Scenes** node, any scene node and the **Displayers** node. Expand any displayer which displays surfaces (geometry or scalar displayers).



Expand the **Expert** category in the *Properties* window and activate the checkbox of the *Smooth Shade* property.



Culling Polygons

You can control the display of polygons depending on their orientation.

Each face in a polygonal surface has two sides: outward-facing, and inward-facing, determined by the vertex ordering of the polygon. A displayer's *Face Culling* expert property tells the display to not draw (to cull) faces which have the orientation selected in the property's drop-down list. This results in faster display because only half the faces are actually drawn.



No culling	Front faces culled

The best uses for this are mainly with baffles, where turning on this feature speeds up display in some cases. Culling can also be used to check the surface for defects; that is, turn on backface culling and you can see holes, which are possible signs of an improperly oriented surface.

Copying Displayers

A displayer can be copied. When a displayer is copied, all properties are copied, resulting in an exact duplicate of the original.

You can also copy the properties of one displayer onto another, provided they are of the same type.

Copying Properties of a Displayer

When properties are copied from one displayer to another, all of the properties are applied so that the destination displayer matches the source displayer.

Note: It is not possible to copy properties between different types of displayers; for example, the properties of a geometry displayer cannot be copied onto a vector displayer.

Pop-Up Menu of the Displayers Manager

These pop-up menu items are available when you right-click the **Displayers** manager node.

New Displayer	Creates a displayer of the type you select in the submenu.	
	Geometry	Executes the add displayer action to add a geometry displayer to the scene. With more than one geometry displayer in the same scene, you can show various parts in different modes. For example, a rotor blade can be

		shown as a solid (shaded) part with no grid, while the rotor hub can be shown with grid cells (all edges) only.
	Scalar	Executes the add displayer action to add a new scalar displayer to the scene. See Creating Geometry Displayers. Creating a scalar displayer allows you to select a different part on which to show scalar data, either the same type of data or a different type.
	Vector	Executes the add displayer action to add a vector displayer to the scene. See Creating Scalar Displayers. Creating a vector displayer allows you to select a different part on which to show vectors, either the same type of vector or a different type.
	Streamline	Executes the add displayer action to add a new streamline displayer to the scene. See Creating Streamline Displayers. No parts other than streamline derived parts can be selected by this displayer. If your simulation has no streamline derived parts, you can still

		create this displayer. However, you have nothing to add to its Parts manager.

Working with Volume Rendering

Volume rendering is used to display (semitransparent) volumetric objects that are defined in a 3D domain.

Defining a Resampled Volume Part

The resampling process generates an approximation of the result values that reside in the underlying 3D finite volume regions.

The resampled volume displays results in a 3D domain. For each position in the 3D domain color and opacity values are specified.

The **Resampled Volume** node has its own properties and pop-up menu, and is contained in the **Derived Parts** node of the simulation object tree.



Note: The resampled volume part can be rendered only in the scalar displayer. To display the resampled volume data, set the *Representation* expert property of the scalar displayer to `Volume Mesh`.

Resampled Volume Properties

Origin	A Cartesian position vector that defines the origin of the resampled volume. You can include units in the entry, or select the units in the dialog that is activated by clicking ... (**Property Customizer**).
Rotation Axis	Sets the rotation axis for the resampled volume.
Size	Sets the size of the resampled volume in three dimensions. You can include units in the entry, or select the units in the dialog that is activated by clicking ... (**Property Customizer**).

Parts	Selects the input parts as the source for deriving this resampled volume.
Rotation Angle	Specifies the angle of rotation along the specified rotation axis in radians.

Resampled Volume Expert Properties

Cell Size to Voxel Ratio	Determines how fine the resulting structured grid will be.
Minimum Voxel Size	Sets the minimum size of a structured cell.

Resampled Volume Pop-Up Menu

Edit Part in Current Scene	Executes the define resampled volume action allowing you to set the properties of this part interactively. This option is only available if a scene is open.

Creating a Resampled Volume Part

You can create a resampled volume with the pop-up menu.

To create a resampled volume, right-click the **Derived Parts** node and select **New Part > Resampled Volume...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new resampled volume part to appear in the simulation tree immediately. Define the part using the properties.

Editing a Resampled Volume Part

Multiple techniques are available for editing a resampled volume.

You can edit a resampled volume using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Resampled Volume In-Place Dialog

This dialog facilitates the definition of the resampled volume with interactive features.

If certain basic conditions are met when you start the process to create a resampled volume part or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.


truck_cab
simulation
scene/plot
edit
Uniformly Resample 3D Region
Input Parts
[fluid]
Select...
Coordinates
Origin
Size
X 0.28799998760223
0.57599997520446
Y 0.06480000168085
0.12960000336170
Z 0.0 m
0.0 m
Rotation Axis
Rotation Angle
X 1.0 m
0.0 radian
Y 0.0 m
Z 0.0 m
Cell Size to Voxel Ratio 0.1
Minimum Voxel Size 0.001 m
Preview
Voxel Count:
Snap to part
Display
No Displayer
New Scalar Displayer
Existing Displayer
Scalar 1
Create
Close
Help

To select the input parts that you want to include in the resampled volumes, click the **Input Parts** button.

To specify the dimensions and position of the bounding box, set the following properties in the *Coordinates* group box:

- Origin
- Size
- Rotation Axis
- Rotation Angle

To control the number of voxels in the resampled volume, set the following properties:

- Cell Size to Voxel Ratio
- Minimum Voxel Size

Before you create the resampled volume, you can get a graphical preview of it and a preliminary estimate of the number of voxels. To do this, click **Preview**.

To position the bounding box along a part, activate **Snap to part**. This feature is available when the resampled volume tool is active.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Working with Volume Rendering Settings

Additional controls let you choose a rendering mode, activate the lighting mode, and specify accuracy.

When you add a resampled volume part to a scalar displayer, that displayer acquires a new sub-node, **Volume Rendering Settings**. Use the properties of this sub-node to control how the data is displayed.



Volume Rendering Settings Properties

Rendering Mode	Allows you to select a mode of rendering.	
	Direct	Every scalar sample value is mapped to the displayer opacity plus color, and is then attenuated along the view direction.
	Maximum Opacity Projection	A transparent mode with high contrast to highlight volume areas that contain scalars with high opacity. An example is when velocity magnitude is mapped to a high-pass color map. The transparent mode reveals additional points of interest, regardless of the opacity attenuation which is depending on the viewpoint. However, the projective nature of

		this mode removes clues for depth perception.
	Level Surface	Allows you to choose between single or multiple surfaces. Activates a Level Surface Mode sub-node.
Lighting Mode	Allows you to apply displayer lighting to the volume rendering. The primary application for the Local Lighting mode is when the *Rendering Mode* property is set to `Level Surface`. In the Level Surface mode, standard lighting will be applied to the level surface exactly as if it were a proper surface. Additionally, local lighting is used in the Direct rendering mode for a more "surface-like" image without actually reverting to simple surfaces.	
	None	Does not use displayer lighting.
	Local Lighting	Applies the settings of the displayer surface lighting. The advantage of this technique is that details in the scalar field can be discerned more easily in many cases. The following Marschner-Lobb examples of volume rendering illustrate the benefit of local lighting. • Local lighting:

		 • No local lighting: 
Quality	Determines how accurately the resampled volume is processed during ray-casting.	

Working with the Level Surface Mode

Volume rendering in STAR-CCM+ is based on a ray-casting approach that is implemented in graphics hardware. The level-surface volume rendering mode tries to capture user-defined level values of interest during ray-casting through the resampled volume derived part.



The **Level Surface Mode** node allows you to choose between single and multiple surfaces for volume rendering in its properties.



Level Surfaces Compared with Isosurfaces

Level surfaces are in principle similar to isosurfaces. However, level surfaces provide a better visual accumulation of partly transparent

surfaces than isosurfaces with the same opacity settings. The data inputs, data extraction approaches, and storage locations of these types of objects are different. The effects of the inputs and extraction depend on the mesh, the native data location of the scalar field, and the accuracy settings for the resampled volume part.

	
Level Surfaces	Isosurfaces

Feature	Data Input	Extraction Approach	Hardware Employed	Storage	Location
Level Surfaces	Resampled Volume derived part	Ray-Casting	Graphics Hardware / GPU	Image buffer in graphics memory	Client
Isosurfaces	Discrete meshed region or part	Polygonalization of intersected mesh edges	CPU	Polygonal surface mesh in main memory	Client and Server

See also Examples of Effects of Volume Rendering Settings on Level Surfaces.

For an interactive exploration of simulation results, level surfaces have the following advantages over isosurfaces:

- Typically they update very fast, at the speed of volume rendering.
- They provide better results that are based on the opacity that you specify in the displayer and the colormaps.

In comparison to isosurfaces, the level surfaces have the following limitations:

- They cannot show a secondary scalar field on their surface.
- They cannot serve as an input to other derived parts or reports.
- They support a maximum of 32 level values.
- Their accuracy depends on the resolution of the underlying resampled volume. An example is bumpy disruptions of an otherwise flat surface. If such artifacts occur, use isosurfaces.
- They do not work with software-based OpenGL implementations such as Mesa OpenGL, or Microsoft OpenGL (remote desktop).

Specifying Level Surfaces

- To work with a single level surface, select the **Level Surface Mode > Value** node and enter a value for *Level Value*.
- To work with multiple level surfaces that are equally spaced in a range:
  - Select the **Level Surface Mode** node and set *Level Surface Mode* to **Multiple (range)**.
  - Select the **Level Surface Mode > Values (range)** node and set the following properties:
    - *Number of Levels*
    - *Minimum Level Value*
    - *Maximum Level Value*
- To work with multiple level surfaces that you arbitrarily set:
  - Select the **Level Surface Mode** node and set *Level Surface Mode* to **Multiple (arbitrary)**.
  - Select the **Level Surface Mode > Values (arbitrary)** node and set the properties using one of the following techniques:
    - Enter settings for *Level Values*.
    - Click ... (**Property Customizer**) of the *Level Values* property.



In the dialog that appears, enter values as follows:

- Set the number of level surfaces in the spin box.
- If you know the exact values to display, enter them manually in the text box and click **OK**.
- If you want STAR-CCM+ to generate values, click **Query**, then click **Auto Generate Values**. Click **OK**.

  The **Query** action retrieves the scalar range from the server. The **Auto Generate Values** action produces an even set of values in the text box according to the number of level surfaces that you specify.

Examples of Effects of Volume Rendering Settings on Level Surfaces

Lighting mode:



Local lighting on



Local lighting off

Opacity:



Displayer Opacity = 0.05



Displayer Opacity = 0.75

Colormaps:



Blue-red (default)



Blue-red-band

Level Surface Mode Properties

Level Surface Mode	The mode for level surface volume rendering, specifying whether a single surface or multiple surfaces are created.	
	Single	Uses one surface (default).
	Multiple (range)	Selects a range technique for creating multiple surfaces.

	Multiple (arbitrary)	Selects an arbitrary technique for creating multiple surfaces.

Working with Volume Rendering Colormaps

Some colormaps are designed for use with volume rendering. These colormaps provide a mapping of the colored scalar to opacity.



If existing color maps are chosen for volume rendering, then they apply full opacity to all scalar values, visually resulting in a compact volume block. The volume rendering colormaps should not be used for coloring in other scalar displayers that show surfaces or streamlines.

The colormap types are as follows:

- Blue-red-band -- rainbow color map with the opacity decreasing from center towards minimum and maximum (band pass filter)
- Blue-red-highpass -- rainbow color map with the opacity increasing from minimum to maximum (high pass filter)
- Blue-red-lowpass -- rainbow color map with the opacity increasing from maximum to minimum (low pass filter)
- Cool-warm-band -- cool/warm color map with the opacity increasing from center towards minimum and maximum (band elimination filter)

Viewing the Final Result

When you finish creating your resampled volume, you can view it in the assigned scene.

Once you have finished creating a resampled volume part, you can rotate and zoom in the display to get a clear view of the resampled volume, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display. The resampled volume part can be rendered only in the scalar displayer.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Creating a Resampled Volume Part

You can create a resampled volume with the pop-up menu.

To create a resampled volume, right-click the **Derived Parts** node and select **New Part > Resampled Volume...**



Typically, an in-place dialog would be used for this process, and it would appear as soon as you execute this pop-up menu action.

However, if the conditions for activating this dialog are not met, then this action causes the new resampled volume part to appear in the simulation tree immediately. Define the part using the properties.

Editing a Resampled Volume Part

Multiple techniques are available for editing a resampled volume.

You can edit a resampled volume using the in-place dialog. Right-click it (either in the *Graphics* window, or in the object tree), and select **Edit Part in Current Scene**.



In the dialog, click **Apply** to implement your changes.

However, if certain conditions are not met, this menu item is disabled (grayed out). In such a situation, you can simply click the node of the derived part in the simulation tree and use the properties.

Using the Resampled Volume In-Place Dialog

This dialog facilitates the definition of the resampled volume with interactive features.

If certain basic conditions are met when you start the process to create a resampled volume part or edit an existing one, a new in-place dialog appears in the *edit* tab of the simulation window.


truck_cab
simulation
scene/plot
edit
Uniformly Resample 3D Region
Input Parts
[fluid]
Select...
Coordinates
Origin
Size
X 0.28799998760223
0.57599997520446
Y 0.06480000168085
0.12960000336170
Z 0.0 m
0.0 m
Rotation Axis
Rotation Angle
X 1.0 m
0.0 radian
Y 0.0 m
Z 0.0 m
Cell Size to Voxel Ratio 0.1
Minimum Voxel Size 0.001 m
Preview
Voxel Count:
Snap to part
Display
No Displayer
New Scalar Displayer
Existing Displayer
Scalar 1
Create
Close
Help

To select the input parts that you want to include in the resampled volumes, click the **Input Parts** button.

To specify the dimensions and position of the bounding box, set the following properties in the *Coordinates* group box:

- Origin
- Size
- Rotation Axis
- Rotation Angle

To control the number of voxels in the resampled volume, set the following properties:

- Cell Size to Voxel Ratio
- Minimum Voxel Size

Before you create the resampled volume, you can get a graphical preview of it and a preliminary estimate of the number of voxels. To do this, click **Preview**.

To position the bounding box along a part, activate **Snap to part**. This feature is available when the resampled volume tool is active.

The Display group box and the Create and Close buttons are common functions that are described elsewhere.

Working with Volume Rendering Settings

Additional controls let you choose a rendering mode, activate the lighting mode, and specify accuracy.

When you add a resampled volume part to a scalar displayer, that displayer acquires a new sub-node, **Volume Rendering Settings**. Use the properties of this sub-node to control how the data is displayed.



Volume Rendering Settings Properties

Rendering Mode	Allows you to select a mode of rendering.	
	Direct	Every scalar sample value is mapped to the displayer opacity plus color, and is then attenuated along the view direction.
	Maximum Opacity Projection	A transparent mode with high contrast to highlight volume areas that contain scalars with high opacity. An example is when velocity magnitude is mapped to a high-pass color map. The transparent mode reveals additional points of interest, regardless of the opacity attenuation which is depending on the viewpoint. However, the projective nature of

		this mode removes clues for depth perception.
	Level Surface	Allows you to choose between single or multiple surfaces. Activates a Level Surface Mode sub-node.
Lighting Mode	Allows you to apply displayer lighting to the volume rendering. The primary application for the Local Lighting mode is when the *Rendering Mode* property is set to `Level Surface`. In the Level Surface mode, standard lighting will be applied to the level surface exactly as if it were a proper surface. Additionally, local lighting is used in the Direct rendering mode for a more "surface-like" image without actually reverting to simple surfaces.	
	None	Does not use displayer lighting.
	Local Lighting	Applies the settings of the displayer surface lighting. The advantage of this technique is that details in the scalar field can be discerned more easily in many cases. The following Marschner-Lobb examples of volume rendering illustrate the benefit of local lighting. • Local lighting:

		 • No local lighting: 
Quality	Determines how accurately the resampled volume is processed during ray-casting.	

Working with the Level Surface Mode

Volume rendering in STAR-CCM+ is based on a ray-casting approach that is implemented in graphics hardware. The level-surface volume rendering mode tries to capture user-defined level values of interest during ray-casting through the resampled volume derived part.



The **Level Surface Mode** node allows you to choose between single and multiple surfaces for volume rendering in its properties.



Level Surfaces Compared with Isosurfaces

Level surfaces are in principle similar to isosurfaces. However, level surfaces provide a better visual accumulation of partly transparent

surfaces than isosurfaces with the same opacity settings. The data inputs, data extraction approaches, and storage locations of these types of objects are different. The effects of the inputs and extraction depend on the mesh, the native data location of the scalar field, and the accuracy settings for the resampled volume part.

	
Level Surfaces	Isosurfaces

Feature	Data Input	Extraction Approach	Hardware Employed	Storage	Location
Level Surfaces	Resampled Volume derived part	Ray-Casting	Graphics Hardware / GPU	Image buffer in graphics memory	Client
Isosurfaces	Discrete meshed region or part	Polygonalization of intersected mesh edges	CPU	Polygonal surface mesh in main memory	Client and Server

See also Examples of Effects of Volume Rendering Settings on Level Surfaces.

For an interactive exploration of simulation results, level surfaces have the following advantages over isosurfaces:

- Typically they update very fast, at the speed of volume rendering.
- They provide better results that are based on the opacity that you specify in the displayer and the colormaps.

In comparison to isosurfaces, the level surfaces have the following limitations:

- They cannot show a secondary scalar field on their surface.
- They cannot serve as an input to other derived parts or reports.
- They support a maximum of 32 level values.
- Their accuracy depends on the resolution of the underlying resampled volume. An example is bumpy disruptions of an otherwise flat surface. If such artifacts occur, use isosurfaces.
- They do not work with software-based OpenGL implementations such as Mesa OpenGL, or Microsoft OpenGL (remote desktop).

Specifying Level Surfaces

- To work with a single level surface, select the **Level Surface Mode > Value** node and enter a value for *Level Value*.
- To work with multiple level surfaces that are equally spaced in a range:
  - Select the **Level Surface Mode** node and set *Level Surface Mode* to **Multiple (range)**.
  - Select the **Level Surface Mode > Values (range)** node and set the following properties:
    - *Number of Levels*
    - *Minimum Level Value*
    - *Maximum Level Value*
- To work with multiple level surfaces that you arbitrarily set:
  - Select the **Level Surface Mode** node and set *Level Surface Mode* to **Multiple (arbitrary)**.
  - Select the **Level Surface Mode > Values (arbitrary)** node and set the properties using one of the following techniques:
    - Enter settings for *Level Values*.
    - Click ... (**Property Customizer**) of the *Level Values* property.



In the dialog that appears, enter values as follows:

- Set the number of level surfaces in the spin box.
- If you know the exact values to display, enter them manually in the text box and click **OK**.
- If you want STAR-CCM+ to generate values, click **Query**, then click **Auto Generate Values**. Click **OK**.

  The **Query** action retrieves the scalar range from the server. The **Auto Generate Values** action produces an even set of values in the text box according to the number of level surfaces that you specify.

Examples of Effects of Volume Rendering Settings on Level Surfaces

Lighting mode:



Local lighting on



Local lighting off

Opacity:



Displayer Opacity = 0.05



Displayer Opacity = 0.75

Colormaps:



Blue-red (default)



Blue-red-band

Level Surface Mode Properties

Level Surface Mode	The mode for level surface volume rendering, specifying whether a single surface or multiple surfaces are created.	
	Single	Uses one surface (default).
	Multiple (range)	Selects a range technique for creating multiple surfaces.

	Multiple (arbitrary)	Selects an arbitrary technique for creating multiple surfaces.

Working with Volume Rendering Colormaps

Some colormaps are designed for use with volume rendering. These colormaps provide a mapping of the colored scalar to opacity.



If existing color maps are chosen for volume rendering, then they apply full opacity to all scalar values, visually resulting in a compact volume block. The volume rendering colormaps should not be used for coloring in other scalar displayers that show surfaces or streamlines.

The colormap types are as follows:

- Blue-red-band -- rainbow color map with the opacity decreasing from center towards minimum and maximum (band pass filter)
- Blue-red-highpass -- rainbow color map with the opacity increasing from minimum to maximum (high pass filter)
- Blue-red-lowpass -- rainbow color map with the opacity increasing from maximum to minimum (low pass filter)
- Cool-warm-band -- cool/warm color map with the opacity increasing from center towards minimum and maximum (band elimination filter)

Viewing the Final Result

When you finish creating your resampled volume, you can view it in the assigned scene.

Once you have finished creating a resampled volume part, you can rotate and zoom in the display to get a clear view of the resampled volume, which is shown in the following screenshot. If you did not assign the part to the current scene, you can drag and drop its node into the visualization display. The resampled volume part can be rendered only in the scalar displayer.



You can take this part out of the display using one of the following techniques:

- Deleting the derived part
- Removing the displayer from the scene
- Removing the derived part from the scene
- Hiding the derived part

Saving and Printing Scene Data

In STAR-CCM+ you can produce output of the contents of a scene display in a variety of ways.

The options are:

Exporting an Interactive Three-Dimensional Scene

You can export the contents of a scene to a `.sce` or VRML file from STAR-CCM+.

You can open the `.sce` file in STAR-View+ and work with the same three-dimensional image that you had in STAR-CCM+. The VRML file can be viewed using a suitable 3rd party viewer such as Cortona 3D Viewer. In both cases you can interact with the image to rotate, pan, and zoom.

To save the scene contents to these formats:

1. Right-click the scene node or the scene display, and select **Export 3D Visualization File...** from the pop-up menu.



This will activate the *Save Scene File* dialog, in which the Files of type: filter is used to select the format of the saved file. This is set to the **STAR-View+ (*.sce)** format by default.

To export the scene as a VRML file:

2. Select the **VRML (*.wrl)** format from the drop-down list.



3. Navigate to the location where you want the file saved, enter the file name in the File Name: text box, and click **Save**.



This action is disabled if more than one scene node is selected.

Saving Scene Image Files

The **Hardcopy...** pop-up menu item lets you save one or more scenes to image files.

To save a scene to an image file:

1. Right-click the scene node and select **Hardcopy...**

   To save multiple scenes use `<Shift><Click>` or `<Ctrl><Click>` to select the scene nodes, then right-click one of the selected nodes and select **Hardcopy...**.



This pop-up menu also appears if you right-click the background of a scene in the *Graphics* window, from which a single image file of the current scene can be saved.

2. Use the *Save To File* dialog to save the selected scene(s) to file.

   The saved image(s) can then be incorporated into a printed report, a slide presentation, a Web page or another commonly used medium.



The *Save To File* dialog includes a **Magnification:** text box, which lets you enlarge the picture by an integer value. (STAR-CCM+ will ignore fractional and negative entries by reading them as 1.)

A magnification value of 1 means the size and resolution of the resulting file will be the same as that of the screen (typically 72 dpi), so this is essentially a screenshot of what you see in the visualization display. A larger magnification creates a larger file at the same resolution, giving the effect of a higher resolution image if the file is imported into a document and scaled down. For example, suppose you want an image approximately 300 dpi, suitable for printing in an office document. A magnification of 4 would produce an image four times as large as the original, still at 72 dpi. However, by specifying a 25% scale when importing the image into a presentation program or image editor, you can obtain an effective resolution of 288 dpi.

It is possible to set specific values for the image resolution, in pixels, by clearing the **Use Screen Resolution** checkbox and entering values for the **X Resolution** and **Y Resolution**.

The image file can be written in any of the following formats, available in the dialog's Files of type: drop-down list:

- PostScript (`.ps`)
- Bitmap (`.bmp`)

- Netpbm superformat (`.pnm`)
- Joint Photographic Experts Group (`.jpg`)
- Portable network graphics (`.png`)

When saving multiple scenes the settings specified in the *Save To File* dialog will be applied to all scenes. The image files are saved under the name provided in the **File Name** text box and numbered incrementally according to their position in the simulation tree. In the example shown in the above screenshots the image files would be named `intake-scene_0.png`, `intake-scene_1.png` and `intake-scene_2.png`, which correspond to **Geometry Scene 1**, **Mesh Scene 1** and **Scalar Scene 1** respectively.

Saving a hardcopy of a scene that is not open or active in the *Graphics* window will trigger the scene to be opened.

Saving Scene Images in Batch

You can generate scene images offscreen, that is, without having the STAR-CCM+ workspace active on your monitor.

This option may be necessary if you are running the client remotely, and/or if you need to generate a large number of these images. STAR-CCM+ lets you do this easily by running a macro in batch.

1. Start recording a macro. See Recording the Macro.
2. Create a new scene. For example, to create a new scalar scene, right-click on the **Scenes** node and select **New Scene** > **Scalar** from the pop-up menu.



3. Right-click on the scene node and select **Hardcopy...**as shown below.



4. Save (stop recording) the macro.
5. Run the macro in batch mode:
   - Batch mode is run from the command line using the `-batch` option:

     ```
     % starccm+ -batch macroname.java
     ```

   - If the simulation file is not open:

     ```
     % starccm+ -batch macroname.java
     simulationname.sim
     ```

   If all clients are disconnected from a running server, no graphics scenes will be saved to files.

Printing a Scene

You can print scene images on paper.

To send a scene image to a printer connected to your system, select the scene node, then use the **File** > **Print** menu item to invoke the print action.

Using Annotations

The **Annotations** node is the main manager object for all annotations that have been created in the simulation. By default it contains five annotations named **Iteration**, **Logo**, **Scene**, **Solution Time** and **Time Step** when the simulation is created. The annotations manager can hold any number of annotations up to the computer's memory limits.



This section defines what is meant by an annotation and describes how to use annotations in scenes. This section also explains how to work with the various types of annotations and describes the operations that are common to multiple annotations.

Annotations share a common pop-up menu. See Annotation Pop-Up Menu.

What Is an Annotation?

An annotation is an object that you can use to add information to the scene display in the *Graphics* window.

Five annotations named **Iteration**, **Logo**, **Scene**, **Solution Time**, and **Time Step** are created by default with the simulation, and you can create new annotations with the pop-up menu of the main annotations manager. The properties of an annotation become active when you select its node.

The **Annotations** manager node contains template annotations. When you apply one of these annotations to a scene, a copy of the original node is created, and stored within the equivalent manager node

belonging to the scene. Annotation objects that belong to a scene are termed 'applied annotations' to distinguish them from template annotations. Scenes can have their own annotation objects that you can customize separately from the original annotations. These types of annotations are called applied annotations, and they are stored in the **Attributes** > **Annotations** node within the scene node. For more information, see What Are Applied Annotations .



Note: The properties of the applied annotation (located in the **Scene** > **Attributes** > **Annotations** node) are different from the properties of the template annotation (in the **Tools** > **Annotations** node). The properties of the applied annotation override the properties of the original annotation.

You can add annotations to any number of scenes, where they may be placed anywhere in the *Graphics* window. You can change the properties of the text such as the size, font, and color.



You can use the mouse to relocate an annotation in a scene, provided that the annotation has the *Interactive* expert property and the checkbox of that property is ticked.



You may add user-defined annotations or default annotations where the text is pre-defined. Pre-defined annotations are provided to enable you to use common and useful annotations more easily.

The process by which a scene accesses an annotation is very similar to the way a scene's part displayer accesses parts. Annotations are stored in a manager separate from the scenes just as parts are stored separately as boundaries, regions or derived parts. Also, annotations can be added and removed through the use of an object selector.



All user-defined annotations may be renamed. Pre-defined annotations cannot be renamed.

See Also

Adding Annotations to Scenes
Deleting and Removing Annotations
Changing the Position of the Annotation

Annotation Pop-Up Menu

This pop-up menu becomes available if you right-click an annotation in the scene display or in the object tree.

Remove	Removes the annotation from the active scene.

See Also

Common Pop-Up Menu Items

Working with a Simple Text Annotation

A simple text annotation is used to place a customized text label in a scene display.

It is represented by a node with a default name of **User**. This object has its own properties and shares a pop-up menu with other types of annotations.



This section explains the following:

See Also

Annotation Pop-Up Menu

Simple Text Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

Text	The text of the annotation. The property customizer opens a dialog with a larger space for more text. If you enter the function name of any report prepended by '`$`', STAR-CCM+ will replace this by the numerical value provided by the report. For example, `$InletPressureReport` will return the numerical value given by the report **Inlet Pressure**. Further information is given in the next section.	
Background	Displays the text with or without a colored background box.	
	Ticked	Shows a box
	Cleared	No box
Font	The font and style of the text. These settings can be entered directly, or you can click the property customizer to activate a dialog. To	

	change the point size, use the *Default Height* property.	
Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.	
Background Color	The color of the background text box (visible when the **Background** checkbox is ticked), which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Shadow	Displays the text with or without a shadow effect.	
	Ticked	Shadow
	Cleared	No shadow

Simple Text Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the text in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Creating a Simple Text Annotation

To create a simple text annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New** > **Simple Text** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **User**.



If multiple annotations are created, the new annotation names will be concatenated with a number to ensure the name is unique.

Once the annotation has been created, it can be renamed at any time, modified to specify the required text, or added to any scene.

See Also

Modifying a Simple Text Annotation
Adding Annotations to Scenes

Displaying a Report Value in a Simple Text Annotation

You can include the value of any report within the text of a simple text annotation. To do this:

1. First find the field function that corresponds to the report within the **Tools > Field Functions** manager and record the name given in the *Function Name* property.
   The example below shows the field function corresponding to a report called **Inlet Pressure**. The *Function Name* is `InletPressureReport`.



2. Use this function name in the text of the simple text annotation by adding '`$`' to the front of the name.
   This would be `$InletPressureReport` for the example shown above. Note that this text can be surrounded by any other text if desired.



The function name will be replaced by the reported value, and will be updated whenever the report updates.

Modifying a Simple Text Annotation

To change the properties of a simple text annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Any changes to these properties will apply in all scenes where this annotation appears.

To change the font and font style of the text:

2. Click ... (**Custom Editor**) in the *Font* property to activate the dialog.



3. Use the *Height* property in the main *Properties* window to change the point size.
4. In the *Properties* window, you can change the color of the text one of two ways:
   - You can select the new color from the drop-down list in the right column of the *Color* property.



- To choose a color visually, you can use the *Color Selector* by clicking ... (**Custom Editor**) in that same property.

The text may also be given a shadow if the box under the *Shadow* property is ticked.

The content of the text can only be changed if the annotation is user-defined.

5. To change the text of the annotation, you can use one of two methods:

   - Directly in the *Text* property of the *Properties* window



- By clicking the ... (**Custom Editor**) of that property to open a separate text input dialog.



The results will appear in the scene display if you have it open.



Working with a 2D Image Annotation

A 2D image annotation is used to import a graphic into the background of the scene display.

It is represented by a node with a default name of **2D Image**. This object has its own properties and shares a pop-up menu with other types of annotations.



This section explains the following:

See Also

Annotation Pop-Up Menu

2D Image Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

File	The path to the imported graphic. The formats that can be imported are PNG, JPEG, TIFF, PNM and BMP.

2D Image Annotation Expert Properties

Default Height	The default height of the annotation as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the annotation in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.

Creating a 2D Image Annotation

To create a 2D image annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New > 2D Image** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **2D Image**. Each additional annotation of this type will be concatenated with a number to ensure the name is unique.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

Modifying a 2D Image Annotation

To change the properties of a 2D image annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Any changes to these properties will apply in all scenes where this annotation appears.

2. To specify the graphic file used by this annotation, you can type the path directly into the *File* property, or click ... (**Custom Editor**) to activate the dialog.



The formats that can be imported are PNG, JPEG, TIFF, PNM and BMP.

Once you have added this annotation to any scene, the picture will appear in the background. Its lower left corner will be set at the center of the display by default.



To shift its position:

3. You can enter coordinates in the *Default Position* property. Alternatively, you can click and drag the image provided the checkbox of the *Interactive* expert property is ticked.



To reduce the image:

4. Navigate to the scene annotation properties and enter a decimal fraction in the *Height* property to adjust the size as a percentage of the scene.



For example, setting the *Height* to `0.45` scales down the image to 45% of its original size.



While this property allows you to enlarge the image beyond its original size, this step is not recommended since it may give the image a grainy appearance.

5. To control whether the image appears behind or in front of the model, use the drop-down list in the *Location* property.
   In this illustration, `Background` is selected so that the picture serves as a backdrop for the geometry model.



The results will appear in the scene display if you have it open.



See Also

Adding Annotations to Scenes

Working with a 3D Image Annotation

The 3D image annotation lets you import an image and position it within three-dimensional space in the scene, rather than in the background of the scene display. This feature facilitates more precise placement of an image near a part of the model.

This object is represented by a node with a default name of **3D Image**. It has its own properties and shares a pop-up menu with other types of annotations.



This section explains the following:

See Also

Annotation Pop-Up Menu

3D Image Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

File	The path to the imported graphic. The formats that can be imported are PNG, JPEG, TIFF, PNM and BMP.

3D Image Annotation Expert Properties

Default Height	The default height of the annotation as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the annotation in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width	The default width of the annotation as a fraction of the visualization display width, where 1.0 equals the full width.
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.

Creating a 3D Image Annotation

To create a 3D image annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New > 3D Image** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **3D Image**. Each additional annotation of this type will be concatenated with a number to ensure the name is unique.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes
Modifying A 3D Image Annotation

Modifying A 3D Image Annotation

To change the properties of a 3D image annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Any changes to these properties will apply in all scenes where this annotation appears.

The procedure for importing a graphic into this annotation is very similar to that used when modifying a 2D image annotation. See Modifying a 2D Image Annotation.

2. Once the graphic is imported, you can adjust its appearance in a scene with the following properties:
   - The *Height* and *Width* properties set the scale.
   - The *Position* properties adjust the position relative to the model.
   - The *Opacity* property lets you add transparency to the image so you can use it as an overlay.

The results will appear in the scene display if you have it open.



Working with a Background Annotation

A background annotation is used to import a graphic into the scene display.

It is represented by a node with a default name of **Background**. This object has its own properties and shares a pop-up menu with other types of annotations. It functions in a manner similar to that of a 2D image annotation, except with fewer settings. This annotation is useful when you simply want to import a ready-made image into the background of a scene display.



This section explains the following:

See Also

Annotation Pop-Up Menu

Background Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

File	The path to the imported graphic. The formats that can be imported are PNG, JPEG, TIFF, PNM and BMP.

Background Annotation Expert Properties

Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.

Creating a Background Annotation

To create a background annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New** > **Background Image** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **Background**. Each additional annotation of this type will be concatenated with a number to ensure the name is unique.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes
Modifying a Background Annotation

Modifying a Background Annotation

This image annotation lets you select the file and change its opacity in its properties, just as you would for a 2D image annotation. However, the image is always shown behind other scene elements, and is resized vertically and/or horizontally to fill the scene display.

See Also

Background Annotation Properties
Modifying a 2D Image Annotation

Working with the Iteration Annotation

The iteration annotation identifies the number of completed iterations.

One of five pre-defined types created with each new simulation file, this annotation's read-only text is updated automatically as more iterations are completed. This object is represented by a node with a name of **Iteration**. It has its own properties and shares a pop-up menu with other types of annotations.



See Also

Annotation Pop-Up Menu

Iteration Annotation Properties

You can use the properties listed below to modify the specifications of the annotation.

Text	Displays the number of iterations (read-only).	
Time Step	Controls whether to report the time-step in the annotation. This is useful for unsteady simulations.	
	Ticked	Includes the time-step.
	Cleared	Omits the time-step.
Solution Time	Controls whether to report the solution time in the annotation. This is useful for unsteady simulations.	
	Ticked	Includes the solution time.
	Cleared	Omits the solution time.
Background	Displays the text with or without a colored background box.	
	Ticked	Shows a box
	Cleared	No box

Font	The font and style of the text. These settings can be entered directly, or you can click the property customizer to activate a dialog. To change the point size, use the *Default Height* property.	
Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.	
Background Color	The color of the background text box (visible when the **Background** checkbox is ticked), which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Shadow	Display the text with or without a shadow effect.	
	Ticked	Shadow
	Cleared	No shadow

Iteration Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the text in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Modifying the Iteration Annotation

To change the properties of the iteration annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Generally you can work with the iteration annotation as you would with a simple text annotation, since these types use virtually identical properties. The main exception is that the iteration annotation's text content cannot be changed, since it is designed to indicate the number of iterations in the solution.



See Also

Modifying a Simple Text Annotation

Working with the Logo Annotation

The logo annotation displays the CD-adapco logo in the scene.

It is one of five pre-defined types created with each new simulation file. This object has its own properties and shares a pop-up menu with other types of annotations.



See Also

Annotation Pop-Up Menu

Logo Annotation Properties

You can use the properties listed below to modify the specifications of the annotation.

Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.

Logo Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the logo in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Modifying the Logo Annotation

The CD-adapco logo appears in the background of every scene display by default.



To change the properties of the logo annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



2. To shift the annotation's position, you can enter coordinates in the *Position* property.

   Alternatively, you can click and drag the logo provided the checkbox of the *Interactive* expert property is ticked.
3. To adjust the transparency of the annotation, change the value of the *Opacity* property.

   The settings range from 0.0 (completely invisible) to 1.0 (completely opaque).
4. To control whether the logo appears behind (`Background`) or in front of (`Foreground`) the simulation object(s), use the drop-down list in the *Display Location* property.
5. To change the logo size, enter decimal fractions in the *Height* property to represent percentages of the visualization display height (1.0 equals the full height).
6. Clear the checkbox of the *Interactive* expert property if you want to ensure that mouse movements will not accidentally move the logo.

Any changes to these properties will apply in all scenes where this annotation appears.



Working with the Scene Annotation

The scene annotation, one of five pre-defined annotations created with each new simulation, displays information about the simulation in the scene display.

This object is represented by a node with a name of **Scene**. It has its own properties and shares a pop-up menu with other types of annotations.



See Also

Annotation Pop-Up Menu

Scene Annotation Properties

You can use the properties listed below to modify the specifications of the annotation.

Product Name	Controls whether to include the name of the program in the scene information.	
	Ticked	Shows the name.
	Cleared	Omits the name.
Product Version	Controls whether to include the version of the program in the scene information.	
	Ticked	Shows the version.
	Cleared	Omits the version.
Date	Controls whether to include the date and time (timestamp) of the last save in the scene information.	
	Ticked	Shows the timestamp.
	Cleared	Omits the timestamp.
User	Controls whether to include the STAR-CCM+ user name in the scene information.	
	Ticked	Shows the user name.
	Cleared	Omits the user name.
Simulation	Controls whether to include the simulation title in the scene information.	
	Ticked	Shows the title.
	Cleared	Omits the title.
Background	Displays the text with or without a colored background box.	
	Ticked	Shows a box
	Cleared	No box
Font	The font and style of the text. These settings can be entered directly, or you can click the property customizer to activate a dialog. To change the point size, use the *Default Height* expert property.	
Color	The color of the text, which can be selected from a drop-down list. The property	

	customizer executes the define color action. See Defining Colors.	
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.	
Background Color	The color of the background text box (visible when the **Background** checkbox is ticked), which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Shadow	Display the text with or without a shadow effect.	
	Ticked	Shadow
	Cleared	No shadow

Scene Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the text in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Modifying the Scene Annotation

To change the properties of the scene annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Any changes to these properties will apply in all scenes where this annotation appears. The *Font* property and the expert properties work the same as those you would use for modifying a simple text annotation.

However, the other properties are unique to this annotation, allowing you to display read-only information about the simulation file in a scene display.



By default, when you add this annotation to a scene, all but the product version information items are activated. You can include or omit any combination of them by using the checkbox in the corresponding property.

See Also

Modifying a Simple Text Annotation

Working with the Solution Time Annotation

The solution time annotation identifies the elapsed solution time for unsteady simulations.

One of five pre-defined types created with each new simulation file, this annotation's read-only text is updated automatically as the solution proceeds. This object is represented by a node with a name of **Solution Time**. It has its own properties and shares a pop-up menu with other types of annotations.



See Also

Annotation Pop-Up Menu

Solution Time Annotation Properties

You can use the properties listed below to modify the specifications of the annotation.

Text	Displays the elapsed physical time of the solution (read-only).	
Background	Displays the text with or without a colored background box.	
	Ticked	Shows a box
	Cleared	No box
Font	The font and style of the text. These settings can be entered directly, or you can click the property customizer to activate a dialog. To change the point size, use the *Default Height* expert property.	
Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.	

Background Color	The color of the background text box (visible when the **Background** checkbox is ticked), which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Shadow	Display the text with or without a shadow effect.	
	Ticked	Shadow
	Cleared	No shadow

Solution Time Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the text in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Modifying the Solution Time Annotation

To change the properties of the solution time annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Generally you can work with the solution time annotation as you would with a simple text annotation, since these types use virtually identical properties. The main exception is that the solution time annotation's text content cannot be changed, since it is designed to indicate the elapsed time in the solution.



Working with the Time-Step Annotation

The time-step annotation identifies the number of completed time-steps (for unsteady simulations).

One of five pre-defined types created with each new simulation file, this annotation's read-only text is updated automatically as the solution proceeds. This object is represented by a node with a name of **Time Step**. It has its own properties and shares a pop-up menu with other types of annotations.



See Also

Annotation Pop-Up Menu

Time-Step Annotation Properties

You can use the properties listed below to modify the specifications of the annotation.

Text	Displays the current completed time-step of the solution (read-only).	
Background	Displays the text with or without a colored background box.	
	Ticked	Shows a box
	Cleared	No box
Font	The font and style of the text. These settings can be entered directly, or you can click the	

	property customizer to activate a dialog. To change the point size, use the *Default Height* expert property.
Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.
Background Color	The color of the background text box (visible when the **Background** checkbox is ticked), which can be selected from a drop-down list. The property customizer executes the define color action.
Shadow	Display the text with or without a shadow effect.
	Ticked — Shadow
	Cleared — No shadow

Time-Step Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the text in the visualization display based on the coordinate system specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Modifying the Time-Step Annotation

To change the properties of the time-step annotation, bring its node into focus.

1. Open the **Tools** > **Annotations** node and select the annotation node.



This will make the annotation's properties accessible.



Generally you can work with the time-step annotation as you would with a simple text annotation, since these types use virtually identical properties. The main exception is that the time-step annotation's text content cannot be changed, since it is designed to indicate the number of time-steps in the solution.



Working with the Report Annotation

A report annotation is used to display the result of a report as formatted text in the scene display.

The text is refreshed whenever the report itself updates. This annotation has its own properties and shares a pop-up menu with other types of annotations. In the screenshot below, **Inlet Mass Flow** and **Outlet Mass Flow** are report annotations.



This section explains the following:

See Also

Annotation Pop-Up Menu

Report Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

Text	The text of the annotation (read-only). This is fixed as the title and value of the associated report.	
Number Format	A text field that defines the numeric formatting of the value associated with the report. Details on accepted entries are given in the section What Is the Number Format.	
Report	The name of the associated report (read-only).	
Font	The font and style of the text. These settings can be entered directly, or you can click the property customizer to activate a dialog. To change the point size, use the *Height* property.	
Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Position	The position of the text in the visualization display based on the coordinate system specified in the *Coordinate System* property.	
Coordinate System	Selects the type of coordinates to use for position:	
	Display	Shows x-y values in pixels in the display of the *Graphics* window. The origin is the lower left-hand corner of that display.
	Normalized Display	Shows x-y values as fractions of the length and height of the *Graphics* window display, respectively. The origin is the lower left-hand corner of that display.
	Viewport	Shows x-y values in pixels in the display of the *Graphics* window. The origin is the lower left-hand corner of that display.

	Normalized Viewport	Shows x-y values as fractions of the length and height of the *Graphics* window display, respectively. The origin is the lower left-hand corner of that display.
	View	Shows x-y values as view coordinates (-1,1).
	World	Shows x-y-z values as global coordinates. As you rotate the object, the annotation keeps its position relative to that object.
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.	
Shadow	Display the text with or without a shadow effect.	
	Ticked	Shadow
	Cleared	No shadow
Display Location	Controls the visibility of the annotation relative to the object(s) in the display.	
	Background	Shows the object(s) in front of the annotation.
	Foreground	Shows the annotation in front of the object(s).
Height	The character height as a fraction of the visualization display height, where 1.0 equals the full height.	

Report Annotation Expert Properties

Interactive	Controls whether the annotation can be moved in the scene display with the mouse.

	Ticked	Allows you to move the annotation in the display.
	Cleared	Locks the position of the annotation in the display.

What Is the Number Format?

The **Number Format** property of the report annotation is used to convert the value returned by the report to a suitable display format. The property is entered as text and should follow the pattern:

```
%[flags][width][.precision]specifier
```

where the elements in brackets are optional, and the specifier is a character that indicates the style of format required. The table below provides the accepted values for each element. Note that the measurement unit appearing on the display is defined by the report itself.

flags	An optional element that controls some overall format behavior.	
	-	An optional character that will cause the converted value to be left-justified within the width of the output.
	+	An optional character that will place + or – in front of the output value.
	0	An optional zero that will pad unused digits in front of the output with 0's.
	#	An optional character that will cause the decimal point to be printed even when no nonzero digits follow.
width	The minimum number of characters to be printed in the output value. If the output value	

	has fewer characters than the field width, it will be blank-padded on the left (or right, if the `-` flag has been given) to make up the width. If the `0` flag has been given, 0's will be used for padding instead of blanks.	
.precision	An optional period followed by a number that specifies the number of digits to appear after the decimal point. For example, `.3` would mean that three digits will appear after the decimal point in the output value.	
specifier	A character that indicates the style of output required.	
	f	Report value will be output in the form [-]m.ddd where the mantissa m will have as many digits as it needs, and the number of d's after the decimal point is given by the precision. Six d's will be used if the precision has not been set.
	e	Report value will be output in scientific notation, that is [-]m.ddde+xxx or [-]m.ddde-xxx. The mantissa, m, consists of one digit, and the number of d's is set by the precision (six if not given).
	E	Same as e, but with an uppercase E to identify the exponent.
	g	Use either f or e depending on which style will give the greatest precision in minimum space.

	G	Same as g, but with an uppercase E used to identify the exponent (if present).

Some examples of the effect of these formats on a report value are as follows:

Report value = `467699.65625 Pa`

Number Format	Output	Comment
`%f`	`467699.656250 (Pa)`	Default number of d's after decimal point (six).
`%6f`	`467699.656250 (Pa)`	Width of 6 has no effect as more than six digits are present
`%10.2f`	`467699.66 (Pa)`	One blank inserted at start to pad number to full width of 10 characters.
`%+012.2e`	`+004.68e+005 (Pa)`	+ symbol inserted for positive number; 0's used to pad output to full width of 12 characters.
`%8.3E`	`4.677E+005 (Pa)`	Three digits after decimal point; no padding required.

Creating a Report Annotation

To create a report annotation:

1. Right-click the relevant report node and choose **Create Annotation from Report**.

   This will add an annotation node with the same name to the **Tools** > **Annotations** node.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes

Modifying a Report Annotation

The report annotation can generally be treated as a simple text annotation, except that the text itself cannot be changed directly. Changing the name of the report node associated with the annotation will change the text.

The *Number Format* property is unique to the report annotation, and changes the formatting of the report value in the scene display. Full details are provided in the section What Is the Number Format? As an example, the scene below shows the value of a pressure report annotation having the *Number Format* set to `%10.2f`. This format means that the number will be presented in decimal notation, with a minimum of ten characters overall and two characters after the decimal point. As only nine characters are required to display the number in this format, a space is inserted before the leading digit.



If the *Number Format* is set to `%10.3e`, the report value will be presented in scientific notation where there is one digit before the decimal point and the number of digits following is given by the precision. In this case, three digits will be presented after the decimal point and a small `e` will be used to identify the exponent.



Working with the Scene Image Annotation

A scene image annotation is used to import a scene as an image in another scene display.

The annotation is updated in real-time as the solver is running or when changes are made to the scene that is used for the annotation. Furthermore, the annotation is refreshed whenever any of its properties are modified.

However, when the solver is running, the update frequency setting of the source scene (the scene shown as the annotation) is ignored. When the update frequency criterion is met for the destination scene (the scene containing the annotation) both the source scene and destination scene get updated.

This annotation shares a pop-up menu with other types of annotations. In the screenshot below, **Scalar Scene 1** is a scene image annotation.



This section explains the following:

See Also

Annotation Pop-Up Menu
Working with Scenes

Scene Image Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.
Scene	Selects a scene from a drop-down list of the existing scenes in the simulation.

Scene Image Annotation Expert Properties

Default Width	The default width of the annotation as a fraction of the visualization display width, where 1.0 equals the full width.
Default Height	The default height of the annotation as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the annotation in the visualization display based on the coordinate

	system specified in the *Coordinate System* property.

Creating a Scene Image Annotation

To create a scene image annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New > Scene Image** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **Scene Image 1**.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes

Modifying a Scene Image Annotation

This image annotation lets you select the scene image and change its opacity in its properties, just as you would for a 2D image annotation. The scene image is shown in the foreground by default, and can be scaled to appear alongside the object(s) displayed in the scene.

See Also

Modifying a 2D Image Annotation
Scene Image Annotation Properties

Working with the Plot Image Annotation

A plot image annotation is used to import a plot as an image in a scene display.

The annotation is refreshed whenever any of its properties are modified. This annotation shares a pop-up menu with other types of annotations. In the screenshot below, **XY Plot 1** is a plot image annotation.



This section explains the following:

See Also

Annotation Pop-Up Menu

Plot Image Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.
Plot	Selects a plot from a drop-down list of the existing plots in the simulation.

Plot Image Annotation Expert Properties

Default Width	The default width of the annotation as a fraction of the visualization display width, where 1.0 equals the full width.
Default Height	The default height of the annotation as a fraction of the visualization display height, where 1.0 equals the full height.
Default Position	The default position of the annotation in the visualization display based on the coordinate system specified in the *Coordinate System* property.

Creating a Plot Image Annotation

To create a plot image annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New** > **Plot Image** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **Plot Image**. Each additional annotation of this type will be concatenated with a number to ensure the name is unique.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes
Modifying a Plot Image Annotation

Modifying a Plot Image Annotation

This image annotation lets you select the plot and change its opacity in its properties, just as you would for a 2D image annotation. The plot

image is shown in the foreground by default, and can be scaled to appear alongside the object(s) displayed in the scene.

See Also

Modifying a 2D Image Annotation
Plot Image Annotation Properties

Working with the Scene Grid Annotation

A scene grid annotation lets you add a grid to a scene, allowing easier interpretation of visual data.

The annotation is represented by a node with a default name of **Scene Grid n**.



This object has properties for customizing this grid, including subnodes that let you set the grid spacing with display units or model units. It also shares a pop-up menu with other types of annotations.

The grid is perpendicular to your view of the scene display.



This section explains the following:

See Also

Annotation Pop-Up Menu

Scene Grid Annotation Properties

You can use the properties listed below to adjust the specifications of the annotation.

Opacity	Adjusts the transparency of the text. The settings range from 0.0 (completely invisible) to 1.0 (completely opaque). This allows you to use the annotation as an overlay.	
Grid Color	The color of the grid, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Text Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Spacing Mode	Selects the method for setting the grid spacing.	
	Display	Sets the grid spacing in display units. Activates the **Display**

		Spacing subnode. See Display Spacing.
	Model	Sets the grid spacing in model units. Activates the **Model Spacing** subnode. See Model Spacing.
Line Thickness	Controls the grid line width in pixels.	
Precision	Controls the grid coordinate precision. The default value is 3; this will display three digits after the decimal place.	

Display Spacing and Model Spacing

What Is the Display Spacing?

The **Display Spacing** subnode of the scene grid annotation represents the use of scene display units to specify the spacing of the grid, using entries in this object's properties.



Display Spacing Properties

Horizontal Spacing	Sets spacing in the horizontal direction.
Vertical Spacing	Sets spacing in the vertical direction.

What Is the Model Spacing?

The **Model Spacing** subnode of the scene grid annotation represents the use of model units to specify the spacing of the grid, using entries in this object's properties.



Model Spacing Properties

Horizontal Spacing	Sets spacing in the horizontal direction.
Vertical Spacing	Sets spacing in the vertical direction.

Creating a Scene Grid Annotation

To create a scene grid annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New** > **Scene Grid** from the pop-up menu.



A new annotation will appear under the **Annotations** node with a default name **Scene Grid n**, with **n** being the total number of annotations of this type in the simulation. This numbering ensures that the default name is unique.



Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes
Modifying a Scene Grid Annotation

Modifying a Scene Grid Annotation

You can modify a scene grid annotation in its properties. Additionally, its subnodes let you choose which units to use for spacing the grid: display or model.

See Also

Scene Grid Annotation Properties
Display Spacing and Model Spacing

Working with the Derived Data Annotation

A derived data annotation is used to display information about derived data objects that are associated with data set functions.

This annotation has a node with the default name **Derived Data Annotation n**, with **n** being the total number of annotations of this name in the simulation. This annotation has its own properties and shares a pop-up menu with other types of annotations.



See Also

Annotation Pop-Up Menu

Derived Data Annotation Properties

You can use the properties that are listed below to adjust the specifications of the annotation.

Background	Displays the text with or without a colored background box.

	Activated	Shows a box
	Deactivated	No box
Background Color	The color of the background text box (visible when the **Background** checkbox is ticked), which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Color	The color of the text, which can be selected from a drop-down list. The property customizer executes the define color action. See Defining Colors.	
Opacity	Adjusts the transparency of the annotation. The settings range from 0.0 (invisible) to 1.0 (opaque). This setting allows you to use the annotation as an overlay.	
Shadow	Displays the text with or without a shadow effect.	
	Activated	Shadow
	Deactivated	No shadow
Derived Data	The derived data object, which is selected from a drop-down list	
Text	The text of the annotation, read only. For time histories, the text displays the name of the derived data object, the display time, and the name of the field function that is recorded in the history. For spectra, the text displays the name of the derived data object and object properties and their values, indicating how the data was processed.	
Font	The font and style of the text. These settings can be entered directly, or you can click the property customizer to activate a dialog. To change the point size, use the *Height* property.	

Derived Data Annotation Expert Properties

Default Height	The default character height as a fraction of the visualization display height, where 1.0 equals the full height.

Default Position	The default position of the text in the visualization display, based on the coordinate system that is specified in the *Coordinate System* property.
Default Width Stretch Factor	The default amount by which the annotation is stretched in the horizontal direction.

Creating a Derived Data Annotation

To create a derived data annotation:

1. Expand the **Tools** node and right-click the **Annotations** node.
2. Select **New > Derived Data Annotation** from the pop-up menu.



A new annotation appears under the **Annotations** node with a default name **Derived Data Annotation n**, with **n** being the total number of annotations of this name in the simulation. This numbering ensures that the default name is unique.

Once the annotation has been created, it can be renamed at any time, modified as needed, or added to any scene.

See Also

Adding Annotations to Scenes
Modifying a Derived Data Annotation

Modifying a Derived Data Annotation

You can modify a derived data annotation in its properties. The derived data annotation can generally be treated as a simple text annotation.

See Also

Derived Data Annotation Properties
Modifying a Simple Text Annotation

Adding Annotations to Scenes

Annotations can be added to a scene by any of the following methods:

- Using drag-and-drop to drag the annotation onto the scene node.



- Using drag-and-drop to drag the annotation into the visualization display.



Wherever you drop the annotation in the scene, it will appear at the coordinates specified in the *Default Position* expert property.



- Using the *Select Objects* dialog.

To use the *Select Objects* dialog:

1. Select the **Attributes** > **Annotations** node within the scene node.



2. Click the right-hand column of the *Annotations* property.



This will activate the *Select Objects* dialog.



3. Use it to select annotations for the scene.

After you add the annotation in the scene, it will be represented in the object tree by an applied annotation. See Customizing an Annotation in a Scene.



After you add this object to the scene, editing the properties of the original annotation does not affect the applied annotation. You must remove the applied annotation from the scene(s) and then put a revised annotation there. Alternatively, you can edit that scene's applied annotation directly. (The exception is the text content of a simple text annotation, which can only be edited from the original.)

Annotating Scenes with Solution History Data

You can display information regarding recorded solution data in a scene using annotations. The following information can be displayed in a scene:

- Iteration
- Solution time
- Time-step

The *Representation* property of the applied annotations sets the source of the data for these annotations. In order to have the data come from a recorded solution view, you must update the *Representation* property, that is located in the applied annotation *Properties* window.



The *Representation* property is updated automatically for applied annotations that already exist in the scene.

See Also

What Are Applied Annotations?
Retrieving Results from the Simulation History Files

Changing the Position of the Annotation

To change the position of a particular annotation, you can click and drag it within the scene display, provided that the annotation has the *Interactive* expert property and the checkbox of that property is ticked.



Alternatively, you can enter coordinates in the properties of the applied annotation -- the object in the scene representing the annotation that you added. (Generally once an applied annotation is in a scene, it can only be changed by modifications to its own properties.)

1. Open the **Scenes** node, the node of the particular scene, the **Attributes** node, and the **Annotations** node, and then select the node of the applied annotation.
2. Enter the position vector directly under the *Position* property.



For a 3D image annotation, this type of entry is in two properties: *Position* and *Normal*.

3. Alternatively, click the ... (**Custom Editor**) of that property to open a separate text input dialog.



For most annotations, the position can be entered using a choice of visualization display-specific coordinate systems:

- Normalized Viewport: x,y (0, 1) normalized values in viewport; use this to attach an annotation to a fixed location in the visualization display.
- World: x,y,z global coordinate values; use this to attach an annotation to a point in the model space. This will result in text that moves with the attachment point when the view is changed.

These options can be selected from the drop-down list of the *Coordinate System* property of the appropriate applied annotation.

In the example below an annotation has been used with the **Normalized Viewport** coordinate system to add a title into the *Scalar Scene 1* display. A label has also been added using the **World** coordinate system.



As the scene is manipulated the title remains in the same location, but the label remains near the same part of the symmetry plane as it moves.



See Also

What Are Applied Annotations?
Adding Annotations to Scenes

Changing the Font of the Annotation

To change the font of an annotation:

1. Select an annotation from the feature tree.
2. Within the *Properties* window, click the ellipsis button to the right of the *Font* property.



This opens the *Font* dialog, where a **Font** and **Font Style** can be selected.



Clicking **OK** will apply the font to the currently selected annotation.

To change the font color:

3. Click the ellipsis button to the right of the *Font Color* property.



This opens the *Color* dialog, where colors can be selected.



Changing the Background Color of the Annotation

To change the background color of an annotation:

1. Select an annotation from the feature tree.
2. Within the *Properties* window, tick the *Background* property checkbox.



3. Then, click the ellipsis button to the right of the *Background Color* property.



This opens the *Background Color* dialog, where colors can be selected. See Defining Colors.



Copying Properties of Annotations

When properties are copied from one annotation to another, all of the properties are applied so that the destination annotation matches the source annotation. This applies only to user-defined annotations of the same type. For example, the properties of a 2D image annotation can be copied to another 2D image annotation, but not to a 3D image one.

The drag-and-drop technique for annotations is to drag the annotation node into the visualization display of the destination scene.



See Also

Copying Properties

Deleting and Removing Annotations

All user-defined annotations may be deleted. Pre-defined annotations cannot be deleted.

If you want to remove an annotation from a particular scene without eliminating it, you can use either of the following techniques:

- Access the *Annotations* property of the scene node.
- Bring into focus the scene from which you want to remove the annotation, then right-click the annotation and select **Remove**. The pop-up menu with this item is available when you right-click either the annotation in the scene display or its node in the object tree.

Pop-Up Menu of the Annotations Manager

These are the pop-up menu items available when you right-click the **Annotations** manager node.

New	Adds a new annotation of the selected type to the **Annotations** node.	
	Simple Text	Creates a new simple text annotation. See Working with a

		Simple Text Annotation.
	2D Image	Creates a new 2D image annotation. See Working with a 2D Image Annotation.
	3D Image	Creates a new 3D image annotation. See Working with a 3D Image Annotation.
	Background Image	Creates a new background image annotation. See Working with a Background Annotation.
	Plot Image	Creates a new plot image annotation. See Working with the Plot Image Annotation.
	Scene Grid	Creates a new scene grid annotation. See Working with the Scene Grid Annotation.

See Also

What Is an Annotation?

Using Colormaps

The **Colormaps** node is the main manager object for all colormaps that have been created in the simulation.

By default this node contains several pre-defined colormaps including **blue-red**, **cool-warm**, **grayscale**, **high-contrast** and **spectrum** when the simulation is created.



Many of the controls for colors in a scene are in the color bar of a displayer. However, these objects let you create custom color configurations.

What Is a Colormap?

A colormap is a table of color values that specifies certain colors and shades for data in the scene display.

Each colormap object in the simulation resides in the main colormaps manager, and is accessed by scenes in which you select that colormap.



User-defined colormaps can be edited in the colormap editor.

All user-defined colormaps may be renamed. Pre-defined colormaps cannot be renamed. When you rename a colormap that is being used in a scene, the new name appears automatically in the drop-down list of the displayer color bar's properties.

Creating Colormaps

This section explains how to create colormaps.

To create a colormap:

1. Expand the **Tools** node and right-click the **Colormaps** node. Select **New** from the pop-up menu.



   A new colormap appears under the **Colormaps** node with a default name. If you have a visualization scene open in the *Graphics* window, the *Create Color Map* dialog opens in the *Edit* panel.

2. Proceed to define the colormap.



Once the colormap has been created, it can be renamed at any time, modified to specify the required colors, or selected in any scene.

Activating the Colormap Editor

3. To activate the colormap editor, use one of the following techniques:
   - Select the colormap node and click the right half of the *Color Map* property.
   - Right-click the colormap node and select **Edit with Colormap Editor**.

   When you use the **Edit with Colormap Editor** menu item with a scene open, the dialog appears in the *Edit* tab of the simulation window.

Defining Colormaps

When you create a colormap, by default it has the properties of a blue-red colormap which is 100% opaque. The colormap editor lets you change those properties.

With the colormap editor, you can do the following:

- Modify the colors at the end of the color bar
- Add specific colors anywhere along the spectrum of the color bar
- Control the opacity of the colors along any part of the color bar



As you hover the mouse pointer over the top and bottom edges of the color bar, STAR-CCM+ displays prospective control point positions. Control points above the color bar define opacity; control points below the color bar define color. Click existing control points to edit their properties.

Adding a Color to the Spectrum

1. Click the position beneath the color bar where you want to place the color point.

   The *Properties* of the color point appear in the dialog.
2. Click the **Custom Editor** button to reveal the *Change Color* dialog.
3. In the *Change Color* dialog, select the color that you require and click **OK**.

Adding Variation in Opacity to the Spectrum

4. Click the position above the color bar where you want to place the control point for opacity.

The *Properties* of the opacity control point appear in the dialog.

5. Adjust the slider or enter a specific value. Its values range from 0.0 (100% transparent) to 1.0 (100% opaque).

Modifying Existing Control Points

6. To access the properties of a control point, do one of the following:
   - Click the control point.

     The control point takes the shape of a teardrop, indicating that it has the focus. The properties of that control point appear in the *Properties* group box.



   - Double-click the control point.

   A separate dialog appears for adjusting color or opacity.
7. To change the position of a control point, click and drag it.
   The *Position* property of the control point changes as you move it.

Colormap Dialog Controls

The *Edit Colormap* dialog exposes several properties for the color spectrum and control points along it.

Preview Mode	Controls how the dialog previews your changes.	
	On Change	Shows the effect as soon as you complete a change.
	Manual	Shows the effect after you complete a change and click **Preview.**

	Continuous	Shows the effect as you are making a change.
Color Space	Sets the color space that is used for interpolation.	
	RGB (Red, Green, and Blue)	The color model that is commonly used in video displays.
	HSV (Hue, Saturation, and Value)	Works like mixtures of paint.
	Lab	The CIE-L*a*b system.
	Diverging	Involves a special mode where colors pass through white when interpolating between two saturated colors.
	HSV-Wrapped	Uses HSV, but takes the shortest path in Hue (going back through 0 if that is the shortest way around the hue circle).

Deleting and Removing Colormaps

All user-defined colormaps may be deleted, provided they are not used by any scene displayer. Pre-defined colormaps cannot be deleted.

If you want to remove a colormap from a particular scene without eliminating it, access the *Colormaps* property of the scene displayer color bar node and select a different colormap.

Pop-Up Menu of the Colormaps Manager

Use this pop-up menu to create colormaps.

New	Adds a new colormap to the **Colormaps** manager node.

Using Transforms

The **Transforms** node, which has a pop-up menu, is a manager object for all transforms that have been created in the simulation.

By default it contains one transform named **Identity**, which cannot be deleted or renamed, when the simulation is created. You can create additional transforms directly through the pop-up menu. The transforms manager can hold any number of transforms up to the memory limits of the computer.



This section defines what is meant by a transform and describes how to create, use, and manage transforms in STAR-CCM+.

What Is a Transform?

Displayers use transforms to modify the default position, orientation, or size of the part in the displayer.

They are represented by nodes in the object tree which have properties and are subject to common menu items.



Transforms can repeat or mirror parts, or reposition, resize, or reorient them. The transform feature is for viewing only.

In STAR-CCM+, a transform is used with a displayer to allow for the repositioning, reflection, or repeat of parts in a scene. Transforms also work with coordinate systems (the Laboratory coordinate system by default).

Caution: When you switch the coordinate system of a transform, you change the spatial position of that transform. To compensate for this change, adjust spatial properties (such as *Rotation Axis* and *Rotation Origin*) of the transform. For details, see How Transforms Work with Coordinate Systems.

STAR-CCM+ provides the following types of transforms:

- Periodic transform—is automatically generated from a periodic interface once you initialize the solution, and is deleted if the interface is deleted. This transform gets most of its values automatically from the interface, such as the rotation angle, axis of rotation and translation vector, so the only property you can change is the number of repeats (copies). The graphics periodic transform lets you edit everything, but you must set all of its properties from scratch.

  Applying a periodic transform to a displayer shows the specified number of copies of the parts in the displayer, at periodic intervals based on information from the interface. Hence, the portion of the device that gets displayed depends on the following:

  - The parts included in the displayer
  - The number of copies
  - The intervals between the copies

- Symmetry transform—is an automatically generated transform from a symmetry boundary. This transform automatically extracts the plane of symmetry from the boundary. This transform is tied to the boundary and hence cannot be edited, and is deleted if the boundary type is changed, or the boundary is deleted. Applying a symmetry transform to a displayer shows the mirror image of the parts in the displayer about the symmetry plane.
- Graphics periodic transform—this transform behaves the same as a periodic transform, but it is your responsibility to define the rotation angle, rotation axis, or translation vector.
- Graphics symmetric transform—this transform behaves the same as a symmetric transform, but it is your responsibility to define the plane of symmetry. The plane is of the form ax + by + cz + d = 0, where a, b, and c are the x, y, and z components of the plane normal and d = -ax0 - by0 -cz0 where x0, y0, and z0 are the origin of the plane.

  For an axisymmetric problem, the axis is always the x-axis. Therefore you should enter `0,1,0,0` in the *Plane* property of this transform.
- Meridional projection transform—this transform is a simple projection of the points of the part into the (z,r) plane.
- Simple transform—this transform does not copy the original part, but translates, rotates, or scales the part accordingly. Applying a simple transform to a displayer transforms the parts in the displayer accordingly.
- Identity transform—this transform is automatically generated, and is the default transform for all displayers. There are no properties associated with this transform, and when applied to a displayer, shows all parts in the displayer at their natural position, size, and orientation.

A transform can be applied to a displayer to view the entire model for a solution in which only a symmetric or periodic repeating portion was solved. For example, when modeling a compressor stage, one often only solves for a single blade row. It becomes useful to view adjacent blade rows, or even the entire repeating solution. Using a periodic transform as the displayer's transform provides a way to view the complete geometry. Similar arguments apply for a problem with mirror symmetry.

A simple transform can be used in a displayer to view parts which would normally occlude each other by separating them.

All transforms can be renamed, except for the identity transform. All transforms can be copied provided they are of the same type.

Transform Properties

These properties reflect the types of boundaries you created.

The transform named **Identity** does not have any of these properties.

For the properties that include units in the values, you can change the units. For details on transforms and units, see How Transforms Work with Coordinate Systems.

Coordinate System	A list of coordinate systems in the simulation for use by the transform.
Boundary	The read-only name of the boundary whose type is reflected in the particular transform.
Plane	Indicates the coordinates of the plane along which the symmetry transform is created.
Number of Repeats	Controls the number of times the image is repeated in the display of the *Graphics* window. This is for a periodic repeat only.
Interface	Indicates which interface is used automatically for a periodic repeat (read-only).
Rotation Angle	Sets the rotation angle for a simple transform; is set automatically and is read-only for a periodic repeat.
Rotation Axis	Sets the rotation axis for a simple transform and meridional projection transform; is set automatically and is read-only for a periodic repeat.
Rotation Origin	Sets the rotation origin for a simple transform and meridional projection transform; is set automatically and is read-only for a periodic repeat.
Scale	Scales the part (simple transform only). By saving this setting in a transform, you can access it readily by switching to this transform in a displayer.
Translation	Shifts the Cartesian coordinates of the part in the display of the *Graphics* window; is set

	automatically and is read-only for a periodic repeat.	
Symmetry Mode	Selects the use of a symmetry transform.	
	Repeat	Displays mirror image of the parts in the displayer about the symmetry plane.
	Flip	Facilitates displaying different scalar quantities on either side of a symmetric transform.

How Transforms Work with Coordinate Systems

Transforms have units associated with each position property and axis property depending on the selected coordinate system.

The following example shows a simple transform with its properties.



This transform has three properties whose specification and units depend on the chosen coordinate system: *Rotation Origin*, *Rotation Axis*, and *Translation*. Each of these properties is represented as a coordinate

in the specified coordinate system, but the interpretation of the components depends on the type of property:

- The *Rotation Origin* property specifies a position in the chosen coordinate system which defines the rotation origin relative to the referencing coordinate system origin. In the above example, the position is `[0.0, 5.0, 0.0] m, m, m`, indicating the three components each have units of meters.
- The *Rotation Axis* property is a vector, but is also represented as a position in the chosen coordinate system. It is the vector formed from the origin of the chosen coordinate system to the specified position in that coordinate system.
- The *Translation* property is a vector, and shifts the Cartesian coordinates of the part in the display of the *Graphics* window relative to the origin of the chosen coordinate system.

When the coordinate system is changed, each property that depends on the coordinate system has its components recalculated with the units pertaining to that system. Consider the following example of a cylindrical coordinate system.



The Cylindrical 2 coordinate system has the Laboratory coordinate system as its reference system. The *Origin* property reads `1.0,1.0,1.0` because that is the position of the Cylindrical 2 coordinate system relative to the origin of the Laboratory coordinate system. Changing the

Coordinate System property of Simple Transform 1 to `Cylindrical 2` results in a change of its properties as shown in the following screenshot.



The components of coordinate system-dependent properties are recomputed to reflect the appropriate units for the cylindrical coordinate system.

In this example, the position of Simple Transform 1 has shifted relative to the Laboratory coordinate system. This shift is because the transform now uses the Cylindrical 2 coordinate system for its origin, and the Cylindrical 2 coordinate system is positioned at 1, 1, 1 relative to the Laboratory coordinate system.

Using Automatically Created Transforms

In STAR-CCM+, there are three types of automatically created transforms:

- A symmetric transform—if a boundary is of type symmetry, then a corresponding symmetry transform node is created and stored under the **Tools** > **Transforms** node, its name being that of the boundary.
- A periodic transform—if there is a periodic interface in the simulation, then a corresponding periodic transform node is created

and stored under the **Tools** > **Transforms** node, its name being that of the interface.

- The identity transform—this is the default transform, and lives under the **Tools** > **Transforms** node.

Creating a New Transform

To create a transform, simply right-click the **Tools** > **Transforms** node and create the desired type of transform.



A new transform node of the selected type appears in the tree under the **Tools** > **Transforms** node. The default name of the transform is based on its type. However, renaming the transform is possible.

You can edit the properties of the new transform to suit your needs. The new transform can now be used in any displayer, and appears in the list of available transforms in the *Transform* property of all displayers.

Transforms are shared, so changing the properties of a transform appropriately changes every displayer using that transform. For example, consider a scene with a part displayer and a scalar displayer

both using a graphics symmetry transform. If you change the symmetry plane definition in the graphics symmetry transform, then both the part displayer and scalar displayer update and change the scene accordingly.

Deleting Transforms

All user-defined transforms can be deleted by right-clicking on the transform node and selecting **Delete**.

Automatically created transforms are managed by the simulation, and are only deleted when deemed necessary by the simulation. For example, a periodic transform is tied to an interface and is deleted if the interface is deleted or the interface type is changed from periodic. Similarly, a symmetry transform is tied to a symmetry boundary and is deleted if the boundary is deleted or the boundary type changes. An identity transform cannot be deleted.

Using a Symmetry Transform

If you want to avoid changing the view that you might have set up in an existing geometry scene, right-click the **Scenes** node and select **New Scene** > **Geometry** from the pop-up menu.



If you do create a geometry scene, remove some default parts to isolate the ones in which you are interested. Expand the node of the new scene, the **Displayers** node, and the **Geometry 1** node. Select the **Parts** node.



Use the *Parts* property to remove all parts except for the ones you want to view.



To see the design without the outline, select the **Geometry 1** node within the **Displayers** node.



Deactivate the *Outline* property.



The transform lets you show the complete image in the *Graphics* window. With the **Geometry 1** node selected, set the **Transform** expert property to one of the `Symmetry` items in the drop-down list.



The model appears in full view in the *Geometry Scene* display. (The illustration below shows the model after rotating and zooming.)



Unwrapping, Unrolling, and Projecting with Transforms

Some visualization requires unwrapping results to a flat plane.

Consider, for example, a turbomachinery case that has a blade-to-blade view for a radial compressor that was created with the block-mapped coordinate system. The flow enters the domain in the axial direction and exits in the radial direction. As a result the blade-to-blade view sweeps around from the axial to the radial direction. Viewing results on such a curved surface would be difficult. Additionally for a meridional view, it would be useful to project the results on the radial-axial plane. For these reasons STAR-CCM+ provides the Meridional Projection transform.

In addition, the (m', theta) warp is available as a derived part.

Using a Meridional Projection Transform

This transform is a simple projection of the points of the part into the (z,r) plane. It requires only two vectors in its properties: *Rotation Origin* and *Rotation Axis*.

The following example uses the meridional projection transform:



- The rotation axis is [0, 0, 1] (aligned with the z-axis).
- The origin is upstream of the surface at [0, 0, 0].
- The following screenshot shows the view down the z-axis. The transform maps the part to the (x, y) axis.



Displaying Different Scalar Quantities on Either Side of a Symmetry Transform

The symmetry transform's **Repeat** option allows you to create a reflection of a scalar displayer's parts, but it shows the same parts and scalars in both their actual and transformed locations.

The following screenshot shows an example of a displayer using this option with temperature as the scalar.



By using this transform's **Flip** option, you can display a different scalar quantity on the mirrored set of parts. For example, to show temperature on the original and pressure on the mirror, do as follows:

- Create a graphics symmetry transform, make sure that the *Plane* property conforms to the symmetry plane that you need, and then set the *Symmetry Mode* property to `Flip`.
- Set up one scalar displayer showing temperature with no transform (that is, the default `Identity` selection).
- Set up a second scalar displayer as a copy of the first, except it shows pressure and uses the symmetry transform.

You have a mirrored displayer with different scalar quantities.



Pop-Up Menu of the Transforms Manager

These pop-up menu items are available when you right-click the **Transforms** manager node.

New Graphics Transform	Creates a transform of the type selected in the submenu:	
	Simple Transform	Executes the create simple transform action to add a simple transform to the **Transforms** node.
	Graphics Symmetry Transform	Executes the create graphics symmetry transform action to add a graphics symmetry transform (mirror image) to the **Transforms** node.

	Graphics Periodic Repeat	Executes the create graphics periodic repeat action to add a graphics periodic repeat transform (multiple repeats of the model that comprise a larger object) to the **Transforms** node.
	Meridional Projection Transform	Executes the create meridional projection transform action to add a meridional projection transform to the **Transforms** node.

Working with Views

The **Views** node, which has a pop-up menu, is the manager object for all views that have been saved in the simulation.



It exists even when empty to allow creation of the first view. The views manager can hold any number of views up to the computer's memory limits.

This section defines what a view is and describes how to use it.

What Is a View?

In STAR-CCM+, every scene has a camera to convert the three-dimensional objects in a scene into the two-dimensional image visible on the screen. The information defining the orientation, position and projection of the three-dimensional objects is stored in a **View** object in STAR-CCM+.

When a specific view is created or saved a **View** node is added to the simulation tree, which has its own properties and is subject to common menu items.



Information on the current view of a scene is displayed by selecting the **Attributes > View** node for the scene. The current view can be controlled interactively by the changing the property values of this node, by using the mouse, or by defining specific views using the methods described in this chapter.



There are no restrictions on renaming views. Saved views can be deleted.

View Properties

The properties of a view consist of the position, focal point (target), orientation (view-up) and projection mode of the image in a scene display. The focal point and position vector are displayed under their respective properties in the *Properties* window. Another vector is displayed under the *View Up* property which dictates the vertical vector in the scene window. The *Projection Mode* property controls how the camera maps world coordinates into view coordinates. The two options are:

- Perspective - implies a mapping of world coordinates into view coordinates that roughly approximates a camera lens. The center of projection is a finite distance from the view plane. As a result, the

scene will have a vanishing point, and closer objects appear larger than more distant objects.

- Parallel - implies a mapping of world coordinates into view coordinates that preserves all parallel lines. In this mode, an object appears the same size regardless of how far away it is from the viewer. This is equivalent to having a center of projection that is infinitely far away in perspective mode.

The view properties are summarized below.

Focal Point	Numerical coordinates of the focal point. A comma or space must separate the numeric values. To see the effect, select the **Restore View** submenu of the (**Save-Restore-Select Views**) menu in the toolbar. The `<H>` key sets the focal point of the camera to the position where you hover with the mouse.	
Position	Numerical coordinates of the position. A comma or space must separate the numeric values. To see the effect, select the **Restore View** submenu of the (**Save-Restore-Select Views**) menu in the toolbar.	
View Up	The coordinates for viewing. A comma or space must separate the numeric values.	
Coordinate System	A list of coordinate systems in the simulation for use by this view.	
Projection Mode	Provides options for projecting the object in the display:	
	Perspective	Displays the object with a vanishing point.
	Parallel	Displays the object with edges parallel to each other.

Setting a Predefined View

STAR-CCM+ provides a number of predefined views which allow scenes to be displayed relative to common orientations of the xyz axes.

These are useful when you wish to orientate the scene with an axis or when the orientation of the scene has become difficult to determine and you wish to reset the scene relative to an axis.

From the *Scene* toolbar, click on the camera icon and select **Views** from the drop down menu. This reveals another drop-down menu from which you must choose the axis to look down. For instance, if the +X axis is selected the resulting view will have you looking along the x-axis with positive x pointing out of the screen. If you choose the -X axis the resulting view will have you looking along the x-axis with negative x pointing out of the screen. Once an axis and direction have been chosen you must then select a second axis and direction from the drop-down menu which will define the view up.



Another set of predefined views consists of standard views, accessed in the submenu by that name. These options correspond to aspects of the object in the scene.



An additional option to view the scene along the vector `[1,1,1]` is also provided in the view's properties.

See Also

Using the Distance Measuring Tool

Setting the Coordinate System of the View

You can select one of the coordinate systems in the simulation to be used by a view. One benefit of doing this is to follow mesh motion.

Selecting a coordinate system can be done using the same techniques as other modifications:

- In the *Properties* window



- In the *Edit View* dialog



- In the (**Save-Restore-Select Views**) menu



This menu setting changes, and can be changed by, the *Coordinate System* property setting for the scene's current view node. It does not affect the properties of existing saved view objects in the **Tools > Views** node.

If you save a view in a scene, the new view object is saved with the coordinate system currently used in that scene. At the same time, if you restore a view (that is, access a saved view object), you will replace the scene's current coordinate system selection with that of the saved view object.

Copying Properties of Views

When properties are copied from one view to another, all of the properties are applied so that the destination view matches the source view.

This has the same effect as restoring the view.

The drag-and-drop technique in particular can be used to copy properties of views in two ways. The typical way is to drag the source node onto the destination node, which requires navigating the tree. But

an additional technique for views is to drag the view node into the visualization display of the destination scene.

Saving a View

One of the most useful features of views is the ability to save them. This allows you to manipulate a scene and save an interesting view and then continue manipulating the scene.

The view may then be recovered without additional manipulation of the scene. However, only the viewing position of the object(s) displayed in the scene is saved; positions of two-dimensional annotations, scalar bars or other display features are not saved.

Once you have manipulated the scene so that interesting details are displayed, click on the **View** toolbar button and select **Store Current View**.



Alternatively, right-click the **Attributes** > **View** node in the object tree and select **Store View**.



A new view node will appear in the tree under **Tools** > **Views**. Expand the **Tools** node and then expand the **Views** node. The default name of the new view node will be **View** along with a number equal to 1 plus the number of view nodes already existing. However, it is possible to rename the view node.

Creating a View

To create a new view object, right-click the **Views** manager node and select **New**.

This has the same effect as saving a view.

Restoring a View

Each of the saved views can be restored.

Click on the **View** icon in the Visualization toolbar and select the

Restore View submenu of the (**Save-Restore-Select Views**) menu. A drop-down menu with image previews of each of the views will be shown. Once you have chosen a view the scene will orientate itself to that view. The view is panned, rather than snapped, into place so that you keep your visual context relative to the view you had. However, only the viewing position of the object(s) displayed in the scene is restored; positions of two-dimensional annotations, scalar bars or other display features are not saved in views.



Views can also be restored by copying them into the current view.

A view can be restored into any scene even though it was not saved from that scene.

See Also

Using the Distance Measuring Tool

Modifying Views

To modify a view simply select its node, right-click and select **Edit...**

This action makes the *Edit View* dialog available.



The *Parallel Scale* property in the dialog is used to scale the image in parallel projection mode (when the **Parallel Projection** checkbox is activated). The parameter works as an inverse scale, in that larger numbers will produce smaller images. Any user-defined value should be greater than zero. When you activate parallel projection mode, a parallel scale subnode is added to the **View** node.

The diagram below shows how these settings relate to the scene. Change the settings in the fields as desired and click **Apply**, then click **Close**.



Note: STAR-CCM+ may change some of the settings entered here in order to resolve the view in the scene.

You can also modify a view by selecting the **View** node and editing its properties. This technique can only be used to edit individual properties, and is not recommended for changing the *Position*, *Focal Point* or *View Up* properties; in this case the method described above should be used instead.

Changing the Projection Mode

Each scene can be viewed in two projections modes; parallel projection, and perspective projection.

The perspective projection will provide a realistic view of the scene whereas the parallel mode will look less realistic and will be similar in style to technical diagrams.

	
Perspective projection mode	Parallel projection mode

By default the scene is viewed using the perspective projection mode. To change to parallel projection mode in the current view expand the **Scenes** node, expand the scene node and click on **View**. Under the *Projection Mode* in the *Properties* window choose **Parallel** under the pull-down menu. The view of the scene will automatically update.

Another approach involves the **Projection Mode** submenu of the (**Save-Restore-Select Views**) menu. This setting changes, and can be changed by, the *Projection Mode* property setting for the scene's current view node.



These actions also add a parallel scale subnode to the **View** node.

The projection mode of saved views may also be changed. Expand the **Tools** node, expand the **Views** node and click on the chosen view. You may then change the mode using the same pull-down menu. However, if that view is used by a scene and you want the scene to reflect the mode change, then you must refresh the scene's current view by restoring the view.

You can also change the default projection mode. Select **Tools** > **Options**. This activates the *Options* dialog.



Open the STAR-CCM+node in the **Options** node. Select the **Visualization** node and use the drop-down list of the *Default Camera Projection Mode* property to change the default projection of the object in the display.

Setting the Parallel Scale

The **Parallel Scale** subnode of the **View** node is used to scale the image in parallel projection mode.

It is made available when you set the *Projection mode* property of **View** node to `Parallel`. You can specify the scale in this object's properties.



Parallel Scale Properties

Parallel Scale	The user-defined value of the scale.

The parameter works as an inverse scale, in that larger numbers will produce smaller images. Any user-defined value should be greater than zero.

	
Setting of 0.2	Setting of 1.2

Following Mesh Motion in a View

If you are using DFBI motion, you can make a view follow the motion of the mesh.

This is done by linking a view to a 6-DOF body through its coordinate system. During iteration, the 6-DOF body's movements update that coordinate system, which in turn keeps the view centered on the 6-DOF body.

To make this work, first create a 6-DOF body. This object will have a body-fitted coordinate system created automatically along with it.



This type of coordinate system is updated by the movements of the 6-DOF body with each iteration.

In the *Coordinate System* property of the view object, select that coordinate system. This selection will associate the view with the 6-DOF body.

Finally, add the view to a scene by dragging the view node into the visualization display of the destination scene. That will cause the display to be updated with the 6-DOF motions.

Pop-Up Menu of the Views Manager

These are the pop-up menu items available when you right-click the **Views** manager node.

New View	Executes the create view action to add a view to the **Views** node.

See also: common menu items.

Recording and Viewing Solutions

STAR-CCM+ can save selected scalar and vector functions to a simulation history file while a simulation is running.

You can control the choice of functions, file location, and frequency of writing, in the simulation history nodes that are located within the Solution Histories manager node.



To access the data that is stored within a simulation history file, create and use a solution view. Solution views are located within the Solution Views manager node. You can create and view snapshots of the current solution in steady state simulations, which is useful for comparing the results between different design iterations. You can animate the recorded solution data and record this animation to file. You can also use recorded solution data in reports and plots. Information from the simulation history file can be used to annotate scenes to display the

iteration number, the solution time, and the time step of the displayed state.

Understanding the Simulation History File

The solution data that you select is saved to an external file, called the simulation history (`.simh`) file, at the specified time intervals.

The simulation history file grows in size as more solution data is saved to it. Simulation history files are designed to work with the simulation (`.sim`) file that created them; simulation history files do not work with different or empty simulation files.

Each simulation history file contains a copy of the simulation's volume mesh in addition to the solution data. When recording solution history, the mesh is only written when it changes. If only the vertex coordinates change for the same mesh connectivity, such as during morphing or rigid body motion, only the modified coordinate array is written in the subsequent states. This behavior is true for recorded snapshots and automatic recording to a simulation history file.

In its current implementation, the simulation history file does not contain a copy of the surface data. If you want to perform surface post-processing, use the volume data that is projected onto the surfaces. A diagram demonstrating how the data is handled in STAR-CCM+ is shown below.



Additionally, no solution settings, such as solver parameters, are contained in this file.

Moving or deleting the simulation history file breaks the link between it and the simulation file. Any scenes that display data from the simulation history file are not rendered. If the link between the simulation history file and the simulation file is broken, you can re-establish it right-clicking on the **Solution Histories** node and selecting **Load**. You can then navigate to the simulation history file and load it into the simulation.

Recording Solution Data

Use the **Solution Histories** node to create and configure simulation history files. Solution data that you select is stored in these simulation history files.

The simulation history files that you create are located under the **Solution Histories** manager node. You can select what data is stored in these files and how often the data is saved to file.



Each simulation history file node contains two sub-nodes: the **Update** node, that controls how the update frequency is defined, and the **Update Frequency** node, where you can specify how often the updates take place.

See Also

Understanding the Simulation History File

Creating a Simulation History File

You can store specified simulation data in simulation history files.

To create a simulation history file:

1. Right-click on the **Solution Histories** node and select **New...**



   The *Save* dialog appears and asks you where you would like to save the simulation history file.

2. Choose a location that has adequate space and give the file an appropriate name.

   Note that you must have write access to the location you choose, otherwise the solution data cannot be saved to the file.

3. Click **Save** to create the simulation history file and close the dialog.

A new sub-node with the name of the simulation history file appears in the **Solution Histories** node. In the following example, the simulation history file was given the name `pressureAndVelocityData`.



The red asterisk next to this node indicates that STAR-CCM+ is set to write data to the simulation history file when you run the simulation.

You can interact with the simulation history file by right-clicking on its node.



The following actions are available:

Create Snapshot	Creates a snapshot of the current solution and adds it to the simulation history file. This option is useful for comparing different design iterations.
Create Recorded Solution View	Creates a recorded solution view sub-node in the **Solution Views** node. From here, you are able to control the time of the solution that is stored in the simulation history file.
Rescan	Rescans the simulation history file to determine whether there were any states added to it since you last used the file.
Clear	Clears the solution history data from the simulation history file. The simulation history file is deleted and recreated the next time data is written to it, either through a manual capture (snapshot) or automatically (auto-record).
Edit...	Brings up an Edit dialog from where you can modify all the parameters of the simulation history file in one operation.

Delete	Removes the simulation history file from the STAR-CCM+ object tree; this action does not delete the simulation history file from the disk.
Rename	Brings up a window from where you can rename the node in the STAR-CCM+ object tree. This action does not affect the simulation history file on disk; the original name of the file is retained.

Configuring the Simulation History File

You can configure the simulation history file by adjusting its properties through the *Properties* window.



Here, you can select what scalars and vectors you want to save, along with what regions in the geometry to save them from.

To record the desired scalar or vector data:

1. Click the ... (**property customizer**) button located next to the *Scalar Functions* and *Vector Functions* properties respectively.

   A selection manager dialog appears.



2. Use the **>** (**Add selected**) button to select the desired scalars or vectors-the selections are added to the right-hand side column.



3. Click **OK** to close the dialog.

You can select the vector functions, and the regions this data comes from in the same manner.

Further information on simulation history file properties is provided in the following table:

Auto-record	Controls whether STAR-CCM+ saves the solution data to the simulation history file or not.	
	Activated	The selected solution data is written to the simulation history file at the specified time intervals. When this option is enabled, an asterisk appears besides the simulation history file node in the object tree.
	Deactivated	No solution data is written to the simulation history file.
Scalar Functions	Allows you to select the desired scalar field functions you wish to save to the simulation history file.	
Vector Functions	Allows you to select the desired vector field functions you wish to save to the simulation history file.	
Regions	Allows you to select the regions that you wish to save to the simulation history file.	
Path	Displays the relative path to the simulation history file where the solution data is written. Clicking ... (**Property Customizer**) displays the absolute path to the simulation history file. Note that you cannot edit this path.	
State	Displays the number of saved states that are present in the simulation history file.	

Selecting the Update Method and Setting the Update Frequency

Select the trigger (the method that defines how often the solution data is saved to the simulation history file) in the **Update** node. See Trigger Types.



To control how often the solution data is saved to the simulation history file, use the **<Trigger> Frequency** node. See Trigger Properties.



For more information on selecting the trigger type and defining the update frequency, refer to Setting the Trigger for an Update.

Clearing the Solution

Clearing the solution from a simulation file also clears the data that is stored in the simulation history file. Only data that was stored by the *auto-record* property is cleared. If a simulation history file contains saved snapshots, no solution data is cleared.

Caution: If you do not want to lose the solution data that is stored in the simulation history file, deactivate the *Recorded Solution History* checkbox in the *Clear Solution* dialog window. If this option is activated, the simulation history file will be removed from your disk—this action is irreversible.

The solution history settings are retained in the simulation file. Next time you run the solver, the simulation history file will be created with the same settings as before it was deleted.

Retrieving Results from the Simulation History Files

Use the **Solution Views** node to create solution views. Solution views allow you to retrieve and view recorded solution data that is stored on a simulation history file.

You can control the display of recorded solution data using these solution views.



You can create multiple solution views that reference the same simulation history file. This capability is useful for performing side-by-side comparisons of different case set-ups or design iterations. In each solution view, you can choose to display the transient data that is stored in each state, or the snapshots that you created.

Caution: Simulation history files created in STAR-CCM+ v7.04 or earlier are missing data at interfaces for cases run in parallel. In order to restore the missing data, clear the simulation history files and re-run the simulation in a later version of STAR-CCM+. No action is required for cases run in serial.

See Also

Using Solution Snapshots

Creating a Solution View

Solution views allow you to retrieve and view recorded solution data that is stored on a simulation history file.

To create a solution view:

1. Right-click on the simulation history file sub-node (in the **Solution Histories** node), and select **Create Recorded Solution View**



A new node that is named after the simulation history file is added to the **Solution Views** manager node.



You can apply the recorded solution to scenes, reports, and plots. You can control the solution that is displayed through its Properties window. The **Current Solution** node is automatically generated and contains the data for the current solution (the last iteration) in the simulation file.

When you create solution view, a new representation is also added to the **Representations** node.



This representation contains a copy of the volume mesh which is stored in the simulation history file. A solution view is associated with its representation. The properties of the solution view determine which data is taken from the simulation history file and stored in the representation. Changing the properties in the solution view node updates the representation with the new data from the simulation history file.

See Also

Applying Recorded Solution Views to Scenes
Using Recorded Solution Data with Plots
Using Recorded Solution Data with Reports
What Are Representations?

Applying Recorded Solution Views to Scenes

Recorded solution views contain stored data from your simulation. You can view this stored data by applying it to scenes.

To apply a recorded solution view to a scene:

1. Do one of the following:

   - Drag the recorded solution view from the object tree into the scene.



   - Drag the solution view representation from the **Representations** node into the scene.



   - Select the **Displayers > Scalar 1** node in the scene and set the *Representation* property in the *Properties* window to the representation of the recorded solution view.



Note: Selecting the representation in the *Properties* window only applies the recorded solution to the selected displayer. Using one of the drag and drop methods applies the recorded solution to all displayers and annotations in the scene.

Visualizing Recorded Solution Data

You can change the time of the displayed solution in the solution view *Properties* window.



When you first apply a recorded solution view to a scene, the scene displays the data for the first time-step that is stored in the simulation history file. In this example, the time-step is 0.02 s.



To change the solution time, adjust one of the time properties, such as *Solution Time*. A slider bar appears for this property--drag this slider to the desired time.



Alternatively you can enter the value that you want. The scene updates to display the relevant solution data.



To display the data that is contained in a snapshot, select the desired snapshot from drop-down list next to the *State Name* property. You can also select specific states from this property.

More information about each property is provided below:

Solution History	Selects which simulation history file to read the saved data from.
State Index	Displays the index of the currently loaded state. Click to the right of this property to reveal a slider where the state can be changed by moving the slider.
State Name	Displays the name of the currently loaded state. Click to the right of this property to bring up a drop-down menu from which to select states.
Iteration	Displays the iteration number of the currently loaded state. Click to the right of this property to reveal a slider for changing the iteration number.

Time Step	Displays the currently loaded time step. Click to the right of this property to reveal a slider for changing the time step.
Solution Time	Displays the currently loaded time. Click to the right of this property to reveal a slider for changing the solution time.

See Also

Applying Recorded Solution Views to Scenes

Using Recorded Solution Data with Reports

You can use recorded solution data with reports. This ability is useful for looking at engineering quantities at a particular instance in time.

To use the recorded solution data in a report:

1. Click to the right of the *Representation* property in the *Properties* window of the report.
2. Select the appropriate recorded solution view representation from the drop-down list.



Note: You can only use data that is stored in the simulation history file with reports. A list of common reports and their required field functions is provided at the end of this section.

3. Run the report to display the data.

The recorded solution data is controlled using the solution view Properties window—run the report after changing the solution time to display the corresponding data in the *Output* window.

Required Field Functions for Common Report Types

A list of common reports and required field functions is provided below.

Report Type	Required Field Functions
Single Stream Heat Exchanger/ Heat Exchanger	Face Flux Pressure

Report Type	Required Field Functions
	Temperature
Thrust	Pressure Velocity Face Flux Density
Mass Average	Density Cell Relative Velocity
Mass Flow Average	Density Cell Relative Velocity
Mass Flow	Face Flux Density Cell Relative Velocity
Momentum/ Momentum Coefficient	Pressure Wall Shear Stress
Force	Pressure Wall Shear Stress

Note: For any other type of report, make sure that you record the appropriate field functions to the simulation history file.

See Also

What Is a Report?
Visualizing Recorded Solution Data

Using Recorded Solution Data with Plots

You can use recorded solution data with XY Plots. This feature allows you to plot two-dimensional profiles of solution data.

To use recorded solution data in a plot:

1. Click to the right of the *Representation* property in the *Properties* window of the plot.
2. Select the appropriate recorded solution view representation from the drop-down list.



Note that only data that is stored in the simulation history file can be used with plots. The recorded solution data is controlled using the solution view Properties window. Changing the solution time, for example, causes the plot to update automatically and display the data for that moment in time.

See Also

What Is an XY Plot?

Using Solution Snapshots

A snapshot is essentially a copy of the selected solution data (scalars/vectors for that moment in time) in a simulation file. Snapshots can be used to compare multiple steady state solutions in a single simulation file, such as different design iterations.

To create a snapshot:

- Rright-click on the simulation history file in the object tree, and select **Create Snapshot**.



The *Create Snapshot* dialog appears. Here you can give the snapshot a suitable name. Click **OK** to close the dialog and add the solution data to the simulation history file.

A suggested procedure to follow for comparing steady state simulations is described below:

1. Create a simulation history file.
2. In the *Properties* window, disable the Auto-Record property so that no data is saved at each iteration.
3. Continue by selecting the scalars/vectors you want to save.
4. Set up a case and run it to convergence.
5. Create a snapshot of the solution, as described above, giving it an appropriate name.

6. Clear the solution.
7. Set up the case with the new settings you want to test.
8. Run the solver to convergence.
9. Create a snapshot of the new results.
10. Repeat steps 6 - 9 as many times as you want.
11. Create a new Recorded Solution View. This action provides access to all the saved snapshots.
12. Apply the recorded solution view to a scene and switch between the snapshots.

You can perform a side-by-side comparison by using multiple recorded solution views, where each solution view displays the data from different snapshots.



You can then create multiple scenes to compare the results, where each scene is assigned a different recorded solution view.



See Also

Recording Solution Data
Configuring the Simulation History File
Applying Recorded Solution Views to Scenes

Animating Solution Views

You can animate the transient and steady state solution data that is stored in a simulation history file.

For transient solutions, you can view the solution data against time. For steady state solutions, you can view a slideshow of the different snapshots.

Each recorded solution view node contains an Animation sub-node.



The *Animation Mode* can be selected from the *Properties* window.



Select one of the options that are shown above. This action adds a sub-node with the name of the selected *Animation Mode* to the solution view. For example if you select **State Index** as the *Animation Mode*, the new sub-node is named **State Index Animation**.



You can define the recorded solution view animation parameters in the Properties window.

Recorded solution view animations are controlled in the same way as standard scene animations. Similarly, you can record a copy of the animation to disk as you would with a normal scene.

For more details on how the parameters in the *Write animation* dialog affect the playback of the recorded movie, see Defining Playback Settings for Recorded Movies.

More details on what each property controls is provided below. Note that the properties that are shown below are for the **State Index Animation** node. The name of the *Start* and *End* properties vary depending on what *Animation Mode* you selected.

Cycle Time	Specifies the time that it takes to loop through all the state indexes.	
Auto-Range	Controls whether the values of the *Start Index* and *End Index* properties are automatically populated when playing the animation.	
	Activated	The values of *Start Index* and *End Index* are automatically filled in when playing the animation.
	Deactivated	The values of *Start Index* and *End Index* are not automatically filled in. These values must be entered manually.

Start Index	The state index for which the animation starts.
End Index	The state index for which the animation ends.

See Also

Animating Scenes
Controlling the Animation
Recording the Animation onto Separate Files

Animating Vectors and Sweeping Section Planes in Recorded Solution View Scenes

You can view animated vectors or sweeping section planes in the scene at any of the saved states.

The sweeping rate for plane sections is specified with a *Cycle Time*, like in the solution view animation. For both types of animation, the animation attribute of the scene specifies a target frame rate in frames per second (fps). It is called the target frame rate because a scene can be too complicated to draw at the rate you specified.

Together, the *Cycle Time* and the *Target frame rate (fps)* determine how many frames of the animation are used to animate the effect and the total number of frames in the animation. For example, if the *Cycle Time* is `10` seconds and the *Target frame rate (fps)* is `25` fps, the total number of frames in one cycle is 250 frames (10 * 25).

Plane sweeps and vector animations interpolate the positions of their objects at each frame to build a smooth animation. Solution state changes do not interpolate between states that are stored in the simulation history file. Scenes can therefore change abruptly when switching from one state to the next.

Controlling the Animation Playback

If the animation settings are such that the number of frames in the animation is less than the number of states between the *Start Index* and *End Index*, some states are skipped. On the other hand, if the animation settings are such that the number of frames in the animation is greater

than the number of states between the *Start Index* and *End Index*, states are repeated on multiple frames of the animation.

When an applied solution view has the animation enabled, you can also use the **Animation** toolbar to animate through the solution states.

As with vector and plane sweep animations, the **Animation** toolbar can be used to pause an animation that is playing. When paused, the frame stepping buttons are enabled so that you can step through the animation forward or backward one frame at a time. The following behavior can be observed:

- In the case where the number of frames in the animation is less than the number of states between the *Start Index* and **End Index**, pressing the forward button skips over one or more states.
- In the case where the number of frames in the animation is greater than the number of states between the *Start Index* and *End Index*, you need to press the forward button several times to advance to the frame so that the next state appears. If there are vector animations or plane section sweep animations that are enabled in the scene, their animations change from one frame to the next. If there are no vector or plane section sweep animations, the frame stepping buttons appears to be doing nothing, especially when the same state is repeated over several frames.

Defining Playback Settings For Recorded Movies

The playback settings for recorded movies are controlled in the *Write animation* dialog.



The properties that affect the recorded movie are:

- **Frame Rate** - this property controls the playback of the recorded movie. Higher frame rates take longer to write, and result in a larger file size for the movie.
- **Start Time** - this property sets the time at which the recorded solution data starts writing to the simulation history file. For example if the fluid interaction you are interested in starts after 5 s of recorded solution data, then it is recommended that you set the movie to start at 5 s.
- **Animation Length** - this property defines the length of the recorded movie. This parameter must match the recorded *Cycle Time* if you want the movie to play once, otherwise the animation loops, and plays the same content multiple times. For example if the *Cycle Time* is 20 s, and the *Animation Length* is set to 60 seconds, the animation plays the same content 3 times (60/20).

If you want the recorded movie to play back at the normal speed (not slowed down/speed up), the *Animation Length* must match the *Cycle Time*. Some examples are provided below:

Cycle Time (s)	Number of States	Animation Length (s)	Frame Rate (fps)	Playback Speed
10	250	10	25	x1

Cycle Time (s)	Number of States	Animation Length (s)	Frame Rate (fps)	Playback Speed
10	250	20	25	x1 (looped twice)
10	500	20	25	x0.5
10	500	10	25	x1
10	500	5	50	x2

Reporting Results

Reports provide summaries of solution data.

The **Reports** manager node, which has its own pop-up menu, contains all of the report nodes in the simulation tree.



A reference section on report types is also provided.

What Is a Report?

A report presents a computed summary of the current simulation or CPU data. A node represents each report in the simulation tree, each with its own pop-up menu. Report properties, however, vary by report type.

The report summaries are useful for post-processing, and enable engineering quantities such as drag, lift, torque, or mass flow to be computed. They are also useful for computing diagnostic quantities such as minimum cell volume or average wall y^+.

There are three basic categories of report types:

- Types that require definition of the scalar on which you want to report, such as minimum and maximum
- Types where the scalar is inherently defined, such as force
- System-type reports such as solver CPU time

Additionally, an expression report is available that depends on a combination of reports other than system reports.

The first two categories of report require a selection of parts over which the report is calculated, and they are available on:

- Regions only
- Surfaces only
- Boundary surfaces only
- Derived parts only
- Lagrangian parcels only
- A combination of the above

The description of the pop-up menu of the **Reports** manager node, which lists choices for report creation, includes information on which of these part types a particular report accesses.

The reports that are made available to a simulation depend on the models that are activated. However, five reports (solver CPU time, solver elapsed time, maximum, minimum and sum) are always available.

Report Pop-Up Menu

These pop-up menu items are available when you right-click a report node.

Run Report	Runs the report once, generating the results in the *Output* window.
Create Monitor and Plot from Report	Creates a monitor for this report and then also creates a plot from that monitor. The new plot takes the name of this report suffixed with " Plot". The plot is not automatically displayed on creation, but the new report node is selected in the tree.
Create Monitor from Report	Creates a monitor for this report. The monitor takes the name of the report suffixed with " Monitor" and its node is automatically selected in the tree.
Create Annotation from Report	Creates an annotation for this report. The annotation takes the name of the report and its node is automatically selected in the tree.

Adding Reports

To create any report, right-click the **Reports** node and select the report type that you wish to create from the **New Report** pop-up menu.



What is added to the **Reports** node, the default naming of the new objects and any further actions that are required depend on whether the report is a system report, a scalar-based report, a specific report, or an expression report. However, reports can also be renamed.

Once a report is created, it is saved with the simulation.

Adding a System Report

System reports (labeled **... Time** in the pop-up menu) are added in the same way as any other report.

This immediately adds a single child object (**Solver Elapsed Time per Iteration** in the example below) to the **Reports** node without any properties.



The new report object is named the same as the submenu item that you selected, and it is also removed from the pop-up menu once it has been added to the **Reports** node.



Adding a Scalar-Based Report

Scalar-based reports (those labeled **Maximum**, **Minimum**, **Sum**, **...Averaged** and **...Integral** in the pop-up menu) are added in the same way as any other report.

Once the selection has been made from the pop-up menu, the new report is added in the **Reports** node (**Surface Average 1** in the example below).



The default name of the new report is the name that appeared on the pop-up menu, suffixed with a number to make this report unique.

The scalar-based reports all have two properties to allow for the following:

- Definition of the scalar that is operated on
- Selection of the parts that act as the source for the report

The *Smooth Values* expert option, when selected, uses interpolated values at the vertices instead of the default cell center values.



The *Scalar* property also allows you to select the individual components of certain functions, such as positions and centroids, in local coordinate systems.



Only parts relevant to the report are displayed in the part selector. For example, the surface average report that is used in the previous example only allows surfaces or lines to be selected, whereas the volume averaged report only allows regions to be selected.

Note: When the chosen scalar quantity is not defined on a face, and yet the parts for the report include face parts, STAR-CCM+ provides the report with scalar values on the cells adjacent to the face parts. For example, you could have a Sum report that requires Volume for the Scalar Field Function, and the outer boundaries only of the domain for the input Parts. Instead of returning a zero value for the volume of the boundary faces, STAR-CCM+ returns the sum of the volumes of the cells adjacent to the boundary faces.

Adding Specific Reports

Specific reports (labeled with a specific scalar, such as **Force Coefficient**, in the pop-up menu) are added in the same way as other reports.

Once the selection has been made from the pop-up menu, the new report is added to the **Reports** node.



The default name of the new report is the name that appeared on the pop-up menu, suffixed with a number to make this name unique. These reports have various properties to allow for the specification of the reference values required by the report, and also a property to allow for the definition of the parts that should be operated on. Most of these specific reports only operate on boundaries and certain derived parts.

Adding Expression Reports

Expression reports (those labeled **Expression** in the pop-up menu) are added in the same way as other reports.

Once the selection has been made from the pop-up menu, the new report is added to the **Reports** node.



The default name of the new report is the name that appeared on the pop-up menu, suffixed with a number to make this name unique. The expression reports all have three properties to allow for the following:

- Physical dimensions (dimensionless by default)
- Definition of the expression from the function names of report field functions
- Units that are used for specifying the quantity

Renaming Reports

Any of the reports may be renamed using the standard mechanisms, but be careful that the name you apply still reflects the type of report.

If it does not (as in the example below), identify the type of the report from the description area at the bottom of the *Properties* window.



Running Reports

A report is run manually by right-clicking the report node and selecting **Run Report** from the pop-up menu.



The results appear in the *Output* window sorted alphabetically, first by region, then by boundary. This order is the same as the order used in the parts selector for regions and boundaries, as shown in the following screenshot.



With specific report types, the result of the report is printed to a new tab in the *Output* window, which you can work with like any *Output* tab.



If you have already run the report once before, then subsequent runs send the results to the same tab that was created with the first execution of the action.

It is also possible to run all reports at once.

Running All Reports

You can run all the reports in your simulation file with a single command:

- Right-click the **Reports** node and select **Run all reports** from the pop-up menu.



The results of all the reports appear in the *Output* window, in a separate tab for each report.



Monitoring Reports

Reports can be monitored while the solution is iterating by creating a monitor.

This is useful when you are interested in the history of the quantity as the solution evolves, or if you want to use this report as one of your stopping criteria. Report monitors can also be plotted.

Using Reports as Field Functions

When you create a report, a field function corresponding to that report is added automatically.



For example, when you add a report named `Force Coefficient 1`, a drop-down list with field function selections, wherever it is used, includes an item named `Report: Force Coefficient 1`.



If you are using anything other than alphanumeric characters to name a report, use curly brackets (`{}`). For example, if a report happens to have the name `Mass Flow * 2`, its corresponding field function can be interpreted by STAR-CCM+ to multiply the mass flow by 2 in the calculations. Therefore the report would need to be renamed `{MassFlow} * 2`. These brackets are also necessary for report names that include characters from non-Western alphabets.

Using Coordinate Systems with Force and Moment Reports

You can link a report to a Cartesian coordinate system, and thus simplify your direction calculation by changing the frame of reference.

This can be useful if, for example, you are modeling a certain part that is oriented differently from standard (laboratory) coordinates. Additionally, you can change the coordinates for more than one report by editing a single object, rather than modifying each report individually.

The following types of reports use Cartesian coordinate systems:

- Force
- Force coefficient
- Moment
- Moment coefficient
- Reaction force
- Reaction moment

Other than the default laboratory coordinate system, which is suitable for simpler problems, this feature uses the Cartesian coordinate system only. Even if you have various coordinate systems in your simulation, the report's *Coordinate System* property uses a filter that recognizes only the Cartesian coordinate systems in the simulation (in addition to the laboratory coordinate system).

To use a Cartesian coordinate system in a report, first add it to the simulation. Then select the node of the report and choose that coordinate system from the drop-down list of the *Coordinate System* property.


Servers
intake
Reports
Force Coefficient 1
Monitors
Plots
Scenes
Representations
Tools
Annotations
Coordinate Systems
Block-mapped 1
Cartesian 1
Cartesian 2
Laboratory
Spherical 1
Field Functions
Force Coefficient 1 - Properties
Properties
Reference Density
1.0 kg/m^3
Reference Velocity
1.0 m/s
Reference Area
1.0 m^2
Coordinate System
Laboratory
Direction
Force Option
Reference Pressure
Cartesian 1
Cartesian 2
Laboratory
Parts
[Fluid: Free_Stream]
Units

Using Bands with Force and Moment Reports

With advanced Java techniques, you can add flexibility in calculating force and moment by dividing the object into bands.

The following types of reports use bands:

- Force
- Force coefficient
- Force (harmonic)
- Moment
- Moment coefficient

STAR-CCM+ starts creating bands by locating the beginning and end of the object along the direction that you specify. (Specify the direction with the *Direction* property of the force, force (harmonic) and force coefficient reports, or the *Axis* property of the moment and moment coefficient reports.)



The distance between these two locations is divided based on your number of bands. Then values are calculated for each band.

Set this up as follows:

- Complete your entries in the report's properties and its *Number of Bands* expert property.
- Access the Java API. Use the following methods from the Method Summary of the `ForceAndMomentBaseReport` class:
  - `getBandPositions()`
  - `getBandPressureValues()`
  - `getBandShearValues()`
  - `getBandValues()`

Copying Properties of a Report

It is possible to copy properties among similar reports of all categories except for system reports, which do not have properties.

Among reports of the same type, this operation causes all of the properties of the source report to be applied except for *Force Option*, which must be set separately for the destination report.

It is not possible to copy properties between different types of reports.

Deleting Reports

Reports may be deleted using the standard methods, unless a monitor is using the report.

In this case, the **Delete** menu item is grayed and the `<Delete>` key has no effect (no warning is issued).

Report Types Reference

This section details the available reports along with their formulations if necessary.

An additional type, the expression report, exists to combine specific and/or statistical reports.

System Reports

No scalar is selected, and no input parts need be defined. When one of these reports is added to the report manager node, it is removed from the pop-up menu of that node.

System Report Properties

Units	Units that are used for specifying the quantity.

Solver Iteration CPU Time

The solver CPU time is the accumulated CPU time for all processes for the last iteration executed.

Solver Iteration Elapsed Time

The solver elapsed time shows how long it took to execute the last iteration.

Solver CPU Time Per Time-Step

The solver CPU time is the accumulated CPU time for all processes, per time-step (for unsteady simulations), for the last time-step executed.

Solver Elapsed Time Per Time-Step

The solver elapsed time shows how long it took to execute the last time-step (for unsteady simulations).

Total Solver CPU Time

The total solver CPU time is the accumulated CPU time for all processes. Clearing the solution resets the total CPU time to zero.

Total Solver Elapsed Time

The total solver elapsed time is the cumulative time taken to run the entire solution. Clearing the solution resets the total solver elapsed time to zero.

Statistical Reports

Statistical Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The property customizer (...) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Scalar	Scalar field function value ϕ upon which the report is based. The particular formulation for the report can be found in the links above.
Units	Units that are used for specifying the quantity.

Statistical Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

DarsCFD Reports

The DarsCFD Reports have their own properties and use a common report pop-up menu.

There are three types of DarsCFD Report:

- DarsCFD Quantity Maximum
- DarsCFD Quantity Minimum
- DarsCFD Quantity Sum

These reports take a particular value and perform a global operation (maximum, minimum or sum) over all of the processors. For the serial

run, all three types of reports report the same value. Several of the DarsCFD Variable options only produce reports with values when the ISAT model is available. If the ISAT model is not available, these reports have a zero value.

DarsCFD Reports Properties

DarsCFD Variable	Choice of *DarsCFD Variables* available to plot.	
	None	This is the default option.
	Cpu Time	The time that the CPU takes with complex chemistry calculations.
	Elapsed Time	The time that has elapsed during complex chemistry calculations.
	Adds	The number of points that have been added to the table. (ISAT only).
	Grows	The number of points in the table that have increased their region of accuracy. (ISAT only).
	Retrieves	The number of points that are retrieved from the ISAT table; these retrieves avoid expensive direct ODE solver calculations (ISAT only).
	Direct Evaluations	The number of points that are calculated by calling the ODE solver. (ISAT only).
	Table Size	The ISAT table size.
	Queries	The total number of queries given to ISAT.

	Direct ODE Queries	The number of queries that result in direct ODE solver calculations. Low values in this report reflect a good setup of ISAT parameters.
Variable Accumulation Type	Format of displayed value for those *DarsCFD Variables* that report a value of time.	
	Cumulative	The time taken for the complex chemistry solver calculations in all iterations. This is the default option.
	Per Iteration	The time taken for the complex chemistry solver calculation for each iteration
Units	Units that are used for specifying the quantity.	

When you select the properties you require, the name of the report node in the simulation tree updates to reflect your chosen properties.

Element Count

The element count report has its own properties and uses a common report pop-up menu.

It counts the number of elements (cells, faces, parcels, tracks) in the data source.

Element Count Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Units	Units that are used for specifying the quantity.

Frontal Area

The frontal area report has its own properties and uses a common report pop-up menu.

It calculates the combined projected area of all parts in the report onto a user-specified plane. A rasterization technique is used to first generate an image in memory, then sum pixels resulting from the rasterization, and finally scale the result by the characteristic pixel area. You can link this report to a local coordinate system and thus simplify your direction calculation by changing the frame of reference. This can be useful if, for example, you are modeling a certain part that is oriented differently from standard (laboratory) coordinates.

To use a local coordinate system in a report, first add it to the simulation. Then select the node of the report and choose that coordinate system from the drop-down list of the *Coordinate System* property.

Frontal Area Report Properties

Coordinate System	Specifies the coordinate system in which the *Normal* and *View Up* vectors are defined.
Normal	Specifies the normal vector for the plane from which the projection is calculated.
View Up	Sets the up vector for the plane.
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Units	Units that are used for specifying the quantity.

Harmonic Mass Average

The mass average scalar quantity is computed over a volume as:

$$Mass\ \ Average = \frac{\hat{q}_{\phi 0}}{q_0} \tag{2711}$$

where:

$$\hat{q}_{\phi 0} = \left[\mathbf{C} \cdot \mathbf{q}_{\phi}^{*}\right]_{0th\ harmonic} \tag{2712}$$

$$\hat{q}_{0} = \left[\mathbf{C} \cdot \mathbf{q}^{*}\right]_{0th\ harmonic} \tag{2713}$$

$$\mathbf{q}_{\phi}^{*} = \sum_{C} \rho_{C}^{*} \phi_{C}^{*} V_{C} \tag{2714}$$

$$\mathbf{q}^{*} = \sum_{C} \rho_{C}^{*} V_{C} \tag{2715}$$

ρ_C^* is the time-level cell density; ϕ_C^* is the time-level cell scalar quantity, where the scalar could be one of the quantities listed under the *Scalar Option* property; V_c is the cell volume; **C** is the complex Fourier operator that converts time-level solutions into Fourier coefficients; and the subscript "0th harmonic" denotes the time-mean or the 0th harmonic of the Fourier coefficients.

The harmonic mass average report has its own properties and uses a common report pop-up menu.

Harmonic Mass Average Report Properties

Scalar Option	Selects one of the following values: • `Pressure` • `U_Velocity` • `V_Velocity` • `W_Velocity` • `Density` • `Temperature` • `Total Enthalpy` • `Total Pressure` • `Total Temperature` • `Entropy`
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.
Units	Units that are used for specifying the quantity.

Harmonic Mass Average Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Harmonic Mass Flow Average

The mass flow average scalar quantity through a surface is computed as:

$$Mass\ Flow\ Average = \frac{\hat{q}_{\phi 0}}{\hat{q}_0} \tag{2716}$$

where:

$$\hat{q}_{\phi 0} = \left[\mathbf{C} \cdot \mathbf{q}^*_\phi\right]_{0th\ harmonic} \tag{2717}$$

$$\hat{q}_0 = \left[\mathbf{C} \cdot \mathbf{q}^*\right]_{0th\ harmonic} \tag{2718}$$

$$\mathbf{q}^*_\phi = \sum_f \rho^*_C \phi^*_C \left| \mathbf{v}^*_f \cdot \mathbf{a}_f \right| \tag{2719}$$

$$\mathbf{q}^* = \sum_f \rho^*_f \left| \mathbf{v}^*_f \cdot \mathbf{a}_f \right| \tag{2720}$$

ρ^*_f is the time-level face density; ϕ^*_f is the time-level face scalar quantity, where the scalar could be one of the quantities listed under the *Scalar Option* property; $\mathbf{v}^*_f$ is the time-level face velocity vector; $\mathbf{a}_f$ is the face area vector; **C** is the complex Fourier operator that converts time-level solutions into Fourier coefficients; and the subscript "0th harmonic" denotes the time-mean or the 0th harmonic of the Fourier coefficients.

The harmonic mass flow average report has its own properties and uses a common report pop-up menu.

Harmonic Mass Flow Average Report Properties

Scalar Option	Selects one of the following values: • `Pressure`

	• U_Velocity • V_Velocity • W_Velocity • Density • Temperature • Total Enthalpy • Total Pressure • Total Temperature • Entropy
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.
Units	Units that are used for specifying the quantity.

Harmonic Mass Flow Average Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Line Integral

The line integral report is represented in the simulation tree by a node. Right-click this node to access the pop-up menu of the line integral report.

As a statistical report, the line integral largely has the same properties as the others of this type. However its *Field Function* property selects a scalar or vector field function value upon which the report is based.

A line integral refers to an integral of a function, whether it is a scalar or a vector field, along a curve. One example of using a line integral is calculating the amount of work done on a particle by a force field as it moves along an arbitrary curve *C* in that vector field.

To compute a line integral, a parameterization of the curve C is used, that is, C is expressed by a vector function $\mathbf{r}(t)$, with $a \le t \le b$. Here, $\mathbf{r}(a)$ and $\mathbf{r}(b)$ are the initial and end point of the curve, respectively.

The line integral of a scalar field f along a parametric curve C is:

$$\int_C f \; ds = \int_a^b f(\mathbf{r}(t)) \, |\mathbf{r}'(t)| \, \square t \approx \sum_i f(\mathbf{r}(t_i)) \Delta s_i \tag{2721}$$

where Δs_i is the distance between points on the curve.

The line integral of a vector field $\mathbf{F}(\mathbf{r})$is:

$$\int_C \mathbf{F}(\mathbf{r}) \cdot d\mathbf{r} = \int_a^b \mathbf{F}(\mathbf{r}(t)) \cdot \mathbf{r}'(t) \, \square t \; \approx \sum_i \mathbf{F}(\mathbf{r}(t_i)) \cdot \Delta \mathbf{s}_i \tag{2722}$$

where $\Delta \mathbf{s}_i$ is the distance vector between points on the curve.

Reversing the orientation of curve C does not change the value of the line integral of a scalar field but reverses the sign of the line integral of a vector field.

The following limitations exist when you work with vector quantities:

- The start and end points for the integration cannot be user-specified.
- The display does not indicate the direction of integration of the line.

Mass Averaged

The mass averaged scalar quantity is computed over a volume as:

$$Mass \;\; Average \;\equiv \frac{\int_V \rho \, \phi \, dV}{\int_V \rho \; dV} = \frac{\sum_c \rho_c \phi_c V_c}{\sum_c \rho_c V_c} \tag{2723}$$

where ϕ_c is the cell value of the selected scalar, ρ_c is the cell density and V_c is the cell volume.

The mass average over parcels in a Lagrangian phase is computed as:

$$Mass\ Average \equiv \frac{\int_V \rho\,\phi\,dV}{\int_V \rho\ dV} = \frac{\sum_\pi \phi_p m_p n_\pi}{\sum_\pi m_p n_\pi} \tag{2724}$$

where m_p is the mass of a single particle represented by a parcel and $m_p n_\pi$ is the mass of the parcel.

See also the properties of this statistical report.

Mass Flow Averaged

The mass flow averaged scalar quantity through a surface is computed as:

$$Mass\ Flow\ Average \equiv \frac{\oint \rho\,\phi\ \left|\mathbf{v}\cdot d\mathbf{a}\right|}{\oint \rho\ \left|\mathbf{v}\cdot d\mathbf{a}\right|} = \frac{\sum_f \rho_f \phi_f \left|\mathbf{v}_f \cdot \mathbf{a}_f\right|}{\sum_f \rho_f \left|\mathbf{v}_f \cdot \mathbf{a}_f\right|} \tag{2725}$$

where ϕ_f is the face value of the selected scalar, ρ_f is the face density, $\mathbf{v}_f$ is the face velocity vector, and $\mathbf{a}_f$ is the face area vector.

See also the properties of this statistical report.

Maximum

The maximum value report returns the maximum cell value in the selected regions, the maximum face value in the selected boundaries, or the maximum value of ϕ for each parcel in a Lagrangian phase.

See also the properties of this statistical report.

Minimum

The minimum value report returns the minimum cell value in the selected regions, the minimum face value in the selected boundaries, or the minimum value of ϕ for each parcel in a Lagrangian phase.

See also the properties of this statistical report.

Particle Average

This report type is available only when the Lagrangian Multiphase model is activated.

The particle average of ϕ for reports that use Lagrangian phases is represented by:

$$\frac{\sum_{\pi} \phi_p n_\pi}{\sum_{\pi} n_\pi} \tag{2726}$$

where the sum is over all the parcels in the data source and n_π is the number of particles represented by a parcel, the "particle count".

As a statistical report, the particle average has the same properties as the others of this type. However its *Parts* property only selects from among the Lagrangian phases in the simulation.

Sum

The sum report is computed over surfaces as:

$$Area \quad Sum = \sum_{f} \phi_f \tag{2727}$$

The sum report is computed over volumes as:

$$Volume \quad Sum = \sum_{c} \phi_c \tag{2728}$$

The sum report is computed over parcels in a Lagrangian phase as:

$$Parcel \quad Sum = \sum_{\pi} \phi_p \tag{2729}$$

See also the properties of this statistical report.

Surface Area Averaged

This report type is available only when the Lagrangian Multiphase model is activated.

The surface area average over parcels in a Lagrangian phase is computed as:

$$Area \quad Average \equiv \frac{1}{a}\int \phi \square a = \frac{\sum_{\pi} \phi_p A_s n_\pi}{\sum_{\pi} A_s n_\pi} \tag{2730}$$

where A_s is the surface area of a single particle represented by a parcel.

See also the properties of this statistical report.

Surface Average

The surface average of a scalar quantity ϕ is computed on a surface as:

where ϕ_f is the face value of the selected scalar and A_f is the face area.

See also the properties of this statistical report.

Surface Integral

The surface integral of a scalar quantity is computed as:

$$Surface \quad Integral \equiv \int \phi \square a \square = \sum_f \phi_f A_f \tag{2732}$$

where ϕ_f is the face value of the selected scalar and A_f is the face area magnitude.

See also the properties of this statistical report.

Surface Standard Deviation

The standard deviation of a scalar quantity is computed across a surface as:

$$Standard \quad deviation \quad of \ \phi = \sqrt{\frac{\sum_f \left(\phi_f - \bar{\phi}\right)^2 A_f}{\sum_f A_f}} \tag{2733}$$

where $\bar{\phi}$ is the surface average of ϕ, ϕ_f is the face value of the selected scalar and A_f is the face area.

Multiple Parts

When you select multiple parts for this report:

- The average value $\bar{\phi}$ is the collective mean for all the parts—not for each part individually.
- The *Output* window reports the standard deviation of ϕ for the surfaces of each of the component parts individually, as well as for all parts combined. In all cases, however, the reported values use the collective average of ϕ.

See also the properties of this statistical report.

Surface Uniformity

The uniformity of a scalar quantity is computed on a surface as:

$$Uniformity\ index\ \ of\ \phi = 1 - \frac{\sum_f \left(\phi_f - \bar{\phi}\right) A_f}{2\left|\bar{\phi}\right| \sum_f A_f} \tag{2734}$$

where $\bar{\phi}$ is the surface average of ϕ, ϕ_f is the face value of the selected scalar and A_f is the area of a face [444].

The uniformity index describes the distribution of a certain quantity on a surface. If the quantity is distributed equally, the resulting number is 1. This report is useful in applications where a uniform flow rate is desired across a whole area. Heat exchangers, catalysts, and filters are examples of such applications.

Multiple Parts

When you select multiple parts for this report:

- The average value $\bar{\phi}$ is the collective mean for all the parts—not for each part individually.
- The *Output* window reports the uniformity of ϕ for the surfaces of each of the component parts individually, as well as for all parts

combined. In all cases, however, the reported values use the collective average of ϕ.

See also the properties of this statistical report.

Volume Average

The volume average of a scalar quantity is computed as:

$$Volume\ \ Average \equiv \frac{1}{V}\int \phi\ dV\ = \frac{\sum_c \phi_c V_c}{\sum_c V_c} \tag{2735}$$

where ϕ_c is the cell value of the selected scalar and V_c is the cell volume.

The volume average over parcels in a Lagrangian phase is computed as:

$$Volume\ \ Average \equiv \frac{1}{V}\int \phi\ dV\ = \frac{\sum_\pi \phi_p V_p n_\pi}{\sum_\pi V_p n_\pi} \tag{2736}$$

where V_p is the volume of a single particle represented by a parcel.

See also the properties of this statistical report.

Volume Integral

The volume integral of a scalar quantity is computed as:

$$Volume\ \ Integral \equiv \int \phi\ dV\ = \sum_c \phi_c\ V_c \tag{2737}$$

where ϕ_c is the cell value of the selected scalar and V_c is the cell volume.

See also the properties of this statistical report.

Volume Standard Deviation

The standard deviation of a scalar quantity is computed in a volume as:

$$Standard\ \ deviation\ \ of\ \phi = \sqrt{\frac{\sum_c \left(\phi_c - \bar{\phi}\right)^2 V_c}{\sum_c V_c}} \tag{2738}$$

where $\bar{\phi}$ is the volume average of ϕ, ϕ_c is the value of the selected scalar in a cell and V_c is the cell volume.

Multiple Parts

When you select multiple parts for this report:

- The average value $\bar{\phi}$ is the collective mean for all the parts—not for each part individually.
- The *Output* window reports the standard deviation of ϕ for the volumes of each of the component parts individually, as well as for all parts combined. In all cases, however, the reported values use the collective average of ϕ.

See also the properties of this statistical report.

Volume Uniformity

The uniformity of a scalar quantity is computed in a volume as:

$$Uniformity\ \ index\ \ of\ \phi = 1 - \frac{\sum_c \left(\phi_c - \bar{\phi}\right) V_c}{2\left|\bar{\phi}\right| \sum_c V_c} \tag{2739}$$

where $\bar{\phi}$ is the volume average of ϕ, ϕ_c is the value of the selected scalar in a cell and V_c is the cell volume.

This uniformity index describes the distribution of a certain quantity in a volume. If the quantity is distributed equally, the resulting number is 1. This report is useful in applications where a uniform flow rate is desired throughout an entire volume. Heat exchangers, catalysts, and filters are examples of such applications.

See also the properties of this statistical report.

Specific Report Types

6-DOF Report Shared Properties

Coordinate System	The coordinate system in which the direction **d** is specified.
Direction	A vector defining the direction **d** in which the acceleration is reported.
Body	The 6-DOF body for which the acceleration is reported.
Units	Units that are used for specifying the quantity.

Force

The force report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu. It also has an expert property for obtaining values using bands.

The force on a surface is computed as:

$$f = \sum_{f} \left(\mathbf{f}_f^{pressure} + \mathbf{f}_f^{shear} \right) \cdot \mathbf{n}_f \tag{2740}$$

where $\mathbf{f}_f^{pressure}$ and $\mathbf{f}_f^{shear}$ are the pressure and shear force vectors on the surface face f, and $\mathbf{n}_f$ is a user-specified direction vector that indicates the direction in which the force should be computed.

The pressure force vector on surface face f is computed as:

$$\mathbf{f}_f^{pressure} = \left(p_f - p_{ref} \right) \mathbf{a}_f \tag{2741}$$

where p_f is the face static pressure (Eqn.(255)), $\mathbf{a}_f$ is the face area vector, and p_{ref} is the reference pressure. As defined, this is the force that the fluid exerts on the surface. The shear force on surface face f is computed as:

$$\mathbf{f}_f^{shear} = -\mathbf{T}_f \cdot \mathbf{a}_f \tag{2742}$$

where $\mathbf{T}_f$ is the stress tensor at face f. This shear force is exerted on the surface by the fluid.

Note: The report reference pressure p_{ref} is defined relative to the simulation reference pressure. This means that a p_{ref} of 0 Pa does not indicate

absolute zero pressure, but is rather the value of the simulation reference pressure. Further, when the surface on which the force is computed is a closed surface, the choice of p_{ref} affects only the overall accuracy of the result: therefore set it to a value that is close to the average pressure on the surface. For an unclosed surface, p_{ref} is effectively the pressure that is applied on the back-side of the surface, and therefore directly affects the result.

Note: The force report calculates the force on the boundaries (or boundary threshold parts). The boundary of a Porous Baffle Interface consists of an open portion and a closed portion: the force report calculates the total force on both portions. Whereas pressure acts on both the open and closed portions, wall shear acts only on the closed portion. Therefore, you should take care when interpreting the results of force reports on these interfaces.

Force Report Properties

In addition to constants, you can enter field function syntax directly.

Coordinate System	A list of coordinate systems in the simulation for use by this report.	
Direction	A Cartesian vector that indicates the direction in which the force is computed.	
Force Option	Which of the forces (shear and pressure) are considered in the computation of force.	
	Shear	Only the force due to shear, $\mathbf{f}_f^{shear}$ in Eqn. (2654) is considered.
	Pressure	Only the force due to pressure, $\mathbf{f}_f^{pressure}$ in Eqn. (2654) is considered.
	Pressure & Shear	Both pressure and shear forces, $\mathbf{f}_f^{pressure}$ and $\mathbf{f}_f^{shear}$ in Eqn. (2654) are considered.
Parts	A list of parts that are used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector	

	dialog (it lets you interact with other portions of the workspace).
Reference Pressure	The reference pressure for the computation of Eqn.(2655).
Units	Units that are used for specifying the quantity.

Force Report Expert Properties

Number of Bands	A text entry for the number of bands into which to divide the model.
Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Force (Harmonic)

The force (harmonic) report is similar to the Force report and is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu. It also has an expert property for obtaining values using bands.

This report computes the real or imaginary component of the Fourier coefficient of the surface pressure force for a user-specified mode n as:

$$\hat{\mathbf{f}}_n = \frac{1}{N}\sum_{i=0}^{N}\mathbf{f}_j e^{-i\frac{2\pi n j}{N}} \tag{2743}$$

where N is the number of time levels and $\mathbf{f}_j$ is the surface force at time level j which is computed as:

$$\hat{\mathbf{f}}_j = \sum_f \mathbf{f}_{jf}^{pressure} \cdot \mathbf{n}_f \tag{2744}$$

where is $\mathbf{f}_{jf}^{pressure}$ the pressure force vector at time level j and surface face f, and $\mathbf{n}_f$ is a user-specified direction vector that indicates the direction in which the force should be computed.

The pressure force vector $\mathbf{f}_{jf}^{pressure}$ at time level j and surface face f is computed as:

$$\mathbf{f}_{jf}^{pressure} = \left(p_f - p_{ref}\right)\mathbf{a}_f \tag{2745}$$

where p_f is the face static pressure (Eqn.(255)), $\mathbf{a}_f$ is the face area vector, and p_{ref} is the reference pressure. The note on reference pressure provided for the Force report applies to this report.

Force (Harmonic) Report Properties

Mode	A list of coordinate systems in the simulation for use by this report.	
Complex Option	Which component of the force is considered in the computation.	
	Real	Selects the real component.
	Imaginary	Selects the imaginary component.
	Magnitude	Includes the complete force in the computation rather than one of the components.
Coordinate System	A list of coordinate systems in the simulation for use by this report.	
Direction	A Cartesian vector that indicates the direction in which the force is computed.	
Reference Pressure	The reference pressure for the computation of Eqn.(2659).	
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.	
Units	Units that are used for specifying the quantity.	

Force (Harmonic) Report Expert Properties

Number of Bands	A text entry for the number of bands into which to divide the model.
Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.

Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Force Coefficient

The force coefficient report calculates a dimensionless force in a user-specified direction, and can be used for lift and drag calculations.

It is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu. It also has an expert property for obtaining values using bands.

The force coefficient of a surface is defined as follows:

where $\mathbf{f}_f^{pressure}$ and $\mathbf{f}_f^{shear}$ are the pressure and shear force vectors on the surface face f, and ρ_{ref}, v_{ref} and a_{ref} are the reference density, velocity, and area, respectively. $\mathbf{n}_D$ is a user-specified direction vector. To compute drag, for example, specify the direction vector as the direction of travel, and for lift, specify the direction as normal to the direction of travel.

The pressure force vector on surface face is computed as:

where p_f is the face static pressure (Eqn.(255)), $\mathbf{a}_f$ is the face area vector, and p_{ref} is the reference pressure. As defined, this is the force that the fluid exerts on the surface. The note on reference pressure provided for the Force report applies here.

The shear force on surface face f is computed as:

where $\mathbf{T}_f$ is the stress tensor at face f. This shear force is exerted on the surface by the fluid.

Force Coefficient Report Properties

In addition to constants, you can enter field function syntax directly.

Coordinate System	A list of coordinate systems in the simulation for use by this report.
Direction	A Cartesian vector that indicates the direction in which the force is computed.
Force Option	The type of force to be computed.

	Shear	Only the force due to shear, $\mathbf{f}_f^{shear}$ in Eqn. (2660) is considered in the computation.
	Pressure	Only the force due to pressure, $\mathbf{f}_f^{pressure}$ in Eqn. (2660) is considered in the computation.
	Pressure & Shear	Both pressure and shear forces, $\mathbf{f}_f^{pressure}$ and $\mathbf{f}_f^{shear}$, are considered in the computation of Eqn. (2660).
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).	
Reference Area	The reference area for the computation of Eqn. (2660).	
Reference Density	The reference density for the computation of Eqn.(2660).	
Reference Pressure	The reference pressure for the computation of Eqn.(2661).	
Reference Velocity	The reference velocity for the computation of Eqn.(2660).	
Units	Units that are used for specifying the quantity.	

Force Coefficient Report Expert Properties

Number of Bands	A text entry for the number of bands into which to divide the model.
Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.

Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Generalized Vibratory Force

This report type is available only when Harmonic Balance Flutter motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

This report computes the modal force of the user-specified flutter displacement mode on the structure. It is computed as:

where $\mathbf{x}_f^1$ is the first Fourier mode of the displacement of face f, $\mathbf{a}_f^0$ and $\mathbf{a}_f^1$ are, respectively, the zeroth and first Fourier modes of the face areas, and p_f^0 and p_f^1 are, respectively, the zeroth and first Fourier modes of the pressure at the face. The generalized vibratory force F is a complex number and the report outputs the real and imaginary components of this force.

You can use the *Monitor Option* property of this report to monitor the real and imaginary parts of this quantity individually, as well as the magnitude of the complex number F.

Generalized Vibratory Force Report Properties

Monitor Option	Which component of the force is considered in the computation.	
	Real	Selects the real component.
	Imaginary	Selects the imaginary component.
	Magnitude	Includes the complete force in the computation rather than one of the components.
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.	
Units	Units that are used for specifying the quantity.	

Generalized Vibratory Force Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Harmonic Mass Flow

The total mass flow through a boundary is computed as:

where $\hat{\rho}_{0f}$ is the time-mean face density, $\hat{\mathbf{v}}_{0f}$ is the time-mean face velocity vector, and $\mathbf{a}_f$ is the face area vector. The sign convention is such that a positive value of $\dot{m}$ implies that the mass flow is out of the domain.

The harmonic mass flow report has its own properties and is subject to the report node's pop-up menu.

Harmonic Mass Flow Report Properties

Units	Units that are used for specifying the quantity.
Parts	A list of parts that is used as the source for this report. Uses the in-place part selector for editing.

Harmonic Mass Flow Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Heat Exchanger

There are two types of heat exchanger reports: single and dual stream:

- The single stream heat exchanger report uses calculations with one region. By this method, one stream is assumed to have a uniform

temperature and the other stream is modeled by specifying the heat exchanger enthalpy source.

- The dual stream heat exchanger report uses calculations with an interface between two regions. This technique involves the heat exchange between a cold fluid stream and a hot fluid stream. These streams are modeled as two different physics continua (each having different material properties).

  For this report, the monitor values Mass Flux, Inlet Temperature, Outlet Temperature, and Pressure Drop may be selected for both the regions that comprise the dual stream heat exchanger.

  This report also gives efficiency for the two streams as follows:

Each of these reports is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu.

Heat Exchanger Report Properties

Monitor Units	Units in which the Monitor Value is monitored and plotted, but the preferred units are used in the report output for the heat transfer rate, pressures and temperatures.	
Monitor Value	Selects the quantity to be monitored.	
	Mass Flux	Selects mass flux $\dot{m}$.
	Heat Transfer	Selects heat transfer Q.
	Inlet Temperature	Selects inlet temperature T_{in}.
	Outlet Temperature	Selects outlet temperature T_{out}.
	Pressure Drop	Selects pressure drop Δp.
Heat Exchanger	Selects a region (single stream) or interface (dual stream) from a drop-down list of applicable objects in the simulation.	
Cold Downstream Boundaries	Downstream boundary interfaces of the cold stream (dual stream only)	
Hot Downstream Boundaries	Downstream boundary interfaces of the hot stream (dual stream only)	

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).

Heat Exchanger Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Heat Transfer

The heat transfer report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu.

The heat transfer report gives the total heat transfer in Watts at a boundary. It is defined as:

where $\mathbf{q}_f''$ is the face heat flux vector and $\mathbf{a}_f$ is the face area vector. The sign convention is such that a positive value of q implies that heat transfer is from the fluid to the boundary.

It is possible to calculate the heat passing through a portion of a boundary by creating a boundary threshold derived part and running a heat transfer report on that part.

Heat Transfer Report Properties

Units	Units that are used for specifying the quantity.
Parts	A list of parts that is used as the source for this report. Uses the in-place part selector for editing.

Li-Ion Cell

The Li-ion cell report is represented in the simulation tree by a node which has properties and is subject to the report node's pop-up menu.

For more information about setting up Li-ion cell reports, refer to Reporting Li-Ion Cell Parameters.

Li-Ion Cell Report Properties

Monitor Value	Sets the value that the report monitors.	
	State of Charge	Monitors the state of charge *SOC*.
	Open Circuit Potential	Monitors the open circuit voltage *OCV*.
Voltage at 100% SOC	The voltage at which the battery is fully charged.	
Voltage at 0% SOC	The lower cut-off voltage.	
Positive Electrode Parts	A list of parts that is used as the source for positive electrodes for this report. Uses the in-place part selector for editing.	
Negative Electrode Parts	A list of parts that is used as the source for negative electrodes for this report. Uses the in-place part selector for editing.	
Units	The units of the monitored value. This property is updated automatically when the monitor value is changed.	

Li-Ion Cell Report Expert Properties

SOC Step Increment	The state of charge step increment in the state table.	
Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.	
Smooth Values	Whether to calculate the report using node (interpolated) values of the chosen quantity or cell values.	
	Activated	Use node (interpolated) values.
	Deactivated	Use cell values.

Mass Flow

The mass flow report is represented in the simulation tree by a node which has properties and is subject to the report node's pop-up menu.

The total mass flow through a boundary is computed as:

where ρ_f is the face density, $\mathbf{v}_f$ is the face velocity vector and $\mathbf{a}_f$ is the face area vector. The sign convention is such that a positive value of $\dot{m}$ implies that mass flow is out of the domain.

Mass Flow Report Properties

Units	Units that are used for specifying the quantity.
Parts	A list of parts that is used as the source for this report. Uses the in-place part selector for editing.

Mass Flow Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Phase Mass Flow

The phase mass flow report is represented in the simulation tree by a node which has properties and is subject to the report node's pop-up menu.

The mass flow rate of a phase i through a boundary B is computed as:

where $a_{i,f}$ is the volume fraction of phase i at boundary face f, $\rho_{i,f}$ is the density, $\mathbf{v}_f$ is the velocity vector, and $\mathbf{a}_f$ is the area vector. The sign convention is for a positive value of $\dot{m}$ to indicate mass flow out of the domain.

Phase Mass Flow Report Properties

Phase	The phase for which the mass report is run.
Units	Units that are used for specifying the quantity.
Parts	A list of parts that is used as the source for this report. Uses the in-place part selector for editing.

Phase Mass Flow Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Moment

The moment report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu. It also has an expert property for obtaining values using bands.

The moment of force on a surface about an axis is defined:

where $\mathbf{f}_f^{pressure}$ and $\mathbf{f}_f^{shear}$ are the pressure and shear force vectors and$\mathbf{a}$ is a vector defining the axis through point$\mathbf{x}_0$ about which the moment is taken and $\mathbf{r}_f$ is the position of face f relative to $\mathbf{x}_0$.

The pressure force vector on surface face is computed as:

where p_f is the face static pressure (Eqn.(255)), $\mathbf{a}_f$ is the face area vector, and p_{ref} is the reference pressure. As defined, the fluid exerts this force on the surface. The note on reference pressure provided for the Force report applies here.

The shear force on surface face is computed as:

where $\mathbf{T}_f$ is the stress tensor at face f. This shear force is exerted on the surface by the fluid.

Moment Report Properties

In addition to constants, you can enter field function syntax directly.

Axis	A Cartesian vector specifying the axis $\mathbf{a}$ of Eqn. (2670) about which the moment is computed.
Axis Origin	A Cartesian position vector specifying the point $\mathbf{x}_0$ on the axis $\mathbf{a}$ about which the moment is taken.
Coordinate System	A list of coordinate systems in the simulation for use by this report.

Force Option	Which of the forces (shear and pressure) are considered in the computation of the moment.	
	Shear	Only the force due to shear, $\mathbf{f}_f^{shear}$ in Eqn. (2670), is considered.
	Pressure	Only the force due to pressure, $\mathbf{f}_f^{pressure}$ in Eqn.(2670), is considered.
	Pressure & Shear	Both pressure and shear forces, $\mathbf{f}_f^{pressure}$ and $\mathbf{f}_f^{shear}$ in Eqn.(2670) are considered.
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).	
Reference Pressure	The reference pressure for the computation of Eqn.(2671).	
Units	Units that are used for specifying the quantity.	

Moment Report Expert Properties

Number of Bands	A text entry for the number of bands into which to divide the model.
Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Moment Coefficient

The moment coefficient report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu. It also has an expert property for obtaining values using bands.

The moment coefficient of force on a surface about an axis is defined as:

where $\mathbf{f}_f^{\text{pressure}}$ and $\mathbf{f}_f^{\text{shear}}$ are the pressure and shear force vectors and **a** is a vector defining the axis through point $\mathbf{x}_0$ about which the moment is taken and $\mathbf{r}_f$ is the position of face f relative to $\mathbf{x}_0$. ρ_{ref}, v_{ref}, a_{ref} and r_{ref} are the reference density, velocity, area, and radius, respectively.

The pressure force vector on surface face f is computed as:

where p_f is the face static pressure (Eqn.(255)), $\mathbf{a}_f$ is the face area vector, and p_{ref} is the reference pressure. As defined, the fluid exerts this force on the surface. The note on reference pressure provided for the Force report applies here.

The shear force on surface face is computed as:

where $\mathbf{T}_f$ is the stress tensor at face f. This shear force is exerted on the surface by the fluid.

Moment Coefficient Report Properties

In addition to constants, you can enter field function syntax directly.

Axis	A Cartesian vector specifying the axis **a** of Eqn. (2673) about which the moment is computed.	
Axis Origin	A Cartesian position vector specifying the point $\mathbf{x}_0$ on the axis **a** about which the moment is taken.	
Coordinate System	A list of coordinate systems in the simulation for use by this report.	
Force Option	Which of the forces (shear and pressure) are considered in the computation of the moment.	
	Shear	Only the force due to shear, $\mathbf{f}_f^{\text{shear}}$ in Eqn. (2673), is considered.
	Pressure	Only the force due to pressure, $\mathbf{f}_f^{\text{pressure}}$ in Eqn.(2673), is considered.
	Pressure & Shear	Both pressure and shear forces, $\mathbf{f}_f^{\text{pressure}}$

		and $\mathbf{f}_f^{shear}$ in Eqn.(2673) are considered.
Parts	A list of parts that are used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).	
Reference Area	The reference area for the computation of Eqn. (2673).	
Reference Density	The reference density for the computation of Eqn.(2673).	
Reference Pressure	The reference pressure for the computation of Eqn.(2674).	
Reference Radius	The reference radius for the computation of Eqn.(2673).	
Reference Velocity	The reference velocity for the computation of Eqn.(2673).	
Units	Units that are used for specifying the quantity.	

Moment Coefficient Report Expert Properties

Number of Bands	A text entry for the number of bands into which to divide the model.
Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Pressure Drop

The report calculates the mass-flow averaged absolute total pressure difference between two groups of boundary surfaces or interfaces representing the high and low pressure boundaries of the system.

Pressure Drop Report Properties

Units	Units that are used for specifying the quantity.

High Pressure	Specify the boundary surfaces or interfaces with the high pressure value for which to calculate the pressure drop. Click ... (**custom editor**) to open the object selection dialog.
Low Pressure	Specify the boundary surfaces or interfaces with the low pressure value for which to calculate the pressure drop.

Pressure Drop Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Reaction Force

The reaction force report is available only when the Solid Stress model is activated. It calculates the reaction force and applied force on a surface in a user-specified direction.

As well as the reaction and applied forces, the report also displays the body forces acting on the body. The reaction force report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu.

The reaction force on a surface is the summation of the forces due to prescribed displacements, and fixed and elastic foundation boundary conditions. It is computed as:

where $\mathbf{f}_f^{disp.}$, $\mathbf{f}_f^{n.\ disp.}$, $\mathbf{f}_f^{c.\ found.}$, $\mathbf{f}_f^{fixed}$ and $\mathbf{f}_f^{mixed-disp.}$ are the force vectors due to displacement, normal displacement, an elastic foundation, a fixed condition, and the displacement part of a mixed boundary condition on the surface face f, respectively. The variable $\mathbf{n}_f$ is a user-specified direction vector that indicates the direction in which the force should be computed.

The displacement force vector f on surface face is computed as:

where σ_f is the stress tensor and $\mathbf{a}_f$ is the face area vector. The force vectors due to an elastic foundation, a fixed boundary condition, and

the displacement part of the mixed boundary condition, $\mathbf{f}_f^{c.\ found.}$, $\mathbf{f}_f^{fixed}$ and $\mathbf{f}_f^{mixed-disp.}$ are also computed using this equation.

The normal displacement force vector on surface face f is computed as:

where $\mathbf{n}_f$ is the unit vector normal to the cell face.

The applied force on a surface is the summation of pressures, tractions, and forces, and it is computed as:

where $\mathbf{f}_f^{pressure}$, $\mathbf{f}_f^{traction}$, $\mathbf{f}_f^{force}$ and $\mathbf{f}_f^{mixed-traction}$ are the force vectors due to pressure, traction, force, and the traction part of the mixed boundary condition on the surface face f. The variable $\mathbf{n}_f$ is a user-specified direction vector that indicates the direction in which the force is computed.

The pressure force vector on surface face f is computed as:

where p_f is the applied face pressure, or the fluid gauge pressure if the surface is a fluid-solid interface, and $\mathbf{a}_f$ is the face area vector.

The traction force vector on surface face f is computed as:

where $\mathbf{t}_f$ is the traction vector. The force vector due to the traction part of the mixed boundary condition, $\mathbf{f}_f^{mixed-traction}$, is also computed using this equation.

Reaction Force Report Properties

Coordinate System	A list of coordinate systems in the simulation for use by this report.
Direction	A Cartesian vector that indicates the direction in which the force is computed.
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Units	Units that are used for specifying the quantity.

Reaction Force Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Reaction Moment

The reaction moment report is available only when the Solid Stress model is activated.

It is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu.

The reaction moment due to the reaction forces on a surface about an axis is defined:

where $\mathbf{f}_f^{disp.}$, $\mathbf{f}_f^{n.\ disp.}$, $\mathbf{f}_f^{c.\ found.}$, $\mathbf{f}_f^{fixed}$ and $\mathbf{f}_f^{mixed-disp.}$ are the force vectors due to displacement, normal displacement, an elastic foundation, a fixed condition, and the displacement part of a mixed boundary condition on the surface face f, respectively. The variable $\mathbf{a}$ is a vector defining the axis through point $\mathbf{x}_0$ about which the moment is taken and $\mathbf{r}_f$ is the position of face f relative to $\mathbf{x}_0$.

The displacement force vector on surface face is computed as:

where σ_f is the stress tensor and $\mathbf{a}_f$ is the face area vector. The force vectors due to an elastic foundation, a fixed boundary condition, and the displacement part of the mixed boundary condition, $\mathbf{f}_f^{c.\ found.}$, $\mathbf{f}_f^{fixed}$ and $\mathbf{f}_f^{mixed-disp.}$ are also computed using this equation.

The normal displacement force vector on surface face is computed as:

where $\mathbf{n}_f$ is the unit vector normal to the cell face.

The reaction moment due to the applied forces on a surface about an axis is defined:

where $\mathbf{f}_f^{pressure}$, $\mathbf{f}_f^{traction}$, $\mathbf{f}_f^{force}$ and $\mathbf{f}_f^{mixed-traction}$ are the force vectors due to pressure, traction, force, and the traction part of the mixed boundary condition on the surface face f. The variable $\mathbf{n}_f$ is a user-specified

direction vector that indicates the direction in which the force is computed.

The pressure force vector on surface face f is computed as:

where p_f is the applied face pressure, or the fluid gauge pressure if the surface is a fluid-solid interface, and $\mathbf{a}_f$ is the face area vector.

The traction force vector on surface face f is computed as:

where $\mathbf{t}_f$ is the traction vector. The force vector due to the traction part of the mixed boundary condition, $\mathbf{f}_f^{mixed-traction}$, is also computed using this equation.

Reaction Moment Report Properties

Axis	A Cartesian vector specifying the axis **a** of Eqn. (2683) and Eqn.(2686) about which the moment is computed.
Axis Origin	A Cartesian position vector specifying the point $\mathbf{x}_0$ on the axis **a** about which the moment is taken.
Coordinate System	A list of coordinate systems in the simulation for use by this report.
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Units	Units that are used for specifying the quantity.

Reaction Moment Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Rigid Body Acceleration

This report type is available only when 6-DOF Motion is activated.

See also the properties of this 6-DOF report.

This report calculates the linear acceleration of a body in a given direction. It is determined as follows:

Here, $a_{\mathbf{d}}$ is the acceleration of the body in the direction **d**, **v** is the velocity vector of the body, and t is the time.

Rigid Body Angle

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

This report calculates the angle of rotation of a body relative to the initial body orientation. The angle of rotation is either constrained in the interval $[-\pi, \pi]$, or unconstrained to reflect the full rotation of the body:

- Using unconstrained angles is advantageous in cases with large angular motions of the body (for example, a body undergoing multiple revolutions around an axis). When choosing the unconstrained angle option, the reported rotation angle reflects the full rotation of the body. For example, if the body has made two revolutions, the reported angle is . The default option is to report the unconstrained angle.
- The constrained angle is clamped to the interval $[-\pi, \pi]$, or $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ depending on the selected *Angle Report Option* and *Euler Angle Convention*. The constrained angle is provided when backward compatibility is required.

Euler Angle Convention

The *Euler Angle Convention* is interpreted in the following manner. Consider **Rotation X-Y-Z Axis** as an example. For this convention, three rotations describe the orientation of the body. The first rotation is around the x axis of the lab coordinate system (giving the first Euler angle). The second rotation is around the y axis of the lab system (giving the second Euler angle). The third is a rotation around the z axis of the lab system (giving the third Euler angle). The Euler angles are given with respect to fixed rotation axes.



It is also possible to interpret the above Euler angles with respect to moving rotation axes that result from previous rotations. For this purpose, read the axes from back to front. For example, **Rotation X-Y-Z Axis** means that the first rotation is about the z axis of the lab system. The second is around the y axis of the new coordinate system created by rotating the lab system around the z axis. The third rotation is about the x axis of another new coordinate system created by rotating the previous system around its y axis.



The order of rotation is read in the reverse order for the following reason: Let R_X, R_Y and R_Z denote the rotation matrices around the X, Y and Z axes of the laboratory system respectively. For the **Rotation X-Y-Z Axis** option, the overall matrix is $R_Z R_Y R_X$. The same matrix is obtained for moving axes with the order of rotations reversed. If the rotation is first around Z, then Y' (obtained after rotation about Z) and X'' (obtained after a double rotation) - where the primes indicate moved coordinate system axes - then

$$R_{X'}{}'R_{Y'}R_Z = R_Z R_Y R_X \tag{2776}$$

Gimbal Locks

A gimbal lock occurs when the second Euler angle reaches $\frac{\pi}{2}$ or $-\frac{\pi}{2}$. In a gimbal lock, only the sum of the first and the third Euler angles is determined. The reported Euler angles describe the orientation of the body correctly even in a gimbal lock. However, plots of Euler angles over time can show discontinuities. Therefore, try to avoid gimbal locks by choosing an Euler Angle Convention in which the Euler angle reaching a magnitude of $\frac{\pi}{2}$ is the first or the third angle.

Rigid Body Angle Report Properties

Angle Report Option	The axis of the coordinate system relative to which rotation is measured. For fixed axes, the axes of the lab coordinate system are used. For moving axes, the corresponding rotated axes of the lab system are used.	
	Rotation Z Axis	Select rotation about the z-axis.
	Rotation Y Axis	Select rotation about the y-axis.
	Rotation X Axis	Select rotation about the x-axis.
Euler Angle Convention	Select the order of the rotations describing the orientation of the body. For fixed axes, the first rotation is around the first axis of the lab system in the option (giving the first Euler Angle). The second rotation is around the second mentioned axis of the lab system (giving the second Euler Angle). The third rotation is around the last axis of the lab system (giving the third Euler angle). For moving axes, read the order of rotation from back to front.	
	Rotation Y-Z-X Axis	The order of rotation is first Y, then Z and X for fixed axes.
	Rotation X-Y-Z Axis	The order of rotation is first X, then Y and Z for fixed axes.

	Rotation Y-X-Z Axis	The order of rotation is first Y, then X and Z for fixed axes.
	Rotation Z-X-Y Axis	The order of rotation is first Z, then X and Y for fixed axes.
	Rotation X-Z-Y Axis	The order of rotation is first X, then Z and Y for fixed axes.
	Rotation Z-Y-X Axis	The order of rotation is first Z, then Y and X for fixed axes.
Constrained Angle Range	Whether to report the unconstrained angle or the constrained angle in the interval $[-\pi, \pi]$.	
	Activated	Report the constrained angle.
	Deactivated	Report the unconstrained angle. (default)
Body	The 6-DOF body for which the angle of rotation is reported.	
Units	Units that are used for specifying the quantity.	

Rigid Body 1-DOF Angle

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

This report calculates the angle of rotation of a body undergoing One-DOF Rotating Motion relative to the initial body orientation. It reports zero for all other types of motions.

Rigid Body 1-DOF Angle Report Properties

Body	The 6-DOF body for which the angle of rotation is reported.
Units	Units that are used for specifying the quantity.

Rigid Body Angular Acceleration

This report type is available only when 6-DOF Motion is activated.

See also the properties of this 6-DOF report.

This report calculates the angular acceleration of a body with respect to a given direction. It is determined as follows:

Here, $a_{\mathbf{d}}$ is the angular acceleration of the body with respect to the direction **d**, $\vec{\omega}$ is the angular velocity vector of the body, and t is the time.

Rigid Body Angular Momentum

This report type is available only when 6-DOF Motion is activated.

See also the properties of this 6-DOF report.

This report calculates the angular momentum of a body with respect to a given direction. It is determined as follows:

Here, $L_{\mathbf{d}}$ is the angular momentum of the body with respect to the direction **d**, I is the moment of inertia tensor of the body, and $\vec{\omega}$ is the angular velocity vector of the body.

Rigid Body Angular Velocity

This report type is available only when 6-DOF Motion is activated.

See also the properties of this 6-DOF report.

The result of this report is the angular velocity of a body with respect to a given direction. It is calculated as follows:

Here, $\omega_{\mathbf{d}}$ is the angular velocity of the body with respect to the direction , and is the angular velocity vector of the body.

Rigid Body Component Distance

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

The report calculates the distance between two points along a user-specified direction, that is,

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are the two user-specified positions and $\mathbf{s}_d$ is the direction vector.

Rigid Body Component Distance Report Properties

Coordinate System	A list of coordinate systems in the simulation for use in defining the *Direction*.
Direction	Direction vector $\mathbf{s}_d$ on which the distance is projected.
Position 1	The coordinates of the first point, $\mathbf{x}_1$, with respect to *Coordinate System 1*.
Position 2	The coordinates of the second point, $\mathbf{x}_2$, with respect to *Coordinate System 2*.
Coordinate System 1	A list of coordinate systems in the simulation for use in defining *Position 1*.
Coordinate System 2	A list of coordinate systems in the simulation for use in defining *Position 2*.
Units	Units of length in which the distance is reported.

Rigid Body Force

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

The resultant force $\mathbf{f}$ acting on the body (Eqn.(2137)) is computed as:

where f_r is the ramping factor and $\mathbf{f}_g$ is the gravity force acting on the body, computed as:

where $\mathbf{g}$ is the gravitational vector and m is the mass of the body.

The fluid force $\mathbf{f}_{ff}$ acting on the body is calculated as specified in Eqn. (2654). You specify the external force $\mathbf{f}_{ext}$.

Rigid Body Force Report Properties

Coordinate System	A list of coordinate systems in the simulation for use by this report.
Direction	A Cartesian vector that indicates the direction in which the force is computed.
Forces	Specifies which of the body forces to include in the report. The reported value is the sum of the selected forces. Clicking the custom editor button opens the *Forces* selection dialog.
Body	A list of 6-DOF bodies in the simulation for use by this report.
Units	Units that are used for specifying the quantity.

Rigid Body Force Report Expert Properties

Use Raw Forces	Provides option to report resultant force without smoothing or ramping.	
	Activated	The instantaneous force acting on the body (without smoothing or ramping) is reported.
	Deactivated	The instantaneous force as applied in the solver (after any smoothing and ramping) is reported.

Rigid Body Moment

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

The resultant force **n** acting on the body (Eqn.(2138)) is computed as:

where f_r is the ramping factor.

The fluid moment $\mathbf{n}_{ff}$ acting on the body is calculated as specified in Eqn.(2670) with respect to the current center of mass of the body. You specify the external moment $\mathbf{n}_{ext}$.

Rigid Body Moment Report Properties

Coordinate System	A list of coordinate systems in the simulation for use by this report.
Direction	A Cartesian vector that indicates the direction in which the force is computed.
Moments	Specifies which of the body moments to include in the report. The reported value equals the sum of all the moments. Clicking the custom editor button opens the *Moments* selection dialog.
Body	A list of 6-DOF bodies in the simulation for use by this report.
Units	Units that are used for specifying the quantity.

Rigid Body Moment Report Expert Properties

Use Raw Moments	Provides option to report moments without smoothing or ramping.	
	Activated	The instantaneous moment acting on the body (without smoothing or ramping) is reported.
	Deactivated	The instantaneous moment as applied in the solver (after any smoothing and ramping) is reported.

Rigid Body Rotational Energy

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

This report calculates the rotational energy of a body. It is determined as follows:

Here, K_{rot} is the rotational energy of the body, I is the moment of inertia tensor of the body, and $\vec{\omega}$ is the angular velocity vector of the body.

Rigid Body Rotational Energy Report Properties

Body	The 6-DOF body for which the rotational energy is reported.
Units	Units of length in which the distance is reported.

Rigid Body Spring Elongation

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

The report calculates the elongation of a linear spring that is held between a 6-DOF body and the environment, or between two 6-DOF bodies, as defined in a linear spring coupling. If the *Total Length* property is activated, the reported length is:

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are the end-points where the spring is attached. If the *Total Length* property is not activated, the reported length is:

where l_{eq} is the relaxation length of the spring.

Rigid Body Spring Elongation Report Properties

Linear Spring Coupling	A list of linear spring couplings for use by this report.	
Total Length	Sets whether to include the relaxation length of the spring in the reported distance or not.	
	Activated	The total length of the spring including its relaxation length is reported.
	Deactivated	The relaxation length of the spring is subtracted from the total length.
Units	Units of length in which the distance is reported.	

Rigid Body Catenary Length

This report type is available only when 6-DOF Motion is activated. It has its own properties but is subject to the report node's pop-up menu.

The report calculates the length of a catenary that is held between a 6-DOF body and the environment, or between two 6-DOF bodies, as defined in a catenary coupling. If the *Total Length* property is activated, the reported length is:

where L is the total length of the catenary, y(x) is the catenary curve defined in Eqn.(2199), and $\mathbf{x}_1$ and $\mathbf{x}_2$ are the end-points where the catenary is attached. If the *Total Length* property is not activated, the reported length is:

where l_{eq} is the relaxation length of the catenary.

Rigid Body Catenary Report Properties

Catenary Coupling	A list of catenary couplings for use by this report.	
Total Length	Sets whether to include the relaxation length of the catenary in the reported distance or not.	
	Activated	The total length of the catenary including its relaxation length is reported.
	Deactivated	The relaxation length of the catenary is subtracted from the total length.
Units	Units of length in which the catenary length is reported.	

Rigid Body Total Distance

This report type is available only when 6-DOF Motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

The report calculates the absolute value of the distance r_p between two given points, that is,

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are two user-specified positions.

Rigid Body Total Distance Report Properties

Position 1	The coordinates of the first point, $\mathbf{x}_1$, with respect to *Coordinate System 1*.
Position 2	The coordinates of the second point, $\mathbf{x}_2$, with respect to *Coordinate System 2*.
Coordinate System 1	A list of coordinate systems in the simulation for use in defining *Position 1*.
Coordinate System 2	A list of coordinate systems in the simulation for use in defining *Position 2*.
Units	Units of length in which the distance is reported.

Rigid Body Translation

This report type is available only when 6-DOF Motion is activated.

See also the properties of this 6-DOF report.

This report calculates the translational displacement of a body with respect to its initial position in a given direction. It is determined as follows:

Here, $\delta_{\mathbf{d}}$ is the translational displacement of the body in the direction $\mathbf{d}$, and $\vec{\delta}$ is the translational displacement vector of the body.

Rigid Body Velocity

This report type is available only when 6-DOF Motion is activated.

See also the properties of this 6-DOF report.

This report calculates the velocity of a body in a given direction. It is determined as follows:

Here, $v_{\mathbf{d}}$ is the velocity of the body in the direction $\mathbf{d}$, and $\mathbf{v}$ is the velocity vector of the body.

Sauter Mean Diameter

This report type is available only when the Lagrangian Multiphase model is activated.

It has its own properties but is subject to the report node's pop-up menu.

The Sauter mean diameter (SMD) applies to parcels, with the Material Particles model.

in which D_p is the particle diameter and n_π is the number of particles that is represented by a parcel: the "particle count".

Sauter Mean Diameter Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Units	Units that are used for specifying the quantity.

Sauter Mean Diameter Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Thrust

The thrust report integrates the momentum flux and the pressure force over a specified boundary or derived part.

The thrust report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu. This report is active for flow models.

The thrust is computed as:

where ρ_f is the fluid density, $\mathbf{v}$ is the velocity vector, p_f is the exit pressure, p_{atm} is the atmospheric pressure, and $\mathbf{a}_f$ is the surface area vector of the face.

The governing equation that is being solved is the reaction force at the boundary:

This equation is illustrated in the following diagram:



Thrust Report Properties

Direction	Specifies the Cartesian vector that is used to compute the scalar product with the reaction thrust vector. This lets you calculate the component of the thrust that is acting in a particular direction.
Atmospheric Pressure	The atmospheric pressure for the computation of Eqn.(2700).
Coordinate System	A list of coordinate systems that this report can use.
Parts	A list of parts that are used as the source for this report. Opens the in-place part selector for editing. The ... (**property customizer**) button activates a modeless version of the selector dialog (it lets you interact with other portions of the workspace).
Units	Units that are used for specifying the quantity.

Thrust Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.

Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Tumble/Swirl Report

This report is typically used for ICE applications and calculates the swirl/tumble ratio (dimensionless)

defined by:

where:

X_i, Y_i, Z_i	Centroid coordinate of cell (m)
X_m, Y_m, Z_m	Center of mass/volume/origin of the cylinder (m)
u_i, v_i, w_i	Velocity at cell (m/s)
N	Rotational speed (rev/min)
M_x, M_y, M_z	Angular momentum about the x-, y-, and z-axes (kg m^2/s)
I_x, I_y, I_z	Mass moment of inertia about the x-, y-, and z-axes (kg m^2/s)

Swirl/Tumble Report Properties

Units	Units that are used for specifying the quantity.	
Tumble/Swirl	**Tumble**	Select the tumble report. Tumble is defined as the fluid rotation about the plane that contains the *Tumble Origin* and that is orthogonal to the cylinder *Axis*. For example, if the cylinder axis is Z, tumble corresponds to S_x and S_y as given by Eqs. (2702) and (2703).
	Swirl	Select the swirl report. Swirl is defined as the

		fluid rotation around the cylinder **Axis.** For example, if the cylinder axis is Z, swirl corresponds to S_z as given by Eqn.(2704).
Tumble Origin	Point within the cylinder. This point is usually the center of the cylinder and corresponds to X_m, Y_m, Z_m in Eqs. (2702) to (2704). Defines the plane upon which the principal tumble axis lies. Tumble is calculated about this plane.	
Rotational Rate	Rotational speed of the ICE engine.	
Axis	Defines the orientation, that is, the axis of the cylinder with respect to **Coordinate System**.	
Coordinate System	A list of coordinate systems that this report can use.	
Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.	

Swirl/Tumble Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Uniformity Deviation Report

This report calculates the deviation of the local normal velocity from the mass-flow averaged value on a specified set of boundary surfaces or on interfaces.

The uniformity deviation of the velocity normal to the surface (at each face $u_f = \mathbf{u}_f \cdot \mathbf{a}_f$) is:

where the variance of normal speed is:

$$\sigma^2 = \frac{\sum_f \left(u_f - \bar{u}\right)^2 \|\mathbf{a}_f\|}{\sum_f \|\mathbf{a}_f\|} \tag{2799}$$

and the average normal speed is:

$$\bar{u} = \frac{\sum_f u_f \cdot \mathbf{a}_f}{\sum_f \|\mathbf{a}_f\|} \tag{2800}$$

Uniformity Deviation Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.
Units	Units that are used for specifying the quantity.

Uniformity Deviation Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Work Per Cycle

This report type is available only when Harmonic Balance Flutter motion is activated.

It has its own properties but is subject to the report node's pop-up menu.

This report computes the work done by the fluid on the structure (or the boundary surfaces selected). The work per cycle is computed as:

where the summation is over all the faces of the boundary surface, T is the time period, and $\hat{q}_f^0$ is the mean of the integrand $\hat{q}_f^*$ (which

represents the power at each time level) over all the time levels. The power at each time level is computed as:

where p_f^* is the time level pressure at the face, D is the pseudo-spectral time derivative operator, $\mathbf{x}_f^*$ is the time level displacement of the face, $\mathbf{v}_f^* = D\mathbf{x}_f^*$ is the time level velocity of the face, and $\mathbf{a}_f^*$ is the time level face area.

A negative value of the converged work-per-cycle shows that the blade is dynamically stable and does not flutter.

Work Per Cycle Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.
Units	Units that are used for specifying the quantity.

Work Per Cycle Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Coal Volatile Matter Yield

This report type is available only when any of the Raw Coal Devolatilization models (see Understanding the Coal Combustion Phase Models) is activated.

It has its own properties but is subject to the report node's pop-up menu.

The Coal Volatile Matter (VM) Yield report applies to parcels with the Coal Combustion Model and is calculated as follows:

Coal Volatile Matter Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.

Units	Units that are used for specifying the quantity.

Coal Volatile Matter Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Char Burn Out

This report type is available only when any of the Raw Char Oxidation models (see Understanding the Coal Combustion Phase Models) is activated.

It has its own properties but is subject to the report node's pop-up menu.

The Char Burn Out report applies to parcels with the Coal Combustion Model and is calculated as follows:

Char Burn Out Report Properties

Parts	A list of parts that is used as the source for this report. Opens the in-place part selector for editing.
Units	Units that are used for specifying the quantity.

Char Burn Out Report Expert Properties

Representation	Selects a representation. Additional selections appear as representations that are added to the simulation.
Smooth Values	When this property is activated, the report is calculated using node (interpolated) values of the chosen quantity rather than cell values.

Expression Report

The expression report depends on a combination of reports, except for system reports. It works with the field function generated for each report you want to use.

This report is represented in the simulation tree by a node, which has its own properties but is subject to the report node's pop-up menu.

For example, suppose that you want to plot the pressure drop between two boundaries, such as the inlet and outlet of a pipe. Assuming you created a report of surface average pressure for each boundary, you can directly create a report of those reports by entering the following function names in the *Definition* property:

```
$AveragePressure1Report - $AveragePressure2Report
```

The field functions that can be used to define expressions cannot be a function of space; they have to evaluate to a single number as the expressions are constants in space. Variables such as `$Time`, `$TimeStep` and `$Iteration`, as well as the field functions associated with reports, are valid. Variables such as `$Temperature` or `$Pressure` have different values at every different point in space and cannot be used in expressions.

You can also use periodicity with an expression report.

Expression Report Properties

Periodicity	Controls whether to make the report periodic.	
	Non-periodic	Leaves periodicity deactivated.
	Periodic	Causes the report to loop between two specified values. Activates the Periodic Range sub-node.
Dimensions	The physical dimensions of this report, which you can customize.	
Definition	The formulation of this report using the function name of the STAR-CCM+ field functions corresponding to the reports selected. Entering an invalid definition causes an `Unable to compile` error message.	
Units	Units that are used for specifying the quantity.	

Using Periodicity

The **Periodic Range** sub-node of the expression report defines the window over which the report values repeat. It instructs STAR-CCM+ that it is possible for the underlying function to loop over the interval bounds. With this instruction, the difference between two consecutive report values must be computed as a shortest-distance comparison, considering the potential for looping.

- To activate the **Periodic Range** sub-node, select the expression report node and set *Periodicity* to `Periodic`.
- To specify a rotation angle of, for example, 2π:
  - Select the **Periodic Range** sub-node.
  - Set *Periodic Minimum* to `0.0`.
  - Set *Periodic Maximum* to `6.283185`.

Periodic Range Properties

Periodic Maximum	The highest value in the looping range for this expression report.
Periodic Minimum	The lowest value in the looping range for this expression report.

Pop-Up Menu of the Reports Manager

These pop-up menu items are available when you right-click the **Reports** manager node.

New Report	Creates a report of the type you select in the submenu.	
	6-DOF Body Angular Acceleration	Creates a rigid body angular acceleration report. This report type is available only when 6-DOF Motion is activated.
	6-DOF Body Angular Momentum	Creates a rigid body angular momentum report. This report type is available only when 6-DOF Motion is activated.
	6-DOF Body Angular Velocity	Creates a rigid body angular velocity report. This report type is available only when 6-DOF Motion is activated.
	6-DOF Body Force	Creates a rigid body force report. This report type is available only when 6-DOF Motion is activated.
	6-DOF Body Linear Acceleration	Creates a rigid body acceleration report. This report type is available only when 6-DOF Motion is activated.
	6-DOF Body Moment	Creates a rigid body moment report. This report type is available only when 6-DOF Motion is activated.

6-DOF Body Orientation	Creates a rigid body angle report. This report type is available only when 6-DOF Motion is activated.
6-DOF Body Rotational Kinetic	Creates a rigid body rotational energy report. This report type is available only when 6-DOF Motion is activated.
6-DOF Body Translation	Creates a rigid body translation report. This report type is available only when 6-DOF Motion is activated.
6-DOF Body Velocity	Creates a rigid body velocity report. This report type is available only when 6-DOF Motion is activated.
Surface Average	Creates a surface average report on any surfaces.
Dars-CFD Quantity Maximum	Creates a Dars-CFD Quantity Maximum report. This report is available only if the Homogeneous Reactor model is chosen to model reacting multi-component fluid.
Dars-CFD Quantity Minimum	Creates a Dars-CFD Quantity Minimum report. This report is available only if the Homogeneous Reactor model is chosen to model reacting multi-component fluid.

Dars-CFD Quantity Sum	Creates a Dars-CFD Quantity Sum report. This report is available only if the Homogeneous Reactor model is chosen to model reacting multi-component fluid.
Element Count	Creates an element count report on any surfaces.
Expression	Creates an expression report on any surfaces.
Force	Creates a force report on boundary surfaces.
Force (Harmonic)	Creates a force (harmonic) report.
Force Coefficient	Creates a force coefficient report on boundary surfaces.
Frontal Area	Creates a frontal area report on boundary surfaces.
Generalized Vibratory Force	Creates a generalized vibratory force report. This report type is available only when Harmonic Balance Flutter motion is activated.
Harmonic Mass Average	Creates a harmonic mass average report.
Harmonic Mass Flow	Creates a harmonic mass flow report.
Harmonic Mass Flow Average	Creates a harmonic mass flow average report.
Heat Exchanger (Dual Stream)	Creates a dual stream heat exchanger report.

Heat Exchanger (Single Stream)	Creates a single stream heat exchanger report.
Heat Transfer	Creates a heat transfer report on boundary surfaces.
Line integral	Creates a line integral report which can be used for one-dimensional parts.
Mass Averaged	Creates a mass averaged report for regions and Lagrangian phases.
Mass Flow	Creates a mass flow report on boundary surfaces.
Mass Flow Averaged	Creates a mass flow averaged report on boundary surfaces.
Maximum	Creates a maximum report on regions, any surfaces, and Lagrangian phases.
Minimum	Creates a minimum report on regions, any surfaces, and Lagrangian phases.
Moment	Creates a moment report on boundary surfaces.
Moment Coefficient	Creates a moment coefficient report on boundary surfaces.
Particle Average	Creates a particle average report on Lagrangian phases. This report type is available only when the Lagrangian

		Multiphase model is activated.
	Reaction Force	Creates a reaction force report on boundary surfaces. This report type is available only when the Solid Stress model is activated.
	Reaction Moment	Creates a reaction moment report on boundary surfaces. This report type is available only when the Solid Stress model is activated.
	Sauter Mean Diameter	Creates a Sauter mean diameter report on Lagrangian phases. This report type is available only when the Lagrangian Multiphase model is activated.
	Solver CPU Time per Iteration	Creates a solver CPU time per iteration report. This item is removed from the pop-up menu once one has been added to the simulation.
	Solver CPU Time per Time-Step	Creates a solver CPU time per time-step report. This item is excluded from the pop-up menu once one has been added to the simulation, or when a Steady time model is selected.
	Solver Elapsed Time per Iteration	Creates a solver elapsed time per

		iteration report. This item is removed from the pop-up menu once one has been added to the simulation.
	Solver Elapsed Time per Time-Step	Creates a solver elapsed time per time-step report. This item is excluded from the pop-up menu once one has been added to the simulation, or when a Steady time model is selected.
	Sum	Creates a sum report on regions, any surfaces, and Lagrangian phases.
	Surface Area Average	Creates a surface area average report on Lagrangian phases. This report type is available only when the Lagrangian Multiphase model is activated.
	Surface Integral	Creates a surface integral report on any surfaces.
	Surface Standard Deviation	Creates a surface standard deviation report on any surfaces.
	Surface Uniformity	Creates a surface uniformity report on any surfaces.
	Total Solver CPU Time	Creates a total solver CPU time report. This item is removed from the pop-up menu once one has been added to the simulation.

	Total Solver Elapsed Time	Creates a total solver elapsed time report. This item is removed from the pop-up menu once one has been added to the simulation.
	Volume Averaged	Creates a volume averaged report for regions and Lagrangian phases.
	Volume Integral	Creates a volume integral report for regions.
	Volume Standard Deviation	Creates a volume standard deviation report for volume parts.
	Volume Uniformity	Creates a volume uniformity report for volume parts.
	Work Per Cycle	Creates a work per cycle report. This report type is available only when Harmonic Balance Flutter motion is activated.
Run all reports	Runs all reports at once, generating the results in separate tabs in the *Output* window.	
New Group	Facilitates the grouping of objects by creating a group node within the manager node.	
Group By	Allows you to group the objects according to any of the properties of the objects in the manager node, listed in a submenu of this menu item.	
Ungroup	Reverses the grouping process.	

See also: common menu items.

Reports Bibliography

Weltens, H., Bressler, H., Terres, F., Neumaier, H. et al., "Optimisation of Catalytic Converter Gas Flow Distribution by CFD Prediction," SAE Technical Paper 930780, 1993.

Monitoring the Solution

Monitors sample, save, and plot solution data.

The **Monitors** manager node, which has its own properties and pop-up menu, contains all of the monitor nodes in the simulation tree.



Residual monitors are a specific kind of monitor created automatically.

This section explains the following about STAR-CCM+ monitors:

What Is a Monitor?

Monitors enable summary information from the simulation to be sampled and saved during the solution.

This information can then be displayed in an XY plot while the solution is progressing, or post-processed after iteration or unsteady time-stepping is complete. They are useful for watching the behavior of numerical or engineering quantities as the solution evolves, and can be used to define stopping criteria that can stop the solution iteration.

Specific monitors, represented in the simulation tree by nodes that have properties and pop-up menus, may be added one of two ways:

- Automatically, based on the residuals computed by solvers in the simulation (which, in turn, are driven by the model selection that you made in the continua)
- Manually from existing reports

The residual monitors that are added by solvers are listed in the solvers section of the relevant modeling chapter.

Monitor Properties

These properties are the same for residual monitors and report monitors, with the exception of one property.

Because of the similarity, the information for both types is documented here, and differences between the two monitor types are noted where relevant.

Report	This property appears only if the monitor was created from a report. It indicates the name of that report (read-only).	
Active	This property appears only if this is a residual monitor (read-only).	
	Activated	The monitor is active in the current simulation.
	Deactivated	The monitor is inactive. This would happen if you changed models in the middle of the simulation and

		you had not deleted any monitors specific to the previous models.
Trigger	Defines what action will trigger collection of a sample. For report monitors, the default setting depends on whether you are using the Steady model (iterations) or Implicit Unsteady model (time-steps).	
	Iteration	A sample will be taken every `N` iterations, where `N` is defined by the *Frequency* (default for residual monitors).
	Time-Step	A sample will be taken every `N` time-steps, where `N` is defined by the **Frequency**.
	Inactive	No sample is taken.
Maximum Plot Samples	Limits the number of sample points used in the monitor plot.	
Frequency	Determines how often the sample will be taken, based on the setting in *Trigger*. Samples are taken when the number of samples is evenly divisible by the frequency.	
Normalization Option	Defines how the monitor data will be normalized.	
	Auto	Lets STAR-CCM+ do the normalization as derived in the Automatic Normalization node that is added as a sub-node when you choose this option.
	Manual	Lets you define a specific normalization factor in the Manual Normalization node that is added as a sub-

		node when you select this option.

Monitor Pop-Up Menu

This pop-up menu is the same for residual monitors and report monitors.

Create Plot from Monitor	Creates a monitor plot by adding a node to the **Plots** node named `Monitor Plot N` where `N` is the total number of monitor plot nodes.
Create Stopping Criterion from Monitor	Creates a monitor-based stopping criterion by adding a node to the **Stopping Criteria** node named `Aaa Criterion` where `Aaa` is the name of this monitor node.
Export...	Exports the data to a `.csv` file. The exported data retains the preferred units.

Customizing Monitors

The way in which monitors sample data and send data to an XY plot can be customized to suit the particular quantity being monitored.

This section describes how to change the following:

Changing the Trigger Event and Sampling Frequency

Trigger events control how often the monitor samples data.

See Trigger Types.



The frequency of trigger events in which a monitor samples data can be modified from the sub-node properties. See Trigger Properties.



For more information on selecting the trigger type and defining the update frequency, refer to Setting the Trigger for an Update.

If you wish to average results per iteration, it is recommended that you first read the section on choosing between steady and unsteady models.

Note: The *Trigger* property of a monitor does not control the frequency at which the data is plotted. It specifies the frequency at which the monitor collects the data. To control the frequency at which the data is queried and plotted, use the Update feature of the plot node itself. If you want to plot data against physical time, make sure that the same **Physical Time** monitor exists throughout the solution. Otherwise the plot data appears incomplete. If you are running a steady solution but expect to need a plot with physical time, start by running the unsteady solver then switch to the steady one; this ensures that Physical Time exists in the simulation.

Setting Maximum Plot Samples

The *Maximum Plot Samples* property sets an upper limit on the number of points that will be displayed in the monitor plot.

When this number of points has been sampled, STAR-CCM+ will begin to reuse plot storage for points obtained earlier in the simulation. This will result in fewer data points being allocated to earlier parts of the simulation. This allocation applies to the plot only; all of the monitor's solution data points are retained in the simulation file.



In a long simulation STAR-CCM+ might generate many thousands of points, for example one per iteration. After a certain number of points, however, additional ones do not provide any visual improvement to the graph and may result in performance degradation. Limiting the number of points may improve performance in drawing the graph (specifically graphics performance, client-server data exchange and memory consumption).

The monitor tends to retain more data toward the end of the run when the number of samples has exceeded the maximum specified in this property.

Normalizing the Monitor Data

The data in a monitor can be normalized one of two ways:

Normalizing Manually

To normalize the data using a specified value, click the monitor node and select **Manual** from the pull-down menu under the *Normalization Option* property in the *Properties* window.

The **Manual Normalization** node, with its own properties, will be added as a sub-node of the monitor node.



The normalization factor can be specified by expanding the monitor node, clicking on the **Manual Normalization** child node, and then modifying the value in the *Normalization Factor* property in the *Properties* window.



Manual Normalization Properties

Normalization Factor	The normalization factor for the monitored quantity.

Normalizing Automatically

To normalize the data automatically, click the monitor node and select **Auto** from the pull-down menu under the *Normalization Option* property in the *Properties* window.

The **Automatic Normalization** node, with its own properties, will be added as a sub-node of the monitor node.



You must enter the number of samples from the start from which the monitor will search for a maximum normalization factor. Expand the monitor node, click the **Automatic Normalization**child node and modify the *Normalization Number* property in the *Properties* window. The normalization factor found is displayed under the read-only *Normalization Factor* property.

Automatic Normalization Properties

Normalization Factor	The normalization factor for the monitored quantity. Computed as the maximum monitor value over the first N iterations or time-step events of the monitor, where N is the *Normalization Number* (read-only).
Normalization Number	The number of iterations or time-step events, from the start of the solution, used to compute the *Normalization Factor*.

Copying Properties of a Normalization Node

When properties are copied from one normalization node to another, all of the properties are applied so that the destination normalization node matches the source one.

However, it is not possible to copy properties between different types of normalization nodes.

Using Report-Based Monitors

This section explains:

See Also

Using Maximum and Minimum Reports

What Is a Report-Based Monitor?

Report-based monitors are useful for monitoring quantities of engineering interest during a simulation. They provide a dynamic record of these quantities during solution iteration.

They may be viewed, either as the solution iterates, or after the solution has converged, using a monitor plot.

Any report can be used as the basis of a monitor. Report-based monitors can be important for judging convergence, and may be used as stopping criteria.

Similar to other objects, report-based monitors can be renamed within STAR-CCM+. Report-based monitors can only be deleted if they are not displayed in an XY plot or used to define a stopping criterion.

When a report-based monitor is copied within a simulation, its properties are copied. You can only copy this object between simulation files if the destination simulation has a report of the same name.

Creating Report-Based Monitors

To create a monitor from a report, right-click on the report node and choose **Create Monitor from Report**.



The default name of the monitor will be based on the report type. However, renaming the monitor is possible.

Creating Report-Based Monitors and Plots

To create a monitor plot simultaneously with a report-based monitor, right-click on the report node and choose **Create Monitor and Plot from Report**.



Using Field Function Monitors

This section explains:

What Is a Field Function Monitor?

There are six varieties of field function monitor. During a simulation run, they collect data samples for a selected part at a frequency that you specify.

You can control the sample count by specifying the start time and stop time, and by using a sliding sample window.

- Field co-variance -- computes the co-variance of two field functions.
- Field max -- obtains the highest value of a field function for the selected part.
- Field mean -- obtains the average value of a field function for the selected part.
- Field min -- obtains the lowest value of a field function for the selected part.
- Field sum -- provides the total value of a field function for the selected part.
- Field variance -- computes the variance of a field function.

These monitor nodes generally share common properties and a pop-up menu.



Once you set their *Field Function* and *Parts* properties, these monitors can be displayed as scalar quantities just like other field functions.

Field Function Monitor Properties

Publish Mean	For the Field Variance and Field Co-Variance monitors, determines whether STAR-CCM+ will generate a separate field function of the mean of the function being sampled. Two separate properties exist for the Field Co-Variance monitor's field function selections: *Publish Mean 1* and *Publish Mean 2*.	
	Activated	Enables this feature.
	Deactivated	Leaves this feature disabled.
Field Function	One or more field functions which the monitor is to sample in order to perform its calculations. Two separate properties exist for the Field Co-Variance monitor's field function selections: *FieldFunction 1* and *FieldFunction 2*.	
Trigger	Provides choices for triggering the action:	
	Iteration	Calculates based on the specified number

		of iterations in the *Frequency* property. If you wish to average results per iteration, it is recommended that you first read the section on choosing between steady and unsteady models.
	Time-Step	Calculates based on the specified number of time-steps in the *Frequency* property (default).
	Inactive	No sample is taken.
Parts	The model parts over which the field function monitor collects samples. Opens the in-place part selector for editing.	
Sample Count	The number of data samples collected for each cell within the monitor (read-only).	
Enable sliding sample window	Calculates a set of the most recent samples, with the number of samples in that set specified by the user. Activates the Sliding Sample Window Size node.	
	Activated	Enables this feature.
	Deactivated	Leaves this feature disabled.
Title	The name of the field function which contains the values computed by the field function monitor (read-only).	
Collection Start	The solution step at which to begin collecting data. Samples will be collected at time-steps or iterations after this start time, based on the *Frequency* property.	
Enable collection stop	Specifies the final sample time or iteration after which to stop collecting data. Activates the Stop Specification node.	
	Activated	Enables this feature.
	Deactivated	Leaves this feature disabled.

Frequency	Determines how often the sample will be taken, based on the setting in *Trigger*. Samples are taken when the number of samples is evenly divisible by the frequency.

Field Function Monitor Pop-Up Menu

Reset Monitor Data	Clears all samples and accumulated data from the field function monitor.

What Is Field Co-Variance?

The Field Co-Variance is the statistical co-variance of two field function values as the solution progresses, computed for each cell or boundary face in the selected parts.

Two separate properties exist for the Field Co-Variance monitor's field function selections: *FieldFunction 1* and *FieldFunction 2*.

The co-variance is based on data samples collected at a user-specified frequency for each cell of the parts being monitored. The sample co-variance, $\mathrm{Cov}(x,y)$, is computed for each cell as:

$$\mathrm{Cov}\left(x,y\right)=\frac{1}{N-1}\left[\sum_{i=1}^{N}x_i y_i-\frac{1}{N}\left(\sum_{i=1}^{N}x_i\right)\left(\sum_{i=1}^{N}y_i\right)\right]$$

where:

- $x_i$is the Field Function 1 value at collection step i.
- y_i is the Field Function 2 value at collection step i.
- N is the number of time-steps or iterations for which data has been collected.

What Is Field Max?

The Field Max is the maximum value of the specified field function which occurred within each cell or boundary face in the selected parts.

The field function is sampled at a user-specified frequency for each cell of the parts being monitored and the maximum value retained as:

$$\mathrm{Cell}_{\max}=\max\left(\mathrm{Cell}_{\max},\mathrm{Cell}_{current}\right)$$

where $\mathrm{Cell}_{current}$ indicates the field function value at the current iteration or time-step.

What Is Field Mean?

The Field Mean is the average of a field function value as the solution progresses, computed for each cell or boundary face in the selected parts.

This arithmetic mean is based on data samples collected at a user-specified frequency for each cell of the parts being monitored. The arithmetic mean, $\bar{x}$, is computed for each cell as:

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i$$

where x_i is the field function value at collection step i, and N is the number of time-steps or iterations for which data has been collected.

What Is Field Min?

The Field Min is the minimum value of the specified field function which occurred within each cell or boundary face in the selected parts.

The field function is sampled at a user-specified frequency for each cell of the parts being monitored and the minimum value retained as:

$$\mathrm{Cell}_{\min} = \min\left(\mathrm{Cell}_{\min}, \mathrm{Cell}_{current}\right)$$

where $\mathrm{Cell}_{current}$ indicates the field function value at the current iteration or time-step.

What Is Field Sum?

The Field Sum is the total of a field function value as the solution progresses, computed for each cell or boundary face in the selected parts.

This sum is composed of data samples collected at a user-specified frequency for each cell of the parts being monitored. The total, x_N , is computed for each cell as:

$$x_N = \sum_{i=1}^{N} x_i$$

where x_i is the field function value at collection step i, and N is the number of time-steps or iterations for which data has been collected.

What Is Field Variance?

The Field Variance is the sample variance of a field function value as the solution progresses, computed for each cell or boundary face in the selected parts.

This statistical variance is based on data samples collected at a user-specified frequency for each cell of the parts being monitored. The sample variance, $s^2(x)$, is computed for each cell as:

$$s^2\left(x\right) = \frac{1}{N-1}\left[\sum_{i=1}^{N} x_i^2 - \frac{1}{N}\left(\sum_{i=1}^{N} x_i\right)^2\right]$$

where x_i is the field function value at collection step i , and N is the number of time-steps or iterations for which data has been collected.

Using a Sliding Sample Window

The sliding sample window is the N most recent samples of data during the simulation.

To evaluate this set of samples for a field function monitor, you need to activate the *Enable sliding sample window* property of that monitor first, as shown in the following screenshot.



This adds the **Sliding Sample Window Size** node to the object tree under the parent field function monitor node. To specify the number of samples (the "window"), select this node and set the *Sliding Sample Window Size* property to the desired value in the properties window.



Sliding Sample Window Size Properties

Use this property to specify the window size.

Sliding Sample Window Size	Specifies the number of most recent samples to keep.

Specifying When Data Collection Stops

You can set a time or iteration at which the monitor stops collecting data. This can let you get results for specific intervals within the simulation run.

For instance, if you are simulating an impeller and you want to obtain the mean pressures for seven revolutions lasting 0.08 seconds each, you would create seven field mean monitors. The monitors would have settings that increase by 0.08, for example:

- **Field Mean 1** -- collection start = 0 and collection stop = 0.08
- **Field Mean 2** -- collection start = 0.08 and collection stop = 0.16
- **Field Mean 7** -- collection start = 0.48 and collection stop = 0.56

To apply this setting to a particular field function monitor, first you need to select the trigger type (time-step or iteration) in that monitor's *Trigger* property as shown in the following screenshot.



Next, make an entry in the *Collection Start* property, and activate the *Enable collection stop* property.



This adds the **Stop Specification** node to the object tree under the parent field function monitor node. To specify when data collection stops, select this node and set the *Collection stop* property to the desired value in the properties window.



The values you enter in the *Collection stop* property and the parent node's *Collection Start* property will be for a time-step or iteration, depending on what you selected in the parent node's *Trigger* property.

Stop Specification Properties

Use this property to specify when the data collection stops during the simulation run.

Collection stop	Specifies the number of time-steps or iterations to collect data.

Creating Field Function Monitors

To create any new field function monitor, right-click the **Monitors** node and select the type you wish to create from the **New Monitor** pop-up menu.



Deleting Field Function Monitors

Field function monitors can only be deleted if they are not displayed in an XY plot or used to define a stopping criterion.

Renaming Field Function Monitors

Similar to other objects, field function monitors can be renamed within STAR-CCM+.

If you are using anything other than alphanumeric characters to name a field function monitor, you need to use curly brackets (`{}`). For example, if a field function monitor happens to have the name `Mass Flow * 2`, its corresponding field function can be interpreted by STAR-CCM+ to multiply the mass flow by 2 in the calculations. Therefore the field function monitor would need to be renamed `{MassFlow} * 2`. These brackets are also necessary for names that include characters from non-Western alphabets.

Copying Field Function Monitors

When a field function monitor is copied within a simulation, its properties are copied.

If you copy this object between simulation files, some or all of the properties may get copied, depending on the model selections of the destination simulation.

Properties of the Monitors Manager

These properties allow you to control how the monitor data appears in the output.

Monitors to Print	Selects the monitors to print in the output using a standard selection dialog. By default all residual and report monitors in the simulation are selected.
Output Direction	Selects a horizontal or vertical direction for displaying the output.
Heading Print Frequency	Exposes the number of iterations to print between the header lines in the console output. The default value is 10. Only entries greater than zero are valid. You can enter a very large number to get the header printed only once, at the beginning.

Pop-Up Menu of the Monitors Manager

These are the pop-up menu items available when you right-click the **Monitors** manager node.

New Monitor	Creates a new field function monitor of the type you select in the submenu.	
	Field CoVariance	Creates a field co-variance monitor.
	Field Max	Creates a field max monitor.
	Field Mean	Creates a field mean monitor.
	Field Min	Creates a field min monitor.
	Field Sum	Creates a field sum monitor.
	Field Variance	Creates a field variance monitor.
New Group	Facilitates the grouping of objects by creating a new group node within the manager node.	
Group By	Allows you to group the objects according to any of the properties of the objects in the manager node, listed in a submenu of this menu item.	
Ungroup	Reverses the grouping process.	

Plotting Results

The **Plots** node, which has a pop-up menu, is the manager object for all monitor plots and XY plots that have been created in the simulation.

It exists even when empty to allow creation of the first plot, but will typically contain the **Residuals** plot, which is created automatically once you start iterating.

The manager object can contain any number of plots up to the computer's memory limits and the plots are saved in the database when the simulation is saved.

The plotting features in STAR-CCM+ allow you to create three kinds of two-dimensional plots:

- Monitor plots that can display data from monitors, data set functions and/or tables. STAR-CCM+ uses two varieties of monitor plots:
  - Monitor plots are user-created plots.
  - The residual plot is generated and modified automatically by the system.
- XY plots that display solution data from the simulation and/or tables
- Histogram plots that display solution data from the simulation and/or tables

As with many other simulation objects, plots can be deleted, renamed and copied, and can be extensively customized. If the plot is open in the *Graphics* window while you are customizing properties, the relevant object on the plot display will change as you enter the new value. This is a useful way to become familiar with the various properties. All plots, along with any customization, are saved in the simulation file. You can also customize how an XY plot or histogram plot is updated as the simulation iterates.

The plot display is highly interactive, having its own toolbar, zoom, pan and pop-up operations as well as being linked to the simulation tree through drag-and-drop operations. These operations include plotting data from a table and plotting data from a data set, which can also be done using the *Properties* window.

The data being plotted can be viewed in a spreadsheet-like dialog or it can be exported to file for manipulation in other programs. The plot can also be printed to file or to a printer.

A sample XY plot produced by STAR-CCM+ is shown below. In this plot, simulation data is shown on the same plot with data from an external file.



Any type of plot can be opened by double-clicking its node in the simulation tree, or by right-clicking that node and selecting **Open**.



Either operation opens the selected plot to display it in a new document in the *Graphics* window, or brings the document to the foreground if the plot is already open.

You can open all the plots of a simulation at once by right-clicking the **Plots** manager node and selecting **Open All Plots**.



Since there is a lot of similarity between monitor plots and XY plots, parts of this reference section cover the nodes that are specific to monitor plots and XY plots before covering all nodes that are common to both.

See Also

Refreshing the Display

Working with Monitor Plots

This section explains the use of plots that display data from monitors.

What Is a Monitor Plot?

Monitor plots display the values collected by report-based monitors as the solution is advanced.

The monitor plot node has properties and a pop-up menu. An example of a monitor plot is shown below.



Monitor plots can be renamed and deleted when they are no longer needed. Data from tables and data set functions may be plotted with the monitored data.

Monitor Plot Properties

Title	The overall title of the plot. This label appears centered at the top of the plot. Its position and

	format attributes cannot currently be modified. To remove the title set this value to blank.
x-Axis Monitor	Selects a monitor for the x-axis. This allows plotting one monitor against another.
x-Axis Monitor Units	Specifies the units that are applied to the x-values of the monitor data series.

Monitor Plot Expert Properties

Antialiasing	Controls appearance of labels to visual quality.	
	Activated	Uses antialiasing (default), which blends the edge of the characters with the background, creating a smoother appearance on screen.
	Deactivated	Leaves antialiasing off. This setting causes edges of characters to remain sharp, allowing for better printing.
Hover Interactions	Controls whether the display responds to mouse hovering.	
	Activated	Provides interactive responses based on the information under the mouse pointer.
	Deactivated	Leaves the display at its default, with no response to hovering.
Title Font	The font family, font style and font size of the title. Use the customizer to activate a dialog to set this property (changing the text will have no effect).	

	 The property value is displayed using the selected font family, style and size.
Footer	A text box in which you can type a footer for the plot.
Footer Font	The font family, font style and font size of the footer. Use the customizer to activate a dialog to set this property (changing the text will have no effect).  The property value is displayed using the selected font family, style and size.

Monitor Plot Pop-Up Menu

Open	Executes the open plot action to display this plot in the *Graphics* window or make it active if it is already open.
Tabulate	Executes the tabulate action to display the data in a spreadsheet-like dialog. If there is no data in the series, or if there are no series added to the plot, then this dialog will be empty.
Export...	Executes the export plot data action allowing you to save the series data in a .csv file.
Hardcopy...	Saves the plot to a file in one of several formats.
Create Plot Image Annotation	Creates a plot annotation under the Annotation Manager that uses the selected plot.
Apply Plot Options	Sets preferred plot attributes.

Creating a Monitor Plot

To create a plot displaying the collected values of a monitor, right-click on the monitor node and select **Create Plot from Monitor**.



If more than one monitor is selected, then a plot is created based on the monitor that was selected first.

Alternatively, you can right-click on the **Plots** node and select **New Plot > Monitor** and then add the desired monitors to the plot.



See Also

Using Maximum and Minimum Reports

Working with Monitor Data Sets

A monitor data set, represented in the simulation tree by a node with properties and a pop-up menu, defines the source of data and the style of plotting the data.

Monitor Data Set Properties

Legend Name	Changes the legend text in the plot display.
Lock Legend Name	When activated, prevents any automatic changes to the legend name. The legend name is linked to the node name which in turn reflects the name of the monitor that the data series represents. Therefore, name changes to the monitor are automatically reflected in the name of the data series node and the legend name. If this property is activated, the current legend name is locked,

	effectively unlinking the property from the node name. Setting the legend name automatically activates this lock.
Monitor	The name of the monitor that is being used for this series. The monitor linked to this data set cannot be deleted while a link is in place (read-only).
Y-Units	Specifies the units that are applied to the y-values of the data series.

Monitor Data Set Pop-Up Menu

Tabulate	Executes the tabulate data action to view the data in a two-column spreadsheet-style view of the data in this data set. The abscissa (horizontal axis) data will be in the first column and the ordinate (vertical axis) data will be in the second column. If there is no data in the data set, the dialog will be empty.
Export...	Executes the export plot data action allowing you to save the series data in a .csv file.
Apply Plot Options	Sets preferred plot attributes.
Delete	Removes the data from the plot. It does not delete the associated monitor.

Normalizing the Plot

To normalize the values in a monitor plot, use the Normalization Option of the monitor being plotted.

Plotting One Monitor against Another

A monitor plot lets you plot any monitor against any other.

In a default monitor plot, the y-axis uses any one of various monitors while the x-axis represents solution progress. (The x-axis uses the **Iteration** monitor in a steady simulation and the **Physical Time** monitor in an unsteady one.)

However, monitor plots have the added flexibility of plotting monitors against one another with the *x-Axis Monitor* property, which contains a list of all monitors in the simulation. This feature helps you determine correlations between monitors, especially over time. As the simulation run progresses, some monitors define non-monotonic functions: periodic ones such as crank angle, and oscillating ones such as wave height.

Each monitor has an independent data sampling policy based on its *Trigger* and *Frequency* properties. It is not guaranteed that the x- and y-axis monitors sample data at the same rate. Therefore, while the y-axis monitor collects a sample at a given iteration, the plot does not update until the x-axis monitor has also collected a sample.

See the example in the next section.

Note: The *Trigger* property of a monitor does not control the frequency at which the data is plotted. It specifies the frequency at which the monitor collects the data. To control the frequency at which the data is queried and plotted, use the Update feature of the plot node itself. If you want to plot data against physical time, make sure that the same **Physical Time** monitor exists throughout the solution. Otherwise the plot data appears incomplete. If you are running a steady solution but expect to need a plot with physical time, start by running the unsteady solver then switch to the steady one; this ensures that Physical Time exists in the simulation.

Example: The Pitch of a Boat

This example uses an unsteady case from a tutorial simulating the motion of a boat in head waves.

The purpose is to show a correlation analysis between different monitored variables, in this example y-rotation against z-translation:

- Right-click the monitor node (**Z Translation Monitor** in this example) and select **Create Plot from Monitor**.

Note: It is possible to set up a monitor plot using two other techniques: creating an empty monitor plot node and adding the monitors manually, or creating a monitor plot simultaneously with a report-based monitor.

- By default this plot has the following characteristics:
  - The title of the plot has the name of the monitor.

  - The monitor that was used to create this plot is selected for the y-axis.
  - The x-axis is `Physical Time` because this case is unsteady. (In a steady case, the x-axis would be `Iteration`.)
- To reflect the content of the plot, change the *Title* property of the plot node. In this example, the new name is `BoatInHeadWaves (Y Rotation and Z Translation)`.
- To add a monitor to the x-axis of the monitor plot, select the node of that plot and choose the monitor in the drop-down list of the *x-Axis Monitor* property. (In this example it is the `Y Rotation Monitor`.)

The resulting plot is shown in the following screenshot.



Creating a Custom Residual Plot

Since residual monitors are monitor objects, they can be added to any monitor plot in order to create custom residual plots.

Working with XY Plots

This section explains XY plots, which can display scalar values or table data.

What Is an XY Plot?

XY plots are two-dimensional graphs that can be used to display graphs of scalar values from the set of available solution variables or from table data.

The XY plot node has properties and a pop-up menu. An example of an XY plot is shown below.



See Also

Refreshing the Display

XY Plot Properties

You can control some of the characteristics of the XY plot with the properties of the node.

Title	The overall title of the plot. This label appears centered at the top of the plot. Its position and format attributes cannot currently be modified. To remove the title set this value to blank.
Parts	A list of the model and derived parts from which the data are extracted. Editing the property opens the in-place object selector.
Representation	Selects a representation. Additional selections will appear as representations are added to the simulation.

XY Plot Expert Properties

Antialiasing	Controls appearance of labels to visual quality.	
	Activated	Uses antialiasing, which blends the edge of the characters with the background, creating a smoother appearance on screen.
	Deactivated	Leaves antialiasing off. This setting causes edges of characters to remain sharp, allowing for better printing.
Hover Interactions	Controls whether the display responds to mouse hovering.	
	Activated	Causes the display to respond to hovering.
	Deactivated	Leaves the display at its default, with no response to hovering.
Title Font	The font family, font style and font size. Use the customizer to activate a dialog to set this property (changing the text will have no effect).	

	 The property value is displayed using the selected font family, style and size.
Footer	A text box in which you can type a footer for the plot.
Footer Font	The font family, font style and font size of the footer. Use the customizer to activate a dialog to set this property (changing the text will have no effect).  The property value is displayed using the selected font family, style and size.

XY Plot Pop-Up Menu

Open	Executes the open plot action to display this plot in the *Graphics* window or make it active if it is already open.
Tabulate...	Executes the tabulate action to display the data in a spreadsheet-like dialog. If there is no data in the series, or if there are no series added to the plot, then this dialog will be empty.
Export...	Executes the export plot data action allowing you to save the series data in a .csv file.
Refresh	Forces the plot display to refresh, overriding the update policy that is defined on the plot.
Hardcopy...	Saves the plot to a file in one of several formats.
Create Plot Image Annotation	Creates a plot annotation under the Annotation Manager that uses the selected plot.
Apply Plot Options	Sets preferred plot attributes.

Creating an XY Plot

To create an XY plot, right-click on the **Plots** node and select **New Plot > X-Y**.

The XY plot will initially be empty, so you will need to add solution variables and/or table data to be displayed.

Plotting a Solution Variable

The **X Type** and **Y Type** nodes let you display a solution variable in an XY plot.

These nodes function in a similar fashion in that they share a property. As the labels indicate, the **X Type** node sets the type of the abscissa (horizontal axis), while the **Y Type** node sets the type of the ordinate (vertical axis). Both these nodes contain a sub-node **Vector Quantity** or **Scalar Function** that maintains properties about the units and plotted variables.



Other differences between the **X Type** and **Y Type** nodes are as follows:

- The **Y Type** node contains the *Smooth* property to define where data is evaluated for its datasets. If this property is activated, data is evaluated at vertices; otherwise the data is evaluated at cells.
- The **Y Type** node contains dataset sub-nodes that correspond directly to the parts that are defined in the properties of the parent plot node. The dataset sub-nodes are similar to the ones in the **Data Series** sub-node of the monitor plot, and maintain the plotting style information.

In addition, the **Y Type** node is one of possibly multiple Y-type nodes within the **Y Types** manager node. The label of the **Y Type** node includes a number, and one node is created by default with the XY plot, having the label **Y Type 1**.

There are three steps to plotting a solution variable in an XY plot:

- Select a part or parts on which to plot the data.
- Select an axis and a scalar function to display along that axis.
- Specify a direction vector or another scalar function for the other axis.

It is also possible to have a plot with multiple Y types.

Once the solution variable or variables to plot have been specified, the appearance of the plot can be customized. If you wish to compare the solution variable values to reference data, you may plot table data in addition to the solution variable. If an XY plot is open while the simulation is iterating, the plot display updates in accordance to the policy defined by its Update sub-node. See Trigger Types.

X Type and Y Type Properties

Type	Specifies what data will be used for this axis. Typically, one axis will be *Scalar* and the other *Direction* or *Scalar*. Although *Direction* is permitted on both axes, it makes no real sense to do this.	
	Direction	The axis data corresponds to point/cell position vectors dotted with the normalized direction vector specified in the **Vector Quantity** sub-node.
	Scalar	The axis data corresponds to a scalar field function at the point/cell defined in the **Scalar Function** sub-node.
Smooth Values	Defines where the data is evaluated on the associated part (Y Type only).	
	Activated	Interpolates the variables to the vertices of the parts.
	Deactivated	Evaluates the variables at the cells of the parts.

Selecting Parts

XY plots can display data from model parts, such as boundaries, or derived parts, such as line and point probes.

A part data set (or internal data set) node, which has its own properties and pop-up menu, defines the source of data and the style of plotting the data.



For example, if you have reference data at a set of locations, you might create point probes at the corresponding coordinates in your simulation and add the point probe parts to your XY plot.

To add an existing part to an XY plot, select the node of that XY plot and click the *Parts* property. This will bring up the object selector dialog, which allows you to add parts to the plot.

A part can also be copied into the current plot in two ways using drag-and-drop. The first is to drag the part node into the plot window. The second is to drag the part node to the XY plot node, which requires navigating the tree.

Part Data Set Properties

Legend Name	Changes the legend text in the plot display.
Lock Legend Name	When activated, prevents any automatic changes to the legend name. The legend name is linked to the node name which in turn reflects the name of the part that the data series represents. Therefore, name changes to the part are automatically reflected in the name of the data series node and the legend name. If this property is activated, the current legend name is locked, effectively unlinking the property from the node name. Setting the legend name automatically activates this lock.

Part	The associated part (read-only).
Y Type	The associated Y type object (read-only).

Expert Properties

Sort Plot Data	Determines whether the data will be sorted--this ensures that the abscissa (horizontal axis) values run in an increasing order. If you are simply creating a scatter plot, then this will not affect the result, but if you draw a line chart, and the source data is not sorted, this could affect how the data is plotted.	
	Activated	Sort the data
	Deactivated	Do not sort the data
X Offset	A scalar offset that will cause all values of the x-axis data to be shifted by this amount. The units of the offset are implicitly equivalent to those associated with the x-axis data.	
Y Offset	A scalar offset that will cause all values of the y-axis data to be shifted by this amount. The units of the offset are implicitly equivalent to those associated with the y-axis data.	
X Scale Factor	A scaling factor that will cause all values of the x-axis data to be scaled by this amount.	
Y Scale Factor	A scaling factor that will cause all values of the y-axis data to be scaled by this amount.	

Part Data Set Pop-Up Menu

Tabulate	Executes the tabulate data action to view the data in a two-column spreadsheet-style view of the data in this series. The abscissa (horizontal axis) data will be in the first column and the ordinate (vertical axis) data will be in the second column. If there is no data in the series, the dialog will be blank.
Export...	Executes the export plot data action allowing you to save the series data in a .csv file.
Apply Plot Options	Sets preferred plot attributes.
Delete	Removes the data from the plot. It does not delete the associated part.

Selecting a Scalar Function on an Axis

Each XY plot node contains objects within it to specify what should be plotted along the X and Y axes of the graph, respectively.

To select a solution variable to be plotted along one of the axes, open the corresponding node (either **X Type** or **Y Types** > **Y Type**) and select **Scalar** for the *Type* property. This will activate the **Scalar Function** node, which has its own properties.



This node defines the data for the axis of this node's parent to come from a scalar field function. It is only visible if the *Type* property of the parent node is set to **Scalar**.

To specify the solution variable, select the **Scalar Function** sub-node and use the *Field Function* property. Select units with the *Units* property.

Scalar Function Node Properties

Field Function	The specific scalar field function to be plotted as the data value for this axis.

Units	Defines the units that are applied to the selected function.

Specifying a Direction Vector on an Axis

Each XY plot node contains **X Type** and **Y Type** sub-nodes to specify what should be plotted along the X and Y axes of the graph, respectively. To plot position information along one of the axes, select the corresponding node (either **X Type** or **Y Type**) and select **Direction** for the *Type* property.

This activates the **Vector Quantity** node, which has its own properties.



This node defines the data for the axis of this node's parent to come from a direction vector. It is only visible if the *Type* property of the parent node is set to **Direction**.

Since simulation parts may exist in three-dimensional space while each axis of an XY plot represents only a single dimension, the three-dimensional coordinates of the part being plotted must be reduced to one. This objective is accomplished by specifying a direction vector to which the position of sample locations is projected.



To specify the direction in STAR-CCM+, select the **Vector Quantity** node under the **X Type** or **Y Type** nodes. For example, in order to plot X coordinates of the selected parts, you would specify a direction vector of [1,0,0].

Vector Quantity Node Properties

Value	A Cartesian vector that specifies the orientation (in the space of the simulation) of the coordinate axis.

Adding Multiple Y Types to a Plot

It is possible to plot multiple scalars in a single plot by using a separate Y type object for each scalar.

The **Y Types** manager node, which has its own pop-up menu, contains all of the report nodes in the simulation tree.



By default a single **Y Type** node exists within the **Y Types** manager node. To add an extra **Y Type** node, right-click the **Y Types** manager node in the XY plot node and select **New**.

Select the scalar for each Y type, then initialize the solution. The Y types will be plotted in the plot display.

By default the y-axis title concatenates the different field functions that are plotted, corresponding to the different Y types.

Y Types Pop-Up Menu

New	Creates a new **Y Type** node.

Changing the Variable Plotting Style

You can make XY plots appear smoother by using interpolation.

To do this, open the **Plots** node, the node of the particular XY plot, and the **Y Types** node. Select the specific **Y Type** node.

Activate the *Smooth Values* property in the *Properties* window.

This produces a smoother contour of the plot.

The smooth plot is generated using vertex values rather than cell centroid-based values and appears smoother than the default style. However, for most scalar fields vertex values are interpolated from cell values, and thus data extraction with the smooth-filled option is more expensive in terms of computing power.

Working with Histogram Plots

This section explains what histograms are and how to analyze data with them.

What Is a Histogram Plot?

Histogram plots are two-dimensional bar charts that can be used to display data distribution.

The data population is divided into discrete categories, called “bins”. The frequency of occurrences in each bin is given on the vertical scale. The **Histogram Plot** node has properties and a pop-up menu. An example of a histogram plot is shown below.



Histogram Plot Properties

Title	The overall title of the plot. This label appears centered at the top of the plot. Its position and

	format attributes cannot currently be modified. To remove the title set this value to blank.	
Parts	A list of the model and derived parts from which the data are extracted. Editing the property opens the object selector dialog.	
Representation	Selects a representation. Additional selections will appear as representations are added to the simulation.	
Aggregate Parts	Controls how the histogram plot works with parts.	
	Activated	Treats parts as a single entity.
	Deactivated	Works with each part as a separate entity (default).

Histogram Plot Expert Properties

Antialiasing	Controls appearance of labels to visual quality.	
	Activated	Uses antialiasing, which blends the edge of the characters with the background, creating a smoother appearance on screen.
	Deactivated	Leaves antialiasing off. This setting causes edges of characters to remain sharp, allowing for better printing.
Hover Interactions	Controls whether the display responds to mouse hovering.	
	Activated	Causes the display to respond to hovering.
	Deactivated	Leaves the display at its default, with no response to hovering.
Title Font	The font family, font style and font size. Use the customizer to activate a dialog to set this	

	property (changing the text will have no effect). The property value is displayed using the selected font family, style and size.
Footer	A text box in which you can type a footer for the plot.
Footer Font	The font family, font style and font size of the footer. Use the customizer to activate a dialog to set this property (changing the text will have no effect). The property value is displayed using the selected font family, style and size.

Histogram Plot Pop-Up Menu

Open	Executes the open plot action to display this plot in the *Graphics* window or make it active if it is already open.
Tabulate	Executes the tabulate action to display the data in a spreadsheet-like dialog. If there is no data in the series, or if there are no series added to the plot, then this dialog will be empty.
Export...	Executes the export plot data action allowing you to save the series data in a .csv file.
Refresh	Forces the plot display to refresh, overriding the update policy that is defined on the plot.
Hardcopy...	Saves the plot to a file in one of several formats.
Create Plot Image Annotation	Creates a plot annotation that uses the selected plot.
Apply Plot Options	Sets preferred plot attributes.

Creating a New Histogram Plot

To create a new histogram plot, right-click on the **Plots** node and select **New Plot > Histogram**.

The histogram plot will initially be empty, so you will need to add data to be displayed by accessing the plot node's *Parts* property.

Setting up a Histogram

Use the contents of the **Histogram** manager node, which resides in the **Histogram Plot** node, to specify the details of how your data is to be displayed.

This node has properties and is subject to common menu items.



In particular, you can set the bin, and optionally control the weighting and the range.

Setting the Bin of the Histogram

Although the number of bars in the histogram plot is set in the properties of the parent node, the **Bin** node contains additional controls in its properties.

This object is also subject to common menu items.



Bin Properties

Field Function	The specific scalar field function to be plotted as the data value for this histogram.
Units	Selects the type of units to use.

Setting the Weighting

When the *Weighting* property of the parent node is set to `User Specified`, the **Weighting** node becomes available so that you can control the weighting using this object's properties.

This object is also subject to common menu items.



Weighting Properties

Field Function	The specific scalar field function to be weighted.
Units	Selects the type of units to use.

Setting the Range

When the *Range Mode* property of the parent node is set to `Manual`, the **Range** node becomes available so that you can specify the range using this object's properties.

This object is also subject to common menu items.



Range Properties

Range	The minimum and maximum values to be plotted. Use the property customizer button (...) to activate a dialog to set this property.

Histogram Node Properties

Weighting	Controls whether to customize the weighting or leave it at the default.	
	Default	No weighting.
	User specified	Lets you specify the weighting function in the properties of the Weighting node, which is added to the object tree when this setting is selected.
Bins	Sets the number of intervals along the x-axis.	
Smooth Values	Selects a technique for calculating plot values.	
	Activated	Interpolates values to mesh vertices.
	Deactivated	Uses values in the closest cell centroid.

Range Mode	Determines whether or not to automatically set the range of values to plot.	
	Manual	Uses the values set in the Range node, which is added to the object tree when this setting is selected.
	Auto	Automatically determines the range from the data.

Updating Plots

You can control how open XY, monitor, and histogram plots update by changing the trigger type that is located in the **Update** node.

This node, which has its own properties and is subject to common menu items, customizes how plots in the display are updated as the simulation iterates.



The **Update** node has two sub-nodes:

- A trigger node to control the frequency of updates. The name of this node reflects the choice of trigger type. See Trigger Types and Trigger Properties.
- A **Hardcopy** node to control the specifications of the images that are generated by the plot updates. See Update and Hardcopy Properties.

For more information on selecting the trigger type and defining the update frequency, refer to Setting the Trigger for an Update.

Note that:

- The update mechanism is only active whilst the simulation is iterating.
- Multiple Y types within a XY plot displaying different scalar quantities will update simultaneously whenever an **Update Frequency** setting triggers an update. For example, the scalar

quantities displayed in a Y types are updated to represent the new solution. Closed XY plots are not updated until they are opened.

- Hardcopies will only be written to disk at the intervals dictated by the **Update Frequency** setting for a particular plot.

The plot may also be refreshed manually. Expand the **Plots** node and right-click on the chosen plot. Choose **Refresh** from the pop-up menu.

You can also generate an image of the plot after each iteration or time-step to compile an animation of the plot. This is different from the single image generated when you save a hardcopy of the plot.

If all clients are disconnected from a running server, no plots will be saved to files.

Update and Hardcopy Properties

Update Properties

File Format	The graphics format that is used to save images of the visualization display. You have the same options as provided in the save plot to file action.
Enabled	Uses a checkbox with which you can activate or deactivate this feature.
Save to File	Determines whether or not multiple images of the plot display are saved as the solution progresses. This can be used to create an animation showing how the data changes while the simulation runs. STAR-CCM+ will bring the selected plot to the front when it produces a screenshot of that frame.
Output Directory	Path to where the files specified in the *Base Filename* property will be saved.
Base Filename	The prefix of the file name to use when saving plots to file. The full file name will be `prefixNNNNN.ext` where `prefix` is this property, `NNNNN` is a number to indicate the sequence of the files and `ext` is an extension based on the graphics format. If you leave the text box blank, STAR-CCM+ will restore the name that was there previously once that text box no longer has the focus.

Trigger	Determines how the display should be updated to reflect changes in the simulation. Change the trigger type to control the method that updates the scene. See Trigger Types.

Hardcopy Properties

The properties of the **Hardcopy** node let you control the specifications of the images generated by the plot updates. However, if all clients are disconnected from a running server, no graphics plots will be saved to files.

Output Magnification	Produces a larger image than the image in the original display, based on the factor you enter.
Offscreen Width	Sets the width of the image of the plot display when the save plot to file action is executed.
Offscreen Height	Sets the height of the image of the plot display when the save plot to file action is executed.

Deleting Plots

Monitor, XY, and histogram plots can be deleted when no longer needed. The residual plot, which is created automatically, cannot be deleted.

Deleting a plot does not delete the underlying objects that are associated with the datasets of the plot. For instance, a monitor plot deletion will not impact the underlying monitor, which will continue to collect data.

Renaming Plots

Plots may be renamed as desired. Renaming the plot node will not change the displayed title of the plot.

Copying Plots

When a plot is copied, its properties and sub-nodes are copied. You can also copy the plot's information separately between simulation files, depending on whether the destination simulation has the same dependent objects.

For example, if you want to copy a plot that uses a file table, you need to follow these steps:

- Make sure that the destination simulation contains a table with the same presentation name.
- Copy and paste the plot.

Customizing Plots

The appearance of all plots can be modified in a number of ways.

See Also

Plot Display

Changing the Title of the Plot

The title of a plot may be modified by changing the *Title* property.

The title text font may also be changed in the expert properties of the plot node. Use the property customizer button (...) to activate a dialog to set this property (changing the text will have no effect).

Changing the Data Display Style

The style and color of data can be modified with the line style, symbol style, and fill style nodes.

Each style is associated with the source of the data: monitors, parts, or tables. Expand the **Plots node** and expand the chosen plot.

To reveal the **Line Style** and **Symbol Style** nodes for a monitor plot, expand the **Data Series** node and expand the chosen monitor.

To reveal the **Line Style** and **Symbol Style** nodes for an XY plot, do either of the following:

- Expand the **Y Types** node, the chosen Y Type node, and expand the chosen part.
- Expand the **Data Series** node, and expand the chosen data set node.

To reveal the **Line Style** and **Fill Style** nodes for a histogram plot, do either of the following:

- Expand the **Histogram** manager node and the histogram data set node.
- Expand the **Data Series** node, and expand the chosen data set node.

Modifying the Line Style

The **Line Style** node, through its properties, lets you control the appearance of the line of a particular set of data.

The appearance of the line can be customized in several ways:

- The style of the line can be changed (none, solid, dashed).
- The color of the line can be changed.
- The width of the line can be changed.

To modify the color or style of the line click on the **Line Style** node. To change the color of the line either enter the RGB value of the color under the *Color* property in the *Properties* window or open the color selector by clicking on the ellipses button. To change the style of the line choose from the list of styles in the drop-down menu under the **Style** property in the *Properties* window. To change the width of the line modify the *Width* property.

Line Style Properties

Style	Specifies the line style.	
	None	No line displayed.
	(**Solid**)	Displays a solid line.
	(**Dash**)	Displays a dashed line.
	(**Dot**)	Displays a dotted line.
	(**Dash-Dot-Dash**)	Displays a line that alternates between dashes and dots.
		Displays a line that alternates between

	(**Dash-Dash-Dot**)	two dashes and one dot.
Width	Controls the line width in pixels.	
Color	Determines the color used to draw the line. The color can be input directly in RGB format (e.g., red would be entered as `255,0,0`). A comma must separate the numeric values. The property customizer executes the define color action.	

Modifying the Symbol Style

The **Symbol Style** node, through its properties, lets you control the appearance of the symbol of the points in a particular set of data.

The symbols can be customized in several ways:

- The shape of the symbol can be changed.
- The color of the symbol can be changed.
- The size of the symbol can be changed.
- The width of the stroke of the symbol can be changed.
- The spacing between symbols can be changed.

To modify the color or the shape of the symbol click on the **Symbol Style** node. To change the color of the symbol either enter the RGB value of the color under the *Color* property in the *Properties* window or open the color selector by clicking on the ellipses button. To change the shape of the symbol select from the drop-down menu under the *Shape* property in the *Properties* window. The size of the symbol can be changed by modifying the *Size* property.

Symbol Style Properties

Shape	The shape of the symbol.	
	None	Makes the symbol invisible.
	Horizontal Line	Displays the symbol as a short horizontal line.
	Vertical Line	Displays the symbol as a short vertical line.

	Cross	Displays the symbol as a cross.
	Plus	Displays the symbol as a plus sign.
	Star	Displays the symbol as an asterisk type symbol (as opposed to a star shape).
	Filled Square	Displays the symbol as a filled square.
	Empty Square	Displays the symbol as a square with the outline only.
	Filled Circle	Displays the symbol as a filled circle.
	Empty Circle	Displays the symbol as a circle with the outline only.
	Filled Triangle	Displays the symbol as a filled triangle.
	Filled Diamond	Displays the symbol as a filled diamond.
Size	Sets the size of the symbol in pixels.	
Color	The color of the symbol. The color can be input directly in RGB format (e.g., red would be entered as `255,0,0`). A comma must separate the numeric values. The property customizer executes the define color action.	
Width	Sets the width of the stroke of the symbol in pixels.	
Spacing	Sets the interval between symbols.	

Customizing the Appearance of the Histogram

The histogram's appearance, as with other types of plots, is controlled in part by the Axes and Legend objects.

However, other characteristics are affected by the **Fill Style** and **Line Style** nodes, which are under the data node.



Fill Style Properties

Color	Determines the color of the bars. The color can be input directly in RGB format (e.g., red would be entered as `255,0,0`). A comma must separate the numeric values. The property customizer executes the define color action.
Secondary Color	Controls the color of the background. The property customizer executes the define color action.
Pattern	Selects a pattern with which to fill the bars.

For detailed information about the **Line Style** node, see Modifying the Line Style.

Customizing the Legend

The **Legend** node, with its properties, lets you change various aspects of the legend, including the position, orientation and visibility.

By default a legend is generated with the data point symbols and the names of the sources of the data whether that be parts, monitors or external data. The legend can be customized in several ways:

- It can be moved to different positions around the plot.
- The format, font and orientation of the text can be changed. (The text itself can be changed in the display using other objects.)

- It can be removed.

Expand the **Plots** node and expand the chosen plot node. Click on the **Legend** node. The legend location can be customized by selecting an option in the drop-down menu under the *Position* property in the *Properties* window. The legend may be located outside of the plotted area at eight points of the compass, or overlaid on the plot at a **Custom** location. To remove the legend deactivate the checkbox of the *Visible* property.

The orientation of the legend can either be horizontal or vertical and can be changed by selecting from the drop-down menu under the *Layout* property. It is also possible to control the number of columns or rows which affects the layout of the legend.

The format of the legend text can also be modified in the *Expert* properties window. The spacing around legend entries can be customized by modifying the *Item Spacing* property.

The legend text font may also be changed in the *Font* expert property.

Use the property customizer button ( ) to activate a dialog to set this property (changing the text will have no effect).

Renaming the Plot Legend for a Data Set

The legend text that is associated with a dataset can be modified using the properties of that dataset. This may be useful if the default text happens to be fairly long.

Select the node of the dataset and enter the new text in the *Legend Name* property. Note that the *Lock Legend Name* property will automatically become activated.

Your changes will appear in the plot display.

Legend Properties

Visible	Controls visibility of the legend.

	Activated	Shows the legend.
	Deactivated	Hides the legend.
Position	Controls the position of the legend relative to the chart area:	
	North	Top center
	North East	Top right
	East	Middle right
	South East	Lower right
	South	Lower center
	South West	Lower left
	West	Middle left
	North West	Top left
	Custom	Puts the legend into a floating box that you can reposition anywhere in the display by clicking and dragging.
Layout	Controls the position of the legend objects relative to each other. This feature requires more than one data set to be represented in the plot display.	
	Horizontal	Displays the objects in a row.
	Vertical	Displays the objects in a column.

Expert Properties

Font	The font family, font style and font size. Use the property customizer button ( ) to activate a dialog to set this property (changing the text will have no effect).

	 The property value is displayed using the selected font family, style and size.
Item Spacing	Controls the spacing around items in the legend.
Number Of Rows	Specifies the number of rows for the legend.
Number Of Columns	Specifies the number of columns for the legend.
x-Position	Relative x-position on the screen, defined over the interval [0,1].
y-Position	Relative y-position on the screen, defined over the interval [0,1].

Customizing Axes

The **Axes** node contains several sub-nodes that let you customize the axes of the plot.



Customization is available in a variety of features.

Changing the Axis Orientation

The orientation of the axes determines the position and direction of increasing values for each axis.

The default is the axis origin to be based in the bottom left-hand corner of the plot with the x-axis along the base of the plot. To change the orientation of the axes expand the **Plots** node, expand the chosen plot node, and click the **Axes** node. Click on the property customizer button (...) under the *Axes Orientation* property in the *Properties* window which opens the *Axis Orientation* dialog. Click on the required orientation and click **OK** to apply the change.



Axes Properties

Axis Orientation	Executes the define axis orientation action to change the position and direction of increasing values of the two axes. Typing into the property editor will not change the setting -- click on the property customizer button (...) to access the dialog.

Setting the Scale

Each axis can be customized to have either a linear or logarithmic scale, or can be removed from the plot. It is also possible to focus on a region of the plot by customizing the minimum and maximum axis values.

From the **Axes** node click on either the **X Axis** or **Y Axis** node.

To change from the linear to logarithmic scale activate the checkbox of the *Logarithmic* property. To remove the axes from the plot deactivate the checkbox of the *Visible* property.

X Axis and Y Axis Properties

Note: Zero and negative values are not defined for logarithms, and the plot skips data at such points.

Visible	Determines whether the axis is visible. This property affects the sub-nodes--title, ticks and

	grid. The sub-nodes can be individually made visible or invisible.	
	Activated	Makes the title, ticks and grid visible.
	Deactivated	Hides the title, ticks and grid.
Logarithmic	Determines whether the axis is logarithmic or linear.	
	Activated	Makes the axis logarithmic.
	Deactivated	Makes the axis linear.

Expert Properties

Minimum	The minimum value shown on the axis.
Lock Minimum	When activated, locks the axis' minimum to the defined value.
Maximum	The maximum value shown on the axis.
Lock Maximum	When activated, locks the axis' maximum to the defined value.

Customizing the Title on the Axis

It is also possible to change the title of the axes by:

- Changing the text of the title
- Changing the font of the title
- Changing the orientation of the text
- Removing the axes title from the plot

To change the title of an axis:

- Open the **Axes** node.
- Open either the **X Axis** or **Y Axis** node.
- Select the **Title** node.
- Modify the *Title* property.

The title can be taken out of the plot altogether by clearing the checkbox of the *Visible* property. The title may be rotated through 90 degree

intervals by selecting the rotation angle from the drop-down menu under the *Rotation* property.

The axis title text font may also be changed in the *Font* expert property.

Use the property customizer button ( ) to activate a dialog to set this property (changing the text will have no effect).

You can make this font change for multiple plots at the same time using the table editor dialog (in multi-edit mode).

Title Properties

Visible	Controls visibility of the title.	
	Activated	Shows the title.
	Deactivated	Hides the title.
Title	The text of the axis title. An empty string will hide the title. No additional space is gained.	
Lock Title Name	Controls whether to lock the title to the specified value.	
	Activated	Locks the specified title.
	Deactivated	An automated naming convention builds an axis title based on the datasets that are plotted.
Rotation	Selects an angle of rotation for the title.	
	0 degrees	Leaves the title at its normal position.
	90 degrees	Rotates the title 90 degrees clockwise.
	180 degrees	Rotates the title 180 degrees (upside down).
	270 degrees	Rotates the title 270 degrees clockwise (= 90 degrees counterclockwise).

Expert Properties

Font	The font family, font style and font size. Use the property customizer button ( ) to activate a dialog to set this property (changing the text will have no effect).  The property value is displayed using the selected font family, style and size.

Customizing the Major Labels on the Axis

The labels, ticks, and grid associated with each axis can be customized.

Options for customization include:

- Changing the spacing of the ticks to determine how many total ticks are displayed
- Changing the spacing of the labels to avoid over-crowding the axes
- Independent customization of grid lines associated with labeled (major) and non-labeled (minor) tick marks

To access the **Major Labels** node, open either the **X Axis** or **Y Axis** node.

Major Labels Properties

Visible	Determines whether the labels are visible.	
	Activated	Makes the labels visible.
	Deactivated	Hides the labels.

Grid Visible	Determines whether the background grid lines of the labels are visible.	
	Activated	Makes the background grid visible.
	Deactivated	Hides the background grid.

Expert Properties

Font	The font family, font style and font size. Use the property customizer button ( ) to activate a dialog to set this property (changing the text will have no effect).  The property value is displayed using the selected font family, style and size.	
Precision	Controls the maximum number of digits that appear after the decimal place in the axis labels.	
Auto Precision	Determines whether or not to automatically determine the axis label precision.	
	Activated	Automatically sets the precision.
	Deactivated	Uses the value set in the *Precision* property.

Pad Decimals	Determines whether or not to pad labels with zeroes to fill the precision of the labels.	
	Activated	Applies the label precision to annotations.
	Deactivated	Leaves annotations separate.
Spacing	Sets the spacing between axis labels in the units of the axis. For example, with a *Minimum* of `0.0` and a *Maximum* of `8.0`, a Label Spacing of `4.0` will place labels at 0.0, 4.0 and 8.0. Setting this value to `0.0` hides the tick labels.	
Auto Spacing	Determines whether or not to set the spacing automatically between axis labels.	
	Activated	Automatically sets the spacing.
	Deactivated	Uses the value set in *Label Spacing*.
Grid Style	Specifies the grid line style.	
	None	No line displayed.
	(**Solid**)	Displays a solid line.
	(**Dash**)	Displays a dashed line.
	(**Dot**)	Displays a dotted line.
	(**Dash-Dot-Dash**)	Displays a line that alternates between dashes and dots.
		Displays a line that alternates between

	(**Dash-Dash-Dot**)	two dashes and one dot.
Grid Color	Determines the color used to draw the grid line. The color can be input directly in RGB format (e.g., red would be entered as `255,0,0`). A comma must separate the numeric values. The property customizer executes the define color action.	
Grid Width	Controls the grid line width in pixels.	

Customizing the Minor Ticks on the Axis

The style and color of the minor ticks can be modified, or turned off altogether, using the properties of the **Minor Ticks** node.

The lines associated with each axis are modified independently of each other.

To access the **Minor Ticks** node, open either the **X Axis** or **Y Axis** node.

In the *Properties* window, you can change the color, style and width of the line. If you do not wish to see the minor ticks then deactivate the checkbox of the *Visible* property. If you do not wish to see the grid lines then deactivate the checkbox of the *Grid Visible* property.

Minor Ticks Properties

Visible	Determines whether the minor ticks are visible.	
	Activated	Makes the minor ticks visible.
	Deactivated	Hides the minor ticks.
Grid Visible	Determines whether the background grid lines of the minor ticks are visible.	
	Activated	Makes the background grid visible.
	Deactivated	Hides the background grid.

Expert Properties

Count	Sets the number of minor ticks between the major labels.

Grid Style	Specifies the grid line style.	
	None	No line displayed.
	(**Solid**)	Displays a solid line.
	(**Dash**)	Displays a dashed line.
	(**Dot**)	Displays a dotted line.
	(**Dash-Dot-Dash**)	Displays a line that alternates between dashes and dots.
	(**Dash-Dash-Dot**)	Displays a line that alternates between two dashes and one dot.
Grid Color	Determines the color used to draw the grid line. The color can be input directly in RGB format (e.g., red would be entered as `255,0,0`). A comma must separate the numeric values. The property customizer executes the define color action.	
Grid Width	Controls the grid line width in pixels.	

Hiding the Axes

To hide either the X axis or Y axis, select its corresponding node and deactivate the checkbox of the *Visible* property in the *Properties* window.

Setting Preferred Attributes for Plots

You can create a preferred set of attributes for the plots in a simulation and apply them easily.

- To create a preferred set of plot attributes:
  - From the menu bar, select **Tools** > **Options**.
  - In the *Options* dialog, select the **STAR-CCM+** > **Plotting** node.
  - In the *Properties* window, specify attributes such as title fonts and plot line thickness.
- To apply your preferred attributes to a plot, right-click either on the plot node or in the plot display and select **Apply Plot Options**.
- To apply the attributes to multiple plots at once, multi-select the plot nodes, right-click, and select **Apply Plot Options**.
- To have a plot show more recent changes in plotting properties in the *Options* dialog, use the **Apply Plot Options** action again.

Note: STAR-CCM+ retains the attributes that you specify in the *Options* dialog, even in subsequent simulations. However, plots -- whether they are created before or after you set the attributes -- do not show those attributes until you apply them with the right-click action.

Adding and Removing Data in a Plot

You can make various types of simulation data available for plotting: monitors, tabular data, and derived data.

The **Data Series** manager node contains the objects representing the data sets that are used in the plot.



1. To add a data set, right-click on the **Data Series** node and select **Add Data**.
2. In the object selection dialog that appears, make your selections and click **OK**.

Note: The objects that appear in the object selector dialog are of all available data objects in the simulation, not necessarily the ones that are used by the plot.

XY and histogram plots can receive tabular data sets. Monitor plots, however, can receive monitors, tabular data sets, and derived data sets.

3. To customize the appearance of the data in the plot, See Changing the Data Display Style.

You can add multiple data sets at once with the drag-and-drop technique.

4. Multi-select the nodes in the object tree and drag the group onto the Data Series node.

You can also remove any data sets from a plot.

5. To remove a data set, right-click on its node and select **Delete**.

What Is a Tabular Data Set?

A tabular data set defines the source of data to be plotted and the style of plotting the data.

A tabular data set represents the contents of a table made available to a plot. This can be useful for plotting the data alone or against other data for comparison. However, you must have read a table into STAR-CCM+ beforehand. Once this is done, you can plot more than one set of table data in a plot, provided all the tables have identical column names. Simply follow the same steps for each table.

Tabular Data Set Properties

Legend Name	Changes the legend text in the plot display.
Lock Legend Name	When activated, prevents any automatic changes to the legend name.
Table	Selects a table from among the ones represented by nodes in the **Tables** node of the **Tools** node.
Columns	Identifies the columns in the table file.
X-Column	Selects a column in the table to be plotted on the x-axis.
X-Units	Units of the plotted x-values.
Y-Column	Selects a column in the table to be plotted on the y-axis.
Y-Units	Units of the plotted y-values.

Expert Properties

Sort Plot Data	Determines whether the data will be sorted--this ensures that the abscissa (horizontal axis) values run in an increasing order. If you are simply creating a scatter plot, then this will not affect the result, but if you draw a line chart, and the source data is not sorted, this could affect how the data is plotted.	
	Activated	Sorts the data.
	Deactivated	Does not sort the data.

X Offset	A scalar offset that will cause all values of the x-axis data to be shifted by this amount.
Y Offset	A scalar offset that will cause all values of the y-axis data to be shifted by this amount.
X Scale Factor	A scaling factor that will cause all values of the x-axis data to be scaled by this amount.
Y Scale Factor	A scaling factor that will cause all values of the y-axis data to be scaled by this amount.

Tabular Data Set Pop-Up Menu

Swap Columns	Switches the axis to which either column of data is assigned.
Tabulate	Executes the tabulate data action to view the data in a two-column spreadsheet-style view of the data in this series. The abscissa (horizontal axis) data will be in the first column and the ordinate (vertical axis) data will be in the second column. If there is no data in the series, the dialog will be blank.
Export...	Executes the export plot data action allowing you to save the series data in a .csv file.
Delete	Executes the remove data set action to remove this data set from the plot. It does not delete the associated file.

Swapping Data Columns between Axes

You can control which column of data is represented on a particular axis by using the object tree.

Open the **Plots** node, the individual plot (monitor or XY) node and the **Tabular** node. Right-click the node of the tabular object and select **Swap Columns**.

Assuming you have the plot display open, the effect will appear immediately.

Your next step involves changing the titles of the axes.

What Is a Derived Data Set?

A derived data set represents a monitor plot's access to the contents of a derived monitor data set or derived tabular data set within a data set function, such as a Fast Fourier Transform.

Its node in the object tree has its own properties and pop-up menu.



Derived Data Set Properties

Legend Name	Changes the legend text in the plot display.
Lock Legend Name	When activated, prevents any automatic changes to the legend name.
X-Units	Units of the plotted x-values.
Y-Units	Units of the plotted y-values.
Input Data	The monitor or tabular data set function that is the source of this object's data. This read-only value reflects the selection in the **Data Series** manager node. The data set function linked to this data set cannot be deleted while a link is in place.
Percentage RMS Variation	Percentage variation in the sum of RMS from the original signal (read-only).

Derived Data Set Pop-Up Menu

Tabulate	Executes the tabulate data action to view the data in a two-column spreadsheet-style view of the data in this data set. The abscissa (horizontal axis) data will be in the first column and the ordinate (vertical axis) data will be in the second column. If there is no data in the data set, the dialog will be empty.
Export...	Executes the export plot data action allowing you to save the series data in a `.csv` file.
Apply Plot Options	Sets preferred plot attributes.
Delete	Executes the remove data series action to delete the data from the plot. It does not delete the associated data set function.

Tabulating Data

To display the data from a plot in tabular form, right-click on the plot node and select **Tabulate...**

The data from the plot will be presented in the *Tabular Data* dialog, displaying the X and Y data values for one or more data series in a spreadsheet-like view.



If this action is triggered from a plot, all data series in the plot will be displayed in pairs of X and Y value columns (as shown above). If triggered from a dataset, only the data contained by that object is shown.

Features of this dialog include:

- Horizontal and vertical scroll bars
- Column and row resizing by dragging when the mouse cursor changes to an east-west icon or a north-south icon over the column and row dividers
- Hiding columns, showing columns and jumping to a particular row number by right-clicking anywhere on the spreadsheet and selecting from the pop-up menu



- Copying a single value or set of selected values from the tabulated data to the clipboard -- one use for this is to copy a value to another object in the simulation tree. Specifically, select the particular value in the table. Right-click and select **Copy**.



Paste the value into the relevant constant in the simulation tree.

Exporting Data

To export data from one of the curves displayed in a plot, right-click on the dataset node corresponding to the data that you wish to export and select **Export...**

You can export all of a plot's data at once with the same type of operation by right-clicking on the node of that plot.

This action uses an operating system-specific standard *Save* dialog, via the save as action, to store the data points of the selected monitor or data set to file. The exported file is written as comma separated values (`.csv`)--a format supported by many spreadsheets.

The headers of the `.csv` file contain useful information for each column of data:

- For plot data, the header includes the scalar and the part name:

```
"PressureRatio-plane section 1"
```

- For monitor data, the header includes the monitor name and scalar:

```
"Heat Transfer 1 Monitor: Heat Transfer of Boundary
Heat Transfer (W)"
```

Printing and Saving Plot Images

Plot images can be saved as image files or printed.

To send a plot image to a printer connected to your system, select the plot node, then use the **File > Print** menu item to invoke the print action.

Saving a Plot Image File

To save a plot image to a file, right-click on the plot node and select **Hardcopy...** as shown below**.**

A customized, operating system-specific *Save To File* dialog will be presented, allowing you to write the selected plot(s) to file. The saved image(s) can then be incorporated into a printed report, a slide presentation, a Web page or another commonly used medium.



It is possible to set specific values for the image resolution, in pixels, by clearing the **Use Screen Resolution** checkbox and entering values for the **X Resolution** and **Y Resolution**.

The image file can be written in any of the following formats:

- Joint Photographic Experts Group (`.jpg`)
- PostScript (`.ps`)
- Portable document format (`.pdf`)
- Extended PostScript (`.eps`)
- Portable network graphics (`.png`)

This action is disabled if more than one plot is selected.

Plot images can also be saved in batch mode. The procedure is similar to that for saving scene images in batch.

Plot File Formats Reference

This section describes external file formats for plotting.

plt Format

This text-based file format can hold multiple data series against the same abscissa (horizontal axis) data in an array format.

There are four types of lines:

- A definition line containing four space separated fields:
  - The keyword ARRAY
  - The plot title (in single quotes)
  - A count of how many series are stored
  - The number of points per series
- The abscissa data as space separated fields holding the axis title followed by the number of data points specified in the first line (field 4)
- Multiple ordinate data lines (as specified in field 3 of the first line), each as space separated fields holding the series title followed by the number of data points specified in the first line (field 4)
- Optional comment lines beginning with the '#' character

An example of a file is shown below.

```
# The data is in array format

# It has one set of x-axis data, multiple sets of y-axis
data

# It can have names for the data view and for the x and y
series

# Each x point can have a point label

ARRAY 'Reference Plot' 1 3

'' -0.035750001668930054 -0.046149998903274536
-0.04582050070166588

'Lower Temperature' 371.477294921875 327.20318603515625
327.66650390625
```

xy Format

This format is a simple two column file with the abscissa (horizontal axis) data in the first column and the ordinate (vertical axis) data in the second column, with the columns separated by one or more spaces.

There should only be one data pair per line. An example of this file is shown below.

```
-4.035 0.8604

-3.386 0.8584

-2.487 0.84544

-2.017 0.81856

-1.695 0.78999

-0.937 0.67503

-0.761 0.64874

-0.455 0.5848

-0.288 0.56893

-0.11 0.56395

0 0.55086
```

Pop-Up Menu of the Plots Manager

These are the pop-up menu items available when you right-click the **Plots** manager node.

New Plot	Offers a choice of plot type to create.	
	Monitor	Creates a new monitor plot with a default name of `Monitor Plot N`, and adds a new plot tab in the *Graphics* window for that plot. The plot is created with the following defaults: • No data series • An x-axis label of `X Axis` • A y-axis label of `Y Axis` • A title of `Monitor Plot N`
	X-Y	Creates a new XY plot with a default name of `XY Plot N`, and adds a new plot tab to the *Graphics* window for that plot. The plot is created with the following defaults: • No data series • An x-axis label of `Position [1,0,0,] (m)` • A y-axis label of `<Select Function>` • A title of `XY Plot`

	Histogram	Creates a new histogram plot with a default name of `Histogram Plot N`, and adds a new plot tab to the *Graphics* window for that plot. The plot is created with the following defaults: • No data series • An x-axis label of `<Select Function>` • A y-axis label of `Frequency` • A title of `Histogram Plot`
Open All Plots	Displays all plots in the *Graphics* window.	
New Group	Facilitates the grouping of objects by creating a new group node within the manager node.	
Group By	Allows you to group the objects according to any of the properties of the objects in the manager node, listed in a submenu of this menu item.	
Ungroup	Reverses the grouping process.	

Triggers and Update Events

Several features of STAR-CCM+ require that you define a trigger for when an action takes place. STAR-CCM+ provides three built-in trigger types, and one custom trigger type.

The custom trigger type is defined by an update event.

Features for which you set triggers are as follows:

- Solution Histories—control how often the solution writes to the simulation history file.
- Scenes—control how often the scene updates with the calculated solution.
- Plots—control how often the plot updates with the calculated solution.
- Tables—control how often the table updates with the calculated solution
- Auto Exports—control how often the solution data exports to external files.
- Auto Saves—control how often the simulation file saves.
- Monitors—control how often the monitor extracts new data from the calculated solution.

Setting the Trigger for an Update

In general, features that require you to define a trigger expose the **Update** node. Use the properties of this node, and any child nodes, to define when an update action occurs.

To set the trigger type:

1. Navigate to the **Update** node (this example shows the **Update** node in a scene).



2. In the *Properties* window, select the trigger type you want from the *Trigger* property. See Trigger Types.



Note: Changing a monitor trigger during (or after) a simulation run will not result in the clearing of the previously collected data.

The child node changes according to the trigger type.

3. Select the child node.



4. Set the trigger properties. See Trigger Properties.



Trigger Types

There are four methods of triggering an update event:

Trigger Type	Description
Iteration	Triggers an update every n iterations. Specify the update frequency in the *Number of Iterations* property field.
Time Step	Triggers an update every n time-steps, where the *Number of Time Steps* property specifies the update frequency. This trigger type is only appropriate for transient simulations.
Delta Time	Triggers an update when the simulation time advances by the amount you specify in the *Delta Time* property field. This trigger type is only appropriate for transient simulations.
Update Event	Triggers an update based on a user-defined update event. User-defined update events are based on monitor values and provide you with more flexibility in controlling the updates. See User-Defined Update Events.

Note: For each trigger type, apart from **Update Event**, you can control the start, end, and update interval in the properties panel of the child node. See Trigger Properties. For **Update Event**, you can specify these settings in the relevant update event, located under the **Tools** > **Update Event** node. See User-Defined Update Events.

Trigger Properties

Triggers control when an update event occurs.

The properties for the various trigger types are listed below:

[Trigger] Frequency	Specifies the frequency of the update event. (Built-in trigger types).
Start [Trigger]	Specifies from what point the update event starts updating. (Built-in trigger types).
Enable Stop	Specifies whether the *Stop [Trigger]* property is active or not. (Built-in trigger types).
Stop [Trigger]	Specifies from what point the update event stops updating. (Built-in trigger types).
Update Event	Specifies which user-defined update event to use for the selected object. This property only appears for the **Update Event** trigger.

User-Defined Update Events

User-defined update events are custom triggers that provide you with more flexibility compared to the built-in trigger types. With a user-defined update event, you can trigger an update based on monitor values.

You create and manage user-defined update events under the **Tools** > **Update Events** node.



When the conditions for a user-defined update event are satisfied, the event triggers an update and the icon next to the corresponding update event flashes green.



Creating User-Defined Update Events

You can create user-defined update events using various update event parameters.

To create a user-defined update event:

1. Right-click **Tools** > **Update Events** and select **New Event** > **[Update_Event_Type]**.



A new event node appears.



2. Set up the update event parameters. Refer to the individual update event documentation for more information:
   - Logic Event
   - Monitor Asymptote
   - Monitor Delta
   - Monitor Frequency
   - Monitor Range
   - Monitor Frequency Range
   - Monitor Waypoint
3. In the feature of your choice, navigate to the Update node, and set *Trigger* to **Update Event**. In this example, the **Update** node of a scalar scene is selected.



4. Select **Update** > **Event Frequency** and set *Update Event* to the update event you created, in this case **Logic Event**.



Logic Event

The logic event allows you to apply logic arguments to multiple update events. Three logical expressions are available: AND, OR, and XOR.

Logic Event Properties

Enabled	Specifies whether the update event is active or not.	
Satisfied	Indicates whether the update event has been satisfied. This property is read only.	
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.	
Logic Operator	Specifies which logic argument to use for the event.	
	AND	All of the events that you specify must be satisfied for the logic to be true.
	OR	Any one event must be satisfied for the logic to be true. If all events are satisfied, the logic is also true.
	XOR	Any odd number of events must be satisfied for the logic to be true.

Creating a Logic Event

Logic events require multiple update events to function.

In the following example, a simple OR logic event is set up to trigger when one or all of the following update events is satisfied:

- the simulation time is 0.5 seconds
- at every ten iterations
- the user-defined report value (pressure) is greater than or equal to 10 Pascal.

To set up the logic event:

1. Right-click the **Tools > Update Events** node and select **Logic Event**.



A logic event is added to the **Update Events** node.

2. Select the logic event node and set *Logic Operator* to **Or**.



The icon next to the logic event changes to reflect the new logic.



3. Right-click **Logic Event** > **Update Events** and select **New Event** > **Monitor Delta**.



A monitor delta event is added to the logic event node.

4. Select the **Monitor Delta** node and set:
   - *Monitor* to **Physical Time**
   - *Delta Threshold* to `0.5 s`



5. Using the same technique, create a **Monitor Frequency** event and set:

   - *Monitor* to **Iteration**
   - *Sample Frequency* to `10`



6. Using the same technique, create a **Monitor Range** and set:
   - *Monitor* to [**Pressure Report**]
   - *Range Value* to `10 Pa`



The logic event is ready to be applied.

7. In the feature of your choice, navigate to the Update node, and set *Trigger* to **Update Event**. In this example, the **Update** node of a scalar scene is selected.



8. Select **Update** > **Event Frequency** and set *Update Event* to the update event you created, in this case **Logic Event**.



Monitor Asymptote

When the calculations reach a specified threshold, the monitor asymptote event triggers an update. This threshold is passed when the

maximum difference among a set of samples drops below a user-specified value.

The set is a user-specified number of the most recent samples.

This feature is useful when the exact trigger value is unknown, but you do know the maximum change in the monitored value that you consider sufficient for an update.

Monitor Asymptote Properties

Enabled	Specifies whether the update event is active or not.	
Satisfied	Indicates whether the update event has been satisfied. This property is read only.	
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.	
Monitor	Specifies the monitor for this update event.	
Delta Threshold	The threshold value that triggers the update event.	
Sample Count	The number of most recent samples that apply to the asymptotic calculation.	
Inner Samples	Controls whether only the iterations from the current time-step are observed. Leave this option deactivated if your monitored data transcends time-steps.	
	Activated	Disregards the values that are collected during the previous time-steps.
	Deactivated	Considers the values that are collected during the previous time-steps.

Monitor Delta

The monitor delta event observes the change in values that are collected by a monitor.

For example, observing the physical time monitor results in the event being updated every `n` seconds.

Monitor Delta Properties

Enabled	Specifies whether the update event is active or not.	
Satisfied	Indicates whether the update event has been satisfied. This property is read only.	
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.	
Monitor	Specifies which monitor to use for this update event.	
Difference Mode	Specifies how the difference (or delta) is calculated.	
	Accumulated	The difference is calculated as the accumulation of differences between each consecutive pair of collected monitor values. When the delta is exceeded, the difference between the accumulated delta and the threshold is carried over to the next accumulation.
	Non- Accumulated	The difference is calculated between the current monitor value and the value at the last satisfied iteration. This is an absolute difference, ignoring the intermediate values. When the delta is exceeded, the last satisfied value is updated, resulting in a new comparison point from which the next delta is evaluated.
Delta Threshold	The threshold value that triggers the update event.	

Monitor Frequency

The monitor frequency event observes the number of samples collected by a monitor.

For example, observing the physical time monitor results in the event being updated when the monitor collects a number of samples that is divisible by `f`, where `f` is a frequency count.

Monitor Frequency Properties

Enabled	Specifies whether the update event is active or not.
Satisfied	Indicates whether the update event has been satisfied. This property is read only.
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.
Monitor	Specifies which monitor to use for this update event.
Sample Frequency	Specifies the number of samples that trigger the event.
Start Count	Specifies the number of samples that are required before the event triggers.

Monitor Range

The monitor range event updates when the monitor value is within the specified range.

There are two types of ranges:

- Greater or equal to the specified monitor value.
- Less than or equal to the specified monitor value.

Monitor Range Properties

Enabled	Specifies whether the update event is active or not.
Satisfied	Indicates whether the update event has been satisfied. This property is read only.
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.

Monitor	Specifies which monitor to use for this update event.	
Range Operation	Specifies the range for the update event.	
	>=	The event is satisfied when the monitor value is greater or equal to the `Range Value`.
	<=	The event is satisfied when the monitor value is less than or equal to the `Range Value`.
Range Value	Specifies the value that is used in the *Range Operation*.	

Monitor Frequency Range

The monitor frequency range event triggers an update every time that the number of samples collected by a monitor are within a specified range.

For example, observing the physical time monitor results in the event triggering an update every time-step until the number of time-steps exceeds the limit set by the frequency range.

Monitor Frequency Range Properties

Enabled	Specifies whether the update event is active or not.	
Satisfied	Indicates whether the update event has been satisfied. This property is read only.	
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.	
Monitor	Specifies which monitor to use for this update event.	
Range Operation	Specifies the comparison operator made in determining the range of satisfiable values.	
	>=	The event triggers an update when the monitor collects `n` or

		more samples, where `n` is the `Range Value`.
	<=	The event triggers an update when the monitor collects `n` samples or less, where `n` is the `Range Value`.
Range Value	Specifies the value that is used in the *Range Operation*.	
Start Count	Specifies the number of samples that are required before the event triggers.	

Monitor Waypoint

The monitor waypoint event assumes that the monitor is sampling a continuous function over time, which could be periodic.

The waypoint event triggers an update when the values from two consecutive samples cross the `Waypoint Value`. That is, for two consecutive samples, `S1` and `S2`, the monitor waypoint event trigger when:

- `S1` is less than the `Waypoint Value` and `S2` is greater than the `Waypoint Value`.

or

- `S1` is greater than the `Waypoint Value` and `S2` is less than the `Waypoint Value`.

Monitor Waypoint Properties

Enabled	Specifies whether the update event is active or not.
Satisfied	Indicates whether the update event has been satisfied. This property is read only.
Iteration	Displays the iteration at which the event was last evaluated. This property is read only.
Monitor	Specifies which monitor to use for this update event.
Waypoint Value	Specifies at which monitor value the event updates.

Using Data Set Functions

For signal processing in particular, any signal can be represented either in the time or the frequency domain.

These representations are equivalent, that is, $S[T] \equiv S[F]$. Fourier transforms are used to transform a signal from the time domain to the frequency domain, and inverse Fourier transforms are used to do the reverse. Transforming a signal to the frequency domain makes it possible to identify the dominant modes in the signal.

Data set functions are used to apply signal processing to sets of data, which can come from STAR-CCM+ simulations or experiments. The **Data Set Functions** node, which has a pop-up menu, is the manager object for working with data set functions in the simulation tree. The data set functions manager can hold any number of data set functions up to the memory limits of the computer.



Signal processing functions can be used in two contexts in STAR-CCM+:

- They can be applied to data stored in monitors, such as would be obtained from a report of a particular value at a point. The results of this analysis of data at a point can be visualized in monitor plots.

- They can be applied to data stored on selected surfaces at specified time intervals in the `.trn` file. The results of this analysis of data on a surface can be visualized in scenes.

One example of using data set functions for frequency analysis would be a study in which the goal was to reduce the noise that a room fan generates. The initial part of the study would involve a validation exercise, in which you would take experimental measurements of pressure against time at a location set some distance from the fan. These measurements would be used in a frequency analysis to identify the dominant modes for this signal and the frequency bands containing these modes. Large modes within the range of frequency audible to human hearing would suggest a noisy environment. Following the experimental analysis, a transient simulation would be run for the same configuration, with a monitor being set at the same location as the pressure probe used in the experiment. Frequency analysis would be applied to the monitored data to obtain the dominant modes predicted by the simulation. Comparing the predicted modes to the measured modes gives an understanding of the uncertainty that is associated with the simulation. If uncertainty is within acceptable levels, further simulations could be performed for varying configurations of the fan. These simulations would investigate ways of reducing the dominant modes, or at least moving them out of the range of human hearing.

This section explains the following:

Creating Data Set Functions

To add a data set function, right-click the **Data Set Functions** node and select a type from the **New** submenu of the pop-up menu.



The node of the new object appears in the simulation tree.


Servers
intake
simulation
Tools
Annotations
Colormaps
Coordinate Systems
Data Set Functions
C(p) 1
Imported
Tabular
Monitor
Field Functions
Layouts
Material Databases
Motions
Reference Frames
Tables
Transforms

What Are Point Time Histories?

Point Time Histories are time histories at a point in the input signal which are then processed using a Fourier transform.

Point Time Histories can also be used to apply filters to signals.



The **h(p)** node (with a numerical suffix) represents this type of data set function in the object tree. This node has properties and is subject to common menu items.



Point Time History Properties

Window Function	Specifies the window function that you require for this data set function. These options are all standard options for this type of analysis.

	• **None** - This default setting leaves the data set function without a window function until you select one of the other options in this list. • **Hamming** • **Hann** (Hanning) • **Bartlett** • **Blackman**	
Update Interval	Sets the number of time-steps after which to update the time history.	
Filter Type	Specifies the type of filter that is applied.	
	None	Does not use a filter.
	Band pass filter	Keeps frequencies between low and high cut-off levels. Adds a FIR Filter sub-node in which you can set these levels.
	Band stop filter	Discards frequencies between low and high cut-off levels. Adds a FIR Filter sub-node in which you can set these levels.
Start Time	The start time for the time sampling of this object.	
Cut-off time	The end time for the time sampling for this object. Clicking the property customizer button in the right-hand column activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min** or **s**). If you have multiple data set functions, you can vary the settings of this property to sample multiple time intervals from the simulation run.	

What Is the FIR (Finite Impulse Response) Filter?

Sometimes, using the Fourier transform with its Ideal Filter can slow the calculations since this filter has infinite length in time, while the time length of a time history is finite. A solution can be reached faster by applying a FIR (Finite Impulse Response) filter to the time history in the time domain.

The FIR filter is constructed from a Fourier series in 2N+1 coefficients. The first and last N points of the filtered signal are discarded.

The **FIR Filter** node, which has properties and is subject to common menu items, represents this feature. It is made available when the *Filter Type* property of the time history node is set to either `Band pass filter` or `Band stop filter`.



The same type of sub-node can be activated in the Time History:Pressure field function.



FIR Filter Properties

Filter Coefficients	(N), 2N+1 coefficients are used to construct the filter. N points are discarded from the start/end of filtered data.
Filtered Signal Start Time	The start of valid filtered data.
Filtered Signal Cut-off Time	The end of valid filtered data. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify the appropriate time units (**hr**, **min**, or **s**)
Low Cut-off Frequency	Retains frequencies above f_1 (see Eqn. (168)) when the band pass filter is used, and discards them when the band stop filter is used.
High Cut-off Frequency	Retains frequencies below f_2 (see Eqn. (169)) when the band pass filter is used, and discards them when the band stop filter is used. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify both the desired frequency and the appropriate time units (**/s** or **Hz**)

What Are Point Time Fourier Transforms?

Point Time Fourier Transforms are Fourier transforms for one time point in the input signal.



Choosing this type of data set function creates a **G(p)** node with a numerical suffix in the object tree. This node has properties and is subject to common menu items.



The time Fourier Transform (FT) (see Eqn. (125)) lets you transform time signals that are stored in STAR-CCM+ to the frequency domain so that you can identify dominant modes in the signals.

Point Time Fourier Transform Properties

Update Interval	Sets the number of time-steps after which to update the FT.	
Filter Type	Specifies which type of filter is applied. Digital filters allow you to filter out parts of a signal which correspond to certain frequencies.	
	None	Does not use a filter.
	Band pass filter	Keeps frequencies between low and high cut-off levels. Adds an Ideal Filter sub-node in which you can set these levels.
	Band stop filter	Discards frequencies between low and high cut-off levels. Adds an Ideal Filter sub-node in which you can set these levels.
Frequency Function	Specifies the frequency function that is generated. These functions are all standard functions for this type of analysis.	
	Frequency	Frequency in Hz.
	Octave Bands	Discrete frequency bands corresponding to different octaves within hearing range. Each successive band represents twice the range of the previous one.
	3rd Octave Bands	Discrete frequency bands corresponding to one-third of each octave within hearing range.
	Fourier Mode	The index in the Fourier calculations.
	Strouhal Number	The non-dimensional version of the

		`Frequency` option. Selecting this option adds a Strouhal Number sub-node to the FT node.
	Custom Bands	Lets you define a set of custom frequency bands by adding a Custom Bands sub-node to the FT node.
Start Time	The start time of the time sampling for this object.	
Cut-off time	The end time of the time sampling for this object. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min**, or **s**). If you have multiple FTs, you can vary the settings of this property to sample multiple time intervals from the simulation run.	
Amplitude Function	Specifies the amplitude function that is generated. These functions are all standard functions for this type of analysis.	
	Power Spectral Density	Frequency-domain description of the signal power.
	Sound Pressure Level	Sound pressure in decibels. Adds the Reference Sound Pressure node.
	Root Mean Square	The square root of the `Power Spectral Density`.
	A-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the A-scale function, for loudness less than 55 phons (at 1000 Hz, 1 phon = 1 dB). This

		function is typically for problems that involve the frequency range for human hearing. Adds the Reference Sound Pressure node.
	B-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the B-scale function, for loudness of 55-85 phons. Adds the Reference Sound Pressure node.
	C-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the C-scale function, for loudness greater than 85 phons. Adds the Reference Sound Pressure node.
	Re{FFT Coefficients}	The real parts of the FFT coefficients.
	Im{FFT Coefficients}	The imaginary parts of the FFT coefficients.
	Phase{FFT Coefficients}	The phases of the FFT coefficients ($[Im / \mathrm{Re}]$).
	Abs{FFT Coefficients}	The absolute values of the FFT coefficients ($\sqrt{\lvert \mathrm{Re} \rvert^2 + \lvert Im \rvert^2}$).
Analysis Blocks	Specifies the number of analysis blocks N_b (see Equation) for this data set function to use.	
Overlap Factor	Sets the overlap factor a (see Equation) between blocks, if any, in a range from `0` to `0.9`. This feature allows for smooth transitions between blocks. This property can be set with a slider.	
Window Function	Specifies the window function that you require for this FT.	

	Fourier transforms assume a periodic signal; to enforce this condition apply a window function. Window functions smear the dominant modes in the frequency domain. If the signal is broken into several analysis blocks, the window function is applied to each block. There are a number of commonly used window functions. All taper a signal at the beginning and/or end. • **None** - This default setting leaves the data set function without a window function until you select one of the other options in this list. • **Hamming** • **Hann** (Hanning) • **Bartlett** • **Blackman**

What Is the Ideal Filter?

The **Ideal Filter** node represents this feature.

This node has properties and is subject to common menu items. It is made available when the *Filter Type* property of the Fourier transform is set to either `Band pass filter` or `Band stop filter`.



The same type of sub-node can be activated in the Fourier Transform field function.



Ideal Filter Properties

Low Cut-off Frequency	Retains frequencies above f_1 (see Eqn. (168)) when the band pass filter is used, and discards them when the band stop filter is used.
High Cut-off Frequency	Retains frequencies below f_2 (see Eqn. (169)) when the band pass filter is used, and discards them when the band stop filter is used. Clicking the ... property customizer button in the right-hand column activates a property

	customizer that lets you specify both the desired frequency and the appropriate time units (**/s** or **Hz**)

What Are Custom Bands?

The **Custom Bands** object appears in the node of the Fourier Transform when the *Frequency Function* property is set to `Custom Bands`.



By clicking the ... property customizer button of the *Custom Bands* property of this object, you can activate a dialog that makes the settings more accessible.



What Is the Strouhal Number?

The **Strouhal Number** object, which has properties, appears in the node of the Fourier Transform when the *Frequency Function* property is set to `Strouhal Number`.



Strouhal Number Properties

These properties let you adjust the Strouhal number settings, expressed as $l_{ref} \bullet p / u_{ref}$.

Reference Length	The reference length l_{ref}, which can be specified with a text entry. The default is 1.0.
Reference Velocity	The reference velocity u_{ref}, which can be specified with a text entry. The default is 1.0. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify both the desired value and the appropriate units.

What is Reference Sound Pressure?

Reference sound pressure is a constant, used to calculate the Sound Pressure Level.

It specifies the reference ambient pressure p_{ref} that is used when calculating A-, B- or C-Weighted Sound Pressure Level.

The default value of this constant depends on the medium:

- 2.0E-5 Pa for air
- 1.0E-6 Pa for water

Using Point Time Fourier Transforms

You typically run the Point Time Fourier Transform on time-history data that the simulation generates.

Therefore, before you start running the simulation:

1. Identify the variables for which time-history data is required (for example, pressure reported at a point probe).
2. Activate the transient time model.
3. Set up reports, monitors, and monitor plots that use the variables of interest.

The monitors control the trigger and frequency of the calculations. You can create plots that are based on those monitors alone, or based on the FTs that you add. The latter option is possible through the derived data set objects of a monitor plot.

4. Create Fourier Transforms and set their properties:
   a) Right-click the **Tools > Data Set Function**s manager node and choose the required transform from the pop-up menu.
   b) Select from among sets of standard options for amplitude, frequency, and window functions.
   c) Break the signals into analysis blocks, working with each block separately. An overlap factor can be applied to analysis blocks for smooth transitions.

You can create any number of FTs with different combinations of settings, including different sampling time intervals. To track your combinations of settings, simply rename each FT node according to its settings.



Run the simulation and examine the results.

If applicable, validate the simulation results with your own experimental data, which the FT can access through its derived tabular data set.

What Are Point-to-Point Time Fourier Transforms?

You use Point-to-Point Time Fourier transforms to determine the correlation between two input signals. This data set function applies a Fourier transform to each of the input signals, then performs a convolution on both resulting transforms.



Choosing this type of data set function creates the **G(p,q)** node (with a numerical suffix) in the object tree. This function has double the number of properties of a 1-point transform because a set of properties is defined for each point. The corresponding node has properties and is subject to common menu items



Point-to-Point Time Fourier Transform Properties

Update Interval	Sets the number of time-steps after which to update the Fourier transform.	
Filter Type	Specifies the type of filter that is applied. Digital filters allow you to filter out parts of a signal which correspond to certain frequencies.	
	None	Does not use a filter.
	Band pass filter	Keeps frequencies between low and high cut-off levels. Adds an Ideal Filter sub-node in which you can set these levels.
	Band stop filter	Discards frequencies between low and high cut-off levels. Adds an Ideal Filter sub-node in which you can set these levels.
Frequency Function	Specifies the frequency function that is generated. All these functions are standard functions for this type of analysis.	

	Frequency	Frequency in Hz.
	Octave Bands	Discrete frequency bands corresponding to different octaves within hearing range. Each successive band represents twice the range of the previous one.
	3rd Octave Bands	Discrete frequency bands corresponding to one-third of each octave within hearing range.
	Fourier Mode	The index in the Fourier calculations.
	Strouhal Number	The non-dimensional version of the `Frequency` option. Selecting this option adds a Strouhal Number sub-node to the FT node.
	Custom Bands	Lets you define a set of custom frequency bands by adding a Custom Bands sub-node to the FT node.
Start Time	The start time of the time sampling for this object.	
Cut-off time	The end time of the time sampling for this object. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min**, or **s**). If you have multiple FTs, you can vary the settings of this property to sample multiple time intervals from the simulation run.	
Amplitude Function	Specifies the amplitude function that is generated. All these functions are standard functions for this type of analysis.	

	Cross Spectral Density	The power spectral density of the cross-spectrum signal, $G(x_1, x_2; \omega)$ (see Eqn. (136)). This option is comparable to the auto-spectral density of 1-point Fourier transforms, $G(x; \omega)$ (see Eqn. (135)).
	Cross Root Mean Square	The root mean square of the `Cross Spectral Density`. This option is comparable to the RMS of 1-point Fourier transforms. See Frequency-Space Parseval's Relation.
	Coherence	γ in Eqn. (142).A measure of the correspondence between the magnitude of two signals at different frequencies. Its range is 0 - 1.
	Phase	θ in Eqn. (141). A measure of the offset (phase difference) between two signals. Its range is -p to +p.
Analysis Blocks	Specifies the number of analysis blocks N_b (see Equation) for this data set function to use.	
Overlap Factor	Sets the overlap factor a (see Equation) between blocks, if any, in a range from `0` to `0.9`. This property allows for smooth transitions between blocks.	
Window Function	Specifies the window function that you require for this FT.	

	Fourier transforms assume a periodic signal; to enforce this condition apply a window function. Window functions smear the dominant modes in the frequency domain. If the signal is broken into several analysis blocks, the window function is applied to each block. There are a number of commonly used window functions. All taper a signal at the beginning and/or end. • **None** - This default setting leaves the data set function without a window function until you select one of the other options in this list. • **Hamming** • **Hann** (Hanning) • **Bartlett** • **Blackman**

What Are Point-to-Point Time Correlations?

You use Point-To-Point Time correlations to determine the phase differences between the dominant modes of two input signals.

Point-to-Point Time Correlations are time cross-correlations at any two points x_1 and x_2 in space (see Time Cross-Correlations for the formulation)..



Choosing this type of data set function creates the **R(p,q)** node (with a numerical suffix) in the object tree.



Point-to-Point Time Correlation Properties

Amplitude Function	Specifies the function that is used to represent the cross-correlation. • **Correlation Value** • **Correlation Magnitude** • **Correlation Coefficient** - The magnitude is scaled to be in the range -1 to +1.	
Update Interval	Sets the number of time-steps after which to update the time history.	
Filter Type	Specifies the type of filter that is applied.	
	Band pass filter	Keeps frequencies between low and high cut-off levels. Adds a FIR Filter sub-node in which you can set these levels.
	Band stop filter	Discards frequencies between low and high cut-off levels. Adds a FIR Filter sub-node in

		which you can set these levels.
Start Time	The start time for the time sampling of this object.	
Cut-off time	The end time for the time sampling for this object. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min**, or **s**). If you have multiple data set functions, you can vary the settings of this property to sample multiple time intervals from the simulation run.	

What Are Point Inverse Time Fourier Transforms?

This inverse transform allows data to be transformed from the frequency to the time domain.

This transform is given by Eqn. (127)



The **h'(G(p))** node (with a numerical suffix) represents this type of data set function in the object tree. This node has properties and is subject to common menu items.



Point Inverse Time Fourier Transform Properties

Amplitude Function	Selects the method by which an Inverse Fourier Transform can be visualized as its magnitude or using its real or imaginary parts.

	The transform can be attached to both auto- and cross-correlation spectra.	
	Real	The inverse of auto-correlation spectra has only a real part
	Imaginary	The inverse of cross-correlation spectra has both real and imaginary parts
	Magnitude	The magnitude of the spectra.
Added Mean Type	Selects the method that is used to generate the mean value added to the inverse Fourier transform.	
	Calculated	Generated by STAR-CCM+.
	User specified	Entered manually. A User Specified Added Mean node is added to the object tree.
Signal Repeats	Selects the number of times the signal is repeated (without introducing discontinuities). It is set to zero by default.	

What Is the User Specified Added Mean?

The **User Specified Added Mean** node represents this feature.

This node has properties and is subject to common menu items. It is made available when the *Added Mean Type* property of the point inverse time Fourier transform is set to **User specified**.



User Specified Added Mean Properties

Added Mean	Specifies the mean value that is added to the inverse Fourier transform. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify the added mean.

Exporting Audio

You can export the time signal that the point inverse time Fourier transform generates as an audio file using the **Export Audio...** action on the derived data node pop-up menu.

STAR-CCM+ currently supports the Windows Media Player format (`.wav`) for exported audio files.

The **Audio export control** node allows you to scale all audio exports from a point inverse time Fourier transform to the same maximum amplitude. This feature gives similar volume levels across all the signals.



Audio Export Control Properties

Normalization	Specifies whether normalization of exported audio data is required.	
	No normalization	Normalization is not required.
	User specified normalization	Activates normalization of all audio data to the same value. This option creates an additional User specified normalization node in the simulation tree.

User Specified Normalization

This node specifies a *Normalization comparison value* that is used to scale all exported audio signals to the same amplitude range.

Set this value to the peak value of the derived data set that has the highest-amplitude signal. The highest amplitude signal can be detected by subtracting the mean value from all of the derived data sets belonging to the point inverse time Fourier transform and comparing the resulting traces on a plot.

To subtract the mean value from the derived data sets:

1. Select option **User Specified** for the Added Mean Type property on the Inverse Fourier Transform node.
2. Enter a value of `0.0` for the User specified added mean.

When you have estimated the peak amplitude across all the derived data sets, enter this value in the *Normalization comparison value* property.

User Specified Normalization Properties

Normalization comparison value	All exported signals are normalized against this value.

What Are Line Time Histories?

A Line Time History is a collection of Point Time Histories that are drawn from the points that make up a line probe.

A line time history can then be processed using a Line Time Fourier Transform.



The **h(l)** node (with a numerical suffix) represents this type of data set function in the object tree. This node has properties and is subject to common menu items.



Line Time History Properties

Start Signal	The start time for the mapped signal.

End Signal	The end time for the time sampling for this object. Clicking the ... property customizer button in the right-hand column activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min** or **s**). If you have multiple data set functions, you can vary the settings of this property to sample multiple time intervals from the simulation run.

Exporting Fields

You can export derived transient field data as a Matlab `.mat` file.

To specify the name and location of the `.mat` file, right-click the function node, select **Export Field...**, and complete the dialog.



h(l1) Node Properties

Part	The line part containing the positions from which *Field Function 1* is extracted.	
Data Surfaces 1	The imported surface on which *Field Function 1* is extracted.	
Field Function 1	The signal is extracted from this scalar field function.	
Update Active	Whether the update is active for the imported history.	
	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

Exportable Field Data

VERSION	Version string for the file.
X, Y, Z	Position components of the point samples in the file.
X_UNITS, Y_UNITS, Z_UNITS	Units of position components for the point samples in the file.
S	Samples for series that is associated with the point.
S_UNIT	Units of samples for series that is associated with each point.
M_TITLES	Titles of each field that is stored in the file.
M_UNITS	Units for each field that is stored in the file.
Field<i>_<j>	The jth sample of ith field for each point sample.

What Are Line Time Fourier Transforms?

Line Time Fourier Transforms let you calculate Fourier transforms at various point along a line probe.

They allow the creation of either auto-spectra (see Eqn. (135)) or cross-spectra (Eqn. (136)) along the line probe.

$$h(\underline{x};t) \rightarrow G(\underline{x};\omega)$$

Choosing this type of data set function creates a **G(l)** node (with a numerical suffix) in the object tree. This node has properties and is subject to common menu items.



Line Time Fourier Transform Properties

Start Signal	The start time for the sampling of this object.
End Signal	The end time for the sampling for this object. Clicking the ... property customizer button activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min**, or **s**). If you have multiple data set functions, you can vary the settings of this property to sample multiple time intervals from the simulation run.

Amplitude Function	Specifies the amplitude function that is generated. These functions are all standard for this type of analysis.	
	Power Spectral Density	Frequency-domain description of the signal power.
	Cross Spectral Density	The power spectral density of the cross-spectrum signal. This option is comparable to the auto-spectral density of 1-point Fourier transforms.
Analysis Blocks	Specifies the number of analysis blocks N_b (see Equation) for this data set function to use.	
Overlap Factor	Sets the overlap factor a (see Equation) between blocks, if any, in a range from `0` to `0.9`. This property allows for smooth transitions between blocks.	
Window Function	Specifies the window function that you require for this FT. Fourier transforms assume a periodic signal; to enforce this condition, apply a window function. Window functions smear the dominant modes in the frequency domain. If the signal is broken into several analysis blocks, the window function is applied to each block. There are a number of commonly used window functions. All taper a signal at the beginning and/or end.	
	None	This default setting leaves the FT without a window function until you select one of the other options in this list.
	Hamming	Uses the Hamming function.
	Hann	Uses the Hann (Hanning) function.

	Bartlett	Uses the Bartlett function.
	Blackman	Uses the Blackman function.

Exporting Fields

You can export derived transient field data as a Matlab `.mat` file.

To specify the name and location of the `.mat` file:

1. Right-click the function node and select **Export Field...**
2. In the *Save* dialog, navigate to a location, enter a file name, and click **Save**.



Multi-Point Imported Properties

Input time history	The line time history **G(In)** to use as input, selected from pull-down menu.
X Units	Displays the calculated units for the x-axis data.
Y Units	Displays the calculated units for the y-axis data.

Update Active	Whether the update is active for the imported line.	
	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

Multi-Point Imported Expert Properties

X Surface Axis Units	The units in which the X axis is specified.	
Y Surface Axis Units	The units in which the Y axis is specified.	
Z Surface Axis Units	The units in which the Z axis is specified.	

Exportable Field Data

Re{FFT Coefficients}	The real parts of the FFT coefficients.
Im{FFT Coefficients}	The imaginary parts of the FFT coefficients.

Using Line Time Fourier Transforms

You can set up a line time Fourier transform for a transient property and visualize the transform.

To set up a line time Fourier transform:

1. Create a line probe in the region of interest.

To minimize inaccuracies due to visualization interpolation, adjust the probe properties such that data is extracted at cell centers.



2. Create a derived line data set using the line probe.
3. Set the properties of the derived line data set and calculate the transform.
4. Set the properties for the transform plots.
5. Create a monitor plot and plot the derived line data set.
6. Visualize the spectra in a scene.

What Are Line Spatial Fourier Transforms?

Line Spatial Fourier Transforms allows the creation of a matrix of energy densities dependent on wavenumber and frequency along a line probe.

These data set functions map signals from the spatial and frequency domains to wavenumber and frequency domains. This type of transform is available only with the `aero-vibro` feature license.

$$G(\underline{x};t) \rightarrow F(\underline{k};\omega)$$

Choosing this type of data set function creates an **F(G(l))** node in the object tree, which has properties and is subject to common menu items.



Line Spatial Fourier Transform Properties

Start Signal	The start length for the length sampling of this object.
End Signal	The end length for the length sampling for this object. Clicking the ... property customizer button activates a property customizer that lets you

	specify both the time value and the appropriate length units. If you have multiple data set functions, you can vary the settings of this property to sample multiple length intervals from the simulation run.	
Padding Factor	Specifies the degree by which the signal is zero-padded. The factor is a proportion of the original signal strength.	
Window Function	Specifies the window function that you require for this FT. Fourier transforms assume a periodic signal; to enforce this condition, apply a window function. Window functions smear the dominant modes in the frequency domain. If the signal is broken into several analysis blocks, the window function is applied to each block. There are a number of commonly used window functions. All taper a signal at the beginning and/or end.	
	None	This default setting leaves the FT without a window function until you select one of the other options in this list.
	Hamming	Uses the Hamming function.
	Hann	Uses the Hann (Hanning) function.
	Bartlett	Uses the Bartlett function.
	Blackman	Uses the Blackman function.

Using Line Spatial Fourier Transforms

To set up a line spatial Fourier transform for a transient property:

1. Create and set up a line time Fourier transform.
2. Create a line derived data set using a derived line data set from a line time Fourier transform.
3. Set the properties of the line derived data set to calculate the transform.
4. Set the properties for the transform plots.
5. Create a monitor plot and plot the derived line data set.
6. Visualize the spectra in a scene.

What Are Surface Time Histories?

A Surface Time History is a collection of Point Time Histories that are drawn from the points that make up a presentation grid.



The **h(s)** node (with a numerical suffix) represents this type of data set function in the object tree. This type of data set uses CAE data mappers.



This node has properties and is subject to common menu items.

Surface Time History Properties

Start Signal	The start time for the mapped signal.
End Signal	The end time for the time sampling for this object. Clicking the property customizer button in the right-hand column activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min** or **s**). If you have multiple data set functions, you can vary the settings of this property to sample multiple time intervals from the simulation run.
Interpolation Method	Interpolation method when mapping data from source to target surface.
Export Data Location	Data location on target surface for mapped data.
Export Frequency	Frequency of time steps for which data is mapped.
Max. Distance To Target	Maximum distance to target for which a source point is projected to the target surface. Otherwise the source point is left unmoved.

Surface Time History Expert Properties

Use Original Mesh	Allows to use the original mesh for mapping.	
	Activated	The original mesh is used.
	Deactivated	The current mesh is used.

The Surface Time History node has three child nodes. The Imported node allows to create derived transient surface data sets, and has no properties. Mapping constraints can be set through the **Proximity Constraint** node properties and **Normal Direction Constraint** node properties.

Exporting Fields

You can export derived transient field data as a Matlab `.mat` file.

The settings for automatic mapping can be entered through the surface time history node properties, whereas the mapping and export process is carried out through the derived transient surface data set node. For more information about the derived transient surface data set node and its properties, see Creating Derived Transient Surface Data Sets.

To export the derived transient field data as a Matlab `.mat` file:

1. Select the surface time history node, **h(s)**, and set its properties to specify the automatic mapping settings.
2. Expand the **h(s)** > **Imported** node.
3. Select the derived transient surface data set node,**[Multi-Point Imported]**, and activate *Update Active* to enable automatic mapping.
4. Set the remaining properties to specify the field function, data surfaces and target surfaces for mapping. You can select more than one data surface, as STAR-CCM+ allows transient mapping from multiple data surfaces.
5. Set the **Proximity Constraint** node properties and **Normal Direction Constraint** node properties.
6. Right-click the **[Multi-Point Imported]** node and select **Export Field...**



7. In the dialog, specify the name and location of the `.mat` file and click **Save**.

Transient surface data map and export operations can also be carried out through the transient imported model node.

What Are Surface Time Fourier Transforms?

Surface Time Fourier Transforms let you calculate Fourier transforms at various point within a presentation grid.

They allow the creation of auto-spectra (Eqn. (136)) over the presentation grid.

$$h(\underline{x};t) \rightarrow G(\underline{x};\omega)$$

Choosing this type of data set function creates a **G(s)** node (with a numerical suffix) in the object tree.



Surface Time Fourier Transform Properties

Start Signal	The start time for the sampling of this object.
End Signal	The end time for the sampling for this object. Clicking the ... property customizer button activates a property customizer that lets you specify both the time value and the appropriate time units (**hr**, **min**, or **s**). If you have multiple data set functions, you can vary the settings of this property to sample multiple time intervals from the simulation run.

Analysis Blocks	Specifies the number of analysis blocks N_b(see Equation) for this data set function to use.	
Overlap Factor	Sets the overlap factor a (see Equation) between blocks, if any, in a range from `0` to `0.9`. This property allows for smooth transitions between blocks.	
Window Function	Specifies the window function that you require for this FT. Fourier transforms assume a periodic signal; to enforce this condition, apply a window function. Window functions smear the dominant modes in the frequency domain. If the signal is broken into several analysis blocks, the window function is applied to each block. There are a number of commonly used window functions. All taper a signal at the beginning and/or end.	
	None	This default setting leaves the FT without a window function until you select one of the other options in this list.
	Hamming	Uses the Hamming function.
	Hann	Uses the Hann (Hanning) function.
	Bartlett	Uses the Bartlett function.
	Blackman	Uses the Blackman function.

Exporting Fields

You can export derived transient field data as a Matlab `.mat` file.

To specify the name and location of the `.mat` file:

1. Right-click the function node and select **Export Field...**

2. In the *Save* dialog, navigate to a location, enter a file name, and click **Save**.



Multi-Point Imported Properties

Input time history	The surface time history **G(sn)** to use as input, selected from pull-down menu.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported surface.	
	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

Multi-Point Imported Expert Properties

X Surface Axis Units	The units in which the X axis is specified.
Y Surface Axis Units	The units in which the Y axis is specified.
Z Surface Axis Units	The units in which the Z axis is specified.

Exportable Field Data

Re{FFT Coefficients}	The real parts of the FFT coefficients.

What Is Derived Data?

Derived data objects allow a data set function to access data from various other objects in the simulation, as well as from external files.

These objects are added under the **Imported**, **Tabular**, and **Monitor** child nodes of a data set function.



The child nodes of the **Imported**, **Tabular**, and **Monitor** nodes represents the derived data set. The input data is specified in the properties of these nodes. Once the properties are defined, the derived data (that is, the transform, convolution or filter function) are computed. The sub-nodes of a data set function can hold as many derived data sets as the memory of the computer permits.

This section explains the following:

What Are Derived Point Imported Model Data Sets?

Derived point imported model data sets hold a time history, FIR Filter function or Fourier transform.

They are used to visualize data at a point on a surface that is imported from a `.trn` file (see also Importing Surface Spectral Data) or the Fourier transform. The **point imported** node represents a derived point imported model data set in the simulation tree.



This object cannot be deleted while a link is in place from derived data sets in a plot. This object shares pop-up menu items with other derived data sets.

The function definition requires:

- Importing the surface from an existing `.trn` file.
- Extracting a field function from that surface.
- Selecting a point part as a probe.

The above quantities can be specified using the properties of this object. If the *Update Active* property is activated, the point time history or point time Fourier transform is calculated after the last property assignment. Alternatively, you can calculate the transform by right-clicking the **point imported** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Derived point imported model data sets are not updated automatically when you change a part position. A modified part can have many associated derived point imported model data sets. An automatic update would take some time and cause the STAR-CCM+ user interface to lock until the process is complete.

When any part is modified, all of its derived point imported model data sets are shown as requiring an update. The *Update Active* property for the affected derived point imported model data sets is automatically deactivated.

- To update a derived point imported model data set, do one of the following:
  - In the *Properties* window, manually activate the *Update Active* property.
  - Right-click the derived point imported model data set and then click **Update Output**.

What Are Derived Cross Imported Model Data Sets?

Derived cross imported model data sets hold correlations from two input signals.

The **cross imported** node represents a derived cross imported model data set in the simulation tree.



This object cannot be deleted while a link is in place from derived data sets in a plot. This object shares pop-up menu items with other derived data sets.

The function definition requires:

- Importing the surface from an existing `.trn` file.
- Extracting a field function from that surface.
- Selecting a point part as a probe.

The above quantities can be specified using the properties of this object. If the *Update Active* property is activated, the point-to-point time Fourier transform or the point-to-point time correlation is calculated after the last property assignment. Alternatively, you can calculate the transform by right-clicking the **cross imported** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Derived cross imported model data sets are not updated automatically when you change a part position. A modified part can have many associated derived cross imported model data sets. An automatic update would take some time and cause the STAR-CCM+ user interface to lock until the process is complete.

When any part is modified, all of its derived cross imported model data sets are shown as requiring an update. The *Update Active* property for the affected derived cross imported model data sets is automatically deactivated.

- To update a derived cross imported model data set, do one of the following:
  - In the *Properties* window, manually activate the *Update Active* property.
  - Right-click the derived cross imported model data set and then click **Update Output**.

Cross Imported Properties

Field Function 2	Defines the field function that is extracted. Only field functions added by a `.trn` file can be selected.
Imported surface 2	Defines the imported surface from which data is extracted. Only surfaces from a `.trn` file can be selected.

Point part 2	Defines the point x' on the surface where data is extracted. This point is the point on the surface closest to a selected point derived part.	
Point part 1	Defines the point x on the surface where data is extracted. This point is the point on the surface closest to a selected point derived part.	
Field Function 1	Defines the field function that is extracted. Only field functions added by a `.trn` file can be selected.	
Imported surface 1	Defines the imported surface from which data is extracted. Only surfaces from a `.trn` file can be selected.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported point. When a part is modified, this property is automatically deactivated for all of the affected derived cross imported model data sets.	
	Activated	The update is active.
	Deactivated	The update is inactive.

Creating Derived Point or Cross Imported Model Data Sets

Derived point or cross imported data sets are represented as child nodes of the **Imported** manager.

This manager has a pop-up menu.



To create a derived point or cross imported data set, right-click the **Imported** node and select **New derived data from imported model**.



A point imported or cross imported node appears as a child to this node.



Imported Properties

None.

Imported Pop-Up Menu

New derived data from imported model	Creates a point imported or cross imported child node in which data sets can be imported.
Paste	Places the contents of the clipboard into the selected node.
Refresh	Refreshes the sub-nodes of the selected node in the simulation tree.

See also common menu items.

Point Imported Properties

Data Directory	Specifies a file storing the derived data for later use. Data is saved in Mat7.3 format and is HDF5-compatible.
Point part 1	Defines the point x on the surface where data is extracted. This point is the point on the surface closest to a selected point derived part.

Field Function 1	Defines the field function that is extracted. Only field functions added by a `.trn` file can be selected.	
Imported surface 1	Defines the imported surface from which data is extracted. Only surfaces from a `.trn` file can be selected.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported point. When a part is modified, this property is automatically deactivated for all of the affected derived point imported model data sets.	
	Activated	The update is active.
	Deactivated	The update is inactive.

What Are Derived Imported Surface Data Sets?

Derived imported surface data sets hold surface time histories.

The **Point imported** node represents a derived imported surface data set in the simulation tree. The **Multi-Point imported** node represents a data set derived from imported surface data and a derived part.



This object cannot be deleted while a link is in place from derived data sets in a plot. This object shares pop-up menu items with other derived data sets.

The function definition requires:

- Importing the surface from an existing `.trn` file.
- Extracting a field function from that surface.
- Selecting a derived part or a point part as a probe.

The above quantities can be specified using the properties of this object. If the *Update Active* property is activated, the point-to-point time Fourier transform or the point-to-point time correlation is calculated after the last property assignment. Alternatively, you can calculate the transform by right-clicking the **Point imported** or **Multi-Point imported** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Derived imported surface data sets are not updated automatically when you change a part position. A modified part can have many associated derived imported surface data sets. An automatic update would take some time and cause the STAR-CCM+ user interface to lock until the process is complete.

When any part is modified, all of its derived imported surface data sets are shown as requiring an update. The *Update Active* property for the affected derived imported surface data sets is automatically deactivated.

- To update a derived imported surface data set, do one of the following:
  - In the *Properties* window, manually activate the *Update Active* property.
  - Right-click the imported surface data set and then click **Update Output**.

Creating Derived Imported Surface Data Sets

Derived imported surface data sets are represented as child nodes of the **Imported** manager.

This manager has a pop-up menu.



To create a derived imported surface data set, right-click the **Imported** node and select **New derived data from imported surface** or **New derived data from imported surface & derived part**



A **Point Imported** or **Multi-Point Imported** node appears as a child to this node.



Imported Properties

None.

Imported Pop-Up Menu

New derived data from imported surface	Creates a Point Imported or Multi-Point Imported child node in which data sets can be imported.

See also common menu items.

Multi-Point Imported Properties

Part	Specifies the derived part on the surface where data is extracted.	
Field Function 1	Defines the field function that is extracted. Only field functions added by a `.trn` file can be selected.	
Data surfaces 1	Defines the imported surfaces from which data is extracted. Only surfaces from a `.trn` file can be selected.	
Update Active	Whether the update is active for the imported point. When a part is modified, this property is automatically deactivated for all of the affected derived imported surface data sets.	
	Activated	The update is active.
	Deactivated	The update is inactive.

Display slider control Properties

Slider Value	Determines the current display time. The slider can vary from the maximum to the minimum time of the current signal.

Point Imported Properties

Point Part 1	Defines the point x on the surface where data is extracted. This point is the point on the surface closest to a selected point derived part.
RMS Magnitude	Specifies the size of the root mean square value of the signal.
Field Function 1	Defines the field function that is extracted. Only field functions added by a `.trn` file can be selected.

Data Surfaces 1	Defines the imported surfaces from which data is extracted. Only surfaces from a `.trn` file can be selected.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported point. When a part is modified, this property is automatically deactivated for all of the affected derived imported surface data sets.	
	Activated	The update is active.
	Deactivated	The update is inactive.

What Are Derived Transient Surface Data Sets?

Derived transient surface data sets hold fields that can be mapped onto unstructured surfaces.

Using them, you can:

- Map results on `.trn` file meshes at all time steps onto the `.trn` file mesh at a single time step.
- Map results on a mesh from a `.trn` file onto an external mesh.

The `.mat` file format can store unstructured surface meshes as well as structured meshes.

Creating Derived Transient Surface Data Sets

Derived transient surface data sets are represented as child nodes of the **Imported** manager.

This manager has a pop-up menu.



To create a derived transient surface data set, right-click the **Imported** node and select **New derived data from imported surface** or **New derived data from imported surface & derived part**.



A **Multi-Point Imported** node appears as a child to this node.



Imported Properties

None.

Imported Pop-Up Menu

New derived data from imported surface	Creates a Multi-Point Imported child node in which data sets can be imported.

Multi-Point Imported Properties

Target Surfaces	Specifies the surfaces to which the field function is mapped.	
Field Function 1	Defines the field function to be mapped. Only field functions added by a `.trn` file can be selected.	
Data surfaces 1	Defines the list of surfaces from which data is mapped. Only surfaces from a `.trn` file can be selected.	
Update Active	Specifies whether the update is active for the imported point. When a part is modified, this property is automatically deactivated for all of the affected derived transient surface data sets.	
	Activated	The update is active.
	Deactivated	The update is inactive.

Display Slider Control Properties

Slider Value	Determines the current display time. The slider can vary from the maximum to the minimum time of the current signal.

What Are Derived Tabular Data Sets?

Derived tabular data sets allow a Point Time Fourier Transform or Point Time History to access tables that you read into this simulation.

The tables can contain experimental data that you would use to validate simulation results.



This object cannot be deleted while a link is in place from derived monitor data sets in a plot. This object shares pop-up menu items with other derived data sets.

To access external data, carry out the following steps:

1. Read the data into a table in STAR-CCM+.
2. Select that particular table in the properties of this object.
3. In the same properties, assign the data columns of the table to the x- and y-axes.

The above quantities can be specified using the properties of this object. If the *Update Active* property is activated, the point time history or point time Fourier transform is calculated after the last property assignment. Alternatively, you can calculate the transform by right-clicking the **tabular** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Tabular Properties

Input Data	Specifies the table to use. The available tables are represented as nodes in the **Tables** node of the **Tools** node.	
X Column 1	Selects a column in the table to plot on the x-axis.	
Y Column 1	Selects a column in the table to plot on the y-axis.	
RMS Magnitude	Specifies the size of the root mean square value of the signal.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported table.	
	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

What Are Derived Cross Tabular Data Sets?

Derived cross tabular data sets allow correlations to access tables that you read into this simulation.

The tables can contain experimental data that you would use to validate simulation results.



This object cannot be deleted while a link is in place from derived monitor data sets in a plot. This object shares pop-up menu items with other derived data sets.

To access external data, carry out the following steps:

1. Read the data into a table in STAR-CCM+.
2. Select that particular table in the properties of this object.
3. In the same properties, assign the data columns of the table to the x- and y-axes.

The above quantities can be specified using the properties of this object. If the *Update Active* property is activated, the point-to-point time Fourier transform or the point-to-point time correlation is calculated after the last property assignment. Alternatively, you can calculate the transform by right-clicking the **cross tabular** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Cross Tabular Properties

Input Data 2	Selects a table from among the ones that are represented by nodes in the **Tables** node of the **Tools** node.
X Column 2	Selects a column in the table to plot on the x-axis.

Y Column 2	Selects a column in the table to plot on the y-axis.	
Input Data 1	Selects a table from among the ones that are represented by nodes in the **Tables** node of the **Tools** node.	
X Column 1	Selects a column in the table to plot on the x-axis.	
Y Column 1	Selects a column in the table to plot on the y-axis.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported table.	
	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

Creating Derived Tabular or Cross Tabular Data Sets

Derived tabular and cross tabular data sets are represented as child nodes of the **Tabular** manager.

This manager has a pop-up menu.



To create a derived tabular or cross tabular data set:

- Right-click the **Tabular** node and select **New derived data from table**.



A tabular or cross tabular node appears as a child to this node.



Tabular Properties

None.

Tabular Pop-Up Menu

New derived data from table	Creates a tabular or cross tabular child node in which tables can be accessed.
Paste	Places the contents of the clipboard into the selected node.
Refresh	Refreshes the sub-nodes of the selected node in the simulation tree.

What Are Derived Monitor Data Sets?

Derived monitor data sets allow a Fourier Transform or time history function to access the monitors that you create for this simulation.



This object cannot be deleted while a link is in place from derived data sets in a plot. This object shares pop-up menu items with other derived data sets.

To access a particular monitor, simply select it in the properties of this object. If the *Update Active* property is activated, the point time history or point time Fourier transform is calculated after the assigning a particular monitor. Alternatively, you can calculate the transform by right-clicking the **monitor** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Monitor Properties

Data Directory	Specifies a file storing the derived data for later use. Data is saved in Mat7.3 format and is HDF5-compatible.
Input Data 1	The monitor that is the source of the data of this object.
X Units	Displays the calculated units for the x-axis data.
Y Units	Displays the calculated units for the y-axis data.
Update Active	Whether the update is active for the imported monitor.

	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

What Are Derived Cross Monitor Data Sets?

Derived cross monitor data sets allow a correlation to access the monitors that you create for this simulation.



This object cannot be deleted while a link is in place from derived data sets in a plot. This object shares pop-up menu items with other derived data sets.

To access a particular monitor, simply select it in the properties of this object. If the *Update Active* property is activated, the point-to-point time Fourier transform or the point-to-point time correlation is calculated after the monitor assignment. Alternatively, you can calculate the transform by right-clicking the **cross monitor** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Cross Monitor Properties

Input Data 2	The monitor that is the source of the data of this object.	
Input Data 1	The monitor that is the source of the data of this object.	
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported monitor.	
	Activated	The update is active.
	Deactivated	The update is inactive. This state is the default.

Creating Derived Monitor or Cross Monitor Data Sets

Derived monitor and cross monitor data sets are represented as child nodes of the **Monitor** manager.

This manager has a pop-up menu.



To create a derived monitor or cross tabular monitor set:

- Right-click the **Monitor** node and select **New derived data from monitor**.



A monitor or cross monitor node appears as a child to this node.



Monitor Properties

None.

Monitor Pop-Up Menu

New derived data from monitor	Creates a monitor or cross monitor child node in which monitors can be accessed.
Paste	Places the contents of the clipboard into the selected node.
Refresh	Refreshes the sub-nodes of the selected node in the simulation tree.

What Are Point Derived Data Sets?

Derived data sets allow a point inverse time Fourier transform data set function to access any of the existing derived data objects.



This object cannot be deleted while a link is in place from derived data sets in a plot. This object shares pop-up menu items with other derived data sets.

To access a particular derived data object, simply select it in the properties of this object. If the *Update Active* property is activated, the

point inverse time Fourier transform is calculated after the derived data assignment. Alternatively, you can calculate the inverse transform by right-clicking the **point derived** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Point Derived Properties

Input Data	Accesses a particular type of derived data object from a drop-down list of the derived data objects that currently exist in the simulation.	
	None	Does not use a derived data object.
	Derived imported model	Accesses a derived imported model data set.
	Derived monitor	Accesses a derived monitor data set.
	Derived tabular	Accesses a derived tabular data set.
X Units	Displays the calculated units for the x-axis data.	
Y Units	Displays the calculated units for the y-axis data.	
Update Active	Whether the update is active for the imported data.	
	Activated	The update is active.
	Deactivated	The update is inactive.

Derived Data Pop-Up Menu

These objects cannot be deleted while a link to a plot data set is in place.

Update Output	Updates the derived data set.
Export...	Exports the data to a `.csv` file.
Export Audio...	Exports an audio (.wav) file. Applies only to derived data nodes for Inverse Fourier Transforms.
Tabulate	Executes the tabulate data action. This action presents the data in a two-column

	spreadsheet-style view of the data in this data set. The abscissa (horizontal axis) data is in the first column and the ordinate (vertical axis) data is in the second column. When there is no data in the data set, the dialog is empty.

Creating Point Derived Data Sets

A point derived data set is represented as a child node of the **Derived** manager, which has a pop-up menu.



To create a point derived data set:

- Right-click the **Derived** manager node and select **New derived data from derived data**.



A point derived node appears as a child to this node.



Derived Properties

None.

Derived Data Pop-Up Menu

New derived data from derived data	Creates a point derived child node in which monitors can be accessed.

Paste	Places the contents of the clipboard into the selected node.
Refresh	Refreshes the sub-nodes of the selected node in the simulation tree.

Calculating the Line Time Fourier Transform

Before calculating a line time Fourier transform from a line probe, create a derived data set under the **Imported** child node of the transform:

1. Right-click **Tools > Data Set Functions > G(l) 1 > Imported** and select **New derived data from imported model**.



A **line imported** node appears in the simulation tree.



To calculate the line time Fourier transform, assign the line probe and imported surface to an existing line derived data set and select the field function from which the signal is extracted:

2. Select the **G(l) 1 > Imported > line imported** node



3. In the *Properties* window, set *Line part 1* to the line probe you created for the field function extraction.

4. Set *Field function 1* to the field function from which the signal is extracted.
5. Set *Imported surface 1* to the surface on which the field function is extracted.

If the *Update Active* property is activated, the line time Fourier transform is calculated after the last property assignment. Alternatively, you can calculate the transform by right-clicking the **line imported** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Line Imported Properties

Input time history	The line time history **h(ln)** to use as input, selected from pull-down menu.	
X Units	The units of the X axis.	
Y Units	The units of the Y axis. Equals the square of the units of *Field function 1* divided by the units of frequency for power spectral density plots.	
Update Active	Whether the update is active for the imported line.	
	Activated	The update is active.
	Deactivated	The update is inactive.

Line Imported Expert Properties

X Surface Axis Units	The units in which the X axis is specified.
Y Surface Axis Units	The units in which the Y axis is specified.
Z Surface Axis Units	The units in which the Z axis is specified.

Calculating the Line Spatial Fourier Transform

Before calculating the line spatial Fourier transform from an existing line time Fourier transform, import the data under the **Derived** child node of the line spatial Fourier transform:

1. Right-click **Tools > Data Set Functions > F(G(l)) 1 > Derived** and select **New derived data from derived data**.



A **line derived** node appears in the simulation tree.



To calculate the line spatial Fourier transform, select the line derived data that was created for the line time Fourier transform:

2. Select the **F(G(l)) 1 > Derived > line derived** node.



3. In the *Properties* window, set *Input derived data* to the derived data you created for the line time Fourier transform.

If the *Update Active* property is activated, the line spatial Fourier transform is calculated after the property assignment. Alternatively, you can calculate the transform by right-clicking the **line derived** node and selecting `Update Output`, irrespective of the state of the *Update Active* property.

Line Derived Properties

Input derived data	The derived data that was created for the line time Fourier transform.	
X Units	The units of the X axis.	
Y Units	The units of the Y axis.	
Update Active	Whether the update is active for the derived line.	
	Activated	The update is active.
	Deactivated	The update is inactive.

Line Derived Expert Properties

X Surface Axis Units	The units in which the X axis is specified.
Y Surface Axis Units	The units in which the Y axis is specified.
Z Surface Axis Units	The units in which the Z axis is specified.

Post-Processing Data Set Functions

STAR-CCM+ allows you to visualize Fourier transforms and surface spectra.

The follow sections explain the following:

Plotting Fourier Transforms

You can plot a Fourier transform from a monitor containing periodic data.

To plot a Fourier Transform:

1. Ensure that the *Trigger* property of the monitor is set to `Time Step`.
2. Create a Fourier Transform and set its properties.
3. Right-click the **Monitor** node under this transform and choose **New derived data from monitor**.
4. Give the derived monitor a meaningful name and in the derived monitor *Properties* window, assign the *Input data 1* property to the monitor containing the periodic data.



5. Create a monitor plot after the simulation has run.
6. Click the **Derived** node under this monitor plot, and in the *Properties* window, select the derived data monitor to assign it to the *Derived data* property.



7. Double-click the monitor plot node to display the Fourier Transform plot.



Setting the Transform Plot Properties of Line Data Sets

You can set the plot properties of a Line Time Fourier Transform or a Line Spatial Fourier Transform using the **Output control** node.



Output control Properties

3D profile type	Selects the type of three-dimensional display.	
	F-X spectra surface	The auto- or cross-spectra plotting frequency vs. spatial dimension.
	3D spectra surface	The auto-spectra on 3D geometry for a given frequency.
Output field	Selects which kind of FFT information to visualize.	
	Re{FFT Coefficients}	The real parts of the FFT coefficients.
	Im{FFT Coefficients}	The imaginary parts of the FFT coefficients.
	Power Spectral Density	Frequency-domain description of the signal power.
	Phase{FFT Coefficients}	The phases of the FFT coefficients ($atan[Im / Re]$).

	Abs{FFT Coefficients}	The absolute values of the FFT coefficients ($\sqrt{\lvert \mathrm{Re} \rvert^2 + \lvert Im \rvert^2}$).
	Sound Pressure Level	Sound pressure in decibels. Adds the Reference Sound Pressure node.
	A-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the A-scale function, for loudness less than 55 phons (at 1000 Hz, 1 phon = 1 dB). This function is typically for problems that involve the frequency range for human hearing. Adds the Reference Sound Pressure node.
	B-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the B-scale function, for loudness of 55-85 phons. Adds the Reference Sound Pressure node.
	C-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the C-scale function, for loudness greater than 85 phons. Adds the Reference Sound Pressure node.
	Root Mean Square	The square root of the `Power Spectral Density`.
Log data	Displays the log of current data in the current scene or plot.	
	Activated	Log displayed.
	Deactivated	Log not displayed. The default.

2D profile range	Defines the two-dimensional profile.	
	All	Uses all signal frequencies.
	Value	Uses a single value frequency. Adds the Single Value child node to this node.
	Band	Uses a band of frequencies. Adds the Single Band child node to this node.
2D profile axis	Defines the two-dimensional profile.	
	Spatial axis	The abscissa is in length units.
	Frequency axis	The abscissa is in frequency units.

Display slider control Properties

Slider Value	Determines the current display time. The slider can vary from the maximum to the minimum time of the current signal.

Setting a Two-Dimensional Plot of the Transform Using a Single Frequency

You can use a single frequency to plot a Line Time Fourier Transform or a Line Spatial Fourier Transform in two dimensions.

To plot a Line Time Fourier Transform or a Line Spatial Fourier Transform:

1. Select the **Plot output node** and set its *2D profile range* to **Value**.



The **Single Value** node appears as a child node of **Plot output**.



2. Select the **Plot output** > **Single Value** node and set its *Value* to the required frequency.

Single Value Properties

Value	The value of the frequency for which the Fourier transform is plotted in two dimensions. Clicking the ... property customizer button activates a property customizer that lets you specify both the frequency value and the appropriate time units (**/s** or **Hz**).

Setting a Two-Dimensional Plot of the Transform Using a Frequency Band

You can use a frequency band to plot a Line Time Fourier Transform or a Line Spatial Fourier Transform in two-dimensions.

To plot a Line Time Fourier Transform or a Line Spatial Fourier Transform:

1. Select the **Plot output node** and set its *2D profile range* to **Band**.



The **Single Band** node appears as a child node of **Plot output**.



2. Select the **Plot output** > **Single Band** node and set its *Value* to the required frequency.

Single Band Properties

Low Range	The value of the lowest frequency for which the Fourier transform is plotted in two dimensions.
High Range	The value of the highest frequency for which the Fourier transform is plotted in two dimensions. Clicking the ... property customizer button activates a property customizer that lets you specify both the frequency value and the appropriate time units (**/s** or **Hz**).

Post-Processing Surface Spectra

The process of setting up the post-processing of surface spectra involves several steps. The .trn file is used to display data set functions.

The procedure to post-process and visualize surface spectra is as follows:

Exporting Surface Spectral Data

To export surface spectral data, from the menu bar select **File > Auto Export...** opening the *Auto Export* dialog.



The available property settings are detailed below:

- *Trigger Type* -- The `.trn` file can be exported for any transient case setting, with the following options available:
  - **Iterations** -- Data is exported every n iterations, where n is specified in the *Frequency* property.

  - **Time Steps** -- Data is exported every n time steps, where n is specified in the *Frequency* property.
  - **Delta Time** -- Data is exported as defined by the *Delta Time* property.
- *Enabled* -- Must be activated in this instance, to allow the auto export of data.
- *Base Name* -- The file name of the `.trn` file.
- *Frequency* -- Specifies the number of *Iterations* or *Time Steps* between the exports. Neglected if *Delta Time* is set as the *Trigger Type*.
- *Delta Time* -- Specifies the duration of solution time between exports. Neglected if *Iterations* or *Time Steps* is set as the *Trigger Type*.
- *Export Directory* -- The directory to which the `.trn` file is exported.
- *File Format* -- Must be set to **Surface fft files (*.trn)**.
- *Append to File* -- Must be activated to add data to the file defined by the *Base Name* and *Export Directory* properties.
- *Export Solution Data Only* -- If activated the mesh is exported once only. Therefore, effects such as a moving mesh are not captured.
- *Scalar Functions* -- Define the scalar functions to export.
- *Vector Functions* -- Define the vector functions to export.
- *Regions* -- Neglected when exporting to a `.trn` file.
- *Boundaries* -- Define the boundaries to export.
- *Derived Parts* -- Define the derived parts to export.

When the case is run, a `.trn` file is generated, which can then be imported.

Importing Surface Spectral Data

To import surface spectral data, from the menu bar select **File > Import > Import CAE Model...** and navigate to the generated `.trn` file.

You can import surface spectral data for a pre-existing simulation or a new one. When you try this feature for the first time, you can use a new simulation so it is evident what has been imported.

After the import process, a new **Imported Models** manager node appears in the object tree. This node contains imported surfaces corresponding to the previously exported surfaces.



Data set functions can only be performed on surfaces whose discretization does not change with time. For example, if a data set function is calculated on an exported isosurface, values of zero are returned.

Working with Surface Spectra Field Functions

Importing a `.trn` file registers the following field functions:

- Time History:



- Fourier Transform:



The time history and field functions can be displayed in a scene, by setting the appropriate field function as a scalar displayer.

Using the Time History Field Functions

The `.trn` import registers a number of time history field functions relating to the exported field data.



These field functions have mostly read-only properties, except that the **Time History:Pressure** field function lets you select a filter type, which activates a FIR Filter sub-node.

Using the Fourier Transform Field Function

The import registers an FFT field function, relating to the exported scalar data. This field function has a number of properties that allow you to control the Fourier transform, in addition to the base properties of a typical field function.



Fourier Transform Field Function Properties

Fft Input Start Time	The start time (in seconds) for the time sampling for this object. If you have multiple FTs, you can vary the settings of this property to sample multiple time intervals from the simulation run.
Fft Input Cut-off time	The end time (in seconds) for the time sampling for this object. If you have multiple FTs, you can vary the settings of this property

	to sample multiple time intervals from the simulation run.	
Amplitude Function	Specifies the amplitude function that is generated. These functions are all standard for this type of analysis.	
	Power Spectral Density	Frequency-domain description of the signal power.
	Sound Pressure Level	Sound pressure in decibels. Adds the Reference Sound Pressure node.
	Root Mean Square	The square root of the `Power Spectral Density`.
	A-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the A-scale function, for loudness less than 55 phons (at 1000 Hz, 1 phon = 1 dB). This option is typically for problems that involve the frequency range for human hearing. Adds the Reference Sound Pressure node.
	B-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the B-scale function, for loudness of 55-85 phons. Adds the Reference Sound Pressure node.
	C-Weighted Sound Pressure Level	Weights the `Sound Pressure Level` with the C-scale function, for loudness greater than 85 phons. Adds the Reference Sound Pressure node.

Frequency Range Type	Specifies the frequency function that is used. The sub-node that is contained within the FT field function is used to define the input for each frequency range type.	
	Frequency Value	Frequency in Hz. You specify a single value.
	Frequency Range	Frequency in Hz over a range of values. You specify the lower and upper limits.
	Octave Bands	Discrete frequency bands corresponding to different octaves within hearing range. Each successive band represents twice the range of the previous one. The current frequency band defines the hearing range.
	3rd Octave Bands	Discrete frequency bands corresponding to one-third of each octave within hearing range. The current frequency band defines the hearing range.
Analysis Blocks	Specifies the number of analysis blocks used.	
Overlap Factor	Sets the overlap factor between blocks, if any, in a range from `0` to `0.9`. This overlap allows for smooth transitions between blocks.	
Field update	Controls whether the FT field function is updated when the display time or when a field function property is changed.	
	Activated	The FT field function updates when a property is changed or when the display time is changed.

	Deactivated	The FT field function does not update when the display time is changed or when a property in the field function is changed. This option can reduce the amount of time that is taken to update the results.
Window Function	Specifies the window function that you require for this FT. These options are all standard for this type of analysis. • **None** - This default setting leaves the FT without a window function until you select one of the other options in this list. • **Hamming** • **Hann** (Hanning) • **Bartlett** • **Blackman**	
Nyquist Frequency	A read-only value in Hz.	
Filter Type	Specifies the type of filter that is applied.	
	None	Does not use a filter.
	Band pass filter	Keeps frequencies between low and high cut-off levels. Adds an Ideal Filter sub-node in which you can set these levels.
	Band stop filter	Discards frequencies between low and high cut-off levels. Adds an Ideal Filter sub-node in which you can set these levels.
Dimensions	A read-only value of the dimensions that STAR-CCM+ calculates for the FT field function.	

Selecting Units for FT Field Functions

STAR-CCM+ calculates the dimensions for the FT field functions and provides a list of suggested units that are based on these dimensions.

The units for the FT field function are located in *Properties* window of the **Displayers > Scalar 1 > Scalar Field** node of the scene.



In the example that is shown above, STAR-CCM+ calculated the units **using the** dimensions **for the Fft: Pressure** field function (**Pressure** and **Time)**. The suggested units are shown below.



Select the most relevant units from this list. In this particular example, **(Pa)/Hz** is the most suitable choice.

Using the Display Time

Display time is controlled from a slider on the transient model, which updates the scene automatically.



This slider does not affect FT field functions which are controlled entirely from their own properties.

When the *Update data set functions* property of the **Imported Models > trn:fileName** node is activated, the data set functions that contain data from a `.trn` file recalculate whenever the *Transient Display Time* is changed. As this calculation can take some time, you can deactivate the update option until all the times are set to the values you require.



Visualizing the Spectra for Line Data Set Functions

Spectra contours give useful clues on the dominant frequencies in a problem.

To visualize spectra contours:

1. Create a scalar scene.
2. Hide the scene outline.
3. Select the derived data using the *Parts* property of the scalar displayer.
   a) Select the **Scalar 1 > Parts** node.
   b) In the *Properties* window, click the ... property customizer button for the *Parts* property.
      The *Parts* dialog appears.
   c) Select the derived data set.



d) Click **OK**.



4. Select the **Scalar Field** node and set its *Function* to the derived data.



The spectra contour is displayed in the scalar scene window. In the following example, the dominant frequency mode is around 900 Hz.


Geometry Scene 1
Monitor Plot 1
Scalar Scene 1
Power Spectral Density (Pa^2/Hz)
Frequency (Hz)
Distance (m)
Y
Z
X
line imported (Pa^2/s)
-19.028
-14.457
-9.8863
-5.3152
-0.74418
3.8269

Pop-Up Menu of the Data Set Functions Manager

New	Creates a data set function of the type you select in the submenu:	
	Point Time History (h(p))	Creates a point time history.
	Point Time Fourier Transform (G(p))	Creates a point time Fourier transform.
	Point-to-point Time Fourier Transform (G(p,q))	Creates a point-to-point time Fourier transform.
	Point-to-point Time Correlation (R(p,q))	Creates a point-to-point time correlation.
	Point Inverse Time Fourier Transform (h′ (G(p)))	Creates a point inverse time Fourier transform.
	Line Time Fourier Transform (G(l))	Creates a line time Fourier transform.
	Line Spatial Fourier Transform (F(G(l)))	Creates a line spatial Fourier transform.
	Surface Time History(h(s))	Creates a surface time history.
	Surface time Fourier Transform (G(h(s)))	Creates a surface time Fourier transform.

Data Set Functions Formulation

This section includes the following formulation:

Basic Relations

The time Fourier transform is defined as

$$FT_t[h(x;t)] = \int_{-\infty}^{\infty} h(x:t)e^{-i\omega t}\,\square t = g(x;\omega) \quad \forall\,\omega \in [-\infty,\infty] \tag{2805}$$

where $h(x;t)$ is a real function which depends on the position x and time $t \cdot \omega$ is angular frequency.

The integration limits are over the whole time domain $t \in [\infty,\ -\infty]$. However, the integration is over a given period $t \in [-T/2,\ T/2]$ in practice. The function $h(x;t)$ is periodic in time (or is altered to be periodic).

More commonly, a "finite" Fourier transform for a specific time block T is used:

$$FT_t[h(x;t)] = \int_{-T/2}^{T/2} h(x:t)e^{-i\omega t}\,\square t = g(x;\omega,T) \quad \forall\,\omega \in \left[-\frac{m\pi}{T},\frac{m\pi}{T}\right] \tag{2806}$$

where m is the number of sampling points.

If the finite transform is conducted on a kth record of block T, then the transform is:

$$FT_t[h(x;t)] = \int_{-(t_k+(T/2))}^{(t_k+(T/2))} h(x:t)e^{-i\omega t}\,\square t = g(x;\omega,T) \quad \forall\,\omega \in \left[-\frac{m\pi}{T},\frac{m\pi}{T}\right] \tag{2807}$$

The inverse time Fourier transform is

$$FT_t^{-1}[g(x;\omega)] = \frac{1}{2\pi}\int_{-\infty}^{\infty} g(x;\omega)e^{-i\omega t}\,\square\omega = h(x;t) \quad \forall\,t \in [\square -\infty,\infty] \tag{2808}$$

Similarly, although the integration domain is infinite, it is understood that this integration is over a finite frequency space.

The space Fourier transform is defined as

$$FT_s[g(x;\omega)] = \int_x g(x;\omega)e^{-ik\cdot x} \,|\, \square x \,| = f(k;\omega) \quad \forall k \in K; \quad \omega \in [\square -\infty, \infty] \tag{2809}$$

The function is periodic is space (or is altered to be periodic in space).

Using a compact notation

$$g(x;\omega)e^{-ik\cdot x} \,|\, dx \,| \equiv g(x(\xi);\omega)e^{-ik\cdot s(\xi)} \left| \frac{\partial x}{\partial \xi_1} \times \frac{\partial x}{\partial \xi_2} \right| \partial \xi_1 \partial \xi_2 \tag{2810}$$

The space Fourier transform is actually in the parametric space instead of the three-dimensional space of the coordinates of a point on a surface. This parametric representation maps to a two-dimensional space which is both planar and regularly sampled.

The inverse space Fourier transform is

$$FT_s^{-1}[f(x;\omega)] = \frac{1}{(2\pi)^n}\int_K f(x;\omega)e^{-ik\cdot x} \,|\, \square k \,| = g(x;\omega) \quad \forall x \in S;\ \omega \in [\square -\infty, \infty] \tag{2811}$$

Here, n is the dimensionality of the wavenumber space K. For transforms on a surface $k = \{k_1, k_2\} \in K$, $n = 2$, . For transforms on a curve, $n = 1$.

Time-Space Fourier Transforms

The time-space Fourier transform is defined as

$$FT_{ts}[h(x;t)] = \int_s \left(\int_{-\infty}^{\infty} h(x,t)e^{-i(\omega t + k\cdot s)} \square t \right) \,|\, dx = f(k;\omega) \quad \forall k \in K; \quad \omega \in [-\infty, \infty] \tag{2812}$$

The inverse time-space Fourier transform is

$$FT_{ts}^{-1}[f(k;\omega)] = \frac{1}{(2\pi)^3}\int_K \left(\int_{-\infty}^{\infty} f(k;\omega)e^{-i(\omega t + k\cdot s)} \square \omega \right) \,|\, dx \,| = h(x;t) \quad \forall x \in S;\ t \in [-\infty, \infty] \tag{2813}$$

The following relationships hold between the transforms:

$$FT_{ts}[h(x;t)] = f(k;\omega) = FT_s[g(x;\omega)] = FT_s[FT_t[h(x;t)]] \tag{2814}$$

$$FT_{ts}^{-1}[f(k;\omega)] = h(x;t) = FT_t^{-1}[g(x;\omega)] = FT_t^{-1}[FT_s^{-1}[f(k;\omega)]] \tag{2815}$$

Frequency-Space Auto-Spectrum and Cross-Spectrum

The frequency-space (fs) auto-spectrum is defined as

$$G(x;\omega) = g^*(x;\omega)g(x;\omega) \tag{2816}$$

where g^* is the complex conjugate of g.

The frequency-space (fs) cross-spectrum is

$$G(x_1,x_2;\omega) = g^*(x_1;\omega)g(x_2;\omega) \quad \forall x_1 \in S, \forall x_2 \in S \tag{2817}$$

Defining the "finite Fourier transform" for the kth sample time block of length T:

$$g(x;f,T) = \int_0^1 h(x,t)e^{-i\omega t}\,dt \tag{2818}$$

The one-sided cross-spectral density between random processes at any two points in space x_1, x_2 is defined as

$$G(x_1,x_2;f) = \lim_{T\to\infty}\frac{2}{T}E\left[g_k^*(x_1;f,T)g_k(x_2;f,T)\right] \tag{2819}$$

The corresponding one-sided auto-spectral density of the random signal at a single point is:

$$G(x,x;f) = \lim_{T\to\infty}\frac{2}{T}E\left[g_k^*(x;f,T)g_k(x;f,T)\right] = \lim_{T\to\infty}\frac{2}{T}E\left[\left|g_k(x;f,T)\right|^2\right] \tag{2820}$$

where $E[\;]$ represents the expectation operator averaging over all indices k. The expectation is important to obtain the coherence between the two random processes, which are defined below.

Since the cross spectral density is in general complex, it is convenient to convert to magnitude and phase. Magnitude is defined as:

$$\left|G(x_1,x_2;f)\right| = \sqrt{\left[Re\ G(x_1,x_2;f)\right]^2 + \left[Im\,G(x_1,x_2;f)\right]^2} \tag{2821}$$

and phase is defined as:

$$\theta(x_1,\ x_2;f) = atan\left(\frac{Im\,G(x_1,x_2;f)]}{Re\,G(x_1,x_2;f)]}\right) \tag{2822}$$

The coherence is the measure of how much of the random signal at position x' exhibits that time-averaged phase correlation with the random signal at position x. Coherence is defined as

$$\gamma^2(x_1,x_2;f)=\frac{|G(x_1,x_2;f)|^2}{G(x_1,x_1;f)G(x_2,x_2;f)} \tag{2823}$$

Time Cross-Correlations

The time cross correlation is

$$R(x_1,x_2;t)=\int_{-\infty}^{\infty}h(x_1;\tau)h(x_2;t+\tau)\,d\tau \quad \forall x_1\in S,\ \forall x_2\in S \tag{2824}$$

This convolution goes over the finite period. Periodic boundary conditions must be applied. The Fourier transform makes the same assumption of periodicity. When a Fourier transform is calculated, all notions of causality are lost. You can no longer refer to driver and response signals, but only signals which are correlated. That is, it is impossible to say which signal drives which.

The relationship between the time cross-correlation and the time cross-spectrum is

$$R(x_1,x_2;t)=FT_t^{-1}\left[G(x_1,x_2;\omega)\right] \tag{2825}$$

Equivalently,

$$G(x_1,x_2;\omega)=FT_t\left[R(x_1,x_2;t)\right] \tag{2826}$$

Frequency-Space Parseval's Relation

Parseval's relation states that the total square of pressure is

$$p^2(x)=\int|h(x;t)|^2\,dt=\frac{1}{2\pi}\int G(x;\omega)\,d\omega \tag{2827}$$

where the square root of p^2 gives the root mean square pressure p_{rms}.

Frequency-Wavenumber Auto-Spectrum

The frequency-wavenumber (fk) auto-spectrum is defined as

$$F(k;\omega) = f^*(k;\omega) f(k;\omega) \quad \forall k \in K; \quad \omega \in [-\infty, \infty] \tag{2828}$$

Spatial Cross-Correlations

The spatial cross-correlation is defined as

$$C(x_1, x_2; \omega) = \int_s g(x_1;\omega) g(x_2;\omega) \,|\, dx \,| \tag{2829}$$

Assuming that the cross-correlation is weakly homogeneous and relabeling,

$$C(x_1, x_2; \omega) \approx G_{xy}(r;\omega) \tag{2830}$$

where $r = x_1 - x_2$

The convolution (or cross-relation) in physical space is equal to the frequency-wavenumber (fk) auto-spectrum in wavenumber space. Therefore,

$$F_{xy}(k;\omega) = FT_s\left[G_{xy}(r;\omega)\right] \tag{2831}$$

Equivalently,

$$G_{xy}(r;\omega) = FT_s^{-1}\left[F_{xy}(k;\omega)\right] \tag{2832}$$

Frequency-Wavenumber Parseval's Relation

Parseval's relation states that the total square of pressure is

$$p^2 = \iint |\, h(x;t) \,|^2 \,|\, dx \,|\, dt = \frac{1}{(2\pi)^{n+1}} \iint F_{xx}(k;\omega) \,|\, dk \,|\, d\omega \tag{2833}$$

where the square root of p^2 gives the root mean square pressure p_{rms}.

Wave Group Velocity

The group velocity v_g of a wave is the velocity of propagation of the envelope (or modulation) of the wave in space.

The relationship between the phase velocity , wavelength and frequency of a wave is

$$v_\theta = f\lambda \tag{2834}$$

k is defined as

$$k = \frac{2\pi}{\lambda} \tag{2835}$$

In STAR-CCM+, the wavenumber k is plotted in the range $[-\pi, \pi]$. Therefore,

$$v_\theta = \frac{\pi f}{k} \tag{2836}$$

The magnitude of the group velocity is equal to π times the slope of a frequency-wavenumber (fk) plot:

$$v_g = \pi \frac{\partial f}{\partial k} \tag{2837}$$

Spectra Averaging

Defining the time series,

One method of temporal spectral averaging is considered here: blocking.

Blocking is the process of smoothing a spectra. The initial signal is broken into a number of subsignals of the same length. At the same time, sampling granularity is held constant.

Additional blocks can be overlapped using an overlap factor, which can have values in the range $0 < a < 0.9$. In this case, the number of blocks is

and they are defined as

Then, treating each subsignal as an independent signal, define the blocking periods

over which the mean time history due to blocking can be calculated.

Defining a temporal mean spectrum due to the mean of the Fourier transforms of a set of subsignals,

From the Fourier transform addition theorem, you can see that:

Blocking reduces the length of time series. As a result, low frequencies are less well characterized.

It is also possible to average spectra spatially. The spatial mean spectra of the points $\{\boldsymbol{P}_i,\ i=1,N_p\}$ can be expressed as

The temporal-spatial mean is defined as

Spectra Filtering

It is possible to filter time histories using an ideal filter in the frequency domain.

The band pass filter is defined as follows

$$T^{bp}_{\omega_1,\omega_2}(\omega)=\begin{Bmatrix}1,\ \omega\in[\omega_1,\ \omega_2]\\ 0,\ \omega\notin[\omega_1,\ \omega_2]\end{Bmatrix} \tag{2848}$$

with the band stop filter defined as the inverse of Eqn. (167).

$$\omega_1=2\pi f_1 \tag{2849}$$

where f_1 is the low cut-off frequency.

$$\omega_2=2\pi f_2 \tag{2850}$$

where f_2 is the high cut-off frequency.

The filtering operation is then a product in the frequency domain. A frequency filtered time history $h'(x;t)$ can then be defined as follows

$$h'(x;t)=FT_t^{-1}\left[T^{bp}_{\omega_1,\omega_2}(\omega)g(x;\omega)\right] \tag{2851}$$

Sound Pressure Level

The sound pressure is the local change in pressure from the reference ambient pressure that a sound wave causes.

The Sound Pressure Level (SPL) is a logarithmic measure of the effective sound pressure relative to a reference value, and is measured in decibels (dB):

$$SPL = 10\ \log\left(\frac{p_{rms}^2}{p_{ref}^2}\right) \tag{2852}$$

where:

- p_{rms} is the root mean squared sound pressure.
- p_{ref} is the reference sound pressure.

Wizard Guide

This part of the documentation illustrates how STAR-CCM+ is used for special application tools and projects.

The application tools are:

- The Turbo wizard
- The Fire and Smoke wizard
- The Thermal Comfort wizard
- The Wing Mesher wizard
- The SPEED Motor Tool Import wizard

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Modeling physics—the concepts, theory, and implementation details of the models
- Using tools in STAR-CCM+ to facilitate your work
- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format
- A guide to using STAR-View+, a viewer for three-dimensional STAR-CCM+ scenes

See Also

Using STAR-CCM+
Modeling Geometry
Meshing
Best Practices

Running
Analyzing
Tutorial Guide

Using the Turbo Wizard

The turbo wizard is used in the setup of turbomachinery simulations.

The turbo wizard provided by STAR-CCM+ is used in the setup of turbomachinery simulations to:

- generate a high-quality mesh for the fluid domain surrounding a blade, or blades
- set boundary and initial conditions for the fluid domain (optional)

This section includes information on the capabilities of the turbo wizard and describes how to use the wizard to set up a turbomachinery case.

What is the Turbo Wizard?

The turbo wizard is an automatic meshing tool for creating high-quality turbomachinery meshes quickly.

Because these meshes are multi-block and comprised of hexahedral cells, they are highly accurate and efficient to solve. Its capabilities include:

- Elliptic and Algebraic grid generation system
- Moving mesh setup
- Direct blade definition from section data
- Boundary surface monitoring (forces, pressures...)

Because the mesh near the blade is so critical to the accuracy of the analysis, the resolution requirements on the mesh are demanding. In order to avoid creating meshes with unnecessarily high cell counts, the turbo wizard creates a mesh using one of the elliptic (”O”) grid types.

While most problems can be meshed in several ways, the best choice of mesh topology depends upon your particular geometry and analysis requirements.

The process for analyzing turbomachinery problems is divided into three steps:

1. Meshing the geometry
2. Solving the flow problem
3. Post-processing the flow solution

The geometry and demands put upon the resolution of the model near a typical blade/impeller setup make step 1 challenging. The turbo wizard simplifies the creation of a high-quality mesh and application of appropriate boundary conditions to it.

A typical analysis starts with either a blade design or a CAD model of the blade. This can be imported into the turbo wizard by creating an input file that defines the blade surface as a set of coordinates. By default, the flow is assumed to move in the +z direction.

The fluid mesh is then created around it, boundary conditions are applied to the fluid domain, and the model is solved in STAR-CCM+.

See Also

Working with Turbomachinery Input Files

Turbomachinery Related Capabilities of STAR-CCM+

The turbo wizard is the primary tool for turbomachinery problems, but there are several other STAR-CCM+ features that can help with the CFD solution.

The summary of the capabilities of these tools includes :

- Modeling Options:
  - Multiple (rotating) Reference Frames
  - Chemically Reacting Flows
  - Free Surface Flows
  - Conjugate Heat Transfer
- Turbulence Models:
  - Spalart-Allmaras
  - K-Epsilon, K-Omega...
- Boundary Conditions:
  - Inlet
  - Stagnation
  - Radial Equilibrium
  - Non-reflecting (transient)
  - Pressure
- Solver:
  - Second-Order
  - Segregated
  - Algebraic Multigrid Acceleration
- Automatic Meshing Tools:
  - Tetrahedral, hexahedral, and polyhedral meshing with surface extrusion layers
  - Elliptic mesh generator for hexahedral meshes near the blade with the turbo wizard

A typical example can involve generating a single-sector mesh of the near-blade region using the turbo wizard and then including this portion in a more complicated model of a piping, ducting, or another

fluid handling system. This approach leverages the high-quality hexahedral meshing capabilities of the turbo wizard with the general, automatic meshing capabilities of STAR-CCM+. The result is an easy-to-create mesh that gives you control over fine details, while minimizing effort and solution time.

Setting Up the Turbo Wizard

The turbo wizard is used to generate a high-quality mesh for turbomachinery blades, and can also be used to set up boundary and initial conditions.

These mesh and physics properties are specified using the *Import turbo wizard geometry file options* dialog, which appears when the turbo wizard is activated.

Clicking **OK** in this dialog always triggers mesh generation. If the information entered in the dialog is invalid or incomplete, this may result in an error. Clicking **Cancel** exits the dialog but does not clear the settings and data that have already been entered. This is discussed in greater detail below.



There are two broad categories of data that are entered in the dialog:

- Meshing data—specifies the properties of the mesh and is defined in the fields outside of the tabbed panel, then in the first four tabs:
  - Number of blade rows
  - Rotation of the current row
  - Flow axis
  - Mesh arrangement

  - Geometry tab
  - Meshing tab
  - Flowpath tab
  - Grid settings tab

  This data must be entered to generate a mesh.

- Physics data—sets the boundary and initial conditions and is defined remaining two tabs:
  - Boundary conditions / turbulence tab
  - Initial conditions tab

  You do not have to enter values in these tabs, but the turbo wizard uses the default values to set up the models and conditions in the simulation if no data is provided.

The turbo wizard is an options dialog, so all settings and values entered in the dialog are retained when it is closed and reopened. When the simulation is closed, the turbo wizard dialog is cleared, and the only way of retaining the information is to save the turbo wizard data separately. Turbo wizard data is related to the active simulation, and cannot be applied across several simulations unless the settings are saved in one simulation and loaded into the next.

See Also

Saving and Loading Turbo Wizard Data

Activating the Turbo Wizard

This section describes how to activate the Turbo Wizard.

To activate the Turbo Wizard:

1. In an open simulation, select **Mesh > Import Turbo Blades...** in the menu bar.



The *Import turbo wizard geometry file options* dialog opens.



Defining the Number of Blade Rows

The turbo wizard can generate a mesh for multiple blade rows, and the first step in using the turbo wizard is to specify the number of blade rows in the model.

For cases involving multiple blade rows, the turbo wizard assumes the convention indicated in the diagram below, where blade row 1 is the inlet, and subsequent blade rows are numbered incrementally.



1. The number of blade rows is specified by increasing or decreasing the integer value for the *Number of blade rows* property at the top of the turbo wizard dialog.

   The default value is 1.



2. Mesh data is entered separately for each blade row, so the next step is to indicate for which blade row the first set of data is entered.
3. The blade row is selected using the *Select current blade row* property.

   For example, setting the value to 2 allows you to enter data for blade row 2.



4. To switch to a different blade row, change the value of the *Select current blade row* property.

Data entered in the *Geometry*, *Meshing*, *Flowpath*, and *Grid Settings* tabs, as well as the rotation option, is specific to each blade row.

Defining Rotation

Rotation is defined per blade row.

Rotation is defined per blade row and is enabled by activating the **Current blade row is rotating** checkbox.



When the **Current blade row is rotating** checkbox has been activated, the following fields are made available:

- Rotational speed of the blade row, in `rpm`
- **Hub surface is rotating**—Activate this checkbox to indicate that the hub is rotating with the blade.
- **Shroud surface is rotating**—Activate this checkbox to indicate that the shroud is rotating with the blade.

These values help the turbo wizard define the sliding interfaces correctly in the simulation.

Default turbulence models are activated automatically when a blade row is specified as rotating.

See Also

Defining the Number of Blade Rows
Setting Up the Boundary Conditions

Defining the Flow Axis

The flow axis defines the direction of flow over the blade.

For impeller blades, this direction is at the inlet.



The **Select Flow Axis** menu is used to specify the **X**, **Y**, or **Z** axis as the flow axis. Use this menu only if you want to define a flow axis other than the default position (flow direction +z).

To apply your selection in this menu, make sure that you activate the checkbox **File geometry in actual location**. By default this option is deactivated.

Selecting Mesh Arrangement

The **Mesh Arrangement** option indicates whether the blades are arranged around an annulus or in a linear pattern.

This setting influences the mesh generated by the turbo wizard and changes the options available when setting up the blade geometry.



Annular

In an annular mesh arrangement, the blade geometry is rotated around an axis. This is the default option and is suitable for a typical turbo wizard analysis. In the following screenshot, a transform has been used to show the full rotation of the blade geometry (blue).



This option can be specified by selecting **Annular** from the **Mesh Arrangement** menu. This is the default option and has an additional setting:

- **Number of blades around annulus**

  This option indicates the total number of blades around the annulus and different values can be applied to each row. For an impeller-splitter case, this number refers to the total number of impeller-splitter pairs around the annulus.



Linear

In a linear mesh arrangement, blades are laid out side-by-side in a linear pattern. The example below shows a single blade in the linear mesh arrangement; a transform has been used to shown several repetitions of the blade geometry (blue).



This option is specified by selecting **Linear** from the **Mesh Arrangement** menu, which activates two more settings:

- **Select Blade Axis**

  This is the direction in which the blade pattern is repeated. In the example above, the y axis is the blade axis.

- **Linear Spacing Between Blades**

  This distance is measured from a point on one blade to the same point on the next blade (either in the repeating pattern or between blade rows). It can also be thought of as the width of the flow domain extending laterally from the first blade. The same spacing distance is applied to all blade rows.



See Also

Setting Up the Blade Geometry

Setting Up the Blade Geometry

The *Geometry* tab is used to specify the source, type, and characteristics of the blade geometry.



- The Blade Geometry Definition/Surface Option determines the source of the blade geometry, and offers the choice of importing the blade definition from an external file, or using an existing surface representation.
- The Select Blade Grid Type option identifies the blade type, and influences the shape of the grid generated around the blade.

The remaining options indicate the characteristics of the blade geometry:

- **Use open bottom of blade surface as hub section**

  If the blade definition does not intersect the hub, or if it just touches the hub, the turbo wizard is not able to calculate a section at that point and mesh generation fails. You should check this box unless the blade definition extends past the hub.

- **Use open top of blade surface as shroud section**

  If the blade definition does not intersect the shroud, or if it just touches the shroud, the turbo wizard is not able to calculate a section at that point and mesh generation fails. You should check this box unless the blade definition extends past the shroud.

- **Number of blades around annulus** or **Linear spacing between blades**

  This option depends on whether the mesh arrangement has been specified as annular or linear.

- **Blade Geometry Units**

  This option identifies the units in which the blade geometry was defined. The units must be the same for each import file, and must match those used in the simulation if the blade geometry is defined by an existing surface representation. The choice of units affects the units for the streamline spacing options in the *Flowpath* tab.



See Also

Selecting Mesh Arrangement

Defining the Blade Geometry Source

The source of the blade geometry for each blade row can be either an external file that you create, or an existing surface representation in STAR-CCM+.

Each option is described below.

Importing Geometry from an External File

Blade geometry can be imported from an external .estg file by selecting **Import blade geometry from external file** for the **Blade Geometry Definition/Surface Option**.



When you define a flow axis other than the default position (flow direction +z), make sure that you activate the checkbox **File geometry in actual location**. By default this option is deactivated.

To select the `.estg` file, click the **Browse...** button next to the **Blade geometry file for blade row x** option. This action opens an operating system-specific standard dialog that can be used to locate and open the file.



Using an Existing Surface Representation

Blade geometry defined in an existing surface representation can be used by selecting **Use surface representation from STAR-CCM+** for the **Blade Geometry Definition/Surface Option**. A suitable surface representation can be obtained by importing the blade surface data into STAR-CCM+. Each blade must be defined in a separate region, so once the surface data has been imported some region manipulation may be required.

The region in which the blade is defined is selected from those listed in the Select Surface Representation for Blade Row x From Region List box.



Consider the following when choosing this approach:

- Selection is on a region-wise basis, rather than a boundary-wise one. So all surfaces in the region are used as the airfoil geometry.
- Only the airfoil surface itself should be in the region. If the hub/shroud section areas are closed surfaces, the surfaces should be removed in order to allow the turbo wizard to determine the sections properly at these locations. That is, there should be an open hole at the bottom (hub) and top (shroud) locations of the airfoil surface.
- If you plan on meshing a radial blade or capped Open-TE C-grid blade, these grid types require that two (2) points at the trailing edge location of the airfoil are segments (that is, corner points). This can be done when using surface representations by creating feature curves (edges) which lie along the corner edges in these blade types.
- Feature curves should only be on these segmented corner locations. If they are anywhere else on the model, the turbo wizard understands that those locations should also be segmented.

See Also

Defining the Number of Blade Rows
Understanding the .estg File Format
Defining the Flow Axis

Selecting the Blade Type

The **Select Blade Grid Type** option indicates the type of blade defined in the source file or representation, and specifies the corresponding grid that is generated by the turbo wizard.

For each blade grid type, the resulting mesh is characterized by the pattern of grid lines around the blade. The options available are described below.



O-Grid

The O-Grid is suited to meshing axial blades that have tapered leading and trailing edges. It generates closed loops of grid lines around the blade, “O”, and an orthogonal grid upstream and downstream of the blade, known as an “H”-grid.



Capped Open (TE) C-Grid

The Capped Open (TE) C-Grid is useful for meshing blades that have a blunt trailing edge. For example:

It generates a C-Grid around the leading edge of the blade, and a fine orthogonal grid projecting from the trailing edge of the blade.



Provided below is a close-up of the orthogonal grid from the trailing edge of the blade.



To successfully create a Capped Open (TE) C-Grid blade, ensure you use feature curves to mark the corners of the blunt trailing edge.

Impeller Blade

This option should be selected if the blade is an impeller. The impeller grid takes into account the blunt trailing edge, and the mesh generated around the blade has characteristic C-shape grid lines. An H-grid is generated upstream and downstream of the blade.



Impeller Blade with Splitter Blades

This option allows a grid to be generated for an impeller blade and a splitter blade. Blade geometry for the impeller and splitter must be specified in separate `.estg` files or regions, depending on the selected blade geometry source. The **Select Blade Geometry Designation** option is used to switch between the impeller blade and the splitter blade, allowing the geometry for each blade to be selected separately.



It is only possible to model a single splitter blade, so the **Number Of Splitter Blades** value is fixed at `1`.

Flat Plate

This option should be selected if the blade is a flat plate (the leading and trailing edges are blunt). An example is provided below:

An orthogonal grid is projected from the blunt edges of the blade.



Round Trailing Edge Radial Grid

This mesh works the same as the impeller grid type, except that the mesher distributes points along the curve that defines the geometric section.

See Also

Defining the Blade Geometry Source

Selecting the Mesh Generation Method

The options in the *Meshing* tab determine which mesh generation method is used by the turbo wizard to generate the mesh.

The two options available are:

- Use the turbo wizard to automatically generate an elliptical grid.
- Use the turbo wizard to generate a closed surface mesh that can be meshed using the volume meshing models available in STAR-CCM+.

The **Mesh generation method** is applied to all blade rows.

Generating an Elliptical Grid Volume Mesh

The default **Mesh generation method** in the turbo wizard is **Build 3D volume mesh using elliptical gridder**.

This option results in the generation of a volume mesh with an O-grid or C-grid around the blade, depending on the blade type, and an H-grid upstream and downstream of the blade. There are no additional settings required for this option.



See Also

Selecting the Blade Type

Generating a Closed Surface Mesh

The turbo wizard can be used to generate a closed surface mesh for the blade geometry.

From this, a volume mesh can be generated outside of the turbo wizard, using one of the standard volume meshing models available in STAR-CCM+. This technique is useful when the blade geometry forms part of a larger model, as it allows entire flow domain to be meshed with the same shape cells, for example polyhedral.

This option is accessed by setting the **Mesh generation method** to **Create fully enclosed surface and use automatic meshers**. The volume meshing models are selected from the *Volume Mesh* and *Optional Meshing Models* group boxes, and when the surface mesh is generated the turbo wizard automatically sets up a mesh continuum with the selected models that can be used later to generate the volume mesh.



When this option is selected, the turbo wizard does not generate a volume mesh. The volume mesh must be generated outside of the turbo wizard, and the parameters of the volume mesh should be entered in the usual way.

Defining the Flow Path

The flow path geometry is a two-dimensional view of the machine showing the boundaries of the flow in the radial-axial plane.

This path is defined by four cubic splines:

- Hub
- Tip
- Inlet
- Exit



These splines must be distinct (one cannot serve as both hub and inlet for example) and are defined in `.bndy` files. They need not meet at their endpoints, but must certainly cross in order to define a closed, planar area.

In the **Flowpath** tab, the full flow path is defined, and a clearance grid at the tip or hub can also be specified.



For the geometry to be valid all geometry splines must be defined:

- **Inlet/Exit boundary file for blade row X**

  The `.bndy` file for the inlet/exit of each row. Typically the spline data for the exit of one row are the same as the inlet data of the next

row. You can use the same `.bndy` file for each but they must both be specified in each row.

The inlet and exit boundary definitions must go radially outward, that is from hub to shroud.

- **Hub/Shroud boundary file**

  The `.bndy` files for the hub and shroud. These files define the splines for the entire machine, and not for each blade row.

  The hub and shroud boundary definitions must go axially outward, that is from inlet to exit.

In the turbo wizard, the term streamlines refers to a series of splines used to define surfaces that subdivide the flow volume in the radial direction. The mesh is defined such that each streamline (or stream surface) lies on a cell boundary, so specifying the number of streamlines also determines the radial size of the mesh. Each streamline is marked with a red dot in the diagram below.



- **Number of streamlines**

  This specifies how many streamlines are generated from the hub to the shroud (including the hub and shroud streamlines). If a clearance grid is specified, this number also includes the streamlines in the clearance grid.

- **Initial hub/shroud streamline spacing**

  The spacing parameters, shown above, allow you to cluster the streamlines near the hub and shroud walls, resulting in larger spacing between streamlines in the middle of the flow path. The turbo wizard uses the spacing values as guidelines and modifies them based on the blade span and the number of streamlines. It is

not possible to generate streamlines that are clustered around the middle of the flow path, instead the turbo wizard generates streamlines that are uniformly spaced.

See Also

Understanding the .bndy File Format

Specifying a Clearance Grid

For each blade row, a clearance grid can be generated in order to model the blade tip clearance at the hub or shroud.

It is generated into the flow domain from the hub or shroud, and ends at a streamline defined by either a fixed distance from the hub or shroud, or an imported `.bndy` file. The streamlines used by the clearance grid are taken from the total number of streamlines defined in the Flowpath tab, and are uniformly spaced.



A clearance grid can be specified by using the drop-down menu for the **Generate clearance grid for blade row x** option and selecting either **Hub clearance grid** or **Shroud clearance grid**. The clearance grid can be deactivated by selecting **No clearance grid**.

Defining the Clearance Streamline

The clearance streamline limits the extent of the clearance grid. It is counted as one of the total number of streamlines in the flow path, and is defined either by setting a fixed distance from the shroud or hub, or by a `.bndy` file.

To define the clearance streamline by a specifying a fixed distance, select **Define clearance streamline by: Setting fixed distance value**. This option causes the **Clearance distance** box to appear, which can be used to set the distance of the streamline from the hub or shroud. The turbo wizard uses this value as a guideline to generate the clearance streamline.



To define the clearance streamline by importing a `.bndy` file, select **Import clearance boundary file** from the **Define clearance streamline by:** menu. This option causes the **Clearance boundary file** option to appear. The `.bndy` file can be selected by clicking the **Browse...** button and selecting the relevant file.

Setting the Number of Clearance Streamlines

The **Number of Clearance Streamlines** indicates how many of the total number of streamlines are in the clearance grid. This number includes the hub or shroud streamline and the clearance streamline. The minimum number of streamlines in the clearance grid is `5`.



See Also

Defining the Number of Blade Rows
Understanding the .bndy File Format

Controlling the Grid Settings

The *Grid Settings* tab can be used to control the parameters that define the mesh.

They are accessed through three subtabs:

- Distribution

  These settings control the spacing and clustering of the grid around the blade.

- Number of points

  These settings control the dimensions of the mesh in terms of grid points.

- Block boundary settings

  These settings control the relative axial size of the three grid blocks.

Defining the Grid Distribution

The settings in the *Distribution* subtab can be used to control several characteristic cell lengths of the mesh.

The characteristic cell lengths of the mesh are shown in the diagram below. This diagram can also be accessed from within the turbo wizard by clicking the **Parameter Information** button in the *Distribution* subtab.



- **Initial Clustering, Leading Edge (DS_LE)**—the minimum cell length parallel to the blade surface at the leading edge.
- **Initial Clustering, Trailing Edge (DS_TE)**—the minimum cell length parallel to the blade surface at the trailing edge.
- **Near Wall (Airfoil) Spacing (DS_IN)**—the minimum cell length perpendicular to the blade surface for cells adjacent to the blade.
- **Periodic Boundary Spacing (DS_OUT)**—the minimum cell length perpendicular to the blade surface for cells on the outer surface of the mesh segment surrounding the blade.

A custom value for any of these parameters can be specified by activating the **User Defined** checkbox for the relevant parameter. This causes a text box to appear in which the new value can be entered. The units used in the turbo wizard are defined in the Geometry tab.



The last option in the *Distribution* tab is:

- **Allow turbo wizard to move airfoil surface points to preserve grid orthogonality**

  This option adjusts the position of points on the blade surface to ensure that grid lines meet the blade surface perpendicularly. The shape of the blade is not changed. This option is only available for O-grid blade types.

When this option is selected, you are able to specify the **Number of Dynamic Orthogonality Iterations**. The default value is `8`, and this should only be changed if you are an advanced user of the turbo wizard. By increasing this value you can improve the quality of the mesh at the blade surface, however the mesh generation process is much slower. The maximum applicable value is approximately `40`.



See Also

Setting Up the Blade Geometry
Selecting the Blade Type

Defining the Number of Grid Points

The settings in the *Number Of Points* subtab determine the dimensions of the mesh in terms of grid points (or number of cell boundaries).

Mesh sizes can be specified for each section of the mesh, for example in the upstream H-grid. The options available are shown in the diagram below, which can also be accessed from within the turbo wizard by clicking the **Parameter Information** button in the *Number Of Points* subtab.



- **Upstream Cap, Axial Points (IX1)**—the mesh dimension in the flow direction for the upstream H-grid block.
- **Points Around Airfoil Surface (IX2)**—the mesh dimension around the blade for the O-grid or C-grid block.
- **DownStream Cap, Axial Points (IX3)**—the mesh dimension in the flow direction for the downstream H-grid block.
- **Upstream Cap, Circumferential Points (JY1)**—the tangential mesh dimension for the upstream H-grid block.
- **Points Between Airfoil Surface and Periodic Surface (JY2)**—the mesh-to-boundary dimension for the O-grid block.
- **Downstream Cap, Circumferential Points (JY3)**—the tangential mesh dimension for the downstream H-grid block.

A custom value for any of these parameters can be specified by activating the **User Defined** checkbox for the relevant parameter. This causes a text box to appear in which the new value can be entered.



The **Trailing Edge Width Points** option replaces **Downstream Cap, Circumferential Points** when an impeller blade type is selected. In such cases, the dimensions of the mesh in the downstream H-grid are specified by this option and the **Points Between Airfoil Surface And Periodic Surface** option, as shown below.



The **Points Between Camber Line and Airfoil Surface** option becomes available in the *Number Of Points* tab when a clearance grid has been specified. This option specifies the mesh dimension across the thickness of the blade, as shown below. The turbo wizard uses this value as a

guideline, and it can be altered depending on the dimensions used for the rest of the mesh.

The options available for the Flat Plate grid type are shown in the diagram below, which can also be accessed from within the turbo wizard by clicking the **Parameter Information** button in the *Number Of Points* subtab.



- **Upstream Cap, Axial Points (IX1)**—the mesh dimension in the flow direction for the upstream H-grid block.
- **Points On Side Of Flat Plate (IX2)**—the mesh dimension in the flow direction for the side of the flat plate.
- **Front Airfoil Block, Axial Points (IX4)**—the mesh dimension in the flow direction for the leading edge of the flat plate.
- **Back Airfoil Block, Axial Points (IX5)**—the mesh dimension in the flow direction for the trailing edge of the flat plate.
- **Downstream Cap, Axial Points (IX6)**—the mesh dimension in the flow direction for the downstream H-grid block.
- **Top Airfoil Block, Circumferential Points (JY2)**—the tangential mesh dimension for the top of the flat plate.

- **Bottom Airfoil Block, Circumferential Points (JY3)**—the tangential mesh dimension for the bottom of the flat plate.
- **Front Airfoil Block, Circumferential Points (JY4)**—the tangential mesh dimension for the leading edge of the flat plate.
- **Back Airfoil Block, Circumferential Points (JY5)**—the tangential mesh dimension for the trailing edge of the flat plate.

See Also

Selecting the Blade Type
Specifying a Clearance Grid

Defining the Grid Block Boundary Settings

The *Block Boundary Settings* subtab is used to control the relative axial size of the three grid blocks: upstream H-grid, O-grid, or C-grid around the blade, and the downstream H-grid.

These relative sizes are specified by two spacing ratios, shown below. This diagram can also be accessed from within the turbo wizard by clicking the **Parameter Information** button in the *Block Boundary Settings* subtab. When an impeller blade type is specified the **Downstream Cap, Outlet Boundary Spacing Ratio** option is unavailable as it cannot be modified for a C-grid.



UPSTREAM CAP
- INLET BOUNDARY SPACING RATIO (ZINO) = $\frac{B}{A+B}$

DOWNSTREAM CAP
- OUTLET BOUNDARY SPACING RATIO (ZEXO) = $\frac{C}{C+D}$

A custom value for any of these parameters can be specified by activating the **User Defined** checkbox for the relevant parameter. This causes a text box to appear in which the new value can be entered.



See Also

Selecting the Blade Type

Setting Up the Boundary Conditions

The turbo wizard automatically prepares the boundaries, interfaces, and physics models necessary for the simulation.

In the *Boundary Conditions / Turbulence* tab the boundary conditions for the inlet and outlet can be specified, and these are passed on to the relevant objects in the simulation tree. Activating the **Automatically set-up default turbulence models?** checkbox implements the default turbulence models, which are shown in the tab. If a rotation has been specified for the current blade row, this option is already selected and is grayed out.





If the default values in this tab are left unchanged, the turbo wizard prepares the simulation using these values.

See Also

Defining Rotation
Defining the Number of Blade Rows

Setting Up the Initial Conditions

The turbo wizard automatically prepares the boundaries, interfaces, and physics models necessary for the simulation.

In the *Initial Conditions* tab, the initial conditions for the simulation can be specified. If the default turbulence models have been implemented in the boundary conditions / turbulence tab, the initial turbulence intensity and viscosity ratio can also be specified.



If the default values in this tab are left unchanged, the turbo wizard prepares the simulation using these values.

See Also

Setting Up the Boundary Conditions

Working with Turbomachinery Input Files

The turbo wizard uses input files of the formats `.bndy`, `.estg`, and `.trbw`.

The first two you prepare. The most common means for creating these files is a typical spreadsheet program. However, it is also possible to write code in Fortran for this purpose.

Understanding the .bndy File Format

This section describes the `.bndy` file format.

Each spline of the flow path geometry is defined by its (r, z) coordinates in a cylindrical coordinate system. The flow is assumed to move in the +z direction. To make viewing the model more convenient, this flow path is constructed on the q = 0 plane.

The flow path is defined by reading in vertex locations from an external `.bndy` file an example of which is shown below:

```
31  (Centimeters)
-22.860001 -17.779999 -12.700000 -7.620000
-3.070860  -2.435860 -1.800860  -1.165860
-0.264160   0.101600   0.736600  1.371600
2.006600   2.644140  3.279140   3.914140
4.521200   5.153660   5.788660  6.423660
7.058660   7.693660  8.328660   8.963660
9.598660  10.233660  10.868659 11.503659
12.138659  12.646660 15.240000
17.525999 17.525999 17.525999 17.525999
17.525999 17.541241 17.571720 17.602200
17.739361 17.805401 17.957800 18.122900
18.298161 18.440401 18.567400 18.674080
18.768059 18.867119 18.948400 19.029680
19.110960 19.177000 19.235420 19.293840
19.329399 19.364960 19.390360 19.405600
19.420839 19.433540 19.433540
```

The format of this file is:

- Line 1 - the number of points in the file and the units
- The rest of the lines are a space and carriage return separated list of numbers. In the example above, the first 31 points are the z positions of the spline and the second set of 31 numbers is the radial coordinate.

Understanding the .estg File Format

The `.estg` file is an external text file that contains the x, y, and z coordinates of the blade profile at each of two or more sections.

By default, the flow is assumed to move in the +z direction. An example follows:

```
#
# Input File: EZTurbo
#
# Original Data File: From NASA
# Modifications     : None
# Date              : 05 September 2000
#
# Case: Rotor 37
#
#       NPTS   NBLD CW  RPM
#       ====  ===== ==  ===
ROTOR    198    36   T    0
NPTS 198
AIRFOIL   1
    columns do not mean anything, first 198 points z,
then x, y
   0.0281091   0.0289117   0.0298336   0.0308953
   0.0321196   0.0335267   0.0351471   0.0370063
   0.0391449   0.0415985   0.0444127   0.0476459
   0.0513517   0.0555958   0.0604470   0.0659739
   0.0722525   0.0793617   0.0873903   0.0964298
   0.1065792   0.1179478   0.1306549   0.1440579
   0.1589214   0.1753774   0.1935733   0.2136664
   0.2358321   0.2602583   0.2886035   0.3199688
   ...
AIRFOIL   13
   0.6383425   0.6388631   0.6394575   0.6401458
   0.6409383   0.6418501   0.6428990   0.6441055
   0.6454872   0.6470772   0.6489009   0.6509911
   0.6533939   0.6561421   0.6592839   0.6628652
   ...
```

Any line that does not begin with a recognized keyword (NPTS, AIRFOIL) or space is ignored by the importer. So in the preceding example all the lines beginning with `#` and the `ROTOR` lines are ignored. The different lines are:

- `NPTS` - the number of points per section. Each section must have the same number of points.
- `AIRFOIL` - the section number starting with the hub section and ending at the shroud. The importer reads the file until the end, there is no need to specify up front how many sections are in the file.

- `<SPACE>` - a data line. The data is simply a series of space or line feed separated numbers. In the above example, the first 198 points are the z values, the next 198 points are the x values and the next 198 points are the y values of the set of x, y, z coordinates defining the profile. There should be `(NPTS * 3)` values in each section.

Understanding the .trbw File Format

The `.trbw` file is a STAR-CCM+-generated non-compressed XML file which stores all of the turbo wizard user interface settings in an external file.

This file can be generated by clicking the **Save Turbo Wizard Settings** button.

To restore the saved settings in the turbo wizard, simply click the **Load Turbo Wizard Settings** button and navigate to the `.trbw` file. Opening the file loads it back into the turbo wizard interface.

Caution: In typical situations, you should not attempt to edit the XML file, as this might corrupt it. The file is fairly descriptive in what parameters are being stored and saved, and so you can modify this file at your own risk.

Saving and Loading Turbo Wizard Data

Settings and values entered in the turbo wizard dialog are referred to as turbo wizard data.

It is useful to store this data so that it can be used in later simulations, and this is possible by saving it in a `.trbw` file. This file format is specific to the turbo wizard, and files of this type can be loaded into the wizard as described below. The options to save and load turbo wizard data are only accessible if the turbo wizard has been activated.

Note: Saving and loading Turbo Wizard Data is not currently recorded by user macros in STAR-CCM+.

Saving Turbo Wizard Settings

To save the turbo wizard data for use with other turbomachinery simulations:

1. Click the **Save Turbo Wizard Settings** button at the bottom of the *Import turbo wizard geometry file options* dialog.



   This opens the *Save Turbo Wizard Data* dialog with the Files of Type filter set to `*.trbw` format.
2. To save the data, navigate to the preferred location, enter an appropriate file name and click on the **Save** button.



See Also

Activating the Turbo Wizard

Loading Turbo Wizard Settings

To import turbo wizard data into the current simulation:

1. Click the **Load Turbo Wizard Settings** button.

   This opens the *Load Turbo Wizard Data* dialog.
2. To load the data navigate to the appropriate location, select the turbo wizard data file and click the **Open** button.


Load Turbo Wizard Data
File selection:
/simulations/turbo
Save In:
turbo
turboData1.trbw
turboData2.trbw
turboData3.trbw
File Name:
turboData1.trbw
Files of Type:
Turbo Wizard Data (*.trbw)
Open
Cancel

Using the Fire and Smoke Wizard

This section describes the fire and smoke wizard feature of STAR-CCM+.

This section explains the following:

- What is the fire and smoke wizard?
- Other capabilities of STAR-CCM+ that can be used with the fire and smoke wizard for problems of this type
- Setting up the fire properties

What is the Fire and Smoke Wizard?

The fire and smoke wizard is an automatic simulation tool for simulating a building fire.

Its capabilities for this application include :

- Fire regions and sources
- Soot sources
- Velocity inlet boundaries for ventilation
- Fire doors
- Heat and smoke detectors

STAR-CCM+ Capabilities Related to Fire and Smoke

This section lists some of the other STAR-CCM+ features that can help model fire and smoke.

The fire and smoke wizard is the primary tool for these types of problems, but there are several other STAR-CCM+ features that can help with the CFD solution:

- Modeling Options:
  - Stationary
  - Three-Dimensional
- Radiation Models:
  - Discrete Ordinates Method (DOM)
- Turbulence Models:
  - Spalart-Allmaras
  - K-Epsilon, K-Omega...
- Boundary Conditions:
  - Velocity Inlet
  - Stagnation
  - Pressure
- Solvers:
  - Second-Order
  - Segregated
  - Algebraic Multigrid Acceleration
- Reports, monitors, and plots
- Derived parts

Setting Up the Fire Properties

This section outlines the process for setting up the fire properties.

This feature would be used after you generate a mesh. The process involves the following:

- Activating the wizard
- Creating fire volumes
- Setting up fire dynamics
- Specifying fire doors
- Setting up ventilation
- Preparing the soot source
- Adding heat and smoke detectors
- Specifying optical properties
- Using post-processing planes
- Setting graph parameters

Activating the Fire and Smoke Wizard

The properties of the fire are set using the **Fire and Smoke Wizard**.

To begin accessing this feature:

1. Select **Windows** > **Toolbars** > **Customize** from the menu bar.



This opens the *Customize Toolbars* dialog as shown in the following screenshot.



2. Drag and drop the **Fire and Smoke Wizard** icon onto the end of one of the toolbars near the top of the workspace.
3. Click **Close** on the *Customize Toolbars* dialog.
4. Click the (**Fire and Smoke Wizard**) button.

   This opens the *Fire and Smoke Wizard* dialog.



There are two broad categories of detail that this dialog provides:

- Required fire characteristics are provided in the following tabs:
  - *Fire Dynamics*
  - *Soot Source*
  - *Optical Properties*
- Optional characteristics of the fire and building are provided in the following tabs:
  - *Create Fire Volume*
  - *Fire Barrier*
  - *Ventilation*
  - *Detectors*

Post-Processing Planes let you display data in precisely defined sections and output that data to images.

5. Finally, after you click **Apply**, the wizard automatically saves your settings to an external `.fsmk` file.

   This allows you to import those settings into another fire and smoke simulation.

6. To import the file, click the **Load FireSmoke** button near the bottom of the dialog.



This activates a standard file-opening dialog with the Files of Type filter set to the `.fsmk` format.



7. Navigate to the file and click **Open**.
8. Additionally, before you click **Close** in the wizard, you can save the file with a name and path of your choice by clicking the **Save FireSmoke** button.

   This opens a *Save* dialog.



9. Navigate to the preferred location, enter a file name, and click **Save**.

Creating Fire Volumes

The *Create Fire Volume* tab allows you to define a space for the fire that is an alternative to the regions created in the mesh generation process.



Begin by entering a name for the fire volume in the **Fire Volume Name** text box.

Next select the volume creation method from the drop-down list:

- A bounding box with minimum and maximum coordinates
- The absolute size with a midpoint and size bounds

The fields below the drop-down list change with your selection of method. Enter the appropriate coordinates.

Once you enter values, you can create the volume by clicking the **Create Fire Source** button.

The fire volume that you create appears with its pertinent data in the Created Fire Volumes list.



Once you finish with the wizard by clicking **Apply**, the fire volume also appears in the object tree within the **Regions** node.



It is also possible to manage the fire volumes in the list. For example, if you wish to create a second fire volume similar to the first, you need not repeat the entire creation process. Simply select the name of the fire volume in the list and click **Copy Fire Source**. An additional fire volume appears in the list. Select that copy, make any changes you wish, and then click **Edit Fire Source** to apply your changes. You can also delete fire volumes by selecting them in the list and clicking **Delete Fire Source**.

Setting Up Fire Dynamics

Set up the dynamic properties of the fire by selecting the *Fire Dynamics* tab.



The *Turn On Fire Regions* list in the upper left lets you tick the checkboxes of regions that represent the fire, including any fire volumes that you may have created in the previous tab.

Select the pressure boundary in the list in the upper right.

If you wish to use an N-spread model, then tick the checkbox.

This makes the fire sources appear in the list, and you can use the buttons to the right of it to change the sequence. Simply select a fire source in the list and click the button to move it up or down within the list.



Select a fire volume from the drop-down list, and do as follows:

- Specify the maximum simulation time for it.
- Specify the values for alpha, p, and the HRR start time.

- Use the radio buttons to select whether you want to specify the limit for Q or for t, and then enter the value.
- Select another fire volume from the list, if applicable, and repeat the above steps.

The fire and smoke wizard retains your entries for each of the fire volumes.

This wizard also sets up a monitor plot of the heat release rate, which is previewed in the panel.



To control the parameters of this plot, click the **Set Graph Parameters** button.

You can also tick a checkbox to use the DOM (Discrete Ordinates Method) radiation model.

Specifying Fire Barriers

If a fire barrier, such as a fire door, is one of the regions created in your mesh, you can select it in this panel.



You can specify the activation and deactivation times of each fire barrier that you select in the drop-down list.

Setting Up Ventilation

You can select a velocity inlet boundary for ventilation, and then specify ventilation properties.



Begin by selecting a velocity inlet boundary from the list in the upper left.

If you are using multiple velocity inlet boundaries, use the drop-down list **Select Velocity Inlet Boundary For Ventilation** to select each one, and do as follows:

- Decide whether the ventilation direction is blowing or extraction, and select that option from the drop-down list.
- Enter values for initial ventilation rate, activation time, end-of-ventilation ramping, and continuation at maximum ventilation.
- Select another velocity inlet boundary from the list, if applicable, and repeat the above steps.

The fire and smoke wizard retains your entries for each of the velocity inlet boundaries.

This wizard also sets up a monitor plot of the heat release rate, which is previewed in the panel. To control the parameters of this plot, click the **Set Graph Parameters** button.

See Also

Setting Graph Parameters

Preparing the Soot Source

The Soot Source tab lets you specify the soot yield and heat of combustion.

Enter values, or accept the defaults, for the soot yield and heat of combustion.



Adding Heat and Smoke Detectors

The *Detectors* tab lets you incorporate heat detectors or smoke detectors into your fire and smoke simulation.



To incorporate heat detectors or smoke detectors into your fire and smoke simulation:

1. Begin by selecting the type of detector from the **Create Heat Or Smoke Detector** drop-down list.

   The settings for a particular type of detector change with your selection.
2. Enter a name for the detector in the **Detector Name** text box.
3. Enter the coordinates of the detector in the *Detector Position* group box.
4. Enter the settings that are specific to that type of detector:
   - For a heat detector, enter the activation temperature and response time index as shown in the preceding screenshot.
   - For a smoke detector, enter the optical density and characteristic length.



5. Once you finish entering values, you can create the detector by clicking the **Create Detector** button.

   The new detector appears with its pertinent data in the List of Heat and Smoke Detectors.



   When the list contains one or more detectors, an additional set of options becomes available.



The settings are specific to each detector that you highlight in the detector list.

6. After you finish with the wizard by clicking **Apply**, a point derived part is created for each detector and appears in the object tree in the **Derived Parts** node.



Additionally, a report, monitor, and plot is created for each detector. These appear in the object tree in the **Reports**, **Monitors**, and **Plots** nodes respectively.



For heat detectors, the report monitors the maximum temperature at the probe location. The corresponding plot displays the temperature against iteration number.



The report created for smoke detectors monitors the soot mass fraction at the probe location.

Specifying Optical Properties

The Optical Properties tab lets you set the specific extinction coefficient.

To specify optical properties:

1. Enter a value, or accept the default, for the specific extinction coefficient.

Using Post Processing Planes

The Post-Processing Planes tab lets you create plane sections.

You can create plane sections with this tab of the wizard.



Begin by entering a name for the plane section in the **Section Name** text box.

Make your selection of scalar values to display by ticking the appropriate checkbox:

- **Temperature**
- **Extinction Coefficient**
- **Soot Density**
- **Visibility Light-Emitting**
- **Visibility Light-Reflecting**

Enter the appropriate coordinates for the origin and normal.

Once you have finished, you can create the plane section by clicking the **Create Section** button. The name of the section appears in the Created Sections list.



The plane section also appears in the object tree within the **Derived Parts** node.



It is also possible to manage the plane sections in the list:

- To change a section, select its name in the list, make changes, and click **Modify Section**.
- To create another section similar to an existing one, select the name of the section in the list, make changes to it, give it a new name and click **Create Section**. An additional section appears in the list. If you did not name it, a number `2` is added to the original section's name by default.
- To delete a section, select it in the list and click **Delete Section**.

The **Update Frequency** setting lets you control how often the scene containing the data should be updated during the solver run. The default setting is `1`, or each iteration.

Finally, you can select a file format for output of the displayed data to an image file.



Setting Graph Parameters

This feature in the fire and smoke wizard lets you customize the monitor plots that it automatically generates.

To control the parameters of a plot:

1. Click the **Set Graph Parameters** button in the *Fire Dynamics* or *Ventilation* tabs of the wizard.

   This opens the *Graph Parameters* dialog.



2. Enter your specifications and click **OK**.

Using the Thermal Comfort Wizard

This section explains the thermal comfort feature of STAR-CCM+.

- What is the thermal comfort model?
- What is the Thermophysiological Occupant Model?
- What is the thermal comfort wizard?
- What are TCM Passengers?
- What are TCM Boundaries?
- What is the thermal comfort solver?
- Other capabilities of STAR-CCM+ that can be used with the thermal comfort wizard for problems of this type
- Setting up a thermal comfort case
- Performing a TCM simulation
- Post-processing a TCM simulation
- Information on working with TCM input files
- Information on working with TCM output files

What is the Thermal Comfort Model?

The thermal comfort model enhances simulation of passenger comfort inside a vehicle (typically an automobile).

Its capabilities for this application include :

- Passenger characteristics—height, clothing resistance, and metabolic rate
- Relative humidity
- Insulating effects of the vehicle walls
- Radiation and convection for the dashboard and other surfaces exposed to the vehicle powertrain
- Air recirculation, which uses the air humidity to evaluate the energy loads on the air conditioning system. There are certain requirements and considerations that you should note when running air recirculation simulations.

The model provides these capabilities by allowing STAR-CCM+ to work with a special one-dimensional code called TIM (German acronym for "Thermophysiological Occupant Model"). TIM models and quantifies the thermophysiological state of a human passenger seated inside a car, or other enclosed space.

The thermal comfort model is represented with a node by that name among the physics continuum models. Its node is subject to common menu items, and appears when you activate the model.



Activating the thermal comfort model provides access to the thermal comfort wizard, which is used to define key simulation settings. Using the thermal comfort model also adds the following nodes to the object tree, which become placeholders for properties defined in the wizard:

- TCM Passengers
- TCM Boundaries
- Thermal Comfort Solver

Notes on Air Recirculation Simulations

The air recirculation feature requires the Multi-Component Gas model, with water vapor defined as the first component and dry air as the second component. You must select the components in this order, as this is how the Thermal Comfort wizard interprets the components.

The temperature and humidity of the recirculated air supplied to the air conditioning system is calculated with the following settings:

- *Percentage of recirculated air* from the Advanced Settings tab.
- *External Humidity* (relative or absolute) from the Advanced Settings tab.
- *External Bulk Temperature* from the External Convection and Radiation tab.

The amount of heat that needs to be removed from the air to match the temperature at the inlet is calculated as the energy load on the air conditioning system. If there are multiple inlets with different temperatures, the lowest temperature is used. In addition, if the temperature of the water vapor in the recirculated air is below the saturation temperature, the energy due to condensation is also calculated.

The total energy load on the air conditioning system is written to the `air_conditioner_load.dat` file after each iteration.

See Also

Activating the Thermal Comfort Model
What is the Thermal Comfort Solver?
What is the Thermal Comfort Wizard?
What are TCM Passengers?
What are TCM Boundaries?
Adding External Convection and Radiation
Using Advanced Settings

What is the Thermophysiological Occupant Model?

The thermophysiological regulation of the human body can be seen as a regulatory circuit consisting of a passive and an active system.

The thermophysiological occupant model is a one-dimensional code that models the influence of ambient conditions in a passenger cabin on the thermal condition of an occupant. This model comprises the passive and the active system. The passive system describes the internal and external heat transfer mechanisms of the human body. The active system describes the thermophysiological regulation of the human body. The model accounts for the influence of clothing and contact areas--such as contact with the seats.

Due to the temperature difference, the human body transfers heat between the skin and the environment. Further, there are internal heat transfer processes resulting from the inhomogeneous temperature distribution within the human body. Under normal conditions, a human releases heat to the environment. In equilibrium, the heat release equals the heat production of the body.

Endotherms use thermophysiological regulation to keep their bodies in a state of thermal equilibrium, regardless of the surrounding environment.

Passive System

Within the thermophysiological occupant model, the passive system captures thermal effects at global and local levels.

In order to capture local effects, the human body is divided into segments and elements. This system has 14 segments and a central vascular system.

Each segment is composed of four concentric telescoping cylinders (the elements). The cylinders correspond to:

- Core

  This element consists to a large extent of skeleton.
- Muscle
- Fat
- Skin



The head is treated as a sphere with four shells. The central vascular system is a separate element. Through blood circulation, the central vascular system is responsible for the heat transport in the axial direction of the body.

The indices that are used in the formulation of this model are:

- j: Numbering of the segments (from 1 through 14)
- i: Numbering of the elements (from 1 through 4)

Body Heat Storage

This section describes the body heat storage model.

The human body stores heat in the different body parts. For a constant specific heat capacity c_p, the stored body heat flux $\dot{Q}_{ST}$ is directly proportional to the temperature gradient $\frac{\partial T}{\partial t}$:

$$\dot{Q}_{ST_{i,j}} = c_{p_{i,j}} \cdot \frac{\partial}{\partial t}\left(T_{i,j}\right) \tag{2853}$$

The stored body heat flux of the central vascular system $\dot{Q}_{ST_{CVS}}$ is given by:

$$\dot{Q}_{ST_{CVS}} = c_{p_{CVS}} \cdot \frac{\partial}{\partial t}\left(T_{CVS}\right) \tag{2854}$$

where c_p is the specific heat capacity.

The specific heat capacity of the body parts is calculated from the mass of the body elements and the specific heat capacities of the body elements:

$$c_{p_{1,j}} = m_{MUS_{1,j}} \cdot c_{p_{MUS}} + m_{BONE_{1,j}} \cdot c_{p_{BONE}} \tag{2855}$$

$$c_{p_{2,j}} = m_{MUS_{2,j}} \cdot c_{p_{MUS}} \tag{2856}$$

$$c_{p_{3,j}} = m_{3,j} \cdot c_{p_{FAT}} \tag{2857}$$

$$c_{p_{4,j}} = m_{4,j} \cdot c_{p_{SKIN}} \tag{2858}$$

The specific heat capacity of the different tissue types are (Meeh; DuBois):

Tissue Type	Specific Heat Capacity (kJ/kg K)
Bones	2093
Muscle	3768
Fat	2512
Skin	3768
Blood	3768

Metabolic Heat Production

This section describes the metabolic heat production model.

The human body produces heat due to biochemical reactions taking place in the body. This heat production is characterized by the base metabolic rate $\dot{Q}_{MET,BAS}$, which is always present to keep the human alive, and a portion of the metabolic rate due to activity. The activity portion consists of the muscular work $\dot{Q}_{MET,WORK}$ and, if applicable, the shivering due to cold $\dot{Q}_{MET,SHIV}$.

When computing the heat production through muscular work, a mechanical efficiency of zero is assumed. This assumption is valid for most types of physical activity. The core, fat, or skin elements do not influence the heat production, because the metabolic work $\dot{Q}_{MET,WORK,TOT}$ takes exclusively place within the muscles. The metabolic work has a local constant $L_{MUS,WORK}$ that is used to distribute the heat production to the individual muscle elements. The equations for computing the metabolic heat production are as follows:

Core:

$$\dot{Q}_{MET_{1,j}} = \dot{Q}_{MET,BAS_{1,j}} \tag{2859}$$

Muscle:

$$\dot{Q}_{MET_{2,j}} = \dot{Q}_{MET,BAS_{2,j}} + \dot{Q}_{MET,SHIV_{2,j}} + \dot{Q}_{MET,WORK,TOT_{2,j}} + L_{MUS,WORK} \tag{2860}$$

Fat:

$$\dot{Q}_{MET_{3,j}} = \dot{Q}_{MET,BAS_{3,j}} \tag{2861}$$

Skin:

$$\dot{Q}_{MET_{4,j}} = \dot{Q}_{MET,BAS_{4,j}} \tag{2862}$$

Internal Heat Transfer

This section describes the internal heat transfer model.

Generally, the heat flux is directed from inside the body to the peripheral regions. This heat transfer is called internal heat transfer. The border to the environment is the skin. The heat transfer is characterized by heat conduction and convection.

Heat conduction between two elements i and i+1 of a segment j is given by:

$$\dot{Q}_{COND_{i,j}} = R_{COND_{i,j}} \cdot \left(T_{i,j} - T_{i+1,j}\right) \tag{2863}$$

The conductivities $R_{COND}(W / K)$ between the different shells are listed in the following table:

Segment	Core - Muscle	Muscle - Fat	Fat - Skin
Head	4.650	9.070	15.350

Segment	Core - Muscle	Muscle - Fat	Fat - Skin
Torso	1.690	6.450	18.260
Upper arm	0.843	2.790	8.373
Forearm	0.843	2.790	8.373
Hand	2.835	5.580	5.580
Thigh	1.833	4.920	17.619
Lower leg	1.222	3.280	11.746
Foot	2.935	10.000	8.085

Heat is transferred between body parts through the transport of blood in a circulatory system. The blood extracts the heat from the warmer regions of the body and releases it to the colder regions.

Heat transfer from one element of a segment to the corresponding element in the next segment is calculated as:

$$\dot{Q}_{BLOOD_{i,j}} = \dot{V}_{BLOOD_{i,j}} \cdot c_{p,BLOOD} \cdot \rho_{BLOOD} \cdot \left(T_{BLOOD} - T_{i,j}\right) \tag{2864}$$

The density of blood is assumed to be 1111 kg/m^3. The volumetric flow rate of blood varies according to the individual segments. Each body element has a base volumetric flow rate $\dot{V}_{BLOOD,BAS}\left(cm^3/s\right)$ that is listed in the following table:

Segment	Core	Muscle	Fat	Skin
Head	13.330	0.075	0.033	0.400
Torso	64.400	1.777	0.638	0.583
Upper arm	0.048	0.086	0.022	0.035
Forearm	0.048	0.086	0.022	0.035
Hand	0.014	0.007	0.007	0.277
Thigh	0.183	0.308	0.067	0.238
Lower leg	0.122	0.206	0.044	0.158
Foot	0.021	0.004	0.011	0.417

In addition to the constant base values, the blood volumetric flow rates of the skin and muscle shell can change due to physical activity or thermoregulation. The blood volumetric flow rates of the core and fat, however, always have the same base value.

External Heat Transfer

This section describes the external heat transfer model.

The human body exchanges heat with the environment through the following heat exchange mechanisms:

- Heat transfer due to convection and radiation
- Heat losses due to diffusion, transpiration, and respiration.

Each heat loss is consolidated to a single value for the global heat balance.

Convection

The heat flux of an undressed skin element of a segment j is computed as:

$$\dot{Q}_{CONV_j} = a_{CONV_j} \cdot \left(T_{4,j} - T_{AIR_j}\right) \tag{2865}$$

where:

$\dot{Q}_{CONV_j}$	Convective heat transfer coefficient of segment j
$T_{4,j}$	Skin temperature of segment j
T_{AIR_j}	Temperature of air

Radiation

Heat transfer due to radiation takes place between the skin surface or the clothes surface of the human body and the surrounding surfaces. By use of the Stefan-Boltzmann law, a radiation heat balance for each skin segment can be formulated as follows:

$$\dot{Q}_{RAD_j} = \epsilon_{SKIN_j} \cdot \sigma \cdot A_j \cdot \left(T_{4,j}^4 - T_{ENV_j}^4\right) \tag{2866}$$

where:

ϵ_{SKIN_j}	Emissivity of skin of segment j
σ	Stefan-Boltzmann constant
A_j	Skin surface of segment j
T_{ENV_j}	Environment temperature of segment j

For temperature differences between the skin and the environment that are smaller than 20 K, the following linear approximation can be used:

$$\dot{Q}_{RAD_j} = a_{RAD_j} \cdot A_j \cdot \left(T_{4,j} - T_{ENV_j}\right) \tag{2867}$$

where a_{RAD_j} is a linearized heat transfer coefficient for radiation.

Using the equations for convection and radiation above, the two exterior heat fluxes are combined to an overall value — the so-called dry heat flux $\dot{Q}_{DRY}$. The heat transfer coefficients for convection and radiation are combined to a joint coefficient a_{DRY_j}:

$$a_{DRY_j} = a_{RAD_j} + a_{CONV_j} \tag{2868}$$

With this dry heat transfer coefficient, an effective temperature T_{EFF_j} is calculated for each skin element. The effective temperature is formed by using the temperature of the ambient air T_{AIR} and the temperature of the surrounding surfaces T_{ENV}:

$$T_{EFF_j} = \frac{a_{CONV_j} \cdot T_{AIR_j} + a_{RAD_j} T_{ENV_j}}{a_{DRY_j}} \tag{2869}$$

The equation for the heat loss due to convection and radiation can now be written as:

$$\dot{Q}_{DRY_j} = a_{DRY_j} \cdot A_j \cdot \left(T_{4,j} - T_{EFF_j}\right) \tag{2870}$$

With the above equation, all dry heat losses are described for the total heat balance.

Respiration

Heat loss due to respiration occurs because of the difference in humidity and temperature between incoming and outgoing air.

The heat loss due to the humidity difference is proportional to the water vapor pressure gradient between the lung and the ambient air. It also depends on the inhaled air volume and, as such, from the metabolic rate. The heat loss due to respiration is given by:

$$\dot{Q}_{RESP,HUM} = 1.752 \cdot 10^{-5} \cdot \dot{Q}_{MET} \cdot \left(p_{SAT,LUNG} - p_{AIR}\right) \tag{2871}$$

The heat loss due to the temperature difference of the inhaled air is given by:

$$\dot{Q}_{RESP,TEMP} = 0.0014 \cdot \left(34 - T_{AIR}\right) \tag{2872}$$

The total heat loss due to respiration is therefore:

$$\dot{Q}_{RESP} = \dot{Q}_{RESP,HUM} + \dot{Q}_{RESP,TEMP} \tag{2873}$$

The total heat loss due to respiration is extracted one half each from the head core and the torso core.

$$\dot{Q}_{RESP_{1,1}} = \frac{1}{2}\dot{Q}_{RESP,HEAD} \tag{2874}$$

$$\dot{Q}_{RESP_{1,2}} = \frac{1}{2}\dot{Q}_{RESP,TORSO} \tag{2875}$$

Diffusion

Because of perspiratio insesibilis, the diffusion of water through the human skin, there is a permanent heat loss $\dot{q}_{DIFF}$ to the environment. According to [Stolwijk], this heat loss can be specified as 5.6 W/m^2 for the entire body. As the heat loss is distributed to all skin elements, an equal rate of 5.6 W/m^2 can be assumed for all skin elements. The heat flux due to diffusion is thus given by:

$$\dot{Q}_{DIFF_j} = \dot{q}_{DIFF} \cdot A_j \tag{2876}$$

where A_j is the area of a segment.

Transpiration

A further form of heat release is the evaporation of sweat on the surface of the skin during transpiration. During transpiration, water vapor transport takes place through the boundary layer. The driving force is the difference between the vapor pressure at the body surface and the vapor pressure of the ambient air. The heat flux due to transpiration is thus given by:

$$\dot{q}_{TRANS} = \beta \cdot \rho \cdot r \cdot \left(X_{SKIN} - X_{AIR}\right) \tag{2877}$$

where:

X_{SKIN}	Moisture content of air at the skin surface
X_{AIR}	Moisture content of ambient air
β	Mass transport coefficient
r	Heat of evaporation of water

Taking into account the analogy between heat transfer and mass transport:

$$\beta = \frac{a_{CONV}}{\rho \cdot c_p \cdot Le^{0.67}} \tag{2878}$$

The Lewis number Le does not deviate much from a value of one for the considered range of application. Therefore, the following relation can be assumed for the surface heat flux due to transpiration evaporation:

$$\dot{q}_{TRANS} = \frac{a_{CONV}}{c_p} \cdot r \cdot \left(X_{SKIN} - X_{AIR}\right) \tag{2879}$$

Environmental conditions lead to a maximum evaporation rate. When the environment is saturated, sweat can no longer evaporate. Drops of sweat form on the skin that do not contribute to the heat loss. The maximum evaporation rate is given by:

$$Q_{EVAP,\max_j} = k \cdot a_{CONV} \cdot A_j \cdot \left(p_{SAT,SKIN} - p_{AIR}\right) \tag{2880}$$

where $k = 0.0165 K / Pa$ is the evaporation coefficient [Berbner].

The heat losses due to diffusion and transpiration are summed as:

$$\dot{Q}_{EVAP_j} = \dot{Q}_{DIFF_j} + \dot{Q}_{TRANS_j} \tag{2881}$$

Heat Balance

This section describes the Heat Balance model.

The general total heat balance for the human body is:

$$\dot{Q}_{MET} - \dot{Q}_{ST} - \dot{Q}_{DRY} - \dot{Q}_{EVAP} - \dot{Q}_{RESP} - W = 0 \tag{2882}$$

The difference between $\dot{Q}_{MET}$ and W corresponds to the metabolic heat production. This heat balance can be established for all segments. The work W is equal to zero. Solving the above equation for the stored body heat flux $\dot{Q}_{ST}$ for the elements of a body segment and assuming that $T_{Cloth} > T_{Core} > T_{Muscle} > T_{Fat} > T_{Skin}$:

Core:

$$\dot{Q}_{ST_{1,j}} = \dot{Q}_{BLOOD_{1,j}} + \dot{Q}_{MET_{1,j}} - \dot{Q}_{COND_{1,j}} - \dot{Q}_{RESP_{1,j}} \tag{2883}$$

Muscle:

$$\dot{Q}_{ST_{2,j}} = \dot{Q}_{BLOOD_{2,j}} + \dot{Q}_{MET_{2,j}} - \dot{Q}_{COND_{2,j}} + \dot{Q}_{COND_{1,j}} \tag{2884}$$

Fat:

$$\dot{Q}_{ST_{3,j}} = \dot{Q}_{BLOOD_{3,j}} + \dot{Q}_{MET_{3,j}} - \dot{Q}_{COND_{3,j}} + \dot{Q}_{COND_{2,j}} \tag{2885}$$

Skin:

$$\dot{Q}_{ST_{4,j}} = \dot{Q}_{BLOOD_{4,j}} + \dot{Q}_{MET_{4,j}} + \dot{Q}_{COND_{3,j}} - \dot{Q}_{DRY_j} - \dot{Q}_{EVAP_j} \tag{2886}$$

In addition to these heat balance equations, the central cardiovascular system must be taken into account:

$$\dot{Q}_{ST,BLOOD} = \sum_{j=1}^{14} \left(\sum_{i=1}^{4} \left(-\dot{Q}_{BLOOD_{i,j}} \right) \right) \tag{2887}$$

In each segment, a radial heat flux is taken into account through the term $\dot{Q}_{COND}$. All elements are coupled through the central cardiovascular system of the above equation. The central blood temperature is a function of all temperatures and blood volumetric flow rates in the body.

In total, there are 57 differential equations that must be solved. At the beginning of the calculation, the temperatures must be initialized. An equilibrium condition is reached when the element temperatures no longer change $\left(\frac{\partial T}{\partial t} = 0\right)$ and the stored body heat flux becomes zero.

Clothing

This section describes the clothing model.

Clothing is viewed as an additional layer on the skin establishing a direct contact between skin and clothing. The clothing only influences the dry and wet heat fluxes of the passive system. The clothing does not influence the active system.

The influence of clothing on the dry heat flux is taken into account through a thermal efficiency factor F_{TE_j}. This factor quantifies the influence of the insulation through clothing in relation to the dry heat transfer coefficient:

$$F_{TE_j} = \frac{a_{CLOTH_j} \cdot F_{SURF_j}}{a_{CLOTH_j} + a_{DRY_j} \cdot F_{SURF_j}} \tag{2888}$$

The thermal efficiency factor ranges between 0 and 1, where 1 corresponds to a unclothed condition. The factor F_{SURF} takes the increase in size of the surface due to clothing into account. According to DIN ISO 7730, F_{SURF} for a skin element is computed as follows:

$$F_{SURF_j} = 1.00 + 1.29 \cdot I_{CLOTH} \tag{2889}$$

$$F_{SURF_j} = 1.05 + 0.645 \cdot I_{CLOTH} \tag{2890}$$

where I_{CLOTH} is the isolation factor clothing in m^2K/W.

Including F_{TE_j}, the dry heat flux becomes:

$$\dot{Q}_{DRY_j} = a_{DRY_j} \cdot F_{TE_j} \cdot A_j \cdot \left(T_{4,j} - T_{eff_j}\right) \tag{2891}$$

The temperature of the clothing surface T_{CLOTH_j} is:

$$T_{CLOTH_j} = T_{4,j} - \frac{\dot{Q}_{DRY_j}}{A_j a_{CLOTH_j}} \tag{2892}$$

A further effect of clothing is to restrain evaporation of moisture on human skin. To account for this effect, a factor F_{FD} is defined that reflects the water vapor permeability of clothing.

$$F_{FD_j} = \frac{1}{1 + 0.92 \cdot \left(\frac{a_{CONV_j}}{a_{CLOTH_j}}\right)} \tag{2893}$$

$$Q_{EVAP,\max_j} = k \cdot a_{CONV_j} \cdot F_{FD_j} \cdot A_j \cdot \left(p_{SAT,SKIN_j} - P_{AIR_j}\right) \tag{2894}$$

Contact Surface

This section describes the Contact Surface model.

The passenger is permanently in contact with the vehicle seat. Therefore, it is necessary to divide the skin elements at the contact segments into two different surfaces. This leads to two skin temperatures for a segment:

- Skin temperature for the surface in contact with the seat
- Skin temperature for the free surface

The effective temperature at the contact surface is replaced by the temperature of the seat $T_{CONTACT}$. The dry heat flux is calculated only from heat conduction between the contact surface and the seat. Taking clothing into account, the thermal efficiency factor at the contact surface is given by:

$$F_{TE_j} = \frac{a_{CLOTH_j} \cdot F_{SURF_j}}{a_{CLOTH_j} + a_{CONTACT} \cdot F_{SURF_j}} \tag{2895}$$

The dry heat flux then becomes:

$$\dot{Q}_{DRY_j} = a_{CONTACT} \cdot F_{TE_j} \cdot A_{CONTACT_j} \cdot \left(T_{4,j} - T_{CONTACT_j}\right) \tag{2896}$$

The thermal efficiency factor F_{TE} is used to modify all base values. For all further calculations, such as regulation or heat losses, the respective temperatures of the skin elements are used and so two separate values result.

Active System

The thermophysiological regulation is a system whose functioning is still not fully understood today. A description or a model for the thermophysiological regulation applies the relations of cause and effect that are obtained through experimental results.

Such a model cannot describe the direct functions and processes of the central nervous system.

Generally speaking, the used thermophysiological regulation model is a cascaded proportional regulation. A central controller calculates four controlled variables for the entire body based upon the temperature signals. These controlled variables are passed on to all segments where they are processed in a local controller unit. The local controller provides blood flow rates and heat fluxes that influence the thermal condition of the body. This way, the influence of different body parts on the total thermal condition can be investigated.

The skin temperatures of all body segments and the head core temperature are the input quantities of the controller. The head core temperature corresponds in the model to the temperature of the hypothalamus, the center of temperature regulation. This temperature is particularly important for the thermophysiological regulation, because the controller always uses it as a reference. From these segment temperatures, the control deviations, θ, are calculated. The control deviation consists of the difference between the respective segment temperature and a target value that is called reference temperature T_{REF}. The reference temperature is known for each segment. This calculation is performed for the head core temperature and the skin temperatures.

$$\theta_{1,1} = T_{1,1} - T_{REF_{1,1}} \tag{2897}$$

$$\theta_{4,j} = T_{4,j} - T_{REF_{4,j}} \tag{2898}$$

These control deviations are the basis for the computation of the global controlled variables. For the skin temperatures, it is distinguished whether the control deviation is a warm or a cold signal. A positive deviation corresponds to a warm signal and a negative deviation corresponds to a cold signal:

$\theta_{WA_{4,j}} = \theta_{4,j}$ for $\theta_{4,j} \geq 0$

$\theta_{CO_{4,j}} = -\theta_{4,j}$ for $\theta_{4,j} < 0$

The element signals are combined to a total signal. When calculating this total signal, the distribution of the thermal receptors on the skin (cold and warm spots) are taken into account. The values are obtained through multiplication of the number of cold and warm spots of a skin element with the corresponding area fraction of the entire body surface [Aschoff]. The following table gives the distribution of receptors for each skin element:

Segment	Zref (cold signal)	Xref (warm signal)
Head	6.05e-02	0.201
Torso	0.4926	0.38
Upper arm	0.01714	0.02999
Forearm	0.01714	0.01625
Hand	0.09224	0.025
Thigh	0.04514	0.08999
Lower leg	3.01e-02	0.03275
Foot	1.67e-02	0.0155

The total signals of the skin elements for warm and cold impulses are written as:

$$\theta_{WA,TOT} = \sum_{j=1}^{14} \left(\theta_{WA_{4,j}} \cdot X_{REF_j} \right) \tag{2899}$$

$$\theta_{CO,TOT} = \sum_{j=1}^{14} \left(\theta_{CO_{4,j}} \cdot Z_{REF_j} \right) \tag{2900}$$

By using these values together with the control deviation of the head core temperature, the global controlled variables during heat load (vasodilation and transpiration) and during cold load (vasoconstriction and shivering) are calculated.

After calculating the global controlled variables, these values are converted into local controlled variables for each body part. The global

conversion constants are labelled G and the local conversion constants are labelled L. The global controlled variables are denoted S. To characterize the type of controlled variable, the indices V_d for vasodilation, V_c for vasoconstriction, *Transp* for Transpiration, and *Shiv* for shivering are used. The local controlled variables are volumetric flow rate of blood, evaporation, and metabolic heat production.

Target Value

In the regulatory circuit, the target values correspond to the reference temperatures.

The target values have constant values and therefore, the command values are neglected. The head core reference temperature is 37.3 °C. The reference temperatures for each body element T_{REF} are given in the following table:

Segment	Core	Muscle	Fat	Skin
Head	37.3	36.7	36.343	36.116
Torso	37.336	36.965	35.593	34.736
Upper Arm	35.739	35.176	34.629	34.407
Forearm	35.739	35.176	34.629	34.407
Hand	35.684	35.627	35.592	35.543
Thigh	36.561	35.908	35.090	34.814
Lower leg	36.561	35.908	35.090	34.814
Foot	35.647	35.573	35.548	35.502

Under normal conditions, the body is at reference temperature. In such a case, the body has no regulatory activity.

Thermal Transpiration

This section describes the Thermal Transpiration model.

The global controlled variable, transpiration, is calculated as:

$$S_{TRANSP} = G_{TRANSP,KK} \cdot \theta_{1,1} + G_{TRANSP,HA}\left(\theta_{WA,TOT} - \theta_{CO,TOT}\right) \tag{2901}$$

The basis for this global controlled variable and the used coefficients are experiments described in [Stolwijk (1966)].

The following assumptions are made:

- Transpiration is proportional to the head core temperature at a constant average skin temperature.
- Transpiration is proportional to the average skin temperature at a constant head core temperature.

The global constants for calculating the transpiration signal have the following values:

$$G_{TRANSP,KK} = 372.2 \frac{W}{K}$$

$$G_{TRANSP,HA} = 33.7 \frac{W}{K}$$

The above equation means that the influence of the core head temperature is approximately 11 times higher than of the skin temperature.

The local effect of transpiration is accounted for by a factor that depends on the control deviation and consequently on the temperature of each skin element. The factor behaves in such a way that when a local control deviation of the skin temperature occurs, an n-fold amplification of transpiration takes place. The value of this amplification factor can vary in literature. For this model, the value from [Stolwijk (1971)] is used, which is 2.

$$\dot{Q}_{TRANSP_j} = \left(L_{TRANSP} S_{TRANSP}\right) \cdot 2^{\frac{\theta_{4,j}}{10}} \tag{2902}$$

The distribution of sweat formation across the skin surface is assumed to correspond to the distribution of perspiratory glands on the skin surface . The values of this local controlled variable are listed in the following table:

Segment	L_{TRANSP}
Head	0.081
Torso	0.481
Upper arm	0.039
Forearm	0.039
Hand	0.016
Thigh	0.064
Lower leg	0.043
Foot	0.018

The local transpiration rate is limited by the maximum possible evaporation. The evaporation also includes diffusion. This maximum evaporation rate tells how large the total evaporation heat loss of the body to the environment can be at maximum. This limitation results from the vapor pressure difference between the skin surface and the environment and additionally depends on the convective heat transfer coefficient.

If the computed value for the evaporation heat loss exceeds the sum of the constant diffusion heat flux $\dot{Q}_{DIFF}$ and the evaporation through transpiration $\dot{Q}_{TRANSP}$, the maximum possible value is taken. Therefore, the heat loss due to transpiration is reduced:

$$\dot{Q}_{TRANSP_j} = \dot{Q}_{EVAP,MAX_j} - \dot{Q}_{DIFF_j} \tag{2903}$$

The ratio of effective transpiration to maximum possible transpiration yields the degree of moisture w_j, that is, the fraction of moistened skin to total skin surface.

$$w_j = \frac{\dot{Q}_{TRANSP_j} + \dot{Q}_{DIFF_j}}{\dot{Q}_{EVAP,MAX_j}} \tag{2904}$$

Vasomotricity

Vasomotricity enables the human body to react to heat or cold loads. The global controlled variables are vasoconstriction in case of cold load and vasodilation in case of heat load.

These two global controlled variables are calculated separately by the controller and are used for the calculation of the local controlled variables of the elements:

$$S_{VD} = G_{VD,KK} \cdot \theta_{1,1} + G_{VD,HA}\left(\theta_{WA,TOT} - \theta_{CO,TOT}\right) \tag{2905}$$

$$S_{VC} = -G_{VC}\left(\theta_{1,1} + \theta_{WA,TOT} - \theta_{CO,TOT}\right) \tag{2906}$$

For the vasomotricity, the local controlled variable is given by the variation of the volumetric blood flux $\dot{V}_{BLOOD}$ that flows through the various skin parts. The local effect of the skin temperatures of the segments is taken into account when the volumetric blood fluxes in the skin parts are calculated.

A decrease in skin temperature of 10 °C leads to an increase of resistance in the blood vessels [Stolwijk (1971)]. This effect is taken into account

through the factor $2^{(\theta_{4,j}/6)}$. The volumetric blood flux through the skin parts is composed as follows:

$$\dot{V}_{BLOOD_{4,j}} = \left(\frac{\dot{V}_{BLOOD,BAS_{4,j}} + L_{VD_j} S_{VD}}{1 + L_{VC_j} S_{VC}} \right) \cdot 2^{\left(\frac{\theta_{4,j}}{6}\right)} \tag{2907}$$

The above equation connects the local and the global controlled variables. The global controller constants are:

$$G_{VD,KK} = 32.5 \cdot 10^{-6} \frac{m^3}{sK}$$

$$G_{VD,HA} = 2.1 \cdot 10^{-6} \frac{m^3}{sK}$$

The vasoconstriction shows a high sensitivity with respect to the head core temperature and the skin temperature. For the calculation of the vasoconstriction controlled variable signal, a multiplicative approach is chosen. The controller constant is in this case:

$$G_{VC} = 5 \frac{1}{K}$$

The influence of vasodilation on the skin elements is investigated in [Hertzman]. Based on those results, the distribution for the vasodilation is chosen from the values that return the biggest dilation of the blood vessels.

The local constants for vasodilation and vasoconstriction are listed in the following table:

Segment	L_{VD}	L_{VC}
Head	0.132	5.0e-02
Torso	0.322	0.0
Upper arm	2.375e-02	1.25e-02
Forearm	2.375e-02	1.25e-02
Hand	6.049e-02	0.175
Thigh	6.9e-02	5.999e-02
Lower leg	4.6e-02	4.0e-2
Foot	5.0e-02	0.175

Thermoregulatory Heat Production

When the body temperature drops, the skeletal muscles are activated. This process is called shivering and serves to maintain the inner temperature.

The global controlled variable for shivering is given by:

$$S_{SHIV} = -G_{SHIV} \cdot \theta_{1,1} \cdot \theta_{CO,TOT} \tag{2908}$$

Analogous to the other thermo-regulatory controlled variables, the global controlled signal for shivering is transferred to the local controller and is multiplied by a controller constant to obtain the local portion of the heat production:

$$Q_{MET,SHIV_{2,j}} = S_{SHIV} \cdot L_{MUSC,SHIV_j} \tag{2909}$$

According to measurements in [Benzinger] and [Nadel], there is a multiplicative connection between the controlled variable and the local controller constant. From experiments, the global controller constant is:

$$G_{SHIV} = 24.425 \frac{W}{K^2} \tag{2910}$$

The distribution of the heat production in each muscle element $L_{MUSC,SHIV}$ is listed in the following table. Additionally, the distribution of heat production through muscular work in each muscle element $L_{MUSC,WORK}$ is listed:

Segment	$L_{MUSC,SHIV}$	$L_{MUSC,WORK}$
Head	2.0e-02	0.0
Torso	0.86	0.3
Upper arm	1.25e-02	2.0e-02
Forearm	1.25e-02	2.0e-02
Hand	0.0	5.0e-03
Thigh	2.1e-02	0.18
Lower leg	1.4e-02	0.12
Foot	0.0	5.0e-03

The volumetric flux of blood through the muscle elements is given by:

$$\dot{V}_{BLOOD_{2,j}} = \dot{V}_{BLOOD,BAS_{2,j}} + 0.2514 \cdot 10^{-6} \cdot S_{SHIV} L_{MUSC,SHIV} + 0.2514 \cdot 10^{-6} L_{MUSC,WORK} Q_{MET,WORK,TOT} \tag{2911}$$

What is the Thermal Comfort Wizard?

The thermal comfort wizard is a tool that provides a clear and structured approach to setting up a thermal comfort simulation.

The wizard is used to define the key properties required by the TIM program, which are summarized below:

- General settings, which define the TCM passenger properties:
- Advanced settings, which define the solver properties:
- External convection and radiation, which define TCM boundary properties.
- Pre-processing, which creates scenes and reports specific to thermal comfort simulation.
- Post-processing, which creates plots and reports specific to thermal comfort simulation.

A guide to using the thermal comfort wizard is included in the Setting up a Thermal Comfort Case section.

What are TCM Passengers?

Thermal comfort passengers are defined in the General Settings tab of the thermal comfort wizard.

The TCM passengers are represented with a node by that name in the object tree. Its node is subject to common menu items, and appears when you activate the thermal comfort model.



The **TCM Passenger** manager contains subnodes for each individual occupant in the vehicle, as defined by the wizard. These subnodes are placeholders for the passenger characteristics, which are displayed in the *Properties* window.



The managers for each occupant also contain subnodes, which hold information regarding the body segment boundaries, for both solid and air regions.



See Also

Working with General Settings
What is the Thermal Comfort Wizard?
Activating the Thermal Comfort Model
What is the Thermal Comfort Model?

What are TCM Boundaries?

Thermal comfort boundaries are defined in the External Convection and Radiation tab of the thermal comfort wizard.

The TCM boundaries are represented with a node by that name in the object tree. The node appears when you activate the thermal comfort model.



Each TCM boundary has a subnode(s) that represents the boundary conditions. Selecting this node displays the boundary value in the *Properties* window.



See Also

Adding External Convection and Radiation
What is the Thermal Comfort Wizard?
Activating the Thermal Comfort Model
What is the Thermal Comfort Model?

What is the Thermal Comfort Solver?

Thermal comfort solver settings are defined in the Advanced Settings tab of the thermal comfort wizard.

The thermal comfort solver is represented with a node by that name in the object tree. Its node is subject to common menu items, and appears when you activate the thermal comfort model.



The **Solver > Thermal Comfort** node acts as a placeholder for solver properties specific to the TIM program, which are displayed in the *Properties* window.



See Also

Using Advanced Settings
What is the Thermal Comfort Wizard?
Activating the Thermal Comfort Model
What is the Thermal Comfort Model?

STAR-CCM+ Capabilities Related to Thermal Comfort

This section outlines the STAR-CCM+ features that are related to thermal comfort modeling.

The thermal comfort wizard is the primary tool for these types of problems, but there are several other STAR-CCM+ features that can help with the CFD solution:

- Modeling Options:
  - Stationary
  - Three-Dimensional
  - Solid Energy
- Radiation
- Turbulence Models:
  - K-Epsilon, K-Omega...
- Boundary Conditions:
  - Mass Flow Inlet
  - Pressure
- Interfaces:
  - Internal
  - Contact
- Solvers:
  - Second-Order
  - Segregated
  - Algebraic Multigrid

Setting Up a Thermal Comfort Case

This section outlines the process of setting up a thermal comfort case.

This wizard should be used after you generate a mesh, select models, and set up initial conditions and boundary conditions. The process involves the following:

- Activating the thermal comfort model
- Activating the thermal comfort wizard
- Working with general settings
- Using advanced settings
- Adding external convection and radiation
- Pre-processing
- Saving and reloading TCM settings

Activating the Thermal Comfort Model

The Thermal Comfort model is selected from within the *Physics Model Selection* dialog on the relevant continuum.

To select the thermal comfort model, the following models are required:

- Any model from the *Material* group, depending on simulation requirements.
  - However, if you wish to model air recirculation, you must also select Multi-Component Gas from the *Material* group box to model the water vapor and dry air gas components.
- Either the Segregated or Coupled solver.
- An energy model from the *Optional Physics Models* group, which may be one of the following:
  - Segregated solver:
    - Segregated Fluid Temperature
    - Segregated Fluid Isothermal
    - Segregated Fluid Enthalpy
    - Segregated Solid Energy
  - Coupled solver:
    - Coupled Energy

- Coupled Solid Energy

The Thermal Comfort model is then available in the *Optional Physics Models* group.



When the Thermal Comfort model is activated, the TCM Passenger and TCM Boundaries managers become available in the object tree.

See Also

What are TCM Passengers?
What are TCM Boundaries?

Activating the Thermal Comfort Wizard

The properties are set using the Thermal Comfort Wizard.

To begin accessing this feature:

1. From the menu bar, select **Window** > **Toolbars** > **Customize**.



This opens the *Customize Toolbars* dialog as shown in the following screenshot.



2. Scroll down to the **Thermal Comfort Wizard** icon.
3. Drag and drop this icon onto the end of one of the toolbars near the top of the workspace.
4. Click **Close** on the *Customize Toolbars* dialog.
5. Click the (**Thermal Comfort Wizard**) button.

   This opens the *Thermal Comfort Wizard* dialog.



The thermal comfort wizard assumes that the model has been built, and boundaries defined. Initial temperature values for the passenger body segments must also be specified before using the wizard. Data is then entered into the wizard to correctly define the characteristics of the thermal comfort simulation. Information entered in the *Thermal Comfort Wizard* dialog is retained automatically, so you can move between each of the tabs and even close and reopen the wizard without losing any information. The data is only applied to the simulation and added to the relevant managers in the object tree when the Save TCM Settings button is clicked. It is also possible to reload information into the wizard, to match the current TCM settings in the STAR-CCM+ simulation.

The dialog includes a **Close** button that closes the wizard without writing to the file. Data entered in the wizard are cleared only if the STAR-CCM+ client is shut down.

See Also

Saving and Reloading TCM Settings

Working with General Settings

The *General Settings* tab allows you to specify properties for regions and boundaries.



Conventions for Preparing Body Segment Boundaries

Each passenger should be composed of fourteen body segments, which must be defined as separate boundaries. These boundaries correspond to the fourteen drop-down lists in the General Settings tab. Provided below is a diagram illustrating how the body segment menus have been arranged to approximate a human body.



In cases where the passenger touches solid regions, such as the torso and legs touching a seat, there must be a separate boundary for each solid segment as well.

It is important to note that boundary labels need to include the names of their respective body segments, as shown in the following screenshot. The thermal comfort wizard recognizes only the boundaries that are named in accordance with this convention.



If there are multiple people in the simulation, the boundary names should reflect this, for example `Passenger2 LeftUpperArm`.

Entering Passenger Properties

Begin using the wizard by selecting a value in the Number of Passengers spin box. This value should match the total number of people represented in the STAR-CCM+ simulation.



Based on the number you specify, the Person setting contains a numbered list representing passengers. Select a passenger by the number, and complete your settings for that passenger throughout the tab, as described in detail in the rest of this section. Once all the settings have been provided for that passenger, you can change the Person

number and enter the next set of data. The thermal comfort wizard retains your entries for each passenger.

For each passenger, begin by assigning the passenger a generic name, such as `Passenger1`. By default the first passenger has the name `Driver`, which can be changed.



Specify height, clothing resistance, and metabolic rate.

- The passenger **Height** is used by the TIM program to calculate the body mass distributions for the vascular system and the body surface area values.
- The **Clothing Resistance** defines the thermal conductivity of the passenger's clothing in m^2K/W.
- The Metabolic Rate is the rate at which the person releases energy to the surroundings as heat, which depends on their level of activity.

  It is set by selecting one of the **Driver** or **Passenger** radio buttons, or by entering a **Custom** value. When the **Passenger** option is selected the metabolic rate is set to `58 W/m^2`, which is equivalent to a person who is seated and relaxed. When the **Driver** option is chosen the metabolic rate is set to `75.4 W/m^2` (passenger value `x1.3`), indicating that the driver is more active. The **Custom** option allows you to specify an alternative value by entering a scaling factor relative to 58 W/m^2, and is usually in the range of `0.8` to `2` (where `0.8` is equivalent to reclining, and `2` corresponds to medium activity.)



Make sure that the correct fluid region is selected under Select Model Region. The options in this drop-down list, and in similar ones for boundaries, correspond to the regions and boundaries that you created for this simulation file.



Select the appropriate boundary for each body segment.



Repeat this for any solid regions and boundaries that are in contact with the passenger.



For a multiple occupant model, the *Passenger Filter* can be used as a search function to display the relevant boundaries for a particular passenger. For example, before opening the thermal comfort wizard, the driver boundaries can be renamed to contain the word driver. `Driver` can then be applied in the filter to display the required boundaries in the drop-down menu for each body segment.



You can specify an initial body segment temperature in the **Initial Temp** text box. This is the value passed to the TIM program the first time it is called. Thereafter, the body segment temperature updates each time TIM is run. This body segment temperature differs from the boundary temperature specification set in the simulation tree. The boundary temperature specification is used during the initial iterations before calling TIM, and is then overwritten by the body segment temperature. All temperatures in the wizard are specified in degrees Celsius.

The **Reload Boundaries** button causes the *Thermal Comfort Wizard* to re-read the boundary information in the simulation and update options available in the Boundary Selection menus. This is used when boundary names have been changed in the simulation, and resets all boundary menus to `-none-`.



Using Advanced Settings

If you are an experienced user of TCM, you may want to adjust parameters in the *Advanced Settings* tab. Otherwise, leave the settings at their defaults.



This tab includes a host of controls with which you can refine your TCM settings.

Temperature limiter (deg C): 2

Humidity Model: Relative humidity will be computed using vapor fraction data from Star-CCM+

Relative humidity will be held constant at (%) 30.0

TIM output frequency (number of TIM time steps per output) for generating out*.dat files 10

Number of STAR iterations before initial call to TIM (steady-state): 50

Number of STAR iterations after subsequent calls to TIM: 2

Convergence tolerance, defined as Sum[abs(Tnew - Told)/Tnew] / 14 : 0

Output frequency per STAR iteration for updating passenger_tim_??.dat files 5

Update passenger_tim_??.dat files? Yes

No

Percentage of recirculated air (0-100%) 0

External Humidity Absolute (kg/m^3) 0.0

Relative (%)

- **Temperature Limiter**—the maximum amount TIM is allowed to change the skin temperatures it reports to STAR-CCM+ each time it runs. Larger values lead to faster convergence, but smaller values result in a tighter convergence.
- **Humidity Model**—This radio button determines whether the evaporation off of the passengers' skin is based on humidity values calculated in the CFD model, or uses a single fixed value for relative humidity. The default value for a constant humidity level is 30%.
- **TIM Output Frequency**—specifies how many TIM time steps to run each time it is called.
- **Number of STAR iterations before initial call to TIM**—sets how many iterations to run the CFD-only portion of the model. This gives the flow field time to develop before performing the passenger thermal calculations as well.
- **Number of STAR iterations after subsequent calls to TIM**—sets how frequently the passenger thermal calculations are performed.
- **Convergence tolerance**—stops the simulation if the rate of change of passenger skin temperatures falls below this threshold.
- **Output frequency**—specifies how often to update the files with detailed results of the passengers' thermal calculations. There is a separate file for each passenger, numbered `01`, `02`, …
- **Update passenger_tim_??.dat files**—This radio button specifies whether to write to the files containing results of the passengers' thermal calculations.

- **Percentage of recirculated air (0-100%)**—Specifies how much of the cabin air is recirculated and cooled by the air conditioning system. For example, a value of 0 means that the air cooled by the air conditioning system has the same temperature and humidity as the external air. Similarly, a value of 100 means that the air being cooled has the same temperature and humidity as the air at the outlet boundary.
- **External Humidity**—This radio button specifies whether the external humidity used in the air recirculation calculations is relative or absolute. The value of external humidity (relative or absolute) is entered in the text field next to the radio button.

You can use the Maximum Steps and Stop File stopping criteria to stop the simulation running.

See Also

What is the Thermal Comfort Model?

Adding External Convection and Radiation

In the *External Convection and Radiation* tab you can define the properties of any thermal boundaries in the model, as well as specify the mass flow rate and temperature at the inlets.

It lets you apply external convection and non-solar radiation heat fluxes easily, as well as specify changes to thermal properties with time (transient cases only). In order to do this, you must set up the appropriate boundary conditions in advance, for example setting the Thermal Specification to *Convection*. The values set in this panel are used by TIM when the TCM model is run, and overwrite any existing boundary values.

The functionality and different sections of the panel are described in detail below.



Boundary Selection

To specify the thermal properties or flow conditions for a boundary, it must first be selected using the drop-down menus labeled **Region** and **Boundary**.



When you have finished entering data for a particular boundary, the next boundary can be selected. Data is retained for each boundary.

Boundary Type

When the boundary has been selected, as described above, the next step is specifying the boundary type. This is done by selecting the appropriate radio button under Boundary Type, which activates the corresponding text boxes so that values can be entered. The different options are described below.



- **Fixed Wall Temperature**—boundary thermal specification is set to **Temperature**. Activates the **Wall Temperature** text box (shown below) under Wall Thermal Settings in which a constant temperature value can be entered.



- **Fixed Wall Heat Flux**—boundary thermal specification is set to **Heat Flux**. Activates the **Wall Heat Flux** text box under Wall Thermal Settings in which a constant heat flux value can be entered.



- **Inlet**—boundary type is set to **Mass Flow Inlet**. Activates the **Inlet Boundary Settings** text boxes where the mass flow rate and temperature at the inlet can be entered.



- **External Convection**—boundary thermal specification is set to *Convection*. Activates the **External Convection Boundary Condition** text box, allowing a reference heat transfer coefficient to be entered.
- **External Convection with Wall Conduction**—same as the previous setting but has more options for specifying a wall thickness and corresponding conductivity value in the **External Conduction Boundary Settings** text boxes. This allows you to simulate the insulating effects of walls without explicitly modeling them as solid cells.
- **External Convection and Radiation**—boundary thermal specification is set to *Convection*, and Radiation physics models have been enabled. Activates the **External Radiation and Convection Boundary Settings** text boxes, allowing external thermal properties to be entered. When this option is selected STAR-CCM+ performs a linearized approximation of the radiation heat flux in parallel to convective heat transfer.

- **External Convection and Radiation with Wall Conduction**—same as the previous setting but has more options for specifying a wall thickness and corresponding conductivity value in the **External Conduction Boundary Settings** text boxes. This allows you to simulate the insulating effects of walls without explicitly modeling them as solid cells.



The External Convection Global Settings are common to all boundaries, and allow heat transfer coefficients for one velocity (reference) to be applied to a model operating at a different velocity (vehicle). The external bulk temperature for the model is also specified here, which is used when calculating external convection and the temperature of recirculated air.



Because external heat transfer coefficient values may not be available for the speed at which the vehicle is traveling, TCM provides this method to scale the values that are available:

- **Vehicle Velocity** is the actual velocity at which the vehicle is traveling.
- **Reference Velocity** is the velocity at which the external heat transfer coefficients are valid.
- For each external boundary, a **Reference Heat Transfer Coefficient** can be entered. These heat transfer coefficients are scaled by the ratio of the two velocities raised to the **Scaling Exponent** before being applied to the boundaries:

  $$h_{app} = h_{ref} \cdot \left(\mathbf{v}_{vel} / \mathbf{v}_{ref}\right)^n$$

  where h_{app} is the applied heat transfer coefficient, h_{ref} is the reference heat transfer coefficient, $\mathbf{v}_{vel}$ is the vehicle velocity, $\mathbf{v}_{ref}$ is the reference velocity, and n is the exponent.

Adding Time Settings

In transient cases you can specify changes to the thermal properties of a boundary with time, for example at time `0` the wall temperature is `10` deg C, then at `30` seconds the temperature increases to `15` deg C. To do this, you must first specify after how many seconds a change occurs. This is done by entering a value (in seconds) in the **Add Time Setting** text box and clicking **Add Time**.



This value is added as an option in the drop-down menu as shown below. To use the time setting, first ensure that the appropriate boundary has been selected, then use the drop-down menu to select a time setting. When you do this the text boxes containing the boundary's values are cleared, allowing you to enter the new values. Values are scaled linearly between time settings and multiple time settings can be used in a TCM simulation.



See Also

What is the Thermal Comfort Model?

Pre Processing

The thermal comfort wizard allows you to perform more pre-processing steps in order to get useful data from the model and prepare scenes.

Details on preparing scenes for analysis are included in the subsection called Preparing Scenes. This section covers all other pre-processing capabilities available in the *Pre Processing* tab.



Fluid Volume > Get Values

When you click the **Fluid Volume > Get Values** button, the thermal comfort wizard evaluates the fluid volume in the model and report the size of the volume in m^3.

Window Settings > Get Values

Clicking the **Window Settings > Get Values** button reports the area of each window in the simulation. The boundaries are selected automatically based upon the standard naming conventions. A text file called `window_area.dat` is also written to the current working directory, which contains the window area information.

Creating All Scenes

Clicking the **Create All Scenes** button creates all the scenes described in the Preparing Scenes section. Each scene is created with an isometric view (-X, -Y, +Z) and a reverse isometric view (+X, +Y, -Z).

Hardcopy All TCM Scenes

Clicking the **Hardcopy All TCM Scenes** button creates a PNG image file of all scenes that you have created with the thermal comfort wizard. For example, if you have clicked **Specific Heat Exterior** and **Density Exterior**, then only those two scenes are exported. The image file is located in the current working directory.

Preparing Scenes

This section describes how to prepare scenes.

The set of buttons shown below can be used to create scenes displaying contours of various field functions on certain boundaries in the model.



Specific Heat, Thermal Conductivity, and Density

These buttons generate scenes showing contours of specific heat, thermal conductivity, or density for the vehicle exterior or interior solids. The exterior scene includes all solid regions that do not have "`road`" in the name and are on the exterior of the model; it is intended for models in which the windows, doors, and other panels are modeled as solids. The interior scene is used to display contours on solid regions on the interior of the model, which are identified as those with "`seat`" in the name. If there are no solid regions in the model, these buttons are not activated.

Thermal Emissivity and Solar Emissivity

These buttons generate scenes showing contours of thermal emissivity or solar thermal emissivity for the exterior and interior surfaces of the fluid regions (not solid regions). The exterior scene displays contours on all fluid boundaries in the model, whereas the interior scene only includes boundaries that have "`seat`" in their name. If the model does not include a radiation model, these buttons are not activated.

It is assumed that two separate radiation spectral bands have been defined in the model, and that the first band has lower wavelengths (that is, is solar radiation), while the second band has higher wavelengths (that is, is non-solar radiation). For the Solar Emissivity scenes, the emissivity values for the first band are displayed. For the Thermal Emissivity scenes, the emissivity values for the second band are displayed.

Displaying Sun Location

For a case where the sun has been included in the simulation by using the Solar Loads model, the **Sun Location** button can be used to give a visual indication of the location of the sun. The sun is represented in the model by a sphere, and its position is calculated automatically based on the *Azimuth* and *Altitude Angles* you have specified in the solar load properties.

When you click the **Sun Location** button:

- A new sphere part is created;
- A new *Sun Location* scene is created;
- The fluid region and the sun part are displayed in the scene.

You are able to view this scene once you have closed the thermal comfort wizard.



Displaying All TCM-Specified Boundaries

Clicking the **All TCM-specified Boundaries** button creates a geometry scene displaying all boundaries specified within the thermal comfort wizard's *External Convection and Radiation* tab. This is useful for identifying the boundary types assigned in the thermal comfort wizard, and quickly creating scalar or vector scenes with the TCM boundaries by adding scalar or vector displayers.

The TCM boundaries are categorized with a different geometry displayer for each boundary type.



A mesh color is applied to each scene displayer to visually distinguish between TCM boundary types. The scene displayers, equivalent TCM boundary type, and mesh color are summarized below:

Scene Displayer	TCM Boundary Type	Mesh Color
Convection	External Convection	Red
Convection Conduction	External Convection with Wall Conduction	Red
Convection Conduction Radiation	External Convection and Radiation with Wall Conduction	White

Scene Displayer	TCM Boundary Type	Mesh Color
Convection Radiation	External Convection and Radiation	Cyan
Heat Flux	Fixed Wall Heat Flux	Blue
Inlet	Inlet	Magenta
Temperature	Fixed Wall Temperature	Green

Heat Transfer Coefficients

Clicking the **Heat Transfer Coefficients** button creates a scalar scene displaying the heat transfer coefficients for all solid regions and boundaries. This speeds up the solution monitoring and post-processing requirements by automatically creating a heat transfer coefficient scene with typical settings and boundaries.

Saving and Reloading TCM Settings

The **Save TCM Settings** button takes the data currently entered in the thermal comfort wizard, and adds it to the relevant placeholders in the object tree.



You can also load the existing TCM settings into the wizard by clicking the **Reload TCM Settings** button. This replaces the current settings in the wizard with the TCM settings stored in the STAR-CCM+ object tree. Any changes made in the dialog since the last time the TCM settings were saved are overwritten when the data is loaded.

Performing a TCM Simulation

Once the TCM settings have been saved, the simulation can be run.

The solver runs the STAR-CCM+ simulation for the specified number of iterations, then pass the heat flux values for each body segment to the TIM program. TIM uses the heat fluxes to calculate skin and clothing temperatures, which are passed back to STAR-CCM+. The two programs continue this iterative process until the change in body temperature falls below the convergence criterion, until you stop the solver, or until the stopping criteria have been reached.

The results can be displayed in a typical scalar scene.



Using Auto Save in a TCM Simulation

The auto save capability within STAR-CCM+ can be used in a TCM simulation to save the `passenger_state.dat` file at given intervals. This file is used by TIM to store data relating to the thermal state of the passenger, and is updated after each TIM iteration. When the `passenger_state.dat` file exists in the working directory, TIM automatically uses this as the starting point for subsequent iterations. By using the autosave facility, not only can you save backup copies of

the `.sim` file, but you can also save the corresponding `passenger_state.dat` file, allowing you to restore the simulation following a system failure, for example.

When restarting from a solution other than the current one, it is usually a good idea to restore the `passenger_state.dat` file as well as the CFD solution. This is because the internal body temperatures are stored in this file and not in the `.sim` file. To restore a previous solution, replace the `passenger_state.dat` file with the desired `passenger_state.dat.xx` file, by removing the `.xx` from the end of the file name. This has no effect on the TCM settings in the `passenger.dat` file.

See Also

Saving and Reloading TCM Settings

Post Processing a TCM Simulation

The thermal comfort wizard has some additional post-processing capabilities that are specific to thermal comfort simulations.

These are available in the *Post-Processing Calculations* tab and are described below.



Creating LMV Plots

The **Create LMV Plots** button allows you to create plots displaying the LMV (local mean vote) values for occupants in the model. When you click this button, the plot(s) is created and a hardcopy of each plot, such as the one shown below, is saved as a `.png` file. As well as this, a `.csv` file containing the numerical data is generated. LMV calculations require that radiation be active in the model.



In the LMV plot, the upper and lower limits for thermal comfort are indicated by red and blue lines. Each body part is listed on the Y-axis, and the corresponding local mean vote for the part is shown as a black dot on the plot. If a black dot is on the left of the blue line, the corresponding body part may be considered to be too cold. If the black dot is on the right of the red line, the body part may be considered to be too hot. In either of these cases, this would indicate an inadequate level of thermal comfort for the passenger.

Energy Balance in Extruded Regions

Clicking the **Create Plot** button under *Energy Balance in Extruded Regions* creates a plot displaying the heat transfer (measured in Watts) on the interior and exterior surface of extruded regions. Before creating this plot, you must follow the naming conventions below so that the wizard selects the appropriate regions and boundaries:

- The name of the extrusion region must start with `Extrusion` (for example, Extrusion roof).
- The name of the interior boundary must end with `base` (for example, Interior base).

- The name of the exterior boundary must end with `top` (for example, Exterior top).

Calculating Values

There are three boxes in the *Post-Processing* tab that allow you to calculate the following values:

- Total heat transfer through windows

  This is the total heat transfer through each window, in W. A file called `window_heat_transfer.dat` is created in the current working directory, which contains the heat transfer data.

- Energy through each inlet

  This is the amount of energy transferred into the cabin through each inlet, in W.

- Area exposed to solar radiation

  This is the total area of faces where direct solar irradiation is >0 for each window and passenger boundary, in m^2.

Clicking the **Calculate** button for each value causes the thermal comfort wizard to perform the calculations and report the values in the corresponding box.

Working with TCM Input Files

This section describes the format of the external `passenger.dat` file, which is used with the thermal comfort wizard.

This file contains values for each passenger, and is used with the information in the TCM Passenger and TCM Boundaries managers for the thermal comfort calculations. The file is composed of sections, each of which starts with `SECTION <SECTION NAME>` and finishes with `END`. Any lines in the file which begin with `%` are considered comments. Additionally, everything following `!` is considered a comment.

An example of the file format has been provided.

PASSENGER

`SECTION PASSENGER` contains one line, with the number of passengers in the simulation (including the driver, if applicable).

PASSDATA

`SECTION PASSDATA` loops through data of all the passengers. For each passenger, the following items are required (one item per line):

- Passenger name
- Passenger height (in meters)
- Passenger surface metabolic rate (relative to 58 W / m^2)
- Clothing resistance (in m^2K/W)—This is applied to all body segments except the head and hands.
- Initial surface temperatures (in degrees C) for each body segment -- Each passenger is divided into 14 segments, and data for the segments is always presented in the sequence shown in the example.

After all of the data is given for Passenger #1, the sequence is repeated for Passenger #2, #3, and so on, as needed.

See Also

What is the Thermal Comfort Wizard?
What are TCM Passengers?
What are TCM Boundaries?

File Format Example

This section contains an example of a typical `passenger.dat` file.

```
% All lines that begin with a % are considered as
comments
% ======================================================
SECTION PASSENGER
% ======================================================
1
END
%
SECTION PASSDATA
% ======================================================
% Data for passenger 1
% ======================================================
Driver ! Passenger name
1.85 ! Passenger height
1.3 ! Passenger surface metabolic rate (relative to 58
W/m^2)
0.035 ! Clothing resistance (m^2K/W) all segments
except head & hands
%
%
% Initial body segment surface temperatures (deg C)
(must be in this order)
35.0 ! Head
35.0 ! Torso
35.0 ! Upper Left Arm
35.0 ! Upper Right Arm
35.0 ! Lower Left Arm
35.0 ! Lower Right Arm
35.0 ! Left Hand
35.0 ! Right Hand
35.0 ! Upper Left Leg
35.0 ! Upper Right Leg
35.0 ! Lower Left Leg
35.0 ! Lower Right Leg
35.0 ! Left Foot
35.0 ! Right Foot
% ============ END OF PASSENGER 1 DATA
===================
% To add more passenger data, simply repeat the above
lines for each passenger
END
% ======================================================
```

Working with TCM Output Files

This section details the contents of the thermal comfort output file, `passenger_tim_xx.txt`, where `xx` refers to the passenger number in the simulation.

This file contains information regarding the passenger characteristics used in the simulation, and a summary of results. The sections of the output files are listed below:

- Passenger information:
  - Anthropometric Basis Data
  - Initial Thermo-Physiological Quantities
  - Isolation Values: Clothing / Car Seat
  - Basis Values for Body Segments
  - Boundary Conditions at Beginning of Computation
- Results:
  - Final Thermo-Physiological Quantities
  - Values for Body Segments
  - Boundary Conditions at End of Computation

Anthropometric Basis Data

This section of the `passenger_tim_xx.txt` file contains details regarding the passenger anthropometry.

```
ANTHROPOMETRIC BASIS DATA :
============================

          SIZE (M)                    =
          WEIGHT (KG)                 =
          SKIN UPPER SURFACE (M2)     =
```

The passenger weight and skin surface area are defined based upon the passenger height entered in the wizard.

Initial Thermo Physiological Quantities

This section of the `passenger_tim_xx.txt` file contains information relating to the passenger vascular system, heat production, and heat loss, at the start of the run.

```
THERMO-PHYSIOLOGICAL QUANTITIES :
=================================

     CENTRAL VASCULAR SYSTEM :
          MASS (KG)                    =
          HEAT CAPACITY (J/K)          =
         REF.-TEMPERATURE (C)          =

     HEAT PRODUCTION (W) :
          BREATHING                    =
          ADDITIONAL ACTIVITY          =
          ENTIRE CONVERSION            =
          IN MET:                      =

     HEAT LOSS (W) :
         BREATHING
         SKIN      :
           SWEATING AND EVAPORATION    =
           DRY HEAT LOSS               =
           OVERALL                     =
         BODY    :
          SKIN AND RESPIRATION         =

     HEAT BALANCE (W) :
          PRODUCTION - LOSS            =


     BLOOD VOLUME FLOW  (CM3/S) :
          BREATHING                    =
```

The `IN MET` property shown above relates to the metabolic rate entered in the wizard. This value, along with the passenger skin surface area is used to determine the `ENTIRE CONVERSION` property, which is the total passenger heat production.

Isolation Values Clothing / Car Seat

This section of the `passenger_tim_xx.txt` file contains information relating to the passenger clothing and car seats.

```
ISOLATION VALUES: CLOTHING / CAR SEAT :
=========================================

     CLOTHING:
     ----------
     UNCLOTHED:

     TAKING HUMIDITY TRANSFER RESISTANCE
     CLOTHING ACCORDING TO NISHI/GAGGE


     CAR SEAT:
     ------------
           PADS:  HEAT CAPACITY
              HUMIDITY TRANSFER RESISTANC

      THE BACK  : HEAT CAPACITY
              HUMIDITY TRANSFER RESISTANT
```

Basis Values for Body Segments

This section of the `passenger_tim_xx.txt` file contains tabulated data regarding body segments and segment layers.

This section of the `passenger_tim_xx.txt` file summarizes the following properties:

- Mass (kg)
- Heat capacity (J/K)
- Skin surface area (m^2)
- Heat production from breathing (MW)
- Evaporation from skin (W)
- Blood flow for breathing (cm^3/s)
- Heat transfer (W/K)
- Temperature (C)

Provided below is an example of the information found in this section:

```
 MASS PORTION OF BODY SEGMENT (KG) :

                CENTER      MUSCLE      FAT        SKIN
-------------------------------------------------------
HEAD            4.047       0.744       0.372      0.268
TRUNK          11.113      17.549       7.067      1.346
UP-ARM LE       0.562       0.842       0.242      0.121
UP-ARM RI       0.562       0.842       0.242      0.121
LO-ARM LE       0.562       0.842       0.242      0.121
LO-ARM RI       0.562       0.842       0.242      0.121
HAND LE         0.130       0.037       0.074      0.093
HAND RI         0.130       0.037       0.074      0.093
UP-LEG LE       2.078       3.053       0.714      0.359
UP-LEG RI       2.078       3.053       0.714      0.359
LO-LEG LE       1.385       2.035       0.476      0.240
LO-LEG RI       1.385       2.035       0.476      0.240
FOOT LE         0.216       0.037       0.112      0.119
FOOT RI         0.216       0.037       0.112      0.119
-------------------------------------------------------
TOTAL  (KG) :   71.890
```

Boundary Conditions at Beginning of Computation

This section of the `passenger_tim_xx.txt` file contains information about the boundary conditions on each of the passenger body segments at the start of the run.

The following conditions are detailed:

- Heat flux coefficients:
  - Convection
  - Radiation
  - Clothing
  - Heat transfer coefficient
  - Effective temperature (C)
- Temperatures and coefficients:
  - Air temperature
  - Confinement surface temperature
  - Thermal efficiency factor
  - Permeation efficiency factor

Final Thermo-Physiological Quantities

This section of the `passenger_tim_xx.txt` file contains results relating to the passenger vascular system, heat production, and heat loss.

```
THERMO-PHYSIOLOGICAL QUANTITIES  :
================================

     CENTRAL BLOOD  :
         TEMPERATURE (C)              =
           HEAT TRANSPORT             =


     HEAT PRODUCTION (W) :
          BREATHING                   =
          ADDITIONAL ACTIVTIY         =
          ENTIRE CONVERSION           =


     HEAT LOSS (W) :
         BREATHING            :
         SKIN      :
           SWEATING AND EVAPORATION  =
           DRY HEAT LOSS             =
           OVERALL                   =
          BODY   :
          SKIN AND RESPIRATION       =


     HEAT BALANCE (W) :
          PRODUCTION -  LOSS         =


     BLOOD VOLUME FLOW  (CM3/S) :
          BREATHING                  =
          CONTROL PORTION            =
          ENTIRE CONVERSION          =
```

Values for Body Segments

This section of the `passenger_tim_xx.txt` file contains tabulated data regarding body segments and segment layers:

This section of the `passenger_tim_xx.txt` file summarizes the following properties:

- Heat Production (MW)
- Heat flux from skin (W)

- Blood flow (cm^3/s)
- Heat transfer between layers (W/K)
- Temperature difference (K)
- Temperature allocation (C)
  - Mean skin temperature (C)
- Surface temperature (C)
- Heat balance (W)
- Heat flux from skin (W)
- Evaporation heat balance (W)
- Water vapor pressure (Pa)

Provided below is an example of the information found in this section:

```
 SEGMENT HEAT PRODUCTION  (MW) :

                CENTER        MUSCLE       FAT        SKIN
-------------------------------------------------------------
HEAD         14718.80        435.11      130.80      62.90
TRUNK        47916.20      28491.39     2560.80     325.00
UP-ARM LE      184.80       1479.28       97.60      21.80
UP-ARM RI      184.80       1479.28       97.60      21.80
LO-ARM LE      184.80       1479.28       97.60      21.80
LO-ARM RI      184.80       1479.28       97.60      21.80
HAND LE         48.30        279.51       31.60      19.90
HAND RI         48.30        279.51       31.60      19.90
UP-LEG LE      721.98      10895.77      255.60      92.88
UP-LEG RI      721.98      10895.77      255.60      92.88
LO-LEG LE      481.32       7263.85      170.40      61.92
LO-LEG RI      481.32       7263.85      170.40      61.92
FOOT LE         82.40        271.61       46.80      33.50
FOOT RI         82.40        271.61       46.80      33.50
-------------------------------------------------------------
TOTAL  (W)  :  143.29
```

Boundary Conditions at End of Computation

This section of the `passenger_tim_xx.txt` file contains information about the boundary conditions on each of the passenger body segments at the end of the run.

The following conditions are detailed:

- Heat flux coefficients:
  - Convection
  - Radiation

  - Clothing
  - Heat transfer coefficient
  - Effective temperature (C)
- Temperatures and coefficients:
  - Air temperature
  - Confinement surface temperature
  - Thermal efficiency factor
  - Permeation efficiency factor

Bibliography

This section lists the bibliographic references for the Thermal Comfort wizard.

Meeh, K. 1879. "Oberflächenmessung des menschlichen Körpers", J. Biol., 15, pp. 428-458

DuBois, D. and DuBois, E.F. 1915. "Clinical Calorimetry: Tenth paper, A formula to estimate the approximate surface area if height and weight be known", Arch. Intern. Med., 17, pp. 863-871

Stolwijk, J.A.J. 1971. "A mathematical model of physiological temperature regulation in man", NASA Contractor Report, NASA CR-1855

Berbner, D.F., Kerslake, D.McK., and Waddel, J.L. 1958. The relation between the coefficients for the heat exchange by convection and by evaporation in man", J. Physiol., 141, pp. 164-168

Aschoff, J. and Wever, R. 1958. "Kern und Schale im Wärmehaushalt des Menschen", Die Naturwissenschaften, 45, Nr. 20, pp. 477-485

Stolwijk, J.A.J. and Hardy, J.D. 1966. "Temperature regulation in man - A theoretical study", Pflüger Arch., 307, pp. 129-162

Hertzman, A.B. and Randall, W.C. 1948. "Regional differences i the basal and maximal rates of blood flow in the skin", J. Appl. Physiol., pp. 234-241

Benzinger, T.H., Kitzinger, C., and Pratt, A.W. 1963. "The Human Thermostat", In: Temperature. Its Measurement and Control in Science and Industry, 3, Part 3, Biology and Medicine, Hardy, J.D. (Editor)

Nadel, E.R. 1969. "The role of peripheral thermoreceptors in the integrated thermoregulatory response in resting and excercising man", Dissertation, Univ. Calif., Santa Barbara

Using the Wing Mesher

This section introduces the Wing Mesher tool.

The Wing Mesher tool provided in STAR-CCM+ is used for the rapid meshing of aerodynamic wing type geometries which can then be used for analysis within the software. An example of such a wing is shown below:



This section covers how to use the Wing Mesher tool to set up the wing definition and produce a mesh based on a simple example.

What is the Wing Mesher?

The Wing Mesher is an automatic meshing tool for creating high-quality hexahedral meshes for wing type geometries.

The Wing Mesher is an automatic meshing tool for creating high-quality hexahedral meshes for wing type geometries quickly so that they can be included as part of a larger flow domain. It can be used for a single wing geometry or for multiple wings positioned one after the other. An example mesh for a single wing is shown below:



The hexahedral mesh is generated using a multi-block approach and is conformal throughout, resulting in an accurate and efficient grid structure for the solver to use. Automatic clustering of the mesh, both away from the airfoil to the far field and in the vicinity of the trailing edge, is included to further enhance and optimize the mesh characteristics.

The box surrounding the wing(s), can either be explicitly defined (using the simple shape parts within STAR-CCM+ if necessary) or can be implicitly represented based on corner coordinates, which by default are calculated by the tool. Each wing surface however, must have a blunt trailing edge (ie not taper to a single line) and be composed of

reasonable quality triangular faces. The surface repair tools and surface remesher can be used to ensure these conditions if they are not present in the originating surface.

The wing profile imported into STAR-CCM+ can include taper and generally change size, shape, and profile along its span. When two or more wings are present then the downstream wings can overlap in the flow direction, although they should not be touching.

See Also

Creating Simple Shape Parts
Preparing Surfaces
Surface Remesher

Setting Up the Wing Mesher

This section describes how to use the Wing Mesher tool to generate volume mesh based on an imported surface description.

In general, the setup of the Wing Mesher properties is relatively simple and quick to perform. Most of the effort required to produce a successful volume mesh is in the preparation of the wing surface.

The steps undertaken to achieve this are:

Importing the Surface Geometry

A surface representation of some kind must be present before the Wing Mesher tool can be accessed.

The surface can either be imported from an external file or alternatively be created directly within STAR-CCM+ using the 3D-CAD module. If multiple wing geometries are being considered at the same time, then they should all be imported before using the Wing Mesher tool.

Once the surface geometry is imported, it should be prepared appropriately for the Wing Mesher.

See Also

What Is a 3D-CAD Model
Importing Geometry
Surface and Geometry Requirements

Surface and Geometry Requirements

In general, the wing surface and the bounding box (if included) needs to be of a reasonable triangle quality in order for the Wing Mesher to be successful.

If the imported surface is composed for STL or similar style triangulation, then the surface remesher should be first used to improve the quality. An example triangulation for the wing is shown below:



Each wing surface can be composed of one or more unique boundaries. These boundaries cannot be used for the bounding box or other wings. If a bounding box is included, then the wing and box surface must form a single, closed, manifold surface. This means that the ends of the wing sides are “open”, as shown below:



The open wing end condition is also true for wing surfaces that are not contained in a bounding box. In other words, when using coordinates to define the mesh domain bounds, the wing surface is simply the "wetted" surface and is open on both ends.

If you decide to use a bounding box to define the wing exterior domain, then this box should not extend more than a chord length or so upstream, downstream, above and below the wing. In the other words, the domain represents the near field of the wing and does not extend to the far field. Each side of the box should also be represented by a single unique boundary. An example bounding box geometry is shown here:



If the bounding box is not included with the imported wing, then it can be created using the simple shape tools and then intersected with the wing surface. The surface repair diagnostics and tools should then be used for any final required clean up to meet the requirements outlined above.

It is also critical that each wing have a blunt trailing edge and not taper to a single line. An example wing is shown below:



Finally, the wing axis should generally align with one of the laboratory coordinate directions as best as possible.

In summary:

- The wing surface and bounding box (if included) should be of a reasonable quality triangulation.
- One or more unique boundaries can be used for each wing.
- If a bounding box is used, then the wing and box must form a single, closed and manifold surface. The bounding box upper, lower, front, and rear surfaces should also be approximately no more than one chord length away from the wing. Each box side should also be represented using a single unique boundary.
- The wing ends must be open in all instances.
- Each wing must have a blunt trailing edge.
- The wing should be aligned to the laboratory cartesian system as best as possible.

If any of the above criteria are not met, then the wing mesher is likely to fail during meshing or not proceed at all. The next step after preparing the surface is to launch the Wing Mesher tool.

See Also

Defining the Bounding Box
Creating Simple Shape Parts
Preparing Surfaces
Merging Intersecting Boundaries
Surface Remesher

Launching the Wing Mesher Tool

This section describes how to launch the Wing Mesher tool.

The Wing Mesher launch button is not directly accessible in default STAR-CCM+ interface. The tool has to be first added to a toolbar before it is launched. The option to create a toolbar and add the Wing Mesher button to it is provided in the **Window** > **Toolbars** > **Customize** dialog. Select this option as shown below:



The *Customize Toolbar* dialog appears:



The wing mesher button is located near the end, under the **Wizards** node manager and is called **Import Wings**. The button can be dragged and dropped to an existing toolbar in the interface (it is recommended you add it to the end of an existing one) or alternatively a new empty toolbar can be created and named using the **New Toolbar...** button. After dragging and dropping the **Import Wings** button, it will look something like this:

The *Customize Toolbars* dialog can be closed and the **Import Wings** button in the toolbar can be pressed to launch the wing mesher options panel. This panel is shown below:



Should more than one surface representation exist on attempting to launch the Wing Mesher panel, then the following dialog is seen:



In this instance, the desired surface representation to keep should be selected and the **OK** button pressed to proceed and open the panel.

Setting the Number of Wings

The first step in defining the Wing Mesher properties is to set the number of wings.

This is done using the **Number of wings** selector as shown below:



By default, the panel is set up to mesh one wing but can deal with any number. The appropriate number should be selected, equal to the number of wings in the surface representation to be used for the meshing.

The wing boundary selection can then be made.

Wing Boundary Selection

For each wing to be meshed, the wing boundary assignment needs to be made.

The current wing to be specified should first be selected using the **Select current wing** option before the boundary selection is made using the list boxes:



For example, for the first wing, the current wing value should be set to 1 and the appropriate boundary or boundaries moved from the left-hand side list box to the right-hand side one, as shown above. The process can then be repeated for any additional wings.

If there are multiple wings, then they should be numbered incrementally in the flow direction. In other words, the most upstream wing is the first one.

One the boundary selection for the wing(s) has been made then the flow direction and wing axis specification can be made.

Setting the Flow Direction and Wing Axis

This section describes how to set the flow direction and wing axis.

The flow direction and wing axis definitions are shown below:



As shown above, the flow direction is defined as being from the wing leading edge to the trailing edge (positive X-direction in this case) and the wing axis is defined as spanning from the left side to the right side when looking at the wing from the front (positive Z-direction in this case).

The appropriate specifications should be made using the **Select Flow Direction** and **Select Direction of Wing Axis** options in the panel:



The specification of the flow direction and wing axis is important as it provides guidance to the mesher as to where the leading and trailing edges of the wing are. The flow direction does not need to match exactly with the laboratory cartesian ordinates but should generally be aligned to one of them.

The grid resolution around the wing can be set next.

Setting the Grid Resolution

It is generally recommended that the meshing tool automatically determine the grid resolution.

The default behavior of the Wing Mesher tool is to automatically calculate the grid resolution around the wing itself. This state is shown below:



If you want to supply your own parameters, then the **Use User Defined Grid Parameter Settings** option should be turned on:



Currently, only the first three parameters listed in the tool are available to be used. The parameter determining the number of layers wrapping the wing refers to the number of layers both above and below the wing.

Defining the Bounding Box

The Wing Mesher is designed to produce a C-type hexahedral grid within a regular block shape.

The block, or bounding box, size can be determined in one of two ways:

- Implicitly, by supplying the box min and max coordinates.
- Explicitly, by including a set of boundaries to define the box faces.

The default behavior of the panel is to automatically calculate the min and max box coordinates based on the wing geometry size, as indicated by the **XYZ-Coordinates** option for **Define range of the grid bounding box**:



The values can be left as they are or modified accordingly to provide a different sized box. The distance between the box sides must be equal to or smaller than the wing span dimension. In other words, they cannot extend beyond the span of the wing. The box should also not be too small and encroach on the wing surfaces. In general, a box extending one wing chord length above, below, upstream and downstream of the wing is recommended as a good starting point.

As an alternative, the box boundaries can be used (if they are available) to define the bounding box. Select the **Selected Boundaries** option for **Define range of the grid bounding box** to allow the individual boundaries to be selected.



The appropriate boundary selection should be made for each of the six box sides, as shown above. Only one boundary can be selected for each

side of the box and no two boundaries can be shared. The box shape need not align with the coordinate system and can even be curved in the flow direction.

Note that if you are using wing geometries that are highly curved then using the default (implicit) coordinate based box option is recommended.

You are now ready to generate the mesh.

Generating the Mesh

The meshing process creates a volume mesh representation.

It is recommended that the simulation file be saved before starting the mesh generation process since irreversible changes to the surface representation occur. Pressing the **OK** button on the panel starts the meshing process. If successful, a volume mesh representation will be created after a minute or two (depending on the number of wings and grid density) which can then be displayed as normal.

An example wing is shown below:



Various integrity checks are done at the beginning of the process to check that the wing and box alignment is correct. If there are errors in the setup, then these are reported and need to be corrected before retrying the meshing process once more.

Introduction to the SPEED Motor Tool Import Wizard

This section describes the STAR-CCM+ SPEED Motor Tool Import Wizard.

SPEED is a software package that CD-adapco provides for use in the design and analysis of electrical machines. It provides a wide array of machine design options, including a material database and analysis programs to evaluate machine performance. The STAR-CCM+ SPEED Motor Tool Import Wizard lets you import SPEED electrical machine data into STAR-CCM+ for further analysis, such as calculating 3D magnetic fluxes within the electrical machine.

Importing SPEED Data Into STAR-CCM+

You can import SPEED data into STAR-CCM+ using a geometry definition file (`.xGDF`).

In the general SPEED finite element analysis workflow, after you have defined the machine geometry using the SPEED design module, you can use the GoFER tool to set up a finite element analysis. The GoFER tool creates an `.xGDF` file that contains the machine design and the selected material properties.

The `.xGDF` file contains the following information:

- Geometry of the machine
- Material properties and geometry of each subdomain
- Eccentricity details
- Skew details
- Material property tables that describe the magnetic flux density **B**, magnetic field strength **H**, and permeability **B/H** of the motor components.
- Motor type details that are required for modeling the endwindings, and winding information from the SPEED winding editor.

  STAR-CCM+ currently supports importing the following for the stator:

  - Simple windings

    Simple windings are defined in the STAR-CCM+ SPEED Motor Tool Import Wizard. These windings are single-layer (one coil per slot), where the span of each winding is one. Therefore, for a twelve-slot stator, six coils are created. If you choose not to run an electromagnetic simulation, these windings are further simplified by merging the individual windings into one winding body. These merged windings are referred to as "simple-bulk" or "bulk-type" windings. For simple windings, the Import Wizard is primarily used to specify the shape of the endwindings.

  - Complex windings

    Complex windings are defined in the `.xGDF` file. These windings can be single-layer or two-layer (two coils per slot). A slot can also be empty. The span of each winding can be greater than one (the winding does not return to its neighboring slot) and spans can also be negative. For this reason, complex windings are also known as distributed windings.

STAR-CCM+ currently supports single-layer, complex windings, for all spans. STAR-CCM+ also supports two-layer, complex windings but only if the winding spans are all one. STAR-CCM+ currently does not support many realistic two-layer, complex windings for spans greater than one.

The `.xGDF` file contains no information on the exterior flow domain. You can specify this information manually when you import the `.xGDF` file.

You can import a `.xGDF` file that contains a symmetrical sector of the machine design, rather than the complete geometry. The sector can be any appropriate section that utilizes the motor symmetries, such as a half sector, quarter sector, sixth sector, or eighth sector. For more information, see Importing Symmetry-Reduced Models.

Accessing the SPEED Motor Tool Import Wizard

This section describes how to make the SPEED Motor Tool Import Wizard accessible in STAR-CCM+.

To make the SPEED Motor Tool Import Wizard accessible in STAR-CCM+, add the SPEED Motor Tool Import Wizard icon to the toolbar:

1. In the STAR-CCM+ menu, select **Window** > **Toolbars** > **Customize**.



2. In the *Customize Toolbars* dialog, under the **Wizards** node, select SPEED Motor Tool Import Wizard.



3. Drag the SPEED Motor Tool Import Wizard icon onto one of the toolbars at the top of the STAR-CCM+ workspace and drop it at the appropriate position.

The icon is added to the toolbar at the specified location.

4. In the *Customize Toolbars* dialog, click **Close**.

Using the SPEED Motor Tool Import Wizard

This section describes how to use the SPEED Motor Tool Import Wizard.

To use the SPEED Motor Tool Import Wizard:

1. In an open simulation, click the (**SPEED Motor Tool Import Wizard**) button in the STAR-CCM+ toolbar.

   The **SPEED Motor Tool Import Wizard** dialog appears.



2. In the **SPEED Motor Tool Import Wizard** dialog, select the appropriate *SPEED Motor Tool Import Options*:

Electromagnetic Simulation	Specifies the simulation type. This option is activated by default. If you deactivate this option, the simulation is for cooling analysis only. For more information, see Setting the Simulation Type.
2D Simulation	Imports the model for a 2D simulation. When this option is enabled, the geometry is created as a thin slice suitable for 2D meshing. The 3D geometrical features that are specified in the `.xGDF` file, such as stator windings, rotor field windings, tilted axis eccentricity, and skew, are not imported, and no end-windings are generated. For more information on 2D meshing, see Creating a True Two-Dimensional Mesh.
Build Exterior Domain	Builds an exterior domain according to the specified parameters. This option is activated by default and is required for 2D simulations. If you do not want to create an exterior domain, deactivate this option. For more information, see Setting up an Exterior Domain.
Build Rotor Windings / Cages	When this option is activated (the default setting), STAR-CCM+ creates either rotor field windings or a rotor cage according to the type of motor that is imported. For more information, see Setting up Windings.
Build Stator Windings	When this option is deactivated (the default setting), the stator windings are not modeled in the motor geometry. If this option is activated, STAR-CCM+ builds the stator windings according to the **Simplified Stator Windings** option setting. For electromagnetic simulations, windings are required. For cooling simulations, windings are optional: if you want to simplify the geometry, you can ignore the windings. However, if you include the windings, the imported windings are modeled as Bulk-type windings in the

	motor geometry. Under the **Simplified Stator Windings** setting, you can specify either Bulk-Tub or Bulk-Out shaped windings. The Bulk-Tub windings are more common. For more information, see Setting up Windings.
Simplified Stator Windings	When this option is activated (the default setting), STAR-CCM+ builds the motor windings according to the parameters that are specified in the *Stator Winding Endturn Shape Options* group box. If this option is deactivated, the windings information is imported from the `.xGDF` file. For more information, see Setting up Windings.

3. Under *Load Geometry Definition File*, specify the `.xGDF` file that you want to import into STAR-CCM+:
   1. Click **Browse** and then, in the *Open* dialog, select the appropriate `.xGDF` file.



   2. Click **Open** to load the selected file.
4. In the **SPEED Motor Tool Import Wizard** dialog, click **OK** to confirm the settings.

The SPEED Motor Tool Import Wizard imports the specified `.xGDF` file into STAR-CCM+.

For more information, see Using SPEED Data in STAR-CCM+.

Setting the Simulation Type

By default, the SPEED Motor Tool Import Wizard sets up an electromagnetic simulation. However, you can choose to set up a cooling simulation instead.

An electromagnetic simulation is a transient simulation. STAR-CCM+ turns on the relevant electromagnetic solvers in physics continua, invokes sliding mesh for rotor motion, and creates electromagnetic-specific reports and field functions.

A cooling simulation is a steady-state simulation. STAR-CCM+ invokes multiple reference frames for rotor motion.

You can specify that the simulation is a 2D simulation. When this option is enabled, the geometry is created as a thin slice suitable for 2D meshing. The 3D geometrical features that are specified in the `.xGDF` file, such as stator windings, rotor field windings, tilted axis eccentricity, and skew, are not imported, and no end-windings are generated.

Choosing a 2D simulation implies an electromagnetic simulation, so the **Electromagnetic Simulation** option is activated automatically when you select the **2D Simulation** option. You can deactivate **Electromagnetic Simulation** if desired.

The external domain is required for 2D simulations, so the **Build Exterior Domain** option is activated automatically. You can choose **Cylindrical** or **Hexahedral** to specify the domain shape, but you cannot modify the domain size: the **Custom Size** option is disabled for 2D simulations.

Windings cannot be included in 2D simulations. When you select the **2D Simulation** option, the options for **Build Rotor Windings / Cages**, **Build Stator Windings**, and **Simplified Stator Windings** are disabled.

Note: If the stator windings are complex and of the Full, Side-by-Side (two-layer) type, deactivate the **Simplified Stator Windings** option before you select the **2D Simulation** option. This setting is required to split the "Slot" bodies to accommodate the winding configuration.

To set the simulation type:

1. Under *SPEED Motor Tool Import Options*, activate or deactivate **Electromagnetic Simulation**, whichever is appropriate.



2. If you want to specify a 2D simulation, activate **2D Simulation**.

Setting Up an Exterior Domain

The .xGDF file specifies the motor geometry only: it does not include any exterior domain information. You can let the SPEED Motor Tool Import Wizard build an exterior (air flow) domain as part of the import process.

Alternatively, you can build an exterior domain manually in STAR-CCM+ as a separate action.

To set up an exterior domain as part of the import process:

1. Under *SPEED Motor Tool Import Options*, select **Build Exterior Domain**.



2. Specify the shape of the domain by selecting the appropriate option: **Cylindrical** or **Hexahedral**.

   The cross-section of the exterior domain can be circular or square. If circular, the resulting 3D shape is cylindrical. If square, the resulting 3D shape is planar, hexahedral.

   The default dimensions are automatically calculated to enclose the entire motor, including any endwindings. If necessary, you can modify the created domain in STAR-CCM+.
3. If you want to make the exterior domain larger than the default size, select **Custom Size** and then set the appropriate dimensions:

Radius	The radius of the cylinder, or the X and Y dimensions of the hexahedron.
Negative Z-coordinate	The Z dimension in the negative direction. The specified value must include the negative sign.
Positive Z-coordinate	The Z dimension in the positive direction.

   The dimension values must include the units.

The SPEED Motor Tool Import Wizard ensures that the size of the exterior domain is large enough to enclose the motor, including the endwindings. If you specify a dimension that is too small, it is corrected automatically when the .xGDF file is imported into STAR-CCM+. The *Output* window displays details of any corrections that are made during the import process.

Setting Up Windings

The `.xGDF` file can include both rotor and stator winding information from the SPEED winding editor. You can choose to import this information or ignore it. For stator windings you have the option to create simplified windings instead.

There are three classes of windings: Bulk, Simple, and Complex. Bulk and Simple windings are simplified representations of the windings. STAR-CCM+ imports these windings as follows:

- Bulk windings are imported into a single region named **Windings**.

- Simple windings are imported into a separate region for each winding. These regions have simple names such as **Winding 1**, **Winding 2**, and so on.
- Complex windings are imported into a separate region for each winding. These regions have detailed names that include the phase number and the coil number that are taken from the `.xGDF` file winding data.

STAR-CCM+ can import the following types of windings from the `.xGDF` file:

Stator	Simple windings Single coil side type Complex windings Two-layer coil side type Complex windings: • Top-Bottom • Full Top-Bottom • Side-by-Side • Full Side-by-Side
Rotor	For Induction motors, the rotor field windings are imported as Bulk Tub-type windings. For all other motors, the rotor field windings are imported as Simple Tub-type windings (separated windings).

Any other types of complex windings that are specified in the `.xGDF` file cannot be imported into STAR-CCM+.

To include windings in the imported model, activate the appropriate setting in the *SPEED Motor Tool Import Wizard* dialog.

An example of a motor with single coil side type complex windings imported from a `.xGDF` file is shown below:



An example of a motor with two-layer top-bottom coil side complex windings imported from a `.xGDF` file is shown below:



An example of a motor with two-layer side-side coil side complex windings imported from a `.xGDF` file is shown below:



An example of a non-induction type motor sector model with rotor windings imported from a `.xGDF` file is shown below:



If you want to create simplified windings, specify the appropriate parameters. The SPEED Motor Tool Import Wizard lets you specify simple stator windings only: the endturn span is always one, that is, the winding connects adjacent slots. An endturn span of one is typical for PC-SRD motors. Rotor windings and complex stator windings, where the endturn span is greater than one (that is, the winding spans multiple slots), cannot be specified in the SPEED Motor Tool Import Wizard - they must be imported.

If you specify simplified windings, the Slot (PC-BDC) or Coil (PC-SRD) regions are removed from the imported geometry. Winding regions are added in the appropriate location.

For electromagnetic simulations, windings are required. For cooling simulations, windings are optional. However, if you include the windings, the imported windings are modeled as Bulk-type windings in the motor geometry. Under the **Simplified Stator Windings** setting, you can specify either Bulk-Tub or Bulk-Out shaped windings. The Bulk-Tub windings are more common.

To set up windings:

1. Under *SPEED Motor Tool Import Options*, select the appropriate options:
   - **Build Rotor Windings / Cages**

When this option is activated (the default setting), STAR-CCM+ creates either rotor field windings or a rotor cage according to the type of motor that is imported.

For induction motors imported from PC-IMD which have solid bodies designated as “Cage…”, the end-ring geometry for the cage is created. STAR-CCM+ creates the geometry according to the dimensions of the “Cage” bodies, and you cannot modify the shape of the end-rings.

For other motor types, such as those imported from PC-BDC, when motors have solid bodies designated at “Field Coil”, the endturn geometry for the field coils is created. STAR-CCM+ creates the geometry according to the dimensions of the “Field Coil” bodies, and you cannot modify the shape of the rotor field winding endturns.

In both cases, the only information that is imported from the `.xGDF` file is the 2D cross-section starting from the rotor. The `.xGDF` file does not supply any information for the geometry of either the cage end-rings or rotor field winding endturns.

Examples of each case are shown below.

Cage end-rings created from “Cage 1” solid bodies for induction motors (PC-IMD):



Rotor field windings created from “Field Coil” solid bodies for brushless, DC motors (PC-BDC):



- **Build Stator Windings**

  When this option is deactivated (the default setting), the stator windings are not modeled in the motor geometry.

  If this option is activated, STAR-CCM+ builds the stator windings according to the **Simplified Stator Windings** option setting.



If this option is deactivated, the windings are not modeled in the motor geometry.

2. Set the **Simplified Stator Windings** option as appropriate:

   - If this option is deactivated (the default setting), the windings information is imported from the `.xGDF` file.

     STAR-CCM+ automatically calculates the axial extensions (the straight part of the coil between the stator and elbow) so that the windings do not interfere.

- When this option is activated, STAR-CCM+ builds the motor windings according to the parameters that are specified in the *Stator Winding Endturn Shape Options* group box.

  For an electromagnetic simulation, any windings information in the `.xGDF` file is ignored. For a cooling simulation, the **Simplified Stator Windings** setting specifies the general shape of the endwindings: either Bulk-Tub or Bulk-Out shaped windings. In general, bulk-tub shaped windings provide a better description of the true winding geometry. However, to better model the correct volume covered by true winding geometry, the winding data from the `.xGDF` file is used to modify the thickness of the bulk-type winding body (that is, the toroidal thickness).

  For more information, see Cooling Simulation Endwindings.

3. Under *Stator Winding Endturn Shape Options*, specify the winding endturn shape by selecting the appropriate option: **Tub**, or **Out**.

   The default shape is **Tub**.

4. Set the appropriate dimensions:

Elbow 1 Radius	This parameter applies to both winding shapes.
Extension 1 Length	The straight extension from elbow 1 to elbow 2. This parameter applies to **Out** windings only.
Elbow 2 Radius	This parameter applies to **Out** windings only.

   The dimension values must include the units.

The SPEED Motor Tool Import Wizard automatically calculates a minimum and maximum set of values, which are based on the motor geometry.

The minimum dimensions require minimum physical space while ensuring, in most cases, that the endwinding does not interfere with itself. The maximum dimensions ensure that the endturn does not go beyond the adjacent slot, as the span must equal one.

You can obtain the minimum dimensions by entering `0` for each parameter. You can obtain the maximum dimensions by entering a large number, such as `1000000`.

The maximum value for some parameters may be unrealistic. The *Output* window displays details of any corrections that are made during the import process.

Cooling Simulation Endwindings

For a cooling simulation, the imported windings are modeled as Bulk-type windings in the motor geometry.

The **Simplified Stator Windings** setting specifies the general shape of the endwindings: either Bulk-Tub or Bulk-Out shaped windings.

The thickness of the endwindings toroid in the imported model is defined as follows:

- If no windings information is available in the `.xGDF` file, the endwinding thickness is set equal to the thickness of the slot face sketch.
- If windings information is available in the `.xGDF` file, the thickness of the endwindings toroid is calculated as:

  $$\frac{1}{2} \times \text{maximum span} \times \text{coil thickness}$$

  For example, if the maximum span is 10, the endwindings toroid thickness is 5 times the thickness of the slot or coil face sketch from which the winding originated. If no explicit coil face is provided, the slot is assumed to be single layered, meaning one coil per slot.

  The span value is a non-zero integer and can be negative, so the absolute value is used. If the maximum span is 1 or 2, the thickness of the endwindings toroid is set to the coil thickness.

The model that is shown below has the endwindings toroid that is calculated for a maximum span of 3.



The example below shows the same model with the endwindings toroid calculated for a maximum span of 6.



Setting Up Eccentric Motors

The SPEED Motor Tool Import Wizard lets you import eccentricity information into STAR-CCM+.

If you design an eccentric (rotor off-center) motor in SPEED, the eccentricity information is included in the `.xGDF` file. The SPEED Motor Tool Import Wizard lets you import the eccentricity information into STAR-CCM+. However, if you want to modify the eccentricity parameters, you can edit the `.xGDF` file before importing it.

Note: Any changes that you make to the `.xGDF` file are at your own risk. The eccentricity values are imported into STAR-CCM+ as is, and are not checked or corrected.

Eccentricity is specified in the `.xGDF` file as follows:

```
<eccentricity>
     <mode> 0 </mode>
     <dx> 0 </dx>
     <dy> 0 </dy>
     <dz> 0 </dz>
     <dthx> 0 </dthx>
     <dthy> 0 </dthy>
  </eccentricity>
```

The parameters are as follows:

mode	Specifies the eccentricity mode. Valid values are `0` (static) or `1` (dynamic). Static eccentricity is a simple lateral displacement of both the rotor and the axis of rotation of the rotor away from the global Z-axis. Dynamic eccentricity is when the centroid of the rotor is displaced from its axis of rotation. The rotor wobbles around its axis of rotation, which is the global Z-axis.
dx, dy, dz	Specifies the lateral displacement of the rotor about the global origin (0, 0, 0). SPEED motors are centered about the global origin. Non-eccentric motors rotate about the global Z-axis. There are no restrictions on lateral displacements, and they can be used with a tilt displacement.
dthx, dthy	Specifies the X-tilt displacement and Y-tilt displacement of the rotor. The axis of rotation of the rotor is not parallel to the axis of the stator. In this configuration, the air gap profile at one end of the machine is different to the profile at the other end and the

	model includes axial forces as well as radial forces. X-tilt displacement rotates the rotor about the X-axis of the global origin (up or down), using the right-hand rule. Y-tilt displacement rotates the rotor about the Y-axis of the global origin (left or right), using the right-hand rule.

Setting Up Skewed Motors

SPEED lets you design skewed motors, where either or both the rotor and the stator are twisted along the rotation axis. Skew information is included in the `.xGDF` file.

The SPEED Motor Tool Import Wizard supports importing continuous skew settings and stepped skew settings into STAR-CCM+. The imported parts are separated into the appropriate regions, according to the skew mode:

- Skew mode 1 (continuous): One region per part.
- Skew mode 2 (stepped): One unique region per slice (step).
- Skew mode 3 (V-continuous): Two regions per part, split symmetrically.
- Skew mode 4 (V-stepped): One unique region per slice (step).

If you want to modify the skew parameters, edit the `.xGDF` file before you import it.

Note: Any changes that you make to the `.xGDF` file are at your own risk. The skew values are imported into STAR-CCM+ as is, and are not checked or corrected.

Skew is specified in the `.xGDF` file as illustrated in the following example:

```
<skew>
  <rotor>
    <mode> 1 </mode>
    <angle> 30 </angle>
    <slices> 2 </slices>
    <sections> 1 </sections>
    <symmetry> 0 </symmetry>
  </rotor>
  <stator>
    <mode> 0 </mode>
    <angle> 0 </angle>
```

```
    <slices> 1 </slices>
    <sections> 1 </sections>
    <symmetry> 0 </symmetry>
  </stator>
</skew>
```

The parameters are as follows:

rotor	Specifies the skew settings for the rotor.
stator	Specifies the skew settings for the stator.
mode	Specifies the skew mode. Valid values are: • `0` - none • `1` - linear • `2` - stepped • `3` - V-linear • `4` - V-stepped
angle	Specifies the skew angle about the Z-axis of the global origin, using the right-hand rule. The angle can be negative or positive. Note the following: • if mode = `0`: the angle = `0` and is ignored • if mode = `1` or `2`: the angle is across the full length of the section. • if mode = `3` or `4`: the angle is across half the length of the section.
slices	Specifies the number of steps in a stepped skew model: • if mode = `0`, `1`, or `3`: the number of slices = `1` and is ignored. • if mode = `2` or `4`: the number of slices = the number of steps across the length of the section.
sections	Specifies the number of repeatable units of the model, as described by the mode, angle, and slices settings.
symmetry	Specifies the symmetry of the sections when the model contains two or more sections. The valid values are:

	• `0` - each section is a repeat of the previous section. • `1` - each section is a mirror image of the previous section. If sections = `1`, symmetry = `0`.

The following diagrams illustrate the settings that you can make for each skew mode. Each example shows a single section, and the skew angle *a* is the same for each.



1	Skew mode is Linear; skew angle is *a*.
2	Skew mode is Stepped; skew angle is *a*; number of slices is 5.



3	Skew mode is V-Linear; skew angle is *a*.
4	Skew mode is V-Stepped; skew angle is *a*; number of slices is 5.

Importing Symmetry Reduced Models

This section describes the symmetry-reduced models that you can import.

STAR-CCM+ lets you import a `.xGDF` file that contains a symmetrical sector of the machine design, rather than the complete geometry.



The sector can be any appropriate section that utilizes the motor symmetries, such as a half sector, quarter sector, sixth sector, or eighth sector.



As part of the import process, all body faces on the two periodic planes are given a unique face name. An example is shown below.



Importing a sector model lets you exploit axisymmetric periodic geometry to reduce the cell count and the simulation time. STAR-CCM+ automatically creates the appropriate periodic conditions between the two symmetry boundaries.

However, depending on the requirements of the simulation, you may need to use the full geometry. If necessary, you can use the 3D-CAD modeler to create the full model, or use SPEED to create a Whole Motor model that you can import.

If the reduced sector domain cuts through a slot, the windings can be incorrect in the imported model. To resolve this issue, you can expand the domain to include a half span to the adjacent slot.

See Also

Working with 3D-CAD Models

Magnet Coordinate Systems

STAR-CCM+ creates magnet coordinate systems for motors that have permanent magnets of the Cartesian type.

Radially magnetized magnets use the global cylindrical coordinate system that STAR-CCM+ creates automatically.

A region is created for each magnet that is defined in the `.xGDF` file, as shown in the example below.



STAR-CCM+ creates a unique coordinate system for each magnet region. In each magnet coordinate system, the Z axis points in the same direction as the global Z axis. The X axis is typically perpendicular to the base of the magnet (the face nearest to the axis of rotation). An example is shown below:



Skewed Motors

For linearly skewed motors (mode 1), STAR-CCM+ creates one region for each magnet. For V-linear skewed motors (mode 3), STAR-CCM+ creates two regions per magnet. For stepped skewed or V-step skewed motors (modes 2 and 4 respectively), STAR-CCM+ creates one region for each skew slice, per magnet. This convention is repeated for each skew section, regardless of the skew mode.

Illustrated below is an 8-magnet, PC-BDC motor with a stepped skewed rotor. The rotor is composed of three slices and only one section.



Therefore the number of magnet regions, and therefore the number of magnet coordinate systems, is *magnets x slices x sections*: 8 x 3 x 1 = 24.


magnetCoordinateSystemsSkew
Simulation
Scene/Plot
Coordinate Systems
Laboratory
Local Coordinate Systems
Magnetic Field Magnet 1 1
Magnetic Field Magnet 1 2
Magnetic Field Magnet 1 3
Magnetic Field Magnet 2 1
Magnetic Field Magnet 2 2
Magnetic Field Magnet 2 3
Magnetic Field Magnet 3 1
Magnetic Field Magnet 3 2
Magnetic Field Magnet 3 3
Magnetic Field Magnet 4 1
Magnetic Field Magnet 4 2
Magnetic Field Magnet 4 3
Magnetic Field Magnet 5 1
Magnetic Field Magnet 5 2
Magnetic Field Magnet 5 3
Magnetic Field Magnet 6 1
Magnetic Field Magnet 6 2
Magnetic Field Magnet 6 3
Magnetic Field Magnet 7 1
Magnetic Field Magnet 7 2
Magnetic Field Magnet 7 3
Magnetic Field Magnet 8 1
Magnetic Field Magnet 8 2
Magnetic Field Magnet 8 3
Geometry Scene 1
STAR-CCM+
Y
Z
X

Using SPEED Data in STAR-CCM+

This section summarizes how SPEED data is used in STAR-CCM+.

The SPEED Motor Tool Import Wizard can take a few seconds to load the data from the `.xGDF` file into STAR-CCM+. A 3D representation of the SPEED geometry appears in the *Geometry Scene*.

An example is shown below with the transparency activated.



Once the SPEED model has been imported into STAR-CCM+, you can treat it as any CAD geometry and make any necessary modifications.

In addition to importing the geometry, STAR-CCM+ does the following:

1. Generates parts, regions, and appropriate interfaces for the machine, and populates the Parts, Regions, and Interfaces nodes in the Simulation tree.
2. Creates a mesh continuum with the following models:
   - **Extruder**

- **Polyhedral Mesher**
- **Prism Layer Mesher**
- **Surface Remesher**



You can modify these settings if necessary, but it is recommended that you use the default models.

3. Creates the appropriate physics continua, such as Magnet, Rotor Steel, Slot, Stator Steel, and assigns each to the corresponding region. Each of these physics continua includes the Magnetic Vector Potential model.

   The example that is shown below is the Rotor Steel continuum.



4. If you selected the Build Exterior Domain option in the SPEED Motor Tool Import Wizard, STAR-CCM+ creates the appropriate bounding box to represent the solution domain.

   By default, the domain is sized automatically based on the dimensions of the machine. It is made up of Exterior (Outer) and Exterior (Inner) regions to represent the stationary and rotating domains. The machine is subtracted from both the outer and inner domains.



5. Creates plane sections at X = Y = Z = 0 and adds them under the **Derived Parts** node.
6. Generates the appropriate field functions and reports for the electromagnetic analyses.



7. Populates the **Tools** > **Tables** node with property tables based on the materials that are selected in SPEED.



These material property tables can be used as tabular data for the magnetic permeability material property of the Magnetic Vector Potential model. The tables contain columns for the magnetic flux density, magnetic field strength, and permeability. An example table is shown below:

Tabular Data

	Index	B	H	B/H
0	1	0	0	0
1	2	0.1	9.83	0.0102
2	3	0.2	19.66	0.0102
3	4	0.3	29.5	0.0102
4	5	0.4	39.3	0.0102
5	6	0.5	49.1	0.0102
6	7	0.6	58.3	0.0103
7	8	0.65	62.3	0.0104
8	9	0.7	66.18	0.0106
9	10	0.75	70.2	0.0107
10	11	0.8	74.77	0.0107
11	12	0.85	80	0.0106
12	13	0.9	86.3	0.0104
13	14	1	104.2	0.0096
14	15	1.1	130.5	0.0084

Export... Close

See Also

Setting Up an Exterior Domain

Using The IDF Import Wizard

The IDF Import Wizard imports IDF (Intermediate Data Format) files into STAR-CCM+.

IDF files contain design information about circuit boards and electronic components, and are widely used in the electronics cooling industry. The import process reads these files and builds the circuit board in 3D-CAD.

IDF files are neutral text files that typically contain board and library information. Board files contain information about the shape and size of a circuit board, and define the location and size of holes, and the placing of components. Library files contain information about the shape and size of the components. STAR-CCM+ supports the following board file and library file formats:

Board File	Library File
`.emn`	`.emp`
`.brd`	`.lib` `.pro`
`.brf`	`.ldf`

When using the Import IDF Tool, you select the board file. However, the library file must also be in the same directory to load the complete circuit board. If no library file is found, you are asked to locate it manually.

In the electronics cooling industry, the board and its components are all built from the XY plane. The board outline and the components are created from 2D profiles, which are then extruded outwards.

Importing IDF Files

This section describes how to use the IDF Import Wizard.

The IDF Import Wizard is hidden by default from the STAR-CCM+ user interface.

To import IDF files:

1. Add the Import IDF File tool to the STAR-CCM+ user interface. See Customizing a Toolbar.
2. Click the **Import IDF File** icon.

   The *Import IDF Options* dialog appears.



3. Set the desired import options. See Import IDF Options Panel.
   1. Filter out specific components:

      To omit specific components, use **Exclude** filters. Any selected components in the *Available Reference Designator Filter Components* list are excluded from the import. By default, all components are selected.

2. Filter out components based on size:

   To omit all components that are smaller than a specified size, add a Property Size filter. The size that you specify corresponds to the maximum size of a bounding box that fully encloses the 2D profile of the component.

3. Filter out holes on the circuit board:

   To omit holes from the circuit board, add a Hole Radius filter. Any holes which are smaller than the specified value are filtered out.

4. Click **OK**.

The specified components are read into STAR-CCM+ and are subsequently built in 3D-CAD. Each component is built using CAD operations. The benefit of building all components is that you can then modify the circuit board design in 3D-CAD in response to analysis results and re-run the simulation.

When the import completes, a 3D-CAD model named **IDF Import 3D-CAD 1** is added to the **Geometry > 3D-CAD Models** node. Edit this 3D-CAD model to enter the 3D-CAD environment.

Import IDF Options Panel

This section describes the *IDF Options* panel.

Use the following options to specify what components to import into STAR-CCM+.

Available Reference Designator Filter Components	Object selector panel that contains a list of components that are identified in the selected file. Select the components that you want to use with the *Reference Designator Filter Options*. By default, all components are selected.	
Reference Designator Filter Options	**Exclude**	Excludes the components that are selected in the *Available Reference Designator Filter Components* list.
Property Value Filter Option	**Property Size**	Specifies the size of the components to use in the filter. The size is defined as the

		maximum edge in the bounding box that encloses the 2D profile of the component. When added to the filter, components that are smaller than the specified size are not imported.
Geometry Feature Filter Option	**Hole Radius**	Specifies the radius of the hole to use in the filter. When added to the filter, holes which are smaller than the specified value are not imported.
Current Filter	Contains a list of active filters. To delete a filter (or groups of filters), select the filter and click **Delete Filter**.	

Simulation Assistants

A Simulation Assistant provides a visual workflow within the STAR-CCM+ user interface (UI).

By default, this visual workflow appears as a panel to the right-hand side of the Graphics window. The panel is composed of a series of steps, or tasks, that you work through in sequence. You can reposition the Simulation Assistant panel in STAR-CCM+ as if it were any other scene.



STAR-CCM+ provides a Java framework that allows you to create your own Simulation Assistants. You create a Simulation Assistant using a combination of Java classes from the Simulation Assistant framework and Java macros that you record from STAR-CCM+. You require a

working knowledge of Java in order to create a Simulation Assistant. When complete, you can share Simulation Assistant with other users, who can load them directly into STAR-CCM+. You do not require any knowledge of Java to use a Simulation Assistant.

If you are provided with an existing Simulation Assistant, you can load this into the STAR-CCM+ UI.

Benefits of Using a Simulation Assistant

There are numerous benefits to having an assistant in your simulation:

- New users can learn how to use STAR-CCM+ faster by focusing on the topic of interest, and understand the process workflow.
- Existing users can save time by following a pre-defined workflow that only exposes the areas that are relevant to their work.
- Expert users can create workflows that make it easy for non-expert users to accomplish complex tasks.

Assistant Workflow

The workflow of a Simulation Assistant is made up of a hierarchy of tasks. Each task is presented in its own panel in the assistant. The aim of the workflow is to help you set up a simulation in STAR-CCM+ without forcing you to perform tasks in a linear order.

Note: It is recommended that you follow a top-to-bottom approach when working with assistants. Jumping back and forth between tasks, especially if the tasks have no conditions, can lead to errors.

The Simulation Assistant User Interface

Every assistant has various user interface items.



Tasks

Tasks are hierarchical groups of steps that you perform in a simulation.

For example, a task could contain the steps for setting up a physics continuum. In the Simulation Assistant, each task is represented in a separate panel, called the Task Panel.

A task can contain any number of steps which are related to the task. Typically, the steps are list items which contain a combination of plain text and hyperlinks that are linked to an action. These actions are known as Star Actions, and can drive any operation in the STAR-CCM+ User Interface. Clicking on Star Actions typically causes a change in the STAR-CCM+ simulation. A Star Action can also execute code written by the author of the Simulation Assistant. The code can interact with STAR-CCM+ User Interface elements or the STAR-CCM+ API. The task is disabled while a Star Action executes. Star Actions appear as hyperlinks in steps.

Tasks can contain conditions, which determine when a task is ready to use, or when a task is complete. If tasks do not have conditions, you are best to follow a top-to-bottom approach through the workflow, without moving backwards and forwards between the tasks.

You can control which tasks and sub-tasks are visible in the assistant by expanding and collapsing them.



To expand a task:

- Click (or anywhere else) in the task panel that you want to expand. The selected task expands.

To collapse a task:

- Click (or anywhere else) in the task panel that you want to collapse. The selected task collapses.

Sub-Tasks

Sub-tasks are tasks that are nested under a single task. Typically, sub-tasks are groups of steps that are related to the parent task.

Steps and Star Actions

Steps and Star Actions appear in tasks and are the smallest unit in the Simulation Assistant.

A step is typically represented as a bullet point, or a numbered item, and is expected to perform a single action. A step can also contain a description of what the user must do.

See Also

Conditions

Collapsible Information

Collapsible information contains additional information that is hidden by default in the assistant.

Collapsible information is marked by the icon. Typically, information that you store here is not essential for the completion of the task, but is there to provide extra information when necessary.



To view the additional information:

- Click the icon.

  The additional information is displayed in the task.



Object Tree Focus

Simulation assistants can provide an object tree focus—that is, when you click a step, the relevant node in the STAR-CCM+ object tree expands and is highlighted.

This facility is useful for users that are not familiar with the STAR-CCM+ object tree. Object tree focus can also act as a hint to what the corresponding Star Action does, such as expanding the physics continuum models node when the task adds a model to the physics continuum.

There are two types of object tree focus:

- *Inclusive* — the relevant node in the object tree is expanded, leaving all other expanded nodes as they were.



- *Exclusive* — the relevant node in the object tree is expanded, and all other nodes are closed.


Servers
Star 1
simulation
scene/plot
Star 1
Geometry
3D-CAD Models
Parts
Fluid
Descriptions
Contacts
Tags
Filters
Operations
Continua
Servers
Star 1
simulation
scene/plot
Star 1
Geometry
Continua
Physics
Models
Constant Density
Gas
Gradients
Laminar
Segregated Flow
Steady
Three Dimensional
Reference Values
Initial Conditions
Regions

Conditions

Conditions add an extra layer of control and guidance to the assistant, and act as a safeguard when you move between tasks. In the Simulation Assistant, conditions are represented by tabs in each task panel.

There are two types of conditions:

- *Pre-conditions* — are conditions that must be satisfied before you can start the current task.
- *Post-conditions* — are conditions that notify you if the constraints for the current task have been satisfied.

Satisfying Conditions

Conditions control whether you can perform a particular task or not.

If a task has a pre-condition that has not been satisfied, the task is grayed out and you cannot perform the task.



By default, condition tabs are minimized.



To view more information about a particular condition:

- Click the condition tab in the task panel.

  The tab expands to reveal more information, such as what is required in order to satisfy the condition. In the following example, a geometry part must be present in the simulation in order to satisfy the condition.



When a condition is satisfied, a green checkmark appears next to the condition. Tasks that have satisfied their post-conditions also have a green checkmark next to the title in the task panel.



Note: Tasks that have only satisfied their pre-conditions do not have a green checkmark in their task panel.

Changing the Context of a Condition

The context of a condition is the object in STAR-CCM+ that satisfies the condition—for example, a geometry part. You can change which object satisfies a condition by changing the context of the condition.

In the following example, a condition is already satisfied by a geometry part, however the part that satisfied the condition will be changed.

To change the context of a condition:

1. On the STAR-CCM+ object tree, select the object that you want to use to satisfy the condition.



2. Expand the tab of the condition that you want to change and click ⋯.



The selected object is now used to satisfy the condition.



Note: You can only replace the context of a condition with an object of the same type. For example, you can replace a CAD part with a Mesh part, as both are geometry parts—but you cannot replace a CAD part with a 3D-CAD model. The object may also need to match other parameters that are typically specified in the condition description.

Loading a Simulation Assistant into STAR-CCM+

Simulation assistants are stored in precompiled `.java` or `.jar` files.

To load a Simulation Assistant into your simulation:

1. Launch STAR-CCM+ and open or create a simulation.
2. Select **File > Load Simulation assistant...**



The *Open* dialog appear.

3. Navigate to the location where the Simulation Assistant is stored and select it.
4. Click **OK**.

The assistant loads into the simulation. By default, the assistant appears on the right-hand side of the user interface.


Internal Flow Assistant
Import Geometry
Post-Condition
This task imports a geometry part into STAR-CCM+.
Import the surface mesh.
Create Region from Part
Create Physics
Generating the Mesh
Create Scalar Scene
Run the Simulation

Setting Up a Simulation Assistant Project in NetBeans

To develop STAR-CCM+ Simulation Assistants, it is recommended that you use an IDE (Integrated Development Environment).

IDEs offer many benefits compared to standard text editors, such as auto-indenting your code and offering auto-completion for method and class names (when you import the appropriate application libraries). This guide uses NetBeans 7.3.1 IDE. You can get a free copy of NetBeans SE from the official website.

This section contains the following topics:

Creating A Simulation Assistant Project In NetBeans

To set up a NetBeans project:

1. Download and install the latest version of NetBeans from the official website.
2. Launch NetBeans.
3. Select **File > New Project...**.

The *New Project* dialog appears.

4. Select **Java**from the *Categories* list.
5. Select **Java Application** from the *Projects* list.



6. Click **Next**.

The *New Java Application* dialog appears.

7. In the **Project Name** field, enter the project name.
8. Clear the *Create Main Class* option.



9. Click **Finish**.

The project is added to your projects list. Continue by adding the Java packages and Java classes.



Adding Java Packages, Java Classes, and XHTML Files to the Simulation Assistant Project

Typically, Simulation Assistant projects contain two Java packages, one for the assistant code, and the other for the XHTML text that is displayed in the user interface of the assistant.

This topic contains the following sections:

- Adding Java Packages to the Project
- Adding Java Classes to the Project
- Adding XHTML Files to the Project

Adding Java Packages to the Project

To add Java packages to the Simulation Assistant Project:

1. Right-click on **[project_name] > Source Packages** and select **New Java Package...**



   The *New Java Package* dialog appears.

2. In the **Package Name** field, enter the package name.



3. Click **Finish**.

   The Java package is added to the project.



Adding Java Classes to the Project

Java classes contain the code that made up the Simulation Assistant.

To add a Java class to the Simulation Assistant project:

1. Right-click **[project_name] > Source Packages > [package_name]** and select **New Java Class...**



The *New Java Class* dialog appears.

2. In the **Class Name** field, enter the class name.



3. Click **Finish**.

The Java class is added to the Java package. Add as many classes as required for your workflow.



Adding XHTML Files to the Project

Typically, all the text that is displayed in the Simulation Assistant user interface is stored in XHTML files, and is contained in a separate Java class.

To add an XHTML file to the simulation assistant project:

1. Right-click on **[project_name] > Source Packages > [package_name]** and select **New > Other...**



   The *New File* dialog appears.

2. Select **Other** from the *Categories* list.
3. Select **XHTML File** from the *File Types* list.



4. Click **Next**.

   The *New XHTML File* dialog appears.

5. In the **File Name** field, enter the file name.



6. Click **Finish**.

   The XHTML file is added to the Java package. Add as many XHTML files as required for your workflow.


Projects
Files
Services
Simulation Assistant
Source Packages
Assistant
XHTML
Task1Text.xhtml
Libraries

Adding STAR-CCM+ Libraries to the Simulation Assistant Project

You can import the STAR-CCM+ libraries into NetBeans so that you can use the auto-check and auto-complete features of the IDE.

Caution: If you are creating the project on a Windows platform, it is recommended that you copy and paste the STAR-CCM+ libraries to the root directory of your system, for example: `C:\STAR-CCM+ Libraries\`. When building code, Windows has a limitation of how large the path to the libraries can be. In the default STAR-CCM+ installation directory, the STAR-CCM+ libraries would contain paths that are too long. When building the project, these long paths would cause problems.

To import the STAR-CCM+ libraries:

1. Right-click on **[project_name] > Libraries** and select **Add Library...**.



The *Add Library* dialog appears.

2. Click **Create...**.



The *Create New Library* dialog appears.

3. In the *Library Name* field, enter the name of the library.
   It is recommended you name the STAR-CCM+ libraries according to the version number, for example `starccm+_804010`.



4. Click **OK**.

The *Customize Library* dialog appears.

5. Click **Add JAR/Folder...**



The *Browse JAR/Folder* dialog appears.

6. Navigate to your `[STAR-CCM+_install_dir]\star\lib\java` folder.

Note: If you are going to build projects on Windows platforms, it is recommended that you copy this folder to your root directory, for example: `C:\STAR-CCM+ Libraries\`.

7. Add all the `.jar` files that are located in the following subfolders:
   - `/platform/core`
   - `/platform/core/locale`
   - `/platform/lib`
   - `/platform/modules`
   - `/platform/modules/ext`
   - `/platform/modules/locale`

   The libraries that you add are displayed in the *Classpath* tab.



8. If you want to reference the Java documentation, add the following folder:
   - [STAR-CCM+_install_dir]\doc\client\html
9. Click **OK** to close the *Customize Library* dialog.
10. Select the library package that you just created and click **Add Library**.

    The *Add Library* dialog closes.
11. Click **OK** to close the *Project Properties* dialog.

The STAR-CCM+ libraries are added to the project.


Projects
Files
Services
Simulation Assistant
Source Packages
Libraries
STAR-CCM+_806003 - core_star.jar
STAR-CCM+_806003 - openide_star.jar
STAR-CCM+_806003 - core.jar
STAR-CCM+_806003 - org-openide-filesystems.jar
STAR-CCM+_806003 - boot.jar
STAR-CCM+_806003 - org-openide-modules.jar
STAR-CCM+_806003 - org-openide-util-lookup.jar
STAR-CCM+_806003 - org-openide-util.jar
STAR-CCM+_806003 - org-jdesktop-layout.jar
STAR-CCM+_806003 - org-netbeans-api-annotations-common.jar

Developing a Simulation Assistant

A Simulation Assistant is a user-defined panel within the STAR-CCM+ user interface. The panel is intended to provide a series of tasks that a user completes in sequence.

This section of the documentation provides documentation on the Java classes and coding that you use to create a Simulation Assistant. This material assumes that you are familiar with Java coding as well as STAR-CCM+. It also assumes that you are using an Integrated Development Environment (IDE) such as NetBeans 7.3.1 or later.

Overview

The Simulation Assistant uses a simple version of the model-view-controller (MVC) design approach to separate the underlying data (model), how the data is represented (view), and how actions are performed (controller).

At the highest level, the Simulation Assistant, you only create a model with simple data (a derivative of the Java class SimulationAssistant). The view and the controller are handled internally.

For an individual task, the model (the Task class), specifies the type of view (the Content class). The model also specifies the controller (a derivative of the TaskController class) to represent and control itself at runtime. The view and controller are instantiated when the Task is created and displayed. For a task, you almost always create a model (a derivative of Task) and a controller (typically a derivative of FunctionTaskController). The type of view is specified using one of the pre-existing Content classes containing or pointing to some XHTML.



As the API was developed to maximize coding flexibility, the model, view, and controller can be extended in a number of ways. This development section emphasizes the expected approach.

For a rudimentary Simulation Assistant, see the Hello World example at:

`[STAR-CCM+_Installation] > star > examples > assistant > HelloWorld`

See also the tutorial Simulation Assistant: Internal Flow Assistant.

Model

Task objects are the model objects for a task.

They contain the name, description, state (hidden, disabled, enabled, or complete), a list of subtasks (if any), and pointers to the view and controller to use with the Task (the `getContent()` and `getControllerClass()` functions).

Excluding subtasks, the initial values for all of these items is most easily specified using a `@StarAssistantTask` annotation on your Task subclass. When the data in the model changes (for example, the state transitions from disabled to enabled), the model fires a property change, which the view observes.

View

Content objects link a Task to a view of the Task.

The Simulation Assistant API uses subclasses of Content to determine the correct `TaskContentView` for rendering the Task. Typical Content-derivative choices include `HTMLFileContent`, which points to a complete XHTML file to be displayed, and `HTMLFragmentContent`, which contains a snippet of XHTML to display for the task.

In more advanced situations, you can use a Java UI component such as a JPanel to represent the Task. Write a custom Content-derivative and a custom TaskContentView-derivative, and register the new combination in `TaskContentLoader`.

Controller

The `TaskController` handles user-invoked actions.

The actions that the controller receives from the view are string commands in their purest form. Sometimes these have name/value pair arguments, but the most commonly used TaskController-derivative, `FunctionTaskController`, interprets these commands to be no-argument functions of the `FunctionTaskController` subclass that the author has written.

To use `FunctionTaskController`, your subclass must be public, and the function to execute must be public and take no arguments. For example, a command like "`myFunction`" is represented in the HTML as `<a href="staraction:myFunction">`.When clicked, it executes a function in your `FunctionTaskController` subclass with the signature `public void myFunction()`.

Files

Developing a Simulation Assistant uses the following file types:

- Java files, to contain the code for the assistant.
- JAR files, to wrap and export the Simulation Assistant project as a pre-compiled file.

  You can use **File > Load Simulation Assistant** to load assistants directly from either Java or JAR files.
- XHTML files, to contain text and links for the user interface. You can also write this text as inline XHTML, within the Java files of the assistant. To make sure that your XHTML code is well-formed, go to http://validator.w3.org/.
- Image and media files, to augment XHTML text. Use image files for mathematical notation. Supported formats include GIF, JPG, PNG, and SVG.

Tasks and Classes

The Simulation Assistant consists of a list of tasks to be done in order.

Example:

1. Import the geometry
2. Create the regions from the parts
3. Create and set up the physics of the simulation
4. Set up and generating the mesh
5. Visualize the solution
6. Run the simulation

In turn, these tasks often consist of subtasks. Example:

3a. Create physics continuum

3b. Set material properties

3c. Set initial conditions and boundaries

Set up the structure of classes in the Java code of the assistant to parallel the task structure, so that each high-level class performs one high-level class. Classes for subtasks are called by the top level tasks. See Supplied Classes for the classes that STAR-CCM+ provides.

Use the following rules:

- Create a Java file containing the class for your Simulation Assistant, extended from `SimulationAssistant`. See SimulationAssistant Class.
- Within your Simulation Assistant class, create a list of tasks for the top-level outline of the Simulation Assistant.
- Create a series of Java files, one for each task, extended from `Task`. See Task Class.
- Create XHTML code for each task, supplying the text and links for the user interface.
    - For extended text, use a separate XHTML file. Link to it with the `contentPath` parameter of the annotation for the task, or pass an HTMLFileContent object to the `setContent()` method of the task.
    - For small sections of XHTML code, specify them with the `contentPath` parameter of the annotation for the task, or pass

an HTMLFragmentContent object to the `setContent()` method of the task.

- Define the actions for the tasks in extensions of AbstractTaskController.
- Where necessary, define the pre-conditions and post-conditions for the tasks, using Condition, ConditionTrigger, and their extentions.

Star Actions

A Star Action is an action that is performed in the STAR-CCM+ interface.

The Star Action links a word or phrase in XHTML text to a function named in the task. The assistant performs the Star Action when the user clicks the text.

Star Action XHTML code has the format:

```
<a href="staraction:function_name">action_text</a>
```

Examples:

```
<a href="staraction:createPhysaicsContinuum">Create</a>
```

```
<a href="staraction:materialProperties">Define</a>
```

The `action_text` strings (`Create` and `Define` in the examples) appear as the link text presented by the Simulation Assistant:

1. Create a physics continuum
2. Define the material properties

The `function_name` strings (`createPhysicsContinuum` and `materialProperties` in the examples) are the names of Java functions that are used in the task.

Annotations

The Simulation Assistant provides two annotations for generic set methods:

- `@StarAssistant`, for use in extensions of the SimulationAssistant Class class, sets the values of `desc`, `display`, and `outline`.

  See: **Help > Java API > star.assistant.annotation > StarAssistant**

- `@StarAssistantTask`, for use in extensions of the Task Class class, sets the values of `contentFragment`, `contentPath`, `controller`, `display`, `id`, and `state`.

  See: **Help > Java API > star.assistant.annotation > StarAssistantTask**

Client Server Objects and Lookups

Client Server Objects (CSOs) are any of the items that are represented in the STAR-CCM+ simulation tree.

A model, an interface, a mesh, an injector, a property, a field function, are all examples of CSOs. It is a very broad category.

A lookup is a search tool that is assigned to a list of CSOs associated with a task. If you set a task lookup to search for a CSO and it fails to find the CSO in its own list, the lookup searches the lists of related tasks. It checks, in order, preceding siblings, parent, preceding siblings of the parent, grandparent, and so on, up to the lookup list for the whole Simulation Assistant.

Put CSOs on the lookup list when you expect to use them again in the task, or in a parent or child of the task. A CSO that satisfies a Condition is automatically added to the lookup. See Conditions, Triggers, and Predicates.

The lookup-related functions, provided in AbstractTaskController, are:

Function	Description
`addToTaskLookup(`*`cso`*`)`	Puts a CSO on the lookup list, where cso is the object in question.
`lookupObject(`*`class`*`)`	Searches for a CSO on the lookup list hierarchy, where class is the class of the object you are looking for. The lookup returns the first object of the class.
`lookupObjects(`*`class`*`)`	Searches for all CSOs of a class on the lookup list hierarchy, where class is the class of the object you are looking for. The lookup returns all objects in the class.
`getLookup()`	Returns the entire lookup object.
`printLookup()`	Prints the lookup contents to the output window.
`removeFromTaskLookup(`*`cso`*`)`	Removes a CSO from the lookup list, where cso is the object in question.
`clearTaskLookup()`	Clears the lookup list of all CSOs.

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant.ui > AbstractTaskController

Best Practices

Add a CSO to the lookup when you expect to use it with a different action, particularly when that action is in the same task or in a parent or child task. The lookup is also critical for many pre- and post-condition tests, so if a CSO being tested in a condition has been successfully set up, add it to the lookup.

A CSO that satisfies a Condition is automatically added to the lookup.

Tasks

The Simulation Assistant is composed of Tasks, each in a separate Java file.

Each Task controls a separate panel in the Simulation Assistant window. Task panels display text, including links that let users perform Star Actions.

Tasks share data by putting it in their lookups, where other tasks find it when they use the hierarchical search in `lookupObject()`.

A Task can have a list of Pre-conditions, a list of Post-conditions, both, or neither.

- If there is a list of Pre-conditions, all the conditions on the list must be satisfied before the Task is enabled.
- If there is a list of Post-conditions, all the conditions on the list must be satisfied before the Task is marked as complete.

See Also

Task Class
Conditions, Triggers, and Predicates

Branching and Looping

Tasks can branch and loop.

To use branching, make a Task with two or more subtasks. These subtasks remain hidden until a choice is made in the parent Task, usually by clicking a link. The choice activates one of the branches, making the chosen subtask visible and leaving all of the other subtasks hidden.

See `\star\examples\assistant\SimpleBranch`

To use looping, use LoopTask, LoopingCondition, and LoopIterationTrigger.

See `\star\examples\assistant\SimpleLoop`

Conditions, Triggers, and Predicates

Use Conditions to control what tasks are enabled or disabled in the Simulation Assistant, and to make sure that the user creates a valid simulation.

Conditions are classes that track objects meeting certain criteria. When a Condition finds such an object, it *holds* the object (records a reference to it) and sets its *Satisfied* property to `true`. If the object is deleted, or if its properties or the properties of its children change to no longer meet the criteria, the Condition stops holding and sets its *Satisfied* property to `false`.

A Condition hunts for objects using a set of Triggers. A Trigger monitors a set of objects, such as those in the lookup hierarchy. When it finds the type of object that is specified for it, it passes the object to the Condition for evaluation.

A Condition makes the evaluation using a Predicate. If the object meets the criteria for the Predicate, the Predicate returns `true` and the Condition sets its *Satisfied* property to `true`.

A Condition has one Predicate, and the Predicate evaluates one object at a time. However, the object can contain other objects or point to other objects (for example, if the passed object is a manager node). The Predicate can evaluate the object with combinations of logical operations using the CompoundPredicate extension.

See:

- Condition
  - subclass CSOCondition
- ConditionTrigger (abstract class)
  - subclass CSOConditionTrigger
  - subclass LookupConditionTrigger
    - sub-subclass CSOLookupConditionTrigger
  - subclass LoopIterationTrigger
  - subclass TaskConditionTrigger
- Predicate
  - subclass CompoundPredicate

XHTML

The text appearing in a task panel is coded in XHTML.

You can write the text in either of two ways: by using XHTML formatted text strings (XHTML fragments), or by referencing XHTML files in the Java files.

Use formatted text strings for brief messages and referenced XHTML files for extended text.

This example uses the XHTML fragment format:

```
@StarAssistantTask(display="Hello Task",
contentFragment="<ul><li>
<a href=\"staraction:hello\">Hello</a>!</li></ul>",
controller=HelloWorld.HelloWorldTask.
HelloWorldTaskController.class)
```

This example uses an XHTML file:

```
@StarAssistantTask(display="Create A Continuum With
Pre-defined Models"
contentPath="html/CreateContinuaTask.xhtml",
controller=CreateContinuaTask.
CreateContinuaTaskController.class)
```

XHTML text can be used only once in a task panel. However, a task panel can display any one of several alternate passages, depending on conditions. For situations like this one, you can use the HTMLFileContent and HTMLFragmentContent classes to specify the alternate text.

Supplied Classes

This section describes the classes that are supplied by STAR-CCM+.

In most cases, most or all of the classes in your Simulation Assistant are extensions of these classes. The supplied classes are:

High-level classes:

- SimulationAssistant Class
- Task Class
- LoopTask

Task controller classes:

- AbstractTaskController
- TaskController
- FunctionTaskController

Content classes:

- Content
- HTMLFileContent
- HTMLFragmentContent

Condition classes:

- Condition
- CSOCondition
- LoopingCondition

Condition trigger classes:

- ConditionTrigger
- CSOConditionTrigger
- LookupConditionTrigger
- CSOLookupConditionTrigger
- LoopIterationTrigger
- TaskConditionTrigger

Predicate classes:

- Predicate
- CompoundPredicate

Class Relationships

The following diagram shows inheritance relationships between the provided classes:



The following diagram shows usage relationships between the provided classes:



SimulationAssistant Class

Extend this class to create the top-level class for your Simulation Assistant.

The top-level class contains the title of the Simulation Assistant and a list of the constituent tasks, which are given in order of execution. See Task Class.

A SimulationAssistant must contain the following items:

- Imports:
  - star.assistant.SimulationAssistant
  - star.assistant.Task
- A constructor with the name of the assistant.
- A list of tasks

A SimulationAssistant often contains these items:

- The annotation `@StarAssistant`
- Imports:
  - java.util.ArrayList
  - java.util.List
  - star.assistant.annotation.StarAssistant
  - star.assistant.annotation.StarAssistantTask
  - star.assistant.ui.FunctionTaskController

Example:

```
import java.util.ArrayList;
 import java.util.List;
 import star.assistant.SimulationAssistant;
 import star.assistant.Task;
 import star.assistant.annotation.StarAssistant;

@StarAssistant(display = "Internal Flow Assistant")
 public class InternalFlowAssistant extends
 SimulationAssistant {

public InternalFlowAssistant() {
...
    }
}
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > SimulationAssistant

Task Class

Use extensions of this class to create the classes for individual tasks in your Simulation Assistant, one class for each task pane in the assistant. The Task displays text and allows the user to perform specified Star Actions.

The Task can have a list of pre-conditions that must be satisfied before the user can perform the task. The `setPreconditions()` function receives this list.

Independently, the Task can have a list of post-conditions that must be satisfied before the user can complete the function and continue. The `setPostconditions()` function receives this list.

A Task must use import star.assistant.Task

A Task often uses these items:

- @StarAssistantTask
- import star.assistant.annotation.StarAssistantTask
- import star.assistant.ui.FunctionTaskController
- setPreconditions()
- setPostconditions()

Example:

```
package Assistant;

import star.assistant.Task;
import star.assistant.annotation.StarAssistantTask;
import star.assistant.ui.FunctionTaskController;
...

@StarAssistantTask(display = "Run the Simulation",
  contentPath = "HTML/06_RunSimulation.xhtml",
    controller = Task06RunTheSimulation
      .RunTaskController.class)

public class Task06RunTheSimulation extends Task {

  public class RunTaskController extends
    FunctionTaskController {
```

```
...

    }

}
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > Task

LoopTask

LoopTask is an extension of the Task Class.

Use extensions of this class to create a basic task that keeps a counter on a defined action and notifies the user when a predefined condition has been met. Use `LoopingCondition` in `LoopTasks`, to control the number of repetitions. You can also use any other Conditions.

Example:

```
import star.assistant.LoopIterationTrigger;
 import star.assistant.LoopTask;
 import star.assistant.LoopingCondition;
 import star.assistant.annotation.StarAssistantTask;
 import star.assistant.ui.FunctionTaskController;

...

@StarAssistantTask(display="Create Boundaries for a
Region",
                   contentPath="html/
CreateBoundariesTask.xhtml",

controller=CreateBoundariesTask.CreateBoundariesTaskCont
roller.class)
public class CreateBoundariesTask extends
LoopTask<Boundary> {


    public CreateBoundariesTask() {

...

        //loop condition
         LoopingCondition<Boundary> loopCondition = new
LoopingCondition<Boundary>();
         loopCondition.setDesc("At least 3
boundaries.");
```

```
         LoopIterationTrigger<Boundary> loopTrigger =
new LoopIterationTrigger<Boundary>(3);
         loopCondition.setLoopTrigger(loopTrigger);
         loopCondition.setPredicate(new
Predicate<List<Boundary>>() {


            @Override
             public boolean evaluate(List<Boundary>
boundaries) {

...

                return false;
                 }

...


                }

                return true;
             }
         });

        setLoopCondition(loopCondition);

...
    }
...
}
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > LoopTask

Task Controller Classes

These classes handle user input.

AbstractTaskController

FunctionTaskController inherits functions from this class. `AbstractTaskController` is an abstract class implementing TaskController.

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant.ui > AbstractTaskController

TaskController

FunctionTaskController inherits functions from this class. `TaskController` is an interface that is implemented by AbstractTaskController.

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > TaskController

FunctionTaskController

`FunctionTaskController` is an extension of AbstractTaskController. Extend it to create a controller that defines functions that are called when the user clicks a hyperlink. The functions execute code sequentially (including macro code); that is, operations are synchronous.

A FunctionTaskController must use the following items:

- import star.assistant.Task
- import star.assistant.ui.FunctionTaskController
- A function name matching a Star Action name in the associated XHTML file.

A FunctionTaskController often uses these items:

- @StarAssistantTask
- import star.assistant.annotation.StarAssistantTask

Example:

```
public class CreateBoundaryTaskController extends
FunctionTaskController {

public void createBoundary() {
Region region = lookupObject(Region.class);
Boundary boundary =
region.getBoundaryManager().createEmptyBoundary();
boundary.setBoundaryType(InletBoundary.class);
addToTaskLookup(boundary);
}
}
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant.ui > FunctionTaskController

Content Classes

These classes control the XHTML text that the assistant displays.

Content

HTMLFileContent and HTMLFragmentContent inherit functions from this abstract class.

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > Content

HTMLFileContent

Use this class to reference an XHTML file when you cannot use an annotation.

For instance, a task sometimes must display different text depending on varying conditions. In such a case, you can use this class to reference the alternate text. If the amount of text is small, use HTMLFragmentContent.

Example:

```
Content content = new HTMLFileContent();
content.setValue("CCWorkflow1_MeshSetup.xml");
setContent(content);
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > HTMLFileContent

HTMLFragmentContent

Use this class to display text in a task panel when you cannot use the text string in an annotation.

If the amount of text is large, use HTMLFileContent.

Example:

```
Content c = new HTMLFragmentContent();
c.setValue("<ul><li>
<a href=\"staraction:hello\">Hello</a>!</li></ul>");
this.setContent(c);
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > HTMLFragmentContent

Condition Classes

These classes control what tasks are enabled or disabled in the assistant.

Condition

Use extensions of this class to monitor a set of objects for potentially relevant objects.

A Condition uses one or more ConditionTrigger extentions and a Predicate extension. When the type of object that is assigned to a `ConditionTrigger` appears, the Condition evaluates the object using the Predicate. If the Predicate returns the value `true`, the `ConditionTrigger` stops monitoring other objects, the Condition holds the object, and sets its *Satisfied* property to `true`.

If the Predicate returns the value `false`, the Condition sets its *Satisfied* property to `false` and continues monitoring the set of objects.

This class has CSOCondition as a subclass.

Example:

```
Condition<String> cnd = new Condition<String>();
cnd.setDesc("A string containing \"Cat\"");
LookupConditionTrigger<String> cndTrigger =
new LookupConditionTrigger<String>(String.class);
cnd.setTriggers(Collections.singleton(cndTrigger));
cnd.setPredicate(new Predicate<String>() {
@Override
public boolean evaluate(String string) {
if (string.contains("Cat")) {
return true;
}
```

```
return false;
}
});
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > Condition

See Also

Conditions, Triggers, and Predicates

CSOCondition

`CSOCondition` is an extension of Condition. Use extensions of this class to monitor conditions that are associated with CSOs. `CSOCondition` watches for the addition and deletion of the CSO or its children, in addition to property changes.

Example:

```
CSOCondition<PhysicsContinuum> cnd = new
CSOCondition<PhysicsContinuum>();
cnd.setDesc("The physics continuum must have the
Radiation model enabled.");
CSOLookupConditionTrigger<PhysicsContinuum> cndTrigger =
new
CSOLookupConditionTrigger<PhysicsContinuum>(PhysicsConti
nuum.class);
cnd.setTriggers(Collections.singleton(cndTrigger));
cnd.setPredicate(new Predicate<PhysicsContinuum>() {
@Override
public boolean evaluate(PhysicsContinuum continuum) {

if (continuum.getModelManager().hasModel(
RadiationModel.class) != null) {
return true;
}

return false;
}
});
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > CSOCondition

See Also

Conditions, Triggers, and Predicates

LoopingCondition

`LoopingCondition` is an extension of Condition. Use extensions of this class in a LoopTask to monitor the state of a LoopIterationTrigger.

Example:

```
 //loop condition
LoopingCondition<Boundary> loopCondition = new
LoopingCondition<Boundary>();
loopCondition.setDesc("At least 3 boundaries.");
LoopIterationTrigger<Boundary> loopTrigger = new
LoopIterationTrigger<Boundary>(3);
loopCondition.setLoopTrigger(loopTrigger);
loopCondition.setPredicate(new
Predicate<List<Boundary>>() {

    @Override
    public boolean evaluate(List<Boundary> boundaries) {
...
        }

        return true;
    }
});
setLoopCondition(loopCondition);
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > LoopingCondition

See Also

Conditions, Triggers, and Predicates

Condition Trigger Classes

These classes work within Conditions and monitor objects.

ConditionTrigger

Use extensions of this class in Conditions to monitor system events and detect when potentially relevant objects become available. When the type of object for the `ConditionTrigger` appears, the trigger passes the object to the Condition for evaluation. If the Condition is satisfied and holds the object, it deactivates the trigger.

`ConditionTrigger` is an abstract class. It has CSOConditionTrigger, CSOLookupConditionTrigger, and TaskConditionTrigger as subclasses.

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > ConditionTrigger

See Also

Conditions, Triggers, and Predicates

CSOConditionTrigger

`CSOConditionTrigger` is an extension of ConditionTrigger. Use extensions of this class in Conditions to monitor conditions that are associated with a particular CSO and its children.

Example:

```
public CSOCondition<PhysicsContinuum>
createCondition(PhysicsContinuum pc) {
CSOCondition<PhysicsContinuum> cnd = new
CSOCondition<PhysicsContinuum>();
cnd.setDesc("The physics continuum must have the
Radiation model enabled.");
CSOConditionTrigger<PhysicsContinuum> cndTrigger =
new CSOConditionTrigger<PhysicsContinuum>(pc);
cnd.setTriggers(Collections.singleton(cndTrigger));
cnd.setPredicate(new Predicate<PhysicsContinuum>() {
@Override
public boolean evaluate(PhysicsContinuum continuum) {
if (continuum.getModelManager().hasModel(
RadiationModel.class) != null) {
return true;
}
return false;
}
});
return cnd;
}
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > CSOConditionTrigger

See Also

Conditions, Triggers, and Predicates

LookupConditionTrigger

`LookupConditionTrigger` is an extension of ConditionTrigger. Use extensions of this class in Conditions to monitor the task Lookup for the addition and removal of objects of a particular class.

Example:

```
CSOCondition<PhysicsContinuum> cnd = new
CSOCondition<PhysicsContinuum>();
cnd.setDesc("The physics continuum must have the
Radiation model enabled.");
CSOLookupConditionTrigger<PhysicsContinuum> cndTrigger =
new
CSOLookupConditionTrigger<PhysicsContinuum>(PhysicsConti
nuum.class);
cnd.setTriggers(Collections.singleton(cndTrigger));
cnd.setPredicate(new Predicate<PhysicsContinuum>() {
@Override
public boolean evaluate(PhysicsContinuum continuum) {
if (continuum.getModelManager().hasModel(
RadiationModel.class) != null) {
return true;
}
return false;
}
});
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > LookupConditionTrigger

See Also

Conditions, Triggers, and Predicates

CSOLookupConditionTrigger

`CSOLookupConditionTrigger` is an extension of LookupConditionTrigger. Use extensions of this class in Conditions to monitor CSOs in the task Lookup. The trigger watches for the addition and deletion of the CSO or its children, in addition to property changes. The trigger passes the object to the Condition only when these changes occur.

Example:

```
import star.assistant.CSOLookupConditionTrigger;
...
public class ModifyPolynomialDensityTask extends Task {
    public ModifyPolynomialDensityTask() {

    //add a precondition
    //precondition: unavailable unless Polynomial in T
exists.
    CSOCondition<Continuum> continuumCondition = new
CSOCondition<Continuum>();
    continuumCondition.setDesc("A liquid physics
continuum with a \"Polynomial in T\" method assigned
for the Polynomial Density property.");
    CSOLookupConditionTrigger<Continuum>
continuumConditionTrigger = new
CSOLookupConditionTrigger<Continuum>(Continuum.class);
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > CSOLookupConditionTrigger

See Also

Conditions, Triggers, and Predicates

LoopIterationTrigger

`LoopIterationTrigger` is an extension of ConditionTrigger. Use extensions of this class in LoopingConditions to monitor a list of objects whenever an object is added to the list, to determine if a LoopTask repeats.

Example:

```
 //loop condition
LoopingCondition<Boundary> loopCondition = new
```

```
LoopingCondition<Boundary>();
loopCondition.setDesc("At least 3 boundaries.");
LoopIterationTrigger<Boundary> loopTrigger = new
LoopIterationTrigger<Boundary>(3);
loopCondition.setLoopTrigger(loopTrigger);
loopCondition.setPredicate(new
Predicate<List<Boundary>>() {
    @Override
    public boolean evaluate(List<Boundary> boundaries) {
...
        }

        return true;
    }
});
setLoopCondition(loopCondition);
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > LoopIterationTrigger

TaskConditionTrigger

Use extensions of this class in Conditions to monitor a task for property changes. This way, tasks can monitor each other.

Example:

```
public Condition<Task> createCondition(Task task) {
Condition<Task> cnd = new Condition<Task>();
cnd.setDesc("The task must be complete.");
TaskConditionTrigger cndTrigger = new
TaskConditionTrigger(task);
cnd.setTriggers(Collections.singleton(cndTrigger));
cnd.setPredicate(new Predicate<Task>() {
@Override
public boolean evaluate(Task task) {
if (task.getState() == Task.TaskState.COMPLETE) {
return true;
}
return false;
}
});
return cnd;
}
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.assistant > TaskConditionTrigger

See Also

Conditions, Triggers, and Predicates

Predicate Classes

These classes work within Conditions and evaluate objects.

Predicate

Use extensions of this class in Conditions to evaluate objects. The Predicate evaluation function returns `true` or `false`, and all classes that implement Predicate must define this function. Each Condition has one Predicate, which evaluates a single object, but the object itself can contain other objects or point to other objects, and the Predicate can base its evaluation on these objects. The Predicate does not carry over any information from one evaluation to the next. Do not use it to alter any objects or properties in its evaluation code.

Example:

```
physicsContinuumCondition.setPredicate(new
Predicate<PhysicsContinuum>() {

@Override
public boolean evaluate(PhysicsContinuum continuum) {
if (continuum != null &&
continuum.getReferenceValues()
.has(ReferencePressure.class) != null) {
return true;
}
return false;
}
});
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.common.filters > Predicate and its subclasses.

See Also

Conditions, Triggers, and Predicates

CompoundPredicate

`CompoundPredicate` is an extension of Predicate. Use extensions of this class to make evaluations from logical combinations of sub-Predicates. Extensions for `And`, `Or`, and `Not` relations are provided.

Example:

```
CSOCondition<Boundary> cnd = new
CSOCondition<Boundary>();
cnd.setDesc("An inlet or outlet boundary.");
CSOLookupConditionTrigger<Boundary> cndTrigger =
new CSOLookupConditionTrigger<Boundary>(Boundary.class);
cnd.setTriggers(Collections.singleton(cndTrigger));
Predicate<Boundary> inletPredicate = new
Predicate<Boundary>() {
@Override
public boolean evaluate(Boundary boundary) {
if (boundary.getBoundaryType() instanceof
InletBoundary) {
return true;
}
return false;
}
};
Predicate<Boundary> outletPredicate = new
Predicate<Boundary>() {
@Override
public boolean evaluate(Boundary boundary) {
if (boundary.getBoundaryType() instanceof
OutletBoundary) {
return true;
}
return false;
}
};
List<Predicate<Boundary>> predicates = new
ArrayList<Predicate<Boundary>>();
predicates.add(inletPredicate);
predicates.add(outletPredicate);
cnd.setPredicate(new OrPredicate<Boundary>(predicates));
```

Online API:

For coding specifics, see the STAR-CCM+ **Help** menu:

Help > Java API > star.common.filters > CompoundPredicate

Help > Java API > star.common.filters > AndPredicate

Help > Java API > star.common.filters > NonePredicate

Help > Java API > star.common.filters > OrPredicate

See Also

Conditions, Triggers, and Predicates

Using the Simulation Assistant Wizard in NetBeans

The Simulation Assistant Wizard is a NetBeans plugin that automatically generates a Simulation Assistant project in the NetBeans IDE.

This plugin is installed into the NetBeans IDE. It cannot be installed directly within STAR-CCM+. It is not necessary to use the Simulation Assistant Wizard in order to create a Simulation Assistant, however, the Wizard is provided as a helpful tool that automates some of the process common to creating a new Simulation Assistant.

This section describes how to install and run the Simulation Assistant wizard. The wizard creates a skeleton Simulation Assistant project with a single Task and a single Star Action on that task. You can then customize this skeleton in order to create a useful Simulation Assistant. Running the wizard again will create additional Tasks and Star Actions on a project.

Installing The Wizard

1. Select **Help > Steve Portal** and login.
2. On the Steve Portal, go to the *Downloads* tab, locate **NetBeansWizard.nbm**, and download it.
3. Open NetBeans.
4. From the NetBeans **Tools** menu, select **Plugins**.
5. In the *Plugins* dialog, click the *Downloaded* tab.
6. Click **Add Plugins...** , and locate `NetBeansWizard.nbm`.
7. Select `NetBeansWizard.nbm` and click **Open**.
8. Select **Simulation Assistant Helper Module** and click **Install**.
9. Work through the installation dialog, taking the defaults.

Running the Wizard

1. From the Netbeans **File** menu, select **New Project**.
2. In the *New Project* dialog, select STAR-CCM+ from **Categories** and **Simulation Assistant Project** from **Projects**.
3. Click **Next** to activate the wizard.
4. Enter the information in the fields:

a) Enter a title and description, then click **Next**.

b) Enter a project name and, if necessary, browse to a new location, then click **Next**.

c) Enter a package name in lowercase, for example, `star`.

d) Enter a class name or accept the default, then click **Next**.

e) If the STAR-CCM+ library is not on the *Available Libraries* list, create it:

1. Click **Managing Libraries...**

   The *Ant Library Manager* dialog opens.

2. Click **New Library...**

   The New Library dialog opens.

3. In *Library Name*, enter STAR-CCM+, then click **OK**.

4. Click **OK** to return to the *Include Libraries* dialog.

f) Select STAR-CCM+ from the *Available Libraries* list and click **Managing Libraries...**.

   The *Ant Library Manager* dialog opens.

g) Click **New Library...**.

   The *New Library* dialog opens.

h) Enter STAR-CCM+ as the library name and click **OK**.

i) Click **Add JAR/Folder...** and enter the JAR files in the following libraries in `[STAR-CCM+_Installation]/star/lib/java/platform`:

- `.../core/*.jar`
- `.../core/locale/*.jar`
- `.../lib/*.jar`
- `.../modules/*.jar`
- `.../modules/ext/*.jar`
- `.../modules/locale/*.jar`

j) When done, click **OK**.

5. Under *Available Libraries*, select the STAR-CCM+ library you created and click **Next**.
6. Click **Finish**.

Tutorial Guide Contents

Tutorial Guide 1

Using Tutorial Macros and Files 3

Downloading the Tutorial Files from User Services 4
Understanding the Directory Structure and Naming Conventions 5
Playing the Macro 7
Using the Final Simulation Files 9

Introduction 11

Starting a STAR-CCM+ Simulation 13
Saving and Naming a Simulation 19
Importing the Geometry 20
Visualizing the Imported Geometry 23
Defining Boundary Surfaces 25
Assigning Parts to Regions 32
Setting Boundary Types 35
Generating the Mesh 41
Selecting the Physics Models 49
Setting Initial Conditions 55
Defining the Region Continuum 56
Setting Boundary Conditions and Values 57
Setting Solver Parameters and Stopping Criteria 61
Visualizing the Solution 62
Monitoring Simulation Progress 67
Running the Simulation 73
Adjusting Solver Parameters and Continuing 75
Visualizing the Results 80
Plotting Data on a Slice 89
Adding Streamlines 97

Closing and Reopening the Simulation....104
Summary....105

Geometry....106
3D-CAD: Cyclone Separator....107
Geometry Parts: CAD Assembly....160

Mesh....210
Surface Wrapper: Intake Manifold....211
Surface Repair Tools Brake Assembly....254
Parts-Based Meshing External Aerodynamics....285
Region Based Meshing Imported Control Valve....304
Region-Based Meshing Multi-Region Heat Exchanger....324
Directed Meshing: Electrical Machine....338

Incompressible Flow....356
Steady Flow: Lid-Driven Cavity Flow....357
Steady Flow: Channel Flow with Multiple Meshes....373
Steady Flow: Laminar and Turbulent in an S-Bend....405
Steady Flow: Backward Facing Step....440
Adjoint Flow Solver: External Flow over a Dual Element Wing....470
Adjoint Flow: Wing Shape Optimization....487
Steady Multi-Component Flow: Dilution Pipe....504
Unsteady Multi-Component Flow: Dilution Pipe....528
Porous Resistance: Isotropic Media....542
Porous Resistance: Orthotropic Media....574
Solution Recording and Playback: Vortex Shedding....585

Compressible Flow....609
Subsonic Flow: NACA-Type Intake....610
Transonic Flow: RAE2822 Airfoil....648
Transonic Flow: RAE2822 Airfoil Using Overset Mesh....669
Adjoint Flow Solver: Flow through an S-Bend....701
Turbo Wizard: Single Row....717
Turbo Wizard: Multiple Rows....739
Harmonic Balance: Single Stage Periodic Flow....778

Heat Transfer and Radiation....806
Conjugate Heat Transfer: Heated Fin Introduction....807
Multi-Part Solid: Graphics Card Cooling Introduction....828
Natural Convection: Concentric Cylinders Introduction....863

Dual Stream Heat Exchanger: Car Radiator 875
Surface-to-Surface Radiation: Thermal Insulator 897
Multiband Surface-to-Surface Radiation: Solar Collector 921
Thermal Comfort Wizard: Single Occupant in a Cabin 938

Multiphase Flow 968
VOF: Gravity-Driven Flow 969
VOF: Capillary Effects 989
VOF: Cavitation 1014
VOF: Boiling 1037
VOF: Melting-Solidification 1065
Lagrangian: Particle-Laden Flow 1096
Lagrangian: Solid Particle Erosion 1130
Eulerian: Dispersed Phase Acceleration 1177
Eulerian: Hibiki's Bubble Column 1199
Eulerian: Mixture Settling 1238
Eulerian: Bubble Formation in a Fluidized Bed 1268
Eulerian: Particle-Laden Flow in a Rotating Body 1291
Eulerian: Degassing Boundary 1324
Eulerian: Wall Boiling 1347
Eulerian: Conjugate Heat Transfer Wall Boiling 1398
Fluid Film: Liquid Film Flow 1436
Fluid Film: Binary Liquid Film Flow with Evaporation and Edge Stripping 1460
Dispersed Multiphase: Airfoil Icing 1500

Discrete Element Method 1520
DEM Particles in a Conveyor 1521
DEM Particle Settling 1556

Motion 1593
Moving Reference Frames: Rotating Fan 1594
Rigid Body Motion: Rotating Fan 1609
DFBI: Boat in Head Waves 1620
DFBI: Lifeboat with Overset Mesh 1650
Marine Resistance Prediction: KCS Hull with a Rudder 1684
Moving Reference Frames: Marine Propeller in Open Water 1726
Body Force Propeller Method: Marine Self-Propulsion 1748
Morphing: Cylinder with Boundary Motion 1772

Combustion 1800

3-Step Eddy Breakup: Non-Premixed Propane....1801
PPDF Equilibrium: Non-Premixed Propane....1823
PPDF Flamelets: Non-Premixed Hydrogen....1836
Complex Chemistry: Premixed Hydrogen....1846
Surface Chemistry: Methane Deposition....1859
Complex Chemistry: Reduced CPU Time with ISAT....1881
Droplet Combustion: Non-Adiabatic Toluene....1902
Coherent Flame Model: Premixed Propane....1924
Adiabatic PCFM: Partially-Premixed Methane....1941
Progress Variable Model: Sandia Flame D with Participating Media Radiation....1965
Fire and Smoke Wizard: Steckler Room....1988

Solid Stress....2015
Solid Stress: Exhaust Manifold....2016
Fluid Structure Interaction: Vibrating Pipe....2044

Aeroacoustics....2074
Broadband Models: Noise from a Cylinder (Preparation)....2075
DES and FW-H On-The-Fly: Noise from a Cylinder (Unsteady Analysis)....2097
Ffowcs Williams-Hawkings: Sound Propagation....2125
Signal Post-Processing: FFT and Wavenumber....2142
Acoustic Wave Modeling: Noise from a Cylinder....2158

Electromagnetism....2187
Ohmic Heating: Domestic Fuse....2188

Battery....2232
Battery Simulation Module: Cell Thermal Analysis....2233
Cylindrical Cells: Cell Thermal Analysis....2287
Li-Ion Battery Cell Model: Cell Electrochemistry Analysis....2330

Casting....2370
Basic Gravity Sand Casting....2371
STAR-Cast: Investment Casting....2411

Automation....2447
Simple Java Macros: Post-Processing Objects....2448
Intermediate Java Macros: Run Multiple Simulations....2478
Simulation Assistant: Internal Flow Assistant....2529

Coupling with CAE Codes....2574
Abaqus File-Based Coupling: Exhaust Manifold....2575

Abaqus Co-Simulation: Thermal Coupling....2616
Abaqus Co-Simulation: Mechanical Coupling....2658
STAR-CCM+ to STAR-CCM+ Coupling: Heat Transfer in a Chimney....2688
GT-POWER Co-Simulation: 1D Coupling....2720
OLGA Co-Simulation: 1D Coupling....2775
AMESim Co-Simulation: 1D Coupling....2804
Co-Simulation API: Spindle Valve....2820

Tutorial Guide

Tutorials show you how to use STAR-CCM+ for various applications in a step-by-step format with recommendations for setup, initialization and steps of the solution process specific to the application. In addition to these written instructions, macro and simulation files are available for download for a large proportion of cases.

The tutorials are categorized as follows:

- Introduction to STAR-CCM+ presents the important concepts and the general workflow. If you are new to STAR-CCM+, complete this tutorial first.
- Geometry tutorials demonstrate how to create and modify parts using 3D-CAD.
- Mesh tutorials demonstrate how to use the surface wrapper and several of the meshing features.
- Incompressible flow tutorials demonstrate simple liquid flows. Features such as steady and unsteady flows are introduced as well as multi-component flows, porosity and solution recording.
- Compressible fluid flow tutorials demonstrate the solution of subsonic and transonic flows and turbomachinery problems.
- Heat transfer and radiation tutorials demonstrate conjugate heat transfer, radiation and thermal comfort features.
- Multiphase tutorials demonstrate VOF, Lagrangian and Eulerian multiphase and fluid film features.
- Discrete Element Method tutorials demonstrate how to model large numbers of interacting discrete particles.
- Motion tutorials demonstrate the moving reference frames, rigid body motion, dynamic fluid body interaction and mesh morphing.
- Combustion tutorials demonstrate the various combustion models and the fire and smoke wizard.
- Solid stress tutorials demonstrate how to solve for stress and strain and how to couple these interactions in a fluid structure interaction.
- Aeroacoustics tutorials demonstrate the stages in solving a complex transient acoustics problem in the near and far field.
- Electromagnetism tutorials demonstrate analyses that involve an electric field.

- Battery tutorials demonstrate the process of setting up a battery simulation.
- Casting tutorials demonstrate the process of setting up a casting simulation.
- Automation tutorials demonstrate how to use Java macros for repetitive tasks.
- Coupling with CAE codes tutorials demonstrate how to couple STAR-CCM+ simulations.

Additionally, the CD-adapco Knowledge Base contains several video presentations that describe various features and procedures in STAR-CCM+.

Contents

Using Tutorial Macros and Files
Introduction
Geometry
Mesh
Incompressible Flow
Compressible Flow
Heat Transfer and Radiation
Multiphase Flow
Discrete Element Method
Motion
Combustion
Solid Stress
Aeroacoustics
Electromagnetism
Battery
Casting
Automation
Coupling with CAE Codes

Using Tutorial Macros and Files

Macros, input files and final simulation files for a range of tutorials are provided as an optional download package on the customer service center web-site. These macros and final simulation files are provided as an aid to the written tutorials, so that you can check your final results against the downloaded files, or against a simulation that is built and run using the macros.

Where a tutorial has several solution stages, a macro has been recorded for each solution stage. A simulation file for each completed macro is provided. You can load the solution for a preliminary stage and play the macro of the subsequent sections. The macros and sim files use naming conventions to facilitate their use.

Contents

Downloading the Tutorial Files from User Services
Understanding the Directory Structure and Naming Conventions
Playing the Macro
Using the Final Simulation Files

Downloading the Tutorial Files from User Services

To download a compressed package of the tutorial files and macros from the CD-adapco customer service site, http://steve.cd-adapco.com

1. Use your web browser to navigate to the CD-adapco customer service site and log in as normal. Contact your sales representative if you do not have a login.
2. Click the **File Downloads** option on the top bar of the website home page. This presents a tabbed window containing all products available for download.
3. Choose the STAR-CCM+ item in the menu on the left side.
4. Select the version of STAR-CCM+ for which you require tutorial files.
5. The tutorial files are presented in the **Related Files and Documentation** section at the bottom of the page. Click the download link to the right of the `STAR-CCM +TutorialFiles<Version>.7z` entry. Save the file to an appropriate directory on your disk.
6. Using a suitable decompression tool, unzip the `.7z` file. On Linux, to preserve the directory structure, use the following command line:

   ```
   7z x STAR-CCM+TutorialFiles<Version>.7z
   ```

Understanding the Directory Structure and Naming Conventions

The tutorial files are grouped in directories that are named according to the tutorial sets listed in the documentation.

Directory names do not have spaces and use a camelCase convention. For example, the files for the Subsonic Flow: NACA-Type Intake tutorial are located in the `compressibleFlow` folder.



The macro files within a directory are constructed so that each macro produces only a single solution. Macro names are derived from the tutorial name. For example, the Simple Java Macro is recreated using the macro `SimpleJavaMacroPostProcessingObjects.java`

Where more than one solver run is required within a tutorial, macros are created for each stage. Macros that refer to stages beyond the first stage include the documentation step from which they begin in their file name. For example, the introductory tutorial requires two solver runs. Its macros are:

- `Introduction1.java`
- `Introduction2AdjustingSolverParamters.java`

To run the second part of a tutorial, it is necessary to have the solution to the first part located in the same working directory. The download from User Services contains simulation files for each individual macro — that is, for intermediate stages as well as the final stage. These simulation files are named according to the macro file used to create them. Thus, `Introduction1.java` has a related `.sim` file called `Introduction1_final.sim`. This file can be loaded before running the macro for the second stage, `Introduction2AdjustingSolverParameters.java`.

Playing the Macro

STAR-CCM+ uses the tutorial macros to construct the same model that the tutorial instructions describe.

Before you play a macro, it is important to read through the appropriate section of the tutorial guide. There are three ways a tutorial macro can be used interactively:

- Playing an initial or single macro
- Playing a macro dependent on a pre-existing simulation
- Modifying the macro to play up to any section

Macros can also be run in batch mode.

Playing an Initial or Single Macro

Macros that do not have precursor solutions can be run by creating a new simulation, clicking the (**Play Macro**) toolbar button and selecting the appropriate macro file. The file executes dynamically.

Playing a Macro Dependent on an Existing Simulation

For cases that require existing solutions to be present, the recommended procedure is to load the appropriate file from the downloaded tutorial files before running the macro. For example, to execute the `SolarCollector.java` macro it is necessary to first load the `ThermalInsulator_final.sim` file.

It is possible also to use the results of an upstream macro as the start of the downstream solution. So, using a single simulation, you can run `Introduction1.java` and then`Introduction2AdjustingSolverParamters.java` to obtain the solution to the final introductory case.

Modifying the Macro to Play up to Any Page

The automation tutorials provide guidelines on how macros can be modified to suit your needs. The tutorial macros have been structured in a way that allows you to run each case up to a particular section.

To play, for example, the macro for the Domestic Fuse tutorial up to the step, “Visualizing the Volume Mesh”, do the following:

1. Open the file, `DomesticFuse.java`, in a text editor of your choice.
2. Comment out the following pages in the playlist section using block comments as shown below.

```
public class DomesticFuse extends StarMacro {

    @Override
      public void execute() {
        play(new ImportingTheSurfaceMeshAndNamingTheSimulation());
        play(new VisualizingTheImportedGeometry());
        play(new ExaminingTheImportedGeometry());
        play(new IncreasingTheTessellationDensityOfTheFuseWire());
        play(new AssigningGeometryPartsToRegions());
        play(new SettingUpTheMesh());
        play(new AddingAVolumetricControl());
        play(new GeneratingTheMesh());
        /*play(new VisualizingTheVolumeMesh());
        play(new SelectingCopperOhmicHeatingPhysicsModels());
        play(new SpecifyingCopperOhmicHeatingMaterialProperties());
        play(new SelectingCeramicPhysicsModels());
        play(new SpecifyingCeramicMaterialProperties());
        play(new SelectingSiliciumPhysicsModels());
        play(new SpecifyingSiliciumMaterialProperties());
        play(new AssigningPhysicsContinua());
        play(new SettingGlobalInitialConditions());
        play(new DefiningBoundaryConditions());
        play(new SettingUpReportsMonitorsAndPlots());
        play(new SettingUpAPlaneSection());
        play(new SettingUpScalarScene());
        play(new SettingSolverVariables());
        play(new SettingStoppingCriteria());
        if (StarTestEnv.getSingleton().getStartestHome() == null) {
            play (new RunningTheSimulation());
          }
        play(new TemperatureAndElectricPotentialPlot());
        play(new ReportPlots());*/
    }
```

3. Save the macro.
4. Start a simulation and run the modified macro.

Using the Final Simulation Files

For each of the macros that are provided, a corresponding simulation file is also provided. These can be used to check the settings in the model tree of a case you have manually created from the written tutorial instructions.

As tutorial meshes are often coarse, and the solvers tend to be run for a minimum number of time steps, it is not recommended that tutorials are used for verification or benchmarking. They have not been created with this objective in mind.

You can use the supplied simulation files in those tutorials that begin with loading an existing simulation file of a previous tutorial simulation. However, the names of these supplied files differ from the names in the tutorial exercises. The following table lists the title of each applicable tutorial, the name of the simulation file that is opened at the start of that tutorial, and the path and name of the corresponding downloadable simulation file.

Tutorial	Simulation Filename Used	Supplied Simulation File and Path
Unsteady Multi-Component Flow: Dilution Pipe	`dilPipeSteady.sim`	`incompressibleFlow/SteadyMultiComponentFlowDilutionPipe_final.sim`
Porous Resistance: Orthotropic Media	`isotropicPorousMedia.sim`	`incompressibleFlow/PorousResistanceIsotropicMedia_final.sim`
Harmonic Balance: Single Stage Periodic Flow	`singleRow.sim`	`compressibleFlow/TurboWizardSingleRow_final.sim`
Multiband Surface-to-Surface Radiation: Solar Collector	`thermalInsulator.sim`	`heatTransferAndRadiation/SurfaceToSurfaceRadiationThermalInsulator_final.sim`
Rigid Body Motion: Rotating Fan	`fan_mrf.sim`	`motion/MovingReferenceFra`

Tutorial	Simulation Filename Used	Supplied Simulation File and Path
		mesRotatingFan_final.sim
PPDF Equilibrium: Non-Premixed Propane	ebuCombustion.sim	combustion/ThreeStepEddyBreakUpNonPremixedPropane_final.sim
PPDF Flamelets: Non-Premixed Hydrogen	ebuCombustion.sim	combustion/ThreeStepEddyBreakUpNonPremixedPropane_final.sim
Droplet Combustion: Non-Adiabatic Toluene	ebuCombustion.sim	combustion/ThreeStepEddyBreakUpNonPremixedPropane_final.sim
DES and FW-H On-The-Fly: Noise from a Cylinder (Unsteady Analysis)	cylinder.sim	aeroacoustics/BroadbandModelsNoiseFromACylinderPreparation2RefiningVolMesh_final.sim
Adjoint Flow: Wing Shape Optimization	adjointWing_adjoint.sim	incompressibleFlow/AdjointExternalFlowDualElementWing2AdjointFlowSolution_final

Introduction

Welcome to the STAR-CCM+ introductory tutorial. In this tutorial, you explore the important concepts and workflow. Complete this tutorial before attempting any others.

Throughout this tutorial, links to other sections of the online documentation explore important concepts. For example, for a clearer understanding of the changes to typeface in the tutorials, refer to the typographic conventions.

This tutorial is useful in addition to STAR-CCM+ training. For more help, contact your local CD-adapco office. A list of contacts is found at: http://www.cd-adapco.com/about/locations.html.

The case is a transonic flow over an idealized symmetrical blunt body in a wind tunnel.



The tutorial workflow includes:

- Importing the geometry files.
- Generating a polyhedral mesh.

- Setting boundary names and types.
- Defining the models in the continua and applying them to the regions.
- Defining the region conditions, values, and boundary conditions.
- Running the simulation.
- Post-processing the results.

Contents

Starting a STAR-CCM+ Simulation
Saving and Naming a Simulation
Importing the Geometry
Visualizing the Imported Geometry
Defining Boundary Surfaces
Assigning Parts to Regions
Setting Boundary Types
Generating the Mesh
Selecting the Physics Models
Setting Initial Conditions
Defining the Region Continuum
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Monitoring Simulation Progress
Running the Simulation
Adjusting Solver Parameters and Continuing
Visualizing the Results
Plotting Data on a Slice
Adding Streamlines
Closing and Reopening the Simulation
Summary

Starting a STAR-CCM+ Simulation

Start the simulation and familiarize yourself with the STAR-CCM+ user interface.

1. Launch STAR-CCM+ using the appropriate instructions for your operating system.

   After a brief display of the splash screen, the STAR-CCM+ client workspace opens without loading an existing simulation or creating a simulation. This interface to the STAR-CCM+ software is a self-contained graphical user interface (GUI) with panes and subwindows. Some of the GUI terminology is shown in the following screenshot.



   The menu bar provides access to application-wide actions, with some of the more important actions being duplicated in the toolbar.

2. Let the mouse hover over any of the buttons in the toolbar.

   A short description of what that button does appears in a tooltip. In this case, the tool tip shows the name of the bar.

Contents

Creating a Simulation
Working with Objects

Creating a Simulation

Creating a simulation is the first step in a new STAR-CCM+ analysis.

STAR-CCM+ is a client-server application with the client (user interface or batch interpreter) running in one process, and the server (the solver) running in another process. Start a server process on the same machine as the client:

1. Start a simulation by selecting **File > New Simulation** from the menu bar.

   The *Create a New Simulation* dialog appears.



2. In the *Create a New Simulation* dialog, click **OK**.

   A new window containing a simulation object tree is created in the Explorer pane, with the name Star 1. The initial folder nodes for this simulation are shown in the following screenshot. Other nodes are added to the object tree as you progress further.



As the new simulation is created, a window that is named Output appears in the lower-right portion of the STAR-CCM+ workspace. The *Output* window describes the progress of actions in the simulation.

Working with Objects

Much of your interaction with the simulation is through the objects in the simulation tree that was added to the *Explorer* pane.

Familiarize yourself with the objects in the simulation tree:



The tree represents all the objects in the simulation. Nodes are added in later sections when a geometry is imported and models are defined in the continuum. Most of your interaction with the simulation is by selecting nodes in the tree and:

- Right-clicking to expose an action menu.
- Using keys to, for example, copy and delete objects.
- Dragging objects to other tree nodes or onto visualization displays.

The handle next to a node indicates that subnodes exist below that one. To open a node and show the subnodes, click the handle. To close it, do the same for an open node.

Most objects in the tree have one or more properties that define the object. Access to the properties is through the table.



To modify the properties of an object, click its node once to select it. Edit most types of properties in the value cell. Otherwise, when setting values of complex properties, click the **Property Customizer** button to the right of the value. A property-specific dialog opens.

Saving and Naming a Simulation

The simulation file (.sim) contains STAR-CCM+ analysis setup, runtime, and results information.Here, you save the simulation to disk.

To save the simulation:

1. Select **File > Save As**.
2. Navigate to the directory where you want to locate the file.



3. In the *Save* dialog, type `bluntBody.sim` into the **File Name** text box and click **Save**.

   The title of the simulation window in the *Explorer* pane updates to reflect the new name.

It is useful to save work in progress periodically. This tutorial includes reminders to save the simulation at the end of each section. If you would like to return to a stage of the tutorial at a later stage, save it with a unique name.

Importing the Geometry

STAR-CCM+ provides tools to generate the geometry, but for this tutorial you start from a pre-generated Parasolid file of the design.

Before solving this case, you can recognize that the flow is likely to be symmetrical in two planes around the body. Hence, only a quarter of the geometry is modeled. This symmetry condition reduces computational costs and solution time without losing information or accuracy.

The first step in this tutorial is to import the geometry:

1. Select **File > Import > Import Surface Mesh** from the menu bar.
2. In the *Open* dialog, navigate to the `doc/startutorialsdata/introduction/data` subdirectory of your STAR-CCM+ installation directory and select file `bluntBody.x_t`.



3. To start the import, click **Open**.
   The *Import Surface Options* dialog appears.
4. Select **Create new Part** from the *Import Mode* box.



5. Click **OK** to import the geometry.

   STAR-CCM+ provides feedback on the import process in the *Output* window. A new geometry scene is created in the *Graphics* window and shows the imported geometry.

6. In the simulation tree, expand the **Geometry** > **Parts** node to see the new `bluntBody` part.



The imported geometry represents the fluid volume around the body in the wind tunnel.

7. Save the simulation by clicking **Save**.

Visualizing the Imported Geometry

The *Geometry Scene 1* display is already present in the *Graphics* window. Initially, solid, colored surfaces show the surfaces of the geometry.

View the imported geometry in the Graphics window.



A new geometry scene contains:

- A part displayer, **Geometry 1**, which contains all faces in the geometry part and is preset to display a shaded surface.
- A second part displayer, **Outline**, which also contains all faces in the geometry part and is preset to display the mesh outline.

Contents

Panning, Zooming, and Rotating the View

Panning, Zooming, and Rotating the View

Familiarize yourself with the options to change the view of the geometry.

In the *Graphics* window, all three mouse buttons have drag operations that change the view of the geometry:

- To rotate about the selected point, hold down the left mouse button and drag.
- To zoom in or out, hold down the middle mouse button and drag.
- To translate or pan, hold down the right mouse button and drag.
- To rotate around an axis perpendicular to the screen, press the `<Ctrl>` key and hold down the left mouse button while dragging.

There are also several "hot" keys that rotate the view:

- To align with the X-Y plane, press the `<T>` key.
- To align with the Y-Z plane, press the `<F>` key.
- To align with the Z-X plane, press the `<S>` key.
- To fit the view within the *Graphics* window, press the `<R>` key.

1. Adjust the view as shown in the following screenshot.



2. Save the simulation.

Defining Boundary Surfaces

Specify the boundary type for each surface.

The following illustration gives an overview of how you set up this problem.



Index	Description
1	Stagnation Boundary
2	Wall
3	Symmetry Planes
4	Pressure Boundary
5	Slip Walls

Split the surfaces that form specific boundaries:

1. Open the **Geometry** > **Parts** > **bluntBody** > **Surfaces** > **Faces** node.
2. Right-click the **Faces** node and from the pop-up menu, select **Split by Patch**.



The *Split Part Surface by Patch* dialog appears.

3. In *Geometry Scene 1*, select the face in the low X direction, rotating the geometry as necessary.



4. Type `Inlet` in the **Part Surface Name** field.



5. Click **Create**.

   The face is removed from the scene and the new surface appears within the surfaces manager node belonging to the parent part.

6. Repeat the previous steps for the following surfaces:

Location	Corresponding part surface patch number	Name
High X face	3	`Pressure`
Low Z face	4	`Symmetry_plane1`
Low Y face	6	`Symmetry_plane2`

The outside walls of the region can be defined as a single boundary by combining two faces into one part surface using the multi-select functionality.

7. Hold down the `<Ctrl>` key and click both the High Y and High Z faces.
8. Type `Slip_wall` in the **Part Surface Name**.

9. Click **Create**.

   The remaining patches remain as part of the default surface.



10. Click **Close**.
11. Save the simulation.

Contents

Renaming the Surface and Part

Renaming the Surface and Part

Renaming nodes is a common operation in STAR-CCM+. There are two ways a node can be renamed.

Rename the **Faces** node and the **bluntBody** part:

1. Select the **Geometry** > **Parts** > **bluntBody** node.
2. Press <`F2`>.
3. Rename the part `subdomain-1`.



4. Press <`Enter`> or click another node to accept.

To rename the Faces part surface, the previous method can be used. Use the alternative right-click menu option:

5. Right-click the **Geometry** > **Parts** > **subdomain-1** > **Surfaces** > **Faces** node and select **Rename...**.



6. Rename the part surface `Inner_wall`.

   The parts are ready to be converted to a region and boundaries.
7. Save the simulation.

Assigning Parts to Regions

Geometry parts are used to prepare the spatial representation of the model. The computational model to which physics can be applied is defined in terms of regions, boundaries, and interfaces.

Create a region and associated boundaries from the geometry part and its surfaces:

1. Right-click **Geometry > Parts > subdomain-1** and select **Assign Parts to Regions...**.



The *Assign Parts to Regions* dialog appears.

2. Select **Create a Region for Each Part**.
3. Select **Create a Boundary for Each Part Surface**.



4. Click **Apply**.
5. Close the dialog.

   The portion of the object tree below the **Regions** node appears as shown below. All of the surfaces appear as individual boundaries within the region.



6. Save the simulation.

Setting Boundary Types

Having previously given the boundaries sensible names, you can set the boundary types.

Specify the appropriate boundary types:

1. Select the **Symmetry_plane1** node and set **Type** to **Symmetry Plane**.



   The boundary node icon changes to reflect the new type.

2. Using the same technique, change **Symmetry_plane2** to **Symmetry Plane**.

For compressible flows, the most appropriate inflow and outflow types are stagnation inlet and pressure outlet.

3. Change the type of the **Inlet** boundary to **Stagnation Inlet** and the type of the **Pressure** boundary to **Pressure Outlet**.

   The **Slip_wall** and **Inner_wall** boundaries retain the default **Wall** type. Slip walls are boundary conditions and are set up later.

4. Save the simulation.

Contents

Selecting Parts
Adding and Removing Parts from a Scene

Selecting Parts

STAR-CCM+ highlights selected parts in the visualization window.

To highlight parts:

1. Click a part in the *Geometry Scene 1* display, for example the `High X` face.

   The object becomes highlighted and a label appears with the name of the object selected.



In the **bluntBody** object tree, the node that corresponds to this object is also highlighted.



Conversely, if a boundary node is selected in the **bluntBody** tree, the surface corresponding to that boundary is highlighted in the *Graphics* window.

Adding and Removing Parts from a Scene

Use the part selector dialog to add and remove parts from a scene.

To add and remove some parts:

1. Open the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** > **Parts** node and click the ellipsis (**Custom Editor**) for the **Parts** value.



The *Parts* dialog appears.



2. Click (**Column Selection**) to clear the current selections.
3. Expand **Regions > subdomain-1 > Boundaries**.
4. Select **Inner_wall**.



5. Click **OK**.

   This selection removes the solid-colored boundaries representing the wind tunnel from the scene, leaving the outlines and the blunt body.



Add those parts back to the scene using a drag-and-drop method:

6. Open the **Regions** > **subdomain-1** > **Boundaries** node.
7. To extend your selection, hold down the <Ctrl> key and select all the boundary nodes except **Inner_wall**.
8. Release the <Ctrl> key but continue to keep the left mouse button pressed.
9. Drag the nodes onto the display and release the left mouse button.



   A pop-up menu appears.

Use this menu to choose which of the part displayers in the scene receive the parts that you have dragged across.

10. Select **Add to Geometry 1**.



The parts are restored to the scene.

11. Save the simulation.

Generating the Mesh

Several steps are involved in defining and generating a volume mesh.

To generate the mesh:

1. Create a mesh continuum.
2. Choose appropriate meshing models.
3. Specify initial global settings.
4. Modify the boundary-specific settings.

Contents

Creating the Mesh Continuum
Choosing Meshing Models
Specifying the Mesh Settings
Generating the Volume Mesh
Displaying the Volume Mesh
Refining the Mesh

Creating the Mesh Continuum

The mesh continuum is used to specify the required meshing models.

To create the mesh continuum:

1. Right-click **Continua**.
2. Select **New** > **Mesh Continuum**.



A new node, **Mesh 1**, is added to the simulation tree.



Choosing Meshing Models

The meshing continuum has a Models node which you can now populate.

To choose the meshing models:

1. Right-click the **Mesh 1** node and choose **Select Meshing Models...**.



In the *Mesh 1 Model Selection* dialog:

2. Select **Surface Remesher** from the *Surface Mesh* group box.
3. Select **Polyhedral Mesher** and **Prism Layer Mesher** from the *Optional Models* group box.

   The Enabled Models section appears as shown below.



4. Click **Close**.
5. Expand the **Models** and **Reference Values** managers.

   The simulation tree appears as shown below.



Specifying the Mesh Settings

Generating a mesh often requires several iterations to achieve the desired density and distribution of cells. In this tutorial, you generate a volume mesh, add some simple refinement, and regenerate.

To specify the mesh settings:

1. Select the **Continua > Mesh 1 > Reference Values > Base Size** node and set **Value** to `0.01 m`.

   With this base size, approximately 8 cells are created across the width of the region.



2. Select **Continua > Mesh 1 > Reference Values > Number of Prism Layers** and set **Numbers of Prism Layers** to `5`.



Generating the Volume Mesh

Now that you have specified the appropriate settings, you can generate the volume mesh.

1. Click (**Generate Volume Mesh**) in the toolbar or select **Generate Volume Mesh** in the **Mesh** menu.

   The run and progress of the meshers are displayed in the *Output* window.



Displaying the Volume Mesh

Display the mesh to permit its inspection in the existing displayer.

To display the volume mesh:

1. Select the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** node.
2. To display the surface of the mesh, activate **Mesh**.



3. Right-click an empty part of the display and select **Apply Representation** > **Volume Mesh**.



This action reveals the mesh on the boundaries.



4. Save the simulation.

Refining the Mesh

The prism layer has been generated on the inner and the slip walls by default. Because a fluid boundary layer does not form on slip walls, you can customize the prism mesh settings on that boundary.

To refine the mesh:

1. Select the **Regions** > **subdomain-1** > **Boundaries** > **Slip_wall** > **Mesh Conditions** node and set **Customize Prism Mesh** to **Disable**.



To provide better definition of the blunt body, modify the surface mesh at the inner wall:

2. Select the **Regions** > **subdomain-1** > **Boundaries** > **Inner_wall** > **Mesh Conditions** > **Custom Surface Size** node and activate **Custom Surface Size**.

   A new **Mesh Values** node appears.
3. Select the **Surface Size** node and set **Relative/Absolute** to **Absolute**.
4. Select the **Surface Size** > **Absolute Minimum Size** node and make sure that **Value** is `0.0010 m`.

To regenerate the mesh:

5. Click (**Generate Volume Mesh**).
6. Zoom in to the area around the slip walls and inner walls to see the improved mesh.



This mesh is satisfactory for an initial solution.

7. Save the simulation.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature and velocity, and the mathematical formulation.

In STAR-CCM+, the physics models are defined on a physics continuum. In this example, the flow is turbulent and compressible. You use the Coupled Flow model together with the default K-Epsilon Turbulence model.

To select the physics models:

1. Right-click the **Continua** > **Physics 1** node and choose **Select models...**.



The *Physics Model Selection* dialog appears as shown in the following screenshot.



2. Select the **Gas** radio button from the *Material* group box, since this exercise involves an idealized gas (air).

Since the **Auto-select recommended models** checkbox is activated, the *Physics Model Selection* dialog guides you through the model selection process by selecting certain default models automatically as you make some choices.

Certain models, when activated in a continuum, require other models also to be activated in that continuum. For instance, once a continuum contains a liquid or a gas, it also needs a flow model. Once it has a flow model, it needs a viscous model (inviscid, laminar, or turbulent). Once turbulence is activated within a fluid continuum, select a turbulence model. The prompt **Additional model selections are required** alerts you to the fact that you have not completed the model selection.



3. The following selections are required for this simulation:
   a) Select **Coupled Flow** from the *Flow* group box.
   b) Select **Ideal Gas** from the *Equation of State* group box.
      The **Coupled Energy** model is selected automatically.
   c) Select **Steady** from the *Time* group box.
   d) Select **Turbulent** from the *Viscous Regime* group box.
   e) Select **K-Epsilon Turbulence** from the *Reynolds-Averaged Turbulence* group box.
      The **Realizable K-Epsilon Two-Layer** and the **Two-Layer All y+ Wall Treatment** models are selected automatically.

To reverse part or all of the model selection process, simply clear the checkboxes of the models you wish to deactivate. Other active models require the selections that are grayed out. Therefore, deactivate the models that are not grayed out to begin with.

When complete, the *Physics Model Selection* dialog appears as shown in the following screenshot:



No optional models are required for this simulation.

4. Click **Close**.
5. To see the selected models in the simulation tree, open the **Continua** node in the *bluntBody* window of the *Explorer* pane.

   The color of the **Physics 1** node has turned from gray to blue to indicate that models have been selected.
6. To display the selected models, open the **Physics 1** node and then the **Models** node.



You can view the model properties in the object tree.



When you select the **Gas** model, the properties of air such as dynamic viscosity are used by default. Since this problem uses air, the properties are acceptable.

7. Save the simulation.

Setting Initial Conditions

Initial conditions in a continuum specify the initial field data for the simulation.

Examples of initial conditions are:

- Pressure
- Temperature
- Velocity components
- Turbulence quantities

Each model requires sufficient information so that the primary variables of the model can be set. For some models, such as turbulence models, the option of specifying the information in a more convenient form is presented. For example, turbulence intensity and turbulent viscosity ratio instead of turbulent kinetic energy and turbulent dissipation rate.

In steady-state simulations, the solution ought to converge independently of the initial field. However, the initial field still affects the path to convergence, and with it the cost in computing power. Therefore specify initial conditions and values judiciously, particularly when the physics is complex.

Contents

Specifying the Initial Conditions for the Simulation

Specifying the Initial Conditions for the Simulation

The stagnation inlet boundary has conditions that correspond to a Mach number of 0.75. The equivalent freestream velocity is roughly 300 m/s, the value you use to initialize the velocity field.

To specify the initial conditions:

1. Open the **Continua** > **Physics 1** > **Initial Conditions** node.
2. Open the **Velocity** > **Constant** node and set **Value** to `300,0,0`.
3. In the same **Initial Conditions** node, open the **Turbulent Viscosity Ratio** node and select the **Constant** node.
4. Set **Value** to `50`, which is the same as the turbulent viscosity ratio that is set on the stagnation boundary condition.
5. Save the simulation.

Defining the Region Continuum

You can associate a continuum with zero, one, or more regions. Multiple continua can be defined and each region assigned a different one. In this tutorial, only one geometry part is present, which requires a single region to represent it, which was named `subdomain-1`. There is a simple way to determine which set of continua models are allocated to a region.

To define the region continuum:

1. Select the **subdomain-1** node in the **Regions** node.



2. Set the **Physics Continuum** property to **Physics 1**, the name of the continuum where you defined all the models earlier.



During this tutorial, there is one continuum in the above drop-down list, but as you add more (to the **Continua** node in the simulation tree), they are added to the drop-down list. This process determines which continua is used in the analysis, though continua can be defined and left unused.

Setting Boundary Conditions and Values

Although the concept of a boundary type is fairly unambiguous (wall, stagnation inlet, and so on.), models need more information to deal with the type. Conditions provide this information.

For example, with a boundary of type wall, the conditions specify whether this wall is to be a no-slip wall, a slip wall, or a moving wall. The conditions also tell you whether you want to apply a specified temperature (Dirichlet) thermal boundary condition or a specified heat flux (Neumann) thermal boundary condition.

The types and conditions inform models how to deal with a boundary (or region or interface) but they do not specify actual numerical input. Values provide this input. Value nodes are added in response to choices made on conditions nodes.

Contents

Setting Inlet Conditions and Values
Setting the Slip Wall Condition

Setting Inlet Conditions and Values

Set the appropriate conditions and values for the Inlet boundary.

To achieve a Mach number of approximately 0.75 at the inlet, isentropic relations were used to determine the inlet total pressure and outlet static pressure for a given total temperature. For an outlet static pressure equal to one atmosphere (absolute) and a static temperature of 300 K, the inlet total pressure is 164,904 Pa (absolute). The inlet total temperature is 344.8 K.

The boundary values are specified as gauge pressures. Therefore, using the default reference pressure of 101,325 Pa (one atmosphere), set the inlet total pressure to 63,579 Pa (gauge). The outlet static pressure is 0 Pa (gauge), the default value.

To set the Inlet conditions and values:

1. Open the **Regions** > **subdomain-1** > **Boundaries** node.
2. Open the **Inlet**, **Physics Values** and **Total Pressure** nodes, and select the **Constant** node.



3. Set **Value** to `63579`.

The total pressure setting is relative to the operating pressure value of 101,325 Pa. A supersonic static pressure is also required as part of a stagnation inlet condition but is only used if the inlet velocity becomes supersonic at some instance during solution iteration. The default value of 0.0 Pa relative pressure is sufficient as long as supersonic flow does not occur.

The next value to set is the total temperature.

4. In the same **Physics Values** node, open the **Total Temperature** node and select the **Constant** node.



5. Set the **Value** to `344.8`.
6. Select the **Physics Values** > **Turbulent Viscosity Ratio** > **Constant** node.



7. Set **Value** to `50`.

   This value has been determined to give a reasonable decay of turbulence in the channel core (for the default turbulence intensity).

Setting the Slip Wall Condition

To simulate the body in a wind tunnel, the upper boundary is set to a slip wall, thus avoiding the need to resolve the boundary layer on this wall.

To set the Slip Wall condition:

1. Select the **Slip_wall** > **Physics Conditions** > **Shear Stress Specification** node.



2. Set **Method** to **Slip**.

Setting Solver Parameters and Stopping Criteria

Set the appropriate stopping criteria for the simulation.

To set the stopping criterion:

1. Select the **Stopping Criteria** > **Maximum Steps** node.



2. Set the **Maximum Steps** to `300`.

The solution does not converge in this number of iterations. The solution does not run for more than 300 iterations, unless this stopping criterion is changed or deactivated.

3. Save the simulation.

Visualizing the Solution

Create a scalar scene to display the results of the simulation.

To watch the Mach number on the blunt body and vertical symmetry plane solution developing, set up a scalar scene.

To visualize the solution:

1. Right-click the **Scenes** node and select **New Scene > Scalar**.



A new *Scalar Scene 1* display appears.



Hide all parts except the blunt body itself and the vertical symmetry plane.

2. Select the new **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Parts** node.



3. Click the (**Custom Editor**) in the right half of the **Parts** property.
4. In the *Parts* dialog, expand **Regions > subdomain-1**, select **Inner_wall** and **Symmetry_plane1**, and click **OK**.



Define Mach number as the scalar to display:

5. Right-click the scalar bar (near the bottom of the *Graphics* window).
6. Select **Mach Number > Lab Reference Frame** from the pop-up menu.



Since the geometry of this example is symmetric, the mesh was cut into quarters with two symmetry planes to reduce computing costs. However, the symmetric repeat transform lets you create the visual effect of the complete geometry by setting up the mirror image of the model in the Graphics window. In this case, you only do one repeat so that half the model is shown.

7. Select the **Scalar 1** node.



8. Set the **Transform** expert property to **Symmetry_plane2 1**.



To remove the outline:

9. Select the **Displayers > Outline 1** node.
10. Clear the checkbox of the **Outline** property.



The scalar scene appears as below.



11. Save the simulation.

Monitoring Simulation Progress

STAR-CCM+ can dynamically monitor a quantity of interest while the solution develops.

The process involves:

- Setting up a report that defines the quantity of interest and the region parts that are monitored

  See Setting Up a Report.
- Defining a monitor using that report, which controls the update frequency and normalization characteristics
- Setting up an X-Y plot using that monitor

  See Setting Up a Monitor and a Plot.

For this case, force is monitored on the body in the x-direction of the flow, which effectively is the total drag force. This process starts with the report definition.

Contents

Setting Up a Report
Setting Up a Monitor and a Plot

Setting Up a Report

Create a report to collect the simulation results.

To set up a report:

1. Right-click the **Reports** node and select **New Report > Force Coefficient**.



This action creates a **Force Coefficient 1** node under the **Reports** node.

2. Select the **Force Coefficient 1** node and enter the settings for the report in the *Properties* window.

   a) Enter `1.277` for **Reference Density**, the density of the freestream air.

   b) Enter `264.6` for **Reference Velocity**, the velocity at the inlet.

   c) Enter a value of `0.0161269` for **Reference Area**, the projected area of the quarter of the blunt body that is used in the simulation.

   d) Make sure that **Direction** is `[1.0,0.0,0.0]` (for drag).



Use the drag-and-drop method to select the monitored part, namely the **Inner_wall** boundary:

3. Open the **Regions > subdomain-1 > Boundaries** node and then drag the **Inner_wall** node onto the **Force Coefficient 1** node.



The **Inner_wall** entry is now listed in the **Parts** property of the **Force Coefficient 1** node.

The setup for the report is now complete, and a monitor and a plot can be made from that report.

Setting Up a Monitor and a Plot

Create a monitor and a plot to display the report data.

1. Right-click the **Force Coefficient 1** node and select **Create Monitor and Plot from Report**.



This action creates a **Force Coefficient 1 Monitor** node under the **Monitors** node.

The existing monitors in this branch of the simulation tree are the residual monitors from the solvers that the models use. With the new monitor node selected, the default settings for the **Force Coefficient 1** monitor are seen.



These settings update the plot every iteration while the solution is running.

In addition to the new monitor node, a **Force Coefficient 1 Monitor Plot** node is created in the **Plots** node.



2. Select the new **Force Coefficient 1 Monitor Plot** node.

   The property settings for the plot appear.



Additional properties of the plot can be adjusted using the subnodes of the **Force Coefficient 1 Monitor Plot** node.

3. To view the plot display, do one of the following:

   - Double-click the **Force Coefficient 1 Monitor Plot** node.

     or

   - Right-click the **Force Coefficient 1 Monitor Plot** node and select **Open** from the pop-up menu.

   A *Force Coefficient 1 Monitor Plot* appears in the *Graphics* window.



This plot display is automatically updated with the drag force coefficient when the solver starts to run.

4. Save the simulation.

The pre-processing setup is completed. The simulation can now be run to convergence.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Do one of the following:

   - Click (**Run**) in the *Solution* toolbar.
   - Use the **Solution** > **Run** menu item.

   The *Residuals* display is created automatically and shows the progress of the solver. If necessary, click the *Residuals* tab to bring the *Residuals* plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.

The tabs at the top of the *Graphics* window allow you to select any one of the active displays for viewing.

2. To see the results, click the tab of the *Scalar Scene 1* display.
3. Rotate and zoom if desired for a better view.



4. During the run, it is possible to stop the process by clicking (**Stop**) in the toolbar.

   If you do halt the simulation, you can resume it by clicking (**Run**). If left alone, the simulation continues until the stopping criterion of 300 iterations is satisfied.

5. After the run is finished, save the simulation.

Adjusting Solver Parameters and Continuing

The default Courant number of 5 is conservative for this type of problem on a mesh of reasonable quality, so it can be increased to accelerate convergence. When the solution has reached the stopping criterion of 300 iterations, adjust the solver parameters.

To adjust the solver parameters:

1. Select the **Solvers** > **Coupled Implicit** node and set **Courant Number** to `10.0`.



2. Select the **Stopping Criteria** > **Maximum Steps** node and set the **Maximum Steps** to `1500`.

   This setting instructs STAR-CCM+ to iterate 1200 steps in addition to the 300 already done.

The maximum number of 1500 steps is a reasonable setting, in this case, to get to convergence. Normally, you do not know ahead of time what a suitable iteration count is. You can either set a large iteration and watch the residuals or a monitor plot, or add a stopping criterion from a monitor.

3. To continue running the simulation, click (**Run**).

Contents

Disconnecting and Reconnecting
Completing the Run

Disconnecting and Reconnecting

The STAR-CCM+ client-server design permits you to run a simulation, disconnect the client from the running server, and then reconnect later

to check on results. You can also reconnect from a different machine, even a machine of another architecture.

To disconnect the client from the server:

1. Open the *Servers* window (if it is not already open) by selecting **Window > Servers**.



2. To see your simulation, look in the *Servers* window of the *Explorer* pane.
3. Reselect the simulation window in the *Explorer* pane, and select the **File > Disconnect** menu item.



4. In the *OK to disconnect* dialog, click **OK**.

   The client disconnects from the server, but both are still running, although the simulation window no longer appears in the client. If desired, the client can be shut down using **File > Exit**.

   The *Servers* window of the *Explorer* pane lists all server processes (simulations) that are running locally.

5. Selecting the node in the tree displays the server process properties.



6. Continue by right-clicking the server process node in the **Servers** tree and selecting **Connect** from the pop-up menu.



When you reconnect, none of the previously created scenes and plots are displayed in the *Graphics* window.

7. To display a scene of interest, do one of the following:
   - Double-click its node in the simulation object tree.

     or
   - Right-click the node and select **Open**.

Completing the Run

At the end of the calculation, the *Residuals* display show most of the residuals flattening out, which is a good indicator of convergence.

To confirm that the solution has converged:

1. Select the *Force Coefficient 1 Monitor Plot* display to verify that the drag coefficient has also flattened out.



2. Once the run is finished, save the simulation.

Visualizing the Results

View the results of the simulation using the scalar and vector scenes that you prepared earlier.

Once the solution is finished, you can examine the results in scalar scenes and vector scenes.

Contents

Examining Scalars
Examining Vectors

Examining Scalars

View the results of the simulation in a scalar scene.

To see the Mach number results for the finished solution:

1. Make the *Scalar Scene 1* display active.



By default, **Filled** cell values are shown.

To display smooth contours:

2. Select the **Scalar 1** node and set the **Contour Style** to (**Smooth Filled**).

   The contours of the scalar display now appear smooth.



Display two scalar values at the same time: one on the symmetry plane, the other on the inner wall:

3. Add another scalar displayer by right-clicking the **Displayers** node and selecting **New Displayer > Scalar**.



The creation of the second displayer has added a second scalar bar in the display.

To view both scalar bars:

4. Point the mouse in the interior of a scalar bar to drag either one of them to a different location.



5. In the display, right-click the scalar bar (blue) of the new displayer and select **Pressure**.



6. Right-click the inner wall in the display, select **Displayers**, and open its submenu.
7. Among the checkboxes to the right of **in**, clear the checkbox under **Scalar 1** and tick the one under **Scalar 2**.



This single step transfers the part from one displayer to another.

8. Select the **Scenes > Scalar Scene 1 > Displayers > Scalar 2** node.



9. Set the **Contour Style** to (**Smooth Filled**).
10. Set the **Transform** expert property to `Symmetry_plane2 1`.

    The blunt body shows pressure values while the symmetry plane shows Mach number.



11. Save the simulation.

Examining Vectors

This section demonstrates how to examine velocity vectors in the model. In this example, you set up vectors in a reflected symmetry plane around a solid-colored blunt body.

1. First, create a scene by right-clicking the **Scenes** node.
2. Select **New Scene** > **Vector**.



A new *Vector Scene 1* display appears.

You do not need the outlines that the **Outline** part displayer provides by default, so reuse that displayer to show the blunt body as a solid object.

3. Select the **Vector Scene 1** > **Displayers** > **Outline 1** > **Parts** node and click (**Custom Editor**).

4. In the *Parts* dialog, click the column selection button twice.
5. Expand the **Regions** and **subdomain-1** nodes.
6. Select the **Inner_wall** node.
7. Click **OK**.



8. Select the **Outline 1** node in the **Displayers** node.



9. Make the following changes to the **Outline** displayer properties:
   a) Deactivate the **Outline** property.

b) Activate the **Surface** property.

c) Set the **Transform** expert property to `Symmetry_plane2 1`.



10. Select the **Vector 1** > **Parts** node and click the ellipsis (**Custom Editor**) for the **Parts** value.



11. In the *Parts* dialog, expand the **Regions** and **subdomain-1** nodes. Select the **Symmetry_plane1** node and click **OK**.



12. Select the **Vector 1** node.



13. Set the **Transform** expert property to `Symmetry_plane2 1`.

These settings show a solid blunt body with velocity vectors surrounding it.



It is possible to make the arrows shorter to see the vectors more clearly.

14. Select the **Vector 1 > Glyph > Relative Length** node and change **Glyph Length (%)** to `2.0`.

    The arrows now appear shorter.



15. Save the simulation.

Plotting Data on a Slice

Create an x-y plot to visualize the pressure coefficient (on the y-axis) along a slice of the blunt body.

Creating the x-y plot requires the following steps:

1. Create a derived part for the slice.
2. Create a plot using the new derived part.

Contents

Creating a Derived Part
Plotting Simulation Data

Creating a Derived Part

Create a plane section that cuts the blunt body. In STAR-CCM+ terminology, this plane section is called a derived part.

To create the derived part:

1. Right-click the **Scenes** node and select **New Scene** > **Geometry**.



2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane...**.



A new interactive in-place dialog is activated to allow definition of the derived part.

To define the source part for this derived part:

3. In the *Input Parts* group box, click **[subdomain-1]**.
4. In the dialog that appears, click the column selection button twice.
5. Expand the **Regions** and **subdomain-1** nodes.
6. Select the **Inner_wall** boundary and click **OK**.



The plane tool in the scene display allows interactive definition of the plane that cuts the blunt body.

7. Rotate and pan the image as needed so that the inner wall is in front of the plane tool rectangle as shown.



8. In the dialog, leave the *Display* box at its default setting.

The "barbell" represents the normal of the cutting plane, which can be changed by dragging one of the balls to another location. The plane can also be dragged to change the origin of the cutting plane.

A bounding box slightly larger than the regions in the scene is displayed to limit the movement of the plane.

9. Move the mouse over the plane tool, then click-and-drag the rectangle so that it cuts the blunt body close to the symmetry plane.



The properties of the *edit* dialog are similar to those in the following screenshot:



10. Click **Create** and then **Close**.

    In the *bluntBody* window, a node, **Plane Section**, has been added within the **Derived Parts**.



It is now possible to plot data on the new surface.

11. Save the simulation.

Plotting Simulation Data

Set up a plot of pressure coefficients for the new derived part.

To create the plot:

1. In the *bluntBody* window, right-click the **Plots** node and select **New Plot > X-Y**.



An empty plot node is added to the tree and opens in the *Graphics* window.

To select the part on which to plot:

2. Select the **Plots > XY Plot 1** node and click ... (**Custom Editor**) for the **Parts** property.
3. In the *XY Plot 1* dialog, make sure that **Derived parts > plane section** is ticked.
4. Click **OK**.

To set the variable to plot on the y-axis:

5. Select the **XY Plot 1 > Y Types > Y Type 1** node and make sure that the **Type** property is set to **Scalar**.
6. Select the **Y Type 1 > Scalar Function** node and click ... (**Custom Editor**) for the **Field Function** property.
7. In the *Scalar Function - Field Function* dialog, select **Pressure Coefficient**.



8. Click **OK**.
9. Select the **XY Plot 1** > **Axes** > **X Axis** node and ensure that the **Lock Minimum** and **Lock Maximum** properties are deactivated.

Match the reference density, reference pressure, and reference velocity of the **Pressure Coefficient** field function in the **Force Coefficient 1** report. Copy and paste these properties from the report to the field function:

10. Select the **Reports** > **Force Coefficient 1** node.
11. In the *Properties* window, select the value of the **Reference Density** property, and press **<Ctrl+C>**.
12. Select the **Tools** > **Field Functions** > **Pressure Coefficient** node.



13. In the *Properties* window, select the value of the **Reference Density** property, press `<Ctrl+V>`, and then `Enter`.

14. Repeat the above steps for the **Reference Velocity**.
15. Leave the **Reference Pressure** property at `0.0 Pa`.

    The plot appears as shown below.



16. Save the simulation.

Adding Streamlines

Set up a streamline derived part so that it shows the recirculating flow behind the blunt body. Modify the part to show streamlines over the blunt body.

To visualize streamlines in STAR-CCM+, do the following:

- Create a scene.
- Create a streamline derived part.
- Add a streamline displayer, with suitable properties, to the new scene based on the streamline derived part.

To create a scene and add streamlines:

1. Right-click the **Scenes** node and select **New Scene** > **Geometry**.
2. Right-click the **Derived Parts** node.
3. Select **New Part** > **Streamline**.



As with the plane section earlier, an interactive in-place dialog appears.



The entire fluid region is used as the input part for the streamline part.

4. In the *Input Parts* group box, click the **Select** button.
5. In the *Select Objects* dialog, make sure that `subdomain-1` is the only region selected.
6. Click **OK**.
7. In the *Vector Field* group box, make sure that **Velocity** is selected.
8. For the seed mode, select **Point Seed**.



To position the seed point just after the back of the blunt body (in the direction of flow):

9. Enter the following coordinates under **Seed Position**:
   `0.05` for X, `0.01` for Y, and `0.01` for Z
10. Enter `0.002` for **Seed Radius** and `10` for **Number of Points**.



11. Click **Create**.

    After a few seconds, the streamlines are shown in the display.



In the *bluntBody* window, a **streamline** node appears in the **Derived Parts** node.



12. Click **Close** in the interactive in-place dialog.
    A new streamline displayer has been added to the scene.
13. Select the **Displayers > Section Stream 1** node to customize how the streamlines appear:
    a) Set the **Mode** to **Tubes** so that the streamlines are displayed as tubes instead of lines.
    b) Make sure that **Orientation** is set to **Normals**.
    c) Specify **Width** at `2.0E-4`, the width of the tube in meters.



14. In the display, right-click the color bar and select **Velocity: Magnitude**.



Improve the contrast between boundaries and streamlines:

15. Select the **Geometry 1** node and set the **Color Mode** to **Constant**.



You now fine-tune the streamline derived part.

16. Select the**Derived Parts** > **streamline** node and make sure that the **Rotation Scale** expert property is `1.0`.
17. Select the **streamline** > **2nd Order Integrator** node and set the **Integration Direction** to **Both** so that the streamlines are generated both upstream and downstream.

In the Expert properties:

18. Set **Initial Integration Step** to `0.1` to provide more resolution to the streamlines.
19. Set **Maximum Propagation** to `5.0`.
20. Set **Max Steps** to `1000`.

    This option, together with the maximum propagation, provide a stopping criterion to make sure that the streamlines are not calculated endlessly.

    The modified streamlines appear as below.



The next step involves moving the seed point of the streamline derived part to the front of the blunt body.

21. Select the **2nd Order Integrator** node and change **Integration Direction** to **Forward**.
22. Select the **streamline** > **Point Seed** node and set the x-value of the **Center** to `-0.05`, so that the value of the property is `-0.05, 0.01, 0.01`.



The modified streamlines appear in the display.



23. Save the simulation.

Closing and Reopening the Simulation

Close the simulation and reopen it again.

To close a simulation while it is running:

1. Close the *bluntBody* window.

   The following dialog prompts you to choose whether to wait for the operation to finish before closing the window, or close the window and leave the operation to complete in the background.



2. Click **Disconnect**.

It is easy to reopen the file and restore the displays.

3. Select **File** > **Recent Files** > **/bluntBody.sim**.



Summary

This tutorial has introduced the following STAR-CCM+ features:

- Starting STAR-CCM+ and creating a simulation.
- Saving and naming the simulation.
- Importing a geometry file.
- Visualizing the imported geometry.
- Splitting and naming the geometry surfaces.
- Renaming techniques.
- Changing types of boundaries.
- Designing a three-dimensional mesh
- Visualizing the mesh.
- Selecting the physics models.
- Defining the initial conditions.
- Defining the boundary conditions and values.
- Setting the solver parameters and stopping criteria.
- Setting up a monitoring report and plot.
- Running the solver until the residuals are satisfactory.
- Disconnecting from a server and reconnecting during a run.
- Analyzing results using the visualization, monitor-plot, and XY-plot features.
- Creating streamlines.
- Reopening a simulation.

Geometry

The tutorials in this set illustrate various STAR-CCM+ features for creating and working with parts and 3D-CAD.

Contents

3D-CAD: Cyclone Separator
Geometry Parts: CAD Assembly

3D-CAD: Cyclone Separator

This tutorial demonstrates how to build and run a simulation entirely within STAR-CCM+.

You create the geometry for this simulation in 3D-CAD, the modeler available in STAR-CCM+. You then convert the 3D-CAD model to a part and perform a full simulation. The geometry in this simulation is a simplified cyclone separator.



While creating the geometry in 3D-CAD, you set up a design parameter to allow you to modify the geometry from the main simulation tree. The purpose of this simulation is to analyze the effect on the flow pattern when placing the outlet pipe at two different levels within the cyclone chamber.



The following basic concepts of 3D-CAD are introduced in this tutorial:

- Creating a 3D-CAD model
- Creating sketches and applying constraints and dimensions
- Revolving and extruding sketches
- Naming faces
- Using a 3D-CAD model in a simulation
- Modifying the model geometry and rerunning the case

Contents

Prerequisites
Creating the 3D-CAD Model
Creating a Geometry Part
Assigning a Part to a Region
Selecting Physics Models
Generating a Mesh
Preparing a Scalar Scene
Running the Simulation
Creating a Streamlines Scene
Modifying the Geometry and Rerunning
Summary

Prerequisites

The instructions in the 3D-CAD tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Creating the 3D-CAD Model

Create a 3D-CAD model.

1. Launch STAR-CCM+.
2. Save the simulation under the name `cyclone.sim`.

Launch 3D-CAD and create the geometry:

3. Right-click **Geometry** > **3D-CAD Models** and select **New**.



The 3D-CAD environment launches and a 3D-CAD model is created. Within 3D-CAD, an object tree that is specific to the 3D-CAD model is displayed. This tree contains **Bodies**, **Features**, and **Design Parameters** nodes. A *3D-CAD View* scene is also created automatically. Initially, this scene displays three starting sketch planes which are used as a basis for adding features to the model.

4. Rename the **3D-CAD Model 1** node to `Cyclone`.

Contents

Sketching the Cyclone Profile
Revolving the Cyclone Profile
Creating the Outlet Pipe
Creating the Inlet Duct
Creating a Design Parameter
Specifying Inlet and Outlet Faces

Sketching the Cyclone Profile

Create a sketch that contains the profile of the cyclone chamber.

Every 3D-CAD model contains three default sketch planes: **XY**, **YZ**, and **ZX**. Create a sketch on the **YZ** sketch plane:

1. Right-click the **Features** > **YZ** node and select **Create Sketch**.



Creating a sketch activates sketch mode, in which a grid is displayed on the sketch plane in the *3D-CAD View* scene, and the *Sketch* panel opens. This panel contains the tools for creating sketch entities, such as lines, circles, and arcs. The panel also allows you to control the settings for the grid that is displayed in the scene. Each sketch plane has local X and Y axes relative to the position and orientation of the plane in relation to the global coordinate system.

2. Click (**View normal to sketch plane**) to align the sketch plane with the plane of the screen.



3. Click (**Set sketch grid spacing**) to specify the distance between the grid lines on the sketch plane so that the spacing is suitable for the dimensions of the model geometry.
4. In the *Grid Spacing* dialog, set *Grid Spacing* to `0.025 m`.



5. Click **OK**.

Use the mouse to adjust the position, orientation, and zoom level of the sketch plane.

6. In the *3D-CAD View* scene:
   - Left-click and drag the mouse in any direction to rotate the view.
   - While holding down the <`Ctrl`> key, click and drag to roll the view.
   - Right-click and drag to pan the view.
   - Middle-click and drag to zoom in or out. Alternatively, if your mouse has a scroll wheel, turn the mouse wheel to adjust the zoom level.

7. Click (**View normal to sketch plane**) to realign the sketch plane with the screen.
8. Adjust the view in the *3D-CAD View* scene so that the top half of the grid is shown. Leave some white space above the grid as the sketch will cover some of this space.



9. Click (**Create line**) and move the mouse over the grid in the *3D-CAD View* scene.

   The mouse pointer has a blue square that is attached to it, which indicates the position of the point on the grid. The coordinates of this point are displayed in the *Line Properties* box of the *Sketch* panel, and change as you move the mouse. The (**Snap to Grid**) option is enabled by default, so the point jumps to an intersection between two grid lines as it is approaches them.
10. Position the blue square at `[0.0, 0.3] m`.



11. Click to fix the blue square in position.

    Create the first set of lines with approximate lengths. To exit the tool while creating a sketch entity, press the <`Esc`> key. If you want to reposition any sketch entities while you are sketching, exit the tool first, then click-and-drag the relevant sketch entities to the desired position. To resume sketching, simply select the appropriate tool from the *Create Sketch Entities* box and continue.

12. Move the mouse pointer to the left. A blue preview is shown to indicate the position of the line. Position the mouse so that the blue line extends horizontally left of the starting point and click the left mouse button.



The **Line** tool creates continuous lines, so the end point of the first line is the start point of the next. The coordinates that are shown in the *Line Properties* box are automatically updated as each line is defined.

13. Move the mouse pointer into the white space above the grid and notice that it still snaps to grid intersections even though they are not visible.
14. Position the mouse so that the second line is vertical and ends in the white space above the grid.

    You can adjust the zoom level while you are in the process of creating sketch entities.
15. Left-click to fix the end point of the second line.



When you draw sketch entities outside of the grid that is shown in the *3D-CAD View* scene, the grid automatically expands.

16. Create a short horizontal line that extends to the left of this point and left-click to fix the end point in position.



17. Move the mouse vertically downwards to create a longer vertical line and click to fix the point.



18. Press the <Esc> key to exit the line creation tool.

Contents

Adding Constraints and Dimensions
Adding a Construction Line

Adding Constraints and Dimensions

Constraints and dimensions are applied to the sketch to control the position and size of the sketch entities on the sketch plane.

Apply the constraints first to maintain the current orientation of the lines. Following this, apply the dimensions to define the length of each line.

1. Right-click the first horizontal line and select **Apply Horizontal Constraint**.



A glyph is added to the sketch to show that the constraint has been applied.

2. Select one of the horizontal lines, then hold down the <Ctrl> key and select the other horizontal line. Right-click and select **Apply Parallel Constraint**.



Apply similar constraints to the vertical lines:

3. Right-click one of the vertical lines and select **Apply Vertical Constraint.**
4. Using the technique described above, select both of the vertical lines, right-click and select **Apply Parallel Constraint**.



Set the length of each line by applying dimensions:

5. Right-click the first horizontal line and select **Apply Length Dimension**.



6. In the *Dimension* dialog, enter a value of `0.075 m`.



7. Click **OK**.
8. Repeat the previous two steps for the remaining lines. Use the values that are displayed in the following sketch:



The sketch appears as follows:



Notice that when you select each line, a list of the constraints and dimensions that are applied to that line is displayed in the *Constraints/Dimensions* box in the *Line Properties*.

Draw the remainder of the cyclone profile, and use some additional constraints to control its position on the sketch plane:

9. Pan the view down and adjust the zoom level so that the 3D-CAD View scene appears as follows:



10. Click (**Create line**) and move the mouse to the starting point for the previous set of lines.

    The mouse automatically snaps to existing sketch entities within the sketch, so when the blue square is close to the point, it jumps onto that point.
11. Click the top of the existing point.
12. Move the mouse vertically downward until a line *Length* of `0.7 m` is shown in the *Line Properties* box. Left-click to fix the point at the grid intersection.



13. Move the mouse to the left and create a short line that extends horizontally from this point.

Complete the profile by creating a diagonal line to define the cone shape of the cyclone:

14. Move the mouse pointer to the free end of the first set of lines and click to close the sketch profile.
15. Right-click the longest vertical line and select **Apply Vertical Constraint**.
16. Right-click the horizontal line at the bottom of the sketch and select **Apply Length Dimension**.
17. In the *Dimension* dialog, enter a value of `0.05 m` and click **OK**.

Move the axis of the cyclone so that it is aligned with the Y-axis of the sketch plane. To do this you can drag-and-drop the sketch:

18. If necessary, adjust the zoom level. Left-click and hold at any point on the longest vertical line, drag the mouse pointer onto the origin and release the mouse button.



Adding a Construction Line

A construction line defines the axis around which the cyclone profile revolves.

Add a construction line to the profile sketch:

1. Click (**Create line**) and create a line that is aligned with the Y-axis as shown below.



2. Press <Esc> to exit the line creation tool.
3. Right-click the line that you just added and select **Set As Construction**.



The appearance of the line changes from solid to dashed to indicate that it is now a construction line. Upon exiting the sketch, this line is no longer visible, as construction lines are only shown when you are creating or editing a sketch.

Apply a constraint between the construction line and the vertical line at the center of the cyclone to make sure that they are collinear:

4. To select the construction line, hold down the <`Ctrl`> key and select the long vertical line on the Y-axis.
5. Right-click either of the lines and select **Apply Collinear Constraint**.



The cyclone profile is now complete and is ready to revolve.

6. Click **OK** at the bottom of the *Sketch* panel to exit sketch mode.

Revolving the Cyclone Profile

Right-clicking a sketch in the feature tree lets you perform various CAD operations, such as revolving a profile around an axis.

Revolve the sketch of the cyclone chamber profile:

1. Rename the **Sketch 1** node to `Sketch Profile`.

2. Right-click **Sketch Profile** and select **Create Revolve**.



The *Revolve* panel lets you specify the parameters for the revolved feature. The *Direction* option defines the direction in which the sketch is revolved. The *Body Interaction* option specifies how the body that is created from the revolution interacts with any other bodies in the model. In this case, the sketch is revolved through 360 degrees and no other bodies are present. Therefore the *Direction* and *Body Interaction* settings do not affect the outcome and can be left as the default. The *Axis* option lets you choose which construction line to use as the axis of revolution. In this case, the sketch only contains one construction line, so this menu contains a single option, `Line 8`. A blue preview of the selected axis is shown in the *3D-CAD View* scene.

3. In the *Revolve* panel, click set *Angle* to `360.0 deg` and press `<Enter>`.



A preview of the feature is shown in the *3D-CAD View* scene.



4. Click **OK** to confirm the settings and create the revolved feature.

   A **Revolve 1** feature node is added to the feature tree, and the solid body is displayed in the *Graphics* window. Bodies that form a 3D-CAD model are shown under the **Bodies** node.



5. Expand the **Bodies** node and select **Body 1**.

   The corresponding body is highlighted in the *3D-CAD View 1* scene.
6. Rename the **Body 1** node to `Fluid`.

Creating the Outlet Pipe

The outlet pipe has a circular cross-section and is created by sketching a circle and extruding it. The circle that is sketched is slightly smaller in diameter than the hole in the top of the cyclone, and the difference between the two is the thickness of the outlet pipe.

Create the sketch for the outlet pipe on an existing face:

1. Rotate the model in the *3D-CAD View 1* scene so that you can see the sunken face in the top of the model.
2. Right-click the sunken face and select **Create Sketch**.



A sketch plane that is defined by the selected face is displayed on the model in the *3D-CAD View* scene.

3. Click (**View normal to sketch plane**).
4. To see the hole, in the *Vis* toolbar, click (**Make scene transparent**).
5. Click (**Create circle**) in the *Create Sketch Entities* box and click at `[0.0, 0.0] m`.

   The first click positions the circle at the origin of the sketch. The second mouse click lets you specify the radius of the circle.
6. Move the mouse so that the radius is approximately `0.06 m`, using the *Circle Properties* box for reference, and click.



You can define the circle radius more accurately by entering an exact value in the *Circle Properties* box.

7. In the *3D-CAD View 1* scene, click on the circle.
8. Set *Radius* to `0.07 m`



9. Click **OK** to exit the sketch.

10. In the *Vis* toolbar, click (**Make Scene Transparent**).

Extrude the sketch to form the outlet section of the cyclone geometry:

11. Right-click the **Sketch 1** node in the feature tree and select **Create Extrude**.



12. In the *Extrude* panel, make sure that *Body Interaction* is set to the default option of **Merge**.

    This option causes it to merge with the **Fluid** body, resulting in a single body.

13. Set *Distance* to `0.45 m`.



14. Retain the default settings for the other options.
15. Use the mouse to rotate the model and check the preview of the extrusion.
16. Click **OK**.

The outlet pipe is added to the cyclone chamber:



Creating the Inlet Duct

Fluid enters the top of the chamber horizontally through a rectangular inlet duct.

Create the inlet duct:

1. Right-click the **Features > YZ** node and select **Create Sketch**.
2. Click (**View normal to sketch plane**).
3. Click (**Make scene transparent**).
4. Zoom in on the top left-hand corner of the geometry.



5. Click (**Create rectangle**) and click the top left corner of the cyclone.



6. Move the mouse pointer down and to the right until the rectangle *Length* that is shown in the properties box is `0.05 m` and the *Width* is `0.15 m`.



7. Left-click to fix the point in position.
8. Click **OK**.

   Constraints are applied automatically to a rectangle to make sure that the sides are vertical and horizontal.

9. Click (**Make Scene Transparent**).

Extrude the sketch:

10. Right-click the **Sketch 2** node and select **Create Extrude**.
11. In the *Extrude* panel, set the extrusion *Distance* to `0.35 m` and click **OK**.

The inlet duct is added to the cyclone separator:



Creating a Design Parameter

Design parameters allow you to modify an aspect of the geometry from outside 3D-CAD.

Create a design parameter to control the position of the outlet pipe:

1. Right-click the **Sketch Profile** node and select **Edit**.
2. Click (**View normal to sketch plane**).
3. Zoom in on the top of the sketch.
4. Select the vertical line that defines the hole at the top of the cyclone chamber. Right-click the dimension arrow for this line and select **Edit Dimension**.



5. In the *Dimension* dialog, activate the *Expose Parameter* option.
6. In the *Name* textbox enter `PipeDepth` and click **OK**.



To make sure that the changes that are made to the length of this line only affect the depth of the hole into the cyclone chamber, fix the total height of the cyclone using a vertical dimension. Also fix the position of a point at the top of the cyclone.

7. Right-click the point at the top of the **PipeDepth** line and select **Apply Fixation Constraint**.



8. Zoom out so that the whole sketch is visible.
9. Select the point at the top of the **PipeDepth** line, hold down the <`Ctrl`> key and select the point at the bottom-right corner of the sketch.
10. Right-click one of the selected points and select **Apply Vertical Distance**.



11. In the *Dimension* dialog, make sure that the value is `0.9 m` and click **OK**.

    The final sketch appears as follows:



12. Click **OK** to exit the sketch.

The design parameter is displayed under the **Design Parameters** node in the object tree:

13. Expand the **Design Parameters** node and select the **PipeDepth** node.



The value of this design parameter is shown. Editing this quantity has a direct effect on the dimensions of the model.

Specifying Inlet and Outlet Faces

The final step in preparing the model geometry is to specify the inlet and outlet faces of the model by setting face names.

When the 3D-CAD model is imported into the simulation via geometry parts, faces that have been named are defined as separate part surfaces.

Therefore, when the geometry is assigned to a region, these surfaces can easily become separate boundaries.

Name the faces for the inlet and outlet:

1. Rotate the model so that the inlet duct is visible.
2. Right-click the rectangular inlet face and select **Rename**.



3. In the *Rename* dialog, enter `Inlet` and click **OK**.
4. Right-click the circular face at the top of the outlet pipe and select **Rename**.



5. In the *Rename* dialog, enter `Outlet` and click **OK**.
   The geometry is now complete and you can exit 3D-CAD.
6. At the bottom of the object tree, click **Close 3D-CAD**.



When you exit 3D-CAD, the *3D-CAD View 1* scene closes automatically.

7. Save the simulation.

Creating a Geometry Part

Create a geometry part from the 3D-CAD model.

To use the 3D-CAD model in a simulation, create a geometry part:

1. Right-click the **Geometry > 3D-CAD Models > Cyclone** node and select **New Geometry Part**.



2. In the *Parts Creation Options* dialog, click **OK** to close the dialog. The default settings are acceptable.



3. Expand the **Parts** > **Fluid** > **Surfaces** node.

   The inlet and outlet faces that you specified in the previous section are defined as separate surfaces.

Assigning a Part to a Region

To define the computational domain, assign the cyclone separator part to a region.

1. Right-click the **Geometry** > **Parts** > **Fluid** node and select **Assign Parts to Regions**.
2. In the *Assign Parts to Regions* dialog, select:
   a) **Create a Region for Each Part**
   b) **Create a Boundary for Each Part Surface**
3. Click **Apply** then **Close**.

   A new region is created with boundaries that define the inlet, outlet, and main fluid volume.

Check that the region and boundaries are defined correctly:

4. Create a Geometry Scene.
5. Expand the **Regions** > **Fluid** > **Boundaries** node and select each of the boundary nodes to make sure that they are specified correctly.

Selecting Physics Models

Physics models define the physical variables and phenomena in the simulation. In this tutorial, the K-Omega turbulence model is used to limit the simulation run time.

To select the physics models:

1. Create a physics continuum.
2. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Omega Turbulence**

Group Box	Model
Transition	**Turbulence Suppression**

3. Click **Close**.
4. To review the models, open the **Physics 1** > **Models** node.



5. Save the simulation.

Contents

Specifying Boundary Conditions

Specifying Boundary Conditions

Specify conditions at the inlet and outlet boundaries:

1. Edit the **Regions** > **Fluid** > **Boundaries** node.
2. Set the following properties:

Node	Property	Setting
Inlet	*Type*	**Velocity Inlet**
Physics Conditions > Turbulence Specification	*Method*	**Intensity + Length Scale**
Physics Values > Velocity Magnitude > Constant	*Value*	`5 m/s`
Outlet	*Type*	**Flow-Split Outlet**

Generating a Mesh

Use a polyhedral mesh to analyze the flow patterns in the cyclone separator.

1. Create a mesh continuum.
2. For the mesh continuum, **Continua > Mesh 1**, select the following models in order:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**

3. Click **Close**.

Define the mesh settings and generate the volume mesh:

4. Select the **Mesh 1 > Reference Values > Base Size** node and set *Value* to `0.01 m`.
5. Generate the volume mesh.
6. Create a mesh scene.



7. Save the simulation.

Preparing a Scalar Scene

Use a scalar scene to visualize the vertical velocity on a section through the cyclone chamber.

You can use this scene to visualize the solution while the simulation is running.

1. Create a scalar scene.
2. Create a section plane with the following properties:

Property	Setting
Input Parts	**Regions > Fluid**
origin	`[0, 0, 0] m`
normal	`[0, 1, 0] m`
Display	**Existing Displayer > Scalar 1**

3. Rotate the view in *Scalar Scene 1* so that the plane section is visible.



Define the scene settings:

4. Click **scene/plot**.
5. Edit the **Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Blended**
Scalar Field	*Function*	**Total Pressure > Lab Reference Frame**

6. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

1. Click (**Run**) in the *Solution* toolbar.
2. While the simulation is running, select the *Scalar Scene 1* tab at the top of the *Graphics* window to visualize the solution.
3. When the simulation has finished running, save it.

The scalar scene appears as follows:



Creating a Streamlines Scene

Create a scene to display the streamlines inside the cyclone separator.

1. Create a Geometry Scene.
2. Rename the **Geometry Scene 2** node to `Streamlines`.

Use the *Create Streamline* panel to specify the properties of the streamlines. In this case, the streamlines are grown from the cyclone inlet.

3. Create a Streamline derived part with the following properties:

Property	Setting
Input Parts	**Regions > Fluid**
Seed Parts	**Regions > Fluid > Boundaries > Inlet**
Part U-Resolution	2

Property	Setting
Part V-Resolution	8

Define the scene settings:

4. Click **scene/plot**.
5. Edit the **Streamlines > Displayers** node and set the following properties:

Node	Property	Setting
Displayers		
Geometry 1	*Opacity*	0.3
Streamline Stream 1	*Mode*	**Ribbons**
Scalar Field	*Function*	**Velocity > Lab Reference Frame > Laboratory > k**

To extend the streamlines through the fluid domain up to the outlet, increase the maximum propagation property. This defines how far the streamlines are propagated through the fluid domain from the starting points. A higher value causes the streamlines to extend further.

6. Click **Simulation**.
7. Select the **Derived Parts > streamline > 2nd Order Integrator** node and set *Maximum Propagation* to 15.

The streamlines appear as follows:



Contents

Animating the Streamlines

Animating the Streamlines

Animate the streamlines inside the cyclone chamber.

To animate streamlines:

1. Click **scene/plot**.
2. Edit the **Displayers** node and set he following properties:

Node	Property	Setting
Displayers		
Streamline Stream 1 > Animations	*Animation Mode*	**Tracers**
Streamline Settings	*Delay between tracers (sec)*	6

Node	Property	Setting
	Head size	`0.01`
	Tail Length (sec)	`1`

3. In the *Animation* toolbar:

   a) Click (**Play**)

   b) To stop the animation, click (**Stop**)

Modifying the Geometry and Rerunning

Modify the position of the outlet pipe using the design parameter, **PipeDepth**, that you created earlier.

1. Expand the **Geometry** > **3D-CAD Model** > **Cyclone** node.



2. Select the **Design Parameters** > **Pipe Depth** node and set **Value** to `0.35 m`.

   The **Cyclone** node shows an **Out-of-Date** feature state icon to indicate that the design parameter has changed and that the model requires updating.

3. Right-click the **Cyclone** node and select **Update**.



The outlet pipe moves further down inside the chamber, and less of the pipe is visible from the top.

To display the changes to the geometry, update the geometry part:

4. Display **Geometry Scene 1** in the *Graphics* window.
5. Right-click the **Geometry** > **Parts** > **Fluid** node and select **Update**.



The changes are applied to the geometry in the scene. These changes are not yet applied to the volume mesh.

Update the volume mesh:

6. Select the **Mesh Scene 1** tab in the *Graphics* window.
7. Regenerate the volume mesh.



The new volume mesh shows the outlet pipe positioned lower in the cyclone chamber.

The solution from the original geometry is automatically mapped onto the new mesh, decreasing the simulation run time for the second run. As fewer iterations are required to solve the case, alter the stopping criteria:

8. Select the **Stopping Criteria** > **Maximum Steps** node and set **Maximum Steps** to `1400`.
9. Close the *Streamlines Scene*.

   Leaving the *Streamlines Scene* open increases the run time because streamlines are calculated after each iteration.
10. Save the simulation as `cycloneModified.sim`.
11. Run the Simulation.
12. When the simulation has finished running, save it.
13. Open the Streamlines Scene and animate the scene as described previously.

The streamlines for the modified geometry appear as follows:



Summary

This tutorial has introduced the following STAR-CCM+ features:

- Using basic modelling techniques in 3D-CAD: Sketching, Revolving, and Extruding
- Exposing a design parameter
- Specifying inlet and outlet faces within 3D-CAD
- Creating geometry parts from the 3D-CAD model
- Assigning the 3D-CAD part to a region
- Modifying the 3D-CAD model using a design parameter and rerunning the case

Geometry Parts: CAD Assembly

This tutorial demonstrates how to use the geometry parts feature in STAR-CCM+ with an imported CAD assembly.

STAR-CCM+ also offers a number of tools that make it easier to prepare a simulation that is based on geometry which contains multiple parts. Several of these tools are demonstrated in this tutorial, such as: reorganizing the assembly structure, splitting and combining parts, and repairing part surfaces. The process of assigning geometry parts to regions is also covered, as well as the use of tags, filters, and meta data.

The geometry that is used in this tutorial is based on a device that is used to extract debris from high-pressure gas flows. The device operates by spinning the oncoming fluid around a central chamber and using centrifugal force to pull sand or other debris from the fluid. The clean flow media exits through a central pipe while the debris is collected from an outlet at the base. The assembly and its constituent parts were prepared using a standard CAD package, and are provided for this tutorial in Parasolid Transmit format.



Note: the geometry is based on an original design from Wood Group Pressure Control Ltd (WGPC), and is used with their permission.

A computational mesh is generated for the component in three stages. To begin, a closed internal surface mesh is generated using the surface wrapper. This wrapped surface is then retriangulated using the surface remesher. Finally a volume mesh is generated using the trimmer and prism layer meshers, and the extruder is used to extend the inlet and outlet of the device. To capture the complex and intricate features of the geometry the mesh generation process utilizes contact prevention conditions, volumetric controls and the wrapper scale factor.

Contents

Prerequisites
Importing the Surface Mesh
Reorganizing the Assembly
Retessellating Parts
Repairing Part Surfaces
Splitting and Combining Part Surfaces
Creating a Part from Part Surfaces
Creating a Tag
Applying a Tag to a Part
Setting Part Metadata
Setting Up Filters
Assign Parts to Regions
Assigning Part Surfaces to Boundaries
Creating the Mesh Continuum
Setting Global Mesh Properties
Creating Volume Shapes
Creating a Volumetric Control
Setting Wrapper Volume of Interest
Specifying a Contact Prevention Condition
Setting Up Extrusions
Deleting Existing Feature Curves
Initializing the Mesh
Generating the Surface Mesh
Generating the Volume Mesh
Acknowledgements
Summary

Prerequisites

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Surface Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/geometry/data`.
5. Select `sandcatcher.x_b` then click **Open**.
6. In the *Import Surface Options* dialog, click **OK** to accept the default options.
7. Save the simulation as `sandcatcher.sim`.

Reorganizing the Assembly

Reorganize the components in the assembly.

One of the benefits of using geometry parts is that you can reorganize the assembly nodes to suit the needs of the analysis. In this case, the assembly has a number of nuts and studs which are superficial and do not affect the analysis of the flow through the part. Group these nuts and studs into a composite part to simplify the object tree.

1. Select the **Nut(1_12)**, **Nut(1_25)**, **Stud(1.12 X 05.25)**, and **Stud(1.25 X 06.00)** nodes.
2. Right-click the **Nut(1_12)** node and select **Composite**.



A composite part that contains the nuts and studs is created.

3. Expand the **Geometry** > **Parts** > **Assembly 1** > **Composite** node.



Note: To reverse this process, you can explode the composite part.

Give the assembly and composite parts more suitable names:

4. Rename the **Assembly 1** node to `Sandcatcher`.
5. Rename the **Composite** node to `Nuts/Studs`.

Retessellating Parts

It is possible to change the surface triangulation and edge definition of a part. Use this capability to generate a fine tessellated surface on the **FlowCage** part.

First, display only the **FlowCage** in *Geometry Scene 1*:

1. Right-click the **Block** part, which is marked using an X in the screenshot below, and select **Hide**.



The selected part is hidden from view. To restore hidden parts at any time, right-click on the white area of the scene and select **Restore Hidden Parts...**.

Rotate the model so that the internal parts are clearly visible:

2. In the *Graphics* window, right-click the **FlowCage** and select **Show Only This Part**. The **FlowCage** is highlighted in the following screenshot.



3. In the *Vis* toolbar, click the (**Show All Meshes**) button.
4. Position the geometry as follows:



5. Right-click the **Geometry** > **Parts** > **Sandcatcher** > **FlowCage** node and select **Retessellate**.
6. In the *Tessellation Options* dialog, set the *Tessellation Density* to **Fine** and click **OK**.

   After a moment, the part is regenerated with a fine tessellated surface.
7. Click in the white region of the *Graphics* window to deselect the part.

   The scene appears as follows:



Restore the visibility of all parts:

8. Right-click a blank area in the geometry scene and select **Restore All Hidden Parts**.

Repairing Part Surfaces

Fill the holes in the geometry using the Surface Repair mode of STAR-CCM+.

There are three holes in the main block that must be filled: the inlet, outlet, and a hole for debris. Fill the holes by repairing the surface of the block:

1. Right-click the **Geometry** > **Parts** > **Block** node and select **Repair Surface**.



2. In the *Surface Repair Options* dialog, make sure that *Geometry Parts* is set to `[Sandcatcher.Block]`.



3. Click **OK**.
   The following diagram shows the location of the inlet, outlet, and debris holes in the block part. Fill each of the holes in turn.



4. Zoom in on the inlet hole so that the edges are clearly visible.
   Fill the holes by selecting the continuous edge around each hole and creating a patch. First, alter the Selection Controls to select edges:
5. Locate the *Selection Control* box in the *edit* panel. Clear the **Faces** option and activate the **Edges** option.



Edges can be selected by clicking them in the *Graphics* window. To select continuous edges, you can double-click a single edge in the desired range. Use this technique to select the continuous edge around the hole.

6. Double-click one of the edges around the rim of the hole.



The selected edges are marked with a pink line.

Note: To clear any select items at any time, click (**Select None**).

7. In the *Surface Repair* box, click (**Fill holes using selected edges**).



The hole is filled with a tessellated surface.



Note: While using the Surface Repair tool, it is possible to undo previous actions by pressing <`Ctrl+Z`> or by clicking the (**Undo**) button at the bottom of the *edit* panel.

The inlet hole has now been filled. Use the same technique to fill the outlet and debris holes.

8. Click the (**Save-Restore-Select views**) button in the *Vis* toolbar and select **Views** > **-Z** > **Up +Y**.



9. Zoom in on the hole and, using the same technique as the previous step, double-click one of the edges around the hole.



10. Click the (**Fill holes using selected edges**) button in the *Surface Repair* box to fill the hole.
11. Click **Close** in the *edit* panel.

Splitting and Combining Part Surfaces

You can manipulate part surfaces in a similar way to manipulating boundaries. Split and combine various part surfaces to define the inlet and outlet faces in the block part:

1. In the *Vis* toolbar, click (**Hide All Meshes**).
2. In the *Vis* toolbar, click (**Save-Restore-Select views**) and select **Views > +X +Y -Z > Up +Y**.
3. Select the **Sandcatcher > Block > Surfaces > Faces** node and click (**Reset View**) in the *Vis* toolbar.
4. Right-click the **Faces** node and select **Split By Patch**.
5. In the *edit* panel, hold down the `<Ctrl>` key and select patches **5771** and **5777** from the *PartSurface Patches* list.
   The selected patches are highlighted on the model in the *Graphics* window. Notice that it is not just the hole faces that have been highlighted, but some of the faces next to the holes. When the holes are filled, each face is automatically joined onto an existing local patch. Further manipulation of the faces is required to split them apart and make sure that the inlet and outlet are defined using separate part surfaces.



6. Click **Create** and then **Close**.
   A surface, **Faces 2**, is created from the patches that you selected. To display the surface in the *Graphics* window:
7. Select **Surfaces** > **Faces 2**.
8. In the *Vis* toolbar, click (**Reset View**) and use the mouse buttons to position the part as shown in the following screenshot.



9. Right-click the **Faces 2** node and select **Split By Part Curves**.



10. In the *Split Part Surfaces by Part Curves* dialog, set the following:
    a) *Part Surfaces* to **Faces 2**.
    b) *Part Curves* to **Edges**.

    Manipulate the created surfaces to achieve a single surface for the main body of the block, and separate surfaces for the inlet and outlet.
11. Holding down the **<Ctrl>** key, select the **Faces**, **Faces 2**, and **Faces 2 3** nodes. Right-click the **Faces** node and select **Combine**.

Note: The surfaces are combined under the name of the first node that you select, which in this case, is **Faces** .

12. Rename the **Faces 2 2** node to `Inlet`.
13. Rename the **Faces 2 4** node to `Outlet`.

Creating a Part from Part Surfaces

Create a part from the part surfaces.

Before you create the parts, you must first split the existing part it into separate parts. You then create parts from the separated surfaces. Use the following technique to split the **FlowDiverter** into two parts for the benefit of the meshing process.

1. Right-click the **FlowDiverter** > **Surfaces** > **Faces** node and select **Split By Patch**.
2. In the *Graphics* window, zoom in on the highlighted part. In the *edit* panel, hold down the **<Ctrl>** key and select **5124**, **5125**, **5130**, and **5131** from the *PartSurface Patches* list. These patches correspond to the funnel, including the tip.



3. Click **Create** then **Close**.
4. Right-click the **Faces 2** node and select **Create New Part From Surfaces**.



5. Rename the **Sandcatcher** > **Part** node to `FlowDiverterFunnel`.
6. Rename the **FlowDiverterFunnel** > **Surfaces** > **Faces 2** node to `Faces`.
7. Save the simulation.

Creating a Tag

Create a tag that is applied to all parts that are in contact with the fluid.

Tags allow you to have greater control over the selection of parts when setting up a simulation. They allow you to filter "tagged" parts in selection dialogs, a feature that is useful when working with models that have many parts.

To create a tag:

1. Right-click the **Geometry** > **Tags** node and select **New Tag**.
2. Rename the **Tags** > **Tag** node to **FluidContact**.
3. Right-click the **FluidContact** node and select **Set Custom Icon**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/geometry/data`.
5. Select the `tagIcon.png` file and click **Open**.

The icon for the **FluidContact** tag has now changed.



Any number of tags can be created in a simulation, and it is useful to set custom icons for each tag so that they are easily distinguishable.

Applying a Tag to a Part

Tag the relevant parts with the **FluidContact** tag using three different techniques.

1. Expand the **Geometry** > **Parts** > **Sandcatcher** node.
2. Multi-select the **FlowCage**, **FlowDiverter**, and **FlowDiverterFunnel** nodes.
3. Right-click **FlowCage** and select **Apply Tag** > **FluidContact**.



A small **FluidContact** icon has appeared next to each of the tagged parts to indicate that these parts are in contact with the fluid. The **FlowDiverter** is not in contact with the fluid, so remove the tag.

4. Select the **FlowDiverter** node and click **[FluidContact]** to the right of the *Tags* property.
5. In the *FlowDiverter - Tags* dialog, deselect **FluidContact**.
6. Click **OK**.

Use another method to apply the tag to the pipe choke part.

7. Expand the **Sandcatcher** and **Tags** nodes. Drag-and-drop the **FluidContact** node onto the **PipeChoke** node.
   The simulation tree appears as shown below.



Apply the tag to the block part:

8. Select the **Block** node and set *Tags* to **FluidContact**.
   The simulation tree appears as follows:



9. Save the simulation.

Setting Part Metadata

Specify metadata for the location of some of the parts in the model by using the name 'position' and the values 'internal' and 'top'.

Metadata is useful for specifying the properties of parts, and can be used in a similar way to tags. To apply metadata you provide a suitable name, for example 'material' and a suitable value corresponding to this name for the relevant parts, for example 'steel' or 'Al-alloy'.

To apply metadata:

1. Multi-select the **FlowCage**, **FlowDiverter**, **FlowDiverterFunnel**, and **PipeChoke** nodes.
   The common properties of the selected objects are displayed.
2. Click **{}** to the right of the *Metadata* property and type `{position:internal}`.



3. Press <`Enter`> to confirm the value.
4. Select the **HoistRing(3052T91)** node and set *Metadata* to `{position:top}`.



Setting Up Filters

Create two filters.

Filters are based on tags and metadata and they allow more rapid selection of related parts when setting up a simulation.

Create two filters to allow you to select:

- Any parts that are in contact with the fluid
- Only internal parts that are in contact with the fluid

To create the filters:

1. Right-click the **Tools** > **Filters** node and select **New**.



2. In the *Filter* dialog, select **Tag** from the drop-down menu.
3. From the new drop-down menu that appears, select **contains**.
4. Click (**Custom Editor**) .
5. In the *Select Object* dialog, select the **FluidContact** tag.
6. Click **OK** to close the *Select Object* dialog.
7. Click on a blank space in the *Filter* dialog.
   The filter is updated.



8. Click **OK**.
   The filter is added to the **Filters** node.
9. Rename the new **Filter** node to `FluidContact`.
   The first filter is now complete.

To create the second filter, you can make a copy of the **FluidContact** tag filter, which shows only those parts that are in contact with the fluid, and add a metadata filter to show only the internal parts:

10. Make a copy of the **FluidContact** filter and rename it to `InternalContact`.
11. Right-click the **InternalContact** filter and select **Edit...**.
12. In the *InternalContact* dialog, click (**Add a new item**).
    A new entry is added to the filter.
13. From the new drop-down menu, select **Metadata**.
14. From the second drop-down menu, select **Contains**.
15. In the first text field, enter `position`.

16. In the second text field, enter `internal`.



17. Click on a blank space in the *Filter* dialog.
    The filter is updated.



18. Click **OK**.
19. Save the simulation.

Assign Parts to Regions

Assign parts to a region using two methods.

In the first method, create a region and assign a number of parts to that region directly through the **Geometry** > **Parts** node.

1. Expand the **Geometry** > **Parts** > **Sandcatcher** node.

2. Multi-select the **Block**, **FlowCage**, **FlowDiverter**, and **FlowDiverterFunnel** nodes. Right-click the **Block** node and select **Assign Parts To Regions**.
3. In the Assign Parts To Regions dialog:
   a) Select **Create a Boundary for Each Part Surface**.
   b) Leave all other options unchanged.



   c) Click **Apply** then **Close**.
4. In the *Vis* toolbar, click (**Reset View**).
5. Clear the selected parts by clicking anywhere in the white region of the Graphics window.
6. Right-click anywhere in the white region of the *Graphics* window and select **Restore All Hidden Parts**.
7. Expand the **Regions** > **Region** > **Boundaries** node.

A region was created and given the default name, **Region**, and the surfaces that are associated with the parts that were selected have been assigned to boundaries. Again, the part names have been preserved.

8. Click and drag the **Region** node onto the blank area in *Geometry Scene 1*, and select **Part Displayers > Add to Geometry 1**.



Note: When a part is assigned to a region, the color changes from light gray to dark gray.

9. Click a blank area in *Geometry Scene 1* to clear all selections in the scene.

10. In the *Vis* toolbar, click (**Save-Restore-Select views**) and select **View > +X+Y+Z > Up +Y**.

11. Use the mouse buttons to position the part as shown in the following screenshot:



Now use another method to assign an additional part to the region using one of the filters that was created previously.

12. Select the **Region** node and click the ... (**Custom Editor**) button to the right of the **Parts** property.

    The *Region - Parts* dialog appears.



13. At the top left of the *Region - Parts* dialog, click the **Show/Hide Filters** button.

    The filters are shown in the *Filters* box to the left of the dialog.

14. Select the **FluidContact** filter.

    The contents of the *Tree View* window shows only those parts that are in contact with the fluid.

15. Select **PipeChoke**.

    All the parts that are in contact with the fluid are now selected.



16. Click **OK**.

    The **PipeChoke** part has is now assigned to **Region**.

Assigning Part Surfaces to Boundaries

Assign the part surfaces to boundaries.

When the **PipeChoke** part is assigned to **Region**, the associated surfaces are assigned to the default boundary, which in this case is **Block**.

Create a boundary for the **PipeChoke** part surfaces.

1. Select the **Region** > **Boundaries** > **Block** node and note the value of the *Part Surfaces* property. Resize the window if necessary.

   The **PipeChoke** surface is currently associated with the **Block** boundary.



To create a separate boundary for the **PipeChoke** part surface:

Return to the **Geometry** > **Parts** simulation tree and use the **Set Boundary** option on the relevant part surface:

2. Expand the **Geometry** > **Parts** > **Sandcatcher** > **PipeChoke** > **Surfaces** node.
3. Right-click the **Faces** node and select **Set Boundary** > **New**.



In the **Regions** > **Region** node, a boundary, **Faces**, appears.

4. Rename the **Boundaries** > **Faces** node to `PipeChoke`.

   Notice that the value of the *Part Surfaces* property is **Sandcatcher.PipeChoke.Faces**.



5. Select the **Region** > **Boundaries** > **Block** node.

   The **PipeChoke** part surface is automatically removed from the **Block** boundary. The only part surface that is associated with the **Block** is **Sandcatcher.Block.Faces**.



6. Save the Simulation.

Creating the Mesh Continuum

Select the meshing models for the simulation.

To extract the internal surface of the component, use the surface wrapper model.

1. Create a mesh continuum.
2. For the mesh continuum, **Continua** > **Mesh 1**, select the following models in order:

Group Box	Model
Surface Mesh	**Surface Wrapper** **Surface Remesher**
Volume Mesh	**Trimmer**
Optional Models	**Prism Layer Mesher** **Extruder**

3. Click **Close**.

Setting Global Mesh Properties

Set the global mesh properties for the simulation.

Use the surface wrapper to generate a closed surface mesh for the internal fluid region. Use the surface remesher to retriangulate the surface and optimize it for generation of the volume mesh. Use a surface

wrap that is fine enough to resolve the complex features of the geometry, use a coarser remeshed surface to reduce the final mesh cell count. To reconcile these requirements, use the wrapper scale factor. The wrapper scale factor allows you to define global mesh properties for the remeshed surface, and to specify how much finer the wrapped surface is, by setting a percentage factor. The percentage factor determines how much smaller the base size for the wrapped surface is. In this case, the length of the wrapped surface is 20% of the global size that defines the remeshed surface.

Use the proximity refinement option to refine the mesh around faces that are close together. This setting is turned on by default for the Surface Remesher and Trimmer meshing models.

1. Select the **Mesh 1** > **Models** > **Surface Wrapper** node and activate **Do proximity refinement**.



2. Edit the **Mesh 1** > **References** node.
3. Set the following properties:

Node	Property	Setting
Base Size	**Value**	`15 mm`
Prism Layer Thickness > **Relative Size**	**Percentage of Base**	`20`
Surface Proximity	**# Points in Gap**	`20`
Wrapper Scale Factor	**Percentage**	`20`

Creating Volume Shapes

Refine the mesh in two areas of the geometry that contain narrow tubes. Specify these areas by creating volume shapes, and define the mesh size with a volumetric control.

1. Right-click the **Tools** > **Volume Shapes** node and select **New Shape** > **Cylinder**.

2. In the *Create Cylinder Volume Shape* dialog, set the following properties:

Property	Setting
Start Coordinate	`[0, 0, 0.1] m`
End Coordinate	`[0, 0, 0.14] m`
Radius	`0.1 m`

3. Click **Create**.

   Create a second cylinder with different properties in the same *Create Cylinder Volume Shape* panel:

4. In the *Create Cylinder Volume Shape* dialog, set the following properties:

Property	Setting
Start Coordinate	`[0, 0, -0.17] m`
End Coordinate	`[0, 0, -0.35] m`
Radius	`0.05 m`

5. Click **Create** then **Close**.

   To show the position of the two cylinders in the *Graphics* window:

6. Expand the **Volume Shapes** node and multi-select the **Cylinder 1** and **Cylinder 2** nodes.

7. In the *Vis* toolbar, click (**Make Scene Transparent**).



8. In the *Vis* toolbar, click (**Restore Scene Transparency**) button.

Creating a Volumetric Control

Create a volumetric control. The mesh refinement within each cylindrical volume is the same, so only one volumetric control is needed.

1. Expand the **Continua** > **Mesh 1** node.
2. Right-click the **Volumetric Controls** node and select **New**.
3. Select the **Volumetric Control 1** node and click the (**Custom Editor**) button to the right of the **Shapes** property.
4. In the *Volumetric Control 1* dialog, click the **Column Selection** button so that both volume shapes are selected.
5. Click **OK**.

   By assigning the volume shapes to the volumetric control, the **Volumetric Control 1** node is activated.

Note: The volumetric control is not applied to the Surface Wrapper as the wrapped surface mesh has a sufficient level of refinement due to the wrapper scale factor.

Specify the mesh properties within the cylinders:

6. Select the **Volumetric Control 1** > **Mesh Conditions** > **Surface Remesher** node and activate **Customize surface remesher**.



7. Select the **Trimmer** node and activate **Customize isotropic size**.



Define the relative size for both the remeshed surface and trimmed volume.

8. Select the **Mesh Values** > **Custom Size** > **Relative Size** node and set **Percentage of Base** to `15`.
9. Save the simulation.

Setting Wrapper Volume of Interest

The wrapper volume of interest is specified as a Mesh Condition for the Region that is wrapped, which in this case is **Region**.

As the fluid region is the third largest volume in the sandcatcher component, use the Nth largest option to specify the volume of interest.

1. Select the **Regions** > **Region** > **Mesh Conditions** > **Volume of Interest Specification** node and set **Method** to **Nth Largest**.



2. Select the **Mesh Values** > **Volume Number** node and set **Volume Number** to 3.

Specifying a Contact Prevention Condition

Specify a contact prevention condition to capture the geometry between the **FlowCage** and the **FlowDiverterFunnel**.

A contact prevention condition can be used to make sure that two or more surfaces that are close to one another are not joined when the surface is wrapped. This feature can be useful when there is an intricate feature in the geometry that would otherwise be lost, depending on the surface size settings provided.

Use a contact prevention condition to capture the geometry between the **FlowCage** and the **FlowDiverterFunnel** as shown in the diagram below.



When defining the contact prevention, provide a value for the minimum distance between the two surfaces at which contact prevention is guaranteed. This distance is known as the search floor. A small gap exists between the **FlowCage** and **FlowDiverterFunnel/FlowDiverter**. Select a value for the search floor so that the gap is filled, while capturing an adequate level of detail in the geometry.

1. Under the **Region > Mesh Values** node, right click **Contact Prevention** and select **New > One Group Contact Prevention Set**.



2. Select the **Contact Prevention 1** node and set **Search Floor** to `0.55` mm.
3. Click the (**Custom Editor**) button to the right of the **Boundaries** property.
4. In the *Contact Prevention 1 - Boundaries* dialog, select **FlowCage** and **FlowDiverterFunnel**.
5. Click **OK**.

Setting Up Extrusions

Generate an extruded mesh at the inlet and outlet to aid solution convergence.

The Extruder meshing model was selected previously, so now specify the mesh conditions and values at the relevant boundaries.

1. Set the **Block.Inlet** and **Block.Outlet** boundary types to the following:

Boundary	Type
Block.Inlet	**Velocity Inlet**
Block.Outlet	**Pressure Outlet**

2. Expand the **Regions** > **Region** > **Boundaries** node.
3. Select the **Block.Inlet** node and set the following properties:

Node	Property	Setting
Mesh Conditions		
Extrusion Type	*Extrusion Type*	**Constant rate normal**
Mesh Values		

Node	Property	Setting
Normal Extrusion Parameters	*Number of Layers*	`20`
	Stretching	`1.5`
	Magnitude	`0.16 m`
	Specify a New Region	**none**

4. Select the **Block.Outlet** node and set the following properties:

Node	Property	Setting
Mesh Conditions		
Extrusion Type	*Extrusion Type*	**Constant rate normal**
Mesh Values		
Normal Extrusion Parameters	Number of Layers	`80`
	Stretching	`4.0`
	Magnitude	`0.9 m`
	Specify a New Region	**none**

Deleting Existing Feature Curves

It is good practice to delete any existing feature curves before using the surface wrapper to avoid them being imprinted on the wrapped surface. The surface wrapper generates a set of feature-curves that are based on the sharp edge angle that was defined in the mesh continuum.

Delete the feature curves:

1. Expand the **Region** > **Feature Curves** node, right-click **Feature Curve** and select **Delete**.
2. In the *Confirm Object Deletion* dialog, click **Yes**.
3. Save the simulation.

Initializing the Mesh

When working with geometry parts the mesh it is necessary to initialize the mesh before surface or volume meshing to provide the starting surface for the meshing models using the current surface and curve assignments.

In this case, a number of parts have been assigned to **Region**, which is associated with a mesh continuum, so the initial surface is comprised of only these parts.

1. In the *Mesh Generation* toolbar, click the (**Initialize Meshing**).

Note: There are some warning messages in the *Output* window.



The status message **Initializing Meshing: Finished** is displayed in the status bar when mesh initialization is complete. View the initial surface representation.

2. In the *Vis* toolbar, click (**Show All Meshes**).
3. Expand the **Scenes** node. Right-click the **Geometry Scene 1** node and select **Apply Representation** > **Initial Surface**.

The initial surface appears as follows:



Generating the Surface Mesh

The mesh parameters are fully defined and you can proceed to generate the surface mesh:

1. Generate the surface mesh.

   The output from the mesher is shown in the *Output* window and the ongoing status is shown in the status bar. The meshing process takes some time and is finished when the message beginning with `Surface Meshing Pipeline Completed` appears in the *Output* window.

View each of the wrapped surface mesh and remeshed surfaces in the active scene and compare the level of mesh refinement in each:

2. Right-click in the white region of the active scene and select **Apply Representation > Wrapped Surface**.

   The wrapped surface is displayed in the *Graphics* window.
3. Use the mouse buttons to zoom, pan and rotate the view to examine the mesh.



To view the remeshed surface:

4. Right-click in the white region of the active scene and select **Apply Representation > Remeshed Surface**.



5. Save the simulation.

Generating the Volume Mesh

Generate the trimmed volume mesh, including extrusions at the inlet and outlet, and prism layers at wall boundaries.

1. Generate the volume mesh.

   The output from the meshers is shown in the *Output* window and the ongoing status is shown in the status bar. The meshing process takes some time and is finished when the message beginning with `Volume Meshing Pipeline Completed` appears in the *Output* window.

Apply the volume mesh representation to the active scene. Use an alternative method to what you used for the surface mesh representation:

2. Expand the **Representations** node.
3. Drag-and-drop the **Volume Mesh** node from the object tree into the active scene.



Update the selected parts in the geometry displayer for the current scene:

4. Click **scene/plot**.
5. Edit the **Displayers > Geometry 1 > Parts** node.
6. In the *Edit* dialog, select all the boundaries in the region.
7. Click **OK**.
8. In the *Vis* toolbar, click the (**Reset View**).
9. Use the mouse buttons to pan, zoom, and rotate the view to examine the volume mesh.



10. Save the simulation.

Acknowledgements

The assistance of Mr. Stephen McCulloch and Mr. Jim Brinkley of Wood Group Pressure Control Ltd. is gratefully acknowledged, as is the permission of WGPC Ltd to present a modified form of the geometry.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Importing the mesh and naming the simulation.
- Reorganizing the assembly.
- Retessellating parts.
- Repairing part surfaces.
- Splitting and combining part surfaces.
- Creating a part from part surfaces.
- Creating a tag and applying it to a part.
- Setting part meta data.

- Setting up filters.
- Assigning parts to regions.
- Assigning part surfaces to boundaries.
- Creating a mesh continuum and setting up the mesh.
- Creating volume shapes.
- Creating a volumetric control.
- Setting wrapper volume of interest.
- Specifying a contact prevention condition.
- Setting up extrusions.
- Deleting existing feature curves.
- Initializing the mesh.
- Generating and visualizing surface and volume meshes.

Mesh

The tutorials in this set illustrate various STAR-CCM+ features for building CFD meshes.

Contents

Surface Wrapper: Intake Manifold
Surface Repair Tools Brake Assembly
Parts-Based Meshing External Aerodynamics
Region Based Meshing Imported Control Valve
Region-Based Meshing Multi-Region Heat Exchanger
Directed Meshing: Electrical Machine

Surface Wrapper: Intake Manifold

This tutorial demonstrates how to use the surface wrapper to extract the internal fluid surface from a collection of parts that represent an intake manifold assembly.

The surface wrapper is run several times to present the general iterative nature of the tool when dealing with new or unfamiliar geometries. With each iteration, additional features and options are included to improve the overall fidelity of the resultant wrapped surface. With experience, you can utilize the surface wrapper and get the desired result in one or two iterations.

The geometry is shown below:



The manifold assembly is composed of several parts, including the main body, mounting struts, air-flow sensor, control valve, and directional flow baffles. Most parts are in their original CAD form and have not been prepared for use in a CFD analysis.

For a simple case with only a few parts, manual repair to remove unwanted surfaces and stitch together remaining ones to form a closed volume geometry, is a realistic solution. As the number of parts and overall complexity increases, however, manual repair becomes

inefficient. Instead, use the surface wrapper to extract a high-fidelity representation of the geometry. Use the representation in the surface remesher to obtain a high-quality triangulation from which the volume meshing tools can proceed.

The issues that the surface wrapper deals with in this example are:

- Material thickness
- Open flow boundaries
- Disconnected geometry/flanges
- Overlaps/double surfaces
- Intersecting bodies (struts)
- Surface holes
- Channels
- Internal thin wall baffles
- Close proximity bodies (valve, sensor)

These issues are illustrated in the cut-away diagram below:



The surface wrapper can be a fast and efficient process of extracting both internal and external surfaces from CAD type geometries. Successful use of the wrapper requires practice, due to the numerous paths available. For example, tweaking some of the basic input values to deal with a particular feature can sometime have a significant effect on the overall result. Repeated use of the wrapper however builds up a wealth of experience.

Contents

Prerequisites
Importing the Surface Mesh
Identifying Important Boundaries
Activating the Wrapper Model
Setting the Initial Reference Values
Capping the Geometry
Running the Leak Detector
Producing a Basic Wrap
Improving the Manifiold Body Fidelity and Curvature
Improving the Sensor Part Definition
Improving the Channels Definition
Improving the Flange Definition
Improving the Valve Definition
Improving the Baffle Definition
What to Do Next
Summary

Prerequisites

The instructions in the Surface Wrapper: Intake Manifold tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Technique	Tutorial
Parts-Based Meshing	Parts-Based Meshing: External Aerodynamics

Importing the Surface Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/mesh/data`.
5. Select `baffles.x_t`, `manifold_body.x_t`, `sensor.x_t`, `struts.x_t`, and `valve.x_t` then click **Open**.
6. In the *Import Surface Options* dialog, set the following properties:

Property	Setting
Import Mode	**Create New Part**
Tessellation Density	**Fine**

7. To accept the remaining default *Import Surface Options*, click **OK**.
8. Save the simulation as `Star1.sim`.

Identifying Important Boundaries

The import process created 11 new parts which you can combine and rename appropriately.

Merging parts allows you to rename several parts with one name. You can then rename the boundaries on the parts as necessary.

Combine and rename the parts:

1. Expand the **Geometry** > **Parts** node.
2. Multi-select **baffles1**, **baffles2**, and **baffles3**.
3. Right-click one of the selected parts and select **Combine**.
4. Rename the combined part **Baffles**.
5. Combine and rename the remaining parts as shown in the table below:

Parts	New Part Name
manifold_body1, valve2	**Cross pipe**
manifold_body2	**Manifold body**
manifold_body3	**Inlet pipe**
sensor	**Sensor**
struts1, struts2	**Struts**
valve1	**Valve**

Activating the Wrapper Model

Close or "cap off" any open boundaries or large holes in the surface so that the wrapper can determine what is inside the geometry and what is outside.

To assist with the hole detection process, activate the surface wrapper model and specify initial surface sizes so that the leak detector can be invoked to check for large holes. The leak detection tool also contains a subset of the surface repair tools, which you can use initially to close the gaps in the model. The wrapping process can close holes that are smaller than the specified target mesh size for each surface.

Activate the Surface Wrapper:

1. Right-click the **Geometry > Operations** node and select **New > Surface Wrapper**.
2. In the *Create Surface Wrapper Auto Mesh Operation* dialog, select all parts then click **OK**.

   Two **Surface Wrapper** nodes appear in their unexecuted state in the simulation tree, at the following locations:

   - **Geometry > Parts**

     This node represents the part that is produced by the Surface Wrapper operation.

   - **Geometry > Operations**

     This node represents the action of the Surface Wrapper and allows you to execute the Surface Wrapper operation.

Setting the Initial Reference Values

Values for the surface wrapper must be set with consideration of the parameters of the geometry.

To use the leak detector, specify some initial first guess reference values for the surface minimum and target sizes. For this tutorial, use relative size specification with a base size of 0.1 m, which is approximately the width of the main manifold body.

1. Navigate to **Geometry > Operations > Surface Wrapper > Default Controls**.
2. Set the following properties:

Node	Property	Setting
Base Size	*Value*	`0.1m`
Target Surface Size > Relative Size	*Percentage of Base*	`10.0`
Minimum Surface Size > Relative Size	*Percentage of Base*	`5.0`

You can influence the surface wrapper with the following default controls:

CAD Projection	Allows you to turn on or off the *Project to CAD* option. Activating this option allows vertices to be projected back to the imported CAD surface at the end of the wrapping process, resulting in a surface definition closer to the original. By default, the *Project to CAD* property is activated. Retain this setting for this tutorial.
Surface Curvature	Allows the specification of the number of triangles that are produced on curved surfaces. It contains a node, **Basic Curvature**, which sets the initial number of points on a circular edge. Increasing the # *Pts/circle* value, increases the resultant triangulation in curved areas of

	the geometry. For this tutorial, maintain the default value of `36`.
Feature Angle	Provides an angle setting for which feature curves are generated during the wrapping process. In general, pre-existing feature curves in the imported surface are ignored in preference to feature curves that are generated during the wrapping process. The only exception is when the feature curves are on a flat surface and define a surface detail of some kind (such as an imprint). Decreasing the *Angle* value increases the number of feature edges that the wrapping process generates. For this tutorial, maintain the default value of `30` degrees.

Capping the Geometry

As you are extracting the internal volume of the manifold, then a mostly watertight import surface must be supplied as a starting point. Achieve this surface by capping-off the open inlet and outlet boundary areas in the manifold.

Two different tools are available to help close off holes, namely the polygonal patch fill and the hole filler. The polygonal patch filler is a quick and easy way of closing arbitrary shaped holes which do not have a closed loop definition and/or are not planar. The process literally "patches" up the surface by creating faces that cover the hole area completely. The hole filler is a more exact method of filling well-defined holes by using a closed loop feature edge definition around the hole. Both methods are used in this example.

1. To open the *Wrapper Leak Detection* dialog and associated surface repair tools:
   a) Navigate to **Geometry** > **Operations**.
   b) Right-click the **Surface Wrapper** node and select **Run Leak Detection**.
2. In the *Wrapper Leak Detection dialog*, set **Mode** to **Repair Surface**.



3. Change the view so that the manifold inlet pipe part is as shown below:



The inlet boundary is capped off in three stages. Start with the rectangular section on the mid-plane:

4. In the *Surface Repair* group box, click (**Fill Polygonal Patch**).
5. Click once on each of the four opposite corners of the mid-plane of the inlet pipe material thickness surface, as shown below. To complete the patch fill operation, right-click when selecting the fourth corner point.



The surface patch is filled in as shown:



Note: If you make a mistake, use (**Undo**). Having an overlap in the patched surface is perfectly acceptable as the wrapper deals with this overlap automatically.

6. Repeat the process above for the cylindrical portion of the mid-plane and on the inlet plane half-cylinder.

   The patched surfaces are shown below:





For the remaining open boundaries, use the hole filler instead of the polygonal patch filler. The default settings of the wrapper leak detection tool allow you to select only faces. To select edges or vertices, activate the corresponding options:

7. In the *Selection Control* group box, activate **Edges** and deactivate **Faces**.
8. Rearrange the view to look at the open boundary of the cross pipe.
9. Double-click the feature edges around the hole to select them all in one go.
10. To fill the hole, in the *Surface Repair* group box, click (**Fill Holes using Selected Edges**).

    The result of the hole fill is shown below:



11. Repeat the above process and fill the three holes at the end of the outlet pipes of the main manifold body. To select all three feature curve loops, hold the `<Ctrl>` key during your selection.

    The result of the hole fill on the outlet pipes is shown below:



All the major holes in the geometry are now closed.

Running the Leak Detector

To highlight any remaining holes before the first wrap, run leak detection.

Leaks are not desired because they cause an inversion in the wrapping process, which results in the incorrect volumes being wrapped.

To perform a leak detection:

1. In the *Wrapper Leak Detection* dialog, set **Mode** to **Locate Points**.
2. In the **Locate Points** group box, set **Current Point** to **Source Point** and use the point tool on the display to position the source point (small red ball) in the middle of the manifold body (approximate coordinates are X=0, Y=0.12, and Z=-0.01).



3. Set **Current Point** to **Target Point 1**.
4. Use the point tool to position the target point (small red ball) to anywhere outside the manifold geometry (for example, approximately X=0, Y=0.12, Z=0.1).

   The source point is now marked with a small blue ball:



5. To run the leak detector, in the *Leak Detection Actions* group box, click **Recompute Template+Paths**.

   The process takes several seconds to compute the octree mesh and detect the presence of leaks. In this case, the message `No path found between source and target points. These points will be on different sides of the final wrapped surface` is shown in the *output* window, indicating that no leak is present.

   If a hole is detected, use the polygonal patch filler to patch the leak area. To rerun the leak detector, click **Recompute Paths**.
6. At the bottom of the edit dialog, click **Close**.

Producing a Basic Wrap

Before refining the surface wrapper, review the outcome from running the surface wrapper with the initial default settings.

By default, the surface wrapper attempts to extract the largest internal volume from the region definition. You can specify the volume of Interest using **Operations > Surface Wrapper > Default Controls > Volume of Interest**.



Run the surface wrapper based on the initial settings:

1. Right-click **Operations** > **Surface Wrapper** and select **Execute**.

   The Surface Wrapper nodes now appear in their executed state, without the exclamation mark.

   The output of the surface wrapper is shown in the *Output* window. When the wrapper is complete, the message `Surface Wrapper Completed` appears.

View the result:

2. Navigate to **Geometry Scene 1** > **Displayers** > **Geometry 1** > **Parts**.
3. Set **Parts** to **Surface Wrapper**.
4. Navigate to **Outline 1** > **Parts**.
5. Set **Parts** to **Surface Wrapper**.
6. Click **Show All Meshes**.



This initial wrapped surface is coarse and defeatured in many key areas, such as the valve and baffles.

However, the surface wrapper did achieve the following:

- Extracted the largest internal volume and eliminated the material thickness.
- Closed the small holes in the inlet pipe surface automatically.
- Defeatured the channels and flanges.
- Eliminated the overlap/double surface in the cross-pipe area.
- Intersected the struts with the main manifold body.
- Extracted the feature curves for the most important features.

You can easily improve on the first wrap result by tweaking the reference settings and including additional options to deal with feature specific issues. The remainder of this tutorial concentrates on improving each geometry feature in turn, identifying the issue and the solution to rectify the problem. The following items are covered:

- Improving the manifold body fidelity and curvature.
- Sensor definition.

- Capturing the channel details (if desired).
- Flange resolution.
- Valve area resolution in general.
- Baffle intersection and resolution.

Improving the Manifold Body Fidelity and Curvature

Reducing the Manifold Body Fidelity and Curvature values has a significant and immediate impact on the general model fidelity.

1. Navigate to **Operations > Surface Wrapper > Default Controls**.
2. Set the following properties:

Node	Property	Setting
Target Surface Size > Relative Size	**Percentage of Base**	`8.0` (8.0mm)
Minimum Surface Size > Relative Size	**Percentage of Base**	`2.5` (2.5mm)

Note: When changing the target size, it is recommended that you rerun the leak detector.

To re-run the Surface Wrapper:

3. Right-click **Operations > Surface Wrapper** and select **Execute.**.

   Due to the smaller surface sizes, the wrapper takes slightly longer to run this time. When the surface wrapper is complete, the display is updated with the new result:



On inspecting this new wrap, you can see the following improvements in the surface definition:

- Better curvature definition for the main manifold fillets.
- Better curvature for the struts.
- Better overall representation of the three curved outlet pipes.

Improving the Sensor Part Definition

The sensor part is joined to the bottom surface of the manifold body, which is not desirable.

The wrapped surface for the sensor area currently appears as shown:



The clearance between the sensor and the base is approximately 2 mm, which is less than the current target surface size of 8 mm.

To prevent the wrapper from joining the two boundaries, use a contact prevention.

1. Navigate to **Geometry** > **Operations** > **Surface Wrapper**.
2. Right-click **Contact Prevention** and select **New** > **One Group Contact Prevention Set**.

   A **Contact Prevention** > **Contact Prevention** node is created.
3. Select the **Contact Prevention** > **Contact Prevention** node and set the following properties:

Property	Setting
Surfaces	**Manifold Body** and **Sensor**
Minimum Size	`0.5 mm` (note the units)

The **Minimum Size** value indicates the triangle size in the space between the two parts, which determines the octree refinement in the area. The value is smaller than the gap size (typically one half or one quarter of the distance). In this example, four octree triangles are generated in the 2 mm gap between the sensor and the manifold base.

You can still join the surfaces even though a **Minimum Size** value is set. The **Minimum Size** value is simply a stopping size for the surface wrapper refinement, which results in the two surfaces being separated. The value also has to be greater than zero.

4. Rename **Contact Prevention** to `Sensor CP`.

   Since the sensor is also fairly small in comparison to the manifold, increase the triangulation in this area by specifying a localized target and minimum surface size for the sensor boundary.

5. Right-click **Surface Wrapper > Custom Controls** and select **New > Surface Control**.
6. Edit the **Custom Controls > Surface Control** node and set the following properties:

Node	Property	Setting
Surface Control	**Part Surfaces**	**Sensor**
Controls > Minimum Surface Size	**Minimum Surface Size**	**Specify Custom Value**
Controls > Target Surface Size	**Target Surface Size**	**Specify Custom Value**
Values > Minimum Surface Size > Relative Size	**Percentage of Base**	`0.5` (equivalent to 0.5 mm)
Values > Target Surface Size > Relative Size	**Percentage of Base**	`1.0` (equivalent to 1.0 mm)

Rerun the Surface Wrapper:

7. Right-click **Operations > Surface Wrapper** and select **Execute**.

   When the wrapping process is complete, STAR-CCM+ updates the display with the new result:



The sensor part is no longer joined to the manifold body. The surface resolution has also improved such that the sensor is well-defined and intersects the manifold side wall crisply.

Improving the Channels Definition

Based on the current reference target size of 8 mm, the three channels that exist in the manifold cross pipe have been defeatured.

The location of the channels in the current wrap is shown below:



Depending on the requirements of the flow analysis, defeaturing the channels can be a desirable result, as it simplifies the geometry without adversely affecting the flow solution. The defeaturing occurs as the channels themselves are approximately only 2.5 mm wide. The octree triangles in this area are based on the target size, which means the surface wrapper cannot capture the detail. To maintain this feature, a localized target size on the order of the channel width is required.

To improve the resolution of the channel, apply boundary-specific wrapper values. Work on the original parts, not the surface wrapper part. The current scene shows the Surface Wrapper part.

1. Separate the surface of the manifold body that contains the channels, split the surface of the manifold body:
   a) Expand the **Geometry** > **Parts** > **Manifold Body** node.
   b) Right-click **Surfaces** > **Faces** and select **Split by Patch**.
   c) In the *Split Part Surface by Patch* in-place dialog, select patches, **78**, **80**, **82**, and **210** through **218** inclusive.
   d) Set **Part Surface Name** to `Channels`.
   e) Click **Create** then **Close**.

   The new surface appears in the simulation tree.

2. Specify the surface size values using another Surface Control:
   a) Right-click **Operations** > **Surface Wrapper** > **Custom Controls** and select **New** > **Surface Control**.
   b) Select the **Surface Control 2** node and set **Part Surfaces** to **Manifold Body** > **Surfaces** > **Channels**.
   c) Expand the **Surface Control 2** node and set the following properties:

Node	Property	Setting
Controls		
Target Surface Size	**Target Surface Size**	**Specify Custom Value**
Minimum Surface Size	**Minimum Surface Size**	**Specify Custom Value**
Values		
Target Surface Size > Relative Size	**Percentage of Base**	`0.5` (equivalent to 0.5 mm)
Minimum Surface Size > Relative Size	**Percentage of Base**	`0.5` (equivalent to 0.5 mm)

   Using a size of 0.5 mm means that up to five triangles are used for the channel width, resulting in a better definition.

Rerun the Surface Wrapper:

3. Right-click **Operations** > **Surface Wrapper** and select **Execute**.

When the process is complete, display the channel part surfaces:

4. Expand the **Displayers** > **Geometry 1** node.
5. Select the **Parts** node, maintain the current selection of parts, and also select the **Surface Wrapper** > **Manifold Body** > **Channels** part surface.



The channels are now well-defined and have sharp edges in the circumferential direction.

Improving the Flange Definition

Since the surface wrapper has defeatured the original detail of the flange parts, a contact prevention is used to improve the output of the surface wrapper around the flanges.

There are two sets of flanges in the manifold assembly:

- between the inlet pipe and manifold body
- between the cross-pipe and the manifold body

These flanges represent high curvature areas that are close to one another but are not joined. Each wrapping iteration results in defeaturing of the flange areas, as shown below:



The flange features must be eliminated from the wrapper result by isolating the original boundaries that correspond to the flanges. Local surface target and minimum sizes -- which are bigger than the existing global reference sizes, must be set for the flanges . This effectively flattens out the area between the manifold and corresponding pipes.

However, to include the flanges in the volume mesh, a contact prevention option must be used. Since there are three different boundaries in this case (the manifold body, inlet pipe, and cross-pipe) then the two group contact prevention set option can be used to minimize the setup effort involved.

1. Right-click **Operations** > **Surface Wrapper** > **Contact Prevention** and select **New** > **Two Group Contact Prevention Set**.

   A new node, **Contact Prevention** appears under the **Contact Prevention** node.
2. Select the **Contact Prevention** > **Contact Prevention** node and set the following properties:

Property	Setting
Surface set 1	**Manifold Body**
Surface set 2	**Cross pipe** and **Inlet pipe**
Minimum Size	`2.0E-4 m (0.2 mm)`

In this case, a contact prevention set is formed between the parts in the Surface set 1 group and the parts in the Surface set 2 group. Contact prevention is not enforced between the inlet pipe and cross pipe parts.

The *Minimum Size* property determines the cut-off thickness where the two curved surfaces in the flange come together. A reasonably small value is used to capture the high curvature and produce a uniform join in the circumferential direction.

This means that the smallest triangle size used at the interface where the two flange surfaces come together is 0.2 mm. A larger value could be used but it would result in a lower definition.

3. Rename the **Contact Prevention** node to `Flange CP`.

Rerun the Surface Wrapper with the new settings:

4. Right-click **Operations > Surface Wrapper** and select **Execute**.

   When the process is complete, STAR-CCM+ updates the display with the new result.



The flanges are now well-defined and have a sharp edge in the circumferential direction.

Improving the Valve Definition

Since the surface wrapper has defeatured the gap between the valve and the cross pipe surface, another contact prevention is used to improve the output of the surface wrapper.



Since the valve and the cross pipe surface are in close proximity, the surface wrapper defeatures the gap between the two parts. A contact prevention is used to improve the resulting poor definition in those areas. The minimum size is calculated based on the gap between the two parts, which is approximately 1 mm. No additional refinement of the valve part is required in this case since it is planar.

1. Right-click **Operations > Surface Wrapper > Contact Prevention** and select **New > Two Group Contact Prevention Set**.
2. Select the **Contact Prevention** node and set the following properties:

Property	Setting
Minimum Size	`0.5 mm`
Surface set 1	**Cross Pipe**
Surface set 2	**Valve**

The *Minimum Size* value determines the number of triangles that are placed in the gap. A value of half the gap size produces a good result and does not produce too much surface refinement.

3. Rename the **Contact Prevention** node to `Valve CP`.

Rerun the Surface Wrapper with the new settings:

4. Right-click **Operations** > **Surface Wrapper** and select **Execute**.

   When the process is complete, STAR-CCM+ updates the display with the new result:



The valve is now well-defined and does not intersect the side wall of the cross-pipe.

Improving the Baffle Definition

The baffles within the geometry are initially represented as intersecting surfaces which pass through the outlet pipes that emerge from the manifold body.

The surface wrapper can detect these surfaces and automatically "inflate" them, so that a small thickness is introduced into the shape of the baffle. So far, the wrapper has not done a good job at capturing either the surface, or the intersection of the baffles with the pipes.



Three main options exist for improving the quality of the resultant baffle definition:

- Creating a true baffle of zero thickness
- Using contact prevention and local boundary sizes
- Using the offset faces/edges tool to inflate the baffle

By default, the surface wrapper is able to inflate surfaces so that they have a thickness. However, the inflation thickness is dependent on the triangulation of the surface. This method is useful if you simply want to provide an internal wall to impede flow. Alternatively, inflate the baffle using the surface preparation offset faces/edges tool and specify the exact inflation thickness and direction before the wrapper is used. The surface wrapper then treats the inflated surfaces the same as any other closed body and incorporates them into resultant surface.

If you want to generate a mesh with zero-thickness or “true” baffles, you can use the surface wrapper to stitch the baffle surfaces onto the

outlet walls. To achieve this, some surface preparation for the baffles is required to connect the baffles to the outlet walls, and they must be defined as interfaces within the model. When the surface wrapper is run, the mesher identifies the baffles and they are stitched into the final mesh with zero thickness.

Note: To try each of the above methods for comparison, first save the simulation file to be able to revert to the model before the baffles were altered.

Contents

Creating a True Baffle
Using the Built-in Wrapper Baffle Inflation with Contact Prevention
Using the Offset Faces/Edges Tool to Inflate the Baffles Manually

Creating a True Baffle

For cases where a zero-thickness baffle is required, and the original imported surface has the necessary faces, the surface wrapper can stitch in the baffle faces after the main wrap.

First connect the baffle to the repaired input surface. When the wrapped surface is fed into the surface remesher, the result is a high-quality, zero-thickness, interior wall created for the baffles.

1. In the geometry displayer, to view the original parts without the surface wrapper, deselect the surface wrapper part.
2. Enter surface repair mode with the baffle and input surface parts:
   a) Multi-select **Geometry** > **Parts** > **Baffles**, and **Manifold Body**.
   b) Right-click one of the highlighted nodes and select **Repair Surface**.
   c) In the *Surface Repair Options* dialog that appears, click **OK**.
   d) Set *Mode* to **Intersect/Boolean Faces**.

   The **Intersect/Boolean Faces** tool requires you to select two sets of faces, *Face Set 1* and *Face Set 2*.
3. Select the faces for *Face Set 1* using tools from the *Selection Control* group box:
   a) In the *Selection Control* group box, deactivate **Edges** and activate **Faces**.

b) Click (**Select faces/edges on part surfaces/part curves**) to select faces on parts.

c) In the *Select Part Surfaces/Curves* dialog, select the **Manifold Body** part.

d) Click **Apply** then **Close**.

The Manifold Body part surfaces appear selected.

4. Add the selected faces to *Face Set 1*:

   a) In the *Face Set 1* box within the **Intersect/Boolean Faces** group box, click (**Add selected faces to the set**).

   b) In the *Selection Control* group box, click (**Select none**) to clear the existing selection that is displayed in the Graphics window.

5. Repeat the above procedure to add all Baffles Part Surfaces to Face Set 2.

6. In the **Intersect/Boolean Faces** box, click **Intersect**.



The display in the *Graphics* window is updated to show the lines of intersection between the two sets of faces. A clearer view can be obtained by showing only the Baffle surfaces in the scene:



If there are baffle surfaces protruding beyond the wrapped surface, the volume mesher does not operate properly. To avoid this, the protruding portions of the baffle surfaces must be removed. These portions are shown below in beige:



7. Delete the protruding portions of the baffle surfaces:
   a) In the *Graphics* window, double-click one of the faces within the area that you are deleting.

      This action selects the whole portion within any bounding edges.
   b) Hold down `<Ctrl>` and continue to double-click faces within the other portions shown previously, until all the portions are selected.
   c) Set *Mode* to **Repair Surface Mesh**.
   d) In the *Surface Repair* group box, click (**Delete selected faces**).
   e) At the bottom of the *Surface Preparation* panel, click **Close**.

Create contacts between the baffles and the manifold body:

8. Expand the **Geometry** > **Parts** > **Baffles** > **Surfaces** node.
9. Multi-select **Faces**, **Faces 2**, and **Faces 3**, then right-click one of the selected nodes and select **Create Baffle**.

Rerun the Surface Wrapper:

10. Right-click **Operations** > **Surface Wrapper** and select **Execute**.
    When the process is complete, STAR-CCM+ updates the display with the new result:



These baffles are ready for use in the surface remesher. The final surface mesh is then suitable for volume meshing as long as parts are assigned to regions.

Using the Built-in Wrapper Baffle Inflation with Contact Prevention

To improve the surface definition using the default built-in wrapper baffle inflation option, include a contact prevention for the baffle surfaces and increase the surface triangulation on the baffles.

A single group contact prevention can be used in this instance, since the baffles all belong to one part and the manifold body to another. Use a small Minimum Size value due to curvature occurring in two different directions at the intersection of the baffles and manifold pipes. The Minimum Size value also has some influence over the thickness for the baffle inflation.

1. If you completed the steps for creating a true baffle, close the simulation without saving, and open the `.sim` file that you saved earlier.
2. Right-click **Geometry** > **Operations** > **Surface Wrapper** > **Contact Prevention** and select **New** > **One Group Contact Prevention Set**.



3. Select the **Contact Prevention** node and set the following properties:

Property	Setting
Surfaces	**Baffles** and **Manifold Body**
Minimum Size	`1.0 mm`

4. Rename the **Contact Prevention** node to `Baffles CP`.
5. To obtain a more refined surface, change the baffle surface sizes using another Surface Control:
   a) Right-click **Operations** > **Surface Wrapper** > **Custom Controls** and select **New** > **Surface Control**.
   b) Select the **Surface Control 3** node and set *Part Surfaces* to **Baffles**, then set the following properties:

Node	Property	Setting
Controls		
Target Surface Size	*Target Surface Size*	**Specify Custom Value**
Minimum Surface Size	*Minimum Surface Size*	**Specify Custom Value**
Values		
Target Surface Size > Relative Size	*Percentage of Base*	`2.0` (2.0 mm)
Minimum Surface Size > Relative Size	*Percentage of Base*	`1.0` (1.0 mm)

6. To rerun the Surface Wrapper, Right-click **Operations** > **Surface Wrapper** and select **Execute**.

   When the process is complete, STAR-CCM+ updates the display with the new result:



A close up of the baffle intersection area is shown below:



Using the Offset Faces/Edges Tool to Inflate the Baffles Manually

Use the surface preparation offset faces/edges tool to inflate surfaces so that they have a thickness.

The inflation can take place in any specified direction and can be done on more than one boundary at a time. By inflating the surfaces, the wrapper can then deal with the resultant closed body shape more easily. The feature is also incorporated as part of the final surface and a good definition is maintained.

1. If you completed the steps for creating a true baffle or using the built-in wrapper baffle inflation with contact prevention, close the simulation without saving, and open the `.sim` file that you saved earlier.
2. In the geometry and outline displayers, to view the original parts without the surface wrapper, deselect the surface wrapper part.
3. Multi-select the **Geometry** > **Parts** > **Baffles** and **Manifold body** parts.
4. Right-click one of the highlighted Parts and select **Repair Surface**.
5. In the *Surface Preparation* in-place dialog, click **OK**.
   No checks are required as you are not directly repairing the surface and do not need to know the location or number of problems that exist.
6. In the *Surface Repair Options* dialog, set **Mode** to **Offset Faces/Edges**.
7. Select the surfaces to offset:
   a) In the **Selection Control** group box, activate **Faces** only.
   b) Click (**Select faces/edges on part surfaces/part curves**).
   c) Select **Baffles**.
   d) Click **Apply** then **Close**.
   The baffle surfaces are highlighted:



Specify the offset parameters:

8. In the *Offset* group box, set the following properties:

Property	Setting
Offset Mode	**Faces**
Offset Action	**Bi-directional Inflate**
Thickness	`0.003 m`
Direction	**+ Normal Direction**

Perform the inflation process:

9. Click **Offset** then **Close**.

   The scene is updated with the new geometry definition for the baffles, as shown below:



In this instance, the original baffle surface has been expanded in both directions and the sides were filled to form a closed body. The original boundary has been removed and no feature curves are included in the new body as there were no feature curves in the original surface (which is desirable in this case). No additional surface preparation is required.

10. Edit the baffle surface sizes using a custom surface control, as outlined in the earlier section. For reference, the sizes used were:

Node	Property	Setting
Minimum Surface Size > Relative Size	*Percentage of Base*	`1.0` (1.0 mm)
Target Surface Size > Relative Size	*Percentage of Base*	`2.0` (2.0 mm)

Rerun the Surface Wrapper:

11. Right-click **Operations > Surface Wrapper** and select **Execute**.

In this instance, the wrapper has incorporated the inflated body into the main geometry and left a well-defined void to represent the baffle.



What to Do Next

The surface wrapper always provides a closed surface definition. The next step is to improve the overall quality of the triangulation so that you can generate the volume mesh.

Use the **Operations > Automated Mesh > Surface Remesher**. As a starting point, the same surface sizes as used for the wrapper can be used for the remesher, but alter the boundaries sizes to optimize the triangulation for the volume meshing process.

The image below shows the result of the surface remesher using the same exact sizes as for the wrapper:

Once you have obtained a valid, high-quality surface mesh, you can assign parts to regions, then generate a volume mesh using the same process as described in the Parts-Based Meshing: External Aerodynamics tutorial in this manual.

Summary

This tutorial introduced the following:

- Importing multiple CAD files and combining regions.
- Using the fill polygonal patch option to cap open boundaries.
- Setting the wrapper properties and reference values.
- Using the leak detection tool.
- Implementing contact prevention.
- Using surface refinement for a boundary.

Surface Repair Tools Brake Assembly

STAR-CCM+ provides a suite of tools to help you prepare surface geometry so that it is ready for meshing. This tutorial provides an introduction to some of these tools, using the example of a brake assembly.

Volume meshers in STAR-CCM+ take a closed surface mesh as their input. Errors in the surface mesh have a direct impact on the success of the volume meshers, and on the quality of the resulting mesh.

Geometry that is imported from CAD is often not prepared for CFD analysis. The imported geometry can have unconnected parts, zero-thickness surfaces, and other defects such as gaps and holes.

Note: For complex geometries that contain too many problems to repair manually, STAR-CCM+ provides the surface wrapper. See the Surface Wrapper tutorial.



Contents

Prerequisites
Loading the Simulation File
Storing the View

Using the Surface Repair Tools
Fixing Face Proximity: Removing Duplicate Faces
Removing Free Edges: Inflating Thin Surfaces
Removing Free Edges: Filling the Hole
Repairing Features
Closing Gaps
Boolean Operation
Merging and Imprinting Adjacent Parts
Remeshing the Surface
Checking Surface Mesh Quality in manual Repair Mode
Summary

Prerequisites

The instructions in the Surface Repair Tools Break Assembly tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Loading the Simulation File

For this tutorial, you are provided with a simulation file that contains pre-defined objects.

Load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/mesh/data`.
5. Select `surfaceRepairCAD.sim` then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `surfaceRepairTutorial.sim`.

Storing the View

Create a Geometry scene and save a specific view.

Saving a specific view of conveniently orientated geometry is useful for you to access throughout the tutorial.

1. Create a Geometry Scene.

Change the color mode and remove the outline.

2. Expand the **Geometry Scene 1 > Displayers** node and set the following properties:

Node	Property	Setting
Geometry 1	*Color Mode*	**Distinguish Inputs**
Outline 1	*Outline*	deactivated

Adjust the view.

3. In the *Vis* toolbar, click (**Save-Restore-Select views**).
4. Select **Views > +X +Y -Z > Up +Y**.

Store the current view.

5. Click (**Save-Restore-Select views**) and select **Store Current View**.

Using the Surface Repair Tools

The surface repair tools let you find and fix problems with the imported geometry.

Launch the surface repair tools:

1. In the *Mesh Generation* tool bar, click (**Launch surface repair**).
2. In the *Surface Repair Options* dialog, next to **Geometry Parts**, click **Select**.
3. In the **Select Objects** dialog, select all parts and click **OK**.
4. In the *Surface Repair Options* dialog, click **OK**.
5. In the *Surface Diagnostics* group box, click **Manage Thresholds**.
6. Activate all Thresholds then click **OK**.

   The *Surface Diagnostics* group box reports the number of tessellated faces that fail each of the checks.



7. Click each of the threshold **Count** numbers in turn.

   The *Geometry Scene 1* window updates to show all faces with problems that are caused by each of the Threshold problems:



For example:

8. Click the red **Count** number for the **Pierced faces (default)** threshold.

   The *Geometry Scene 1* window updates to show all faces that have problems that are caused by pierced faces.

Fixing Face Proximity: Removing Duplicate Faces

In this case, many of the faces that fail the Face Proximity check are due to duplicate faces on the upper surface of the brake disk.

Check for duplicate faces and merge them:

1. In the *Selection Control* group box, activate **Faces**.
2. Click (**Select None**).
   Any selections are reset.
3. Multi-select a number of the faces on the upper surface of the brake disk.



4. In the *Surface Repair* group box, click (**Delete selected faces**). This action deletes the chosen faces, and reveals duplicate faces beneath.

5. In the *Surface Repair* group box, click (**Merge adjacent vertices/ duplicate faces**).
6. In the *Merge Options* dialog, click **OK** to accept the default settings.

   Input is set to **Active Model**, which means that all duplicate faces in the geometry are merged. The duplicate surfaces are removed. The Surface Diagnostics group box shows the reduction in the face proximity threshold count.

Reset the selection and deselect all faces and vertices:

7. Click (**Select none**), or press `<SHIFT>` and `<N>`.

Removing Free Edges: Inflating Thin Surfaces

Use the offset tool to convert zero-thickness surfaces into a thin volume.

Surfaces that do not enclose a particular volume are a source of free edges. You can inflate these surfaces to avoid errors that are caused by free edges.

1. At the top of the *edit* window, make sure that *Mode* is set to **Repair Surface Mesh**.
2. In the *Surface Diagnostics* group box, click the green **Count** number to the right of the **Free Edges (default)** diagnostic.

   All free edges are highlighted in the *Geometry Scene 1* window.

The free edges that surround the heat shield indicate that it is a thin surface. Here, you inflate the thin surface to form a volume and so remove the free edges.

Take a closer look at the heat shield:

3. In the *Vis* tool bar:

   a) Click (**Save-Restore-Select views**).

   b) Select **Restore View > View 1**.

   c) Click (**Rubberband Zoom**).
   The rubberband zoom selection tool appears as two lines, horizontal and vertical, that intersect each other in the *Geometry Scene 1* window.

4. Start the selection:

   a) Click once in the area that is shown below:



   b) Move the mouse diagonally down and right until you have selected the following area:



c) Click once to complete your selection.

5. Set **Mode** to **Offset Faces/Edges**.
6. Reset the selection and deselect all **Faces**, **Edges**, and **Vertices**, in the *Selection Control* group box, click **Select None**, or press `<SHIFT>` and `<N>`.
7. In the *Selection Control* group box, activate the **Faces** selection.
8. In the *Geometry Scene 1* window, right-click one face of the heat shield and select: **Part (Heat Shield) > Select**.
9. In the *Offset* group box within the *edit* panel, set the following Properties:

Property	Setting
Offset Action	**Inflate**
Thickness	`0.001m`
Direction	**+ Normal Direction**

10. Click **Offset**.

Note: Take care not to click **Offset** more than once. Doing so creates duplicate faces which can cause problems when performing boolean operations.

Since the heat shield is a volume instead of a surface, it has no free edges.



Removing Free Edges: Filling the Hole

Free edges that are caused by holes or missing faces on the surface must also be repaired.

For surfaces that contain holes or have missing faces, you can fill the holes using the surface repair tools: Fill Polygonal Patch, or, Fill holes using selected edges. If the hole is intentional, then use the Flag edges as a feature tool and maintain the hole as a feature of the geometry.

1. In the *edit* tab, set **Mode** to **Repair Surface Mesh**.
2. In the *Surface Diagnostics* group box, click the green *Count* number next to the **Free Edges** (default) diagnostic.

   The remaining free edges are highlighted on the rim of a hole in *Geometry Scene 1*.

3. In the *vis* tool bar, click (**Save-Restore-Select views**) and select **Standard Views > (+ Z) top**.

4. Use the **Rubberband Zoom** selection tool to view the following area:



5. In the *Surface Repair* group box, click (**Fill holes using selected edges**).

   The hole is filled, but the local curvature is not maintained. As a result, the patch is not consistent with the surrounding geometry:



6. At the bottom of the edit tab, click (**Undo**).

Note: If you use the **Undo** button to remove any unwanted surface repairs on existing faces, the faces with those surface repairs disappear from the *Geometry Scene 1* window. The faces are not permanently deleted. To re-display the missing faces, in the *Display Control* group box, click (**Show all faces**).

7. Right-click (**Fill holes using selected edges**) and select **Options**.
8. In the *Hole Filling Options* dialog, activate **Maintain Local Curvature** and click **OK**.
9. Click (**Fill holes using selected edges**).

The hole is filled with a surface that follows the local curvature.



In the *Surface Diagnostics* group box, there are no free edges.

Repairing Features

Ensure that all feature curves that are required to represent the geometry are correctly identified in the surface mesh.

Feature curves are edges that you wish to retain in the surface mesh.

1. Set *Mode* to **Repair Features**.
2. In the *Feature Diagnostics* group box, click **Start Diagnostics**.

   This process identifies feature errors - that is, edges that are features but are not represented as such, and edges that are incorrectly represented as features.



3. In the *Feature Repair* group box, click (**Fix all feature errors**). All feature errors are now corrected.

Closing Gaps

Project vertices onto a set of faces to close a gap.

An inspection of the brake assembly geometry reveals that there are gaps between the brake pads and the disks. For this tutorial, you close the gaps using the (**project selected vertices tool**).

Select the required vertex and faces, and project the vertices:

1. In the *Edit* tab, set *Mode* to **Repair Surface Mesh**.
2. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Restore View > View 1**.
3. In the *Display Control* group box:
   a) Click (**Hide all faces**).
   b) Click (**Show faces on Part Surfaces**).

c) In the *Show Part Surfaces* dialog, select **Brake Disk**, **Lower Pad**, and **Upper Pad**.

d) Click Apply, then click Close.

4. Click (**Save-Restore-Select views**) and select **Views > +X -Y +Z > Up +X**.

5. Use the (**Rubberband Zoom**) tool to view the following area:



Select all the faces on the surface of the brake pad closest to the brake disk:

6. Select a face on the surface of the brake pad closest to the brake disk:



7. In the *Selection Control* group box, click (**Multiple grow selected**).
8. In the *Multiple grow options* dialog, activate **Break at Sharp Angles**.



9. Click **OK** to accept the remaining default selections.

The surface of the brake pad is selected.

Select all of the vertices on the surface of the brake pad closest to the brake disk:

10. In the *Selection Control* group box:

    a) Click (**Select attached vertices**).

    b) Click (**clear selected faces**).
    The **Faces** option must remain active.

11. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Standard Views > (+Y) right**.

Select the faces on the surface of the brake disk:

12. Press and hold `<Ctrl>` then double-click a face on the brake disk surface.
    All vertices of the brake pad and faces of the brake disk on adjacent surfaces are selected:



Project the vertices onto the faces:

13. In the *Surface Repair* group box, click (**Project selected vertices**).
14. In the *Project Vertices* in-place dialog, set the following properties:

Property	Setting
Projection Mode	**Project to a Surface**
Projection Method	**Shortest Distance**

15. Click **Apply** then **Return**.
    The gap between one brake pad and the brake disk is closed.

Proceed to close the gap between the other brake pad and the brake disk:

16. In the *Selection Control* group box, click (**Select none**).
17. In the Vis tool bar, click (**Save-Restore-Select Views**) and select **Views** > **+X +Y +Z** > **Up +X**.
18. Use the (**Rubberband Zoom**) tool to view the following area:



Select all the faces on the surface of the second brake pad closest to the brake disk:

19. Select a face on the surface of the brake pad closest to the brake disk:



20. In the *Selection Control* group box, click (**Multiple grow selected**).

21. In the *Multiple grow options* dialog, maintain the **Break at Sharp Angles** selection along with the remaining default settings and click **OK**.

    The surface of the brake pad is selected.

Select all of the vertices on the surface of the brake pad closest to the brake disk:

22. In the *Selection Control* group box:

    a) Click (**Select attached vertices**).

    b) Click (**clear selected faces**).
    The **Faces** option must remain active.

23. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Views > +X -Y -Z > Up -Y**.

Select the faces on the surface of the brake disk:

24. Press and hold `<Ctrl>` then double-click a face on the brake disk surface.

    All vertices of the brake pad and faces of the brake disk on adjacent surfaces are selected:



Project the vertices onto the faces:

25. In the *Surface Repair* group box, click (**Project selected vertices**).
26. In the *Project Vertices* in-place dialog, set the following properties:

Property	Setting
Projection Mode	**Project to a Surface**
Projection Method	**Shortest Distance**

27. Click **Apply** then **Return**.

    The gap between one brake pad and the brake disk is closed.

    The Pierced faces and Face Proximity thresholds show a deterioration as extra faces are now in close contact without intersecting. This problem is resolved in the step for imprinting surfaces.
28. To exit surface repair mode, click **Close**.

Boolean Operation

There are several components of the brake assembly that overlap each other or do not intersect cleanly. These components require further attention using the boolean operation tools.

Note:

You can click (**Show All Mesh**) or (**Hide All Mesh**) in the *Vis* tool bar depending if you want to view the mesh or not.

Fix the overlapping heat shield and mount:

1. Make sure that *Geometry Scene 1* is active.
2. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Restore View > View 1**.
3. View only the heat shield and mount:
   a) Right-click **Geometry Scene 1 > Displayers > Geometry 1 > Parts** and select **Edit**.
   b) In the *Parts* dialog, select the **Heat Shield** and **Mount** parts then click **OK**.

   Only the selected parts are displayed in the *Geometry Scene 1* window.
4. In the Vis tool bar, click (**Save-Restore-Select Views**) and select **Views > -X +Y +Z > Up +Y**.
5. Use the (**Rubberband Zoom**) tool to view the area where the parts overlap:



6. Select the **Geometry** > **Parts** > **Heat Shield** node.

   The heat shield appears highlighted in the *Geometry Scene 1* window and shows where the Heat Shield intersects the Mount.



7. Create a new Subtract part that does not intersect the Mount:

a) Multi-Select the **Geometry** > **Parts** > **Heat Shield** and **Mount** nodes.

b) Right-click one of the selected nodes and select **Boolean** > **Subtract Parts**.

c) In the *Subtract parts* dialog, select **Heat Shield** as the **Target Part**.

d) Click **OK** to accept **Discrete** as the default **Boolean Mode**.

Since the geometry has been modified, you are operating on a tessellated definition.

The portions of the Heat Shield part that overlap the Mount part are removed, and a new Subtract part is generated which cleanly intersects with the Mount part.

8. Delete the Heat Shield part.
9. Rename the **Subtract** node to `Heat Shield`.
10. Display the new Heat Shield part in the Geometry Scene 1 window:

    a) Click and hold the **Heat Shield** node.

    b) Drag the part to a space in the *Geometry Scene 1* window.

    c) Select **Part Displayers** > **Add to Geometry 1**.

Fix the Intersecting Brake Disk and Rotor Studs, create a new Subtract part for the five Rotor Stud parts and the Brake Disk part.

11. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Views** > **+Y** > **Up +X**.
12. View only the five Rotor Stud parts and the Brake Disk part:

    a) At the top of the simulation tree, click *scene/plot*.

    b) Edit the **Geometry Scene 1** > **Displayers** > **Geometry 1** node and set **Parts** to the five **Rotor Stud** parts and the **Brake Disk** part.

    c) Click **OK**.

    Only the selected parts are displayed in the Geometry Scene window.
13. Select the rotor stud and brake disk parts that intersect:

    a) Expand the **Geometry** > **Parts** node and multi-select the five rotor stud nodes.

    b) While the rotor stud nodes are still selected, hold down `<Ctrl>` and select the **Brake Disk** node.
14. Create the Subtract part:

a) Right-click one of the selected nodes and select **Boolean > Subtract Parts**.

b) In the *Subtract Parts* dialog, set **Target Part** to **Brake Disk**.

c) Click **OK** to accept **Discrete** as the default **Boolean Mode**.

A new **Subtract** part appears under **Geometry > Parts**.

15. Delete the **Geometry > Parts > Brake Disk** part.
16. Rename the **Subtract** node to `Brake Disk`.
17. Display the new Brake Disk part in the *Geometry Scene 1* window:

    a) Click and hold the **Brake Disk** node.

    b) Drag the part to a space in the *Geometry Scene 1* window.

    c) Select **Part Displayers > Add to Geometry 1**.

Merging and Imprinting Adjacent Parts

Generate a surface mesh suitable for CFD, imprint adjacent parts, and merge the faces and nodes.

View all **Parts** in **Geometry Scene 1**:

1. In the *Mesh Generation* tool bar, click (**Launch surface repair**).
2. Set **Geometry Parts** to **All Parts** and click **OK**.
3. In the *Edit* tab, set **Mode** to **Merge/Imprint**.
4. In the *Merge/Imprint Tool* group box, set the **Imprint Mode** to **Multi-Part Imprint**.
5. In the *Part Browser* box, click **Find Pairs**.

   The Part Browser identifies adjacent pairs of parts. Most of the pairs of parts can be imprinted immediately by using Imprint All. However, in the case of the Calliper and Mount parts, surfaces that are next to the contact area meet with a small angle. These parts need imprinting with different settings.

View the surfaces on the **Calliper** and **Mount** parts that meet with a small angle:

6. In the *Part Browser* box within the *Merge/Imprint Tool* group box, select the **Calliper** and **Mount** parts from the drop-down menus.

   The **Calliper** and **Mount** parts are displayed in the *Geometry Scene 1* window.

7. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Views > +X > +Y**.

8. Use the **Rubberband Zoom** tool to view the following area:



The surfaces on the calliper and mount parts that meet with a small angle are shown.



9. In the *Merge/Imprint Tool* group box, activate **Show Advanced Parameters** and set the following properties:

Property	Setting
Merge Angle	`20.0 deg`
Tolerance	`0.1 mm`

   Setting the **Merge Angle** lower for this pairing prevents nodes and faces on non-contacting surfaces from being incorrectly projected.

10. Click **Imprint Pair**.
11. In the Vis tool bar, click **Reset View** .

    The merged surfaces between the calliper and mount appear highlighted in the *Geometry Scene 1* window.

12. In the *Merge/Imprint Tool* group box, set the following properties:

Property	Setting
Merge Angle	`45 deg`
Tolerance	`0.1mm`

13. In the *Part Browser* box, click **Stop** then click **Find Pairs**.

    Pairs with this specified **Merge Angle** are found.

14. Click **Imprint All**.

The surfaces of the remaining contacting pairs are imprinted and highlighted. All parts now connect conformally, and faces at the contact surfaces between parts are merged.



Only parts that were merged and imprinted are displayed in the *Geometry Scene 1* window.

Display all parts:

15. In the *Display Control* group box, click (**Show all faces**).

Remeshing the Surface

Produce a good quality surface mesh that is suitable for volume meshing.

It is important to check for any further issues that require surface repair and produce a good quality surface mesh so that any surface errors are not projected through the volume mesh.

Check for any surfaces that possibly still require surface repair:

1. Set *Mode* to **Repair Surface Mesh**.
2. In the *Surface Diagnostics* group box, click **Manage Thresholds**.
3. Make sure that all thresholds are activated.
4. Click **OK**.

   There are still some issues that are displayed in the *Surface Diagnostics* group box.
5. For **Face quality (default)**, click **Count**.
6. In the *Vis* tool bar, click (**Save-Restore-Select Views**) and select **Views > -Y > Up -Z**.



In the Graphics window, poor quality faces are displayed, especially on the underside of the Brake Assembly. Most of the selected faces appear thin, with sharp edges.

Since the geometry still has a tessellated surface, this potential issue is likely to resolve itself when the surface is remeshed. So there is no need for further manual repair to fix this issue.

7. For **Face Proximity (default)**, click **Count**.

   In the Graphics window, faces that are in close proximity to others appear selected. Most of the selected faces are on the heat shield part of the Brake Assembly. Faces that are in close proximity are a potential issue in some scenarios. However, in this case the heat

shield part is intentionally close to the rest of the geometry. This potential issue is not a problem as long as a fine enough mesh is created.

8. Click **Close**.
9. Click the simulation tab.
10. Create a Surface Remesher Automated Mesh Operation:
    a) Right-click the **Geometry** > **Operations** node and select **New** > **Automated Mesh**.
    b) Select all **Parts** and activate **Surface Remesher**.
    c) Click **OK**.
11. Edit the **Automated Mesh** > **Default Controls** and set the **Base Size Value** to *0.01m*.
12. Right-click the **Automated Mesh** node and select **Execute**.
13. View the remeshed surface:
    a) Right-click a blank space in the *Graphics* window and select **Apply Representation** > **Automated Mesh.Remesh**.
    b) In the *Vis* toolbar, click (**Show All Meshes**).

    The underside of the Brake Assembly no longer has any thin or sharp faces as the tessellated surface has been replaced with a fine surface mesh.

Checking Surface Mesh Quality in manual Repair Mode

Check the surface mesh quality in more detail.

Visual inspection of the remeshed surface shows that it is a good quality surface mesh. Since the quality of the surface mesh has a direct impact on the quality of the volume mesh, it is important to ensure that the surface mesh is free from errors.

Review the remeshed surface in the manual surface repair mode.

1. In the *Mesh Generation* tool bar, click (**Launch surface repair**).
2. In the *Surface Repair Options* dialog, set the following properties:

Property	Value
Geometry Description	**Automated Mesh.Remesh**
Geometry Parts	All Parts

3. In the *Surface Repair Options* dialog, click **OK**.

   The *edit* tab appears.

4. Set *Mode* to **Repair Surface Mesh**.
5. In the *Surface Diagnostics* group box, click **Manage Thresholds**.
6. Activate all Thresholds then click **OK**.

   The *Surface Diagnostics* group box reports that none of the remeshed faces fail the diagnostic checks.



   You now have a geometry with a good quality surface mesh which is ready to use with the STAR-CCM+ volume meshers.
7. Save the simulation.

Summary

This tutorial used a brake assembly to demonstrate some of the typical steps that are involved in geometry preparation and repair using STAR-CCM+.

The steps covered were:

- Using surface diagnostics
- Removing duplicate parts

- Repairing features
- Inflating thin surfaces
- Filling holes
- Carrying out boolean operations on geometry
- Merging and imprinting parts
- Remeshing the surface
- Checking surface mesh quality
- Closing gaps

Parts-Based Meshing External Aerodynamics

This tutorial demonstrates how to set up a simulation using the parts-based meshing approach.

Parts-based meshing is a method for generating surface and volume meshes on geometry parts. It is different to region-based meshing in that it uses a series of mesh operations to define the process from the initial geometry to the final volume mesh. Having a series of operations allows you to modify the initial geometry and repeat the entire process with no additional inputs. This makes parts-based meshing particularly suitable for parametric design studies.

In this tutorial, you generate a volume mesh that is suitable for an external aerodynamics analysis of a vehicle using the parts-based approach. Each step in the tutorial highlights some of the differences between parts-based meshing and region-based meshing.

The following image shows the geometry:



Contents

Prerequisites
Overview
Importing the Surface Mesh
Storing the View
Preparing the Initial Surface
Filling a Hole on the Car Surface
Creating a Composite Assembly
Generating a Wrapped Surface Mesh
Creating the Fluid Volume
Assigning the Fluid Volume to a Region
Setting Boundary Types
Generating the Volume Mesh
Replacing the Wing and Executing the Pipeline
Summary

Prerequisites

The instructions in the Parts-Based Meshing External Aerodynamics tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Overview

There are several stages within the Parts-Based Meshing External Aerodynamics tutorial.

The tutorial stages are:

1. Importing and visualizing the CAD parts.

   Begin by launching STAR-CCM+ and importing CAD files that contain surface meshes for the car, wing, and bounding box. Then, set up the view and hide some surfaces so that you can examine the vehicle.

2. Preparing the initial surface.

   Remove the original wing from the car and import a new wing that is in a different position. This action creates several problems on the geometry surface, which you fix at later stages of the tutorial.

3. Filling a hole on the car surface.

   Use the fill hole mesh operation to fill the hole that you introduced when you removed the original wing. This step demonstrates the process for creating a mesh operation.

4. Creating a composite assembly.

   One of the advantages of parts-based meshing is that you can perform mesh operations on an assembly of several parts. Create an assembly to contain the parts that you are using in this tutorial.

5. Generating a wrapped surface mesh.

   At this stage you generate a surface wrap of the Vehicle Assembly, which fixes the intersections in the surface mesh. You also set up curve controls and contact prevention to maintain important features on the surface mesh.

6. Creating the fluid volume.

   The fluid simulation requires a volume mesh between the surface of the car and the walls of the bounding box. To create the part that defines the outer surface of the volume mesh, you subtract the output of the surface wrapper operation from the bounding box.

7. Assigning the fluid volume to a region.

   Before generating a volume mesh, you must assign the fluid volume to a region. This action allows the mesher to convert the volume mesh to a finite volume representation that the solvers can use.

8. Setting boundary types.

   Certain boundary types affect the volume mesh. In this tutorial, you have a velocity inlet, pressure outlet, and symmetry boundaries that do not require prism layers.

9. Generating the volume mesh.

   Use an automated mesh operation to generate the volume mesh.

10. Replacing the wing and executing the pipeline.

    An advantage of parts-based meshing is that when you have defined the series of mesh operations, you can repeat them on different initial geometries. Here, you replace the wing component and run the meshing pipeline to generate a new volume mesh.

Importing the Surface Mesh

Launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Surface Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/mesh/data`.
5. Select `pbm_car.x_b` and `pbm_boundingBox.x_b` then click **Open**.
6. Set *Tessellation Density* to **Very Fine**.
7. To accept the remaining default *Import Surface Options*, click **OK**.
8. Save the simulation as `pbm_externalAero.sim`.

Storing the View

Set up and store a convenient view of the geometry for use throughout the tutorial.

To visualize the geometry:

1. Expand the **Displayers > Geometry Scene 1 > Parts** node.
2. Select the **Bounding Box.Far Field** node and deactivate the **Visible** property.
3. Set the view as shown below:



4. To save this view for later use, in the *Vis* toolbar, click (**Save-Restore-Select views**) and select **Store Current View**.
5. Save the simulation.

Preparing the Initial Surface

In this stage of the tutorial, you remove the original wing from the car and import a new wing that is in a different position. These steps create several problems on the geometry surface, which you fix at later stages of the tutorial.

To delete the original wing:

1. Expand the **Geometry** > **Parts** > **Car** > **Surfaces** node.
2. Delete the **Strut** and **Wing** part surfaces.

   Deleting the original wing leaves a hole in the surface mesh.



You fix this hole using the fill holes operation in a later stage of this tutorial.

To import the new wing:

3. From the menu bar, select **File** > **Import** > **Import Surface Mesh**.
4. In the *Open* dialog, select `pbm_wing1.x_b` from the `[INSTALL_DIR]/doc/startutorialsdata/mesh/data` directory and click **Open**.
5. In the *Import Surface Options* dialog, deactivate the **Open Geometry Scene After Import** option and click **OK**.

To add the imported wing to the *Graphics* window:

6. Drag and drop the **Geometry** > **Parts** > **Wing** node onto *Geometry Scene 1*.
7. In the **Part Displayers** menu, select **Add to Geometry 1**.
8. Repeat this process and add the **Wing** > **Curves** > **Edges** node to *Geometry Scene 1*.
9. Save the simulation.

The new wing intersects the car body surface, resulting in pierced faces on the surface. You fix the pierced faces by wrapping the external surface of the entire vehicle at a later stage of this tutorial.



Filling a Hole on the Car Surface

Use the fill hole mesh operation to fill the hole in the surface mesh that appeared when you removed the initial wing.

To fill the hole on the car surface:

1. Right-click the **Geometry** > **Operations** node and select **New** > **Fill Holes**.
2. In the *Create Fill Holes Operation* dialog, select **Car** and click **OK**.

   Two **Fill Holes** nodes appear in their pre-executed state in the simulation tree:

   - **Geometry** > **Parts** > **Fill Holes**.
   - **Geometry** > **Operations** > **Fill Holes**.
3. Select the **Operations** > **Fill Holes** node and set *End Surfaces* to **Main Body**.
4. Right-click the **Operation** > **Fill Holes** node and select **Execute**.

   STAR-CCM+ performs the fill holes operation. This operation causes both **Fill Holes** nodes to be displayed in their executed state.

5. Drag and drop the **Parts** > **Fill Holes** node onto *Geometry Scene 1* and select **Part Displayers** > **Add to Geometry 1**.

The fill holes operation has repaired the car surface.



Creating a Composite Assembly

One of the advantages of parts-based meshing is that you can perform mesh operations on an assembly of several parts. Here, you create an assembly to contain the parts that you are using in this tutorial.

To create a composite assembly:

1. Expand the **Geometry** > **Parts** node and multi-select the **Car**, **Fill Holes**, and **Wing** nodes.
2. Right-click one of the selected nodes and select **Composite**.

   STAR-CCM+ creates a composite node and places the original part nodes within it.
3. Rename the **Composite** node to `Vehicle Assembly`.
4. Save the simulation.

Generating a Wrapped Surface Mesh

At this stage you generate a surface wrap of the Vehicle Assembly to fix the pierced faces. You also set up curve controls and contact prevention to maintain important features on the surface mesh.

The steps to generate a wrapped surface in this tutorial are:

1. Create a Surface Wrapper operation and define its properties. See Creating and Setting Up the Surface Wrapper.
2. Apply mesh refinement to part curves. See Refining the Mesh on Part Curves.
3. Apply contact prevention controls to maintain separation between surfaces on the nose and on the main body. See Preventing Contact between Close Surfaces.
4. Execute the Surface Wrapper operation. See Executing the Surface Wrapper Operation.

Contents

Creating and Setting Up the Surface Wrapper
Refining the Mesh on Part Curves
Preventing Contact between Close Surfaces
Executing the Surface Wrapper Operation

Creating and Setting Up the Surface Wrapper

Create a surface wrapper operation and set the properties that apply to all geometry parts on which it operates.

To set up the surface wrapper mesh operation:

1. Right-click the **Geometry > Operations** node and select **New > Surface Wrapper**.
2. In the *Create Surface Wrapper Auto Mesh Operation* dialog, Select **Vehicle Assembly**.
3. Click OK.

   Two **Surface Wrapper** nodes appear in their pre-executed state in the simulation tree:

   - **Geometry > Parts > Surface Wrapper**.
   - **Geometry > Operations > Surface Wrapper**.
4. Expand the **Surface Wrapper > Default Controls** node and set the following properties:

Node	Property	Setting
Base Size	*Value*	`0.02 m`
Minimum Surface Size > Relative Size	*Percentage of Base*	`5.0`
Volume of Interest	*Method*	**External**

Setting the Volume of Interest method to External instructs the surface wrapper to ignore any internal surfaces within the Vehicle Assembly part. This fixes the pierced faces at the intersection between the wing and the vehicle surface.

Refining the Mesh on Part Curves

To make sure that you capture the detailed features in the surface mesh, apply mesh refinement to all the part curves.

To apply refinement to all part curves:

1. Right-click the **Surface Wrapper > Custom Controls** node and select **New > Curve Control**.

   STAR-CCM+ adds a curve control node within the **Custom Controls** node.
2. Edit the **Custom Controls > Curve Control** node and set **Part Curves** to:
   - **Vehicle Assembly > Car > Curves > Car Edges** and **Edges**.
   - **Vehicle Assembly > Wing > Curves > Edges** .
3. Click **OK**.
4. Edit the **Curve Control** node.
   a) Set **Controls > Target Surface Size** to **Specify Custom Value.**
   b) Expand the **Values > Target Surface Size** node.
   c) Set **Percentage of Base** to `5.0`.

Preventing Contact between Close Surfaces

To make sure that these two surfaces do not touch each other in the wrapped surface, apply a contact prevention between the nose and main body.

To apply contact prevention:

1. Right-click the **Operations > Surface Wrapper > Contact Prevention** node and select **New > One Group Contact Prevention Set**.

2. Rename the **Contact Prevention** node to `Nose - Main Body`.
3. Select the **Contact Prevention** > **Nose - Main Body** node and set the following properties:

Property	Setting
Surfaces	**Vehicle Assembly** > **Car** > **Surfaces** > **Main Body** and **Nose**.
Minimum Size	`1 mm`

Executing the Surface Wrapper Operation

Execute the operation to generate a wrapped surface.

To execute and visualize the surface wrapper operation:

1. Right-click the **Geometry** > **Operation** > **Surface Wrapper** node and select **Execute**.

   STAR-CCM+ performs the surface wrap operation. The result of the operation is that both **Surface Wrapper** nodes beneath the **Geometry** > **Parts** node and the **Geometry** > **Operations** node, are displayed in their executed state.
2. When the operation is complete, create an empty scene.
3. Drag and drop the **Geometry** > **Parts** > **Surface Wrapper** node onto *Scene 1*.
4. In the *Vis* toolbar, click (**Save-Restore-Select views**) and select **Restore View** > **View 1**.
5. Save the simulation.

Visually inspecting the surface shows that the wrapper has fixed the surface intersection issues and prevented contact between close surfaces:

Surface Intersections:



Nose and Main Body Contact Prevention:



Creating the Fluid Volume

The fluid simulation requires a volume mesh between the surface of the car and the walls of the bounding box. To create the part that defines the outer surface of the volume mesh, you subtract the output of the surface wrapper operation from the bounding box.

This tutorial uses the subtract operation instead of the extract volume operation. The reason for using the subtract operation is that the car tires (contained within the Surface Wrapper part) intersect the surface of the Bounding Box part. The extract volume operation cannot intersect surfaces; the subtract operation does.

To create the fluid volume:

1. Right-click the **Geometry** > **Operations** node and select **New** > **Subtract**.
2. In the *Create Subtract Operation* dialog, select **Bounding Box** and **Surface Wrapper**, then click **OK**.

   A **Subtract** node appears beneath the **Operations** node, and a **Subtract** part appears beneath the **Parts** node.
3. Select the **Operations** > **Subract** node and set *Target Part* to **Bounding Box**.
4. Rename the **Parts** > **Subtract** node to `Fluid Volume`.
5. Right-click the **Operations** > **Subtract** node and select **Execute**.

To view the fluid volume:

6. In the *Vis* toolbar, click **Create/Open Scenes** and select **Mesh**.
7. Edit the *Mesh* scene and set the following properties:

Node	Property	Setting
Displayers > Mesh 1	**Outline**	activated
Parts	**Parts**	**Fluid Volume**

8. In the *Vis* toolbar, click (**Save-Restore-Select views**) and select **Restore View** > **View 1**.
9. Select the **Displayers** > **Mesh 1** > **Parts** > **Fluid Volume.Bounding Box.Far Field** node and deactivate the **Visible** property.
10. In the *Vis* toolbar, click **Hide All Meshes**.
11. Save the simulation.

    On the under-side of the Fluid Volume part, the subtract operation has intersected and removed the surface that represents the tire

contact patch. The extract volume mesh operation is not capable of making these geometric changes to a surface.



Assigning the Fluid Volume to a Region

Before generating a volume mesh, you must assign the fluid volume to a region. This action allows the mesher to convert the volume mesh to a finite volume representation that the solvers can use.

To assign the fluid volume to a region:

1. Right-click the **Parts > Fluid Volume** node and select **Assign Parts to Regions**.
2. In the *Assign Parts to Regions* dialog, select the following options:
   - **Create a Region for Each Part**
   - **Create a Boundary for Each Part Surface**
3. Click **Apply** and then **Close**.

Setting Boundary Types

Certain boundary types affect the volume mesh. In this tutorial, you have a velocity inlet, pressure outlet, and symmetry boundaries that do not require prism layers. STAR-CCM+ does not grow prism layers on these boundaries by default .

To set the boundary types:

1. Expand the **Regions** > **Fluid Volume** > **Boundaries** node.
2. Set the following properties:

Node	Property	Setting
Bounding box.Far Field	*Type*	**Symmetry Plane**
Bounding Box.Symmetry Plane	*Type*	**Symmetry Plane**
Bounding Box.Inlet	*Type*	**Velocity Inlet**
Bounding Box.Outlet	*Type*	**Pressure Outlet**

Generating the Volume Mesh

Use an automated mesh operation to generate the volume mesh.

The steps required to generate a volume mesh are:

1. Create an automated mesh operation.
2. Refine the mesh for the vehicle assembly surfaces.
3. Generate and view the volume mesh.

Contents

Creating the Automated Mesh Operation
Refining the Vehicle Assembly Surfaces
Generating and Viewing the Volume Mesh

Creating the Automated Mesh Operation

An automated mesh operation determines the meshing models and controls for generating the surface and volume meshes.

To create the automated mesh operation:

1. Right-click the **Geometry** > **Operations** node and select **New** > **Automated Mesh**.
2. In the *Create Automated Mesh Operation* dialog, select the **Parts** > **Fluid Volume part**.
3. Activate the following meshing models:

Group box	Model
Surface Mesher	**Surface Remesher**
Optional Surface Mesher	**Automatic Surface Repair**
Volume Mesher	**Trimmed Cell Mesher**
Optional Volume Mesher	**Prism Layer Mesher**

4. Click **OK**.
5. Expand the **Automated Mesh** > **Meshers** node to see the models.



To set the sizes and controls for the operation:

6. Expand the **Operations** > **Automated Mesh** > **Default Controls** node.
7. Set the following properties:

Node	Property	Setting
Base Size	*Value*	`0.02 m`
Number of Prism Layers	*Number of Prism Layers*	`5`

Refining the Vehicle Assembly Surfaces

To capture the vehicle surface accurately, refine the surface mesh. This refinement also affects the volume mesh as the trimmed mesher aims to produce cells that match the surface size.

To refine the mesh that surrounds the Wrapped Vehicle surfaces:

1. Right-click the **Operation** > **Automated Mesh** > **Custom Controls** node and select **New** > **Surface Control**.
2. Rename the **Custom Controls** > **Surface Control** node to `Bounding Box`.
3. Select the **Bounding Box** node and set **Part Surfaces** to **Fluid Volume** > **Bounding Box**.
4. Expand the **Bounding Box** node and set the following properties:

Node	Property	Setting
Controls		
Minimum Surface Size	*Minimum Surface size*	**Specify Custom Value**
Target Surface Size	*Target Surface size*	**Specify Custom Value**
Values		
Target Surface Size > **Relative Size**	*Percentage of Base*	`500`
Minimum Surface Size > **Relative Size**	*Percentage of Base*	`100`

Increasing the minimum and target surface sizes for the bounding box that surrounds the vehicle assembly, avoids unnecessary refinement on the outer boundaries. The mesh in the bounding box remains refined near the vehicle assembly and grows coarse nearer the external boundaries. It is computationally less expensive to produce a coarse mesh in areas which do not require refinement.

Generating and Viewing the Volume Mesh

Executing the automated mesh operation runs the surface and volume meshers. When the meshers complete, you can view the volume mesh in a mesh scene.

To generate and view the volume mesh:

1. Right-click the **Operations** > **Automated Mesh** node and select **Execute**.

   STAR-CCM+ generates the volume mesh. This operation takes several minutes to complete.
2. When the automated mesh operation is complete, display *Mesh Scene 1* in the *Graphics* window.
3. In the *Vis* toolbar, click (**Show All Meshes**).
4. Drag and drop the **Representations** > **Volume Mesh** node into *Mesh Scene 1*.



5. Save the simulation.

Replacing the Wing and Executing the Pipeline

An advantage of parts-based meshing is that when you have defined the series of mesh operations, you can repeat them on different initial geometries. Here, you replace the wing component and run the meshing pipeline to generate a new volume mesh.

To replace the wing part and repeat the meshing pipeline:

1. Expand the **Geometry > Parts > Vehicle Assembly** node.
2. Right-click the **Wing** node and select **Replace Part**.
3. In the *Open* dialog, select `pbm_wing2.x_b` and click **Open**.

   Operations that require an update are marked with a yellow warning triangle.
4. Right-click the **Geometry > Operations** node and select **Execute All**.

   When complete, the *Graphics* window displays the updated volume mesh.
5. Save the simulation.

Summary

This tutorial demonstrated how to set up a simulation using Parts-Based Meshing.

The steps covered were:

- The parts-based meshing procedure
- Using mesh operations
- Using the surface wrapper
- The relationship between parts and regions
- The effect of boundary types on prism layer settings
- Automated mesh operations
- Mesh controls
- The meshing pipeline

Region Based Meshing Imported Control Valve

This tutorial demonstrates how to use the surface and volume meshing tools within STAR-CCM+ to generate a polyhedral volume mesh with prism layers for a single region. The inputs required to build a trimmed cell mesh with prism layers are also covered.

Start from a geometry from an imported STL file. The geometry is that of a control valve typically used for water flow management.



The control valve is in an almost closed position, with flow entering at the near end and exiting at the far end. You use the advantage of symmetry to model the physical geometry. Define the appropriate surfaces, boundaries, and feature curves as part of the meshing process.

This tutorial is split into the following sections:

- Polyhedral meshing example
- Replacing the surface mesh and remeshing using polyhedrals
- Trimmed cell meshing example

Contents

Prerequisites
Polyhedral Meshing Example
Replacing the Import Surface Mesh
Re-running the Mesh Generation
Trimmed Cell Meshing Example
Summary

Prerequisites

The instructions in the Region-Based Meshing Imported Control Valve tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Polyhedral Meshing Example

To set up the STAR-CCM+ simulation and generate a polyhedral volume mesh, launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh...**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/mesh/data`.
5. Select `control_valve.stl` then click **Open**.
6. In the *Import Surface Options* dialog, set **Units** to **in** (inches).
7. Click **OK** to accept the remaining default *Import Surface Options*.
8. Save the simulation as `control_valve1.sim`.

Contents

Splitting the Surfaces and Renaming
Extracting Feature Curves
Creating the Region and Boundaries
Selecting the Meshing Models
Setting the Global Mesh Reference Values
Setting the Boundary Mesh Values
Running the Surface Remesher
Running the Polyhedral Volume Mesher

Splitting the Surfaces and Renaming

Prepare the part surfaces so that they conform to what is required for the meshing tools. STL files only contain one patch definition which means that a single surface is created on import. The imported surface must be split so that different mesh and physics properties can be applied to different parts of the model.

To split the imported surface:

1. Expand the **Geometry** > **Parts** > **control_valve** node.
2. Right-click the **Surfaces** > **control_valve** node and select **Split By Angle**.
3. In the *Split Part Surfaces by Angle* dialog, set **Angle** to `45`.
4. Click **OK**.

   The surfaces are added to the simulation tree:



To simplify the geometry, combine some of the surfaces. First, combine the exterior wall surfaces:

5. Multi-select **control_valve 2**, **control_valve 4**, **control_valve 5**, **control_valve 6**, and **control_valve 11**.
6. Right-click one of the selected nodes and select **Combine**.

   The selected surfaces are combined into **control_valve 2**.
7. Rename **control_valve 2** to `Wall`.

Combine and rename the remaining surfaces:

8. Multi-select **control_valve 7**, **control_valve 8**, **control_valve 9**, and **control_valve 10**.
9. Right-click one of the selected nodes and select **Combine**.

   The selected surfaces are combined into **control_valve 7**.
10. Rename **control_valve 7** to `Valve Wall`.
11. Rename the remaining surfaces as shown in the table:

Surface Name	New Name
control_valve	`Symmetry Plane`
control_valve 3	`Inlet`
control_valve 12	`Outlet`

Extracting Feature Curves

Feature curves maintain sharp edges and surface feature details during the meshing process. STL surfaces do not contain feature edges so you create new edges based on an angle criteria.

To create feature edges:

1. Right click the **Parts > control_valve > Curves** node and select **Compute Part Curves**.
2. In the *Mark Part Curves Options* dialog:
   a) Select all Parts.
   b) In the *Part Curve Options* group box, de-activate all options other than **Mark Sharp Edges**.
   c) Set **Sharp Edge Angle** to `20.0`.
   d) Click **Apply** then **Close**.

Creating the Region and Boundaries

Create the regions and assign the appropriate boundary types.

To create the region from the parts:

1. Right click the **Parts > control_valve** node and select **Assign Parts to Regions**.
2. In the *Assign Parts to Regions* dialog set the drop down menus to:
   - Create a Region for Each Part.
   - Create a Boundary for Each Part Surface.

   Leave all other settings at their default value.
3. Click **Apply** then **Close**.

   The regions and boundaries are generated. Setting the boundary type at this stage is important as it affects which boundaries have prism layers next to them. Only wall boundaries have prism layers generated next to them by default.

To set the boundary types:

4. Expand the **Regions > control_valve > Boundaries** node and set the following properties:

Node	Property	Setting
Inlet	*Type*	**Velocity Inlet**

Node	Property	Setting
Outlet	*Type*	**Pressure Outlet**
Symmetry Plane	*Type*	**Symmetry Plane**
Valve Wall	*Type*	**Wall**
Wall	*Type*	**Wall**

The completed list of updated boundaries is shown below:



The surface preparation related to the imported geometry is now complete.

Selecting the Meshing Models

Use the surface remesher and the polyhedral volume mesher for this example, in conjunction with the prism layer mesher.

1. Right click the Continua node and select **New > Mesh Continuum**.
2. Right-click the **Continua > Mesh 1 > Models** node and select **Select Meshing Models**.
3. In the *Mesh 1 Model Selection* dialog, select the following models:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**

4. Click **Close**.
5. To see the list of selected models, expand the **Mesh 1 > Models** node:



6. Maintain the default properties for each model in this tutorial.

Setting the Global Mesh Reference Values

Specify appropriate reference values for the mesh.

The base size for this example is based on the size of the diameter of the inlet. All other relative meshing values are scaled from this. However, you can use absolute sizes for these properties. To measure any part of the geometry, use the Measure distance tool and click any two points on the surface. The distance between the points is reported in the Output window in meters.

1. Expand the **Continua** > **Mesh 1** > **Reference Values** node.
2. Set the following Properties:

Node	Property	Setting
Base Size	*Value*	`0.06 m`
Prism Layer Thickness > **Relative Size**	*Percentage of Base*	`4.0`
Surface Size > **Relative Minimum Size**	*Percentage of Base*	`4.0`
Surface Size > **Relative Target Size**	*Percentage of Base*	`8.0`

Setting the Boundary Mesh Values

If no boundary or region values are set then the meshing models default to the global reference values. For this case, set custom values for the surface remesher triangulation (to increase the density) and prism layer thickness for the valve wall boundary.

The valve wall requires additional cells to be created. Also, the gap between the valve and the outside wall is thinner than the rest of the geometry so requires a thinner prism layer.

1. Expand the **Regions > control_valve > Boundaries > Valve Wall** node.
2. Set the following properties:

Node	Property	Setting
Mesh Conditions		
Custom Surface Size	*Custom surface size*	activated
Customize Prism Mesh	*Customize prism mesh*	**Specify Custom Values**
Mesh Values		
Surface Size > Relative Minimum Size	*Percentage of Base*	`4.0`
Surface Size > Relative Target Size	*Percentage of Base*	`4.0`
Prism Layer Thickness > Relative Size	*Percentage of Base*	`2.0`

Running the Surface Remesher

Generate the surface mesh.

Since the surface remesher and volume mesher models have been selected in tandem, you could proceed and generate both meshes. In this tutorial, however, first generate the surface remesher mesh and view it before proceeding with the volume mesh.

1. To generate the surface mesh, either:

- In the **Mesh Generation** toolbar, click (**Generate Surface Mesh**), or
- Select **Mesh** > **Generate Surface Mesh**.

In the *Output* window:

- The output of the surface remesher is displayed.
- When the process is complete, the message `Surface Remesher Completed` appears.

2. To view the remeshed surface:

   a) In the **Vis** toolbar, click (**Show All Meshes**).

   b) Right-click the **Scenes** > **Geometry Scene 1** node and select **Apply Representation** > **Remeshed Surface**.

   The scene is updated to display the triangulated surface for the control valve geometry.

3. Rotate the geometry as shown:



A higher surface triangulation is present in the valve wall area. This translates into a higher volume mesh density.



Refinement based on the surface curvature is also visible on the curved parts of the geometry. To access the properties of the remeshed surface, select the **Representations > Remeshed Surface** node.



Running the Polyhedral Volume Mesher

Run the polyhedral and prism layer meshers.

1. To run the polyhedral and prism layer meshers, either:
   - In the **Mesh Generation** toolbar, click (**Generate Volume Mesh**), or
   - Select **Mesh > Generate Volume Mesh**

   The *Output* window displays a summary of the mesh quality criteria (namely, the cell quality, face validity, skewness angle and volume ratio) for the core and prism meshes.
2. To view the volume mesh representation, right-click the **Scenes > Geometry Scene 1** node and select **Apply Representation > Volume Mesh**.

   The scene updates to show the polyhedral volume mesh with prismatic near wall layers:



3. Save the simulation.

   To access the properties of the volume mesh, select the **Representations > Volume Mesh > Finite Volume Regions > control_valve** node.



Note: the totals for cells, interior faces and vertices may differ slightly depending on your machine type and architecture.

Replacing the Import Surface Mesh

The original imported surface can be replaced by another surface mesh to generate a volume mesh. In this example, the valve position has been changed, so build a mesh using the same parameters as before.

In this tutorial, a surface mesh database file (.dbs) created in pro-STAR is imported. It is based on the original STL data but with the valve in a location.

1. Copy the `control_valve_mod.dbs` file from the `[INSTALL_DIR]/doc/startutorialsdata/mesh/data` directory to your working directory.
2. Right-click the **Geometry > Parts > control valve** node and select **Replace Part**.
3. In the *Open* dialog, select `control_valve_mod.dbs` and click **Open**.
4. In the *Database Import Options* dialog, set the following properties:

Property	Setting
Database ID	1

Property	Setting
Units	`in` (inches)
Match Part surfaces by Patch ID	activated

5. To import the modified surface, click **OK**.
6. To view the modified surface representation, right-click the **Scenes** > **Geometry Scene 1** node and select **Apply Representation** > **Geometry**.



As the patch definitions in the new surface are identical to the original import, the boundaries are transferred from one model to the other. The surface is ready to be remeshed.

Re-running the Mesh Generation

Update the volume mesh based on the updated geometry.

Multiple meshers can be executed in a single action.

1. To execute the surface remesher, polyhedral volume mesher and prism mesher, either:

- In the **Mesh Generation** toolbar, click (**Generate Volume Mesh**), or
- Select **Mesh** > **Generate Volume Mesh**.

When meshing is complete:

2. To view the volume mesh representation, right-click the **Scenes** > **Geometry Scene 1** node and select **Apply Representation** > **Volume Mesh**.



The previous surface remesher and volume mesh representations have been replaced with meshes for the modified valve geometry.

Trimmed Cell Meshing Example

To generate a high-quality hexahedral based mesh for the original control valve geometry, use the trimmed cell mesher instead of the polyhedral mesher. The trimmed cell mesher also improves the distribution of cells which means fewer cells are required.

First, change the meshing models to use the trimmer and prism mesher models.

1. Right-click the **Continua** > **Mesh 1** > **Models** node and select **Select Meshing Models**.
2. In the *Mesh 1 Model Selection* dialog, set the following properties:

Group Box	Property	Setting
Enabled Meshing models	*Prism Layer Mesher*	**deactivated**
Enabled Meshing models	*Polyhedral Mesher*	**deactivated**
Enabled Meshing models	*Surface Remesher*	**deactivated**
Volume Mesh box	*Trimmer*	**activated**
Optional Models	*Prism Layer Mesher*	**activated**

The surface remesher model was cleared since a good starting surface is now available from the previous session.

To access the default properties for the trimmer, select the **Continua** > **Mesh 1** > **Models** > **Trimmer** node.



The Template mesh growth type option is set to Simple, meaning that only a default (core) and surface growth rate are required.

Contents

Setting the Global Mesh Reference Values - Trimmer
Setting the Boundary Mesh Values - Trimmer
Running the Trimmed Cell Mesher

Setting the Global Mesh Reference Values - Trimmer

The trimmer produces the most optimal cells when the trimming surface contains elements of the same size as those being trimmed. For this reason, the same surface remesher values as used in the first part of the tutorial are used for the global trimmer settings in this case.

The Maximum Cell Size node controls the maximum size of the cells allowed in the core. In this instance it has been left at the default value (10000% of the Base Size) meaning that the size is relatively

unconstrained and can grow many times more than the surface size if required. If smaller cells are required in the core then this value should be reduced.

The Template Growth Rate option controls the extent to which the cells are refined. The Default Growth Rate, set to Fast by default, controls the transition rate in the interior volume mesh (away from the surface). Setting it to fast means it transitions as quickly as possible from one layer to the next, using a minimum of one cell for each transition layer. The Boundary Growth Rate, set to None by default, controls the number of layers of cells at the surface. Setting it to none means that the default growth rate value determines the number of cells to grow at the surface. If you wanted additional refined cells then you could choose one of the other options such as medium or slow, which would introduce at least 2 or 4 refined levels of cells respectively.

Setting the Boundary Mesh Values - Trimmer

Use custom surface size values for both the valve wall and the symmetry plane to control the density of the trimmed cell mesh.

In the case of the valve wall, the surface size settings are used to increase the mesh density in the vicinity of the valve area above and beyond any refinements due to proximity effects. For the symmetry plane, there is no requirement to maintain the global surface mesh density, so allow the mesh size to increase up to a maximum size equal to the base size. You could also repeat this for the inlet and outlet boundaries if desired although this is not demonstrated here.

1. Maintain the same **Regions > control_valve > Boundaries > Valve Wall** property settings as before.
2. Expand the **Regions > control_valve > Boundaries > Symmetry Plane** node.
3. Set the following properties:

Node	Property	Setting
Mesh Conditions		
Custom Surface Size	*Custom surface size*	activated
Mesh Values		
Surface Size > Relative Minimum Size	*Percentage of Base*	`100.0`

The Relative Target Size is already set to be 100% of the base size and does not need to be changed. The symmetry plane is now no longer constrained in the same way as the other surfaces in the geometry and larger cell sizes are seen on this plane.

Running the Trimmed Cell Mesher

1. To run the trimmer and prism layer meshers, either:
   - In the Mesh Generation toolbar, click (**Generate Volume Mesh**), or
   - Select **Mesh > Generate Volume Mesh**

   As before, in the *Output* window:
   - The output from the meshers is displayed.
   - When the process is complete, the message `Converting mesh into finite volume representation in Region control_valve` appears.

   The current geometry scene is automatically updated with the new volume mesh:



A close up of the trimmed cell mesh in the valve area is shown below:



A uniform cell size has been maintained around the valve area, resulting from the 4% and 8% of base size specification for the minimum and target surface size, respectively.

Summary

This tutorial demonstrated how to use the surface and volume meshing tools within STAR-CCM+ to build a single region polyhedral volume mesh with prism layers. It also demonstrated use of the trimmed mesher.

This tutorial introduced the following features:

- Importing an STL type surface data file
- Splitting and renaming surfaces
- Adding feature curves for important features
- Selecting meshing models and define reference properties
- Defining local boundary meshing properties

- Running the surface remesher, polyhedral and prism layer meshers
- Displaying representations
- Replacing the import surface with a modified geometry
- Remeshing a modified geometry using one button press
- Running the trimmer mesher

Region-Based Meshing Multi-Region Heat Exchanger

This tutorial demonstrates how to use the surface and volume meshing tools while working with multiple regions within STAR-CCM+ .

Starting from a Parasolid Transmit geometry, build a multi-region polyhedral volume mesh with prism layers, exported from a CAD package. Define the appropriate boundaries, create a mesh for each region, and a conformal mesh at interfaces between regions.

The heat exchanger unit to be meshed contains multiple fluid flow passages connected by two end chambers, one of which is split to force the flow around the system. The flow domain is contained inside a solid cylinder, which is hollow on the inside. Two separate bodies are contained in the single Parasolid Transmit file. These are the fluid and solid parts of the geometry.

The heat exchanger geometry is shown below. The case is symmetrical so import and mesh half the geometry, only.



Contents

Prerequisites
Importing the Surface Mesh
Defining the Region Names and Types
Defining the Boundaries
Defining the Interface
Selecting the Meshing Models
Setting the Reference Values
Including Prism Layers on the Interface
Setting Region Mesh Options
Running the Surface Remesher
Running the Polyhedral Volume Mesher
Summary

Prerequisites

The instructions in the Region-Based Meshing Multi-Region Heat Exchanger tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Surface Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/mesh/data`.
5. Select `hx_unit_asm.x_t` then click **Open**.

6. In the *Import Surface Options* dialog, set *Import Mode* to **Create new Region**.
7. To accept the remaining default settings, click **OK**.
8. Save the simulation as `heatExchanger.sim`.

Defining the Region Names and Types

During the import process, the Parasolid Transmit file geometry is translated into a triangulated surface representation with a total of two regions and one boundary surface per region. Feature curves are included for each region based on the supplied angle.

Using the above combinations, the import process results in two regions that are distinct and separate from one another and do not share any surfaces.

Review the regions created during the import process.

1. Expand the **Regions** node.

   Two regions, **Body 1** and **Body 2**, were created during the import process.
2. Rename **Body 1** node to `Fluid body`.
3. Rename **Body 2** node to `Solid body`.
4. Select the **Solid body** node and set **Type** to **Solid Region**.

Defining the Boundaries

Identify the flow, symmetry, and interface boundaries so that appropriate meshes for each can be generated.

The CAD import resulted in one boundary per region. CAD information distinguishing individual boundary surfaces is stored as patches. A boundary surface can be split into constituent boundaries using patch information.

Contents

Extracting the Fluid Boundaries
Renaming and Setting the Fluid Boundary Types
Extracting the Solid Boundaries
Renaming and Setting the Solid Boundary Types

Extracting the Fluid Boundaries

Split the fluid region boundaries.

1. Expand the **Regions** > **Fluid body** > **Boundaries** node.
2. Right click the **Boundary 1** node and select **Split by Patch**.

   The *Split Boundary by Patch* dialog appears. Surfaces that belong to the fluid region are color-coded according to their patch definition. All other surfaces are translucent:



   you can use the *Split Boundary by Patch* dialog to select a patch or patches and create a boundary from the original. Select the patch directly on the model by clicking the appropriate face; the corresponding patch number is highlighted in the dialog.

Create the following boundaries for the fluid region:

3. In the *Split Boundary by Patch* dialog, select patch number 12.

   The associated boundary is highlighted in the model.
4. Set **Boundary Name** to `Inlet` and click **Create**.

   The newly extracted boundary patch is removed from the list and from the model display.
5. Repeat the above process to create the following boundaries:

Patch Number	Boundary Name
18	`Outlet`
26	`Symmetry`

6. In the *Split Boundary by Patch* dialog, click **Close**.

   Inlet, **Outlet**, and **Symmetry** nodes appear under the **Fluid body** > **Boundaries** node. The original boundary for the fluid region (**Boundary 1**) is now composed of surfaces that belong solely to the interface plane between the fluid and solid regions. This makes the interface definition easier later on by having all the surfaces in one boundary.

Renaming and Setting the Fluid Boundary Types

Set appropriate boundary types for each boundary.

1. Expand the **Regions** > **Fluid body** > **Boundaries** node.
2. Rename the **Boundary 1** node to `Interface`.
3. Set the following properties:

Node	Property	Setting
Interface	*Type*	**Wall**
Inlet	*Type*	**Velocity Inlet**
Outlet	*Type*	**Pressure Outlet**
Symmetry	*Type*	**Symmetry Plane**

Extracting the Solid Boundaries

Extract the solid boundaries using the Split by Patch technique.

1. Right-click the **Regions** > **Solid body** > **Boundaries** > **Boundary 1** node and select **Split by Patch**.

   The *Split Boundary by Patch* dialog appears. Surfaces that belong to the solid region are color-coded according to their patch definition. All other surfaces are translucent, as shown below:



2. In the *Split Boundary by Patch* dialog, create the following boundaries for the solid region:

Patch Number	Boundary Name
151,152,153	`Outer cylinder wall`
173,174	`Symmetry`
175,176,177	`Inner cylinder wall`

3. Click **Close**.

   The boundary patches are removed from the list and from the model display. **Outer cylinder wall**, **Symmetry**, and **Inner cylinder wall** nodes appear under the **Solid body > Boundaries** node. The original boundary for the solid region (Boundary 1) is now composed of surfaces that belong solely to the interface plane between the solid and fluid regions.

Renaming and Setting the Solid Boundary Types

Set appropriate boundary types for each boundary.

1. Expand the **Regions > Solid body > Boundaries** node.
2. Rename the **Boundary 1** node to `Interface`.
3. Set the following properties:

Node	Property	Setting
Interface	*Type*	**Wall**
Inner cylinder wall	*Type*	**Wall**
Outer cylinder wall	*Type*	**Wall**
Symmetry	*Type*	**Symmetry Plane**

Defining the Interface

An in-place interface is required to link together the two regions so that a conformal mesh is created at the common interface planes during the surface and volume meshing process. If you do not create an interface in this example, then the surface and volume meshing does not proceed correctly due to the close proximity of neighboring boundaries.

1. Multi-select the following nodes:
   - **Regions > Fluid body > Boundaries > Interface**
   - **Regions > Solid body > Boundaries > Interface**
2. Right-click one of the selected nodes and select **Create Interface**.

   Three interface nodes appear in the simulation tree:
   - **Regions > Fluid body > Boundaries > Interface [Interface 1]**
   - **Regions > Solid body > Boundaries > Interface [Interface 1]**
   - **Interfaces > Interface 1**



The first two nodes compose the two halves of the in-place internal boundary interface between the fluid and solid regions.

Selecting the Meshing Models

Select the surface remesher, polyhedral volume mesher, and prism layer mesher.

1. Right-click the **Continua** > **Mesh 1** > **Models** node and select **Select Meshing Models**.
2. In the *Mesh 1 Model Selection* dialog, select the following models:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**

3. Click **Close**.

The default properties for these models are acceptable for this simulation.

Setting the Reference Values

Set the reference values to generate the mesh.

For this example, absolute size specification is used for the prism layer thickness and surface size values. Although the model dimensions are in meters (automatically set during the import process of the Parasolid Transmit file), you can use any units to specify the inputs for the meshers.

Increase the # Points in a gap value for Surface Proximity so that an increased surface triangulation density is obtained in the vicinity of the pipes and inner cylinder surface, which in turn results in a higher concentration of volume cells.

1. Expand the **Continua > Mesh 1 > Reference Values** node.
2. Set the following properties:

Node	Property	Setting
Number of Prism Layers	*Number of Prism Layers*	`1`
Prism Layer Thickness	*Size type*	**Absolute**
Absolute Size	*Value*	`2.0 mm`
Surface Proximity	*# Points in gap*	`4.0`
	Search Floor	`0.0 (default)`
Surface Size	*Relative/Absolute*	**Absolute**
Absolute Minimum Size	*Value*	`4.0 mm`
Absolute Target Size	*Value*	`8.0 mm`

Including Prism Layers on the Interface

The generation of an in-place interface between the fluid and solid regions created a node in the Interfaces node in the simulation tree. This node can control the mesh conditions on either side of the interface.

By default, the option to grow prism layers on either side of the interface is set to off. In your case, you want prism layers on the fluid side only, since there is no requirement to include layers for the solid

material region. To achieve this, activate the prism layer generation option for the interface but deactivate the generation of layers on the region level for the solid body.

1. Expand the **Interfaces** > **Interface 1** > **Mesh Conditions** node.
2. Select the **Interface Prism Layer Option** node and activate **Grow Prisms from Interface**.

   Activating this option means that prism layers are grown from both sides of the interface.

Setting Region Mesh Options

The option to activate or deactivate prism layer generation is available on a region wise basis.

To deactivate the prism layers for the solid region:

1. Expand the **Regions** > **Solid body** > **Mesh Conditions** node.
2. Select the **Customize Prism Mesh** node and set **Customize Prism Mesh** to **Disable**.

   The generation of prism layers for the boundaries in the solid region is now deactivated. However, if you decided to specify custom values for a boundary belonging to this region then it would override the region level specification. There is no requirement to do this for this example.

Running the Surface Remesher

Since the surface remesher and volume mesher models have been selected in tandem, you could proceed and generate both meshes in one go. However, for training purposes, first generate the surface remesher mesh and view it before proceeding with the volume meshing.

1. To generate the surface remesher mesh, either:
   - In the Mesh Generation toolbar, click (**Generate Surface Mesh**), or
   - Select **Mesh** > **Generate Surface Mesh**.

   The output of the surface remesher is shown in the Output window and status shown in the status bar. When the process is complete, the message Surface Meshing Pipeline Completed is shown in the Output window.

To view the remeshed surface representation:

2. Expand the **Scenes** node.
3. Right-click the **Geometry Scene 1** node and select **Apply Representation > Remeshed Surface**.

   The scene updates to display the triangulated surface for the heat exchanger unit geometry.



Surface curvature and proximity-based refinement is visible on the curved parts (that is, the pipes) of the geometry. The properties of the remeshed surface can be seen by clicking the appropriate node in the Representations node.

Running the Polyhedral Volume Mesher

Run the polyhedral and prism layer volume meshers.

1. To run the polyhedral and prism layer meshers, either:
   - In the **Mesh Generation** toolbar, click (**Generate Volume Mesh**), or
   - Select **Mesh > Generate Volume Mesh**.

   As before, the output from the meshers is shown in the *Output* window and status updated in the status bar. When the process is

complete, the message `Volume meshing Pipeline Completed` is shown in the *Output* window. A summary of the mesh generation process for the core and prism meshes can be viewed by scrolling up the *Output* window.

To view the volume mesh representation:

2. Expand the Scenes node.
3. Right-click the **Geometry Scene 1** node and select **Apply Representation > Volume Mesh**

   The scene updates to show the polyhedral volume mesh with a prismatic near wall layer for the fluid region:



A close up of the mesh in the vicinity of the inlet plenum is shown below:



To view the properties of the fluid volume mesh, including the number of cells in the fluid region:

4. Expand the **Representations** > **Volume Mesh** > **Finite Volume Regions** node.
5. Select the **Fluid body** node.



Similarly, to view the properties of the solid body volume mesh:

6. Select the **Solid body** node.



The exact values for cells, interior faces, and vertices differ to the above depending on the machine precision.

Summary

This tutorial demonstrated how to use the surface and volume meshing tools while working with multiple regions within STAR-CCM+.

The steps covered were:

- Importing a Parasolid Transmit type geometry data file.
- Demarcating and renaming boundaries required for multi-region meshing.
- Defining an in-place interface.
- Selecting meshing models and define reference properties.
- Assigning region and interface meshing properties.
- Running the surface remesher, polyhedral, and prism layer meshers.
- Displaying representations.

Directed Meshing: Electrical Machine

This tutorial demonstrates how to use the directed meshing tool in STAR-CCM+. Directed meshing allows you to generate a high-quality structured mesh by creating quadrilateral elements on a surface of a body and then sweeping them through its volume.

In this tutorial, you are given a section of an electrical machine comprising five parts: rotor, magnet, air gap, slot, and stator. The parts are all connected and have regular section profiles, making this geometry a good case to use with directed meshing.



Contents

Prerequisites
Loading the Simulation
Checking Suitability for Directed Meshing
Assigning Parts to Regions
Using Patch Meshing
Using Existing Meshes
Summary

Prerequisites

The instructions in the Directed Meshing tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Loading the Simulation

For this tutorial, you are provided with a simulation file that contains pre-defined simulation objects.

Load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/mesh/data`.
5. Select `directedMeshCAD.sim` then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `directedMeshTutorial.sim`.

Checking Suitability for Directed Meshing

The directed meshing tool takes a surface mesh that is generated on a source surface and sweeps it through the volume to the target surface. The tool uses the remaining surfaces to guide the mesh. This means that each part you choose to use in directed meshing must have at least three distinct surfaces. You can make sure that the parts satisfy this requirement by looking through the defined surfaces of each geometry part.

1. Create a geometry scene.

View the CAD bodies and parts:

2. Expand the **Geometry** > **3D-CAD Models** > **3D-CAD Model 1** > **Bodies** node.
3. To view the parts that are shown in the geometry scene, expand the **Geometry** > **Parts** node.

Each part is associated with a CAD body. To use directed meshing, associate each part with exactly one body.

4. Expand all **Geometry** > **Parts** nodes.
5. Multi-select all **[Part]** > **Surfaces** > **[source surface]** nodes.

   The selected surfaces are highlighted in the *Geometry Scene*.



The subdivision of part surfaces has been chosen to demonstrate the two methods available in the Directed Mesh tool for creating quadrilateral meshes on the source surfaces:

- Using manually created patches
- Using an extracted surface of an existing volume mesh

The **Stator Steel.source**, **Slot.source** and **Air.source** faces are meshed using the patch meshing option. They can be meshed together since the source surfaces are connected and lie on the same plane. The patch mesh is then swept through the volume of the

Stator Steel, Slot, and Air parts to produce a volume mesh for these parts.

The **Magnet.source** and **Rotor.source** faces are meshed using an existing mesh of a part surface they are in contact with. The **Air.bottom** surface mesh is mapped to the **Magnet.source** surface, allowing you to generate a volume mesh for the Magnet part. The **Magnet.target** surface mesh is then mapped to the **Rotor.source** surface. Note the **Air.bottom** surface has been separated out from the other Air surfaces to represent the surface that is in contact with the **Magnet.source** surface. This breakdown of surfaces is a requirement if you want to map an existing mesh to a source surface mesh -- both parts in contact must have a separate part surface for the interface.

Assigning Parts to Regions

Since a volume mesh cannot be generated for parts that are not associated with a region, assign the parts to regions.

1. Multi-select all **Geometry** > **Parts**.
2. Right-click one of the selected parts and select **Assign Parts to Regions**.
3. In the Assign Parts to Regions dialog, select the following options:
   - **Create a Region for Each Part**
   - **Create One Boundary for All Part Surfaces**
4. Click **Apply** then **Close**.
5. Save the simulation.

Using Patch Meshing

You can mesh the Air, Slot, and Stator Steel parts at the same time using patch meshing. Patch meshing requires you to create patches on the source surface manually. STAR-CCM+ uses these patches to generate a quadrilateral surface mesh. The three parts have contacts that are defined between them, allowing them to be used together in one directed mesh operation.

1. Right-click the **Geometry** > **Operations** node and select **New** > **Directed Mesh**.
2. In the *Create Directed Mesh Operation* dialog, select the **Air**, **Slot**, and **Stator Steel** parts.

3. Click **OK**.
4. Right-click the **Geometry > Operations > Directed Mesh** node and select **Edit**.

   The *Directed Mesh* panel that is used for all directed mesh operations appears. Only the selected parts are shown in the *Graphics* window.



5. In the *Directed Mesh 1* window, select the three target surfaces:
   - **Stator Steel.target**
   - **Slot.target**
   - **Air.target**
6. Right-click the highlighted surfaces and select **Part Surfaces > Add To Target** .



7. In the *Directed Mesh 1* window, select the three source surfaces:
   - **Stator Steel.source**

- **Slot.source**
- **Air.source**

8. Right-click the highlighted surfaces and select **Part Surfaces > Add To Source**.

   The source surfaces are grouped together in the Properties of the **Directed Mesh > Source Surfaces** node and are colored green in the *Graphics* window. The target surfaces are grouped in the **Target Surfaces** node and are colored red. The surfaces are shown in the following image with the scene transparency turned on.

Contents

Creating a Patch Topology
Creating a Patch Mesh
Generating a Volume Mesh

Creating a Patch Topology

To use patch meshing, you create patches on the source surfaces. A closed loop of four curves define the patches.

1. Right-click the **Source Meshes** node and select **New Source Mesh** > **Patch Mesh**.
2. In the *Part Collection for this source mesh* dialog, select **Connected Part** and click **OK**.

   The *Edit* panel that is used to create patches and a patch mesh appears.



   The *Graphics* window displays the source surface in gray with feature curves outlined in yellow.
3. Click (**Save-Restore-Select Views**) and select **Views** > **+Z** > **Up +Y**.
4. In the *Create/Edit Patch Entities* box, click (**Auto-populate feature edges with patch curves**).

   This step creates patch curves on feature edges.



This patch topology is not in a valid state for patch meshing. Patches are composed of four curves and appear in green when the topology is valid. To generate proper patches, split some of the current patch curves:

5. In the *Create/Edit Patch Entities* box, click (**Split the selected patch curve**) and maintain the split patch curve option, **At Picked Location**.
6. In the *Graphics* window, following the number points shown below, click the existing feature curves.

Note: you can move vertices at any point in the patch meshing process.

Connect these additional vertices to form complete patches.

7. In the *Create/Edit Patch Entities* box, click (**Create Patch Curve**).
8. Click the vertex point marked “2” in the previous screenshot.

The red lines passing through the point indicate that the point has the “focus”, meaning that the next curve starts from this point.



9. Click the vertex point “6”. A patch curve is created to connect the two points.



10. Continue to create additional patch curves between vertices as shown in the following screenshot.

    To change the focus point as you create the curves, click a blank area in the *Graphics* window to deselect an old focus point before selecting a new focus point.

Note: If you make a mistake in creating a curve, you can delete a patch curve.



Creating a Patch Mesh

Once a patch topology is fully defined, STAR-CCM+ enables the Patch Mesh mode. You define the mesh distribution properties of the patch mesh within this mode.

1. In the *Edit* panel, set **Mode** to **Patch Mesh**.
2. In the *Graphics* window, multi-select the set of curves that are shown below.



3. In the *Mesh Properties* box, set **Number of Divisions** to `14` and click **Apply**.
4. Select a second set of patch curves as shown below. The patch curves for which mesh properties have been defined appear in black.



5. In the *Mesh Properties* box, set **Number of Divisions** to `10` and click **Apply**.
6. Select the remaining short patch curves. Set **Number of Divisions** to `4` and click **Apply**.



The source mesh is shown in the *Graphics* window.

7. **Close** the *Edit* panel.

Generating a Volume Mesh

The next step is to generate a volume mesh using the patch mesh. This step is done by defining a volume mesh distribution for the connected part.

1. Right-click the **Mesh Distributions** node and select **New Volume Distribution**.
2. In the *Parts for creating a mesh distribution* dialog, select **Connected Part** and click **OK**.
3. Select the **Mesh Distributions > Volume Distribution > Number of Layers** node and set **Number of Layers** to `24`.
4. Right-click the **Directed Mesh** node and select **Execute**.

   The meshed parts are shown in the *Graphics* window.



5. To return to the simulation tree, click **Close Directed Mesh** .

Using Existing Meshes

Use the volume mesh that is generated from the patch topology to generate a source mesh for the remaining Magnet and Rotor parts.

Note that the Air, Magnet and Rotor parts have contacts defined between them.

1. As with patch meshing, right-click the **Geometry** > **Operations** > node and select **New** > **Directed Mesh**.
2. In the *Create Directed Mesh Operation* dialog, select the **Magnet** part and click **OK**.

   Only the Magnet and Rotor parts appear in the dialog as the other three parts are already associated with a directed mesh.
3. Right-click the **Operations** > **Directed Mesh 2** node and select **Edit**.
4. In the *Directed Mesh 1* window:

   a) Right-click the top surface of the Magnet part (Magnet:source) and select **Part Surface (source)** > **Add to Source**.

   b) Right-click the bottom surface of the Magnet part (Magnet.target) and select **Part Surface (target)** > **Add to Target**.

   The source surface is displayed in green and the target surface in red.



5. Right-click the **Source Meshes** node and select **New Source Mesh** > **Use Existing Mesh**.
6. In the *Part Collection For This Source Mesh* dialog, select **Magnet**.

   STAR-CCM+ maps the base of the volume mesh for the Air part to the source surface of the Magnet part. It adds an Existing Mesh node, **Magnet**, to Mesh Settings.
7. Right-click the **Mesh Distributions** node and select **New Volume Distribution**.
8. In the *Parts for creating a mesh distribution* dialog, select **Magnet** and click **OK**.
9. Select the **Mesh Distributions** > **Volume Distribution** > **Number of Layers** node and set **Number of Layers** to `12`.
10. Right-click the **Directed Mesh 2** node and select **Execute**.



11. Click **Close Directed Mesh** to return to the simulation tree.

Repeat the process for the Rotor part:

12. Create a Directed Mesh operation.
13. Edit the **Directed Mesh 3** node and select the source and target surfaces.
14. Create a source mesh using an existing mesh. Select the **Rotor** part as the part collection for the source mesh.
15. Create a volume distribution and set the **Number of Layers** to `12`.
16. Execute the directed mesh operation.
17. When the volume mesh for the Rotor part is generated, click **Close Directed Mesh**.
18. Create a Mesh Scene.

The final mesh is displayed below:



Summary

This tutorial demonstrated how to use the directed meshing tool in STAR-CCM+.

The steps covered were:

- Importing and checking CAD geometry for the directing mesh operation.
- Launching a directed mesh operation.
- Identifying source and target surfaces.
- Creating a patch topology for the source surface mesh.
- Using a previously defined volume mesh for the source surface mesh.
- Defining a volume mesh distribution and generating a volume mesh.

Incompressible Flow

The tutorials in this set illustrate various STAR-CCM+ features for incompressible fluid flows as well as porosity and solution recording

Contents

Steady Flow: Lid-Driven Cavity Flow
Steady Flow: Channel Flow with Multiple Meshes
Steady Flow: Laminar and Turbulent in an S-Bend
Steady Flow: Backward Facing Step
Adjoint Flow Solver: External Flow over a Dual Element Wing
Adjoint Flow: Wing Shape Optimization
Steady Multi-Component Flow: Dilution Pipe
Unsteady Multi-Component Flow: Dilution Pipe
Porous Resistance: Isotropic Media
Porous Resistance: Orthotropic Media
Solution Recording and Playback: Vortex Shedding

Steady Flow: Lid-Driven Cavity Flow

This tutorial demonstrates the performance of STAR-CCM+ in solving a traditional square lid-driven cavity flow. The problem geometry consists of a two-dimensional 1 m^2 cavity. The cavity is covered by an impermeable wall that moves in the x-direction with constant velocity of 1 m/s. A stretched quadrilateral mesh with 6400 cells is used.



Material properties:

Density (kg/m^3)	1.0
Dynamic viscosity (Pa.s)	0.0002

Boundary conditions:

Top wall	U = 1.0 (m/s), V=0.0
Bottom wall	U = V = 0.0
Left wall	U = V = 0.0

Right wall	U = V = 0.0

Initial conditions:

U, V velocity components (m/s)	0.0
Operating pressure (Pa)	0.0

The Reynolds number based on the width of the cavity, top wall velocity, fluid density and fluid dynamic viscosity is 5000. The calculation is performed using the Segregated Flow model. Both the U- and V-velocity profiles along lines passing through the center of the cavity are compared with data from the literature [Ghia].

Contents

Importing the Mesh and Naming the Simulation
Visualizing the Imported Mesh
Setting Up the Physics Models
Setting Fluid Properties
Setting Boundary Conditions and Values
Visualizing the Solution
Setting Up Lines for Plotting
Plotting Simulation Data
Plotting Reference Data
Preparing Solver Parameters and Stopping Criteria
Initializing and Running the Simulation
Analyzing the Results
Summary
Bibliography

Importing the Mesh and Naming the Simulation

Create a new simulation and set up the mesh.

A quadrilateral-cell mesh has been prepared for this analysis and is stored in a `.ccm` format file. Import it into the simulation:

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the New Simulation option from the menu bar. .

   STAR-CCM+ creates a window for the new simulation in the *Explorer* pane and give it the default name Star 1

2. Select **File** > **Import** > **Import Volume Mesh**from the menu bar.
3. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select the file `cavityQuad.ccm`.
4. Click the **Open** button to start the import.

   STAR-CCM+ provides feedback on the import process in the *Output* window. A mesh region that is named **cell-2** is created under the **Regions** node to represent the solution domain.
5. Save the new simulation as `cavityQuad.sim`.

Visualizing the Imported Mesh

View the imported mesh, and access various parts that are contained in the region.

Use the following steps:

1. Right-click the **Scenes** node and then select **New Scene** > **Mesh**.

   The mesh appears in a scene in the *Graphics* window.
2. The edges of the square are boundaries. Click one of them.

   The selected boundary is highlighted and a label with its name appears on the display.



In the *cavityQuad* window, the node that corresponds to this boundary is also highlighted.



Similarly, if any other boundary node is selected in the *cavityQuad* window, the part corresponding to that boundary is highlighted in the *Graphics* window.

Setting Up the Physics Models

The default continuum (**Physics 1**) is edited so that the appropriate physical models for the simulation are selected.

To select physical models for this continuum:

1. For the physics continuum **Continua** > **Physics 1**, select the following models, in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected)
Material	**Liquid**
Flow	**Segregated Flow**
	Gradients (Selected automatically)
Equation of State	**Constant Density**
Time	**Steady**
Viscous Regime	**Laminar**

2. Click `Close`.

The color of the **Physics 1** node has turned from gray to blue to indicate that models have been activated.



3. Save the simulation.

Setting Fluid Properties

Set the density and dynamic viscosity fluid properties.

Use the following steps:

1. Expand the **Continua** > **Physics 1** > **Models** > **Liquid** > **H2O** > **Material Properties** > **Density** node.
2. Select the **Constant** node and set the *Value* to `1.0 kg/m^3`.
3. Select the **H2O** > **Material Properties** > **Dynamic Viscosity** > **Constant** node and set the *Value* to `2.0E-4 Pa-s`.

Setting Boundary Conditions and Values

This problem requires only two boundary conditions: moving and stationary. In STAR-CCM+, it is possible to combine two or more boundaries into one.

Use the following steps:

1. Open the **Regions** > **cell-2** > **Boundaries** node.
2. Multi-select the boundary nodes **wall-3**, **wall-4**, and **wall-7**.



3. Right-click and select **Combine**.

   This creates a single boundary, wall-3, that consists of all sides of the square except the top.
4. Rename the **wall-3** node to `Stationary Wall` .
5. Rename the **wall-5** node to `Moving Wall`.

For the moving wall boundary, set the velocity of 1 m/s in the x-direction:

6. Select the **Moving Wall** > **Physics Conditions** > **Tangential Velocity Specification** node and set the *Method* to be `Vector`.
7. Select the **Physics Values** > **Relative Velocity** > **Constant** node and enter `1,0,0` in the *Value* property.
8. Save the simulation.

Visualizing the Solution

Post-processing scenes can be updated every iteration (or time-step) in STAR-CCM+. This allows for dynamic monitoring of various quantities, such as the velocity that you want to visualize in this tutorial. A vector scene is therefore created before running.

Use the following steps:

1. Right-click the **Scenes** node and then select **New Scene** > **Vector**.

   A new *Vector Scene 1* display appears.

The vector bar shows the default setting of `Velocity`.

2. Click and drag the vector bar toward the right side of the display until it switches to a vertical orientation.



3. The edges of the bar can be dragged to make the label larger. Move the square to make more room for the vector bar.
4. To verify that the scene has the entire region selected for the data, select the **Scenes** > **Vector Scene 1** > **Displayers** > **Vector 1** > **Parts** node.

The *Parts* property shows `cell-2`, the name of the region.

Setting Up Lines for Plotting

Data are plotted for both the U- and V-velocity profiles along line probes passing through the center of the cavity.

To create these lines:

1. Make sure that at least one scene display is active in the *Graphics* window.
2. Right-click the **Derived Parts** node and select **New Part** > **Probe** > **Line**

   The *Create Line Probe* dialog appears in the *Explorer* pane to specify the desired line.



3. Set the following properties for the U line:

Property	Value
Input Parts	`cell-2`
Point 1	`[0.5, 0.0, 0.0]`
Point 2	`[0.5, 1.0, 0.0]`
Resolution	`50`
Display	**No Displayer**

4. Click **Create**, then click **Close**.

   A new derived part named **line-probe** is created in the **Derived Parts** manager node.
5. Rename the node to `Simulation (U)`

Next, create the V line probe. This can be done simply by making a copy of the existing derived part and modifying it.

6. Right-click the **Simulation (U)** node and select **Copy**.
7. Right-click the **Derived Parts** node and select **Paste**.
8. Rename the new node to `Simulation (V)`
9. Enter `0,0.5,0` for the *Point 1* property.
10. Enter `1,0.5,0` for the *Point 2* property.
11. Save the simulation.

Plotting Simulation Data

This part of the tutorial produces the plots of the U- and V-velocity profiles. In a subsequent step, experimental data are imported from files and plotted alongside the numerical data. The development of the solution can be observed throughout the run by viewing these plots.

To create the U-velocity plot:

1. Right-click the **Plots** node and select **New Plot > X-Y**.
2. Rename the **XY Plot 1** node `U-Velocity Profile`
3. Set the following properties:

Node	Property	Value
U-Velocity Profile	*Parts*	**Derived Parts > Simulation (U)**
X Type	*Type*	**Direction**
Vector Quantity	*Value*	`[0,1,0]`
Y Types > Y Type 1	*Type*	**Scalar**
Scalar Function	*Field Function*	**Velocity > Lab Reference Frame > Laboratory > i**

Node	Property	Value
Simulation (U) > Line Style	*Style*	(**Solid**)
Symbol Style	*Shape*	`None`
Axes > Y Axis	*Minimum*	`-0.6`
	Maximum	`1.0`
Major Labels	*Spacing*	`0.2`

To create the V-velocity plot:

4. Right-click the **Plots** node and select **New Plot > X-Y**.
5. Rename this **XY Plot 1** node `V-Velocity Profile`
6. Select the `V-Velocity Profile` node and click the right-hand side of the *Parts* property.

Node	Property	Value
V-Velocity Profile	*Parts*	**Derived Parts > Simulation (V)**
Y Types > Y Type 1	*Type*	**Scalar**
Scalar Function	*Field Function*	**Velocity > Lab Reference Frame > Laboratory > j**
Simulation (V) > Line Style	*Style*	(**Solid**)
Symbol Style	*Shape*	`None`
Axes > Y Axis	*Minimum*	`-0.7`
	Maximum	`0.5`
Major Labels	*Spacing*	`0.2`

7. Save the simulation.

Plotting Reference Data

To compare the experimental velocity data with the velocity plots, it is necessary to import the data in supplied files into STAR-CCM+ tables.

Use the following steps:

1. Open the **Tools** node.
2. Right-click the **Tables** node and select **New Table > File**.
3. In the *Open* dialog navigate to the `doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory, set the File Type drop-down list to `(*.xy)`, and select the file `UVelocity.xy`.
4. Click **Open**.

   The data in that file is now stored in a table named **UVelocity**.

To display the data on the U-velocity plot:

5. Return to the **U-Velocity Profile** node.
6. Right-click the **U-Velocity Profile > Data Series** node and select **Add Data**.
7. In the *Add Data Providers to Plot* dialog, select **UVelocity** and click **OK**.
8. Select the **U-Velocity Profile > Data Series > UVelocity** node and set *Legend Name* to `Reference Data (U)`.
9. Make sure that `column 0` is set for the *X column* and `column 1` is set for the *Y column*.
10. Select the **UVelocity > Line Style** node and make sure that the *Style* property is set to `None`.
11. Select the **Symbol Style** node and set the following properties:

Property	Value
Size	`6`
Color	`blue`
Shape	`Filled Triangle`

    The supplied experimental data appears in the U-velocity plot. Numerical data is shown once the simulation run has started.

Use the same techniques for the V-velocity plot:

12. Right-click the **Tables** node and select **New Table > File**.

13. In the *Open* dialog, select the file `VVelocity.xy` and click **Open**.
14. In the **V-Velocity Profile > Data Series** node, add data and select **VVelocity** in the *Add Data Providers to Plot* dialog.
15. Select the **V-Velocity Profile > Data Series > VVelocity** node and set *Legend Name* to `Reference Data (V)`.
16. Make sure that `column 0` is set for the *X column* and `column 1` is set for the *Y column*.
17. Select the **VVelocity > Line Style** node and make sure that the *Style* property is set to `None`.
18. Select the **Symbol Style** node and set the following properties:

Property	Value
Size	`6`
Color	`blue`
Shape	`Filled Triangle`

Both plots show experimental data alongside the numerical data.

19. Save the simulation.

Preparing Solver Parameters and Stopping Criteria

Before running the analysis, look at the velocity solver, as well as the maximum number of iterations to be performed.

Use the following steps:

1. Select the **Solvers > Segregated Flow > Velocity** node and make sure that the *Under-Relaxation Factor* is `0.7`.



2. Select the **Stopping Criteria** > **Maximum Steps** node and make sure that the *Maximum Steps* value is `1000`.

Initializing and Running the Simulation

Initialize the solution and run the analysis.

Use the following steps:

1. Click (**Initialize Solution**) in the Solution toolbar or use the **Solution** > **Initialize Solution** menu item. The initial zero velocity condition is shown in the plot displays.
2. Click (**Run**) in the Solution toolbar to run the analysis.

   The Residuals display is created automatically and shows the progress of the solver. If necessary, click the *Residuals* tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide.

   During the run, it is possible to stop the process by clicking (**Stop**) in the toolbar. The simulation can be continued later by clicking (**Run**). If left alone, the simulation continues until the stopping criteria are satisfied.
3. After the run is finished, save the simulation.

Analyzing the Results

Once the calculations are finished, the displays show the results based on the pre-processing settings.

Use the following steps:

1. Click the *Vector Scene 1* display tab to view the vectors.



The plots also contain numerical results that can be compared to the experimental data:

2. Click the *U-Velocity Profile* display tab:



3. Click the *V-Velocity Profile* display tab:



Summary

This tutorial solved a traditional square lid-driven cavity flow.

The tutorial demonstrated the following STAR-CCM+ features:

- Defining physical models.
- Defining material properties that are required for the selected models.
- Defining boundary conditions.
- Reviewing solver parameters and stopping criteria.
- Monitoring the solution progress.
- Running the solver until the residuals are satisfactory.
- Analyzing results using STAR-CCM+'s visualization and XY-plot facilities.

Bibliography

Ghia, U., Ghia, K.N., and Shir, C.T. 1982. "High-Re solution for incompressible flow using the Navier-Stokes equations and a multigrid method", *J. Comp. Phys.*, **48**, p. 387.

Steady Flow: Channel Flow with Multiple Meshes

This tutorial describes a steady airflow simulation over an obstacle that is mounted at the bottom wall of a three-dimensional channel.

The problem geometry is shown below.



The case is based on an investigation from the Danish Maritime Institute, whose experimental data is compared with the STAR-CCM+ predictions. In the experimental set-up, the flow velocity and turbulence characteristics were non-uniform across the inlet section. However, to simplify the modeling procedure, the case that is presented here is based on a parabolic inlet velocity profile with fixed turbulence characteristics.

The boundaries consist of an inlet, an outlet, a symmetry plane on one of the two 'side' faces and default no-slip walls. The fluid is air and its physical properties are assumed constant and equal to the STAR-CCM+ default values.

Air enters the solution domain at standard pressure and temperature (1 bar and 293 K) with a free-stream velocity of 1.17 m/s. Based on the height of the obstacle, the Reynolds number is 3,115. The turbulence kinetic energy and dissipation rate at the inlet are set to 0.024 m^2/s^2 and 0.07 m^2/s^3, respectively, these values having been derived from experimental data. All fluid mass entering the solution domain exits through the outlet. The flow is isothermal, turbulent, and

incompressible; turbulence is simulated using the standard linear k-e model that is combined with the Wolfstein two-layer model.

The same problem is modeled three times, each variant using a different mesh, as described in:

- Hexahedral Mesh
- Tetrahedral Mesh
- Polyhedral Mesh

Comparisons are then drawn between the accuracy and computational speed of each type of mesh. All three meshes are based on the same initial shell surface defining the problem geometry but each was generated using a different technique.

Contents

Hexahedral Mesh
Tetrahedral Mesh
Polyhedral Mesh
Summary

Hexahedral Mesh

The first mesh for this tutorial is hexahedral. Create a new simulation and import the mesh.

Use the following steps:

1. Create a directory for the tutorial called `meshComparison`.
2. Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the **New Simulation** option from the menu bar.

Continue by importing the hexahedral mesh and naming the simulation. The grid has been saved in the `.ccm` file format, which contains all necessary cell, vertex, and boundary information for the problem geometry.

3. Select **File > Import > Import Volume Mesh** from the menu bar.
4. In the *Open* dialog, navigate to the `[INSTALL_DIR]/doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select file `hex.ccm`.



5. Click **Open** to start the import.

STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.

6. Finally, save the new simulation under file name `hexMesh.sim`.

Contents

Creating a Macro
Setting Up the Physics Models
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Stopping Criteria
Visualizing the Solution
Setting up Line Surfaces for Plotting
Plotting Simulation Data
Initializing and Running the Simulation
Visualizing the Results
Examining Plots
Exporting Simulation Data

Creating a Macro

Record a macro file to avoid having to repeat the steps. Use this macro later to set up the cases for the other two meshes.

The Tetrahedral Mesh and Polyhedral Mesh parts of this tutorial that follow require the same problem set-up.

1. Begin recording the macro by clicking ● (**Start Recording**) in the Macro toolbar.
2. In the *Save* dialog, enter `setup.java` as the **File Name** and click **Save**.

All actions you perform from now until the recording is stopped are recorded in the macro.

Note: Do not save the simulation while the macro is being recorded to prevent later simulations being saved under incorrect names.

Setting Up the Physics Models

Select the physics models for the gas.

The boundaries read in from file `hex.ccm` already have suitable names. Rename the region:

1. Rename the **Regions** > **Default_Fluid** node to `Fluid`.
2. Open the **Fluid** > **Boundaries** node.
3. Select each boundary node in turn.

   The corresponding boundary location is shown in the geometry scene.

Suitable models must be chosen for the simulation, as follows:

4. Open the **Continua** node.
5. Right-click the **Physics 1** node and select the following models, in order:

Group Box	Model
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Time	**Steady**
Viscous Regime	**Turbulent**

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**

In the `hexMesh` window, the color of the **Physics 1** node has turned from gray to blue to indicate that models have been activated.

6. Click **Close**.
7. Open the **Physics 1** > **Models** node to see the selected models.



The default material properties for air and the reference values are suitable for this simulation.

Setting Initial Conditions

Set some initial conditions for this simulation.

1. Open the **Continua** > **Physics 1** > **Initial Conditions** node.
2. Select the **Turbulence Specification** node and set the *Method* to **K + Epsilon**.



3. Select the **Turbulent Dissipation Rate** > **Constant** node and set the *Value* to `0.02 m^2/s^3`
4. Similarly, change the **Turbulent Kinetic Energy** to `0.01` J/kg.

Setting Boundary Conditions and Values

The default boundary conditions for the wall, outlet, and symmetry boundaries are all suitable so the only changes required are to the inlet boundary.

Use the following steps:

1. Select the **Regions** > **Fluid** > **Boundaries** > **Inlet** > **Physics Conditions** > **Turbulence Specification** node and set the *Method* to `K + Epsilon`.
2. Open the **Physics Values** > **Turbulent Dissipation Rate** node and then select the **Constant** node and set the *Value* to `0.07 m^2/s^3`



3. Similarly, change the turbulent kinetic energy at the inlet to `0.024` J/kg.

Give the velocity at the inlet a parabolic profile, using a text file called `inlet.xy`. Before such data can be used as a boundary condition, it must be imported as a table.

4. Open the **Tools** node.
5. Right-click the **Tables** node and select **New Table > File Table**.
6. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` directory and select file `inlet.xy`.
7. Click **Open**.

   The data is now stored in a table called `inlet`.

To assign the data as an inlet velocity magnitude boundary condition:

8. Select the **Regions > Fluid > Boundaries > Inlet > Physics Values > Velocity Magnitude** node and set the *Method* to **Table (x,y,z)**.



9. Select the **Velocity Magnitude** > **Table (x,y,z)** node and set the *Table* to **inlet** and the *Table: Data* to **U**.



Setting Stopping Criteria

The default stopping criterion is set to end the calculation after 1000 iterations. Such a value is excessive, so reduce the maximum number of iterations.

- Select the **Stopping Criteria** > **Maximum Steps** node and set the *Maximum Steps* to `500` .

Visualizing the Solution

Display velocity field vectors on the symmetry plane and contours of y+ values on the walls.

Start by creating a vector scene.

1. Right-click the **Scenes** node and then select **New Scene** > **Vector**.

As you only want to see vectors on the symmetry plane, make sure that this is the only part selected for both the vector and outline displayers.

2. Select the **Scenes** > **Vector Scene 1** > **Displayers** > **Vector 1** > **Parts** node and click the right-hand side of the *Parts* property.
3. Select **Regions** > **Fluid** > **Symplane** in the *Parts* dialog. Click **OK**.
4. Select the **Parts** node within the **Outline 1** node.
5. Edit the *Parts* property so that only **Fluid: Symplane** is selected.
6. Zoom in on the outline in the display window.

Create a plot showing y+ contours on the walls of the geometry:

7. Right-click the **Scenes** node and select **New Scene** > **Scalar**.

8. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Outline 1** > **Parts** node and edit the *Parts* property so that the parts **Fluid: BottomWall**, **Fluid: Obstruction**, **Fluid: SideWall** and **Fluid: TopWall** are selected.
9. Select the same parts for the **Scalar 1** displayer.
10. Right-click the scalar bar at the bottom of the *Scalar Scene 1* display and select **Wall Y+**.
11. Use the mouse buttons to change the viewpoint to one similar to that shown below.



Setting up Line Surfaces for Plotting

Create two parts by intersecting a plane with the bottom wall and symmetry plane boundary surfaces.

This results in the two line segments that are shown below.

To create a section:

1. Click the *Geometry Scene 1* tab in the *Graphics* window to activate it.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**

   A new interactive in-place dialog (*Create Section*) appears.
3. Set the following values:

Box	Property	Value
Input Parts	**Fluid: BottomWall**	Selected
Plane Parameters	**Origin**	[`0.0, 0.0, 1e-8`] m
	Normal	[`0, 0, 1`] m
Display	**No Displayer**	Selected



4. Click **Create**, then click **Close**.



A new derived part named **plane section** is created.

5. Rename this node to `Z = 0`

The section is offset slightly from the symmetry plane to make sure that realistic wall shear stress values are displayed.

Monitor the velocity along a line segment on the symmetry plane at x = 0.38 m. To create a section for this segment:

6. For a second time, right-click the **Derived Parts** node and select **New Part > Section > Plane**

The in-place dialog appears.

7. Set the following values:

Box	Property	Value
Input Parts	**Fluid: Symplane**	Selected
Plane Parameters	**Origin**	[`0.38, 0.0, 0.0`] m
	Normal	[`1, 0, 0`] m
Display	**No Displayer**	Selected

8. Click **Create**, then click **Close**.
9. Rename the new **plane section** node to `X = 0.38`

Plotting Simulation Data

Create plots of velocity and wall shear stress.

For the velocity plot, import experimental data from a file and plot it alongside the numerical data. The development of the solution can be observed throughout the run by viewing these plots.

To create the wall shear stress plot:

1. Right-click the **Plots** node and select **New Plot > X-Y**.
2. Rename the **XY Plot 1** node to `WSS Plot`.
3. Select the **WSS Plot** node and set *Parts* to **Derived Parts > Z = 0**.

Recirculation zones occur upstream and downstream of the obstacle and are visible on the graph as regions where the wall shear stress is negative. To determine the length of the downstream recirculation zone, change the graph formatting.

4. Expand the **WSS Plot > Y Types > Y Type 1** node and set the following properties:

Node	Property	Value
Scalar Function	*Field Function*	**Wall Shear Stress > Laboratory > i**
Z = 0	*Sort Plot Data*	Activated
Line Style	*Style*	**Solid**
Symbol Style	*Shape*	**None**

5. Expand the **WSS Plot > Axes** node and set the following properties:

Node	Property	Value
X Axis	*Maximum*	`0.9`
	Minimum	`0.3`

Node	Property	Value
Major Labels	*Spacing*	`0.05`
Y Axis	*Maximum*	`0.005`
	Minimum	`-0.0075`
Major Labels	*Spacing*	`0.0025`

To create the velocity data plot:

6. Right-click the **Plots** node and select **New Plot > X-Y**.
7. Rename the **XY Plot 1** node to `Velocity Plot at X = 0.38`.
8. Expand the **Velocity Plot at X = 0.38** node and set the following properties:

Node	Property	Value
Velocity Plot at X = 0.38	*Parts*	**Derived Parts > X = 0.38**
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Velocity > Lab Reference Frame > Laboratory > i**
X Type > Vector Quantity	*Value*	`[0.0, 1.0, 0.0] m`

To compare the experimental velocity data with the velocity plot, import the data in the supplied file into a table:

9. Right-click the **Tools > Tables** node and select **New Table > File Table**.
10. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` directory and select file `exp380.xy`.
11. Click **Open**.

The data in that file are now stored in a table named 'exp380'.

To display it on the velocity plot:

12. Return to the **Velocity Plot at X = 0.38** node.
13. Right-click the **Velocity Plot at X = 0.38 > Data Series** node and select **Add Data**.
14. Select **exp380** and click **OK**.

    The supplied experimental data appear in the plot. Numerical data is shown once the simulation has started.

Reverse the axes on the velocity data plot to improve readability:

15. Select the **Axes** node and click the ellipsis on the right of the *Axis Orientation* property.
16. In the *Axis Orientation* dialog, select the option that is shown below:



17. Click **OK**.

To rename the axes:

18. Expand the **Axes** node and set the following properties:

Node	Property	Setting
X Axis > Title	*Title*	`Y (m)`
Y Axis > Title	*Title*	`Velocity (m/s)`

19. To view the solution while it is being generated, look at the scalar and vector scenes.

The pre-processing set-up for the simulation is complete.

20. Stop the macro recording by clicking (**Stop Recording**) in the Macro toolbar.
21. Save the simulation.

Initializing and Running the Simulation

Initialize the solution and run the analysis.

1. Click the (**Initialize Solution**) button in the *Solution* toolbar or use the **Solution > Initialize Solution** menu item. The initial zero velocity condition is shown in the *Velocity Plot at X = 0.38* display.

2. Click the (**Run**) button in the *Solution* toolbar to run the analysis.

   The *Residuals* display is automatically created and shows the progress of the solver. The tabs at the top of the *Graphics* window make it possible to select any active display for viewing. Select the *Velocity Plot at X = 0.38* display if you wish to observe the numerical solution development.

   During the run, it is possible to stop the process by clicking the (**Stop**) button in the toolbar. If you do halt the simulation, it can be continued later by clicking the (**Run**) button. If left alone, the simulation continues until the stopping criteria are satisfied.
3. After the run is finished, save the simulation.

Visualizing the Results

View the vector and scalar scenes.

1. Activate the *Vector Scene 1* display to view the velocity vector field:



2. Click the tab for the *Scalar Scene 1* display to see the y+ distribution along the walls.



The maximum y+ value is 6.9288, at the top of the obstruction. But for the vast majority of the flow regime the y+ values are below 3.0, the maximum value considered acceptable for a rigorous examination of wall data.

Examining Plots

View the velocity and shear stress plots.

Use the following steps:

1. Activate the *Velocity Plot at X = 0.38* display.

The plot now appears as shown below.



The experimental results compare reasonably well with the numerical data. You can see that the greatest discrepancy is found in the recirculation zone downstream of the obstacle.

Given that the y+ contours indicate that the wall data is also reasonably accurate, you can estimate the length of the recirculation zone from the wall shear stress plot.

2. Select the *WSS Plot* tab at the top of the *Display* window.

The plot now appears as shown below:



The above result shows that the downstream recirculation zone extends from approximately X = 0.38 m to X = 0.9 m and hence is roughly 0.5 m in length.

Exporting Simulation Data

Export velocity plot data to a file for comparison with similar data produced using the other two meshes.

1. Open the **Plots > Velocity Plot at X = 0.38 > Y Types > Y Type 1** node.
2. Right-click the **X = 0.38** node within and select **Export**.



3. In the *Save* dialog, enter `hex380` as the File Name and then click **Save**.
4. Close the simulation.

The section dealing with the hexahedral mesh is now complete.

Tetrahedral Mesh

This part of the tutorial models the same problem as the Hexahedral Mesh part, shown in the introductory Steady Flow: Channel Flow with Multiple Meshes section. However, the mesh used in this case is tetrahedral. All problem conditions and properties are defined using the java macro file created when setting up the Hexahedral Mesh case.

Contents

Importing the Mesh and Naming the Simulation
Running the Macro
Opening the Plots
Initializing and Running the Simulation
Visualizing the Results
Examining Plots
Exporting Simulation Data

Importing the Mesh and Naming the Simulation

Create a new simulation and import the tetrahedral mesh.

1. Select **New Simulation** from the menu bar.

Continue by importing the tetrahedral mesh and naming the simulation. The grid has been saved in the `.ccm` file format, which contains all necessary cell, vertex and boundary information for the problem geometry.

2. Select **File > Import > Import Volume Mesh** from the menu bar.
3. In the *Open* dialog, simply navigate to the `[INSTALL_DIR]/doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select file `tet.ccm`
4. Click **Open** to start the import. A geometry scene is created in the *Graphics* window.
5. Finally, navigate back to the `meshComparison` directory and save the new simulation to disk under file name `tetMesh.sim`.

Running the Macro

Replay the macro you created to set up the entire problem and generate the necessary scenes and plots.

1. Click (**Play Macro**) in the Macro toolbar.
2. In the *Open* dialog find and select file `setup.java` created for the Hexahedral Mesh, then click **Open**.

The macro takes a few minutes to run. When it is complete, the vector scene 1 and scalar scene 1 is created in the *Graphics* window.

Opening the Plots

Open the required plots.

1. Open **Plots** node and right-click the **Velocity Plot at X = 0.38** and **WSS Plot** nodes in turn and select **Open**.
2. Save the simulation.

Initializing and Running the Simulation

Initialize and run the simulation.

1. Click (**Initialize Solution**) in the toolbar.
2. Click (**Run**) in the toolbar to run the analysis.
3. After the run is finished, save the simulation.

Visualizing the Results

View the vector and scalar scenes.

1. Activate the *Vector Scene 1* display to view the velocity vector field.



2. Click the tab for the *Scalar Scene 1* display to see the y+ distribution on the walls.



The maximum y+ value is 6.3311, at the top of the obstruction. however for the vast majority of the flow regime the y+ values are below 3.0, the maximum value considered acceptable for a rigorous examination of wall data.

Examining Plots

View the velocity and shear stress plots.

1. Activate the *Velocity Plot at X = 0.38* display.

   The plot now appears as shown below.



The experimental results generally compare well with the numerical data. Again the greatest discrepancy is found in the recirculation zone downstream of the obstacle.

2. Select the *WSS Plot* tab at the top of the *Display* window.

   The plot now appears similar to that shown below.



The above plot shows that the downstream recirculation zone extends from approximately X = 0.37 m to X = 0.82 m and hence is roughly 0.45 m in length.

Exporting Simulation Data

Export data from the velocity plot to a file for comparison with similar data produced in the other parts of this tutorial:

1. Select the **Plots > Velocity Plot at X = 0.38 > Y Types > Y Type 1 > X = 0.38** node and enable the *Sort Plot Data* option.
2. Right-click the **X = 0.38** node within and select **Export**.



3. In the *Save* dialog, enter `tet380` as the File Name and then click **Save**.

The tetrahedral mesh part of the tutorial is now complete.

4. Close the simulation.

Polyhedral Mesh

This part of the tutorial models the same problem as the previous two parts. However, the mesh used in this case is polyhedral. All problem conditions and properties are defined using the java macro file created when setting up the Hexahedral Mesh case.

Contents

Importing the Mesh and Naming the Simulation

Running the Macro
Opening the Plots
Initializing and Running the Simulation
Visualizing the Results
Examining Plots
Comparing Mesh Types

Importing the Mesh and Naming the Simulation

Create a new simulation and import the polyhedral mesh.

Use the following steps:

1. Select **New Simulation** from the menu bar.

Continue by importing the polyhedral mesh and naming the simulation. The grid has been saved in the `.ccm` file format, which contains all necessary cell, vertex and boundary information for the problem geometry.

2. Select **File** > **Import** > **Import Volume Mesh** from the menu bar.
3. In the *Open* dialog, simply navigate to the `[INSTALL_DIR]/doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select file `poly.ccm`.
4. Click **Open** to start the import. A geometry scene is created in the *Graphics* window.
5. Finally, navigate back to the `meshComparison` directory and save the new simulation to disk under file name `polyMesh.sim`.

Running the Macro

Run the setup macro to create the vector and scalar scenes.

Use the following steps:

1. Click (**Play Macro**) in the **Macro** toolbar.
2. In the *Open* dialog find and select file `setup.java` created in the Hexahedral Mesh, and click **Open**.

   The macro takes a few minutes to run. When it is complete, the *Vector Scene 1* and *Scalar Scene 1* is created in the *Graphics* window.

Opening the Plots

Open the required plots.

- Open the **Plots** node and right-click the **Velocity Plot at X = 0.38** and **WSS Plot** nodes in turn and select **Open**.

  The pre-processing setup for the case is complete.

Initializing and Running the Simulation

Initialize and run the simulation.

1. Click (**Initialize Solution**) in the toolbar.
2. Click (**Run**) in the toolbar to run the analysis.

   The *Residuals* display is automatically created and shows the progress of the solver. The tabs at the top of the *Graphics* window make it possible to select any active display for viewing. Select the *Velocity Plot at X = 0.38* display if you wish to observe the numerical solution development.

   During the run, it is possible to stop the process by clicking (**Stop**) in the toolbar. If you do halt the simulation, it can be continued later by clicking (**Run**). If left alone, the simulation continues until the stopping criteria are satisfied.
3. After the run is finished, save the simulation.

Visualizing the Results

View the vector and scalar scenes.

1. Activate the *Vector Scene 1* display to view the velocity vector field.



2. Click the tab for the *Scalar Scene 1* display to see the y+ distribution on the walls.



The maximum y+ value is 6.1811, at the top of the obstruction. However, for the vast majority of the flow regime the y+ values are

below 3.0, the maximum value considered acceptable for a rigorous examination of wall data.

Examining Plots

View the velocity and shear stress plots.

1. Activate the *Velocity Plot at X = 0.38* display.

   The plot appears as shown below:



   The experimental results generally compare well with the numerical data. Again the greatest discrepancy is found in the recirculation zone downstream of the obstacle.

2. Select the *WSS Plot* tab at the top of the *Display* window.

   The plot now appears as shown below.



The above plot shows that the downstream recirculation zone extends from approximately X = 0.37 m to X = 0.77 m and hence is roughly 0.4 m in length. This value is the same as the value the tetrahedral mesh analysis predicts but shorter than the value the hexahedral mesh analysis predicts.

Comparing Mesh Types

You have already examined the velocity plot and seen that the numerical solution compares well with the experimental data. Compare the velocity data that is produced on this mesh with that produced on the Hexahedral Mesh and the Tetrahedral Mesh.

First, import the data from the previous two simulations:

1. Open the **Tools** node.
2. Right-click the **Tables** node and select **New Table > File Table**.
3. Select file `hex380.csv` and click **Open**.
4. Do the same to import the data saved in file `tet380.csv`.

To plot the imported data:

5. Open the **Plots > Velocity Plot at X = 0.38** node.
6. Right-click the **Data Series** node and select **Add Data**.
7. In the *Add Data Providers to Plot* dialog, select **hex380** and click **OK**.
8. Expand the **Data Series > hex380** node and set the following properties:

Node	Property	Value
hex380	*Legend Name*	**HexMesh**
	Sort Plot Data	Activated
Line Style	*Style*	Solid
Symbol Style	*Shape*	**None**

9. Once again, right-click the **Data Series** node and select **Add Data**.
10. In the *Add Data Providers to Plot* dialog, select **tet380** and click **OK**.
11. Select the **Data Series** > **tet380** node and set the following properties:

Node	Property	Value
tet380	*Legend Name*	**TetMesh**
	Sort Plot Data	Activated
Line Style	*Style*	Solid
Symbol Style	*Shape*	**None**

12. Expand the **Data Series** > **exp380** node and set the following properties:

Node	Property	Value
exp380	*Legend Name*	`Experimental`
Symbol Style	*Color*	**Black**

The node representing the internal data set cannot be renamed directly as it inherits its name from the plane section part. Therefore, rename the part instead:

13. Open the **Derived Parts** node.
14. Rename the **X = 0.38** node to `PolyMesh`.
15. Expand the **Plots** > **Velocity Plot at X = 0.38** > **Y Types** > **Y Type 1** node and set the following properties:

Node	Property	Value
PolyMesh	*Sort Plot Data*	Activated
Line Style	**Style**	**Solid**
Symbol Style	**Shape**	**None**

The plot appears as shown below.



There is little difference in the quality of the results from the three meshes, although the result from the polyhedral mesh is slightly better than the other two in the recirculation zone. The polyhedral mesh has only 60,838 cells. The hexahedral mesh has 67,699 cells and the tetrahedral mesh has 140,141 cells. The analysis on the polyhedral mesh only required 121 iterations to reach convergence which compares to 221 for the hexahedral mesh and 181 for the tetrahedral mesh. The results therefore demonstrate the ability of polyhedral meshes to produce good quality results with lower cell counts and shorter run times than other mesh types.

16. Save the simulation.

Summary

This tutorial described a steady airflow simulation, using three different meshing methods.

The tutorial introduced the following STAR-CCM+ features:

- Creating and replaying a java macro

- Renaming a region
- Defining models for single-phase, single-component incompressible fluid flow
- Defining boundary conditions
- Setting a stopping criterion from residuals
- Setting up a line segment on which solution results are plotted
- Plotting wall and flow field data
- Specifying plot options
- Initializing and running the solver until the residual magnitudes are sufficiently small
- Analyzing results using STAR-CCM+'s visualization facilities
- Examining XY-plots
- Exporting simulation data
- Comparing results from different simulations using XY-plots

Steady Flow: Laminar and Turbulent in an S-Bend

This tutorial demonstrates the flow of an incompressible gas through an s-bend of constant diameter (2 cm), for both laminar and turbulent flow. The first part of the tutorial covers the laminar flow, with a Reynolds number (Re) of 500, and the second part covers the turbulent flow, Re = 50,000. The same pipe geometry is used in both cases.

The geometry is created using 3D-CAD, which is the parametric solid modeler available within STAR-CCM+. The necessary region is created and a polyhedral mesh is obtained using the generalized cylinder meshing model. This model extrudes polyhedral cells in the direction of the flow, which results in a mesh that is suited for pipe flows.

Contents

Prerequisites
Creating the S-Bend Geometry
Creating a Geometry Part
Assigning a Part to a New Region
Setting Up the Case

Generating a Surface Mesh
Generating a Volume Mesh
Selecting Physics Models
Modifying Material Properties
Setting Initial Conditions and Boundary Settings
Preparing a Scalar Scene
Setting Up Stopping Criteria
Running the Simulation
Visualizing the Results
Changing to a Turbulent Flow
Suggestions for Further Practice
Summary

Prerequisites

The instructions in the Steady Flow: Laminar and Turbulent in an S-Bend tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Creating geometry in 3D-CAD	Cyclone Separator
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Creating the S-Bend Geometry

Create a new simulation and use 3D-CAD to set up the geometry.

Use the following steps:

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment and create a New Simulation.
2. Save the new simulation to disk with the file name `sBend.sim`.

The geometry is created using 3D-CAD, which is the parametric solid modeler available within STAR-CCM+.

3. To activate 3D-CAD, right-click the **Geometry > 3D-CAD Models** node and select **New**.

Start by renaming the 3D-CAD model.

4. Rename the **3D-CAD Model 1** node to `S-Bend`.

Contents

Creating the Geometry
Specifying Inlet and Outlet Faces

Creating the Geometry

Create the geometry of the s-bend pipe, including the circular inlet face.

Use the following steps:

1. Create a sketch on the YZ plane by right-clicking the **Features > YZ** node and selecting **Create Sketch**.

2. Click (**Set sketch grid spacing**) and change the *Grid Spacing* to `0.0025 m`.

3. Click (**View normal to sketch plane**) to bring the sketch plane into view.

4. Use the (**Create circle**) tool to draw a circle with a *radius* of `0.01 m` and whose center is on the origin, position [0, 0].



5. Click **OK** to exit the sketch.

To create the profile of the pipe along its length:

6. Create a sketch on the XY plane.
7. Click (**View normal to sketch plane**).
8. Use the (**Create line**) tool to draw a line of length `0.035 m` starting at the origin and extending in the positive x direction.



9. Press `Esc` to exit the line tool.

10. Use the (**Center point circular arc**) tool to draw an arc with a radius of 0.02 m. Three mouse clicks are required:
    a) Click position `[0.035 m, 0.02 m]` to define the center point.
    b) Click position `[0.035 m, 0.0 m]` to define the start point.
    c) Click position `[0.055 m, 0.02 m]` to define the end point.



To create a second arc to complete the “S” shape:

11. Use the (**Center point circular arc**) tool to create the second arc. The coordinates for the three clicks are listed below:
    a) Center: `[0.075 m, 0.02 m]`.
    b) Start point: `[0.075 m, 0.04 m]`.
    c) End point: `[0.055 m, 0.02 m]`.



Draw the final section of the pipe.

12. Draw a line with start point `[0.075 m, 0.04 m]` that extends `0.065 m` in the positive x direction.



The completed sketch is shown below.



13. Click **OK** to exit the sketch.

Create the solid body of the pipe using the sweep feature. A sweep requires two sketches: one to act as the profile being swept; the other to act as the path for the sweep.

14. Select both sketches from the 3D-CAD object tree, holding down the Shift key to select multiple items.
15. Right-click one of the highlighted nodes and select **Create Sweep**.



The *Sweep* dialog appears.

16. Accept the default settings and click **OK**.

A new body node, **Body 1**, is added below the **Bodies** node in the object tree.

17. Rename the **Bodies** > **Body 1** node to `Fluid`.
18. Position the geometry as shown below.



Specifying Inlet and Outlet Faces

The final step in preparing the model geometry is to specify the inlet and outlet faces of the model by setting face names.

Use the following steps:

1. Zoom in to the pipe end, right-click the circular face and select **Rename**
2. Enter `Inlet` in the *Rename* dialog and click **OK**.

The benefit of renaming the face in the 3D-CAD model is that it retains its unique identity when the 3D-CAD model is converted to a geometry part.

To rename the outlet:

3. Position the geometry so that the outlet face on the other end of the pipe is visible.
4. Right-click the circular face at the other end of the pipe, select **Rename**, and enter Outlet in the *Rename* dialog. Click **OK** to complete the action.

The geometry is now complete and you can exit 3D-CAD.

5. Click the **Close 3D-CAD** button at the bottom of the object tree.

Note: The *3D-CAD View* scene is closed automatically upon exiting 3D-CAD.

6. Save the simulation.

Creating a Geometry Part

A new geometry part is created using the 3D-CAD model.

Use the following stesp:

1. Right-click the **Geometry** > **3D-CAD Models** > **S-Bend** node and select **New Geometry Part**.



The *Part Creation Options* dialog appears.



2. Click **OK** to accept the default settings.
3. To see the surfaces, expand the **Parts** > **Fluid** > **Surfaces** node.

The inlet and outlet faces in Specifying Inlet and Outlet Faces are defined as separate surfaces.



Assigning a Part to a New Region

Assign the geometry part to a region. The inlet and outlet surfaces have been renamed in 3D-CAD, and are automatically assigned to separate boundaries.

1. Right-click the **Parts > Fluid** node and select **Assign Parts to Regions**.



2. In the *Assign parts to region* dialog set:
   - **Region Mode** (top drop-down) to **Create a Region for Each Part**.
   - **Boundary Mode** (middle drop-down) to **Create a Boundary for Each Part Surface**.



3. Click **Apply** and then **Close**.

A new region has been created with boundaries defining the inlet, outlet, and main fluid volume.

4. Create a geometry scene. Use the mouse to rotate the model in the *Graphics* window to set up the view shown below.



5. Expand the **Regions** > **Fluid** > **Boundaries** node and select each of the boundary nodes to make sure that they have been specified correctly.

Setting Up the Case

In this first part of the tutorial, run the simulation using a steady state laminar flow (Re = 500).

By default, all boundaries are set to be walls. Modify the boundary type for the **Inlet** and **Outlet**:

1. Select the **Regions** > **Fluid** > **Boundaries** > **Inlet** node and set the *Type* to **Velocity Inlet**.



The icon next to the **Inlet** boundary changes, indicating that it is now a velocity inlet boundary.

2. Select the **Outlet** boundary node and set the *Type to* **Pressure Outlet**.

The modified **Boundaries** node is shown below.



Generating a Surface Mesh

Generate a polyhedral mesh using the generalized cylindrical mesher.

This meshing model is suited to pipe flow. Define a prism layer mesh along the wall boundary, to assure that near wall effects are adequately resolved.

1. Right-click the **Continua** node and select **New > Mesh Continuum**.
2. Right-click the **Mesh 1 > Models** node and select the following models, in order:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Generalized Cylinder**
	Prism Layer Mesher

3. Click **Close**.

To define the mesh settings:

4. Select the **Mesh 1 > Reference Values > Base Size** node.



5. Set the *Value* to `0.0008` m.



6. Continuing with the reference values section, set the following properties:

Node	Property	Value
Number of Prism Layers	*Number of Prism Layers*	5
Prism Layer Stretching	*Prism Layer Stretching*	1.8
Prism Layer Thickness > Relative Size	*Percentage of Base*	45

Generate and check the surface mesh before running the generalized cylinder mesher. A good quality surface mesh is needed before the volume mesh can be generated.

7. Click (**Generate Surface Mesh**).
8. Create a mesh scene and examine the surface mesh.



Generating a Volume Mesh

Before generating the volume mesh, define the boundaries that form the cylindrical geometry.

Use the following steps:

1. Right-click the **Mesh 1 > Models > Generalized Cylinder** and select **Manage Cylinders**.



The *Select Cylinders* dialog appears.



2. Make sure the **Default** cylinder is selected, then click **OK**.
3. Select the **Regions > Fluid > Boundaries > Default > Mesh Conditions > Generalized Cylinder Extrusion Type** node.



Verify that the *Extrusion law* property was set to **Constant.**

4. Within the same **Default** boundary, select the **Default > Mesh Values > Generalized Cylinder Parameters** node.

Notice that the *Number of Layers* property was set to `69`. STAR-CCM+ calculates this value from the geometry model.

The volume mesh can now be generated.

5. Click the (**Generate Volume Mesh**) button.
6. Right-click a blank space in *Mesh Scene 1* and select **Apply Representation > Volume Mesh**.
7. Examine the final volume mesh.



8. Save the simulation.

Selecting Physics Models

Select the models for the incompressible fluid.

A physics continuum was automatically created during the mesh generation process. To select the physics models:

1. Right-click the **Continua** > **Physics 1** > **Models** node and select the following models, in order:

Group Box	Model
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Laminar**

2. Click **Close**.

The selected models now appear.



Modifying Material Properties

Set the values for the dynamic viscosity and density of air.

Use the following steps:

1. Select the **Models** > **Gas** > **Air** > **Material Properties** > **Density** > **Constant** node and set the *Value* to `1 Kg/m^3`.
2. Select the **Dynamic Viscosity** > **Constant** node and set the *Value* to `1.716E-5 Pa-s`.

Setting Initial Conditions and Boundary Settings

To achieve a Reynolds number of 500 for this gas and pipe diameter, you require a mean velocity of 0.429 m/s.

To assist solution convergence, define the initial condition for velocity:

1. Select the **Physics 1** > **Initial Conditions** > **Velocity** > **Constant** node and set the *Value* to `[0.429, 0.0, 0.0] m/s`.

Set the same velocity on the **Inlet** boundary.

2. Select the **Regions** > **Fluid** > **Boundaries** > **Inlet** > **Physics Values** > **Velocity Magnitude** > **Constant** node and set the *Value* to `0.429 m/s`.

Preparing a Scalar Scene

A scalar scene is created to show the velocity magnitude on a section plane through the center of the s-bend. Use this scene to visualize the solution while the simulation is running.

Use the following steps:

1. Create a scalar scene.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**
3. In the *Create Section* panel, set the following values:

Property	Value
Input Parts	**Fluid**
normal	[`0, 0, 1`] m
Display	**Existing Displayer**
Dropdown menu	**Scalar 1**

The completed dialog is shown below.



4. Click **Create** and then **Close**.

   Scalar Scene 1 is shown below.



5. Click the **scene/plot** button.
6. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set the *Function* to **Velocity** > **Magnitude**.
7. Select the **Displayers** > **Scalar 1** node and set the *Contour Style* to **Smooth Filled**.
8. Click the **Simulation** button to return to the STAR-CCM+ simulation object tree.

   A vector plot displaying the velocity magnitude is created.
9. Create a vector plot.
10. Using the drag-and-drop method, add the section plane from the **Derived Parts** node to the scene.



Setting Up Stopping Criteria

Limit the number of iterations for the solver to 500.

500 iterations are sufficient for the solution to converge. Use the following steps:

1. Select the **Stopping Criteria** > **Maximum Steps** node and set the *Maximum Steps* to `500`.
2. Save the simulation.

Running the Simulation

The simulation is now ready to be run.

Use the following steps:

1. Click (**Run**).

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view each of the scalar scenes.

During the run, it is also possible to stop the process by clicking (**Stop**) on the toolbar. If you do halt the simulation, it can be continued later by clicking (**Run**). If left alone, the simulation continues until 500 iterations are complete.

The Residuals display is created automatically and shows the progress of the solvers.

2. While the simulation is running, select the *Scalar Scene 1* tab at the top of the *Graphics* window to visualize the solution.
3. When the simulation has finished running, save it.

Visualizing the Results

Examine the scalar and vector scenes.

The scalar scene is shown below.



The vector scene is shown below.



Changing to a Turbulent Flow

Edit the models, initial conditions, and boundary conditions to simulate a turbulent flow (Reynolds number = 50,000).

First, save the state of the simulation as developed so far:

- Save the simulation as `sBendTurbulent.sim`.

Contents

Modifying the Physics Continuum
Modifying Initial Conditions and Boundary Settings
Extending the Stopping Criteria
Clearing the Solution and Running the Simulation
Visualizing the Results for the Turbulent Flow

Modifying the Physics Continuum

Edit the physics models to represent turbulent flow.

Disable the laminar model and enable the turbulent and K-epsilon models.

1. For the physics continuum **Continua** > **Physics 1**, deselect **Laminar** in the *Enabled Models* box.
2. Select the following models, in order:

Group Box	Model
Viscous Regime	**Turbulent**
	Reynold-Averaged Navier-Stokes (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
	Realizable K-Epsilon Two-Layer (Selected automatically)
	Two-Layer All y+ Wall Treatment (Selected automatically)

3. Click **Close**.

The selected models now appear.



Modifying Initial Conditions and Boundary Settings

To achieve a turbulent flow (Re = 50,000) in the existing pipe, increase the velocity of the air to 42.9 m/s at the **Initial Conditions** node and at the **Inlet** boundary condition. Change the **Turbulence Specification** so that the turbulence intensity and length scale can be defined.

Use the following steps:

1. Select the **Continuum > Physics 1 > Initial Conditions > Velocity > Constant** node and set the *Value* to `[42.9, 0.0, 0.0] m/s`.
2. Select the **Turbulence Specification** node and set its *Method* to **Intensity + Length Scale**.

   The **Turbulent Length Scale** node is added to the object tree.
3. Set the following values:

Node	Property	Value
Turbulent Length Scale > Constant	*Value*	`0.0014 m` The turbulent length scale that is used in this case is

Node	Property	Value
		approximately 7% of the hydraulic diameter, which for circular pipes is the same as the pipe diameter.
Turbulent Velocity Scale > Constant	*Value*	`4.29 m/s` The turbulent velocity scale that is used in this case is 10% of the free-stream velocity.
Turbulence Intensity > Constant	*Value*	`0.12` The turbulence intensity is approximately 12% and was calculated using the following equation: $I = \sqrt{\frac{\frac{3}{2}v_t^2}{U^2}}$ where I is the turbulence intensity, v is the turbulent velocity scale, and U is the free-stream velocity.

Use the same values at the **Inlet** boundary. Start by setting the **Turbulence Specification** to use **Intensity + Length Scale**.

4. Select the **Regions > Fluid > Boundaries > Inlet > Physics Conditions > Turbulence Specification** node and set the *Method* to **Intensity + Length Scale**.
5. Set the following values:

Node	Property	Value
Physics Values > Turbulent Length Scale > Constant	*Value*	`0.0014 m`

Node	Property	Value
Regions > Fluid > Boundaries > Inlet > Physics Values		
Velocity Magnitude > Constant	*Value*	`42.9 m/s`
Turbulence Intensity > Constant	*Value*	`0.12`

Extending the Stopping Criteria

The turbulent flow requires a longer time to converge. Increasing the number of iterations to 1000 gives the solver sufficient time to converge.

1. Select the **Stopping Criteria > Maximum Steps** node and set the *Maximum Steps* to `1000`.
2. Save the simulation.

Clearing the Solution and Running the Simulation

Clear the solution from the laminar flow before running the solver.

Use the following steps:

1. Select **Solution > Clear Solution** from the menu.
   The *Clear Solution* dialog appears.
2. Accept the default settings and click **OK**.
   The simulation is now ready to be run.
3. Click (**Run**).
4. While the simulation is running, select the *Scalar Scene 1* tab at the top of the *Graphics* window to visualize the solution.
5. When the simulation has finished running, save it.

Visualizing the Results for the Turbulent Flow

Examine the scalar and vector scenes.

The scalar scene is shown below.



The vector scene is shown below.



Suggestions for Further Practice

Try running the simulation with different initial and boundary conditions.

To see how the laminar and turbulent models differ, you could try running the Re = 50,000 case with the laminar physics model.

1. Open the `sBend.sim` file.
2. Increase the flow velocity to `42.9 m/s` in the **Initial Conditions** and at the **Inlet** boundary.
3. Clear the solution and run the simulation.

The residuals do not converge, but instead start to oscillate which indicates that the model setup is incorrect.

Summary

This tutorial examined the flow of an incompressible fluid through an s-bend of constant diameter, at different Reynolds numbers.

This tutorial covered the following topics:

- Creating a geometry model in 3D-CAD
- Using the generalized cylinder mesher
- Solving a laminar flow through the pipe
- Solving a turbulent flow through the pipe
- Plotting results of each case

Steady Flow: Backward Facing Step

This tutorial demonstrates the flow of an incompressible gas over a backward facing step (with height, H = 1 m).

The turbulent flow has a Reynolds number of 28,000 (based on the step height, H, and the free stream velocity). The K-Epsilon turbulence model with High y+ wall treatment is used as it allows you to use a relatively coarse mesh without prism layers. The bottom wall is also heated with a heat flux of 500 W/m^2. The geometry of the backward facing step is shown below.

Contents

Prerequisites
Setting Up the Geometry
Setting Up the Backward Facing Step Simulation
Summary
Backward Facing Step Bibliography

Prerequisites

The instructions in the Backward Facing Step tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction

Techniques	Associated Tutorial
Creating geometry in 3D-CAD	3D-CAD: Cyclone Separator
Using visualization tools, scenes and plots	Introduction

Setting Up the Geometry

Create the model for the backward-facing step.

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment and create a New Simulation.
2. Save the new simulation with the file name `backwardFacingStep.sim`

The geometry is created using 3D-CAD.

3. Right-click the **Geometry** > **3D-CAD Models** node and select **New**. This creates a model and activates 3D-CAD.

Start by renaming the 3D-CAD model.

4. Select the **3D-CAD Model 1** node at the top of the object tree. Rename the model to `Fluid Domain`.

Contents

Creating the Geometry
Specifying Boundary Faces on the Backward Facing Step Geometry

Creating the Geometry

Create the geometry of the backward facing step with the sketch plane and **Create line** tool.

Use the following steps:

1. Create a sketch on the XY plane by right-clicking the **Features** > **XY** node and selecting **Create Sketch**.
2. Click (**Set sketch grid spacing**) and change the *Grid Spacing* to `0.25 m`.
3. Click (**View normal to sketch plane**) to bring the sketch plane into view.

4. Use (**Create line**) to draw a line of length `1 m` starting at the origin (position `[0, 0]`) and extending in the positive y direction.



5. Right-click the vertex that lies on the origin and select **Apply Fixation Constraint**. This prevents the sketch from moving.
6. Right-click the line and select **Apply Vertical Constraint**.
7. Right-click the line and select **Apply Length Dimension**. In the Dimension dialog, enter `1 m` as the dimension and click **OK**.
8. Continue by creating the other lines as shown below.

It is not necessary to obtain the exact dimensions as you draw the sketch--these can be added later. It is recommended, however, that you add vertical and horizontal constraints to each appropriate line before applying the dimension constraints. This makes sure that the overall shape of the sketch does not change.

9. Click **OK** to exit the sketch.

Extrude the sketch profile to obtain a solid body:

10. Right-click **Sketch 1** and select **Create Extrude**.
    The *Extrude* dialog appears.
11. Set the *Distance* to `0.5 m`
12. Click **OK** to perform the extrusion.
    A new body node, **Body 1**, is added below the **Bodies** node in the object tree.
13. Select the **Bodies** > **Body 1** node and rename it to `Backward Facing Step`.



Specifying Boundary Faces on the Backward Facing Step Geometry

The next step in preparing the model geometry is to identify the boundary faces of the model by setting face names. The benefit of renaming the face in the 3D-CAD model is that it retains its unique

identity when the 3D-CAD model is converted to a geometry part in the next stage of the tutorial.

Use the following steps:

1. Position the geometry as shown below.



2. Right-click the face shown below and select **Rename**.



3. Enter `Inlet` in the *Rename* dialog and click **OK**.
4. Rename the other faces on the geometry to match those shown below.



The renamed faces are listed in the **Bodies** > **Backward Facing Step** > **Named Faces** node, as shown below.



Clicking a node in the object tree highlights the relevant face in the *3D-CAD view* scene.

The backward facing step geometry is now complete. The next step is to create the geometry parts and regions.

5. Click **Close 3D-CAD**.
6. Save the simulation.

Setting Up the Backward Facing Step Simulation

Create a new geometry part using the 3D-CAD model.

Use the following steps:

1. Right-click the **Geometry > 3D-CAD Models > Fluid Domain > Bodies > Backward Facing Step** and select **New Geometry Part**.

   The *Parts Creation Options* dialog appears.
2. The default settings are acceptable, so you can click **OK** to close the dialog.

3. Expand the **Parts** > **Backward Facing Step** > **Surfaces** node and note that the boundary faces specified in 3D-CAD are defined as separate surfaces.



Contents

Assigning a Part to a New Region
Setting the Boundary Types
Setting Up the Mesh Models
Creating the Volume Shapes and Defining the Volumetric Controls
Generating the Volume Mesh and Converting From 3D to 2D
Selecting Physics Models
Modifying Material Properties and Setting Initial Conditions
Setting Up Boundary Conditions
Preparing Scenes for Visualization
Preparing the Plot
Setting Up Stopping Criteria
Modifying Solver Settings
Running the Simulation
Visualizing the Results
Validating the Results

Assigning a Part to a New Region

Assign the geometry part to a region and create boundaries.

1. Right-click the **Parts > Backward Facing Step** node and select **Assign Parts to Regions**.
   The *Assign Parts to Regions* dialog appears.
2. Set the upper dropdown value to **Create a Region for Each Part**.
3. Set the lower dropdown value to **Create a Boundary for Each Part Surface**.



4. Click **Apply** and then **Close**.

A new region is created. Each labeled surface is now a separate boundary in the region.

5. Create a geometry scene.

Setting the Boundary Types

Set the boundary types for the inlet and outlet boundaries.

1. Select the **Regions > Backward Facing Step > Boundary > Inlet** node and set the *Type* to **Velocity Inlet**.



2. Select the **Outlet** node and set the *Type* to **Pressure Outlet**.

Setting Up the Mesh Models

Generate a volume mesh using the trimmer mesher.

To make sure that the y+ value calculated by the turbulence model is around 50, do not refine the mesh near the walls.

1. Right-click the **Continua** node and select **New > Mesh Continuum**.
2. Right-click the **Mesh 1 > Models** node and choose **Select Meshing Models**
3. In the *Meshing Model Selection* dialog select the following models:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Trimmer**

4. Click **Close**.

Modify the mesh model parameters:

5. Select the **Mesh 1** > **Models** > **Trimmer** node and activate the *Do mesh alignment* property.

Edit the mesh reference values:

6. Select the **Mesh 1** > **Reference Values** node and set the following properties:

Node	Property	Value
Base Size	*Value*	`0.2 m`
Mesh Alignment Location	*Location*	`[0.0, 1.0, 0.0] m`
Surface Size > **Relative Minimum Size**	*Percentage of Base*	`100`

7. Save the simulation.

Creating the Volume Shapes and Defining the Volumetric Controls

Create a volume shape to refine the mesh near the bottom wall.

Use the following steps:

1. Right-click the **Tools** > **Volume Shapes** node and select **New Shape** > **Block**.

   The *Create Block Volume Shape* in-place dialog appears.
2. In the *Maximum and Minimum Coordinates* box, set *Corner 1* to `[-3.8 m, 0.0 m, 0.0 m]`
3. In the same box, set *Corner 2* to `[30 m, 1.5 m, 0.5 m]`

A preview of the block shape is visible in the open scene.

4. Click **Create** and then **Close**.

The new shape, **Block 1**, is added to the **Tools** > **Volume Shapes** node. To create the volumetric control:

5. Right-click **Continua** > **Mesh 1** > **Volumetric Controls** and select **New**.

A new node, **Volumetric Control 1**, is added below the **Volumetric Controls** node.

6. Select the **Volumetric Control 1** node and edit the *Shapes* property.
7. Select **Block 1**.



8. Click **OK**.
9. Select the **Volumetric Control 1 > Mesh Conditions > Trimmer** node and activate *Customize isotropic size*.

   A new node, **Mesh Values**, is added to the **Volumetric Controls** node.
10. Select the **Mesh Values > Custom Size > Relative Size** and set the *Percentage of Base* to `50`.

Generating the Volume Mesh and Converting From 3D to 2D

The volume mesh is now ready to be generated and converted to 2D.

Use the following steps:

1. Click (**Generate Volume Mesh**).
2. Create a mesh scene and examine the mesh.

The mesh is now ready to be converted to 2D.

3. Select **Mesh > Convert to 2D**.

The *Convert Regions to 2D* dialog appears.

4. Accept the default settings and click **OK**.
5. Click (**Reset View**) to bring the geometry back into view.

The 2D mesh is shown below.



Delete the region and physics continuum associated with the 3D mesh:

6. Right-click the **Regions** > **Backward Facing Step** node and select **Delete**.
7. Click **Yes** in the dialog that appears.
8. Rename the **Backward Facing Step 2D** region to `Backward Facing Step`
9. Right-click the **Continua** > **Physics 1** node and select **Delete**.
10. Click **Yes** in the dialog that appears.
11. Rename the **Physics 1 2D** continuum node to `Fluid`
12. Save the simulation.

Selecting Physics Models

A physics continuum was automatically created during the mesh generation process. Select the physics models for the first part of the tutorial.

As you have a fine prism layer it is recommended to use the Realizable K-Epsilon Two-Layer turbulence model.

1. For the physics continuum **Continua** > **Fluid**, select the following models, in order:

Group Box	Model
Space	**Two-Dimensional** (Pre-selected)
Time	**Steady**
Material	**Gas**
Flow	**Coupled Flow**
Gradient Metrics	**Gradients** (Selected automatically)
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**

2. Deactivate **Auto-select recommended models**. The appropriate models are selected manually.
3. Select the following models, in order:

Group Box	Model
Reynolds Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Realizable K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**
Optional Models	**Coupled Energy**

4. Click **Close**.

The selected models now appear.



Modifying Material Properties and Setting Initial Conditions

Modify the material properties of air.

1. Select the **Fluid** > **Models** > **Gas** > **Air** > **Material Properties** node and set the following values:

Node	Property	Value
Density > Constant	*Value*	`2 kg/m^3`
Dynamic Viscosity > Constant	*Value*	`7.142E-5 Pa-s`
Specific Heat > Constant	*Value*	`1000.0 J/kg-K`
Thermal Conductivity > Constant	*Value*	`0.1 W/m-K`



Define the initial conditions.

2. Select the **Fluid > Initial Conditions** node and select the following values:

Node	Property	Value
Turbulence Specification	**Method**	**K + Epsilon**
Turbulent Dissipation Rate > Constant	*Value*	`3.0E-6 m^2/s^3`
Turbulent Kinetic energy > Constant	*Value*	`1.5E-4 J/kg`
Velocity > Constant	*Value*	`[1.0, 0.0, 0.0] m/s`



Setting Up Boundary Conditions

Define the inlet boundary using inlet data from a table.

Use the following steps:

1. Right-click **Tools** > **Tables** and select **New Table** > **File Table**.



2. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select file `backStepInletData.csv`.

   The file is added to the **Tables** node.

3. Select the **Regions** > **Backward Facing Step** > **Boundaries** > **Inlet** > **Physics Conditions** > **Turbulence Specification** node and set the *Method* to **K + Epsilon**.



4. In the same **Physics Conditions** node select the **Velocity Specification** node and set the *Method* to **Components**.
5. Select the **Physics Values** node and set the following properties:

Node	Property	Value
Turbulent Dissipation Rate	*Method*	**Table (x,y,z)**
[redacted]Table (x,y,z)	*Table*	**backStepInletData**
	Table: Data	**Turbulent Dissipation Rate**
Turbulent Kinetic Energy	*Method*	**Table (x,y,z)**
[redacted]Table (x,y,z)	*Table*	**backStepInletData**
	Table: Data	**Turbulent Kinetic Energy**
Velocity	*Method*	**Composite**
[redacted]Composite > X Component	*Method*	**Table (x,y,z)**
[redacted]Table (x,y,z)	*Table*	**backStepInletData**

Node	Property	Value
	Table: Data	**Velocity[i]**

Set the heat flux on the bottom wall:

6. Select the **Boundaries > Step Bottom > Physics Conditions > Thermal Specification** node and set the *Method* to **Heat Flux**.



7. Select the **Physics Values > Heat Flux > Constant** node and set the *Value* to `500 W/m^2`
8. Save the simulation.

Preparing Scenes for Visualization

A scalar scene is created to show the velocity magnitude on the 2D mesh. Use this to visualize the solution while the simulation is running.

Use the following steps:

1. Create a scalar scene.
2. Click the **scene/plot** button above the object tree and select the **Displayers > Scalar 1 > Scalar Field** node and set the *Function* to **Velocity > Magnitude**.
3. Click the **Simulation** button to return to the STAR-CCM+ simulation object tree.

Preparing the Plot

Monitor the wall shear stress in the x direction.

Use the following steps:

1. Create an X-Y plot.
2. Rename the new plot to `Wall Shear Stress (x)`.
3. Select the **Plots > Wall Shear Stress (x)** node and click to the right of the *Parts* property.

   The Wall Shear Stress (x) dialog appears.



4. Expand the **Regions** > **Backward Facing Step** > **Boundaries** node.
5. In the column to the left of the Menu, select the **Select Step Bottom** button.
6. Click **OK**.
7. Select the **Plots** > **Wall Shear Stress (x)** > **Y Types** > **Y Type 1** > **Scalar Function** node and set the *Field Function* to **Wall Shear Stress** > **Laboratory** > **i**.
8. Open the plot by right-clicking its node in the object tree and selecting **Open**.

Copy and paste the Wall Shear Stress (x) plot and modify its parameters so that it displays the Nusselt Number. The Nusselt number is the ratio of the heat transferred between two parallel plates (at different temperatures) by a moving fluid to the heat transfer that would occur by conduction alone.

9. Right-click the **Plots** > **Wall Shear Stress (x)** plot and select **Copy**.
10. Right-click the **Plots** node and select **Paste**.

A new plot, **Copy of Wall Shear Stress (x)**, is added to the Plots node.

11. Rename the new plot to `Nusselt Number`.
12. Select the **Plots** > **Nusselt Number** > **Y Types** > **Y Type 1** > **Scalar Function** node and set the *Field Function* to **Nusselt Number**.
13. Right-click the **Nusselt Number** plot and select **Open**.

In the **Modifying Material Properties and Setting Initial Conditions** section you modified the default thermal conductivity value. As a result, the Nusselt number field function, which is affected by this change, needs to be updated.

14. Select the **Tools** > **Field Function** > **Nusselt Number** node and set the *Reference Thermal Conductivity* to `0.1 W/m-K`.



Setting Up Stopping Criteria

Limit the number of iterations for the solver to 500. This is sufficient for the solution to converge.

1. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `500`.
2. Save the simulation.

Modifying Solver Settings

Set the Courant number to 100 and activate the grid sequencing and Expert Driver models.

These tools assist the coupled solver in coming to a converged solution.

- Select the **Solvers** > **Coupled Implicit** node and set the following properties:

Node	Property	Value
Coupled Implicit	*Courant Number*	`100`
Expert Initialization	*Method*	**Grid Sequencing**
Solution Driver	*Driver*	**Expert Driver**



Running the Simulation

The simulation is now ready to be run.

1. Click (**Run**) tton.

   While the simulation is running you can click the tabs at the top of the *Graphics* window to view each of the scalar scenes.

   During the run, it is also possible to stop the process by clicking (**Stop**) on the toolbar. If you do halt the simulation, it can be continued later by clicking (**Run**). If left alone, the simulation continues until 500 iterations are complete.

   The Residuals display is created automatically and shows the progress of the solvers.
2. While the simulation is running, select the *Scalar Scene 1* tab at the top of the *Graphics* window to visualize the solution.
3. When the simulation has finished running, save the simulation.

Visualizing the Results

Examine *Scalar Scene 1*.

The scalar scene is shown below.



Validating the Results

Import the experimental data intoSTAR-CCM+ and add it to the Wall Shear Stress (x) and Nusselt Number plots.

In this simulation, the wall shear stress and skin friction coefficient are the same.

1. Right-click the **Tools** > **Tables** node and select **New Table** > **File Table**
2. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and multi-select the files `backStepSkinFrictionRefData.csv` and `backStepNusseltNumberRefData.csv`.
3. Click **Open**.

The files are added to the **Tables** node, as shown below.



Add the tabular data to the plots, starting with the **Wall Shear Stress (x)** plot.

4. Right-click the **Data Series** node and select **Add Data**.
5. Choose backStepSkinFrictionRefData Table.
6. Rename backStepSkinFrictionRefData Legend Name to `Experimental`.
7. Select the **Tabular** > **Experimental** node and set the following values:

Property	Value
Table	**backStepSkinFrictionRefData**
X Column	**Position and meter**
Y Column	**Pascal and SkinFrictionCoefficient**

8. Select the **Experimental** > **Symbol Style** node and set the following values:

Property	Value
Color	**black**
Shape	**Filled Circle**

The Wall Shear Stress (x) plot is shown below.



Add the tabular data to the Nusselt Number plot.

9. Right-click the **Data Series** node and select **Add Data**.
10. Choose backStepNusseltNumberRefData Table.
11. Rename backStepNusseltNumberRefData Legend Name to `Experimental`.
12. Select the **Tabular** > **Experimental** node and set the following values:

Property	Value
Table	**backStepNusseltNumberRefData**
X Column	**Position and meter**
Y Column	**Nusselt Number and NusseltNumber**



13. Select the **Experimental** > **Symbol Style** node and set the following values:

Property	**Value**
Color	**black**
Shape	**Filled Circle**

The Nusselt Number plot is shown below.



Summary

This tutorial demonstrated the flow of an incompressible fluid over a backward facing step.

The tutorial covered the following:

- Creating a geometry model in 3D-CAD
- Applying fully developed flow data to the inlet boundary
- Plotting the wall shear stress and Nusselt number, and comparing these with experimental results

Backward Facing Step Bibliography

Vogel, J. C and Eaton, J. K., 1985. "Combined Heat Transfer and Fluid Dynamic Measurements Downstream of a Backward-Facing Step." *AIAA J.* 23(2) pp. 163-171.

Adjoint Flow Solver: External Flow over a Dual Element Wing

This tutorial demonstrates how to run an adjoint flow analysis on a dual element wing for a race car.

The velocity of the flow is 100 kph. The objective of the analysis is to identify the factors that influence the downforce generated by the wing.



The angle of attack of the wing is 16.2 degrees. The fluid domain consists of a velocity inlet, pressure outlet, and symmetry boundaries. The wing is positioned in the domain as follows:



An adjoint flow analysis is split into two stages. In the first stage, you run the primary (or primal) flow solvers until the solution converges to machine precision. In the second stage, you run the adjoint solver that takes the primal flow solution as an input.

There are two aspects of achieving a well-converged solution:

- generating a high-quality volume mesh
- applying appropriate solution controls

In this tutorial, you address both aspects to achieve convergence at the level of machine convergence.

For this tutorial, you are provided with a simulation file that contains the following objects already defined:

- A mesh continuum with mesh settings. To keep the file size to a minimum, the volume mesh is not provided—the mesh is ready for you to generate.
- A physics continuum with physics models required for the initial flow solution. The Adjoint Flow model is activated in a later step.
- Initial conditions.
- Regions and boundaries (including the boundary types).
- A mesh scene.

You require a double-precision version of STAR-CCM+ to complete this tutorial.

Contents

Prerequisites
Loading the Simulation File and Visualizing the Mesh
Checking the Mesh Quality

Modifying the Mesh Parameters and Updating the Volume Mesh
Creating the Downforce Report, Monitor, and Plot
Modifying Solver Parameters
Setting Stopping Criteria
Running the Primal Solver
Activating the Adjoint Model
Creating the Adjoint Cost Function
Running the Adjoint Flow Solver
Visualizing Sensitivity of Downforce with Respect to the Flow
Summary

Prerequisites

The instructions in this tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

You must be familiar with the following techniques before attempting this tutorial:

Technique	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Loading the Simulation File and Visualizing the Mesh

Load `adjointWing_start.sim` and display the volume mesh.

It is best to run this tutorial in parallel, using at least 4 cores. This tutorial case has a large number of cells and can take a long time to run to machine convergence if using only one core.

Load the provided simulation file:

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment. Remember to specify multiple cores.
2. Load `adjointWing_start.sim` from the `[INSTALL_DIR]/doc/startutorialsdata/adjoint/data` directory.

   STAR-CCM+ provides feedback on the loading process in the *Output* window.

The automated mesh operation with all the mesh settings have been created for you already.

To generate the volume mesh:

3. Click (**Generate Volume Mesh**).

To visualize the mesh:

4. Right-click the **Scenes** > **Mesh Scenes 1** node and choose **Open**.
5. Right-click in a blank area of the scene and select **Apply Representation** > **Volume Mesh**.
6. Press the <S> key to display the side view of the geometry.
   The volume mesh appears as follows.



7. Save the simulation as `adjointWing_primal.sim`.

Checking the Mesh Quality

Adjoint flow analyses require a high quality volume mesh in order to get a solution that converges to an acceptable level.

The Cell Quality Remediation model includes a field function called Bad Cell Indicator. This field function identifies cells that have a negative impact on solution accuracy and convergence. Using this field function, you can generate a threshold derived part to display poor quality cells.

Activate the Cell Quality Remediation model:

1. Right-click the **Continua > Physics 1** node and choose **Select Models...**

   The *Physics 1 Model Selection* dialog appears.
2. Select **Cell Quality Remediation** from the *Optional Models* group box.
3. Click **Close**.

Check the volume mesh for bad cells. First, create a geometry scene:

4. In the **Vis** toolbar, click (**Create/Open Scenes**) and select **Geometry**.
5. Click **scene/plot**.
6. Right-click **Displayers > Geometry 1 > Parts** and select **Edit...**.

   The *Edit* dialog appears.
7. Right-click on a blank area in the dialog and select **Deselect All**.
8. Click **OK**.
9. Click **simulation**.

Create a threshold to display the bad cells:

10. Right-click **Derived Parts** and select **New Part > Threshold...**.

    The *Create Threshold* panel appears.
11. Set *Scalar* to **Bad Cell Indicator**.
12. Set *Extraction Mode* to **All Above**.
13. Set *Display* to **Existing Displayer > Geometry 1**.

    The remaining settings are acceptable. The completed panel is displayed below:



14. Click **Create** and then click **Close**.
    The open scene is populated with the bad cells.



Multiple clusters of bad cells reduce solver accuracy and stability. Within each cluster, a small cell is surrounded by large cells. It is this large jump in cell size that negatively impacts the solution.

Modifying the Mesh Parameters and Updating the Volume Mesh

One method for improving the quality of the volume mesh is to increase the number of optimization cycles in the volume mesher. You can also increase the quality threshold for the mesher, so that it targets a minimum cell quality.

To modify the mesh:

1. Select **Geometry** > **Operations** > **Automated Mesh** > **Meshers** > **Polyhedral Mesher** and set *Optimization Cycles* to `3`.
2. Set *Quality Threshold* to `0.7`.
3. Click (**Generate Volume Mesh**).
   The number of bad cells is significantly reduced.

Deactivate the Cell Quality Remediation model:

4. Right-click the **Continua** > **Physics 1** node and choose **Select Models...**.

   The *Physics 1 Model Selection* dialog appears.
5. Deselect **Cell Quality Remediation** from the *Enabled Models* group box.
6. Click **Close**.

Contents

Increasing the Courant Number
Activating Grid Sequencing
Activating the Expert Driver

Increasing the Courant Number

For steady-state simulations, the Courant number controls the local pseudo-time step size. Increasing this value increases the rate of convergence. For this case, a Courant number of 200 accelerates convergence without causing the solver to become unstable. See Coupled Flow Controls and General Setup Recommendations.

To modify the Courant number:

- Select the **Solvers** > **Coupled Implicit** node and set *Courant Number* to `200`.



Activating Grid Sequencing

With the Coupled solver, you can activate grid sequencing to perform initialization followed by the computation of an approximate inviscid solution for the flow. See Using Grid-Sequencing Initialization.

To activate grid sequencing:

1. Select the **Solvers** > **Coupled Implicit** > **Expert Initialization** node and set *Method* to **Grid Sequencing**.

   The **Grid sequencing** sub-node appears.
2. Select the **Grid sequencing** node and set *CFL Number* to `100`.



Activating the Expert Driver

The expert driver is a tool that helps control and accelerate convergence of steady-state simulations that use the Coupled Implicit solver.

For this case, the Courant number is ramped from 0.25 to 200 over 100 iterations. This helps the solution converge during the first 100 iterations. See Using the Expert Driver.

To activate the expert driver:

1. Select the **Solvers** > **Coupled Implicit** > **Solution Driver** node and set *Driver* to **Expert Driver**.

   The **Expert Driver** sub-node appears.
2. Select the **Expert Driver** sub-node and set the following values:

Property	Value
CFL Ramp: End Iteration	`100`
Corrections Control: Minimum Explicit Relaxation	`0.5`
Corrections Control: URF Recovery Rate	`10`
Corrections Control: Relative Residual Cutoff	`1.0E-8`

   The final expert driver properties are displayed below.

Expert Driver - Properties

Properties	
Verbose	☐
Corrections Control	☑
Corrections Control: Update Frequency	5
Corrections Control: URF Recovery Rate	10.0
CFL Control	☑
CFL Limiting	☑
CFL Ramp: Start Value	0.25
CFL Ramp: Start Iteration	1
CFL Ramp: End Iteration	100
Expert-Driver Start Iteration	1
Expert-Driver Cutoff Iteration	10000000
Expert	
Corrections Control: Minimum Explicit Relaxation	0.5
Corrections Control: Maximum Explicit Relaxation	2.0
Corrections Control: Relative Residual Cutoff	1.0E-8
Enhanced Dissipation: Start transition	1
Enhanced Dissipation: End transition	100
CFL Control: Courant Number Recovery Rate	1.1
CFL limiting: Target AMG Cycles	4

Creating the Downforce Report, Monitor, and Plot

The aim of this adjoint flow analysis is to investigate the design with a view to increase the amount of downforce that is generated by the wing.

To monitor the downforce, create a force report:

1. Right-click the **Reports** node and select **New Report > Force**.
2. Rename the report to `Downforce`.
3. Select the **Downforce** node and set *Direction* to `[0.0, 0.0, -1.0]`.
4. Edit the *Parts* property, expand the **Regions > Fluid Domain > Boundaries** node, and select **Wing End Plate**, **Wing Lower Element**, and **Wing Upper Element**.



5. Click **OK**.
6. Right-click **Reports > Downforce** and select **Create Monitor and Plot from Report**.
7. Right-click **Plots > Downforce Monitor Plot** and choose **Open**.

Modifying Solver Parameters

When performing an adjoint flow analysis, aim for convergence levels that approach machine precision.

This precision level means that residuals drop to a level around 1.0E-12. For most industrial cases, convergence to machine accuracy is difficult to achieve. In this tutorial, you apply some common steps for improving the convergence level.

Setting Stopping Criteria

For this simulation, the primal solver is left to run for 350 iterations.

To set the stopping criteria:

- Select **Stopping Criteria** > **Maximum Steps** and set *Maximum Steps* to `350`.



Running the Primal Solver

Run the simulation.

To run the simulation:

1. In the **Solution** toolbar, click (**Run**).
2. When the solver is finished, save the simulation.

   The *Residuals* plot is displayed below:



Note: Your residuals may differ slightly, depending on the accuracy of your machine.

Activating the Adjoint Model

Activate the Adjoint Flow physics model.

To activate the model:

1. Right-click the **Continua** > **Physics 1** node and choose **Select Models...**.
   The *Physics 1 Model Selection* dialog appears.
2. Select **Adjoint Flow** from the *Optional Models* group box.
3. Click **Close**.
4. To view the active physics models, open the **Physics 1** > **Models** node.



Creating the Adjoint Cost Function

The adjoint cost function specifies which engineering quantity is the subject of the sensitivity analysis.

In this tutorial you want to know how to increase the amount of downforce that is generated by the wing. A force cost function provides this information.

To create a force cost function:

1. Right-click the **Adjoint Cost Functions** node and select **New > Force Cost Function**.



   A **Force** sub-node is created.
2. Rename the force cost function to `Downforce`.
3. Select the **Downforce** node and set *Force Report* to **Downforce**.



The adjoint force cost function adds the following monitors to the **Monitors** node: **Downforce::P**, **Downforce::U**, **Downforce::V**, and **Downforce::W**.

Running the Adjoint Flow Solver

The adjoint flow solver is run from the Solvers node, as opposed to the primal flow solver which is run using the Run button in the Solution toolbar.

The primal flow solution is frozen while the adjoint flow solver runs. Avoid stepping or running the primal flow solver after running the adjoint flow solver.

To run the adjoint flow solver:

1. Select **Stopping Criteria** > **Maximum Steps** and set *Maximum Steps* to `600`.
2. Right-click the **Solvers** > **Adjoint Flow** node and select **Run**.



The adjoint solver runs.

3. When the solver is finished, save the simulation as `adjointWing_adjoint.sim`.

Visualizing Sensitivity of Downforce with Respect to the Flow

Create a scalar scene to display the sensitivity of downforce with respect to the flow field.

To create a scalar scene:

1. In the **Vis** toolbar, click (**Create/Open Scenes**) and select **Scalar**.
2. Press the `<S>` key to display the side view.
3. Click **scene/plot**.
4. Right-click **Displayers** > **Scalar 1** > **Parts** and select **Edit...**.
   The *Edit* dialog appears.
5. Select **Regions** > **Fluid Domain** > **Boundaries** > **Symmetry**.
6. Click **OK**.
7. Select the **Scalar 1** > **Scalar Field** node and set *Function* to **Adjoint** > **Downforce** > **Adjoint of Downforce w.r.t. X-momentum**.
8. Select the **Scalar 1** node and set *Contour Style* to **Smooth Field + Lines**.
9. In the **Vis** toolbar, click (**Save-Restore-Select Views**) and select **Projection Mode** > **Parallel**.

The scalar scene appears as follows:



This scene displays the sensitivity of the desired adjoint cost function with respect to the flow field. In this case, downforce is the desired adjoint cost function. Directing X-momentum towards areas with positive values will increase the downforce. Similarly, directing X-momentum away from areas with negative values will also increase downforce.

Summary

This tutorial demonstrated how to run an adjoint flow analysis on a dual element wing for a race car.

This tutorial has introduced the following STAR-CCM+ features:

- Loading a simulation file and visualizing a volume mesh.
- Checking the volume mesh for bad cells.
- Increasing the number of times the mesh optimization cycle runs.
- Increasing the mesh quality threshold.
- Running the primal solver to machine precision.
- Setting the stopping criteria.

- Reporting, monitoring and plotting.
- Setting up Grid Sequencing and the Expert Driver.
- Selecting the physics model for adjoint flow analysis.
- Setting up the Adjoint cost function for force.
- Running the Adjoint Flow solver.
- Calculating the Adjoint Mesh.
- Visualizing the adjoint solution.

Adjoint Flow: Wing Shape Optimization

This tutorial shows how to exploit an adjoint flow analysis to set up a shape optimization with respect to a specified cost function.

The tutorial continues from the final solution obtained in the **Adjoint Flow Solver: External Flow over a Dual Element Wing** tutorial. In particular, the tutorial shows how the shape of the lower element can be modified to increase the downforce generated by the wing.



The shape optimization process uses the morphing procedure to deform the geometry of interest. Specified control points are displaced by amounts calculated using mesh sensitivities that the Adjoint Mesh solver provides. The displacement of the control points controls the deformation calculated by the morpher. The displaced positions of the specified control points are set to maximize/minimize the cost function of interest.

For more information about the shape optimization workflow using the Adjoint solvers see the User Guide - Setting up a Shape Optimization.

Contents

Prerequisites
Loading the Simulation
Setting Mesh Deformation Boundary Conditions
Creating the Initial Control Points
Computing the Mesh Sensitivity
Performing the First Shape Optimization

Running The Simulation
Visualizing the Results
Performing Multiple Optimization Cycles
Summary

Prerequisites

The instructions in the Adjoint Flow: Wing Shape Optimization tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+
Running an adjoint flow analysis	Adjoint Flow Solver: External Flow over a Dual Element Wing

Note: You require a double-precision version of STAR-CCM+ to complete this tutorial.

Loading the Simulation

For this tutorial, you are provided with a simulation file that contains pre-defined objects.

Due to the large number of cells, this tutorial simulation should be run in parallel mode on at least 4 cores.

To load the simulation file:

1. Start a double precision version of STAR-CCM+ in a manner that is appropriate to your working environment and select **Load Simulation** from the menu bar.
2. Load the `adjointWing_adjoint.sim` file that you created during the Adjoint Flow Solver: External Flow over a Dual Element Wing tutorial. Make sure that you specify multiple cores.

Setting Mesh Deformation Boundary Conditions

The morpher redistributes mesh vertices according to the design points and their associated displacements. The associated displacements are generated during the optimization cycle on the basis of mesh sensitivity.

The morpher algorithm and its adjoint require each symmetry boundary to have no more than one planar surface:

1. Right-click the **Regions > Fluid Domain > Boundaries > Far Field** node and select **Split by Angle**.

   The *Split Boundaries by Angle* dialog appears.
2. Accept the default settings by clicking **OK**.



This tutorial focuses on the optimization of the shape of the wing lower element, whereas the shape of other wing components is fixed.

To specify the morphing conditions set the appropriate *Method* for the **Morpher** node within the **Physics Conditions** node of each boundary.

3. Select the **Regions** > **Fluid Domain** > **Boundaries** > **Wing Lower Element** > **Physics Conditions** > **Morpher** node and set *Method* to **Floating**.



4. Make sure that for the **Inlet, Outlet**, **Wing End Plate** and **Wing Upper Element** boundaries the *Method* for the **Morpher** node is set to **Fixed**.
5. Save the simulation as `adjointWing_shapeOptimization.sim`.

Creating the Initial Control Points

You create initial control points at the beginning of the shape optimization which the adjoint solver uses to compute the mesh sensitivity.

The displacements of these points to new optimized positions are then used to calculate the corresponding mesh deformation, that is, the optimized shape.

You start by modifying **Geometry Scene 1** in the following way:

1. Right-click the **Geometry Scene 1** node and select **Open**.
2. Right-click the **Geometry Scene 1** > **Displayers** > **Geometry 1** > **Parts** node and select **Edit...**.

   The *Edit* dialog appears.
3. In the *Edit* dialog deselect the **Derived Parts** node.

4. Expand the **Regions** > **Fluid Domain** > **Boundaries** node and select **Wing End Plate, Wing Lower Element** and **Wing Upper Element**.



5. Click **OK**.
6. Right-click the **Geometry Scene1** > **Outline 1** node and select **Toggle Visibility**.

To create the initial control points:

7. Right-click the **Control Points** node and select **New Control Points** > **Line...**.



8. In the *Create Control Points along Line* dialog, set the following properties:

Property	**Value**
Point 1	`[0.1 m, 0.1 m, 0.01 m]`
Point 2	`[0.1 m, 0.6 m, 0.01 m]`
Number	`6`
Display	**Existing Displayer > Geometry 1**

The completed dialog is displayed below.



9. Click **Create**.
10. Create another set of control points along a line and set the following properties:

Property	**Value**
Point 1	`[0.1 m, 0.1 m, 0.15 m]`
Point 2	`[0.1 m, 0.6 m, 0.15 m]`
Resolution	`6`
Display	**Existing Displayer > Geometry 1**

11. Click **Create**, then **Close**.

You can see the locations of the created initial control points in *Geometry Scene 1*.



12. Rename the **Line Control Point Region** node to `Lower` and rename the **Line Control Point Region 2** node to `Upper`.

Computing the Mesh Sensitivity

When the mesh sensitivity computation is performed, the Adjoint Mesh solver calculates the variation of the cost function of interest, in each spatial direction, at each control point position.

To compute the mesh sensitivity:

1. Select the **Solvers > Mesh Deformation** node and make sure that *Enable Automatic Thinning* is activated.

   Although you activate **Automatic Thinning** under the **Mesh Deformation** node, this property also affects the computation of the mesh sensitivity. You thin out the number of control points, thus speeding up the mesh sensitivity computation.
2. Right-click the **Solvers > Adjoint Mesh** node and select **Compute Mesh Sensitivity**.



The mesh sensitivity computation starts and takes several seconds to complete.

To view the mesh sensitivity, visualize the **Adjoint of Downforce w.r.t. Position** field:

3. Right-click the **Geometry Scene 1 > Displayers** node and select **New Displayer > Vector**.
4. Edit the **Diplayers** node and set the following properties:

Node	Property	Setting
Vector 1		
Parts	*Parts*	**(Control Points) [Lower, Upper]**

Node	Property	Setting
Vector Field	*Function*	**Adjoint of Downforce w.r.t. Position**
Glyph	*Vector Scale*	**Absolute**
	Vector Style	**2D Filled**



You can now use the computed mesh sensitivities for each control point to deform the mesh.

Performing the First Shape Optimization

After computing the mesh sensitivity with respect to the control points, you continue by deforming the mesh accordingly. You now perform the first step of the shape optimization.

Generally, you can perform a shape optimization by employing an external optimization tool that uses the mesh sensitivities computed by the Adjoint Mesh solver. This optimization tool calculates the new positions of the control points based on an optimization algorithm of your choice.

For this tutorial, you use a first-order Steepest Descent algorithm with a step size of 0.0001 to update the control point positions. To update the control point positions for the shape optimization cycles, you take into account the total displacement of a control point from its original position. You perform this optimization step completely within STAR-CCM+. You achieve this by associating a user-defined field

function that defines the Steepest Descent deformation with the morpher. This user-defined field function uses a built-in field function for the cumulative morpher displacement.

To perform the shape optimization:

1. Right-click the **Tools > Field Functions** node and select **New > Vector**.
2. Rename this field function to `Displacement` and set the following properties:

Property	Setting
Function Name	`Displacement`
Definition	`$ $CumulativeMorpherDisplacement + 1.e-04*$$ {Adjoint1::CoordAdjoint}`

   The field function is defined such that it multiplies the step size of 0.0001 with the **Adjoint of Downforce w.r.t. Position** field function. You obtain this field by computing the mesh sensitivity of the **Downforce** cost function on the control points (see Computing the Mesh Sensitivity).

3. Multi-edit the **Control Points > Lower** and the **Control Points > Upper** nodes and set the following properties:

Node	Property	Setting
Physics Values		
Morpher: Displacement	*Method*	**Field Function**
Field Function	*Vector Function*	**Displacement**

4. Right-click the **Mesh Deformation** node and select **Deform Mesh**.



To display the improved wing lower element shape, update the **Geometry Scene 1**:

5. Right-click a blank area in the **Geometry Scene 1** and select **Refresh**.

The shape obtained after the first optimization cycle is displayed in the following image.



Running The Simulation

Run the simulation to observe how the downforce generated by the wing increases due to the optimization of the wing lower element shape.

1. Right-click the **Plots** > **Downforce Monitor Plot** node and select **Open**.
2. Select the **Stopping Criteria** > **Maximum Steps** node and set the *Maximum Steps* to `1000`.

To run the primal solver:

3. Click (**Run**) to run the simulation.
   The *Run Simulation* dialog appears.
4. Click **Yes** to confirm.



To run the adjoint solver:

5. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `1300`.

6. Right-click the **Solvers** > **Adjoint Flow** node and select **Run**.
7. When both solvers are complete, save the simulation.

Visualizing the Results

Visualize how the first shape optimization cycle improves the downforce.

1. Right-click the **Plots** > **Downforce Monitor Plot** node and select **Open**.
2. In the toolbar, click the **Zoom Selected Plot** button.
3. In the Downforce Monitor Plot, draw a box around the area where there is an increase in downforce (at Iteration 600).



Performing Multiple Optimization Cycles

You can perform as many optimization cycles as you think appropriate, repeating the process at every cycle.

The following steps outline the process of each optimization cycle:

- Compute the Mesh sensitivity

- Deform the Mesh
- Run the Primal Solver
- Run the Adjoint Solver

Perform a second shape optimization:

1. Compute the mesh sensitivity:
   a) Right-click the **Solvers** > **Adjoint Mesh** node and select **Compute Mesh Sensitivity**.
2. Deform the mesh:
   a) Right-click the **Mesh Deformation** node and select **Deform Mesh**.
3. Run the Primal solver:
   a) Set the **Stopping Criteria** > **Maximum Steps** to `1700`.
   b) Click (**Run**).
4. Run the Adjoint Solver:
   a) Set the **Stopping Criteria** > **Maximum Steps** to `2000`.
   b) Right-click the **Solvers** > **Adjoint Flow** node and select **Run**.

The Downforce Plot history is shown below.

The following image shows the geometry after the second optimization cycle.



Summary

This tutorial presented the general workflow to set up a shape optimization within the adjoint framework.

The following concepts have been introduced:

- Creating control points.
- Computing the mesh sensitivity.
- Deforming the mesh using the morphing procedure.
- Running the solver to observe the cost function improvement.
- Performing additional optimization cycles.

Steady Multi-Component Flow: Dilution Pipe

This tutorial demonstrates the process of setting up and analyzing a turbulent, multi-species flow problem. Constant, steady streams of air and methane gas flow into the solution domain at 5 m/s and 10 m/s, respectively, and mix within the pipe.



Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Geometry
Renaming Regions and Boundaries
Scaling the Mesh
Setting up the Models
Setting Material Properties
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Stopping Criteria

Visualizing the Solution
Running the Simulation
Visualizing the Results
Adding Streamlines
Summary

Prerequisites

The instructions in the Steady Multi-Component Flow: Dilution Pipe tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the New Simulation option from the menu bar.

Continue by importing the mesh and naming the simulation. A hexahedral cell mesh has been prepared for this analysis and is stored in a `ccm` format file.

1. Select **File** > **Import** > **Import Volume Mesh** from the menu bar.
2. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select file `dilPipe.ccm`.
3. Click **Open** to start the import.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
4. Finally, save the new simulation to disk under file name `dilPipeSteady.sim`.

Visualizing the Geometry

In the geometry scene created during mesh import, the internal mesh features are not visible.

To examine these features:

1. Select the **Scenes > Geometry Scene 1 > Displayers > Geometry 1 > Parts** node and click the *Parts* property.

2. In the dialog that appears, click (**List View**), and then clear **Default_Fluid: symmetry plane**.



3. Click **OK.**
4. Rotate the *Geometry Scene 1* display through approximately 180° so that the mesh is displayed in a way that allows you to see both the inner pipe and the orifice plate.

The inner pipe and orifice plate form part of the same boundary as the external wall. This is acceptable in this case because you are not modeling heat transfer through them. If it was necessary to do so, split the mesh into separate regions with baffle interfaces representing the internal features instead of wall boundaries.

Renaming Regions and Boundaries

All regions and boundaries are given default names during the mesh import. Give them more appropriate names.

1. Rename the **Regions** > **Default_Fluid** node to `Fluid`.
2. Rename the **Regions** > **Boundaries** > **Fluid** > **Default_Boundary_Region** node to `Wall`.

Scaling the Mesh

The original mesh was built in inches, but no scaling factor was applied when it was saved to the `.ccm` file. Fix this discrepancy.

The **Mesh** > **Diagnostics** menu item may be used to determine the current mesh dimensions, but this information is already available from the mesh import process and is displayed in the *Output* window.

The mesh region at import extends from 0 to 5.0 in the x-direction, 0 to 5.5 in the y-direction and 0.0 to 1.5 in the z-direction. Without any further action, these units are interpreted by STAR-CCM+ as meters. Scale the region by 0.0254 so that the dimensions are correctly specified.

1. Select **Mesh** > **Scale Mesh** from the menu bar.
2. In the *Scale Mesh* dialog, make sure that the Fluid region is selected and then set *Scale Factor* to `0.0254`.



3. Click **Apply** to scale the domain and then click **Close**.

   The mesh region reduces in size.

4. Click (**Reset View**) in the visualization toolbar to restore the previous viewing distance for the scaled domain in the display.

To verify that the scaling has been applied, you can perform the mesh diagnostics check and review the output values.

Setting up the Models

Models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this example, the flow is turbulent and compressible. The **Segregated Flow** model is used together with the default K-Epsilon turbulence model.

The default continuum is automatically named **Physics 1** when the mesh is imported. To give the continuum a more appropriate name:

1. Rename the **Continua** > **Physics 1** node to `Dilution Pipe`

To select the appropriate physical models for the fluid region:

2. Right-click the **Dilution Pipe** node and choose item **Select models**

The *Physics Model Selection* dialog guides you through the model selection process by showing only options that are appropriate to the choices already made.

3. Make sure that **Three-Dimensional** and **Gradients** are selected.
4. Select **Multi-Component Gas** in the *Material* box.
5. Select **Non-Reacting** in the *Reaction Regime* box.
6. Select **Segregated Flow** in the *Flow* box.
7. Select **Ideal Gas** in the *Equation of State* box.
8. Select **Segregated Fluid Temperature** option in the *Energy* box.
9. Select **Steady** in the *Time* box.
10. Select **Turbulent** in the *Viscous Regime* box.
11. Select **K-Epsilon Turbulence** in the *Reynolds-Averaged Turbulence* box.
12. Click **Close**.

    In the *dilPipeSteady* window, the color of the **Dilution Pipe** node has turned from gray to blue to indicate that models have been activated.
13. To see the selected models, expand the **Dilution Pipe** > **Models** node.



14. Save the simulation.

Setting Material Properties

Material properties must now be assigned to the two fluid components.

For air, the default properties set up by STAR-CCM+ is used. For methane, specify a molecular weight of 16 kg/kmol, a specific heat of 2230 J/kg-K, a thermal conductivity of 0.038 W/m-K and a molecular viscosity of 1.175E-5 Pa-s. The gas mixture density is calculated using the Ideal Gas Law, as specified in Setting up the Models.

1. Right-click the **Models > Multi-Component Gas > Gas Components** node and choose **Select Mixture Components**.
2. In the *Select Mixture Components* dialog, expand the **Material Databases**, **Standard**, and **Gases** nodes to reveal the available gaseous components.
3. Select the components **Air** and **CH4 (Methane)**.



4. Click **Apply** and then click **Close** when complete.

Nodes corresponding to these components appear within the **Gas Components** node.



5. Rename the **CH4** node to `Methane`.
6. If components are not in the order as shown in the above figure, right-click the **Gas Components** node and select **Reorder Mixture Components**. In the *Reorder Mixture Components dialog*, move **Methane** to the bottom of the list.



7. Open the **Material Properties** nodes for Methane and enter the following values:

Node	Property	Value
Dynamic Viscosity > Constant	*Value*	`1.175E-5` Pa-s
Molecular Weight > Constant	*Value*	`16` kg/kmol
Specific Heat > Constant	*Value*	`2230` J/kg-K
Thermal Conductivity > Constant	*Value*	`0.038` W/m-K

8. Save the simulation.

Setting Initial Conditions

Set the initial conditions within the dilution pipe.

1. Select the **Dilution Pipe > Initial Conditions > Species Mass Fraction > Constant** node and click the box next to the *Value* property.
2. In the *Constant - Value* dialog, enter `1.0` for the *Air* mass fraction.



3. Click **OK**.
4. Select the **Static Temperature > Constant** node and change the *Value* to `293` K.



5. Save the simulation.

Setting Boundary Conditions and Values

Specify boundary conditions at the two inlets as well as the pressure boundary.

Start with the inner (methane) inlet:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Inner_inlet** > **Physics Values** node and set the following boundary conditions:

Node	Property	Value
Species Mass Fraction > Constant	*Value*	`0.0, 1.0`
Static Temperature > Constant	*Value*	`323` K
Turbulence Intensity > Constant	*Value*	`0.1`
Turbulent Viscosity Ratio > Constant	*Value*	`100`
Velocity Magnitude > Constant	*Value*	`10` m/s

Define conditions at the air inlet.

2. Edit the **Outer_inlet** > **Physics Values** node and set the following boundary conditions:

Node	Property	Value
Static Temperature > Constant	*Value*	`1.0, 0.0`
Static Temperature > Constant	*Value*	`293` K

Node	Property	Value
Turbulence Intensity > Constant	*Value*	`0.1`
Turbulent Viscosity Ratio > Constant	*Value*	`40`
Velocity Magnitude > Constant	*Value*	`5` m/s

Similar properties must be assigned to the pressure boundary. These are only used at locations where inflow occurs.

3. Edit the **pressure > Physics Values** node and set the following boundary conditions:

Node	Property	Value
Species Mass FractionConstant	*Value*	`1.0, 0.0`
Static Temperature > Constant	*Value*	`308` K
Turbulence Intensity > Constant	*Value*	`0.1`
Turbulent Viscosity Ratio > Constant	*Value*	`40`

The default **Pressure** value is suitable for this tutorial and therefore does not need changing.

This completes the specification of boundary conditions and values.

4. Save the simulation.

Setting Stopping Criteria

Run the calculation for 300 iterations, which is sufficient to achieve a steady-state solution.

1. Select the **Stopping Criteria > Maximum Steps** node and set the *Maximum Steps* to `300`.
2. Save the simulation.

Visualizing the Solution

View the temperature of the fluid region as the solution develops. Start by creating a scalar scene.

1. Create a scalar scene.

The *Scalar Scene 1* display appears.

2. Click (**Save-Restore-Select Views**) on the **Vis** toolbar and select **Views** > **-Z** > **Up +Y**.



3. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node and change the *Contour Style* to **Smooth Filled**.
4. Select the **Scalar 1** > **Parts** node and click the right side of the *Parts* property.
5. In the dialog that appears, click **Select All,** expand the **Regions** and **Fluid** nodes and make sure that the parts belonging to **Fluid** are all selected as shown below.



6. Click **OK**.
7. Select the **Scalar Field** node and click the right side of the *Function* property.
8. In the dialog that appears, select the **Temperature** field function node.
9. Click **OK**.
10. Save the simulation.

Running the Simulation

Run the simulation to obtain the residuals plot and the scalar scene.

To run the simulation:

1. Click (**Run**) in the top toolbar.

   If you do not see this button, use the **Solution** > **Run** menu item. You could also display the **Solution** toolbar by selecting **Tools** > **Toolbars** > **Solution** and then clicking the toolbar button.

   If necessary, click the *Residuals* tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.

The tabs at the top of the *Graphics* window make it possible to select any one of the active displays for viewing.

The Residuals display is created automatically and shows the solver's progress.

2. Click the *Scalar Scene 1* tab to watch the temperature distribution being updated.



During the run it is possible to stop the solution process by clicking the (**Stop**) button in the toolbar. If you do halt the simulation, it can be continued later by clicking the (**Run**) button. If left alone, the simulation continues until 300 iterations are complete.

3. Save the simulation at the end of the run.

Visualizing the Results

View and modify the scalar scene, noting temperature, methane concentration, and pressure.

1. Click the *Scalar Scene 1* tab to see the temperature contours for the converged solution.



This scene can be altered to display contours of other scalars, by right-clicking the scalar bar in the display.

2. In the pop-up menu that appears, select **Mass Fraction** > **Methane**.

A contour plot of methane mass fraction is displayed.



3. Right-click the scalar bar and select **Pressure**.

In this case, however, you are going to produce a pressure plot on the entire external surface of the dilution pipe. This can be done by applying a transformation to the current contour plot, which relates to the pipe's symmetry plane.

4. Select the **Scalar 1** node and set **symmetry plane 1** as the *Transform* property.



This adds a reflection about the symmetry plane of the current scalar plot and make sure that the *Smooth Shade* property is selected.

This removes any shading from the mesh surface display, due to the shape of the mesh rather than the shape of the overall geometry. The edge lines can also be reflected about the symmetry plane by changing the *Transform* property of the **Outline** node.

5. Remove the **Fluid: wall** and **Fluid: symmetry plane** parts from the outline to complete the plot shown below.



6. Save the simulation.

Adding Streamlines

Streamlines can be added into any scene by creating a 'streamline' derived part.

1. Make the *Geometry Scene 1* display active by clicking its tab.
2. In the *dilutionPipeSteady* window, right-click the **Derived Parts** node and select **New Part** > **Streamline**.

An interactive in-place dialog appears to allow definition of the derived part.



3. Make sure that **Fluid** is selected in the **Input Parts** box.
4. In the *Vector Field* box, make sure that **Velocity** is selected.
5. For the *Seed Mode*, make sure that **Part Seed** is selected.
6. Select **Inner_inlet** and **Outer_inlet** as the Seed Parts.



7. Click **OK**.
8. Make sure that **New Streamline Displayer** is selected in the Display section.

The completed dialog appears as shown below.


Create Streamline
Input Parts
[Fluid]
Select...
Snap to Part
Seed Mode
Part Seed
Seed Parts
[Fluid: Inner_inlet, Fluid: Outer_inlet]
Select...
Part U-Resolution
5
Part V-Resolution
5
Vector Field
Velocity
Select...
Direction
Forward
Display
No Displayer
New Geometry Displayer
New Scalar Displayer
New Vector Displayer
New Streamline Displayer
Existing Displayer
Outline 1
Create
Close
Help

9. Click **Create** and then click **Close**.

This procedure creates streamlines originating at the methane inlet and the air inlet.

10. Right-click the scalar bar in the display.
11. In the pop-up menu that appears, select **Velocity: Magnitude**.

The high cell density near the center of the inner inlet means that streamlines starting at this boundary are all bunched together. To rectify this, you can randomize their starting positions.

12. Select the **Derived Parts** > **streamline** > **Source Seed** node and activate *Randomize*.

The final streamline plot is shown below.



13. Save the simulation.

Summary

This tutorial demonstrated the process of setting up and analyzing a turbulent, multi-species flow problem.

It illustrated the following STAR-CCM+ features:

- Defining models for multi-component flow.
- Defining the material properties required for the selected models.
- Running the solver for a set number of iterations.
- Analyzing the results using the built-in visualization facilities.

Unsteady Multi-Component Flow: Dilution Pipe

This tutorial restarts from the results that were obtained in the Steady Multi-Component Flow: Dilution Pipe. It solves the same problem except that, in this case, the velocities at the two inlet boundaries vary with time over a period of 2 seconds.

The velocity-time profiles at the inlets are shown below.



Contents

Prerequisites
Loading an Existing Simulation
Setting up the Unsteady Flow Model
Creating Field Functions
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Autosaving the Simulation
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Unsteady Multi-Component Flow: Dilution Pipe tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading an Existing Simulation

To work with a variation of the case from the steady dilution pipe tutorial, load the simulation file `dilPipeSteady.sim`.

To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file.

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment.
2. To open an existing simulation file, select **File > Load Simulation** from the menus.

   The *Load Simulation* dialog appears.
3. Click **Browse** and then locate the `dilPipeSteady.sim` file that was created in the Steady Multi-Component Flow: Dilution Pipe tutorial.
4. Select file `dilPipeSteady.sim` and then click **Open**.
5. Click **OK** in the *Load Simulation* dialog.

On the left side of the STAR-CCM+ workspace, the simulation object tree that was created in the Steady Multi-Component Flow: Dilution Pipe appears in the *Explorer* pane with the name *dilPipeSteady*.



The name of the simulation that is displayed in the title bars of the *Explorer*, *Properties*, and *Output* windows changes accordingly.

Setting up the Unsteady Flow Model

Select the physics models for the new simulation.

To preserve the original solution, save the simulation with a new name.

1. To create a file for the new simulation, select **File > Save As**.
2. In the *Save* dialog, navigate to the directory where you want to locate the file.
3. Type `dilPipeUnsteady` into the **File Name** text box.
4. Click **Save**

Most of the models have already been selected in the Steady Multi-Component Flow: Dilution Pipe tutorial and so do not need to be specified. However, change the analysis type to unsteady.

5. Right-click the **Continua > Dilution Pipe** node and choose **Select models...**

In the *Physics Model Selection* dialog:

6. Deactivate the **Steady** model from the *Enabled Models* group box.
   The *Time* group box reappears on the left-hand side.
7. Select the **Implicit Unsteady** model.

8. Click **Close**.
9. To see the selected models, expand the **Dilution Pipe** > **Models** node.



10. Save the simulation.

Creating Field Functions

Define the inlet velocity profiles as linear functions of time and enter them into STAR-CCM+ using the field function utility.

The required functions are:

- Inner Inlet: $v = 10 - 4t$
- Outer Inlet: $v = 5 + 2t$

To create the corresponding field functions:-

1. Right-click the **Tools** > **Field Functions** node and select **New** > **Scalar**.
2. Rename the field function **User Field Function 1** to `Inner Inlet Velocity`.
3. Select the **Inner Inlet Velocity** node and set *Definition* to `10-4*$Time`.

4. Create another scalar field function and rename it `Outer Inlet Velocity`.
5. Select the `Outer Inlet Velocity` node and set *Definition* to `5 + 2 * $Time`.
6. Save the simulation.

Setting Boundary Conditions and Values

The only boundary conditions that need changing are the inlet velocity magnitudes. Specify these boundary conditions using the field functions created in section Creating Field Functions.

To set the inner inlet boundary condition:

1. Select the **Regions** > **Fluid** > **Boundaries** > **Inner_inlet** > **Physics Values** > **Velocity Magnitude** node and set *Method* to **Field Function**.



2. Select the **Velocity Magnitude** > **Field Function** node and click the right side of the *Scalar Function* property.
3. In the dialog that appears, select the **Inner Inlet Velocity** field function node.
4. Click **OK**.
5. Open the **Boundaries** > **Outer_inlet** node. Following the same procedure as described above, define the velocity magnitude at the outer inlet by the **Outer Inlet Velocity** field function.



6. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the time-step length and the run time.

The simulation is to restart from the steady-state result that was obtained in the Steady Multi-Component Flow: Dilution Pipe and then run for 2 seconds of physical time. A time-step of 0.025 seconds is used, so the analysis requires 80 steps.

To set the time-step:

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `0.025` s.

To specify the run time:

2. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `2.0`.
3. Save the simulation.

Autosaving the Simulation

Use the Autosave feature to save the entire simulation results every 0.5 s. This feature enables you to examine these results during the run rather than just at the end.

To activate the Autosave feature:

1. Select **File** > **Auto Save**
2. In the *Auto Save* dialog, specify the properties that are shown below.



3. Click **Close**.

Running the Simulation

You can run this simulation without initializing it, because you are restarting the analysis from a previous solution.

To display *Scalar Scene 1*:

1. Right-click the **Scenes** > **Scalar Scene 1** node and select **Open**.

The pressure profile on the external surface of the dilution pipe is displayed in the *Graphics* window.



2. Right-click the legend color scale in the *Graphics* window and change the displayed scalar to **Mass Fraction of Methane**.
3. To remove the reflected half of the problem geometry from the display, select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node and set *Transform* to **Identity**.
4. Select the **Outline 1** node and set *Transform* to **Identity**.
5. To change the view, click in the toolbar and select **View** > **-Z** > **Up +Y**.
6. To run the simulation, click (**Run**) in the top toolbar.

   If necessary, click the *Residuals* tab to bring the Residuals plot into view. The tabs at the top of the *Graphics* window make it possible to select any of the active displays for viewing.

   During the run, it is possible to stop the process by clicking (**Stop**) in the toolbar. If you do halt the simulation, it can be continued later by clicking (**Run**). If left alone, the simulation runs for 2.0 s of physical time.

The Residuals display is created automatically and shows the progress of the solver.

7. To watch the solution as it gets updated, click the *Scalar Scene 1* tab.



8. Save the simulation.

Visualizing the Results

Examine the scalar and vector scenes for the methane mass fraction.

At the end of the run, the methane mass fraction contour plot looks as shown below.



Plot the velocity vectors after 2.0 seconds:

1. Right-click the **Scenes** node and select **New Scene** > **Vector**.
2. Change the view to look down the -Z axis with the +Y axis pointing upwards.
3. Select the **Scenes** > **Vector Scene 1** > **Displayers** > **Vector 1** > **Parts** node and click ... (**Custom Editor**) for the *Parts* property.
4. In the dialog that appears, expand the **Regions** node, select the **Fluid** node, and click **OK**.



The resulting vector plot looks as shown below.



To examine the results after 0.5 s:

5. Select **Load Simulation** from the **File** menu and load the dilPipeUnsteady@5.00000e-001.sim file.



In the new simulation tree:

6. Open the **Scenes** node and double-click the **Scalar Scene 1** node to open the methane mass fraction contour plot in the *Graphics* window.



To view the velocity vectors after 0.5 s, you could create a vector scene as was done earlier. However, it is simpler to copy the vector scene from one simulation to another, as follows:

7. Open the *dilPipeUnsteady@2.00e+00* window and right-click the **Vector Scene 1** node.
8. Select **Copy** from the menu.
9. Return to the simulation tree in the *dilPipeUnsteady@5.00e-01* window and right-click the **Scenes** node.
10. Select **Paste** from the menu.

A new scene named *Vector Scene 1* appears in the tree. To display the scene in the *Graphics* window, double-click the node.



Note how the varying inlet velocity has changed the velocity field between $t = 0.5s$ and $t = 2.0s$.

11. Save the simulation.

Summary

This tutorial re-examined the case presented in the Steady Multi-Component Flow tutorial, but with an unsteady flow.

It illustrated the following STAR-CCM+ features:

- Selecting an unsteady simulation.
- Creating field functions.
- Setting boundary conditions and values.
- Setting solver parameters.
- Using the Autosave feature to store intermediate results during the run.
- Running the solver for a set time.
- Analyzing results.

Porous Resistance: Isotropic Media

This tutorial models flow through a catalyst geometry.

The geometry is shown below. The catalyst material at the center of this geometry is modeled as a cylindrical porous region of length 0.03 m and diameter 0.1 m.



In the porous region, the theoretical pressure drop per unit length can be determined using the equation:

$$\frac{\Delta p}{L} = -\left(P_i \left| v \right| + P_v\right) v$$

Where v is the superficial velocity through the medium and P_i, P_v are coefficients defining the porous resistance, known as the inertial resistance and viscous resistance, respectively.

Values for the resistance coefficients can be measured experimentally or derived using various empirical relationships, depending on the exact nature of the problem. In this case, P_i = 25 kg/m^4 and P_v = 1500 kg/m^{3}s. These values are roughly what you would expect from an isotropic porous catalyst.

Contents

Prerequisites

Importing the Mesh and Naming the Simulation
Visualizing the Imported Mesh
Scaling the Mesh
Setting Up the Models
Setting Boundary Conditions and Values
Specifying Porosity Coefficients
Monitoring the Run Progress
Setting Stopping Criteria
Visualizing the Solution
Running the Simulation
Visualizing Results
Validating the Results
Summary

Prerequisites

The instructions in the Porous Resistance: Isotropic Media tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Import a polyhedral mesh from `doc/startutorialsdata/incompressibleFlow/data`, and save and name the new simulation.

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment.
2. Select the New Simulation option from the menu bar.

Continue by importing the mesh and naming the simulation. A polyhedral cell mesh has been prepared for this analysis and is stored in a `ccm` format file.

3. Select **File > Import > Import Volume Mesh** from the menu bar.

4. In the *Open* dialog, navigate to the `doc/startutorialsdata/incompressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select the file `catalyst.ccm`.
5. Click **Open** to start the import.

   STAR-CCM+ provides feedback on the import process in the *Output* window. Two mesh regions named **Fluid** and **Porous** are created under the **Regions** node to represent the two parts of the solution domain.
6. Save the new simulation as `isotropicPorousMedia.sim`.

Visualizing the Imported Mesh

A geometry scene was automatically created after the import process. Examine the *Geometry Scene 1* display in the *Graphics* window. Initially all parts of the mesh are shown as solid, colored surfaces.

1. Click one of the parts.



   The boundary selected is highlighted and a label with its name appears on the display.

In the *Simulation* window, the node that corresponds to this boundary is also highlighted.



Similarly, if any other boundary node is selected in the *Simulation* window, the surface corresponding to that boundary are highlighted in the *Graphics* window.

2. Left-click anywhere in the *Graphics* window and drag the mouse in any direction. The view rotates about the location where the left mouse button is clicked (marked by a small square).
3. Similarly, right-click and drag the mouse in any direction to pan or shift the display.
4. To zoom in and out, middle-click in the *Graphics* window and move the mouse up and down.
5. Rather than use the mouse facilities to adjust the view

   approximately, click the button at the top of the *Geometry Scene 1* window and select **View** > **+X +Y +Z** > **Up +Z**, as shown below.



6. Use the mouse buttons to zoom in and position the part so that it is similar to that shown below.



Scaling the Mesh

The original mesh was not built to scale and therefore requires scaling so that the dimensions of the porous region are 0.1 x 0.03 x 0.1 meters.

To determine the current dimensions of the porous region:

1. Select **Mesh** > **Diagnostics**.
2. In the *Mesh Diagnostics* dialog, deselect **Fluid** and then click **OK**.



The dimensions of the porous region are displayed in the *Output* window as approximately -0.5 to 0.5 in the x-direction, 0.435 to 0.735 in the y-direction and approximately -0.5 to 0.5 in the z-direction.

Without any further action, these values are interpreted by STAR-CCM+ as specified in meter units. The region therefore needs to be scaled by a factor of 0.1 so that the dimensions are correctly interpreted.

3. Select **Mesh** > **Scale Mesh** from the menu bar.
4. In the *Scale Mesh* dialog, select the **Fluid** and **Porous** regions.
5. Enter a value of `0.1` for the scale factor.



6. To scale the domain, click **Apply**.
   The mesh region reduces in size.
7. Click **Close**.
8. Click (**Reset View**) in the visualization toolbar to restore the previous viewing distance for the scaled domain in the display.

To verify that scaling has been applied, you may perform the mesh diagnostics check and review the output values.

Setting Up the Models

Models define the primary variables of the simulation, including pressure, temperature and velocity, and what mathematical formulation is used to generate the solution. In this example, the flow is turbulent and incompressible. Use the **Segregated Flow** model together with the standard K-Epsilon turbulence model.

The default continuum is automatically named **Physics 1** when the mesh is imported.

1. To give the continuum a more appropriate name, rename the **Continua** > **Physics 1** node to `Air`.
2. For the physics continuum **Continua** > **Air**, select the following models, in order:

Group Box	Model
Space	**Three-Dimensional** (Selected automatically)
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Time	**Steady**
Viscous Regime	**Turbulent**
	Reynolds-Averaged Navier-Stokes (Selected automatically)

3. Deactivate **Auto-select recommended physics models**.

   This is because the standard K-Epsilon model is used in this case instead of the realizable two-layer K-Epsilon model, which is the STAR-CCM+ default.

4. Select the following models, in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

   The color of the **Air** node has turned from gray to blue to indicate that models have been activated.

5. To see the selected models, expand the **Air** > **Models** node.



6. Save the simulation.

Setting Boundary Conditions and Values

The default boundary conditions on the wall and outlet boundaries are suitable for this analysis so the only user specifications required are on the inlet boundary. A fixed velocity of 20 m/s and appropriate turbulence parameters need to be specified there.

1. Select the **Regions > Fluid > Boundaries > inlet > Physics Conditions > Turbulence Specification** node and set the *Method* to **Intensity + Length Scale**.



The boundary values for turbulence can now be set.

2. Select the **Inlet** > **Physics Values** node and set the following boundary conditions:

Node	Property	Value
Turbulence Intensity > Constant	*Value*	0.05
Turbulence Length Scale > Constant	*Value*	0.005 m
Velocity Magnitude > Constant	*Value*	20 m/s

Specifying Porosity Coefficients

As discussed in the introduction to this tutorial, the porous region is defined by its inertial and viscous resistance coefficients, which are 25 kg/m^4 and 1500 kg/m^{3}s, respectively. As the porous region is isotropic,

these values hold in all directions. The coefficients, along with the turbulence parameters in the porous region, must now be specified.

To start with, the **Porous** region must be defined as a porous medium.

1. Select the **Regions** > **Porous** node and set the *Type* to **Porous Region**.



2. Select the **Porous** > **Physics Conditions** > **Turbulence Specification** node.
3. Change the *Method* to **Intensity + Length Scale.**
4. Select the **Physics Values** > **Porous Inertial Resistance** node and make sure that the *Method* property is set to **Principal Tensor**.

   This is suitable because the constant resistance coefficients correspond to the principal (diagonal) components of the porous resistance tensor.
5. Expand the **Porous Inertial Resistance** > **Principal Tensor** node and open the **XX Component**, **YY Component** and **ZZ Component nodes.**
6. Multi-select the **Constant** nodes within the XX, YY and ZZ Components nodes.
7. In the *Multiple Objects Properties* window, change the *Value* to `25` kg/m^4.



The **XX Axis** and **YY Axis** do not need to be changed because the porous resistance components are specified in the directions of the default axes.

8. Open the **Porous Viscous Resistance** > **Principal Tensor** node.
9. Using the same method as for the inertial resistance, specify the three components of the viscous resistance to be 1500 kg/m^3-s.



10. Select the **Turbulence Intensity** > **Constant** node and change the *Value* to `0.1`



The default turbulent length scale of 0.01 m is suitable so it does not need changing.

11. Save the simulation.

Monitoring the Run Progress

As part of the post-processing operations, make sure that the calculated pressure drop across the porous region is close to the theoretical drop given by Eqn.(2538).

To help you do this, monitor the average pressure on both fluid-porous interfaces, the pressure drop across the porous region and the mass flow rate through the porous region throughout the run. Use these monitors to decide when the solution has converged.

To create an average pressure report on the upstream interface:

1. Right-click the **Reports** node and select **New Report > Surface Average**.



2. Rename the **Surface Average 1** node to `Average Upstream Pressure`.
3. Select the **Average Upstream Pressure** node and set the *Scalar Field Function* to **Pressure**.
4. Edit the *Parts* property and select **Porous > Upstream Interface [In-place 2]**.



5. Click **OK**.

   The *Properties* window for the new report should be as shown below.



Create a similar report for the downstream pressure:

6. Right-click the **Reports** node and select **New Report > Surface Average**.
7. Rename the **Surface Average 1** node to `Average Downstream Pressure`.
8. Select the **Average Downstream Pressure** node and set the *Scalar Field Function* to **Pressure** and the *Parts* to **Porous > Downstream Interface [In-place 1]**.

The completed *Properties* window should be as shown below.



Two new field functions corresponding to these reports have been automatically added to the object tree.

9. Open the **Tools > Field Functions** node and scroll down to see the new nodes:



Use these functions to create an expression report for the pressure drop:

10. Right-click the **Reports** node and select **New Report > Expression**.
11. Rename the **Expression 1** node to `Porous Region Pressure Drop`.
12. Select the **Porous Region Pressure Drop** node and click ... (**Custom Editor**) to enter the *Definition* property as:

    ```
    $AverageUpstreamPressureReport -
    $AverageDownstreamPressureReport
    ```

13. Click **OK**.
14. Click ... (**Custom Editor**) for the *Dimensions* property.
15. In the *Porous Region Pressure Drop - Dimensions* dialog, set the **Pressure** dimension to `1`.
16. Click **OK**.

The *Properties* window for the **Porous Region Pressure Drop** report should be as shown below.



Create a monitor and plot from the **Porous Region Pressure Drop** report.

17. Right-click the **Reports > Porous Region Pressure Drop** node and select **Create Monitor and Plot from Report**.



The monitor and plot are added to the object tree:



18. Right-click the **Plots** > **Porous Region Pressure Drop Monitor Plot** and select **Open** from the pop-up menu.

    The empty plot is displayed.

Create a final report for the mass flow rate through the porous region:

19. Right-click the **Reports** node and select **New Report** > **Mass Flow**
20. Rename the **Mass Flow 1** node to `Mass Flow Rate`.
21. Select the `Mass Flow Rate` node and select **Porous** > **Upstream Interface [In-place 2]** in the *Parts* list.



Create a monitor and plot from the **Mass Flow Rate** report.

22. Right-click the **Reports** > **Mass Flow Rate** node and select **Create Monitor and Plot from Report**.

    The monitor and plot are added to the object tree:



23. Right-click the **Plots > Mass Flow Rate Monitor Plot** and select **Open** from the pop-up menu.

    The empty plot is displayed.

Setting Stopping Criteria

Rather than arranging for the analysis to run for a set number of iterations, it is often useful to specify stopping criteria based on residual values and monitored quantities. For this simulation, define a stopping criterion that checks whether the porous region pressure drop monitor has reached a steady value.

1. Right-click the **Monitors > Porous Region Pressure Drop Monitor** node and select **Create Stopping Criterion from Monitor**.



A new node named **Porous Region Pressure Drop Monitor Criterion** appears under the **Stopping Criteria** node.

2. Select the **Porous Region Pressure Drop Monitor Criterion** and change the *Criterion Option* to **Asymptotic**.



3. Select the **Asymptotic Limit** node and change the |*Max* - *Min*| to `0.02`



The analysis runs until the variation in the **Porous Region Pressure Drop Monitor** has a range of less than 0.02 Pa for 10 successive iterations. Since the pressure on the upstream interface is expected to be of the order of 100 Pa, this should only occur when the solution has practically reached convergence.

Visualizing the Solution

View the velocity vector field as the solution develops on a plane bisecting the fluid and porous regions.

Start by creating a vector scene:

1. Right-click the **Scenes** node and then select **New Scene > Vector**. The *Vector Scene 1* display appears.



There are currently no parts in *Vector Scene 1*. You therefore need to create a plane section to display the vectors on.

2. Right-click the **Derived Parts** node and select **New Part > Section > Plane**.

   The *Create Section* dialog appears in the *edit* tab of the simulation window.


isotropicPorousMedia
simulation
scene/plot
edit
Create Section
Input Parts
[Fluid, Porous]
Select...
Plane Parameters
origin
normal
X
2.40802764900002
1.0 m
Y
0.04179999828338
0.0 m
Z
1.52111053464999
0.0 m
Coordinate System
Laboratory
Snap to Part
plane section
Mode
Single
Offset
0.0 m
Display
No Displayer
New Geometry Displayer
New Scalar Displayer
New Vector Displayer
Existing Displayer
Outline 1
Create
Close
Help

3. Click **Select...** to open the *plane section* dialog. Make sure that the **Fluid** and **Porous** regions (without their boundaries) are selected.



4. Click **OK**.
5. In the *Plane Parameters* box, make sure that the normal vector is set to [`1`, `0`, `0`].



6. In the *Display* box, select **New Vector Displayer**.



7. Click **Create** and then **Close**.

8. Click the button at the top of the *Vector Scene 1* window and select **Views** > **-X** > **Up +Z**.

   The section is now clearly visible as part of the display.

There are currently two vector displayers included in the plot. To remove the unnecessary one:

9. Open the **Vector Scene 1** and **Displayers** nodes, right-click the **Vector 1** node and select **Delete**.



10. Click **Yes** in the *Confirm Object Deletion* dialog.
11. Select the **Outline 1** > **Parts** node and click the right half of the *Parts* property.

12. In the dialog that appears, click the column selection button twice to clear all selections.
13. Select the **Derived Parts** > **plane section** node.



14. Click **OK**.

    The vector scene is now ready.
15. Save the simulation.

Running the Simulation

Run the simulation and observe the Residuals plot, the monitor plots, and the vector scene.

To run the simulation:

1. Click (**Run**) in the top toolbar.

   The Residuals display is created automatically and shows the progress being made by the solver. If necessary, click the *Residuals* tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.

   The two monitor plots created in the Monitoring the Run Progress section and the vector scene should also be displayed. The tabs at the top of the *Graphics* window make it possible to select any one of the active displays for viewing. They can also be rearranged and

multiple scenes can even be shown simultaneously, by dragging the tabs around the edges of the *Graphics* window.

The screenshot below shows large fluctuations in the early iterations which are reduced as the simulation converges to a steady state solution.



During the run, it is possible to stop the process by clicking the (**Stop**) button in the toolbar. If you do halt the simulation, it can be continued later by clicking the (**Run**) button. If left alone, the simulation continues until the stopping criterion is satisfied.

2. Save the simulation.

Visualizing Results

Examine the vector scene in detail.

1. Click the *Vector Scene 1* display tab.
2. Using the click and drag technique, expand the width and the height of the color bar. Drag the color bar to a suitable location.
3. Use the mouse buttons to zoom in and position the part and the color bar so that it is similar to that shown below.



To examine the velocity vectors in just the porous region, remove the fluid region from the plane section.

4. Select the **Derived Parts > plane section** node then click the *Parts* property and clear the fluid region.



5. Click away from the **plane section** node and the velocity vectors is displayed within the porous region only.

To alter the orientation and location of the legend:

6. Select the **Scenes** > **Vector Scene 1** > **Displayers** > **Section Vector 1** > **Color Bar** node and change the *Orientation* to **Vertical**.
7. Set the *Width* and *Height* properties to `0.2` and `0.9`, respectively.



8. In the *Graphics* window, click the legend and drag it to a suitable location, such as that shown below.



As a result of the isotropy of the porous region, the vectors are not aligned in any particular direction.

9. Save the simulation.

Validating the Results

As discussed at the start of the tutorial, the calculated pressure drop across the porous region can be compared to the theoretical pressure drop derived from the mass flow rate.

1. Right-click the **Reports** > **Porous Region Pressure Drop** node and select **Run Report** from the menu.



The value produced by this report is displayed in the *Output* window as 138.59 Pa.

To determine the theoretical pressure drop:

2. Right-click the **Mass Flow Rate** report node and select **Run Report**.

The mass flow rate through the porous region is 2.91E-02 kg/s. Since the fluid density is 1.18415 kg/m^3 and the cross-sectional area of the porous region 0.00786 m^2, this corresponds to a superficial velocity through the region of 3.12 m/s. Putting this superficial velocity into Eqn.(2538), with L = 0.03 m, gives a pressure drop of 147.75 Pa across the porous region, which means that the calculated pressure drop is within 10% of the theoretical drop. This is acceptable because the theoretical pressure drop formula assumes a constant superficial velocity in the porous medium. In this comparison, you have used an average velocity value based upon the mass flow rate, implying that the theoretical value is only an approximation to the actual pressure drop.

Summary

This tutorial models flow through the catalyst geometry.

This tutorial illustrates the following STAR-CCM+ features:

- Defining physical models.
- Defining material properties required for the selected models.

- Creating region interfaces.
- Defining boundary conditions.
- Specifying porosity coefficients.
- Monitoring the solution progress.
- Setting up stopping criteria.
- Initializing and running the solver until monitored values approach an asymptotic solution.
- Evaluating the solution results using the visualization and reporting facilities of STAR-CCM+.

Porous Resistance: Orthotropic Media

This tutorial solves the same problem as the Porous Resistance: Isotropic Media, except that the fluid in the porous region cannot flow in any direction other than the bulk flow (y-) direction.

This type of orthotropic porosity is specified by assigning large values to the inertial and viscous resistance components in the cross-flow directions (x- and z- directions in this case). The values of the components in the y-direction are the same as in the Porous Resistance: Isotropic Media, that is:

$$P_i = \begin{bmatrix} 10000 & 0 & 0 \\ 0 & 25 & 0 \\ 0 & 0 & 10000 \end{bmatrix} \mathrm{kg/m^4}$$

and

$$P_v = \begin{bmatrix} 100000 & 0 & 0 \\ 0 & 1500 & 0 \\ 0 & 0 & 100000 \end{bmatrix} \mathrm{kg/m^3s}$$

Contents

Prerequisites
Loading an Existing Simulation
Renaming the Simulation
Setting Initial Conditions
Changing the Porous Resistance Coefficients
Setting Stopping Criteria
Initializing the Simulation
Running the Simulation
Visualizing the Results
Comparing Results
Summary

Prerequisites

The instructions in the Porous Resistance: Orthotropic Media tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading an Existing Simulation

To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online.

The following instructions apply to a locally created simulation file.

1. Select **File > Load Simulation** from the menu bar.

   The *Load Simulation* dialog appears.
2. Click **Browse** and locate file `isotropicPorousMedia.sim`, created in the Porous Resistance: Isotropic Media tutorial.
3. Select this file and click **Open** in the file browser.
4. Click **OK** in the *Load Simulation* dialog.

Renaming the Simulation

Rename and save the simulation.

1. Select **File > Save As** to create a file for the new simulation.
2. In the *Save* dialog, navigate to the directory where you want to place the file.
3. Type `orthotropicPorousMedia` in the **File Name** text box.
4. Click **Save**

The name of the simulation displayed in the title bars of the *Explorer*, *Properties*, and *Output* windows changes accordingly.

Setting Initial Conditions

Set the initial velocity as 3 m/s in the bulk flow (y-) direction.

1. Select the **Continua** > **Air** > **Initial Conditions** > **Velocity** > **Constant** node.



2. Enter `0, -3.0, 0` as the *Value* property.

Changing the Porous Resistance Coefficients

The only model parameters that need to be changed for this problem are the XX and ZZ components of the inertial and viscous porous resistance tensors.

1. Select the **Regions** > **Porous** > **Physics Values** > **Porous Inertial Resistance** > **Principal Tensor** > **XX Component** > **Constant** node.



2. Change the *Value* to `10000` kg/m^4.
3. Do the same for the **ZZ Component** of the inertial resistance tensor. Leave the **YY Component** at the value of `25` kg/m^4.
4. Using the same method, change both the **XX Component** and the **ZZ Component** of the viscous resistance tensor to `100000` kg/m^{3}s.



Keep the **YY Component** value of `1500` kg/m^{3}s.

5. Save the simulation.

Setting Stopping Criteria

Disable the stopping criterion based on the pressure drop monitor and run the analysis for 250 iterations instead.

1. Select the **Stopping Criteria > Porous Region Pressure Drop Monitor Criterion** node and deactivate *Enabled*.
2. Select the **Maximum Steps** node and change the *Maximum Steps* to `250`.
3. Save the simulation.

Initializing the Simulation

Despite the similarity between the two problems, it is preferable in this case to start the analysis from scratch rather than restart from the Porous Resistance: Isotropic Media solution.

1. To clear the existing solution, select **Clear Solution** from the **Solution** menu and click **OK** in the *Clear Solution* dialog.

2. To re-initialize the solution, click (**Initialize Solution**) in the toolbar.
3. Save the simulation.

Running the Simulation

Run the simulation to generate the solution.

1. To run the simulation, click (**Run**) in the top toolbar.

   If you do not see this button, use the **Solution** > **Run** menu item. You could also show the **Solution** toolbar by selecting **Tools** > **Toolbars** > **Solution** and then clicking the toolbar button.

   The tabs at the top of the *Graphics* window make it possible to select any one of the active displays for viewing. You may want to view the Pressure Drop Monitor Plot or the Mass Flow Rate Monitor Plot to observe the solution convergence rate.

   During the run it is possible to stop the process by clicking the (**Stop**) button in the toolbar. If you do halt the simulation, it can be continued later by clicking the (**Run**) button.

   The Residuals display is created automatically and shows the progress being made by the solver. If necessary, click the *Residuals* tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.
2. Save the simulation after the run has finished.

Visualizing the Results

Examine the velocity vectors and magnitude map.

1. At the end of the calculations, open the **Scenes** node and double-click the **Vector Scene 1** node to display the velocity vectors in the porous region.



The flow in this porous medium is essentially unidirectional.

Examine how the porous properties affect the flow downstream of the porous region by comparing velocity magnitude contours at the outlet from the two simulations side-by-side.

2. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
3. Select the **Outline 1** displayer node and deactivate the *Outline* property.
4. Select the **Scalar 1** displayer node and set *Contour Style* to **Smooth Filled**
5. Select the **Scalar 1 Parts** node and add **Fluid: pressure** to the Parts list.
6. In the *Graphics* window, right-click the legend color scale and select **Velocity** > **Magnitude** from the list.



7. Click in the toolbar and select **View** > **-Y** > **Up +Z**.
8. Using the click and drag technique, expand the width and the height of the color bar. Drag the color bar to a suitable location.

   The velocity magnitude contours on the outlet boundary plane are displayed.



9. Save the simulation.

Comparing Results

Compare the results of this tutorial with those of the Porous Resistance: Isotropic Media tutorial.

Open the simulation file from the Porous Resistance: Isotropic Media and create an identical scalar plot for it.

1. Select **File** > **Load Simulation** and open the isotropicPorousMedia.sim file. Its simulation tree is displayed in the *Explorer* pane.



2. Return to the *orthotropicPorousMedia* window, right-click the **Scalar Scene 1** node, and select **Copy**.
3. Open the *isotropicPorousMedia* window, right-click the **Scenes** node and select **Paste** from the menu.

   A scene named *isotropicPorousMedia Scalar Scene 1* is shown in the *Graphics* window.

To display the two scenes side-by-side for comparison:

4. Click the *orthotropicPorousMedia Scalar Scene 1* tab at the top of the *Graphics* window and drag it to the right-hand side. Release the mouse button when you see a red box marking out a strip at the edge of the window, as shown below.



Display the two scalar scenes side-by-side.

5. Click the bar between them and drag it to the left so that they have roughly the same width.



6. Save the simulation.

Summary

This tutorial solved the same problem as the Porous Resistance: Isotropic Media tutorial, but for a fluid that can only flow through the porous region in the direction of the bulk flow.

This tutorial illustrates the following STAR-CCM+ features:

- Loading an existing simulation.
- Renaming and saving the simulation.
- Specifying porous resistance coefficients for an isotropic porous medium.
- Setting the solver parameters and stopping criteria.
- Initializing the simulation.
- Running the solver for a set number of iterations.
- Analyzing results using the built-in visualization facilities.
- Comparing results from different simulations.

Solution Recording and Playback: Vortex Shedding

This tutorial demonstrates how to use the solution recording and playback module for capturing the results of transient phenomena.

The particular scenario being modeled is that of incompressible air flowing over a cylinder with diameter D = 0.01 m. Under the correct conditions, vortices are formed and shed from the cylinder in a regular pattern. The free-stream velocity is 0.15 m/s and the flow is laminar with a Reynolds number (Re) of 75. The predicted Strouhal number and shedding frequency can be determined from a graph of measured lift forces and compared to results obtained by Daily et al.[413].

A 2D volume mesh of a simple cylinder in a fluid domain is provided for this tutorial, the dimensions of which are shown below.



Contents

Prerequisites
Importing the Volume Mesh
Selecting Physics Models
Modifying Material Properties and Setting Initial Conditions
Setting Boundary Conditions
Creating a Scalar Scene
Preparing the Lift Plot
Modifying Solver Settings
Setting Up Stopping Criteria

Setting Up the Solution History File
Running the Simulation
Visualizing the Results
Validating the Results
Creating a Recorded Solution View
Creating an Animation from the Solution View
Summary
Bibliography

Prerequisites

The instructions in the Solution Recording and Playback: Vortex Shedding tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Volume Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `/doc/startutorialsdata/incompressibleFlow/data`.
5. Select `vortexSheddingDomainMesh.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Create a mesh scene and examine the 2D mesh.



7. Save the simulation as `vortexShedding.sim`.

Selecting Physics Models

A physics continuum was added to the object tree when the volume mesh was imported. Select the physics models required to run this case.

The fluid used in this tutorial is air, and the flow is incompressible and laminar. Vortex shedding is a periodic phenomenon and requires the use of a transient solver.

To select the physics models:

1. Rename the **Continua** > **Physics 1** continuum to `Fluid`.
2. For the physics continuum, **Continua** > **Fluid**, select the following models in order:

Group Box	Model
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Constant Density**
Viscous Regime	**Laminar**

3. Click **Close**.

4. To review the models, open the **Fluid** > **Models** node.



Modifying Material Properties and Setting Initial Conditions

Modify the material properties of air, so that the correct Reynolds number is obtained.

1. Within the **Fluid** continuum, select the **Models** > **Gas** > **Air** > **Material Properties** > **Density** > **Constant** node and set Value to `1 kg/m^3`.



2. Select the **Dynamic Viscosity** > **Constant** node and set *Value* to `2.0E-5 Pa-s`.

Set the initial conditions so that the simulation begins with the fluid in a state of motion.

3. Select the **Fluid** > **Initial Conditions** > **Velocity** > **Constant** node and set Value to `[0.15, 0.0, 0.0] m/s`.

Setting Boundary Conditions

Set the required velocity at the inlet boundary.

To set the boundary conditions:

1. Edit the **Regions** > **Fluid_Domain** > **Boundaries** node.
2. Set the following properties:

Node	Property	Setting
Physics Conditions		
Velocity Specification	*Method*	**Components**

Node	Property	Setting
Physics Values		
Velocity > Constant	*Value*	`[0.15, 0.0, 0.0] m/s`

Creating a Scalar Scene

Create a scalar scene displaying vorticity to visualize the solution while the simulation is running.

1. Create a scalar scene.
2. Click the **scene/plot** button located above the object tree.
3. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set Function to **Vorticity** > **Magnitude**.
4. In the same window, set the following properties:

Property	Value
Min	`0`
Max	`28`
Clip	`Off`
Auto Range	`Off`

5. Select the **Displayers** > **Scalar 1** node and set *Contour Style* to **Smooth Filled**.
6. Click the **simulation** button to return to the STAR-CCM+ simulation object tree.

Add an annotation displaying solution time to the scene:

7. Expand the **Tools** > **Annotations** node and drag the **Solution Time** node into the scene.



The Solution Time annotation is added to the bottom left of the scene.

Preparing the Lift Plot

Monitor the lift that the cylinder wall is experiencing to determine the period of oscillation for the vortex shedding.

First create a report:

1. Right-click the **Reports** node and select **New Report > Force Coefficient**.
2. Rename the new plot to `Coefficient of Lift`.
3. Select the **Coefficient of Lift** node and set the following:

Property	Setting
Reference Velocity	`0.15 m/s`
Reference Area	`0.01 m^2`
Force Option	**Pressure**
Direction	`[0.0, 1.0, 0.0]`

4. Set *Parts* to **Regions > Fluid_Domain > Cylinder**.

   The completed **Properties** window is shown below.



Create a monitor and plot from this report.

5. Right-click the **Coefficient of Lift** node and select **Create Monitor and Plot from Report**.

   A new monitor and plot is added to the **Monitors** and **Plots** nodes respectively.



Modify the monitor so that the data is plotted against time; the default setting would plot the data against iterations.

6. Select the **Monitors > Coefficient of Lift Monitor** node and set the *Trigger* to **Time Step**.

7. Right-click the **Coefficient of Lift Monitor Plot** node and select **Open**.

Modifying Solver Settings

Modify the solver settings to more appropriate values.

1. Select the **Solvers** > **Implicit Unsteady** node and set the following properties:

Property	Setting
Time-Step	`0.02 s`
Temporal Discretization	**2nd-order**

2. Select the **Solvers** > **Coupled Implicit** node and set Courant Number to `100`.

Setting Up Stopping Criteria

Reduce the number of inner iterations for each time step and extend the maximum time for the solver to run.

1. Select the **Stopping Criteria** > **Maximum Inner Iterations** node and set *Maximum Inner Iterations* to `15`.
2. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `8 s`.
3. Save the simulation.

Setting Up the Solution History File

Create a solution history, `simh`, file and use it to store selected solution data at specified time intervals.

1. Right-click the **Solution Histories** node and select **New...**.



The *Save* dialog appears.

2. Enter `vortexSheddingData.simh` as the name of the solution history file and click **Save**.

   A new node containing the name of the solution history file is added to the object tree below the **Solution Histories** node.



The red asterisk next to this node means that data is actively written to the file when the simulation runs.

Choose what data to save to the solution history file. As this is a vortex shedding case it is appropriate to store results for pressure, velocity and vorticity.

3. Select the **Solution Histories** > **vortexSheddingData** node and set the following properties:

Property	Setting
Scalar Functions	**Pressure**
Vector Functions	**Velocity**
	Vorticity

Set the frequency with which the selected data is written to the solution history file:

4. Select the **Solution Histories** > **vortexSheddingData** > **Update** node and set *Trigger* to **Time Step**.



5. Select the **Update** > **Time-Step Frequency** node and make sure that *Time-step Frequency* is set to `1`.
6. Select the **Solution Histories** > **vortexSheddingData** node.

   A summary of the properties for the solution history file is shown.



The *Auto-record* property records the data to the solution history file at the required intervals. If you do not want to record the data when the solver is running, deactivate this setting. The *Regions* property is automatically populated with all regions in the simulation. In cases with multiple regions, you may remove regions by clicking to the right of the *Regions* property and clearing the checkbox next to the region you want to remove. The *Path* property displays the relative path to the simulation history file. The *States* property displays the number of saved states stored in the selected solution history file; currently this is displaying 0 as the solver has not yet run.

7. Save the simulation.

Running the Simulation

The simulation is ready to run.

To start the analysis:

1. In the **Solution** toolbar, click (**Run**).

   While the simulation is running you can click the tabs at the top of the **Graphics** window to switch between the plot and the scene. Also, the **Solution Views > Current Solution** node displays the current *Iteration*, *Time Step*, and *Solution Time*. The selected data is saved to the solution history file at every time step.

   The Residuals display is created automatically and shows the progress of the solvers. The simulation continues until the physical time of the simulation reaches 8 seconds.

While the simulation is running,

2. While the simulation is running, select the **Scalar Scene 1** tab at the top of the **Graphics** window to visualize the solution.
3. When the simulation has finished, save the simulation.

Visualizing the Results

Examine the Vorticity scene and the Coefficient of Lift plot.

The scalar scene after 8 s is shown below.



The Coefficient of Lift monitor plot is shown below.



Validating the Results

From the scalar scene and the monitor plot, it is clear that vortex shedding is occurring. The Strouhal number (St) is commonly used when describing oscillating flows and is defined as:

$$St = \frac{fD}{U} \tag{1}$$

Where f is the frequency of vortex shedding, D is the cylinder diameter, and U is the free-stream velocity. In this case, the Strouhal number is given as 0.15 by Daily et al. [413]. The theoretical frequency of vortex shedding is therefore calculated as 2.25 Hz, which gives a period of 0.444 seconds.

The predicted period of shedding can be obtained by zooming into the last two troughs of the monitor plot.

1. Click (**Zoom Selected Plot**) in the **Plot** toolbar.
2. Drag a box around the last two troughs on the plot, as shown below.



The resulting plot is shown below.



3. Click the **scene/plot** button.

4. Select the *Coefficient of Lift Monitor Plot* > *Axes* > *X Axis* > *Grid* node and set the *Spacing* to `0.02`.
5. Click the **simulation** button.

   The enlarged scale on the X-axis makes it possible to measure the period.

The predicted period is shown to be approximately 0.44 seconds.

Note: A relatively large time step was used for this tutorial resulting in a limited number of data points in the plot. If you want more accurate results and a smoother plot, reduce the time step.

There is a difference of less than 1% between the predicted period and the reference period, which is good agreement. The corresponding predicted frequency of 2.27 Hz is also in good agreement with the theoretical frequency of vortex shedding of 2.25 Hz.

Creating a Recorded Solution View

Solution views are used to interrogate this data and make it available for post-processing.

Properties of the solution view set the point in the solution history at which data is read. Data is read into a separate representation linked to the solution view. The solution history file contains all the data that was specified in the previous part of the tutorial.

1. Select the **Solution Histories > vortexSheddingData** node.

   The number of states stored in the solution history file is displayed next to the *States* property.



   There are 400 states stored in the solution history file.

2. Create a solution view to access the states.
3. Right-click the **Solution Histories > vortexSheddingData** node and select **Create Recorded Solution View**.



   A new node, **vortexSheddingData**, is added to the **Solution Views** node. The properties of the solution view node control the data shown by the representation associated with it.



An additional representation linked to this view was added to the object tree under the **Representations** node.



This representation stores solution data from the solution history file.

4. Click the **Scalar Scene 1** tab in the **Graphics** window.
5. Drag and drop the **Solution Views** > **vortexSheddingData** node onto a blank area in the scene window**.**



This action sets the *Representation* property for all displayers (**Outline 1** and *Scalar 1*) and solution history annotations (**Solution Time**) in the scene to **vortexSheddingData**.

The scene is shown below.



This scene corresponds to the data stored in the solution history file for the first time-step. Adjust the solution time to display the solution data at 1.3 seconds.

6. Select the **Solution Views > vortexSheddingData** node and click to the right of the *Solution Time* property.

   A slider bar appears. Here you can drag the slider to change the physical time of the solution being displayed in the scene.

7. Drag the slider so that the time is approximately 1.3 seconds.



The scene updates as shown below.



Values can also be entered rather than using the slider.

8. Click to the right of the *Solution Time* property and enter `3 s`.

   The updated scene is shown below.



Restore the scene to use the most recent solution from the simulation file (stored on the volume mesh representation).

9. Drag the **Solution Views** > **Current Solution** node into the scene as previously described.
10. Save the simulation.

Creating an Animation from the Solution View

Create an animation that shows the development of vortex shedding from the start to a regular periodic state.

1. Drag the recorded solution view, **Solution Views** > **vortexSheddingData**, back into the scene.

2. Select the **Solution Views** > **vortexSheddingData** > **Animation** node and set the *Animation Mode* to **Solution Time**.

   A new node, **Solution Time Animation**, appears below the **Animation** node.



When creating an animation, it is important to set the correct framerate. If the framerate is set too high, the video may reuse identical frames. If the framerate is set too low, the video may appear to play too quickly.

One way to determine the correct framerate is to divide the total number of states (frames) by the total physical time of the simulation. In this case you have 400 states and a physical time of 8 seconds. Knowing this, you can calculate that 50 frames make up one second of simulation time.

Adjust the framerate for the animation.

3. Click the **scene/plot** button.
4. Select the **Scalar Scene 1** > **Attributes** > **Animation** node and set the *Target frame rate (fps)* to `50`.

Record the animation.

5. Click (**Write Movie**) in the **Animation** toolbar.

   The *Write animation* dialog appears.
6. Set the following properties:

Property	Value
Animation Length	`8`
Size	A resolution of your choice

Property	Value
File Name	`vorticityAnimation.avi`

The completed dialog is shown below.



7. Click **Save** to write the animation to disk.
8. Play back the animation using a player of your choice.

Summary

This tutorial demonstrated how to use the solution recording and playback module for capturing the results of transient phenomena.

The tutorial covered the following:

- Creating a Solution History file.
- Selecting appropriate scalars to save to the Solution History file at suitable intervals.
- Plotting the lift coefficient and validating the period of shedding against reference data.
- Creating a Recorded Solution View.

- Viewing recorded solution data in a scene.
- Recording an animation from the Solution History file.

Bibliography

Daily, J.W., and Harleman, D.R.F. Fluid Dynamics, Addison-Wesley, MA, 1966

Compressible Flow

The tutorials in this set illustrate various STAR-CCM+ features for compressible fluid flows as well as turbo wizard and harmonic balance

Contents

Subsonic Flow: NACA-Type Intake
Transonic Flow: RAE2822 Airfoil
Transonic Flow: RAE2822 Airfoil Using Overset Mesh
Adjoint Flow Solver: Flow through an S-Bend
Turbo Wizard: Single Row
Turbo Wizard: Multiple Rows
Harmonic Balance: Single Stage Periodic Flow

Subsonic Flow: NACA-Type Intake

This tutorial illustrates how to solve a three-dimensional compressible flow problem using STAR-CCM+.

The geometry in question is a NACA-type axisymmetric inlet in a freestream air flow of Mach 0.21. The flow that is captured by the inlet is forced into an S-bend before reaching the engine face.



Caution: To keep the cell count low, certain compromises have been made. Even though the case is run turbulent for this exercise, the mesh is more suited for inviscid flow. Because no attempt is made to resolve the boundary layers, the results are thus essentially what one would obtain from an inviscid simulation. This practice is not recommended for most simulations.

To illustrate how STAR-CCM+ "fuses" separate mesh sections into a contiguous mesh, the mesh has been created in three separate non-conformal sections.

In spite of the fact that the mesh is too coarse to produce a mesh-independent solution, the simulation results are shown to compare favorably with experimental data for pressure on the S-duct walls.

The mesh that is used for this intake example is based on an S-duct inlet mesh from NASA.

Contents

Importing the Mesh and Naming the Simulation
Visualizing the Mesh

Manipulating Regions and Boundaries
Setting Up the Physics Models
Setting Fluid Properties
Setting Boundary Conditions
Setting Initial Conditions
Adding Field Functions
Setting Stopping Criteria
Setting Solver Parameters
Running the Simulation
Adjusting Solver Parameters and Continuing
Visualizing the Mach Number
Visualizing the Velocity Vectors
Setting Up a Line Surface for Plotting
Plotting Simulation Data
Plotting Reference Data
Closing and Restarting
Summary

Importing the Mesh and Naming the Simulation

Create a new simulation, import the supplied mesh, and name the simulation `intake`.

Use the following steps:

1. Start up STAR-CCM+ and select the **New Simulation** option from the menu bar.

   Continue by importing the mesh and naming the simulation. The mesh has been created in three separate blocks, and stored in separate `.ccm` files.
2. Select **File > Import > Import Volume Mesh...** .
3. In the *Open* dialog, navigate to the `doc/startutorialsdata/compressibleFlow/data` subdirectory of your STAR-CCM+ installation directory.



4. Multi-select files `sduct-inlet.ccm`, `sduct-core.ccm`, and `sduct-outlet.ccm`.
5. To start the import, click **Open**.
6. To create a file for the new simulation, select **File** > **Save**. Navigate to the directory where you want the file located.
7. In the *Save* dialog, type `intake` into the **File Name** text box.
8. Click **Save**.

Visualizing the Mesh

The *Geometry Scene 1* display is present in the **Graphics** window. Set the display to show the mesh.

Initially all parts of the mesh are shown as solid, colored surfaces.

1. Click one of the parts.

   The object (a boundary in the illustration below) becomes highlighted and a label appears with the name of the object selected.



In the **intake** window, the node that corresponds to this object is highlighted.



Similarly, if a different boundary node is selected, the surface corresponding to that boundary is highlighted in the **Graphics** window.

2. Select the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** > **Parts** node and edit the *Parts* property.



3. Expand all the nodes under the **Regions** node.



This dialog contains a list of the parts to be displayed. By default, all boundary surfaces are displayed.

4. Click **OK**.
5. To display the surface mesh on the selected parts, select the **Geometry 1** node and activate *Mesh*:



The surface mesh appears as shown below:



Manipulating Regions and Boundaries

Combine the separate mesh blocks into a single contiguous mesh, then rename the mesh entities before continuing with the simulation.

Use the following steps:

1. In the **intake** window, open the **Regions** node.

   See three separate regions, corresponding to the three mesh blocks that were imported.

2. Click to select each of the region nodes (**Fluid_1**, **Fluid_1 2**, and **Fluid_1 3**) in turn.

   The selected region is highlighted and labeled in the *Geometry Scene 1* display.



There are two alternative approaches for combining these regions so that information can propagate between them:

- Create internal interfaces.
- Fuse the regions and boundaries.

Using internal interfaces is advantageous if the mesh is to be manipulated further at a later stage, since interfaces can easily be deleted and created. However, there is likely to be some (most likely small) computational overhead that is associated with the interface. In this case, since the mesh does not need to be changed, combine the regions and fuse the adjacent boundaries, creating a single contiguous mesh. The process of "fusing" creates a (possibly non-conformal) interface, and then merges that interface into the mesh.

To combine the separate regions into one, use **Shift Click** to select all three region nodes.

3. Right-click and select **Combine...** .



4. In the *Combine Regions* dialog, make sure that all regions under it are selected.



5. Activate **Combine boundaries with similar names**.

In many situations, the **Fuse adjacent boundaries** option can also be used. However, in this case, the specified tolerance is not suitable for all the boundaries to be fused, so complete the fusing operation later.

6. Click **OK** to combine the regions without fusing the boundaries.

   As a result of this operation, a single region node remains.



Since there is only one region now, you can simplify the region name.

7. Rename the **Fluid_1** node to `Fluid`.

Now that the regions have been combined, fuse the adjacent boundaries that are currently separating the regions. The mesh is not conformal at these boundaries, so a non-conformal intersection needs to be computed as part of the fusing process. Essentially, an interface is created and then merged into the mesh.

The intersection process requires the specification of a tolerance. Use a value of 0.02 (default 0.05) and then perform checks to make sure that this is suitable.

Before creating interfaces, the freestream boundary types must be redefined as walls. This allows the solution to be initialized without activating a flow model, normally a requirement of freestream boundaries.

8. Open the **Fluid** > **Boundaries** node.
9. Select all the freestream boundary nodes: **Free_Stream-4**, **Free_Stream-6**, **Free_Stream-31** and **Free_Stream-32** and set the *Type* to **Wall**.



This bypasses the need to activate any models before initializing, the latter being a necessary step for testing the intersecting interface tolerance.

10. Select the **Free_Stream-4** node, then use the `Ctrl Click` approach to select the **Free_Stream-32** node.

    (It is important to select the boundaries in this order.)

11. Right-click and select **Create Interface**.
12. Open the **Interfaces** > **Interface 1** > **Physics Values** node, and select the **Intersection** node and set the *Tolerance* to `0.02`.



You are now ready to initialize the solution manually.

13. Click (**Initialize Solution**) on the toolbar or select **Solution** > **Initialize Solution**.

To determine whether the test was successful, inspect boundaries **Free_Stream-4 [Interface 1]** and **Free_Stream-32 [Interface 1]** in the *Geometry* scene.

14. To gain access to them, first locate the **Free_Stream-31** part by clicking its node.
15. Right-click that part in the scene and select **Hide**.



16. Next, select the nodes of the two internal interface boundaries, then click the (**Reset View**) button on the toolbar to bring them into focus. (Zoom and rotate if necessary.)



The parts appear intact, also showing that the interface was intersected correctly. This shows that the tolerance value of `0.02` is the right one. Delete the interface so that the boundaries can be fused.

17. Right-click the **Interfaces** > **Interface 1** node and select **Delete**.
18. Click **Yes** in the confirmation dialog.
19. Select the nodes of the **Free_Stream-4** and **Free_Stream-32** boundaries (in that order), right-click and select **Fuse**.



The sequence of boundary selection can affect the results of the fuse operation. When you fuse boundaries, STAR-CCM+ attempts to take the second boundary selected by you and project it onto the first.

Hence, when deciding on the sequence of boundaries to select for fusing, always select the one with the finer mesh grid first. This instructs STAR-CCM+ to project the coarser boundary onto the finer one, leading to better results in the mesh grid. If you were to select the boundaries in the opposite order (trying to project the finer boundary onto the coarser one), this may essentially fail, as you may see only parts of the fused boundaries shown in the scene display.

20. In the *Fuse Boundaries* dialog, make sure that these boundaries are selected:



21. Click **OK**.

    Both these boundaries disappear from the tree, since fusing boundaries incorporates them into the interior mesh grid.
22. Repeat the fusing process for the **Symmetry-1** and **Symmetry-7** boundaries. Again, use the default tolerance, `0.02`.
23. Combine (do not fuse) the **Free_Stream-31** and **Free_Stream-6** boundaries by selecting their nodes, right-clicking, and selecting **Combine**.

    This creates a single freestream boundary entity.
24. Select all the remaining symmetry plane nodes, right-click, and select **Combine**.
25. Select the **Wall-5** node and look in the *Geometry* scene.

26. To see the part more clearly, click (**Make Scene Transparent**).



This boundary includes two non-contiguous surfaces.

27. To correct this, right-click the **Wall-5** node and select **Split Non-Contiguous**.



28. Select the **Wall-5** node and then the new **Wall-5 2** node separately. The split action has made them into two separate boundaries.



29. Rename the following nodes:

Node	New Name
Free_Stream-31	`Free_Stream`
Pressure-2	`Pressure`
Symmetry-3	`Symmetry`
Wall-1	`Cowl`
Wall-2	`Stbd Wall`
Wall-3	`Port Wall`
Wall-5	`Extension_Outer`
Wall-5 2	`Spinner`
Wall-6	`Extension_Inner`

This portion of the simulation tree appears as shown below:



30. Save the simulation.

Setting Up the Physics Models

Models define the primary variables of the simulation, including pressure, temperature and velocity, and what mathematical formulation is used to generate the solution.

In this example, the flow is turbulent and compressible. The Coupled Flow Model is used together with the default K-Epsilon Turbulence model.

1. For the physics continuum **Continua** > **Physics 1**, deactivate **Auto-select recommended models** and select the following models, in order:

Group Box	Model
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
Energy	**Coupled Energy**
Time	**Steady**
Viscous Regime	**Turbulent**
Turbulence	**Reynolds-Averaged Navier-Stokes**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

Group Box	Model
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

2. Click **Close**.
3. In the **intake** window, open the **Continua** > **Physics 1** > **Models** node.

   The selected models now appear within that node.



4. Save the simulation.

Setting Fluid Properties

Since the Ideal Gas model has been chosen, the viscosity, specific heat, and thermal conductivity fluid properties must be set.

1. Select the **Continua** > **Physics 1** > **Models** > **Gas** node, then select the **Air** > **Material Properties** > **Dynamic Viscosity** > **Constant** node, and set *Value* to `1.853E-5 Pa-s`.



2. Within the same **Air** node, select the **Specific Heat** > **Constant** node and set *Value* to `1005 J/kg-K`.

The thermal conductivity can be left at its default value of `0.0260305 W/m-K`.

Setting Boundary Conditions

In this simulation, the freestream Mach number is 0.21, the freestream pressure is 98,258.18 Pa, and the freestream temperature is 300 K. Use the default reference pressure of 101,325 Pa, which means that the freestream gauge pressure is -3066.82 Pa.

1. Select the **Regions** > **Fluid** > **Boundaries** > **Free_Stream** node.
2. Set the *Type* property back to **Free Stream**.
3. In the **Boundaries** > **Free_Stream** > **Physics Values** node, set the following properties:

Node	Property	Setting
Mach Number > **Constant**	*Value*	`0.21`
Pressure > **Constant**	*Value*	`-3066.82 Pa`

The gauge pressure at the duct outlet is specified as -7903.35 Pa.

4. In the **Boundaries** > **Pressure** node, set the following properties:

Node	Property	Setting
Physics Values		
Pressure		
Constant	*Value*	`-7903.35 Pa`

The mesh in the section downstream of the engine face is coarse. For this reason, no attempt is made to resolve the boundary layer. Instead of the

usual no-slip condition, a slip wall is specified on the inner and outer walls.

5. In the **Boundaries** node, set the following properties:

Node	Property	Setting
Extension_Inner		
Physics Conditions		
Shear Stress Specification	*Method*	**Slip**
Extension_Outer		
Physics Conditions		
Shear Stress Specification	*Method*	**Slip**



Setting Initial Conditions

In compressible fluid problems, it is advisable to initialize the flow to a non-zero initial velocity and a meaningful pressure. Here, you use the freestream velocity and the duct outlet pressure.

1. Open the **Physics 1** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Pressure > **Constant**	*Value*	`-7903.35`
Velocity > **Constant**	*Value*	`72,0,0`

2. Save the simulation.

Adding Field Functions

Field functions provide a mechanism by which fields (raw data from the simulation stored in cells and/or on boundaries) may be viewed and defined.

User field functions are field functions that you create to access field data. This section illustrates how to create a field function for post-processing purposes that represents the ratio of absolute pressure to atmospheric pressure.

1. Open the **Tools** node, right-click the **Field Functions** node and select **New > Scalar**.
2. Rename the new field function node to `PressureRatio`.
3. Select the **PressureRatio** node and set *Function Name* to `PressureRatio`.



4. Set *Definition* to `$AbsolutePressure/101325.0`



This entry instructs your new field function to reference the Absolute Pressure field function, and states that the atmospheric pressure value is 101,325.0 Pa.

Setting Stopping Criteria

Stopping criteria allow you to specify the length of the simulation, in this case the number of iterations to be performed.

To set a stopping criterion:

1. Open the **Stopping Criteria** node and select the **Maximum Steps** node.
2. Set the *Maximum Steps* to `100`.



Setting Solver Parameters

In this simulation, the Coupled Flow model is controlled by the Coupled Implicit solver.

This solver enables the local iteration step to be adjusted in terms of the Courant number. The default value of 5 is conservative for a low Mach number problem on a mesh of reasonable quality, so it can be increased to speed convergence.

1. Select the **Solvers** > **Coupled Implicit** node.



2. Set the *Courant Number* to `20`.

This is sufficiently low for the flow field to become established. After 100 iterations, increase the Courant number further to allow a converged solution to be reached more rapidly.



3. Save the simulation.

All necessary material properties, boundary conditions, field functions, and solver parameters for this tutorial have been applied, so the simulation can now be run.

Running the Simulation

Run the simuation and examine the Residuals display.

1. To run the simulation, click the (**Run**) button in the toolbar.
   If you do not see this button, use the **Solution** > **Run** menu item. You could also show the **Solution** toolbar by selecting **Window** > **Toolbars** > **Solution**, and then clicking the toolbar button.

The Residuals display is created automatically and shows the progress of the solver. If necessary, click the **Residuals** tab to bring the Residuals plot into view.

During the run, it is possible to stop the process by clicking the (**Stop**) button on the toolbar. If you do halt the simulation, it can be continued later by clicking the (**Run**) button. If left alone, the simulation continues until the stopping criterion of 100 iterations is satisfied.

2. Save the simulation after the run has finished.

Adjusting Solver Parameters and Continuing

Now that the flow field has become established and the solution is converging well, the Courant number can be further increased.

1. Select the **Solvers** > **Coupled Implicit** node.
2. Set the *Courant Number* to `50`.



3. Select the **Stopping Criteria** > **Maximum Steps** node.
4. Set the *Maximum Steps* to `200`

   This instructs STAR-CCM+ to iterate for a further 100 steps in addition to those already performed.



5. Click (**Run**) to continue running the simulation.
6. Once the run is finished, save the simulation.

Visualizing the Mach Number

As an introduction to data visualization, this section explains how to examine the Mach number field.

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
2. For a clearer view of the model, click (**Save-Restore-Select Views**) on the toolbar. Select **Views** > **+Y** > **Up -Z**.

   This aligns the model along the chosen axes. You can now zoom and pan to adjust the view further.



3. In the simulation tree, select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** node and edit the *Parts* property.

   The *Parts* dialog appears.
4. Select the **Symmetry** object.
5. Click **OK**.
6. Right-click the scalar bar in the display. In the pop-up menu that appears, select **Mach Number** > **Lab Reference Frame**.



The result appears in the *Scalar Scene 1* display.



By default, **Filled** cell values are shown.

7. Select the **Scalar 1** node and set the *Contour Style* to (**Smooth Filled**).

The contours of the scalar display now appear smooth.

8. Save the simulation.

Visualizing the Velocity Vectors

This section describes how to examine velocity vectors.

Use the following steps to produce a vector scene:

1. Right-click the **Scenes** node. Selecting **New Scene** > **Vector** causes a *Vector Scene 1* display to appear.
2. Open the **Regions** > **Fluid** > **Boundaries** node, and then drag the **Symmetry** node into the display.



When your mouse pointer arrives on the display, a pop-up menu appears letting you decide the part displayer to receive the new part.

3. Select **Add to Vector 1**.



At this point, the part is in the scene, though it is difficult to see from the default view.

4. For a clearer view of the model, click (**Save-Restore-Select Views**) on the toolbar to activate the **Save-Restore-Select Views** menu.
5. In that menu, select **View > +Y > Up -Z**. You can also zoom and pan the view as needed.

   The result appears in the *Vector Scene 1* display.



Velocity is the default selection for vector displays. However, if you do not see any vectors after adding a part to the scene, right-click the vector bar in the display. In the pop-up menu that appears, select **Velocity > Lab Reference Frame**.



The vector arrows in this scene are too long. To make them shorter:

6. Select the **Scenes > Vector Scene 1 > Displayers > Vector 1 > Glyph > Relative Length** node and set the *Glyph Length (%)* to `0.2`.



The arrows now appear shorter.



7. Save the simulation.

Setting Up a Line Surface for Plotting

For comparison with the experimental data, a part needs to be created on which to plot the duct wall pressure ratio.

Create the part by intersecting a plane with the surface corresponding to the duct wall, resulting in a line surface.



1. Right-click the **Scenes** node and select **New Scene > Geometry**.
2. In the **intake** window, right-click the **Derived Parts** node and select **New Part > Section > Plane**.

   An interactive in-place dialog appears with default settings to allow definition of the derived part.



3. Click **Input Parts** to activate the *Select Objects* dialog.
4. Expand the **Regions** > **Fluid** nodes and select **Port Wall** only.

Note: The **Fluid** region must be deselected.

5. Click **OK**.
6. In the **Plane Parameters** box, enter values for the y-coordinate: `0.001` for the origin and `1` for the normal. Set all other values to zero.
7. In the **Display** box, select **No Displayer**.

8. Click **Create**.

   In the **intake** window, a node representing the implicit surface created is added under the **Derived Parts** node: **plane section**.



9. Click **Close** in the dialog.

   It is now possible to plot data on the new implicit surface.

10. Save the simulation.

Plotting Simulation Data

In this section, the simulation data is plotted using the user-defined field function for pressure ratio and the line surface created from the S-duct wall.

Use the following steps:

1. In the **intake** window, right-click the **Plots** node and select **New Plot > XY**.

   This creates an empty plot in the tree and displays it in the **Graphics** window.

2. Select the **XY Plot 1** node and click the right half of the *Parts* property.
3. In the *Select Objects* dialog, select the **Derived Parts > plane section**, then click **OK**.
4. Select the **XY Plot 1 > Y Types > Y Type 1** node and make sure that the *Type* property is set to **Scalar**.
5. Select the **Y Type 1 > Scalar Function** node and set the *Field Function* to **PressureRatio**.

   A plot like the one shown below appears:



6. Save the simulation.

Plotting Reference Data

Display an external data set consisting of experimental data together with the simulation data.

First, however, the external data must be loaded into a table.

1. Right-click the **Tools > Tables** node and select **New Table > File Table**.
2. In the *Open* dialog, select `.xy` from the Files of Type drop-down menu.
3. Navigate to the `pressure-ratio-port-wall.xy` file (this is located in the `doc/startutorialsdata/compressibleFlow/data` subdirectory of your STAR-CCM+ installation directory) and open it.
4. Right-click the **Plots > XY Plot 1 > Data Series** node and select **Add Data**.



5. Select **pressure-ratio-port-wall** from the *Add Data Providers to Plot* dialog.
6. Select the **Data Series** > **pressure-ratio-port-wall** node.
7. Set *X Column* to **column0**.
8. Set *Y Column* to **column1**.



The tabular data is displayed on the plot. To view this more clearly:

9. Select the **pressure-ratio-port-wall** > **Symbol Style** node:
   a) Set *Shape* to **Filled Triangle**.
   b) Set *Size* to `12`.
   c) Edit the *Color* property.

      The *Symbol Style - Color* dialog appears.



Make sure that the **Swatches** tab is active, and then select a bright shade of blue to contrast with the existing plot.

d) Click OK.

The external data appears more clearly.



The external data in the preceding illustration is from the following source: Fluid Dynamics Panel Working Group 13. 1991. "Air Intakes for High Speed Vehicles", *Agard Advisory Report 270*, pp. 139-162.

10. Save the simulation.

Closing and Restarting

Close and restart the simulation.

1. Close the file by closing the **intake** window.

It is easy to reopen the file and restore the displays.

2. Select **File > Recent Files > /intake.sim** from the top menus.



Alternatively, you can select the **File > Load Simulation** menu item. This activates the *Load Simulation* dialog.



3. Click the **Browse** button to activate the *Open* dialog. Use it to locate the simulation file and then click **Open**.

   Back in the *Load Simulation* dialog, the name and path of the simulation file appears in the **File** text box.
4. Click **OK**.

   A new window appears for the simulation.

The next step is to restore the displays and plots

5. Expand the **Scenes** node and double-click each scene display's node one at a time.
6. Expand the **Plots** node and double-click the **Residuals** and **XY Plot 1** nodes.

Summary

The Subsonic Flow tutorial illustrated how to solve three-dimensional compressible flow problems.

The tutorial demonstrated the following features of STAR-CCM+:

- Combining and fusing regions and boundaries.
- Selecting the models to be used.
- Defining the material properties that are required for the selected models.
- Defining the region and boundary conditions.
- Setting the solver parameters.
- Running the solver until the residuals are satisfactory.
- Analyzing results using the visualization or XY-plot facilities.
- Reopening the simulation.

Transonic Flow: RAE2822 Airfoil

The tutorial simulates two-dimensional, turbulent, compressible, transonic air flow over an idealized airfoil.

The free-stream Mach number is 0.725 and the angle of attack is 2.92° corresponding to RAE2822 case 6 in [Reference]



The free-stream flow is subsonic, becoming supersonic on the suction side of the airfoil, and subsonic through a shock wave. The lift and drag coefficients are monitored to help determine whether convergence is reached. The final distribution of the pressure coefficient on the airfoil is then compared to experimental data.

Contents

Importing the Mesh and Naming the Simulation
Converting to a Two-Dimensional Mesh
Selecting the Physics Models
Setting Material Properties
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Visualizing and Initializing the Solution
Plotting Graphs
Running the Simulation
Visualizing the Results
Validating the Results
Summary
Airfoil Tutorial Bibliography

Importing the Mesh and Naming the Simulation

Create a new simulation, import the supplied mesh, and name the simulation `aerofoil.sim`.

Use the following steps:

1. Start up STAR-CCM+ and select the **New Simulation** option from the menu bar.

Continue by importing the mesh and naming the simulation. A one-cell-thick, three-dimensional, hexahedral mesh has been prepared for this analysis. The mesh corresponds to an angle of attack of 0° in the default Laboratory coordinate system.

2. Select **File** > **Import** > **Import Volume Mesh** from the menus.
3. In the *Open* dialog, simply navigate to the `doc/startutorialsdata/compressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select file `aerofoil.ccm`, which contains the mesh and boundary definitions.
4. Click the **Open** button to start the import.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.
5. Finally, save the new simulation to disk under file name `aerofoil.sim`.

Converting to a Two-Dimensional Mesh

Convert the mesh region to a two-dimensions.

The existing mesh aligns with the X-Y plane and has a boundary plane at Z = 0 and thus satisfies the requirements to convert three-dimensional meshes to two dimensions:

1. Select **Mesh** > **Convert to 2D**
2. In the *Convert Regions to 2D* dialog that appears, activate **Delete 3D regions after conversion** and click **OK**.



Once you have clicked **OK**, the mesh conversion takes place and the *Geometry Scene 1* display shows the two-dimensional geometry in the *Graphics* window. (If the image does not appear immediately, click (**Reset View**) on the toolbar.)



All the geometry parts are shown, viewed from the positive Z-direction. The mouse rotation option is suppressed for two-dimensional scenes.

3. Right-click the **Physics 1** continuum node and select **Delete**.
4. Click **Yes** in the confirmation dialog.

Selecting the Physics Models

Models define the spatial and temporal solution methods and the physical properties of the flow. In this example, the flow is steady, turbulent, and compressible.

Use the default Spalart-Allmaras turbulence model and the ideal gas model. Use the coupled solver, which is recommended for all supersonic and transonic compressible flows.

By default, a **Physics 1 2D** continuum is created when the mesh is converted to two-dimensional.

1. Rename the **Physics 1 2D** node to `Aerofoil`.

Edit the continuum definition to select appropriate physical models for the fluid.

2. For the physics continuum **Continua** > **Aerofoil**, select the following models, in order:

   Make sure that the **Two Dimensional** model is activated and the **Auto-select** option is selected.

Group Box	Model
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
Time	**Steady**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**Spalart-Allmaras Turbulence**
Enabled Models	**All y+ Wall Treatment**

3. Click **Close**.

Inside the **Continua** node, the color of the **Aerofoil** node has turned from gray to blue to indicate that models have been activated.

4. To review the models, open the **Aerofoil** > **Models** node.



5. Save the simulation.

Setting Material Properties

Set the material properties for gas and air.

Use the following steps:

1. Open the **Gas** and **Air** nodes.

The material properties for air are contained within.

2. Select the **Material Properties > Dynamic Viscosity > Constant** node and change the dynamic viscosity value to `4.58E-5` Pa-s. This setting corresponds to a Reynolds number of 6.5 x 10^6.

Setting Initial Conditions

Specify the coordinate system and the velocity and temperature.

The initial velocity field applies free-stream conditions across the entire domain, that is a velocity of 253 m/s calculated using the equation:

$$u = M\sqrt{\frac{\gamma P_{ref}}{\rho_{ref}}}$$

where: $M = 0.725$, $P_{ref} = 101325$, and $\gamma = 1.4$.

To specify an angle of attack of 2.92° for the initial velocity, create a coordinate system.

1. Open the **Tools** node at the bottom of the simulation tree and right-click the **Coordinate Systems > Laboratory > Local Coordinate Systems** node.
2. Select **New > Cartesian**.
3. In the in-place dialog, in the *Axis Definition* box, change the i Direction to the following:

```
[1, 0.051, 0]
```



4. Click **Renormalize**.

   The j Direction changes automatically to make sure that the axes are perpendicular. The i Direction readjusts itself slightly.

5. Click **Create**, then **Close**.

A **Cartesian 1** node is created within the **Coordinate Systems** node. The resulting *Properties* window is shown below.



This node now defines a coordinate system that, when viewed down the positive Z-axis, has its x- and y-axes rotated anti-clockwise through an angle of 2.92º compared to the laboratory system.

6. Select the **Continua** > **Aerofoil** > **Initial Conditions** > **Velocity** node and change the *Coordinate System* to **Laboratory -** > **Cartesian 1**.
7. Select the **Velocity** > **Constant** node and change the *Value* to `[253, 0.0, 0.0]` m/s.

To specify the initial temperature:

8. Select the **Static Temperature** > **Constant** node and change the *Value* to `300` K.

The default values for the remaining initial conditions are suitable for this problem.

9. Save the simulation.

Setting Boundary Conditions and Values

Set the physical conditions at the domain boundaries.

The geometry for this tutorial has only two boundaries:

- A wall boundary representing the surface of the airfoil.
- A free-stream boundary at the external edge of the solution domain.

1. Rename the **Regions** > **Default_Fluid 2D** node to `Fluid`.
2. Select the **Fluid** > **Boundaries** > **freestream** > **Physics Conditions** > **Flow Direction Specification** node and set the *Method* property to **Components**.
3. Select the **Physics Values** node and set the following properties:

Node	Property	Value
Flow Direction	*Coordinate System*	**Laboratory -** > **Cartesian 1**

Node	Property	Value
Mach Number > Constant	*Value*	`0.725`
Static Temperature > Constant	*Value*	`300` K

4. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the Courant number and the maximum number of steps.

The simplicity of this problem allows a rapidly converging solution to be attained using a large Courant number.

In problems involving more complex geometries or physics, attempting to shorten the run time in this way can cause the solution to diverge. To increase the Courant number:

1. Select the **Solvers** > **Coupled Implicit** node and change the *Courant Number* to `100.0`.



2. Select the **Stopping Criteria** > **Maximum Steps** node and set the *Maximum Steps* to `300`.
3. Save the simulation.

Visualizing and Initializing the Solution

View the Mach number profile during the run to monitor the supersonic flow region above the airfoil.

Start by creating a scalar scene:

1. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.

The *Scalar Scene 1* display appears.

2. Right-click the scalar bar at the bottom of the display and select **Mach Number > Lab Reference Frame** from the pop-up menu.



3. Initialize the run by clicking (**Initialize Solution**) in the toolbar, then use the middle mouse button to zoom in on the airfoil in the center of the scalar scene.

To change the style of the Mach number contours:

4. Select the **Scalar Scene 1 > Displayers > Scalar 1** node and set the *Contour Style* to **Smooth Filled**.
5. Save the simulation.

Plotting Graphs

Plot the lift and drag coefficients to determine when the analysis has converged.

Use the following steps:

1. Right-click the **Reports** node and select **New Report > Force Coefficient**.

   A new report node named **Force Coefficient 1** is created.

2. Rename this node `Drag Coefficient` then enter the following information:

Property	**Value**
Units	Empty
Coordinate System	**Laboratory**
Direction	`[0.998187, 0.040306, 0.0]`
Force Option	**Pressure + Shear**
Reference Pressure	`0.0 Pa`
Reference Density	`1.177 kg/m^3`
Reference Velocity	`253.0 m/s`
Reference Area	`1.0 m^2`
Parts	**[Fluid: wall]**
Number of Bands	`0`
Representation	**Volume Mesh**
Smooth Values	Deactivated



3. Right-click the **Drag Coefficient** node and select **Create Monitor and Plot from Report**.

A new plot node appears named **Drag Coefficient Monitor Plot**.

4. Select the **Drag Coefficient Monitor Plot** node and set the *Title* to `Drag Coefficient Monitor Plot`.
5. Double-click the **Drag Coefficient Monitor Plot** node to display the empty plot in the *Graphics* window.

To create a plot for the lift coefficient:

6. Repeat the steps described above to create and display a plot for the lift coefficient. All settings are the same as for the drag coefficient.
7. Rename the report node **Lift Coefficient** and set its *Direction* property to `[-0.0403, 0.9981, 0.0]`.
8. Select the **Lift Coefficient Monitor Plot** node and set the *Title* to `Lift Coefficient Monitor Plot`.

Experimental data for the pressure coefficient on the airfoil are provided in file `aero_exp.xy` in the `doc/startutorialsdata/compressibleFlow/data` directory. Plot the data on a graph alongside the results of the analysis.

To plot the experimental data:

9. Right-click the **Tools > Tables** node and select **New Table > File Table**.
10. Locate and open file `aero_exp.xy`.
11. Right-click the **Plots** node and select **New Plot > X-Y**.
12. Right-click the **XY Plot 1 > Data Series** node and select **Add Data**.



13. In the *Add Data Providers to Plot* window, select **aero_exp**.
14. Select the **Data Series** > **aero_exp** node and set **aero_exp** for the *Table* property.

Make sure that the *X Column* and *Y Column* properties are as shown below:



A graph of the experimental data appears in XY Plot 1.

To add the numerical data to the same graph:

15. Select the **XY Plot 1** node and click the *Parts* property.
16. Select **Fluid: wall** in the *XY Plot 1 - Parts* dialog.
17. Select the **Y Types** > **Y Type 1** > **Scalar Function** node and set **Pressure Coefficient** for the *Field Function* property.

The initial pressure coefficient is shown in the XY Plot 1 as being zero everywhere.

The pressure coefficient requires specification of a reference pressure and a reference velocity.

18. Select the **Tools > Field Functions > Pressure Coefficient** node and enter a *Reference Density* of `1.177 kg/m^3` and a *Reference Velocity* of `253 m/s`, have the settings shown below:



The usual convention in aerodynamics problems is to reverse the y-axis orientation in pressure coefficient plots.

19. Select the **XY Plot 1 > Axes** node and click the *Axis Orientation* property.

The *Axes - Axis Orientation* window appears.

20. Select the option below.



The setup is now complete.

21. Save the simulation.

Running the Simulation

The simulation is ready to run.

1. To run the simulation, click (**Run**) in the top toolbar, or use the **Solution** > **Run** menu item.

The Residuals display is created automatically and shows the progress of the solver. Observe the run progress by selecting one of the tabs at the top of the *Graphics* window.

During the run, it is possible to stop the analysis by clicking (**Stop**) in the toolbar. Click (**Run**) to continue the solution. If left alone, the simulation continues until 300 iterations have been completed.

Once this stage is reached, make sure that the solution has converged by examining the lift and drag coefficient plots.

2. Double-click the **Lift Coefficient Monitor Plot** node to display the results in the *Graphics* window.



3. Similarly, display the drag coefficient plot and adjust its y-axis scale.



Both monitors have reached constant values so it is reasonable to conclude that the solution has converged.

4. Save the simulation.

Visualizing the Results

The *Scalar Scene 1* display shows the Mach number profile at the end of the run. The profile shows the transonic flow around the airfoil, including the shock wave above it.



Validating the Results

Examine the XY plot and life coefficient report.

To view the numerical and experimental data:

1. Select the *XY Plot 1* tab.

This plot appears as shown below with the X-axis scale adjusted.



Other than the shock position, which is in error as a result of the mesh coarseness and choice of turbulence model, the numerical pressure coefficients compare well with the experimental data. To see the calculated lift coefficient:

2. Right-click the **Reports > Lift Coefficient** node and then select **Run Report**.



In the *Output* window, the *Lift Coefficient Report* returns a value of 0.78. Similarly, the drag coefficient report gives a value of 0.0081.

Summary

The Transonic Flow tutorial simulated two-dimensional, turbulent, compressible, transonic flow.

This tutorial introduced the following features of STAR-CCM+:

- Defining models for compressible flow problems.
- Defining the material properties for the selected models.
- Setting solver parameters for a steady-state run.
- Plotting graphs comparing results with experimental data.
- Initializing and running the solver to a specified stopping criterion.
- Analyzing the results using the built-in visualization facilities.

Airfoil Tutorial Bibliography

Cook, P.H., M.A. McDonald, M.C.P. Firmin "Aerofoil RAE 2822 - Pressure Distributions, and Boundary Layer and Wake Measurements Experimental Data Base for Computer Program Assessment", AGARD Report AR 138, 1979

Transonic Flow: RAE2822 Airfoil Using Overset Mesh

This tutorial demonstrates how to use an overset mesh to simulate transonic flow over the RAE2822 airfoil.

The physics models, initial, and boundary conditions are the same as in the Transonic Flow tutorial. By using an overset mesh for the region surrounding the airfoil, this tutorial shows how the angle of attack is changed without having to generate a mesh. This approach greatly simplifies the case setup for parametric studies.





Contents

Importing the Geometry
Working with Regions
Preparing the Meshes
Specifying Angle of Attack
Coupling the Overset Region
Setting up Physics Models
Setting Material Properties
Setting Initial Conditions
Setting Boundary Conditions
Setting Solver Parameters
Preparing a Scalar Scene
Monitoring Convergence
Running the Simulation

Visualizing the Results
Changing Angle of Attack
Clearing and Rerunning the Simulation
Visualizing the Results
Summary

Importing the Geometry

Load a simulation prepared with the geometry.

The simulation file contains geometry parts for:

- Background region of the mesh
- Overset region that contains the airfoil
- Several parts for use in volumetric mesh controls

To load the file:

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment.
2. Select **File** > **Load Simulation**.
3. In the *Load Simulation* dialog, navigate to the `doc/startutorialsdata/compressibleFlow/data` subdirectory of your STAR-CCM+ installation directory and select `airfoil_oversetParts.sim`.
4. Click **Open** to load the simulation.
5. Select **File** > **Save As** and save the simulation under the name `airfoil_oversetMesh.sim`.

Working with Regions

Create and define regions and boundaries for the domain.

First, assign the Background and Overset parts to regions:

1. Within the **Parts** node, multi-select the **Background** and **Overset** parts.
2. Right-click the selected parts and choose **Assign Parts to Regions...**.



3. In the *Assign Parts to Regions* dialog, select **Create a Region for Each Part** for the **Region Mode** and **Create a Boundary for Each Part Surface** for the **Boundary Mode**.
4. Click **Apply**, and then **Close**.
5. Select the **Regions** > **Background** > **Boundaries** > **Outside** node and set the **Type** to **Free Stream**.



6. Repeat this procedure for the **Symmetry** boundary, but set its type to **Symmetry Plane**.
7. For the **Overset** > **Boundaries** node, set the following properties:

Property	Setting
Airfoil	**Wall**
Overset Boundary	**Overset Mesh**
Symmetry	**Symmetry Plane**

The boundaries for the Overset region are shown below. The Overset Mesh boundary is highlighted.



The Overset Mesh boundary has no effect on the flow solution. It is used with an Overset Mesh interface to couple the overset and background regions.

Preparing the Meshes

Create two mesh continua to represent the Background and Overset regions.

The continua contain the same models, but use different model settings. It is important that, in the overlapping zone, which is defined by the Overset Boundary, mesh cells are of a comparable size in both meshes.

1. Create a mesh continuum and rename it `Background Mesh`.
2. Select the following models for the continuum, in order:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Trimmer**
Optional Models	**Prism Layer Mesher**

3. Click **Close**.



4. Copy and paste the **Background Mesh** continuum node into the **Continua** node.
5. Rename the **Copy of Background Mesh** node to `Overset Mesh`.
6. Edit the **Continua** > **Background Mesh** > **Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	**Value**	`0.5 m`
Template Growth Rate	**Default Growth Rate**	**Slow**
	Boundary Growth Rate	**Slow**

The completed *Edit* dialog looks as follows:



7. Click **Close**.
8. Right-click the **Continua** > **Background Mesh** > **Volumetric Controls** node and select **New**.
9. Select the **Volumetric Control 1** node and set **Parts** to **Overlap**.



10. Edit the **Volumetric Control 1 > Mesh Conditions** node and set the following properties:

Node	Property	Setting
Surface Remesher	**Customize surface remesher**	Activated
Trimmer	**Customize isotropic size**	Activated



11. Click **Close**.
12. Select the **Volumetric Control 1 > Mesh Values > Custom Size > Relative Size** node and set **Percentage of Base** to `12.5`.



Set properties for the overset mesh.

13. Select the **Continua > Overset Mesh > Models > Prism Layer Mesher** node and set **Stretching Mode** to **Wall Thickness**.



14. Edit the **Continua > Overset Mesh > Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	**Value**	`0.5 m`
Number of Prism Layers	**Number of Prism Layers**	`15`
Prism Layer Thickness > Relative Size	**Percentage of Base**	`15.0`
Template Growth Rate	**Default Growth Rate**	**Slow**
	Boundary Growth Rate	**Slow**
Thickness of Near Wall Prism Layer	**Value**	`1.0E-4 m`

The completed *Edit* dialog looks as follows:



15. Click **Close**.
16. As was done for the Background Mesh, create four new Volumetric Controls and give them the following names:
    - `Overlap`.

- `Airfoil`.
- `Leading Edge`.
- `Trailing Edge`.

17. In the **Mesh Conditions** of all four volumetric controls, set the following properties:

Node	Property	Setting
Surface Remesher	**Customize surface remesher**	Activated
Trimmer	**Customize isotropic size**	Activated

18. For each volumetric control, select the following parts and the following values for **Mesh Values** > **Custom Size** > **Relative Size**:

Volumetric Control	Part	Relative Size
Overlap	**Overlap**	`12.5`
Airfoil	**Airfoil**	`5.0`
Leading Edge	**Leading Edge**	`1.0`
Trailing Edge	**Trailing Edge**	`2.0`

With both mesh continua fully defined, assign them to their respective regions:

19. Select the **Regions** > **Background** node and make sure that **Mesh Continuum** is set to **Background Mesh**.
20. Select of the **Regions** > **Overset** node and set **Mesh Continuum** to **Overset Mesh**.
21. To build the volume mesh, click the **Generate Volume Mesh** button in the toolbar. This process can take a few minutes.
22. To view the mesh, create a mesh scene.
23. Zoom in on the Overset boundary to compare the cell size of the overset mesh to the background mesh at this boundary.



24. Save the simulation.

You can convert the mesh from 3D to 2D to simplify the solution.

25. In the menu bar, select **Mesh > Convert to 2D**
26. In the *Convert Regions to 2D* dialog, tick the **Delete 3D regions after conversion** option and click **OK**.

Specifying Angle of Attack

Rotate the Overset region to correspond to a 2.54° angle of attack of the airfoil.

The steps presented here demonstrate how using overset meshes allows you to change the relative position of bodies without the need to generate a mesh.

1. Right-click the **Regions > Overset 2D** node and select **Transform > Rotate**
2. In the **Rotate Regions** dialog, set the **Angle** to `-2.54 deg`.



3. Click **Apply**, and then click **Close**.

Coupling the Overset Region

Overset regions are coupled to the background region using the Overset Mesh interface.

Normally, this interface can be created before generating the background and overset meshes. To activate this interface feature, you must create a physics continuum. Current implementation requires that the background and overset regions use the same physics continuum. In this tutorial, you are creating the Overset Mesh interface after converting the mesh to 2D as this step deletes all previously created interfaces.

To create the interface:

1. Multi-select the **Background 2D** and **Overset 2D** regions.
2. Right-click the selected regions and choose **Create Interface** > **Overset Mesh**.



An **Overset Mesh 1** node appears under the **Interfaces**. Use the default **Distance Weighted Interpolation Option** for the mesh coupling.



When you initialize the Overset Mesh interface, STAR-CCM+ chooses acceptor cells on the overset region. Cells in the background region that are overlapped by the overset region are deactivated.

Contents

Analyzing Overset Mesh

Analyzing Overset Mesh

Display cell status in a scalar scene.

To analyze the mesh before running the simulation, initialize the Overset Mesh interface:

1. Right-click the **Interfaces** node and select **Initialize**.

   Notice that the background mesh within the airfoil part has been removed in *Mesh Scene 1*.



Analyze the overset mesh using the Overset Cell Status field function.

2. Create a scalar scene.

   The *Scalar Scene 1* display appears.

3. Right-click the scalar bar at the bottom of the display and select **Overset Cell Status** from the pop-up menu.



4. Select the **Scalar Scene 1 > Displayers > Scalar 1** node and set *Representation* to **Overset Mesh**.

First look at the cell status in the overset region.

5. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Parts** node and tick the **Overset 2D** region in the **Parts** property.



6. Select the **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Scalar Field** node and set the **Min** to `-2.0` and the **Max** to `1.0`.

   Acceptor cells outline the overset region in blue.



7. To see the background region cell status, change the **Parts** to **Background 2D** within the **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** node.

   The red cells in the background mesh are inactive.



Setting up Physics Models

Models define the spatial and temporal solution methods and the physical properties of the flow.

In this example, the flow is steady, turbulent, and compressible. Use the default Spalart-Allmaras Turbulence model and the ideal gas model. The analysis also uses the coupled solver, which is recommended for all supersonic and transonic compressible flows.

By default, a continuum called **Physics 1** is created when the three-dimensional mesh is generated and a **Physics 1 2D** continuum is created when the mesh is converted to two-dimensional. You can delete the **Physics 1** continuum, as only one is used in this simulation.

1. Delete the **Continua** > **Physics 1** node.

To use a more appropriate name for the **Physics 1 2D** continuum:

2. Rename the **Continua** > **Physics 1 2D** node to `Airfoil`

You can now choose the physics models for the continuum.

3. For the **Airfoil** continuum, select the following models, in order:

Group Box	Model
Space	**Two-Dimensional** (Selected automatically)
Gradient Metrics	**Gradients** (Selected automatically)
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
Time	**Steady**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**Spalart-Allmaras Turbulence**

4. **Close** the dialog.

The selected models now appear within the **Airfoil** node.



Setting Material Properties

Set the Reynolds number to 6.5 x 10^6, changing the dynamic viscosity

Use the following steps:

1. Open the **Continua** > **Airfoil** > **Models** > **Gas** > **Air** node.
2. Select the **Material Properties** > **Dynamic Viscosity** > **Constant** node and set *Value* to `4.61e-5 Pa-s`

Setting Initial Conditions

As in the Transonic Flow tutorial, initialize the velocity field to 247.967 m/s.

1. Select the **Continua** > **Airfoil** > **Initial Conditions** > **Velocity** > **Constant** node and set **Value** to `[247.975,0.0,0.0]` m/s.



To specify the initial temperature:

2. Select the **Static Temperature** > **Constant** node and set *Value* to `291 K`.

The default values for the remaining initial conditions are suitable for this tutorial.

Setting Boundary Conditions

Set the conditions at the boundaries of the airfoil and the domain.

The geometry used for this tutorial has only two boundaries for the physics:

- A wall boundary representing the surface of the airfoil and defined in the Overset region.
- A freestream boundary at the external edge of the entire solution domain defined in the Background region.

Set the freestream conditions on the outer perimeter of the Background region.

1. Select the **Regions** > **Background 2D** > **Boundaries** > **Outside** > **Physics Values** > **Mach Number** > **Constant** node and set **Value** to `0.725`



2. Select the **Static Temperature** > **Constant** node and set *Value* to `291 K`.

All other conditions for the freestream boundary and the default wall are suitable for this problem.

Setting Solver Parameters

The simplicity of this problem allows a rapidly converging solution to be obtained using a larger Courant number.

In problems involving more complex geometries or physics, attempting to shorten the run time in this way may cause the solution to diverge.

To increase the Courant number:

1. Select the **Solvers > Coupled Implicit** node and set *Courant Number* to `20.0`.



2. Save the simulation.

Preparing a Scalar Scene

View the Mach number profile during the run to monitor the supersonic flow region above the airfoil.

Use the following steps:

1. Create a scalar scene.

   A new *Scalar Scene* display appears.
2. Right-click the scalar bar at the bottom of the display and select **Mach Number > Lab Reference Frame** from the pop-up menu.
3. Initialize the run by clicking **Initialize Solution** in the toolbar, then use the middle mouse button to zoom in on the airfoil in the center of the scalar scene.



4. To change the style of the Mach number contours, select the **Scalar Scene 2 > Displayers > Scalar 1** node and set *Contour Style* to **Smooth Filled**.

Monitoring Convergence

To determine when the problem has converged, plot the lift and drag coefficients.

Use the following steps:

1. Right-click the **Reports** node and select **New Report > Force Coefficient**.



A new report node named **Force Coefficient 1** is created.

2. Rename this node `Drag Coefficient`
3. Select the **Drag Coefficient** node and set the following properties:

Property	Setting
Reference Density	`1.2126 kg/m^3`
Reference Velocity	`247.975 m/s`
Parts	**Regions > Overset 2D > Airfoil**



4. Right-click the **Drag Coefficient** node and select **Create Monitor and Plot from Report**.

   A new plot node appears named **Drag Coefficient Monitor Plot**.

5. To display the empty plot in the *Graphics* window, double-click the **Drag Coefficient Monitor Plot** node.
6. Repeat the steps described above to create and display a plot for the lift coefficient. Use the same settings as for the drag coefficient except rename the report node **Lift Coefficient** and set its *Direction* property to `[0.0, 1.0, 0.0]`
7. Save the simulation.

Running the Simulation

You are now ready to run the simulation.

Use the following steps:

1. Click **Run** in the toolbar.
   The Residuals display is created automatically and shows the progress of the solver.

During the run, it is possible to stop the analysis by clicking (**Stop**) in the toolbar. If you do halt the simulation, it can be continued later by clicking (**Run**). If left alone, the simulation continues until 1000 iterations have been completed.

2. Save the simulation after the solver has finished.

Visualizing the Results

The *Scalar Scene* display shows the Mach number profile at the end of the run.

The profile shows the transonic flow around the airfoil, including the shock wave produced behind the supersonic region on the upper surface.



The lift and drag coefficient plots provide a check for convergence. As there is little variation in both monitors, it is reasonable to conclude that the solution has converged.


...e 1
Scalar Scene 2
Drag Coefficient Monitor Plot
Lift Coefficient...
Monitor Plot
0.2
0.1
0
Drag Coefficient Monitor
0
200
400
600
800
1000
Iteration
Drag Coefficient Monitor



Changing Angle of Attack

Rotate the overset mesh.

One advantage of using an overset mesh is being able to transform regions without having to remesh. If, for instance, you want to study the airfoil at a different angle of attack, you can rotate the Overset region by a different angle and rerun the simulation.

1. Right-click the **Regions** > **Overset 2D** node and select **Transform** > **Rotate...** .
2. In the *Rotate Regions* dialog, set **Angle** to `-10 deg`.
3. Click **Apply**.



Note: The transformed overset region must be completely within that refined part of the background region where the background mesh is of the same size as the overset mesh.

Clearing and Rerunning the Simulation

Clear the solution and rerun the simulation using the same physics setup.

Use the following steps:

1. In the menu bar, select **Solution > Clear Solution**.
2. Click **Run** in the toolbar.

Visualizing the Results

Re-examine the scalar scene.

The scalar scene solution is shown below. You can see that if the angle of attack is increased further, the steep adverse pressure gradient on the

upper surface of the airfoil causes the flow to separate close to the leading edge, resulting in a large wake.



Summary

This Transonic Flow tutorial demonstrated how to use overset mesh to simulate transonic flow over an airfoil

This tutorial introduced the following features of STAR-CCM+:

- Creating and analyzing an overset mesh for an airfoil problem.
- Defining models for a compressible flow problem.
- Defining the material properties required for the selected models.
- Setting solver parameters for a steady-state run.
- Analyzing results and convergence.
- Modifying angle of attack.

Adjoint Flow Solver: Flow through an S-Bend

This tutorial demonstrates the basic workflow for setting up an adjoint flow analysis.

The geometry is an s-bend duct with compressible air flowing through it. The flow enters the domain through a stagnation inlet boundary and leaves it through a pressure outlet boundary. A wall boundary condition is imposed on the duct surface. The case is run at a Mach number of 0.2.



An adjoint flow analysis is set up and run in two steps:

1. The setup of the primal flow solution corresponds to the setup of any flow solution simulation using the coupled solver. Ideally, the primal flow solution should be converged down to machine precision. This is important for the adjoint flow solver to subsequently run robustly and produce optimal results.

Note: For industrial cases, this high level of convergence might not be achievable. The adjoint flow solver can still obtain good results as this depends on the individual case.

2. The adjoint solver setup is performed as a separate step after the primal flow solution has been run and converged. Separate solver parameters are available for the adjoint flow solver. The engineering objective to be analyzed by the adjoint method is defined by use of cost functions. Additional field functions are created that can be used to visualize the adjoint flow solution.

Contents

Importing the Volume Mesh
Selecting the Physics Models
Setting Boundary Conditions
Setting the Flow Solver Parameters and Stopping Criteria
Visualizing the Flow Solution
Running the Simulation
Visualization of Flow Results
Selecting the Adjoint Flow Model
Setting up the Adjoint Cost Function
Setting the Adjoint Flow Solver Parameters and Stopping Criteria
Running the Adjoint Flow Solver
Visualizing the Adjoint Flow Results
Summary

Importing the Volume Mesh

A cylindrical mesh has been prepared for this tutorial.

To import the mesh and save the simulation:

1. Create a directory for the tutorial, called `ductAdjoint`.
2. Copy the mesh file `sBend.ccm` to this directory from the `doc/startutorialsdata/adjoint/data` subdirectory of your STAR-CCM+ installation directory.
3. Start up a double precision version of STAR-CCM+ and create a new simulation by clicking **File > New Simulation...**
4. Select **File > Import > Import Volume Mesh**.
5. In the *Open* dialog, browse to the `ductAdjoint` directory, select `sBend.ccm`, and click **Open**.

The imported mesh appears in the *Geometry Scene 1* window.

6. Select **File** > **Save As** and save the simulation under the name `ductAdjointTutorial.sim`.

Selecting the Physics Models

Activate the models for the flow solution.

Use the following steps:

1. For the continuum **Continua** > **Physics 1** select the following models, in order:

Group Box	Model
Enabled Modes	**Three Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
K-Epsilon Two Layer All y + Wall Treatment	**Two-Layer All y+ Wall Treatment** (Selected automatically)

2. Click **Close**.

The selected models are added to the simulation tree.



Setting Boundary Conditions

Set the inflow and outflow boundary conditions:.

Use the following steps:

Set the **Stagnation Inlet** boundary.

1. Select **Regions** > **Duct** > **Boundaries** and set the following properties:

Node	Property	Setting
Stagnation_in	*Type*	**Stagnation Inlet**
Physics Values		
Total Pressure > Constant	*Value*	`2865.6 Pa`
Total Temperature > Constant	*Value*	`275.4 K`

Set the **Pressure Outlet** boundary:

2. Select **Regions** > **Duct** > **Boundaries** and set the following properties:

Node	Property	Setting
Pressure_out	*Type*	**Pressure Outlet**
Physics Values > Pressure > Constant	*Value*	`0.0 Pa.`

Setting the Flow Solver Parameters and Stopping Criteria

Set the Courant number, grid sequencing, and maximum time-steps.

Use the following steps:

To set the numerical parameters for the flow solution:

1. Select the **Solvers > Coupled Implicit** node and set *Courant Number* to `30`.
2. For a good initialization of the flow variables to enhance convergence, select the **Solvers > Coupled Implicit > Expert Initialization** and set *Method* to **Grid Sequencing**.



To set the stopping criteria:

3. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `800`.

Visualizing the Flow Solution

Create scenes for the velocity vectors and pressure contours on a plane bisecting the s-bend duct.

Use the following steps:

Create a section through the s-bend:

1. Right-click the **Derived Parts** node and select **New Part > Section > Plane...**
2. In the edit dialog under the **Plane Parameters** box, set the origin to [0.0,0.0,0.0] and the normal vector to [0.0,1.0,0.0].
3. Select **No Displayer** in the **Display** section.



4. Click **Create**, then click **Close**.

To visualize the velocity vectors, create a vector scene:

5. Right-click the **Scenes** node and select **New Scene > Vector**.
6. Click **Derived Parts > plane section** and drag and drop the **plane section** node into the open vector scene in the **Graphics** window.

7. In the **Displayer Selection** menu, select **Add to Vector 1**.
8. To obtain a side view of the section, press the s key.

To visualize the pressure contours, create a scalar scene:

9. Right-click the **Scenes** node and select **New Scene > Scalar**.
10. Click **Derived Parts > plane section** and drag and drop the **plane section** node into the open scalar scene in the **Graphics** window.
11. In the **Displayer Selection** menu, select **Add to Scalar 1**.
12. Click the **scene/plot** button.
13. Select the **Displayers > Scalar 1 > Scalar Field** node and set the *Function* to Pressure.
14. Select the **Displayers > Scalar 1** node and set *Contour Style* to **Smooth Filled**.
15. To obtain a side view of the section, press the s key.
16. Save the simulation.

Running the Simulation

The simulation is now ready to run.

Click the (**Run**) button.

While the simulation is running, you can click the tabs at the top of the **Graphics** window to view each of the scenes. The Residuals display is created automatically and shows the progress of the solvers.



Visualization of Flow Results

Velocity vectors are shown below.



Pressure contours are shown below.



Selecting the Adjoint Flow Model

Activate the adjoint flow model in addition to the physical flow models already selected:

1. Right-click the **Continua** > **Physics 1** > **Models** node and choose **Select models...**
2. In the *Model Selection* dialog:
   a) Select **Adjoint Flow** from the **Optional Models** group box.
   b) Click **Close**.

An additional **Adjoint Flow** node has been created under **Continua** > **Physics 1** > **Models**.



Setting up the Adjoint Cost Function

A cost function represents the objective of the adjoint simulation with respect to a result.

For this tutorial case, the objective is to reduce the pressure drop between inflow and outflow boundaries along the s-bend duct. Therefore, create a pressure drop cost function.

To set up the pressure drop cost function:

1. Right-click the **Reports** node and select **New Report > Pressure Drop**.
2. Select the **Reports > Pressure Drop 1** node and set the following properties:

Property	Setting
High Pressure	**[Duct:Stagnation_in]**
Low Pressure	**[Duct:Pressure_out]**

3. Right-click the **Adjoint Cost Functions** node and select **New > Pressure Drop Cost Function**.



4. Select the **Adjoint Cost Functions > PressureDrop** node and set *Pressure Drop Report* to **Pressure Drop 1**.

Setting the Adjoint Flow Solver Parameters and Stopping Criteria

You can find the numerical parameters specific to the adjoint flow solver under the **Solvers** > **Adjoint Flow** node.

To view and confirm the current default settings:

1. Select the **Solvers** > **Adjoint Flow** node.
2. Set the *Adjoint Driver Selection* to **Defect Correction**.

   The **Defect Correction** option invokes an iterative defect-correction algorithm that solves the flow adjoint equations. It is the default algorithm.
3. Make sure that the *Courant Number* is set to `100.0`.



Set the stopping criteria for the adjoint solver run. Increase the number of maximum iterations:

4. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `1200`.

Running the Adjoint Flow Solver

The adjoint flow solver is now ready to run.

The actions to step or run the adjoint flow solver are directly attached to the **Adjoint Flow** node. They are separated from the (Run) button that is used to start the primal flow solution.

To run the adjoint flow solver, right-click the **Solvers** > **Adjoint Flow** node and select **Run**.



Note: Do not run the primal flow solver after the adjoint solution has begun.

The adjoint flow solver starts to iterate. To follow the progress of the residuals, click the Residuals tab in the **Graphics** window. The flow variables of the primal flow solution are frozen during the adjoint solver run.

Visualizing the Adjoint Flow Results

In order to analyze the adjoint flow results, additional adjoint field functions are created.

You can use these field functions in the scenes to visualize the results.

To visualize the adjoint of pressure drop with regard to continuity:

1. Create a new scalar scene by right-clicking the **Scenes** node and selecting **New Scene** > **Scalar**.
2. Click the **scene/plot** button.
3. Select the **Displayers** > **Scalar 1** node and set the following properties:

Node	Property	Setting
	Contour Style	**Smooth Filled**
Parts	**Parts**	**Duct**
Scalar Field	*Function*	**Adjoint** > **Pressure Drop** > **Continuity**

A contour plot of the adjoint of pressure drop with regard to continuity is shown below.



The objective is to reduce the pressure drop through the geometry. Directing mass away from areas with positive values decreases the pressure drop. Red areas indicate that directing mass away from here has the most effect on lessening the pressure drop.

To visualize the adjoint of pressure drop with regard to the momentum in x direction:

4. Select the **Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Adjoint > Pressure Drop > X-momentum**.

   A contour plot of the adjoint of pressure drop with regard to the momentum in x direction is shown below.



Directing x-momentum towards areas with negative values decreases the pressure drop. Blue areas indicate that directing x-momentum here has the most effect on lessening the pressure drop.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Importing and visualizing a volume mesh.
- Using the coupled solver Courant number and the Grid Sequencing Initialization to enhance convergence.
- Running the primal solver to machine precision.
- Selecting the physics model for adjoint flow analysis.
- Setting up the adjoint cost function for pressure drop.
- Running the adjoint Flow solver.
- Visualizing the adjoint Flow solution.

Turbo Wizard: Single Row

This tutorial illustrates the simulation of an axial compressor blade in STAR-CCM+ and goes through the process of creating a mesh for a single blade row.

The blade geometry that is used in this tutorial is Rotor37 [415]. The cross-sectional cuts are stored in the `Rotor37.estg` file and imported into STAR-CCM+. The problem geometry is shown below:



The tutorial starts with Importing the Turbo Blades and Generating the Mesh.

Contents

Importing the Turbo Blades and Generating the Mesh
Checking the Models
Visualizing the Mesh
Setting Up Solver and Stopping Criteria
Visualizing the Solution
Running the Simulation
Analyzing Results

Summary
Single Row Tutorial Bibliography

Importing the Turbo Blades and Generating the Mesh

Create a subdirectory for the tutorial called `axialSingleRow`.

Copy the `Rotor37.estg`, `inlet.bndy`, `exit.bndy`, `hub.bndy` and `tip.bndy` files from the `doc/startutorialsdata/compressibleFlow/data/SingleRow` subdirectory of your STAR-CCM + installation into this directory.

`Rotor37.estg` contains geometry data for 13 aerofoil sections between the Hub and Tip in a standard formatted file:

- `inlet.bndy` contains a spline definition for the inlet boundary.
- `exit.bndy` contains a spline definition for the exit boundary.
- `hub.bndy` contains a spline definition for the hub boundary.
- `tip.bndy` contains a spline definition for the tip (shroud) boundary.

Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the **New Simulation** option from the menu bar.

STAR-CCM+ creates an additional tabbed window for the new simulation in the *Explorer* pane and give it the default name `Star 1`. Import the supplied data files into the simulation:

1. Select **Mesh** > **Import Turbo Blades** from the menu bar:



   The *Import turbo wizard geometry file options* dialog appears.

2. Make sure that **Number of blade rows** is set to `1`.

3. Make sure that **Select current blade row** is set to `1`.
4. Activate **Current blade row is rotating**.

   Additional options appear in the dialog.
5. Change the rotational speed to `17,200 rpm`.
6. Activate **Hub surface is rotating**.
7. Set *Select Flow Axis* to **+Z**.



There are six tabs on the *Import turbo wizard geometry file options* dialog. The *Geometry*, *Meshing*, *Flowpath*, *Grid Settings*, *Boundary Conditions / Turbulence* and *Initial Conditions* tabs are used to specify the required files and parameters for the imported blade geometry.

Geometry Tab

8. Set the following properties:

Property	Setting
Blade Geometry Definition/Surface Option	**Import blade geometry from external file**
Select Blade Grid Type	**O-Grid**
Blade geometry file for blade row 1	`axialSingleRow/Rotor37.estg`
Use open bottom of blade surface as hub section	Activated
Use open top of blade surface as shroud section	Activated
Number of blades around annulus	`36`
Blade Geometry Units	`cm` (centimeters)



Meshing Tab

9. Make sure that **Mesh generation method** is set to **Build 3D volume mesh using elliptical gridder**.



Flowpath Tab

10. Set the following properties:

Property	Setting	
Inlet boundary file for blade row 1	`inlet.bndy`	
Exit boundary file for blade row 1	`exit.bndy`	
Generate clearance grid for blade row 1	**Shroud Clearance grid**	This reduces y+ values near the blade tip, which

Property	Setting	
		improves convergence of the turbulence model.
Define clearance streamline by:	**Setting fixed distance value**	
Number of clearance streamlines	`5`	
Clearance distance	`0.0356 cm`	
Hub boundary file	`hub.bndy`	
Shroud boundary file	`tip.bndy`	
Number of total streamlines	`33`	This controls the number of streamlines from the inlet to outlet boundaries, on which the mesh layers between hub and shroud are constructed.
Initial hub streamline spacing	`0.015 cm`	This defines the thickness of the mesh layers next to the hub surfaces.
Initial shroud streamline spacing	`0.015 cm`	This defines the thickness of the mesh layers next to the shroud surfaces.



Grid Settings Tab

11. In the *Distribution* subtab, set the following properties:

Property	Setting
Initial Clustering, Leading Edge (DS_LE)	Activated, `0.003 cm`
Initial Clustering, Trailing Edge (DS_TE)	Activated, `0.003 cm`
Near Wall (Airfoil) Spacing (DS_IN)	Activated, `0.0015 cm`
Periodic Boundary Spacing (DS_OUT)	Activated, `0.03 cm`
Allow turbo wizard to move airfoil surface points to preserve grid orthogonality	Activated

The number of iterations for this option can be left at the default value.



12. Accept the default settings on the *Number Of Points* subtab.
13. On the *Block Boundary Settings* subtab, set the following properties:

Property	Setting
Upstream Cap, Inlet Boundary Spacing Ratio (ZINO)	Activated, `0.15`
Downstream Cap, Outlet Boundary Spacing Ratio (ZEXO)	Activated, `0.15`



Boundary Conditions / Turbulence Tab

14. **Automatically set-up default turbulence models?** is activated and grayed out, since this setting is enabled automatically with the rotation of the hub and/or the shroud.
15. Accept the default settings on the *Stagnation Inlet* subtab.
16. In the *Pressure Outlet* subtab, change the **Constant static pressure** value to `5000 Pa`, leaving the other parameters at their defaults.



Initial Conditions Tab

17. Change the **Initial Velocity Z** value to `150 m/sec`, leaving the other parameters at their defaults.

    This sets a non-zero velocity in the axial direction.



18. Click **Save Turbo Wizard Settings** and, in the *Save Turbo Wizard Data* dialog, navigate to the `axialSingleRow` directory created previously.
19. **Save the data as** `turboWizardData`.

The saved turbo wizard data can then be loaded and modified for the same case or used to set up other cases.

20. To generate the volume mesh, click **OK**.

STAR-CCM+ starts creating the grid plane for all streamlines. Once the mesh is generated, STAR-CCM+ converts it into finite volume representation. See the *Output* window for more details. At the end of the import process, STAR-CCM+ also creates an initial (default) simulation tree in the *Star 1* window.

To create a file for the simulation, and give it a proper name:

21. Select **File** > **Save** from the menu bar.

22. Type `singleRow` in the **File Name** text box then click **Save**.

    The title of the simulation window also changes to *singleRow*.

Checking the Models

By default the Turbo Wizard sets up a steady-state simulation in which:

- A three-dimensional grid is employed.
- The fluid is a compressible gas (air) that obeys the ideal gas equation of state.
- The flow is turbulent and non-isothermal.
- The K-Omega model is used for representing turbulence effects.
- The problem is solved using the coupled flow solver.

These settings can be checked by expanding the Models node of the Physics 1 continuum.

1. Expand the **Regions > Turbo Blade Row 1 > Boundaries > Hub Surface Blade Row 1 > Physics Conditions** node and select **Reference Frame Specification**.

   The *Option* property is set to **Region Reference Frame**.



2. Expand the **Regions > Turbo Blade Row 1 > Physics Values** node and select **Motion Specification**.

   The *Reference Frame* property is set to **Rotating in Turbo Blade Row 1**.



3. The properties of this reference frame can be viewed by expanding the **Tools** > **Reference Frames** node and selecting **Rotating in Turbo Blade Row 1**.



Visualizing the Mesh

To inspect the mesh:

1. Right-click the **Scenes** node and select **New Scene** > **Mesh**.

The *Graphics* window also shows the mesh. Examine the mesh more closely:

2. Produce a *Graphics* window display such as the one shown in the following screenshot.



3. Verify that no malformed or irregular cells are present on the mesh surface.

Setting Up Solver and Stopping Criteria

By default, the solver parameters are set to the Coupled Implicit with a Courant number of 5.

Increase this number to reduce the number of iterations the solver needs to converge. Also uses Grid Sequencing and Expert Driver to make the solver more robust and control the stability of the solution.

1. Select the **Solvers > Coupled Implicit**node and set **Courant Number** to `100.0`.



2. Open the **Solvers > Coupled Implicit > Expert Initialization** node and set **Method** to **Grid Sequencing**.



3. Open the **Expert Initialization > Grid sequencing** node and set **CFL number** to `100.0`.

   This CFL number is specific to the grid sequencing setup and does not need to be equal to the Courant number defined earlier.



4. Open the **Solvers > Coupled Implicit > Solution Driver** node and set **Driver** to **Expert Driver**. Use the Expert Driver's default settings.

Before running the analysis, reduce the maximum number of steps to account for increasing the Courant number and using Grid Sequencing and Expert Driver to control the solver's convergence.

5. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `500`.



Visualizing the Solution

In this section you define three isosurfaces including blade-to-blade and meridional views, to create a Relative Mach number scene.

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.

   A new node that is called **Scalar Scene 1** is created in the simulation tree.

2. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Scalar Field** node and set *Function* to **Relative Mach Number**.



3. Open the **Tools** node.
4. Right-click the **Coordinate Systems** node and select **New Block-Mapped**.

   A new node, **Block-Mapped 1**, appears in the simulation tree.



5. Select the **Block-Mapped 1** node and set the *Interpolation Points* to `[20, 20, 24]`

   This setting gives greater resolution in the meridional direction.



6. Select the **Umin** node and click in the right half of the *Boundaries* property.

   The *Select Objects* dialog opens.

7. Activate **Regions > Turbo Blade Row 1 > Left Blade Cyclic Blade Row 1** and click **OK**.



8. Repeat the above steps for the following nodes:

Node	Property	Setting
Umax	*Boundaries*	**Right Blade Cyclic Blade Row 1**
Vmin	*Boundaries*	**Hub Surface Blade Row 1**
Vmax	*Boundaries*	**Shroud Surface Blade Row 1**

Node	Property	Setting
Wmin	*Boundaries*	**Blade Inlet Blade Row 1**
Wmax	*Boundaries*	**Blade Outlet Blade Row 1**

Create a derived part:

9. Right-click the **Derived Parts** node and select option **New Part > Isosurface...**

   An interactive in-place dialog (*Create Isosurface*) appears on the left of the *Graphics* window to help you define the surface.

   In this dialog:

   a) In the **Scalar** box, select **Block-Mapped 1: Wmin to Wmax**.

   b) Change the isovalue to `0.01`.

   c) From the *Display* box, select **Existing Displayer** and select `Scalar 1`.

   d) Click **Create**, but not **Close**, as you wish to create further iso surfaces.

10. Click the *simulation* tab at the top of the simulation explorer. A new node, **iso**, appears under the **Derived Parts** node in the simulation tree.



11. Rename the **iso** node to `AxialView`.
12. Click the *edit* tab to switch panes, and in the *Create Isosurface* dialog:

    a) Select **Block-Mapped 1: Vmin to Vmax** in the **Scalar** box.

    b) Change the isovalue to `0.5`.

    c) Click **Create** and switch back to the simulation tree by clicking the *simulation* tab at the top.

In the simulation tree:

13. Rename the **iso** node to `Blade2BladeView`.
14. Click the *edit* tab to switch panes, and in the *Create Isosurface* dialog:

    a) Select **Block-Mapped 1: Umin to Umax** in the **Scalar** box.

b) Change the isovalue to `0.025`.

c) Click **Create and then Close**.

In the simulation tree:

15. Rename the **iso** node to `MeridionalView`.
16. Click **(Make Scene Transparent)** to display the above-defined isosurfaces as shown in the following screenshot.



17. Click **(Restore Scene Transparency)** to set the transparency of the scene to its original setting.
18. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1**node and set *Contour Style* to `Smooth Filled`.
19. Save the simulation.

The pre-processing task is now complete. The next step is to run the analysis, as described in Running the Simulation.

Running the Simulation

Run the simulation.

1. Click **(Run)** to start the solver.

The analysis starts and the results, in the form of a graph of residuals vs. iteration number, is displayed automatically in the *Graphics* window.

The solution stops after 500 steps.

2. Save the simulation.

Analyzing Results

This section of the tutorial covers post-processing.

To assess the degree of convergence, it is helpful to look at the inlet and outlet mass flow plots to make sure that the inlet mass flow is equal to the outlet mass flow. For an inlet area, the normal vector introduces a negative sign.

1. Open the **Plots** node.
2. Right-click **Plots** > **Inlet Blade Row 1 Mass Flow Plot** node and select **Open**.

   This command opens the *Inlet Blade Row 1 Mass Flow Plot* scene in the *Graphics* window.



3. Similarly, you can view the *Outlet Blade Row 1 Mass Flow Plot*.



The scalar scene is shown below.



4. Save the simulation.

Summary

This STAR-CCM+ tutorial introduced the following features:

- Starting the code and creating a simulation
- Importing the geometry and flowpath boundary data for an axial compressor blade
- Generating a volume mesh
- Defining the simulation models
- Visualizing the mesh
- Setting up solver and stopping criteria
- Setting up a scalar scene with data from derived parts
- Running the solver for a set number of iterations
- Analyzing the results using the built-in plotting and visualization facilities

Single Row Tutorial Bibliography

Reid, L and Moore, R. D., 1978. "Performance of Single-Stage Axial-Flow Transonic Compressor With Rotor and Stator Aspect Ratios of 1.19 and 1.26, Respectively, and With Design Pressure Ratio of 1.82", NASA TP-1338.

Turbo Wizard: Multiple Rows

Having successfully created a mesh for a single row axial blade and run it in the previous section, add one additional row to the problem geometry in this tutorial.

The two blades that are used are the Rotor37 axial compressor blade, with an additional tip clearance, and an Inlet Guide Vane (IGV) positioned upstream of it. You define blade row 1 to be the IGV and blade row 2 to be the compressor rotor blade. The cross-sectional cuts for Rotor37 and IGV are stored in `Rotor37.estg` and `IGV.estg` files, respectively. These files are imported into STAR-CCM+.

The problem geometry is shown below:



Contents

Importing the Turbo Blades
Checking the Models
Visualizing the Mesh
Setting Up Solver and Stopping Criteria
Visualizing the Solution
Running the Simulation
Analyzing Mass Flow Reports

Summary

Importing the Turbo Blades

Create a subdirectory for the tutorial called `axialMultipleRow`.

1. Copy the following files from the `doc/startutorialsdata/compressibleFlow/data/MultipleRow` subdirectory of your STAR-CCM+ installation into this directory:

File	Description
`Rotor37.estg`	Contains the geometry data of the Rotor for 13 airfoil sections between the Hub and Tip in a standard formatted file.
`IGV.estg`	Contains the geometry data of the IGV for two airfoil sections between the Hub and Tip in a standard formatted file.
`inlet1.bndy`	Contains a spline definition for the inlet boundary for the IGV.
`inlet2.bndy`	Contains a spline definition for the inlet boundary for the Rotor.
`exit1.bndy`	Contains a spline definition for the exit boundary for the IGV.
`exit2.bndy`	Contains a spline definition for the exit boundary for the Rotor.
`hub.bndy`	Contains a spline definition for the hub boundary.
`tip.bndy`	Contains a spline definition for the shroud (tip) boundary.

2. Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the New Simulation option from the menu bar.

   STAR-CCM+ creates an additional tabbed window for the new simulation in the *Explorer* pane and gives it the default name Star 1.

3. Import the supplied data files into the simulation:

   a) Select **Mesh** > **Import Turbo Blades** from the menu bar:



This selection opens the *Import turbo wizard geometry file options* dialog.

b) Change the **Number of blade rows** to 2.

The following tabs are used here to specify the files and parameters for creating a volume mesh for the imported blade geometries:

- *Geometry*
- *Meshing*
- *Flowpath*
- *Grid Settings*
- *Boundary Conditions / Turbulence*
- *Initial Conditions*

Contents

Blade Row 1 (IGV)
Blade Row 2 (Rotor)

Blade Row 1 (IGV)

Set specifications for the first blade row.

1. Set the **Select current blade row** to 1.
2. Make sure **Current blade row is rotating** is de-selected, as the blade row is static.
3. Choose **+Z** from the *Select Flow Axis* drop-down menu.



Geometry Tab

4. Click **Browse** and navigate to the `axialMultipleRow` directory created previously.



5. Select `IGV.estg` and set the following properties:

Property	Setting
Select Blade Grid Type	**O-Grid**
Blade geometry file for blade row 1	`IGV.estg`
Use open bottom of blade surface as hub section	Deactivated
Use open top of the blade surface as shroud section	Deactivated
Number of blades around annulus	`48`

Property	Setting
Blade Geometry Units	**cm**



Meshing Tab

6. Set **Mesh generation method** to **Build 3d volume mesh using elliptical gridder**.

Flowpath Tab

7. Select the following files and set the following properties:

Property	Setting	
Inlet boundary file for blade row 1	`inlet1.bndy`	
Exit boundary file for blade row 1	`exit1.bndy`	
Generate clearance grid for blade row 1	**No clearance grid**	The IGV is static, so no clearance is required.
Hub boundary file for blade row 1	`hub.bndy`	

Property	Setting	
Shroud boundary file for blade row 1	`tip.bndy`	
Number of streamlines	`33`	This controls the number of streamlines from the inlet to the outlet boundaries of the Rotor, on which the mesh layers between hub and shroud are constructed.
Continuous streamlines across blade rows	Activated	
Initial hub streamline spacing	`0.015 cm`	
Initial shroud streamline spacing	`0.015 cm`	



Grid Settings Tab

8. Select the *Distribution* subtab.
9. Activate **Periodic Boundary Spacing** and enter `0.03 cm` in the text box.



10. Select the *Number Of Points* subtab and enter the following properties:

Property	Setting
Upstream Cap, Axial Points	Activated, `12`
Downstream Cap, Axial Points	Activated, `5`



11. Select the *Block Boundary Settings* subtab and enter the following properties:

Property	Setting
Upstream Cap, Inlet Boundary Spacing Ratio	Activated, `0.5`
Downstream Cap, Outlet Boundary Spacing Ratio	Activated, `0.85`



Boundary Conditions / Turbulence Tab

12. Accept the default settings.



Initial Conditions Tab

13. Change the **Initial Velocity Z**value to `150 m/sec`, leaving the other parameters at their defaults.

    This sets a non-zero velocity in the axial direction.



Blade Row 2 (Rotor)

Set specifications for the second blade row.

1. Set the following properties:

Property	Setting
Select current blade row	2
Current blade row is rotating	Activated, `17200 rpm`
Hub surface is rotating	Activated



Geometry Tab

2. Click **Browse**, navigate to `Rotor37.estg`, then click **Open**.
3. Set the following properties:

Property	Setting
Select Blade Grid Type	**O-Grid**
Blade geometry file for blade row 2	`Rotor37.estg`
Use open bottom of blade surface as hub section	Activated
Use open top of the blade surface as shroud section	Activated
Number of blades around annulus	`36`
Blade Geometry Units	**cm**



Meshing Tab

4. Set **Mesh generation method** to **Build 3d volume mesh using elliptical gridder**.

Flowpath Tab

5. Select the following files and set the following properties:

Property	Setting
Inlet boundary file for blade row 2	`inlet2.bndy`
Exit boundary file for blade row 2	`exit2.bndy`
Generate clearance grid for blade row 2	**Shroud clearance grid**
Define clearance streamline by	**Setting fixed distance value**
Number of clearance streamlines	`5`
Clearance distance	`0.0356 cm`
Hub boundary file	`hub.bndy`
Shroud boundary file	`tip.bndy`
Number of total streamlines	`33`

Property	Setting
Continuous streamlines across blade rows	Activated
Initial hub streamline spacing	`0.015 cm`
Initial shroud streamline spacing	`0.015 cm`



Grid Settings Tab

6. Select the *Distribution* subtab and set the following properties:

Property	Setting
Near Wall (Airfoil) Spacing	Activated, `0.0015 cm`
Periodic Boundary Spacing	Activated, `0.03 cm`
Allow turbo wizard to move airfoil surface point to preserve grid orthogonality	Activated

Leave the number of iterations at the default setting.



7. Select the *Number Of Points* subtab and set the following properties:

Property	Setting
Upstream Cap, Axial Points	Activated, `14`
Downstream Cap, Axial Points	Activated, `12`



8. Select the *Block Boundary Settings* subtab and set the following properties:

Property	Setting
Upstream Cap, Inlet Boundary Spacing Ratio	Activated, `0.75`
Downstream Cap, Outlet Boundary Spacing Ratio	Activated, `0.75`



9. Select the *Pressure Outlet* subtab, change the **Constant static pressure** value to `5000 Pa`, and leave the other parameters at their defaults



10. Click **Save** Turbo Wizard Settings and, in the *Save Turbo Wizard Data* dialog, navigate to the `axialMultipleRow` directory created previously.
11. Enter `turboWizardData` for the **File Name** then click **Save**.
12. Click **OK** to generate the volume mesh.

STAR-CCM+ starts creating the grid plane for all the streamlines for row (stage) 1 and then 2. Once it completes building the mesh, it converts both meshes into finite volume representation. See the *Output* window for more details. At the end of the import process, it also creates an initial (default) simulation tree in the *Star 1* window.

To create a file for the simulation, and give it a proper name:

13. Select **File** > **Save** from the menu bar.

14. Type `multipleRow` in the **File Name** text box then click **Save**.

Note: the title of the simulation window also changes to *multipleRow*.

Checking the Models

The Turbo Wizard has generated a steady-state flow model in which:

- A three-dimensional grid is employed
- The fluid is a compressible gas (air) that obeys the ideal gas equation of state
- The flow is turbulent and non-isothermal
- The K-Omega model is used for representing turbulence effects
- The problem is solved using the coupled flow solver

These settings are all made automatically when using the Turbo Wizard, and so no further selections are required. The models can be checked as described in the corresponding section of the Axial Single Row tutorial.

1. Expand the **Regions** > **Turbo Blade Row 1** > **Boundaries** node, go to the **Hub Surface Blade Row 1** > **Physics Conditions** node, and select **Reference Frame Specification**. Set the *Option* property to **Lab Frame**.



2. Expand the **Regions > Turbo Blade Row 2 > Boundaries** node, go to the **Hub Surface Blade Row 2 > Physics Conditions** node, and select **Reference Frame Specification**. Set the *Option* property to **Region Reference Frame**.



Note: the interface between the IGV and rotor is automatically set up as an implicit mixing-plane interface. The presence of this interface can be checked by expanding the Interfaces node.

3. Expand the **Interfaces** node and select node **Mixing Plane 1**



The *Properties* window shows the boundaries that are used in the composition of the interface.



An in-place interface has been created for the tip clearance. Repeat the process above to inspect its property settings.

Visualizing the Mesh

To examine the mesh:

Right-click the **Scenes** node and select **New Scene > Mesh** .

This creates a node, **Mesh Scene 1**, in the simulation tree.



In addition to the **Mesh Scene 1** node, the *Graphics* window also shows the mesh. Examine the mesh more closely by holding and dragging:

- the left mouse button to rotate it
- the right button to pan it
- the middle button to zoom in and ou

and produce a *Graphics* window display such as the one shown below.



Verify that no malformed or irregular cells are present on the mesh surface.

View each of the interfaces to check that they have been created appropriately.

Setting Up Solver and Stopping Criteria

By default, the solver parameters are set to the **Coupled Implicit** with a *Courant Number* of `5.0`.

Increase this number to reduce the number of iterations the solver needs to converge. Also utilize Grid Sequencing and Expert Driver to make the solver more robust and control the stability of the solution.

1. Open the **Solvers** node and then select the **Coupled Implicit** node and set the Courant Number to `50.0`.



2. Open the **Solvers** > **Coupled Implicit** > **Expert Initialization** node and set the **Method** to **Grid Sequencing**.



3. Open the **Grid sequencing** node and set the **Maximum grid levels**to `5`.



Grid Sequencing creates coarse grid levels that are based on the specified maximum grid levels setting. It is necessary to reduce the maximum number of grid levels in this case because you are using few cells in your original mesh. If there are too many grid levels relative to the number of cells in the mesh, the coarsest of the levels

might end up with cell numbers that are too small for the solver to converge.

4. Open the **Solvers** > **Coupled Implicit** > **Solution Driver** node and set **Driver** to **Expert driver**. Use the default settings of the Expert Driver.

Before running the analysis, reduce the maximum number of steps.

5. Open the **Stopping Criteria** node and then select the **Maximum Steps** node and set the *Maximum Steps* value to 500.



Visualizing the Solution

In this section you define three isosurfaces including blade-to-blade and meridional views, to create a Relative Mach number scene.

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.

   A new node that is called **Scalar Scene 1** is created in the simulation tree.

2. Open the **Scalar Scene 1** > **Displayers** > **Scalar 1** node and click the right half of the *Contour Style* property, then select **Smooth Filled** from the pop-up menu.



3. Select the **Scalar Field** node and edit the *Function* property.
4. Select **Relative Mach Number** from the list displayed.



5. Use the mouse buttons to change the direction of the *Scalar Scene 1* to correspond roughly to that of the *Mesh Scene 1* above.



Define a block mapped coordinate system for the IGV.

6. Open the **Tools** node.
7. Right-click the **Coordinate Systems** node and select **New Block-Mapped**.

   A node, **Block-Mapped 1**, appears in the simulation tree as soon as this operation is complete.



8. Select the **Block-Mapped 1** and set *Region* to **Turbo Blade Row 1**.



9. Click **Umin** and click to the right of the *Boundaries* property.

   This opens the *Select Objects* dialog.



10. Select **Regions > Turbo Blade Row 1 > Hub Surface Blade Row 1** and click **OK**.



11. Repeat the above steps for the following nodes:

Node	Property	Setting
Umax	*Boundaries*	**Shroud Surface Blade Row 1**
Vmin	*Boundaries*	**Left Blade Cyclic Blade Row 1**
Vmax	*Boundaries*	**Right Blade Cyclic Blade Row 1**
Wmin	*Boundaries*	**Blade Inlet Blade Row 1**
Wmax	*Boundaries*	**Blade Outlet Blade Row 1**

Now follow the same steps to define a block mapped coordinate system for the Rotor37 blade:

12. Right-click the **Coordinate Systems** node and select **New Block-Mapped**.
13. Select the **Block-Mapped 2** node and set *Region* to **Turbo Blade Row 2**.



14. Set the following properties:

Node	Property	Setting
Umin	*Boundaries*	**Hub Surface Blade Row 2**
Umax	*Boundaries*	**Shroud Surface Blade Row 2**
Vmin	*Boundaries*	**Left Blade Cyclic Blade Row 2**
Vmax	*Boundaries*	**Right Blade Cyclic Blade Row 2**
Wmin	*Boundaries*	**Blade Inlet Blade Row 2**
Wmax	*Boundaries*	**Blade Outlet Blade Row 2**

In this section, Isosurfaces for the IGV are created.

15. Right-click the **Derived Parts** node and select **New Part > Isosurface**.

    A new interactive in-place dialog (*Create Isosurface*) appears on the left of the *Graphics* window to help you define the desired surface.

16. Set the following properties:

Property	Setting
Scalar	**Block-Mapped 1: Wmin to Wmax**
Iso value	`0.01`
Display	**Existing Displayer, Scalar 1**



17. Click **Create**.
18. Switch back to the object tree using the **simulation** button above the *Create Isosurface* dialog.



A node, **Iso**, appears under the **Derived Parts** node in the simulation tree.



19. Switch back to the *Create Isosurface* dialog using the **edit** button.
20. Set the following properties:

Property	Setting
Scalar	**Block-Mapped 1: Umin to Umax**
Iso value	`0.5`

21. Click **Create**.

A new node, **iso 2**, appears under the **Derived Parts** node in the simulation tree.

22. Set the following properties:

Property	Setting
Scalar	**Block-Mapped 1: Vmin to Vmax**
Iso value	`0.01`

23. Click **Create**.

A new node, **iso 3**, appears under the **Derived Parts** node in the simulation tree. All three nodes are shown below.



24. Rename the nodes:

Old Name	New Name
iso	`AxialViewIGV`
iso 2	`Blade2BladeViewIGV`
iso 3	`MeridionalViewIGV`

Follow the above outlined procedure to create the isosurfaces for the Rotor blade.

25. Create the following isosurfaces:

Isosurface Scalar	Iso Value
Block-Mapped 2: Wmin to Wmax	`0.01`
Block-Mapped 2: Umin to Umax	`0.5`
Block-Mapped 2: Vmin to Vmax	`0.01`

26. Click **Close** on the *Create Isosurface* dialog when these operations are complete.

This results in three new nodes with the same default names as before (**iso**, **iso 2**, and **iso 3**).

27. Rename the nodes:

Old Name	New Name
iso	`AxialViewRotor`
iso 2	`Blade2BladeViewRotor`
iso 3	`MeridionalViewRotor`



Note: the reordering that is shown above is achieved by right-clicking the **Derived Parts** node and selecting **Refresh** from the pop-up menu.

To create a plot of blade profile pressure variation, you now create an isosurface interaction at the blade surface.

28. Right-click the **Derived Parts** node and select option **New Part > Isosurface** as before.
29. In the *Create Isosurface* dialog, click in the **Input Parts** box.

    This opens the *Select Objects* dialog as shown in the following screenshot.



30. Use the column selection button to clear the selections.
31. Select **Regions** > **Turbo Blade Row 2** > **Blade Surface Blade Row 2**and click **OK**
32. Set the following properties:

Property	Setting
Scalar	**Block-Mapped 2: Umin to Umax**
Iso value	`0.5`
Display	**Existing Displayer**, **Scalar 1**

33. Click **Create** and then **Close**.

A new node, **iso**, appears under the **Derived Parts** node in the simulation tree.

34. Rename the **iso** node to `Rotor 50% Section`

At the end of a successful completion of creating isosurfaces, the simulation tree contains the following additional nodes.



35. Select the **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Parts** node and click the right half of the *Parts* property.

The *Select Objects* dialog appears.

36. Make sure that only those objects belonging to the above defined isosurfaces are selected, as shown below



37. Click **OK**.

38. Click **(Make Scene Transparent)** to display the above defined isosurfaces as shown below.



In the above picture, the Rotor 50% Section isosurface is superimposed to show its location more clearly by clicking its node under the **Derived Parts** node.

39. Click **(Make Scene Transparent)** to set the transparency of the scene to its original setting.
40. Save the simulation.

The pre-processing task is now complete. The next step is to run the analysis, as described in Running the Simulation.

Running the Simulation

Run the simulation.

1. Click **(Run)** to start the solver.

   The analysis now starts and the results, in the form of a graph of residuals vs. iteration number, is displayed automatically in the *Graphics* window.

   The solution stops after 500 steps.
2. When the run is complete, save the simulation.

Analyzing Mass Flow Reports

To assess the degree of convergence, it is helpful to look at the inlet and outlet mass flow plots for the IGV and Rotor.

1. Open the **Plots** node.
2. Double-click **Inlet Blade Row 1 Mass Flow Plot**.



3. Double-click **Outlet Blade Row 1 Mass Flow Plot**.



4. Double-click **Inlet Blade Row 2 Mass Flow Plot**.



5. Double-click **Outlet Blade Row 2 Mass Flow Plot**.



The mass flows reported for Blade Row 1 are different from those reported for Blade Row 2. This difference is expected as Blade Row 1 has 48 blades and Blade Row 2 has 36 blades. Running the reports manually and multiplying by the appropriate number of blades gives the total mass flow through the system.

6. Right-click **Reports > Inlet Blade Row 1 Mass Flow Report**. Select item **Run Report** from the pop-up menu.



The result of this operation is shown in the output window.



7. Repeat the operation for the **Inlet Blade Row 2 Mass Flow Report** and note the result.

The results from these two reports are multiplied by the respective number of blades to give the overall mass balance.

Mass Flow Report	Reported Mass Flow (kg/s)	Number of Blades	Total Mass Flow (kg/s)
Inlet Blade Row 1	0.42733	48	20.5119
Inlet Blade Row 2	0.56977	36	20.5119

This comparison can be used to assess the overall convergence of the solution between the blade rows for the current computational mesh.

Contents

Analyzing Mach Number Scalar Scene
Analyzing Blade Pressure Profile Plot

Analyzing Mach Number Scalar Scene

Open the scalar scene showing the Relative Mach Number to view the solution.

To help visualization, only plot the scalar on the Blade2BladeViewIGV and Blade2BladeViewRotor isosurfaces.

1. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts**node and click the right half of the *Parts* property.

   The *Select Objects* dialog appears.
2. In the Selected list, de-select all except the **Blade2BladeViewIGV** and **Blade2BladeViewRotor** objects as shown below.



The scene resembles the one shown below.



Analyzing Blade Pressure Profile Plot

You now create an XY plot of blade profile pressure variation on the Rotor 50% Section isosurface.

1. Right-click the **Plots** node and select **New Plot > X-Y**.



A new node, **XY Plot 1**, appears in the simulation tree and XY Plot 1 graph is displayed automatically in the *Graphics* window. The new **XY Plot 1** node is selected automatically.

2. Click the right half of the *Parts* property and select Derived Parts > Rotor 50% Section.



3. Set the following properties:

Node	Property	Setting
XY Plot 1 > X Type > Vector Quantity	*Value*	`[0, 0, 1] cm`
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Pressure**

The *XY Plot* is displayed as shown in the following screenshot.



Summary

This STAR-CCM+ tutorial introduced the following features:

- Starting the code and creating a simulation
- Importing the geometry and flowpath boundary data for an axial compressor and Inlet Guide Van blade geometries
- Generating a volume mesh

- Defining the simulation models
- Visualizing the mesh
- Setting up solver and stopping criteria
- Setting up a scalar scene with data from derived parts
- Running the solver for a set number of iterations
- Analyzing the results using the built-in plotting and visualization facilities, including the XY Plot of a scalar variation

Harmonic Balance: Single Stage Periodic Flow

The tutorial illustrates the steps necessary to modify the single axial compressor case so that the effect of a periodic wake can be simulated using the Harmonic Balance method.



As this case uses the model and simulation resulting from the Single Row Tutorial, it is necessary to complete that tutorial before continuing.

Contents

Loading the Single Row Tutorial Simulation
Selecting the Physical Models for Harmonic Balance
Defining Harmonic Balance Parameters
Defining Boundary Conditions
Setting Solver and Stopping Criteria Parameters
Visualizing the Solution
Running the Simulation
Analyzing Results
Summary

Loading the Single Row Tutorial Simulation

Begin with loading the previously created single row simulation.

Note: To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file.

1. Create a subdirectory for the tutorial called `harmonicBalanceTut`.
2. In this directory, start up STAR-CCM+ in a manner that is appropriate to your working environment, and select **File** > **Load Simulation** from the menu bar.
3. In the *Load Simulation* dialog, use the **Browse** button to navigate to the location of the previously saved `singleRow.sim` file, select it, and click **Open**.
4. The *Load Simulation* dialog shows the path to the selected file in the *File* textbox.
5. Click **OK** to launch a server process and load the selected simulation.



Save the simulation with a different file name before changing anything:

6. Select **File > Save As...** from the menu bar.
7. In the *Save* dialog, navigate to the `harmonicBalanceTut` directory and enter `harmonicBalance` in the *File Name* textbox.
8. Click **Save**.

Selecting the Physical Models for Harmonic Balance

As physical models were previously selected for the single row tutorial, it is necessary to deactivate a number of the models so that the required options are made available.

1. For the physics continuum **Continua** > **Physics 1**, deactivate the following models from the **Enabled Models** list:
   - **Transition Boundary Distance**
   - **Turbulence Suppression**
   - **All y+ Wall Treatment**
   - **SST (Menter) K-Omega**
   - **K-Omega Turbulence**
   - **Reynolds-Averaged Navier-Stokes**
   - **Turbulent**
   - **Coupled Energy**
   - **Ideal Gas**
   - **Coupled Flow**
   - **Steady**
2. Activate **Auto-select recommended models**.
3. Select the following models, in order:

Group Box	Model
Time	**Harmonic Balance**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**Spalart-Allmaras Turbulence** **HB Standard Spalart-Allmaras** (Selected automatically) **All y+ Wall Treatment** (Selected automatically)

4. Click **Close**.
5. To see the selected models, expand the **Continua** > **Physics 1** > **Models** node.



Defining Harmonic Balance Parameters

The Harmonic Balance Flow and Energy model requires definition of blade row parameters.

In this example, a single blade row that consists of 36 blades is required. The number of modes is set to 3, which provides a sufficiently accurate solution and keeps the demand on system resources relatively low.

To specify the required parameters:

1. Expand the **Continua** > **Physics 1** > **Models** > **Harmonic Balance Flow and Energy** manager. right-click the **Blade Rows** node, and select **New**.

   A new node named **Blade Row 1** appears under the **Blade Rows** node and is selected automatically.
2. Select the **Blade Row 1** node and set *Modes* to 3.

3. Select the **Blade Row 1 > Blades Per Pitch node** and set *Number of Blades* to `36` .
4. Leave the *Pitch* property at its default setting of radians.



Defining Boundary Conditions

The Non-Reflecting option is activated for both the inlet and outlet boundaries.

The Wake specification at the inlet is also set to define the periodic perturbation. This consists of defining the characteristics of the wake profile and the upstream blade row. In this example, a Gaussian wake profile is used, with a width of 0.15 and a defect of 2. The upstream blade row has 48 blades and is stationary (0.0 rads/s), as it represents a stator.

1. Expand the **Regions > Turbo Blade Row 1 > Boundaries > Blade Inlet Blade Row 1 > Physics Conditions > Non-Reflecting Option** node and activate *Non-Reflecting Option*.



An additional node, **Wake Specification**, is added to the **Physics Conditions** manager for this boundary when the Non-Reflecting option is activated.

2. Select the **Wake Specification** node and set *Method* to **Gaussian**.



In response to these changes, there are now two extra nodes in the **Physics Values** folder that must be set.

3. Select the **Physics Values** > **Gaussian Wake** node,



set *Rotation Rate* to `0.0 radian/s`, and set *Modes* to `3`.



4. Select the **Physics Values > Gaussian Wake > Blades Per Pitch** node and set *Number of Blades* to `48`.



5. Select the **Physics Values > Non-Reflecting Mode Specification** node and set *Number of modes* to `9`.



Similarly, modify the outlet boundary without the wake specification.

6. Select **Regions > Turbo Blade Row 1 > Boundaries > Blade Outlet Blade Row1 > Physics Conditions > Non-Reflecting Option** and activate *Non-Reflecting Option*.



7. For the same boundary, select the **Physics Values > Non-Reflecting Mode Specification** node and set *Number of modes* to 9



Define the blade row specification:

8. Select the **Regions > Turbo Blade Row 1 > Physics Values > Blade Row Specification** node and set *Blade Row* to **Blade Row 1**.



9. Save the simulation.

Setting Solver and Stopping Criteria Parameters

Set solver parameters:

1. Expand the **Solvers** > **Implicit Harmonic Balance** manager and set *Courant Number* to `2.0`.



To make sure that sufficient steps have been allowed, particularly as the simulation continues from a previously converged solution, modify the stopping criteria:

2. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `6000`.



3. Save the simulation.

Visualizing the Solution

To show the temperature profile on a plane cutting through the center of the mesh, create a scalar scene.

This scene includes multiple transformed sections of the mesh to show the effect of phase lag on the solution.

1. Right-click the **Scenes** manager and select **New Scene** > **Scalar**.

Before proceeding further, define two new transforms that are used to create multiple mesh sections in the scene.

2. Right-click the **Tools** > **Transforms** manager and select **New Graphics Transform** > **Simple Transform**.



3. Rename the **Simple Transform 1** node to `Simple Transform Down`.
4. Select the **Simple Transform Down** node and set *Rotation Angle* to `-10.0 deg`.



5. Right-click the **Transforms** node and select **New Graphics Transform > Simple Transform** as before.
6. Rename this node to `Simple Transform Up`.
7. Select the **Simple Transform Up** node and change the *Rotation Angle* to `10.0 deg`.



Switch back to creating the Scalar Scene.

8. Select the **Scenes** > **Scalar Scene 2** node and click the **scene/plot** button above the object tree.
9. Select the **Scalar Scene 2** > **Displayers** > **Outline 1** node and set the *Opacity* to `0.2`.



10. Select the **Parts** node and edit the *Parts* property. In the *Parts* dialog:
    a) To clear all selections, double-click the **Column Selection** button in the upper left portion of the dialog.



    b) Expand the **Regions** > **Turbo Blade Row 1** manager.
    c) To select **Blade Surface Blade Row 1**, **Hub Surface Blade Row 1**, and **Shroud Surface Blade Row 1**, click once on the selection

button to the left of each of them, as shown in the following screenshot.



d) Click **OK**.

Copy the **Outline 1** node twice: once to provide a transformed outline rotated 10 deg up from the modeled blade, and a second time to provide a transformed outline rotated 10 deg down.

11. Right-click the **Outline 1** node and select **Copy**.
12. Right-click the **Scalar Scene 2** node at the top of the panel and select **Paste**.

This operation pastes the copied node within the **Displayers** folder with the name, **Copy of Outline 1**.

13. To add a second copy of **Outline 1** to the **Displayers** folder, repeat this copy-and-paste operation.

    The second copy is named **Copy of Outline 1 2**.



14. Rename the **Copy of Outline 1** node to `Outline Up`.
15. Select the **Outline Up** node and set *Transform* to **Simple Transform Up**.



16. Rename the **Copy of Outline 1 2** node to `Outline Down`.
17. Select the **Outline Down** node and set the *Transform* to **Simple Transform Down**.
18. Click the (**Save-Restore-Select views**) button in the toolbar and select **Views** > **-X+Y+Z** > **Up +Y**.



The scalar displayers within the scene are set up in a similar manner to the outline displayers.

19. Select the **Scalar 1** > **Parts** manager and edit the *Parts* property.
20. In the *Parts* dialog, expand the **Derived Parts** manager and select **Blade2BladeView**.



21. Click **OK**.
22. To make sure that the Harmonic Balance Functions are activated, click the (**Initialize Solution**) button.
23. Select the **Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Harmonic Balance > Temperature > Time Level > 006 > Temperature at time-level 6**.



24. Also for the **Scalar 1 > Scalar Field** node, set the following properties:

Property	Setting
Min	`250.0`

Property	Setting
Max	`350.0`
Clip	**Off**
Auto Range	**Off**

The final set of properties is as shown below.



Copy the **Scalar 1** node twice for either side of the central position:

25. Right-click the **Scalar 1** node and choose **Copy**.
26. Right-click the **Scalar Scene 2** node (at the top of the subtree) and choose **Paste**.

A copy of the **Scalar 1** node is created with the name **Copy of Scalar 1**.

27. To add a second copy of **Scalar 1** to the **Displayers** folder, repeat this copy-and-paste operation.

    The second copy is named **Copy of Scalar 1 2**.



28. Rename the **Copy of Scalar 1** node to `Scalar Up`.
29. Select the **Scalar Up** node and set *Transform* to **Simple Transform Up**.



30. Expand the **Scalar Up** node, select the **Scalar Field** node, and set *Function* to **Harmonic Balance > Temperature > Time-Level > 006 > Temperature Level 6 Phase Lagged Up**.
31. Also set the same values for the *Min* and *Max* properties that were used previously, that is `250.0` and `350.0` respectively. Make sure that both the *Auto Range* and *Clip* properties are set to **Off**.
32. Select the **Scalar Up > Color Bar** node and deactivate *Visible*.

The color bar is removed from the scene for this particular displayer.

A similar process is required for the final scalar displayer.

33. Rename the **Copy of Scalar 1 2** node to `Scalar Down`.
34. Select the **Scalar Down** node and set *Transform* to **Simple Transform Down**.
35. Expand the **Scalar Down** node, select the **Scalar Field** node, and set *Function* to **Harmonic Balance > Temperature > Time Level > 006 > Temperature Level 6 Phase Lagged Down**.
36. Also set the *Min* and *Max* properties to `250.0` and `350.0` respectively. Make sure that both the *Auto Range* and *Clip* properties are set to **Off**.
37. Select the **Scalar Down > Color Bar** node and deactivate *Visible*.

The steps for **Scalar Scene 2** are complete.

38. To switch back to the main object tree, click the **simulation** button.
39. Save the simulation.

Running the Simulation

Click **(Run)** to start the solver.

The analysis now starts from the previous solution that was obtained in the Single Row tutorial. The residual values in the Residuals plot immediately increase due to the change in the model properties. However, these residuals decrease at a faster rate in the continued solution than they did in the original Single Row tutorial.

Note: The run time for the Harmonic Balance method is longer than it was for the original tutorial. This is because the Harmonic Balance method obtains solutions for several time levels simultaneously.

Analyzing Results

Examine temperature and pressure on the blade surface.

Temperature

1. The temperature at time level 6, shown in **Scalar Scene 2**, is as follows:



The effect of the wake specification at the inlet boundary can be clearly observed on the temperature plot, along with the effect of the phase angle.

Obtaining XY Plot of Pressure on the Blade Surface

To plot the pressure profile at different time levels on the surface of the blade, it is first necessary to create a derived part that defines the location of the profile.

2. Right-click **Scenes > Mesh Scene 1** and select **Open**.
3. Right-click the **Derived Parts** manager and select **New Part > Isosurface**.

   The *Create Isosurface* dialog opens.
4. Click the menu in **Input Parts** box.

   The *Select Objects* dialog appears.
5. To clear all selections, double-click the **Column Selection** button in the upper left portion of the dialog.



6. Expand the **Regions** > **Turbo Blade Row 1** manager.
7. Select **Blade Surface Blade Row 1**.



8. Click **OK**.
9. Click in the **Scalar** box and select **Block-Mapped 1: Vmin to Vmax**.
10. Enter `0.5` in the *iso value* text box.
11. Select **No Displayer** in the **Display** section.

    The final panel is as shown below.



12. Click **Create** and then **Close**.
13. Rename the **Derived Parts > isosurface** node to `Blade 50% Section`.

14. Click (**Hide All Meshes**).
15. Click **(Make Scene Transparent)**.

The derived part is shown in the following image for illustration.



This derived part is now used in an XY plot.

16. Right-click the **Plots** manager and select **New Plot > X-Y Plot**.
    A new node, **XY Plot 1**, appears in the tree below.
17. Rename the **XY Plot 1** node to **Blade Pressure Profile**.
18. Select the **Blade Pressure Profile** node and edit the *Parts* property.
19. In the *Blade Pressure Profile* dialog, expand the **Derived Parts** manager and select **Blade 50% Section**.



20. Click **OK**.
21. Select the **Plots > Blade Pressure Profile > Y Types > Y Type 1 > Scalar Function** node and set *Field Function* to **Harmonic Balance > Pressure > Time-level > 000 > Pressure at time-level 0**.



The profile plot appears as shown in the **Graphics** window.



Summary

This tutorial demonstrated the following features:

- Resuming a previously converged simulation and converting it to run with the Harmonic Balance method
- Setting the necessary boundary conditions including selection of an unsteady wake
- Preparing a scalar scene with transformed parts and phase-shifted scalars
- Running a Harmonic Balance simulation
- Obtaining an XY plot of pressure on a specific section of the blade

Heat Transfer and Radiation

The tutorials in this set illustrate various STAR-CCM+ features for heat transfer, radiation, and thermal comfort.

Contents

Conjugate Heat Transfer: Heated Fin Introduction
Multi-Part Solid: Graphics Card Cooling Introduction
Natural Convection: Concentric Cylinders Introduction
Dual Stream Heat Exchanger: Car Radiator
Surface-to-Surface Radiation: Thermal Insulator
Multiband Surface-to-Surface Radiation: Solar Collector
Thermal Comfort Wizard: Single Occupant in a Cabin

Conjugate Heat Transfer: Heated Fin Introduction

This tutorial incorporates both multi-region conjugate heat transfer and natural convection in a STAR-CCM+ simulation.

The starting point for this tutorial is a one-cell-thick hexahedral grid representing a solid fin within a rectangular enclosure containing air, as shown in the Figure below. Built with pro-STAR, the grid includes a non-conformal grid interface between the fin and surrounding air. The file containing grid data also includes information about the boundary types present in the problem.



The three-dimensional grid is converted to a two-dimensional one and also scaled to the appropriate size before being used for the analysis. All conditions and properties are defined using the STAR-CCM+ GUI and solved using the coupled solver. The results are displayed using the post-processing tools available in STAR-CCM+.

The analysis includes the following heat transfer modes:

- Conduction in the solid fin.
- Diffusion and natural convection in the surrounding fluid.

Radiation effects are not included as part of this analysis.

Contents

Prerequisites

Importing the Mesh and Naming the Simulation
Converting to a Two-Dimensional Mesh
Renaming Regions and Boundaries
Scaling the Mesh
Visualizing the Interior Two-Dimensional Mesh
Setting Up the Models
Setting Material Properties
Setting Initial Conditions and Reference Values
Creating Interfaces
Setting Boundary Conditions and Values
Visualizing the Solution
Reporting and Monitoring with Plots
Initializing and Running the Simulation
Visualizing the Results
Adding Streamlines
Summary

Prerequisites

The instructions in the Conjugate Heat Transfer tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

Use the following steps:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.

4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/heatTransferAndRadiation/data`.
5. Select `fin.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window. Two mesh regions named **AIR** and **AL** are created in the **Regions** node representing the grid domain. A geometry scene is created in the *Graphics* window.
6. Save the simulation as `heatedFin.sim`.

Converting to a Two-Dimensional Mesh

For greater efficiency, convert the mesh region from a one-cell-thick three-dimensional grid to a two-dimensional one.

There are special requirements in the STAR-CCM+ product for three-dimensional meshes that are converted to two-dimensional ones. These requirements are:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The mesh imported for this tutorial was built with these requirements in mind. Were the grid not to conform to the above conditions, it would have been necessary to realign the region using the transformation and rotation facilities in STAR-CCM+.

1. Select **Mesh > Convert To 2D**.
2. In the *Convert Regions To 2D* dialog that appears, make sure that:
   - All regions are selected.
   - The **Delete 3D Regions After Conversion** option is activated.



3. Click **OK**.

   Once you click **OK**, the mesh conversion takes place. The new two-dimensional mesh is shown, viewed from the z-direction, in the *Geometry Scene 1* display which was created after you imported the mesh. The mouse rotation option is suppressed for two-dimensional scenes.
4. Right-click the **Continua** > **Physics 1** continuum node and select **Delete**.
5. Click **Yes** in the confirmation dialog.
6. Repeat these steps for the **Physics 2** continuum node.

Renaming Regions and Boundaries

Change the default region and boundary names to more suitable names.

Use the following steps:

1. Rename **AIR 2D** to `Fluid`.
2. Rename the boundaries within the **Fluid** node:

Original Name	New Name
wall	`Air-Interface`
wall 2	`Vertical-Walls`
wall 3	`Horizontal-Walls`

3. Rename **AL 2D** to `Solid`.
4. Rename the boundaries within the **Solid** node:

Original Name	New Name
Default_Boundary_Region	`Fin-Interface`
wall	`Fin-Bottom`

This portion of the object tree appears as shown below.



Scaling the Mesh

The imported grid uses units that differ from those in STAR-CCM+. To compensate, scale the mesh.

The original grid was built using centimeters as a length unit, but no scaling factor was applied when the grid was saved to the `ccm` file. The current dimensions of the mesh regions are already available from the mesh import process and are displayed in the *Output* window.



The output shows that the mesh regions extend from 0 to 10.0 in the x-direction, 0 to 10.0 in the y-direction and 0.0 to 1.0 in the (now defunct) z-direction. Without any further action, STAR-CCM+ interprets these values as specified in meters. Therefore, scale the region by a factor of 0.01 so that the dimensions are correctly specified.

1. To scale the region, select **Mesh** > **Scale Mesh** from the menu bar.
2. In the *Scale Mesh* dialog, set the **Scale Factor** to `0.01` for the fluid and solid regions.
3. Click **Apply**.

   The mesh region reduces in size.
4. Click **Close**.
5. To restore the previous viewing distance for the scaled domain in the display, click (Reset View) in the toolbar.
6. To verify that scaling has in fact been applied:
   a) Select **Mesh** > **Diagnostics** from the menu bar.
   b) In the *Mesh Diagnostics* dialog, click **OK** for all regions.
7. Check and review the output values.

Visualizing the Interior Two-Dimensional Mesh

This part of the tutorial involves examining the interior of the two-dimensional mesh of the heated fin.

Use the following steps:

1. Select the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** node.
2. Activate the checkbox of the **Mesh** property.

   This shows the mesh in the solution domain interior.

3. To get a clearer view of individual mesh regions, highlight them. For example, select the **Regions** > **Fluid** node.



Setting Up the Models

Models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution. In this example, the flow is turbulent and compressible.

The Coupled Flow model is used together with the default K-Epsilon turbulence model. As there are two materials (air and aluminum) present, two regions are required for the analysis. Define two sets of continuum models, each appropriate to a different region. By default, **Physics 1 2D** and **Physics 2 2D** are automatically created when the mesh is converted to two-dimensional. Use these continua for the fluid and the solid.

1. Rename the continua:

Original Name	New Name
Physics 1 2D	`Air`
Physics 2 2D	`Aluminum Alloy`

2. For the physics continuum, **Continua** > **Air**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically) **Two-Layer All y+ Wall Treatment** (Selected automatically)
Optional Models	**Gravity**

3. Click **Close**.

   The color of the **Air** node turns from gray to blue to indicate that models have been activated.

4. To review the models, open the **Air** > **Models** node.



5. For the continuum of the solid region, **Continua** > **Aluminum Alloy**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Material	**Solid**
Optional Models	**Coupled Solid Energy**
Equation of State	**Constant Density**
Time	**Steady**

6. Click **Close**.

   The color of the **Aluminum Alloy** node turns from gray to shaded gray to indicate that models have been activated.

7. To review the models, open the **Aluminum Alloy** > **Models** node.



8. Save the simulation.

Setting Material Properties

This part of the tutorial involves changing the solid material properties for the aluminum alloy fin.

The fin has a density of 2800 kg/m^3, a specific heat of 880 J/kg-K and a thermal conductivity of 180 W/m-K. The surrounding air is considered compressible based on the Ideal Gas Law. By default it has a variable density, a molecular weight of 28.9664 kg/kg.mol, a specific heat of 1003.62 J/kg-K, a thermal conductivity of 0.0260305 W/m-K, and a molecular viscosity of 1.85508E-5 Pa-s.

1. Navigate to the **Continua** > **Aluminum Alloy** > **Models** > **Solid** > **Al** node.
2. Edit the **Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`2800 kg/m^3`
Specific Heat > Constant	**Value**	`880 J/kg-K`
Thermal Conductivity > Constant	**Value**	`180 W/m-K`

3. Save the simulation.

Setting Initial Conditions and Reference Values

This part of the tutorial involves setting the direction and magnitude of the gravity vector and the reference density for the simulation.

Start with the acceleration force in the negative y-direction.

1. Edit the **Continua** > **Air** > **Reference Values** node and set the following properties:

Node	Property	Setting
Gravity	**Value**	`[0.0, -9.81, 0.0] m/s^2`
Reference Density	**Value**	`1.27588 kg/m^3`

2. Edit the **Continua** > **Air** > **Initial Conditions** > **Static Temperature** > **Constant** node and set **Value** to `293 K`.
3. Edit the **Continua** > **Aluminum Alloy** > **Initial Conditions** > **Static Temperature** > **Constant** node and set **Value** to `343 K`.

   This temperature is close to the final expected temperature of the solid region.
4. Save the simulation.

Creating Interfaces

STAR-CCM+ requires the existence of an interface between all regions to transfer the appropriate mass and energy quantities during the calculation.

For this example, create a contact-type interface between the fluid and solid regions. This procedure involves creating a so-called “in-place interface” between the wall boundaries enclosing each of the two regions, and then defining this interface as being of an appropriate “contact” type for fluid-solid heat transfer.

1. Using the multi-select method, select the following nodes:
   - **Regions** > **Fluid** > **Boundaries** > **Air-Interface**
   - **Regions** > **Solid** > **Boundaries** > **Fin-Interface**
2. Create an in-place interface.

   This step creates the following new nodes:
   - **Fluid** > **Boundaries** > **Air-Interface [Interface 1]**
   - **Solid** > **Boundaries** > **Fin-Interface [Interface 1]**

When you initialize the flow, these boundaries form the connection between the fluid and solid regions, superseding the original wall boundaries.

In addition to the new in-place boundaries, a folder node named **Interfaces** appears in the object tree. It contains an interface node called **Interface 1**.

3. Edit the **Interfaces** > **Interface 1** node and set the following properties:

Property	Setting
Type	`Contact Interface`
Topology	`In-place`

4. Save the simulation.

Setting Boundary Conditions and Values

The heated fin consists of an inverted "T" shape placed in a closed box. The only boundary specifications required for this analysis are wall thermal conditions.

Apply a fixed temperature of 343 K to the base of the solid fin to provide the heat source. Set the vertical sides of the fluid-containing box to a fixed temperature of 293 K. All other wall boundary conditions are considered to be adiabatic. Use a contact-type conducting boundary at the solid-fluid interface.

Start with the wall boundary definitions for the fluid region.

1. Navigate to the **Regions** > **Fluid** > **Boundaries** > **Vertical-Walls** > **Physics Conditions** node.
2. Select the **Thermal Specification** node and set **Condition** to `Temperature`.

   The temperature value for this boundary can now be set.
3. Select the **Vertical-Walls** > **Physics Values** > **Static Temperature** > **Constant** node and set Value to 293 K.

   The conditions for the Horizontal-Walls boundary (upper and lower walls of the fluid region) can remain at the default adiabatic setting.

   You can now set the solid wall boundary conditions and values in a similar way to those of the fluid.
4. Navigate to the **Regions** > **Solid** > **Boundaries** > **Fin-Bottom** > **Physics Conditions** node.

5. Select the **Thermal Specification** node and set **Condition** to `Temperature`.

   The temperature value for this boundary can now be set.
6. Select the **Fin-Bottom > Physics Values > Static Temperature > Constant** node and set Value to 343 K.
7. Save the simulation.

Visualizing the Solution

This part of the heated fin tutorial involves setting up a scene for displaying scalars and vectors of the solution.

To view the temperature of the fluid and solid regions as the solution develops:

1. Create a scalar scene.

   This scene includes both the fluid and solid regions.
2. Set the scalar value to `Temperature`.

   The velocity vectors for the fluid can also be included in the scene by adding a vector displayer to the scalar scene.
3. Create a vector displayer.
4. Edit the **Scenes > Scalar Scene 1 > Displayers > Vector 1 > Parts** node and select the fluid region.
5. Save the simulation.

Reporting and Monitoring with Plots

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops.

To monitor a quantity, set up a report which defines the quantity of interest and the region parts to monitored. Define a monitor based on the report, which controls the update frequency and normalization characteristics. Finally, create an X-Y graph plot can from the monitor.

Use a report monitor to track the heat flux at the interface of the heated fin and surrounding air:

1. Create a heat transfer report.
2. Edit the **Reports > Heat Transfer 1 > Parts** property and select the **Solid > Boundaries > Fin-Interface [Interface 1]** node.

   The setup for the report is now complete, and a monitor can be made from that report.

3. Create a monitor from the report.
4. Create a plot from the monitor
5. Save the simulation.

Initializing and Running the Simulation

To verify that the initial conditions for the regions are correct before you run the simulation, STAR-CCM+ allows you simply to perform solver initialization and then view the results using an appropriate scene.

Use the following steps:

1. Initialize the simulation.
2. To see the initialization result, look at the scalar scene display.



As the colors show, the solid fin is at 343 K while the surrounding air is at 293 K. The velocity value for each cell is 0. The simulation can now be run.

3. Run the simulation.
4. To watch the solution being updated while the simulation is running, look at the scalar scene display.
5. When the simulation has finished running, save it.

Visualizing the Results

After the solution finishes, you can examine results in plots and scenes.

Click the tab of each display that you set up before the run.

1. To view the heat transfer variation at the solid-fluid interface, activate the *Heat Transfer 1 Monitor Plot* display.



2. To see the temperature and velocity vector results, click the tab for the *Scalar Scene 1* display.



By default, filled cell values are shown. To make the display smoother:

3. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node and set Contour Style to `Smooth Filled`.

   The contours of the scalar display now appear smooth as seen in the following zoomed-in screenshot.

Filled:

Smooth Filled:



Additionally, the temperature solution of each region can be viewed by selecting the appropriate part.

4. Select the **Displayers** > **Scalar 1** > **Parts** node and deselect the Fluid region.

   The display appears as shown below.



Adding Streamlines

In this part of the heated fin tutorial, use streamlines to depict flow paths.

Streamlines can be added into any of the scalar or geometry scene plots with the creation of a "streamline" derived part.

1. Make the *Geometry Scene 1* display active.
2. Create a new streamline derived part.

   The *Create Streamline* inset panel appears to allow definition of the derived part.



3. To select input parts, click **Select**.
4. In the object selector, select only the Fluid region.
5. In the *Create Streamline* panel, make the following settings:
   a) For the seed mode, select `Line Seed`.
   b) Activate the display tool option, and set the resolution to `15`.
   The streamline line seed tool appears in the display.
6. Work with the tool in the scene as follows:
   a) Using the mouse buttons, zoom-in and position the part as shown below.
   b) Click the sphere at the end of the tool.
   c) Adjust the line so that it is approximately half a cell above the top of the solid fin. To adjust the line, simply grab the sphere and drag the mouse to its new position.



7. Click **Create** and then **Close**.
8. Click (**Reset View**) in the toolbar.
   The streamlines are generated and shown in the display.
9. Position the part as shown below.



The streamline part can now be added to any suitable scene or be displayed in its own scene if desired.

10. Create a scalar scene.
11. Set the scalar value to `Temperature`.
12. By default, the regions Fluid and Solid are included in the scene. To display the streamlines only:
    a) Edit the **Scenes** > **Scalar Scene 2** > **Displayers** > **Scalar 1** > **Parts** node.
    b) Deselect the regions and select the derived part.
13. Edit the **Scenes** > **Scalar Scene 2** > **Displayers** > **Scalar 1** > **Parts** > **Scalar Field** node:
    a) Set **Min** to `300`.
    b) Set **Max** to `343`.

    Streamlines colored according to the local temperature appear in the *Scalar Scene 2* display.



14. Save the simulation.

Summary

The Conjugate Heat Transfer tutorial used multi-region conjugate heat transfer and natural convection to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Scaling the mesh.
- Plotting geometry and scalar scenes.
- Defining models for multiple continua.
- Defining material properties required for the selected models.
- Initializing and running the solver until the residual magnitudes are sufficiently small.
- Analyzing results using the visualization and monitor plot facilities.
- Creating streamlines.

Multi-Part Solid: Graphics Card Cooling Introduction

This tutorial demonstrates how you can model multiple solid parts in a single continuum using the Multi-Part Solid physics model.

The geometry that is used consists of a simplified graphics card inside a PC tower case.



Index	Description
1	PC Tower Case
2	Inlet 1
3	Inlet 2
4	Graphics Card
5	Outlet 2
6	Outlet 1

The graphics card and its part materials are shown in the following diagram.



Index	Description	Material
1	Heat Sink	Aluminum
2	Capacitor	Alumina
3	Memory Chip	Silicon
4	GPU Chip Casing	Aluminum
5	Printed Circuit Board	FR-4 + Copper
6	Connector Port Block	Aluminum
7	Back Plate	Aluminum

Import the surface mesh of the geometry into STAR-CCM+ as individual parts, before creating the physics continua and regions, to use the Multi-Part Solid model. A polyhedral mesh containing prism layers is used. The embedded thin mesher is used to mesh the thin parts in the geometry.

Contents

Prerequisites
Importing the Geometry

Assigning Parts to Regions
Setting the Boundary Types
Selecting Physics Models
Assigning Continua to Regions
Assigning Materials to the Multi-Part Solid Parts
Setting Up the Mesh Models
Creating Volume Reports for the Heat-Generating Parts
Setting Volumetric Heat Sources
Setting Up Reference Values and Initial Conditions
Modifying the Interface Type
Preparing Scenes for Visualization
Visualizing Assigned Volumetric Heat Sources
Reporting Temperature with a Monitor and a Plot
Modifying Solver Settings and Stopping Criteria
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Multi-Part Solid tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Geometry

The meshing process starts from a CAD geometry in Parasolid format.

Use the following steps:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh**.

4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/heatTransferAndRadiation/data`.
5. Select `graphicsCardAndEnclosure.x_t` then click **Open**.
6. To accept the default *Import Surface Options*, click **OK**.

Note: If you wish to use the Multi-Part Solid model, import all CAD geometry as parts.

A total of 22 parts are imported. Part contacts between coincident parts are automatically generated during the import process.

7. Save the simulation as `graphicsCardCooling.sim`.

Assigning Parts to Regions

Three regions are required for this simulation: one for the fluid and two for the solids.

Anisotropic materials, such as the Printed Circuit Board (PCB), require dedicated physics continua and regions, hence the need for two solid regions in this tutorial. The remaining solid parts are contained in one solid region.

To create the regions:

1. Right-click the **Geometry** > **Parts** > **Air** node and select **Assign Parts to Regions**.
2. In the *Assign Parts to Regions* dialog, select the **Create a Boundary for Each Part Surface** option.

   Selecting this option means that STAR-CCM+ creates one region for the selected part and separate boundaries for each named surface.
3. Click **Apply**.

   A region with the default name **Region** is created.

Create a region for all the solid parts except the PCB:

4. Select all the parts in the dialog.
5. Deselect the **Air** and **PCB** parts.
6. Click **Apply**.

   Another region, **Region 2**, is created.

Create the last region for the PCB:

7. Deselect all the parts in the dialog.
8. Select the **PCB** part.
9. Click **Apply** and then **Close**.

A third region, **Region 3**, is created.

10. Rename the new regions as follows:

Original Name	New Name
Region	`Air`
Region 2	`GPU Components`
Region 3	`PCB`

The following screenshot shows the new regions.



Setting the Boundary Types

Specify boundary types for the graphics card.

Define the boundary types for the inlet and outlet boundaries:

1. Edit the **Regions** > **Air** > **Boundaries** node and set the following properties:

Node	Property	Setting
Air.Inlet 1	**Type**	`Velocity Inlet`
Air.Inlet 2	**Type**	`Velocity Inlet`
Air.Outlet 1	**Type**	`Pressure Outlet`
Air.Outlet 2	**Type**	`Pressure Outlet`

The boundaries are shown in the following image.



Selecting Physics Models

This tutorial requires three physics continua, one for the air, one for the anisotropic material (PCB) and one for the remaining solid parts on the graphics card.

Contents

Creating the Physics Continuum for the GPU Components

Defining Material Properties of GPU Components
Creating the Physics Continuum for the Air Region
Defining Material Properties of Air
Creating the Physics Continuum for the PCB Region
Defining Material Properties of PCB

Creating the Physics Continuum for the GPU Components

For the graphics card cooling tutorial, begin configuring physics models by setting up the continuum for the GPU components.

The first physics continuum that you create is used to model all the solid components mounted on the PCB:

1. Create a physics continuum.
2. Rename the continuum to `GPU Components`.
3. For the physics continuum, **Continua** > **GPU Components**, select the following models in order:

Group Box	Model
Space	**Three Dimensional Gradients** (Selected automatically)
Time	**Steady**
Material	**Multi-Part Solid**
Optional Models	**Segregated Solid Energy**
Equation of State	**Constant Density**

4. Click **Close**.
5. To review the models, open the **GPU Components** > **Models** node.



Defining Material Properties of GPU Components

Multiple solid parts are modeled in the GPU Components continuum. Some parts are made of different materials.

To define these materials in the multi-part solid model:

1. Navigate to the **Continua** > **GPU Components** > **Models** > **Multi-Part Solid** node.
2. Right-click the **Solids** node and choose **Select Mixture Components**.



3. In the *Select Mixture Components* dialog, expand the **Material Databases** > **Standard** > **Solids** node.
4. Select the three materials from the list.

   The materials that are chosen here are not important -- the material properties for these selections are changed in the subsequent steps.



5. Click **Apply** and then **Close**.
6. Rename the materials to `Aluminum`, `Alumina`, and `Silicon`.



7. Edit the **Multi-Part Solid** node and set the material properties as follows:

Node	Property	Setting
Aluminum		
Material Properties		
Density (kg/m^3)		
Constant	**Value**	2700
Specific Heat (J/kg-k)		
Constant	**Value**	896
Thermal Conductivity (W/m-K)		
Constant	**Value**	167

Node	Property	Setting
Alumina		
Material Properties		
Density (kg/m^3)		
Constant	**Value**	3960
Specific Heat (J/kg-k)		
Constant	**Value**	850
Thermal Conductivity (W/m-K)		
Constant	**Value**	30

Node	Property	Setting
Silicon		
▮▮ **Material Properties**		
▮▮▮▮ **Density** (kg/m^3)		
▮▮▮▮▮▮ **Constant**	**Value**	`2330`
▮▮▮▮ **Specific Heat** (J/kg-k)		
▮▮▮▮▮▮ **Constant**	**Value**	`700`
▮▮▮▮ **Thermal Conductivity** (W/m-K)		
▮▮▮▮▮▮ **Constant**	**Value**	`124`

8. Click **Close**.

Creating the Physics Continuum for the Air Region

Create the physics continuum that is used to model the air in the fluid domain.

Constant density is used with the Boussinesq model. The mesh contains four prism layers; the Realizable K-Epsilon Two-Layer turbulence model is used to complement the mesh.

1. Create a physics continuum.
2. Rename the continuum to `Air`.
3. For the physics continuum, **Continua** > **Air**, select the following models in order:

Group Box	Model
Space	**Three Dimensional Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**

Group Box	Model
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically) **Two-Layer All y+ Wall Treatment** (Selected automatically)
Optional Models	**Gravity** **Segregated Fluid Temperature** **Boussinesq Model**

4. Click **Close**.
5. To review the models, open the **Air** > **Models** node.



Defining Material Properties of Air

The Boussinesq model requires you to set the thermal coefficient of expansion for the fluid region.

To define this property:

1. Expand the **Continua** > **Air** > **Models** > **Gas** > **Air** > **Material Properties** > **Thermal Expansion Coefficient** node.
2. Select the **Constant** node and set **Value** to `3.33E-3 /K`.

Creating the Physics Continuum for the PCB Region

This continuum involves working with variations in thermal conductivity.

The PCB is constructed from layers of FR-4 glass-reinforced laminate and copper foil. As a result, the PCB has a different thermal conductivity within the same layer and a different thermal conductivity between the various layers.

1. Create a physics continuum.
2. Rename the continuum to `PCB`.
3. For the physics continuum, **Continua** > **PCB**, select the following models in order:

Group Box	Model
Space	**Three Dimensional Gradients** (Selected automatically)
Time	**Steady**
Material	**Solid**
Optional Models	**Segregated Solid Energy**
Equation of State	**Constant Density**

4. Click **Close**.
5. To review the models, open the **PCB** > **Models** node.



Defining Material Properties of PCB

Rename the Solid node and set its properties.

Edit the material properties of the PCB:

1. Rename the material in the **Continua** > **PCB** > **Models** > **Solid** node to `FR-4 + Copper`.
2. Edit the **FR-4 + Copper** > **Material Properties** node and set the material properties as follows:

Node	Property	Setting
Density (kg/m^3)		
Constant	**Value**	`33`
Specific Heat (J/kg-k)		
Constant	**Value**	`37`
Thermal Conductivity	**Method**	`Anisotropic`

The anisotropic properties of the PCB are defined in a later section.

Assigning Continua to Regions

Since this simulation has multiple regions and continua, it is necessary to associate each region with a particular continuum explicitly.

1. Assign the physics continua to the corresponding regions:

Region Node	Physics Continuum Property Setting
Air	`Air`
GPU Component	`GPU Component`
PCB	`PCB`

Assigning Materials to the Multi-Part Solid Parts

The **GPU Components** region contains multiple solid parts. You need to match these parts with their correct material.

The materials were defined in the **Multi-Part Solid** physics continuum. Assign the parts to materials as follows:

1. Edit the **Regions** > **GPU Components** > **Physics Values** > **Material Part Groups** node.
2. Click (**Custom Editor**) in the **Parts** property of the **Aluminum** node.
3. In the *Aluminum* dialog, select the parts that are made of Aluminum:
   - **GPU Chip**
   - **Heatsink Block**
   - **Plate**
   - **Port Block 1**
   - **Port Block 2**
   - **Port Block 3**

   The following screenshot shows the completed dialog.



4. Click **OK**.
5. Assign the other parts to the remaining materials:
   - All the capacitors, **Capacitor 1** to **Capacitor 6**, to the **Alumina** node
   - All the memory chips, **Memory Chip 1** to **Memory Chip 8**, to the **Silicon** node

   Parts can only be assigned to one particular material. Parts that are already assigned to a material are removed from the list of parts of that material when you assign them to a new material.
6. Save the simulation.

Setting Up the Mesh Models

Generate a volume mesh using the polyhedral mesher.

As the graphics card is built of thin parts, the embedded thin mesher model can also be used.

1. Create a mesh continuum.
2. For the mesh continuum, **Continua** > **Mesh 1**, select the following models in order:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Embedded Thin Mesher**

3. Click **Close**.
4. To review the models, open the **Mesh 1** > **Models** node.



Modify the mesh model parameters:

5. Select the **Mesh 1** > **Models** > **Embedded Thin Mesher** node and activate the **Customize Thickness Threshold** property.

   Enabling this option allows you to set a custom threshold value for the embedded thin mesher. Any part that is thinner than this threshold is meshed using the embedded thin mesher.

Edit the mesh reference values:

6. Select the **Mesh 1** > **Reference Values** > **Base Size** node and set **Value** to `0.02 m`.
7. Select the **Surface Size** > **Relative Minimum Size** node and set **Percentage of Base** to `10`.
8. Select the **Thin Solid Thickness** node and set **Size type** to `Absolute`.
9. Select the **Thin Solid Thickness** > **Absolute Size** node and set **Value** to `4 mm`.
10. Save the simulation.

Contents

Creating the Volume Shapes and Defining the Volumetric Controls
Generating the Volume Mesh

Creating the Volume Shapes and Defining the Volumetric Controls

Use a volume shape to control the mesh refinement around the graphics card.

To create the volume shape:

1. Right-click the **Tools** > **Volume Shapes** node and select **New Shape** > **Block**.
2. In the *Create Block Volume Shape* in-place dialog, make the following settings in the *Maximum and Minimum Coordinates* box:

Property	Setting
Corner 1	`[-0.4 m, -0.15 m, -0.06 m]`
Corner 2	`[-0.07 m, -0.1 m, 0.07 m]`

A preview of the block shape is visible in the scene.



3. Click **Create** and then **Close**.

   The new shape, **Block 1**, is added to the **Tools** > **Volume Shapes** node.

Continue by creating the volumetric control:

4. Right-click **Continua** > **Mesh 1** > **Volumetric Controls** and select **New**.

   A new node, **Volumetric Control 1**, is added below the **Volumetric Controls** node.
5. Select the **Volumetric Control 1** node.
6. Click ... (**Custom Editor**) in the **Shapes** property.
7. In the *Volumetric Control 1* dialog, select **Block 1**.
8. Click **OK**.
9. Select the **Volumetric Control 1** > **Mesh Conditions** > **Polyhedral Mesher** node and activate the **Customize polyhedral mesher** property.
10. Select the **Volumetric Control 1** > **Mesh Conditions** > **Surface Remesher** node and activate the **Customize surface remesher** property.
11. Select the **Mesh Values** > **Custom Size** > **Relative Size** node and set **Percentage of Base** to `25`.

12. Save the simulation.

Generating the Volume Mesh

Generate and visualize a conformal mesh.

The volume mesh is ready to be generated.

1. Click (**Generate Volume Mesh**) in the **Mesh** toolbar.
   A conformal mesh is generated between all regions.
2. To visualize the mesh, create a mesh scene and inspect the fluid domain mesh.



3. Edit the **Mesh Scene 1** > **Displayers** > **Mesh 1** > **Parts** node.
4. Deselect the **Regions** > **Air** region.
   The mesh on the graphics card is shown in the following screenshot.



5. Save the simulation.

Creating Volume Reports for the Heat-Generating Parts

Reports are created to calculate the volume of each heat-generating part.

One report per part type is required.

To create the volume report:

1. Right-click the **Reports** node and select **New Report** > **Sum** report.
2. Rename the report to `GPU Volume`.
3. Select the **Reports** > **GPU Volume** node and set *Scalar Field Function* to `Volume`.
4. Edit the *Parts* property of the **GPU Volume** node.
5. In the *GPU Volume* dialog, select the **Parts** > **GPU Chip** node.

Note: Leave the **GPU Chip** > **Surfaces** > **Faces** node unselected.

6. Click **OK**.
7. Run the report.
   The volume of the GPU chip is `3.2e-06`.

Create another report for one of the memory chips:

8. Copy and paste the **GPU Volume** report.
9. Rename the copy to `Memory Volume`.
10. Edit the *Parts* property of the **Memory Volume** node.
11. In the *Memory Volume* dialog:
    a) Deselect the **Parts > GPU Chip** node.
    b) Select the **Memory Chip 1** node.
12. Click **OK**.
    The remaining memory chips have identical volumes, therefore the same volume report is used with each memory chip.

Setting Volumetric Heat Sources

Heat generated by the electronic components is represented in STAR-CCM+ using volumetric heat sources. Heat sources are set on regions.

You can get the volumetric heat generation for a component from the ratio of its total heat output to its volume. Use the volume reports created earlier in field functions that calculate the ratio.

The volume reports created in the previous section are used in a field function to calculate the volumetric heat source of the heat-generating parts, based on a provided heat source. The GPU chip emits 25 W of heat. Each memory chip emits 4 W of heat. Parts are identified using their *Part Index* properties. The appropriate volume report is selected using its field function name.

1. To check the Part Index number of a part, select its node and look at its *Index* property.
   For example, in the following screenshot, the Part Index of the **Geometry > Parts > GPU Chip** node is `13`.



The Part Index numbers for the memory chips are displayed in the following table:

Part Name	Part Index
Memory Chip 1	12
Memory Chip 2	14
Memory Chip 3	15
Memory Chip 4	16
Memory Chip 5	17
Memory Chip 6	18
Memory Chip 7	19
Memory Chip 8	20

To create the custom field function:

2. Right-click **Tools** > **Field Functions** and select **New** > **Scalar**.

3. Rename the new field function to `Volumetric Heat Source`.
4. Select the **Volumetric Heat Source** node and set the following properties:

Property	**Setting**
Function Name	`VolumetricHeatSource`
Dimensions	**Power/Volume**
Definition	`($CellGeometryPartIndex == 13) ? (25/$GPUVolumeReport) : (($CellGeometryPartIndex == 12) \|\| ($CellGeometryPartIndex == 14) || ($CellGeometryPartIndex == 15) \|\| ($CellGeometryPartIndex == 16) || ($CellGeometryPartIndex == 17) \|\| ($CellGeometryPartIndex == 18) || ($CellGeometryPartIndex == 19) \|\| ($CellGeometryPartIndex == 20) ? (4/$MemoryVolumeReport) : 0)`

The field function units are defined as `W/m^3`.

5. Save the simulation.

Contents

Setting Up Physics Values and Conditions

Setting Up Physics Values and Conditions

This part of the tutorial involves setting physics values and conditions for the regions.

Define the anisotropic thermal conductivity of the PCB region:

1. Expand the **Regions** > **PCB** > **Physics Values** > **Anisotropic Thermal Conductivity** node.
2. Edit the **Principal Tensor** node and set the properties as follows:

Node	Property	Setting
XX Component		
Constant	**Value**	`10 W/m-k`
YY Component		
Constant	**Value**	`0.5 W/m-k`
ZZ Component		
Constant	**Value**	`10 W/m-k`

Set the GPU Component region to use the custom field function you created in the previous section as the energy source:

3. Select the **Regions** > **GPU Components** > **Physics Conditions** > **Energy Source Option** node and set *Energy Source Option* to `Volumetric Heat Source`.
4. Select the **Physics Values** > **Volumetric Heat Source** node and set *Method* to `Field Function`.
5. Select the **Volumetric Heat Source** > **Field Function** node and set *Scalar Function* to `Volumetric Heat Source`.

Setting Up Reference Values and Initial Conditions

Gravity and velocity need to be set.

Adjust the reference gravity and the initial velocity as follows:

1. Select the **Continua** > **Air** > **Reference Values** > **Gravity** node and set *Value* to `[0.0, -9.81, 0.0] m/s^2`.
2. Select the **Air** > **Initial Conditions** > **Velocity** > **Constant** node and set *Value* to `[-1, 0.0, 0.0] m/s`.

Modifying the Interface Type

Contact interfaces are required to have the correct thermal conductivity on either side of the interface.

Make all interfaces in the geometry Contact Interfaces:

1. Expand the **Interfaces** node.

2. Select all the interfaces and set *Type* to `Contact Interface`.

Preparing Scenes for Visualization

Create two scalar scenes: one to display the velocity magnitude on a plane cutting through the tower case, and the other to display the temperature on the graphics card.

Begin by creating the section plane:

1. Make the *Geometry Scene 1* display active.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.
3. In the *Create Section* inset panel, make the following settings:
   a) To select input parts, click **Select**.
   b) In the *Select Objects* dialog, expand the **Regions** node and select all three region nodes.
   c) Set *origin* to `[0 m, 0 m, 0.035 m]`.
   d) Set *normal* to `[0 m, 0m, 1 m]`.
   e) Set *Display* to `No Displayer`.

   The following screenshot shows the completed dialog.



4. Click **Create** and then **Close**.

Set up the scalar scene:

5. Create a scalar scene.
6. Edit the **Displayers** > **Scalar 1** > **Parts** node.

7. In the *Edit* dialog, select the **Derived Parts** > **plane section** node.
8. In the object tree, select the **Scalar 1** > **Scalar Field** node.
9. Set *Function* to `Velocity > Magnitude`.
10. Select the **Scalar 1** node and set *Contour Style* to `Smooth Filled`.

Set up the second scalar scene for displaying the temperature on the graphics card:

11. Create a scalar scene.
12. Edit the **Displayers** > **Scalar 1** > **Parts** node.
13. In the *Edit* dialog, expand the **Regions** node and select the **GPU Components** and **PCB** nodes.
14. Select the **Scalar 1** > **Scalar Field** node.
15. Set *Function* to `Temperature`.
16. Set *Units* to `C`.

Visualizing Assigned Volumetric Heat Sources

Visualize the volumetric heat source.

The volumetric heat source for each heat-generating part can be visualized in a scalar scene. This information is contained within the volumetric heat source field function that you created earlier. Displaying this field function directly in a scalar scene, however, would display the average value from both sides of the boundary interfaces. To display the cell value, use a threshold derived part.

Set up a scalar scene with a threshold derived part:

1. Create a scalar scene.
2. Create a threshold derived part.
3. In the *Create Threshold* inset panel, make the following settings:
   a) To select input parts, click **Select**.
   b) In the *Select Objects* dialog, deselect all objects.
   c) Expand the **Parts** node and select all parts except for the **Air** node and the **Heatsink Block** node. Select only the parts, not the surfaces.

      The completed dialog is shown in the following screenshot.



d) In the *Create Threshold* inset panel, set *Scalar* to `Volumetric Heat Source`.

e) In *Scalar Range*, click **Query**.

This action populates the **Min** and **Max** range with the minimum and maximum values in the selected field function.

f) Set *Display* to `Existing Displayer` and `Scalar 1`.

The following screenshot shows the completed dialog.



4. Click **Create** and then **Close**.

Apply the volumetric heat source field function to the scalar scene:

5. Select the **Displayers > Scalar 1 > Scalar Field** node and set *Function* to `Volumetric Heat Source`.
6. Right-click the **Displayers > Outline 1** node and select **Toggle Visibility**.
7. To reposition the geometry in the scene, press the `<S>` key.

   The volumetric heat sources are shown in the following screenshot.



8. Save the simulation.

Reporting Temperature with a Monitor and a Plot

Create a monitor and plot of maximum temperature to help you judge the rate at which the solution is converging.

The maximum temperature in the simulation is plotted. This plot is useful for assessing solution convergence.

To monitor the temperature:

1. Create a **Maximum** report.
2. Rename the new report to `Maximum Temperature`.
3. Select the **Reports** > **Maximum Temperature** node and set *Scalar Field Function* to `Temperature`.
4. Set *Parts* to `[Air, GPU Components, PCB]`.
5. Set *Units* to `C`.
6. Right-click the **Maximum Temperature** node and select **Create Monitor and Plot from Report**.

   New monitor and plot nodes are added to the Monitors and Plots nodes respectively.

Modifying Solver Settings and Stopping Criteria

Modify the under-relaxation factors for the Segregated solvers to help accelerate convergence.

Use the following steps:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Flow		
Velocity	*Under-Relaxation Factor*	`0.85`
Pressure	*Under-Relaxation Factor*	`0.25`
Segregated Energy	*Solid Under-Relaxation Factor*	`0.9999`

Modify the maximum number of iterations for the solver. 300 iterations are sufficient for the solution to converge.

2. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `300`.
3. Save the simulation.

Running the Simulation

The simulation is now ready to run.

Use the following steps:

1. Run the simulation.
2. To watch the solution being updated while the simulation is running, look at the scalar scene display.
3. When the simulation has finished running, save it.

Visualizing the Results

View the results of the solution.

1. After the run is finished, look at the scalar and plot displays.

   The following screenshot illustrates air flow inside the case.



The temperature on the graphics card is shown in the following screenshot.



The temperature monitor plot is displayed in the following screenshot. The solution has converged.



Summary

The Multi-Part Solid tutorial used simplified geometry of a PC tower to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Importing a CAD geometry as Parts.
- Using the Multi-Part Solid material model to model multiple solids with different materials.
- Creating a custom user field function to calculate the Volumetric Heat Source of multiple heat-generating parts in a single region.

Natural Convection: Concentric Cylinders Introduction

This tutorial demonstrates how to solve a natural convection problem in STAR-CCM+.

An infinitely long cylinder of radius 1.78 cm, heated at a constant temperature of 306.3 K, is placed inside another larger cylinder of radius 4.628 cm at a lower fixed temperature of 293.7 K. These temperatures are chosen to give an even temperature distribution about a bulk temperature (T_b) of 300 K in the annulus. As the flow is expected to be symmetrical about the center-line, it is only necessary to work with half the geometry, as shown below.



Set up the model for an ideal gas with the following material properties, evaluated at bulk temperature of 300 K and pressure of 100,000 Pa.

Property	Setting
Density (kg/m^2)	1.1614
Dynamic viscosity (Pa-s)	1.846E-5

Property	Setting
Kinematic viscosity (m^2/s)	1.589E-5
Specific Heat (J/kg-K)	1007.0
Thermal Coefficient of Volumetric Expansion (1/K)	$1/T_b$ (Ideal Gas) = 0.00333
Thermal Conductivity (W/m-K)	0.0263
Thermal Diffusivity (m^2/s)	2.249E-5

This case repeats the experiments of Kuehn and Goldstein [416]. Set up the physics to correspond to a Rayleigh number of 2.66 X 10^4.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting Up the Physics Models
Setting Boundary Physics Conditions
Setting Solver Criteria
Preparing Scalar and Vector Scenes
Preparing Reports to Monitor Convergence
Running the Simulation
Visualizing the Results
Comparing Results to Experimental Data
Summary
References

Prerequisites

The instructions in the Natural Convection tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

Use the following steps:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/heatTransferAndRadiation/data`.
5. Select `concylMesh.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a geometry scene is displayed in the *Graphics* window.
6. Save the simulation as `conCyl.sim`.

Setting Up the Physics Models

Set up the physics using an implicit coupled flow solver due to its increased robustness over the segregated solver.

Use the following steps:

1. For the physics continuum, **Continua > Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
Viscous Regime	**Laminar**
Optional Models	**Gravity**

2. Click **Close**.
3. To review the models, open the **Physics 1 > Models** node.



4. Expand the **Continua** > **Physics 1** > **Models** > **Gas** > **Air** node.
5. Edit the **Material Properties** node and set the following properties:

Node	Property	Setting
Dynamic Viscosity (Pa-s)		
Constant	*Value*	`1.846E-5`
Specific Heat (J/kg-k)		
Constant	*Value*	`1007.0`
Thermal Conductivity (W/m-K)		
Constant	*Value*	`0.0263`

6. Click **Close**.
7. Edit the **Continua** > **Physics 1** > **Reference Values** node and set the following properties:

Node	Property	Setting
Reference Density (kg/m^3)	*Value*	`1.1614`
Reference Pressure (Pa)	*Value*	`100000.0`

Note: The initial static temperature is by default 300 K, which is the average of the fixed inner and outer wall temperatures. The initial static

temperature is located in the **Continua > Physics 1 > Initial Conditions > Static Temperature** node. The material properties were evaluated at this bulk temperature.

Setting Boundary Physics Conditions

A fixed temperature difference across the annulus is achieved by setting the static temperatures of the inner and outer cylinder walls.

Use the following steps:

1. Expand the **Regions > ConvectionCylinders** node.
2. Edit the **Boundaries** node and set the following properties:

Node	Property	Setting
Inner Wall		
Physics Conditions		
Thermal Specification	*Condition*	`Temperature`
Physics Values		
Static Temperature		
Constant	*Value*	`306.3 K`
Outer Wall		
Physics Conditions		
Thermal Specification	*Condition*	`Temperature`
Physics Values		
Static Temperature		
Constant	*Value*	`293.7 K`

Setting Solver Criteria

Increase the default value of the Courant number of the coupled implicit solver to achieve faster convergence.

You generally want a Courant number that balances the total number of iterations against the time per iteration.

1. Select the **Solvers** > **Coupled Implicit** node and set *Courant Number* to `100.0`.

Because you are using a relatively large Courant number on a simple convection problem, you can expect fast convergence. Therefore decrease the number of iterations for which to run the solver.

2. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `300`.

Preparing Scalar and Vector Scenes

A scalar and a vector scene are created to display the temperature and velocity vectors in the region bounded by the concentric cylinders during the simulation.

Use the following steps:

1. Create a scalar scene.
2. Expand the **Scalar Scene 1** > **Displayers** > **Scalar 1** node.
3. Select the **Scalar Field** node and set *Function* to `Temperature`.
4. Create a vector scene.
   By default, *Vector Field* is set to `Velocity`.

Preparing Reports to Monitor Convergence

To judge convergence for this natural convection problem, monitor the heat balance in your model, which sums to zero for a converged case.

Prepare a report for the inner wall.

1. Create a **Heat Transfer** report.
2. Rename the new report to `Heat Transfer Inner`.
3. Edit the *Parts* property.
4. In the *Heat Transfer Inner - Parts* dialog, select **Regions** > **ConvectionCylinders** > **Inner Wall**.
5. Click **OK**.

Prepare a report for the outer wall.

6. Create another **Heat Transfer** report.
7. Rename this report to `Heat Transfer Outer`.
8. Edit the *Parts* property.
9. In the *Heat Transfer Outer - Parts* dialog, select **Regions > ConvectionCylinders > Outer Wall**.
10. Click **OK**.

Prepare an expression report that calculates the heat balance using the results of the reports. A corresponding field function is generated for each report, allowing you to use them in an expression.

11. Create an **Expression** report.
12. Rename this report to `Heat Balance`.
13. Click ... (**Custom Editor**) in the *Definition* property.
14. In the *Heat Balance - Definition* dialog, enter `${HeatTransferOuterReport} + ${HeatTransferInnerReport}`.
15. Click **OK**.
16. Right-click the **Heat Balance** node and select **Create Monitor and Plot from Report**.

    New monitor and plot nodes are added to the **Monitors** and **Plots** nodes respectively.
17. Open the **Heat Balance Monitor Plot**.

Running the Simulation

The simulation is now ready to run.

Use the following steps:

1. Run the simulation.
2. To watch the solution being updated while the simulation is running, look at the scalar scene and vector scene displays.
3. When the simulation has finished running, save it.

Visualizing the Results

The scalar and vector scenes contain the solution results.

As expected, the temperature decreases between the wall of the inner cylinder and that of the outer cylinder. This decrease corresponds to the

difference in the fixed wall temperatures of the two cylinders. The warmer air rises as it moves from the wall of the inner cylinder towards the outer cylinder, and sinks to the bottom of the annulus as it cools. The air recirculates within the annulus due to natural convection, with the largest velocities seen near the walls of the cylinders, towards the top.

1. Look at these results in the displays.





The heat balance quickly drops to about zero, showing that the problem converged correctly.



Comparing Results to Experimental Data

This part of the tutorial compares results and data using equivalent conductivity.

Kuehn and Goldstein reported their experimental and numerical results in the form of an equivalent conductivity. This equivalent conductivity is defined as the ratio of the actual heat transfer measured (or predicted) across the annulus to the heat transfer that would be obtained by pure conduction alone. That is:

$$k_{eq} = \frac{q_{act}}{q_{cond}} \tag{2}$$

For concentric cylinders, q_{cond} is defined as:

$$q_{cond} = \frac{2\pi k \Delta T}{\ln\left(R_o / R_i\right)} \tag{3}$$

For the case being considered in this tutorial, Kuehn and Goldstein obtained an equivalent conductivity of 2.58.

1. To begin evaluating conductivity, calculate the conductive heat transfer based on Eqn. (4) using values provided earlier in your setup. This comes out to $q_{cond} = 2.177 \ \ W / m$.
2. To retrieve the value of q_{act} for one half of the cylinders, double-click the **Heat Transfer Inner** or **Heat Transfer Outer** node.

   The heat transfer appears in the *Output* window.



The total q_{act} for the complete annulus is therefore 5.56 W.

Inserting this value into Eqn. (3) above calculates an equivalent conductivity of:

$$k_{eq} = \frac{5.56}{2.177} = 2.55 \tag{4}$$

There is only about a 1% difference between this result and the value given by Kuehn and Goldstein, showing that the numerical model has predicted the heat transfer consistent to what has been reported by others.

Note: Because you are working with a 2D representation of the cylinders, the prediction of heat transfer is per meter length, that is, W/m.

Summary

The Natural Convection tutorial used two concentric cylinders to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Importing a volume mesh.
- Defining physics models for laminar flow with natural convection.

- Defining the material properties for the selected models.
- Defining fixed physics conditions at the regions level.
- Setting a convergence condition for the coupled implicit solver.
- Preparing scalar and vector scenes.
- Preparing a heat balance report.
- Running the simulation.
- Examining the temperature, velocity field and heat balance solutions.
- Comparing solution to experimental results.

References

This tutorial repeats the experiments of Kuehn and Goldstein (1978).

Kuehn, T. H and Goldstein, R. J. An experimental study of natural convection heat transfer in concentric and eccentric horizontal cylindrical annuli. J. Heat Trans. (100) pp. 635-640, 1978.

Dual Stream Heat Exchanger: Car Radiator

This tutorial simulates a three-dimensional fluid flow through an idealized water-cooled car radiator.

The purpose of this tutorial is to demonstrate the heat exchanger capabilities of STAR-CCM+, where the radiator transfers heat from the liquid coolant to the outside air.



Two key operations are demonstrated in this tutorial, which uses the trimmed cell method for the volume mesh:

- The mesh alignment option is activated so that a conformal match is obtained across mesh interfaces. STAR-CCM+ does not automatically generate a conformal mesh across interfaces for the trimmed cell mesh type.
- The core region of the radiator is duplicated, which is an essential requirement for the heat exchanger model in STAR-CCM+. Some

modification of the interfaces is required to make sure that the correct regions are connected after the duplication.

Contents

Prerequisites
Importing the Surface Mesh and Naming the Simulation
Assigning to a New Region
Generating the Volume Mesh
Preparing Boundaries and Interfaces
Selecting Physics Models
Assigning Physics Continua
Setting the Boundary Conditions
Specifying Porosity Coefficients
Creating the Actual Heat Exchanger Interface
Reporting and Monitoring with Plots
Setting the Stopping Criteria
Visualization Scenes
Running the Simulation
Visualizing Results
Summary

Prerequisites

The instructions in the Dual Stream Heat Exchanger tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Surface Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

A surface mesh for a car radiator has been prepared for this analysis.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Surface Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/heatTransferAndRadiation/data`.
5. Select `radiator_parts.x_b` then click **Open**.
6. To accept the default *Import Surface Options*, click **OK**.

   STAR-CCM+ provides feedback on the import process in the *Output* window. A geometry scene is also created in the *Graphics* window.
7. Save the simulation as `radiator_hx.sim`.
8. Copy the `uag_table.csv` file to the current working directory.

Assigning to a New Region

Five fluid regions are required for this simulation: Two for the air ducts, two for the coolant, and one for the core.

Assign all the parts to new regions:

1. Right-click all the parts in the **Geometry > Parts** node and select **Assign Parts to Regions**.
2. In the *Assign Parts to Regions* dialog, set the modes as follows:

Mode	Setting
Region mode	**Create a Region for Each Part**
Boundary mode	**Create a Boundary for Each Part Surface**

3. Click **Apply** and then click **Close**.
4. Expand the **Regions** and **Interfaces** nodes.



The required regions, boundaries, and interfaces corresponding to the imported parts are all created. You can see that there are four interfaces corresponding to the four sets of part/core contact data.

Generating the Volume Mesh

Create a volume mesh for the car radiator.

Activate the surface remesher and trimmer meshing models to generate the volume mesh.

1. Create a mesh continuum.
2. For the mesh continuum, **Continua** > **Mesh 1**, select the following models in order:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Trimmer**

3. Click **Close**.

In this tutorial, the meshing is done per region.

4. Select the **Mesh 1** node and activate the *Per-Region Meshing* property.
5. Select the **Mesh 1** > **Models** > **Trimmer** node and activate the *Do mesh alignment* property.
6. Edit the **Mesh 1** > **Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	*Value*	`2.56 cm`
Mesh Alignment Location	*Location*	`[-0.166693, 0.293814, 0.126133] m`
Relative Minimum Size	*Percentage of Base*	`1`

7. Click **Close**.

Create a block volume shape around the core that you want to use to define volumetric controls for mesh refinement.

8. Right-click the **Tools > Volume Shapes** node and select **New Shape > Block**.
9. In the *Create Block Volume Shape* in-place dialog, make the following settings in the *Maximum and Minimum Coordinates* box:

Property	Setting
Corner 1	`[-0.18477, -0.15894, 0.100453]` m
Corner 2	`[0.29612, 0.38635, 0.179]` m

10. Click **Create** and then **Close**.
11. Open the **Tools > Volume Shapes** node and select the **Block 1** node.

    The volume block is displayed in the *Graphics* window as shown below.



Create a volumetric control and select the shape that defines the source.

12. Right-click the **Continua** > **Mesh 1** > **Volumetric Controls** node and select **New**.

    A new node, **Volumetric Control 1**, is added below the **Volumetric Controls** node and is selected automatically.
13. Edit the *Shapes* property.
14. In the *Volumetric Control 1* dialog, select **Block 1**.
15. Click **OK**.
16. Select the **Volumetric Control 1** > **Mesh Conditions** > **Surface Remesher** node and activate *Customize surface remesher*.
17. Select the **Mesh Conditions** > **Trimmer** node and activate *Customize isotropic size*.

    A new node called **Mesh Values** is added to the simulation tree.
18. Select the **Volumetric Controls 1** > **Mesh Values** > **Custom Size** node and set *Size Type* to `Absolute`.

    The **Relative Size** is changed to **Absolute Size** under the **Custom Size** in the simulation tree.
19. Select **Custom Size** > **Absolute Size** and set *Value* to `0.64 cm`.

    The volume mesh can now be generated.

20. To start the mesh generation process, click (**Generate Volume Mesh**) in the **Mesh** toolbar.

    The output from the meshers is shown in the *Output* window. When the meshing process is complete, the message `Volume Meshing Pipeline Completed` is displayed in the *Output* window.
21. Save the simulation.

Preparing Boundaries and Interfaces

STAR-CCM+ requires the existence of an interface between all regions to transfer the appropriate mass and energy quantities during the calculation.

Rename the radiator core and create a duplicate.

1. Rename the **Regions** > **core** node to `air_core`.
2. Create a copy of the **air_core** node and rename it `coolant_core`.
3. Expand the **Interfaces** node and delete the **coolantBase/core** and **coolantTop/core** nodes.

Create the required interfaces.

4. Expand the **Regions** node and create two interfaces:
   a) Create the first interface between the **coolant core** > **Boundaries** > **base_face** node and the **coolantBase** > **Boundaries** > **Coolant_base_interface** node.
   b) Create the second interface between the **coolant core** > **Boundaries** > **top_face** node and the **coolantTop** > **Boundaries** > **Coolant_top_interface** node.

   By default, the **Interface 1** and **Interface 2** interfaces have the *Type* property set to `Internal Interface` and the *Topology* property set to `In-place`.

Selecting Physics Models

Physics models are required to define the behavior of the air and coolant continua.

Use the following steps:

1. Rename the **Continua** > **Physics 1** node to `Air`.

2. For the physics continuum, **Continua** > **Air**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically) **Two-Layer All y+ Wall Treatment** (Selected automatically)
Optional Models	**Segregated Fluid Temperature**

3. Click **Close**.

   The color of the **Air** node turns from gray to blue to indicate that air models have been activated and the continuum is valid.

4. To review the models, open the **Air** > **Models** node.



5. To specify the physics models for the coolant, create a physics continuum.
6. Rename the **Continua** > **Physics 1** node to `Coolant`.
7. For the physics continuum, **Continua** > **Coolant**, select the following models in order:

Group Box	Model
Space	**Three Dimensional** **Gradients** (Selected automatically)
Time	**Steady**
Material	**Liquid**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically)

Group Box	Model
	Two-Layer All y+ Wall Treatment (Selected automatically)
Optional Models	**Segregated Fluid Temperature**

8. Click **Close**.

   The color of the **Coolant** node turns from gray to blue to indicate that coolant models have been activated and the continuum is valid.

9. To review the models, open the **Coolant** > **Models** node.



Assigning Physics Continua

After physics continua have been defined, they can be assigned to their corresponding regions.

Use the following steps:

1. Edit the **Regions** node and set the following properties:

Node	Property	Setting
coolant_core	*Physics Continuum*	`Coolant`
coolantBase	*Physics Continuum*	`Coolant`
coolantTop	*Physics Continuum*	`Coolant`

2. Edit the **Interfaces** node and set the following properties:

Node	Property	Setting
In-place 1	*Type*	`Internal Interface`
In-place 2	*Type*	`Internal Interface`

3. Save the simulation.

Setting the Boundary Conditions

By default the boundaries representing the radiator are all set to no-slip walls. Some of them need to be changed to other boundary types.

Use the following steps:

1. Select the **Regions** > **airDuctUpstream** > **Boundaries** > **Air_inlet** node and set *Type* to `Velocity Inlet`.
2. Edit the **Air_inlet** > **Physics Values** node and set the properties as follows:

Node	Property	Setting
Static Temperature		
Constant	*Value*	`308 K`
Velocity Magnitude		
Constant	*Value*	`20 m/s`

3. Click **Close**.
4. Select the **Regions** > **airDuctDownstream** > **Boundaries** > **Air_pressure_out** node and set *Type* to `Pressure Outlet`.
5. Select the **Regions** > **coolantTop** > **Boundaries** > **Coolant_inlet** node and set *Type* to `Mass Flow Inlet`.
6. Edit the **Coolant_inlet** > **Physics Values** node and set the properties as follows:

Node	Property	Setting
Mass Flow Rate		
Constant	*Value*	`1.8 kg/s`
Total Temperature		
Constant	*Value*	`380 K`

7. Click **Close**.
8. Select the **Regions** > **coolantBase** > **Boundaries** > **Coolant_pressure_out** node and set *Type* to `Pressure Outlet`.

Specifying Porosity Coefficients

Define the air_core and coolant_core regions as porous and set their respective porosity coefficients.

Use the following steps:

1. Edit the **Regions** node and set the properties as follows:

Node	Property	Setting
air_core	*Type*	`Porous Region`
coolant_core	*Type*	`Porous Region`

2. Select the **air_core** > **Physics Values** > **Porous Inertial Resistance** node.

   The *Method* property is set to `Principal Tensor` by default. This setting is suitable because the constant resistance coefficients correspond to the principal (diagonal) components of the porous resistance tensor.
3. Edit the **Porous Inertial Resistance** > **Principal Tensor** node and set the properties as follows:

Node	Property	Setting
XX Component		
Constant	*Value*	`10000` kg/m^4
YY Component		
Constant	*Value*	`1e4` kg/m^4
ZZ Component		
Constant	*Value*	`90` kg/m^4

Note: Do not change the **XX Axis** and **YY Axis** coordinates since the porous resistance components are specified in the directions of the default axes.

4. Edit the **Porous Viscous Resistance** > **Principal Tensor** node and set the properties as follows:

Node	Property	Setting
XX Component		
Constant	*Value*	`1e5` kg/m^3-s
YY Component		
Constant	*Value*	`1e5` kg/m^3-s
ZZ Component		
Constant	*Value*	`450` kg/m^3-s

5. Edit the **coolant_core** > **Physics Values** > **Porous Inertial Resistance** > **Principal Tensor** node and set the properties as follows:

Node	Property	Setting
XX Component		
Constant	*Value*	`1e8` kg/m^4
YY Component		
Constant	*Value*	`2e6` kg/m^4
ZZ Component		
Constant	*Value*	`1e8` kg/m^4

6. Edit the **coolant_core** > **Physics Values** > **Porous Viscous Resistance** > **Principal Tensor** node and set the properties as follows:

Node	Property	Setting
XX Component		
Constant	*Value*	`1e6` kg/m^3-s
YY Component		
Constant	*Value*	`55000` kg/m^3-s
ZZ Component		
Constant	*Value*	`1e6` kg/m^3-s

Creating the Actual Heat Exchanger Interface

Specify the actual heat exchanger interface and use the actual flow dual stream option for the heat exchanger method.

The heat exchanger method uses the flow and energy values that are calculated for the simulation itself, to compute the mean temperature difference and hence the local heat transfer coefficient. You also set the heat exchanger data specification option to UAG Table. This option assumes that the overall heat transfer coefficient (times the heat exchanger area) across the entire heat exchanger is known as a table of the cold stream mass flow rate (say, from an experimental source).

To define the actual heat exchanger interface:

1. Using the multi-select method, select the following nodes:
   - **Regions > air_core**
   - **Regions > coolant_core**
2. Create a heat exchanger interface.

   The previous step creates the **Interfaces > Heat Exchanger 1** node.
3. Edit the **Heat Exchanger 1** node and set the following properties:

Property	Setting
Region-0 (Cold Region)	`air core`
Region-1 (Hot Region)	`coolant core`

4. Select the **Interfaces > Heat Exchanger 1 > Physics Conditions > Heat Exchanger Method** node and set *Option* to `Actual Flow Dual Stream`.

   The previous step creates the following nodes under the **Physics Conditions** node:
   - **Heat Exchanger Data Specification**
   - **Hot Stream Inlet Temperature Specification**
5. Select the **Heat Exchanger Data Specification** node and set *Method* to `UAG Table`.
6. Load the file containing the data:
   a) Create a file table node.
   b) In the *Open* dialog, navigate to your working directory and select the `uag_table.csv` file.
   c) Click **Open**.

   The previous step creates the **Tables > uag_table** node.

7. Select the **Interfaces > Heat Exchanger 1 > Physics Values > UAG Table** node and set the following properties:

Property	Setting
UAG Table	`uag_table`
UAG Table: Mass Flow Rate	`column0`
UAG Table: UAG	`column1`
Cold Upstream Boundaries	`upstream_face [airDuctUpstream/core]`
Hot Upstream Boundaries	`top_face [in-place 2]`

8. Select the **Physics Values > Heat Exchanger First Iteration** node and set *First Heat Exchanger Iteration* to `90`.
9. Save the simulation.

Reporting and Monitoring with Plots

By default STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the solution. You can also create more reports, monitors, and plots to help in assessing convergence.

This part of the tutorial involves creating three reports:

- A surface area report for the outlet air temperature.
- A surface area report for the outlet coolant temperature.
- A heat exchanger report.

Then you create monitors and plots for the above reports and use the plots to monitor the solution during the run.

1. Create a **Surface Average** report.
2. Rename the new report to `Air Temperature Out`.
3. Select the **Air Temperature Out** node and set the following properties:

Property	Setting
Scalar Field Function	`Temperature`
Parts	`Regions > airDuctDownstream > air_pressure_out`

4. Create another **Surface Average** report.
5. Rename this report to `Coolant Temperature Out`.

6. Select the **Coolant Temperature Out** node and set the following properties:

Property	Setting
Scalar Field Function	`Temperature`
Parts	`Regions > coolantBase > coolant_pressure_out`

7. Create a **Heat Exchanger (Dual Stream)** report.

   The **Heat Exchanger (Dual Stream) 1** node is selected automatically.

8. Set the following properties:

Property	Setting
Heat Exchanger	`Heat Exchanger 1`
Cold Downstream Boundaries	`Regions > air_core > downstream_face [airDuctDownstream/core]`
Hot Downstream Boundaries	`Regions > coolant_core > base_face [In-place 1]`

9. Create monitors and plots for the reports:
   a) Select the **Air Temperature Out** and **Coolant Temperature Out** nodes.
   b) Create monitors and plots for them.
   c) In the *Create Plot From Reports* dialog, select **Single Plot**.
   d) Create a monitor and plot for the **Heat Exchanger (Dual Stream) 1** report.
   e) In the *Create Plot From Reports* dialog, select **Single Plot**.
   f) Rename the **Reports Plot** node to `Temperature Plot`.
10. Open the **Heat Exchanger (Dual Stream) 1 Monitor Plot** and **Temperature Plot** plots.

Setting the Stopping Criteria

This is the point at which you would adjust the maximum number of iterations.

To set the stopping criteria:

1. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `1000`.

2. Save the simulation.

Visualization Scenes

Create a scalar scene to visualize temperature at a plane section through the air and coolant parts, and a vector scene to display the velocity vectors at a plane section through the coolant part.

Use the following steps:

1. Create a scalar scene.

Create the two plane sections.

2. Create a plane section derived part.
3. In the *Create Section* inset panel, make the following settings:

Property	Setting
Input Parts	`[air_core, airDuctDownstream, airDuctUpstream]`
Plane Parameters	*origin:* `[0 m, 0 m, 0 m]` *normal:* `[1 m, 0 m, 0 m]`
Display	`No Displayer`

4. Click **Create** and then **Close**.
5. Rename the **Derived Parts > plane section** node to `Section Air`.
6. Create a second plane section with the following properties:

Property	Setting
Input Parts	`[coolant_core, coolantBase, coolantTop]`
Plane Parameters	*origin:* Default values *normal:* `[0 m, 0 m, 1 m]`
Display	`No Displayer`

7. Click **Create** and then **Close**.
8. Rename the **Derived Parts > plane section** node to `Section Coolant`.
9. Edit the **Displayers > Outline 1 > Parts** node.
10. In the *Edit* dialog, open the **Derived Parts** node and select **Section Air** and **Section Coolant**.
11. Click **OK**.

12. Copy the **Outline 1** > **Parts** node to **Scalar 1** > **Parts** node.
13. Select the **Scalar 1** > **Scalar Field** node and set *Function* to `Temperature`.
14. Select the **Scalar 1** node and set *Contour Style* to `Smooth Filled`.

Create a vector scene. For this scene, you leave the outline 1 parts all included.

15. Edit the **Scenes** > **Vector Scene 1** > **Displayers** > **Vector 1** > **Parts** node.
16. In the *Edit* dialog, select the **Derived Parts** > **Section Coolant** node.
17. Click **OK**.
18. Save the simulation.

Running the Simulation

The simulation is now ready to run.

Use the following steps:

1. Run the simulation.
2. To watch the solution being updated while the simulation is running, look at the scalar scene, vector scene, and plot displays.
3. When the simulation has finished running, save it.

Visualizing Results

To examine temperatures, velocity, and other results from the converged solution, look at the scenes and the plots.

1. To see the temperature distribution, look at the *Scalar Scene 1* display.



2. Look at the *Vector Scene 1* display.



3. Open the *Heat Exchanger (Dual Stream) 1 Monitor Plot* display.
4. Edit the **Plots > Heat Exchanger (Dual Stream) 1 Monitor Plot > Axes** node and set the following properties:

Node	Property	Setting
X Axis		
Labels	*Minimum*	50
	Maximum	550
Y Axis		
Labels	*Minimum*	208200
	Maximum	208800

The graph appears similar to the example that is shown below.



5. To see the air and coolant temperatures at the outlet for the converged solution, open the *Temperature Plot* display.



Summary

The Dual Stream Heat Exchanger tutorial used an idealized water-cooled car radiator to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Selecting fluid physics models.
- Creating and assigning a physics continuum.
- Specifying porosity coefficients.
- Creating an actual heat exchanger interface.
- Specifying the heat exchanger data specification using a UAG table.
- Creating reports, monitors, and plots.
- Running the simulation.
- Visualizing the results.

Surface-to-Surface Radiation: Thermal Insulator

This tutorial incorporates gray thermal radiation in a STAR-CCM+ simulation.

All conditions and properties are defined using the STAR-CCM+ GUI and solved using the segregated solver. The results are displayed using the post-processing tools available in STAR-CCM+.

The starting point for the tutorial is a three-dimensional hexahedral grid representing two glass plates within a rectangular enclosure containing air, as shown in the Figure below. Built with pro-STAR, the grid includes interfaces between the glass plates and surrounding air. The imported file contains information about the geometry, mesh, boundaries, and materials present in the problem.



The analysis covers the following heat transfer modes:

- Conduction
- Gray thermal radiation

Contents

Prerequisites
Importing the Mesh and Naming the Simulation

Splitting the Boundary
Renaming Regions and Boundaries
Creating Interfaces
Setting Up the Models
Setting Material Properties
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Reporting and Monitoring with Plots
Initializing and Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Surface-to-Surface Radiation tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

Use the following steps:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/heatTransferAndRadiation/data`.
5. Select `insulator.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and two mesh regions are created under the **Regions** node.

6. Save the simulation as `thermalInsulator.sim`.

Splitting the Boundary

One of the solid boundaries must be split into non-contiguous boundaries.

Use the following steps:

1. Expand the **Regions** > **Default_Solid** > **Boundaries** node.
2. Right-click the **Default_Boundary_Region** node and select **Split Non-Contiguous** from the pop-up menu.

   The boundary is split into four and three new boundary nodes are created as shown.



3. Save the simulation.

Renaming Regions and Boundaries

The simulation currently has the default region and boundary names. Choose more meaningful names.

Use the following steps:

1. Expand the **Regions** node.
2. Rename **Default_Fluid** to `Fluid`.
3. Rename **Default_Solid** to `Solid`.
4. Rename the boundaries within the **Solid** node:

Original Name	New Name
Default_Boundary_Region	`top-solid1`
Default_Boundary_Region 2	`bottom-solid2`
Default_Boundary_Region 3	`top-solid2`
Default_Boundary_Region 4	`bottom-solid1`

This portion of the object tree looks as shown below.



5. Save the simulation.

Creating Interfaces

STAR-CCM+ requires the existence of interfaces between different regions to transfer the appropriate physical quantities during the calculation.

For this example, you create a contact-type interface between the fluid and solid regions. This involves creating a so-called “in-place interface” between the wall boundaries enclosing each of the two regions, and then defining this interface as being of an appropriate “contact” type for fluid-solid heat transfer.

1. Using the multi-select method, select the following nodes:
   - **Regions > Fluid > Boundaries > top-solid1**
   - **Regions > Solid > Boundaries > top-solid1**

2. Right-click and select **Create Interface**.

   The previous step creates the following new nodes:

   - **Fluid > Boundaries > top-solid1 [Interface 1]**
   - **Solid > Boundaries > top-solid1 [Interface 1]**

   These boundaries form the connection between the fluid and solid regions once the flow initialization step has been performed and supersede the original wall boundaries forming the connection.
3. Select the following nodes:
   - **Regions > Fluid > Boundaries > bottom-solid1**
   - **Regions > Solid > Boundaries > bottom-solid1**
4. Right-click and select **Create Interface**.
5. Select the following nodes:
   - **Regions > Fluid > Boundaries > top-solid2**
   - **Regions > Solid > Boundaries > top-solid2**
6. Right-click and select **Create Interface**.
7. Select the following nodes:
   - **Regions > Fluid > Boundaries > bottom-solid2**
   - **Regions > Solid > Boundaries > bottom-solid2**
8. Right-click and select **Create Interface**.

   In addition to the new in-place boundaries, a folder node named **Interfaces** appears in the object tree. Within it, four interface nodes that are called **Interface 1** to **Interface 4** are created.
9. Save the simulation.

Setting Up the Models

Models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution. In this example, the flow is laminar and incompressible.

The Segregated Energy model is used together with the Surface-to-Surface Radiation model to determine the temperature field.

As there are two materials (air and glass) present, two regions are required for the analysis. You define two sets of Physics continuum models, each appropriate to a different region. By default, **Physics 1** and

Physics 2 are automatically created when the mesh is imported. You use these continua for the fluid and the solid.

1. Rename the **Continua** > **Physics 1** node to `Air`.
2. Rename the **Continua** > **Physics 2** node to `Glass`.
3. For the physics continuum, **Continua** > **Air**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Laminar**
Optional Models	**Segregated Fluid Temperature** **Radiation**
Radiation	**Surface-to-Surface Radiation**
Radiation Spectrum	**Gray Thermal Radiation**

4. Click **Close**.

   The color of the **Air** node turns from gray to blue to indicate that air models have been activated and the continuum is valid.
5. To review the models, open the **Air** > **Models** node.



6. For the physics continuum, **Continua** > **Glass**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Solid**
Optional Models	**Segregated Solid Energy Radiation**
Equation of State	**Constant Density**
Radiation	**Surface-to-Surface Radiation**
Radiation Spectrum	**Gray Thermal Radiation**

7. Click **Close**.

   The color of the **Glass** node has turned from gray to shaded gray to indicate that models have been activated for this continuum.

8. To review the models, open the **Glass** > **Models** node.



9. Save the simulation.

Setting Material Properties

Change the default material for the solid continuum to glass.

By default, glass has a density of 2500 kg/m^3, a specific heat of 750 J/kg-K and a thermal conductivity of 1.4 W/m-K. The surrounding air is considered incompressible: by default, it has a density of 1.18415 kg/m^3, a specific heat of 1003.62 J/kg-K, a thermal conductivity of 0.0260305 W/m-K, and a dynamic viscosity of 1.85508E-5 Pa-s.

1. Right-click the **Continua** > **Glass** > **Models** > **Solid** > **Al** node and select **Replace with...**.
2. In the *Replace Material* dialog, expand the **Material Databases** > **Standard** > **Solids**node. Select **Glass (Glass Plate)** and click **OK**.

Setting Boundary Conditions and Values

The thermal insulator consists of a closed heated box with two internal solid glass panels that are surrounded by air. The boundary

specifications that are required for this analysis are wall thermal conditions and radiation properties at the boundaries.

A fixed heat flux of 1000 W/m^2 is applied to the base of the box to provide the heat source. The top of the box is given a fixed temperature of 300 K which corresponds to the temperature of the environment. The vertical sides of the solid panels which are in contact with the environment is set to a fixed heat flux of -62.5 W/m^2. All other wall boundary conditions are considered to be adiabatic, which is the default condition. Contact-type conducting boundaries are also used at solid-fluid interfaces.

Optical radiation properties emissivity, reflectivity, transmissivity, and radiation temperature must also be set at all boundaries. The solver assumes Kirchhoff's Law. Make sure that, for each boundary or interface, emissivity + reflectivity + transmissivity = 1.

Transmissivity of the solid-fluid interfaces is defined at interfaces.

Start with the wall boundary definitions for the fluid region.

1. Edit the **Regions** > **Fluid** > **Boundaries** > **collector** node.
2. Set the following properties:

   STAR-CCM+ adds the decimal point and units automatically, so by simply entering `1000` you get an entry of `1000.0 W/m^2`. A positive value for heat flux indicates that heat is entering the domain while a negative value would indicate that heat is leaving the domain.

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Heat Flux**
Physics Values		
Heat Flux		
Constant	*Value*	`1000` W/m^2
Surface Emissivity		
Constant	*Value*	`1`
Surface Reflectivity	*Method*	**Auto Calculate**

Node	Property	Setting
Surface Transmissivity		
Constant	*Value*	`0`

The conditions for the **side-1**, **side-2**, and **side-3** wall boundaries (side walls of the fluid region) are left at the default adiabatic setting but the optical properties at these boundaries must be set. First, set the **side-1** wall boundary.

3. Edit the **Regions** > **Fluid** > **Boundaries** > **side-1** > **Physics Values** node.
4. Set the following properties:

Node	Property	Setting
Surface Emissivity		
Constant	*Value*	`0.95`
Surface Reflectivity	*Method*	**Auto Calculate**
Surface Transmissivity		
Constant	*Value*	`0`

For this boundary, the application of Kirchhoff's law gives a surface reflectivity of `0.05`.

The conditions for the **side-2** and **side-3** wall boundaries are identical to the conditions for the **side-1** boundary. They can be set up in the same way manually or copied and pasted from the **side-1** boundary.

5. Copy from the **Regions** > **Fluid** > **Boundaries** > **side-1** node and paste onto the **side-2** and **side-3** nodes.

   If you prefer, examine the properties of the **side-2** and **side-3** nodes to verify that the radiative boundary conditions are correct.

Set the conditions for the **top** wall boundary.

6. Edit the **Regions** > **Fluid** > **Boundaries** > **top** node.
7. Set the following properties:

Node	Property	Setting
Physics Conditions		

Node	Property	Setting
Thermal Specification	*Condition*	**Temperature**
Physics Values		
Static Temperature		
Constant	*Value*	`300 K`
Surface Emissivity		
Constant	*Value*	`0.8`
Surface Reflectivity	*Method*	**Auto Calculate**
Surface Transmissivity		
Constant	*Value*	`0.15`

The boundary conditions must be set for the solid-fluid interfaces. The emissivity and reflectivity of the solid-fluid interfaces must be specified at each **Interface [Interface #]** boundary, not the **Interface** boundary, as the latter became defunct when the interface was created (see the previous section Creating Interfaces).

8. Edit the **Regions** > **Fluid** > **Boundaries** node.
9. Set the following properties:

Node	Property	Setting
bottom-solid1 [Interface 2]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`0.775`
Surface Reflectivity	*Method*	**Auto Calculate** The **Auto Calculate** option for the **Surface Reflectivity** results in a reflectivity of `0.075`

Node	Property	Setting
		during the simulation.

Node	Property	Setting
bottom-solid2 [Interface 4]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`0.725`
Surface Reflectivity	*Method*	**Auto Calculate**

Node	Property	Setting
top-solid1 [Interface 1]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`0.8`
Surface Reflectivity	*Method*	**Auto Calculate**

Node	Property	Setting
top-solid2 [Interface 3]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`0.75`
Surface Reflectivity	*Method*	**Auto Calculate**

It is also necessary to set the emissivity, reflectivity, and transmissivity for interface boundaries in the solid region. For this tutorial, it is acceptable to set the emissivity and reflectivity to the same values on the

solid side of the interface as on the fluid side. The copy and paste operation is therefore sufficient.

10. For each of the boundary names listed below, copy from the node in **Regions** > **Fluid** > **Boundaries**, and paste onto the boundary node with the matching name in **Regions** > **Solid** > **Boundaries**:
    - **bottom-solid1 [Interface 2]**
    - **bottom-solid2 [Interface 4]**
    - **top-solid1 [Interface 1]**
    - **top-solid2 [Interface 3]**

Set the solid wall boundary conditions and values in a similar way to the fluid ones.

11. Edit the **Regions** > **Solid** > **Boundaries** > **side-solid1** node.
12. Set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Heat Flux**
Physics Values		
Heat Flux		
Constant	*Value*	`-62.5` W/m^2 The negative value indicates that heat is leaving the domain.
Surface Emissivity		
Constant	*Value*	`0`
Surface Reflectivity	*Method*	**Constant**
Constant	*Value*	`1`
Surface Transmissivity		
Constant	*Value*	`0`

The conditions for wall boundary **side-solid2** are the same as for **side-solid1** and can be copied and pasted from the **side-solid1** boundary.

13. Copy from the **side-solid1** node and paste onto the **side-solid2** node.

Set the surface transmissivity at the interfaces.

14. Expand the **Interfaces** node.
15. In each interface node, select the **Physics Values** > **Surface Transmissivity** > **Constant** node and set *Value* as follows for each interface:

Interface Node	Surface Transmissivity Value
Interface 1	`0.15`
Interface 2	`0.15`
Interface 3	`0.175`
Interface 4	`0.175`

16. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Adjust solver controls such as the under-relaxation factors and the maximum number of iterations.

Since you are only solving for temperature in this tutorial, the segregated flow solver is frozen. In most situations, the Solver Frozen feature is primarily a “debugging” tool and not supported for normal simulation work.

1. Select the **Solvers** > **Segregated Flow** node and activate *Solver Frozen*.
2. To assist solution convergence, select the **Solvers** > **Segregated Energy** node and set the under-relaxation factors for fluid and solid segregated energy as follows:

Property	Setting
Fluid Under-Relaxation Factor	`0.55`
Solid Under-Relaxation Factor	`0.99`

The default stopping criterion is set to end the calculation after 1000 iterations. Such a value is excessive, so reduce the maximum number of iterations.

3. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `500`.

Create a stopping criterion that is based on the energy equation residual. Set the minimum tolerance to 1E-6.

4. Right-click the **Stopping Criteria** node and select **Create New Criterion** > **From Monitor**, then select **Energy** in the *Select Monitor* dialog.
5. Select **Stopping Criteria** > **Energy Criterion** > **Minimum Limit** node and set *Minimum Value* to `1.0E-6`
6. Save the simulation.

Visualizing the Solution

View the temperature of the fluid and solid regions as the solution develops.

Start by creating a plane section part for display.

1. Create a plane section derived part.
2. In the *Create Section* inset panel, make the following settings:

Property	Setting
Input Parts	`[Fluid, Solid]`
Plane Parameters	*origin:* `[0.5 m, 0.5 m, 0.5 m]` *normal:* `[1 m, 0 m, 0 m]`
Display	`Existing Displayer`

3. Click **Create** and then **Close**.

   A node appears within the **Derived Parts** node: **plane section**.

The plane section part can now be added to any suitable scene or be displayed in its own scene if desired. Create a scalar scene and display the plane section and other parts colored by the temperature scalar value.

4. Create a scalar scene.
5. To zoom in and position the part as shown below, use the mouse buttons and click (**Save-Restore-Select views**).



6. Edit the **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Parts** node.
7. In the *Edit* dialog, do the following:
   a) Expand the **Regions > Fluid > Boundaries** node and the **Derived Parts** node.
   b) Select the following parts:
      - **top**
      - **bottom-solid1 [Interface2]**
      - **top-solid2 [Interface3]**
      - **collector**
      - **plane section**
8. Click **OK**.
9. Right-click the scalar bar in the *Graphics* window display. In the pop-up menu that appears, select **Temperature**.
10. Look at the *Scalar Scene 1* display to verify that it contains the thermal insulator geometry along with a color bar for temperature.
11. Save the simulation.

Reporting and Monitoring with Plots

In this case, a report and plot are used for monitoring the irradiation at one of the interfaces between the glass panels and surrounding air.

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. The process involves setting up a report which defines the quantity of interest and the region parts that are monitored. A monitor is then defined based on the report, which controls the update frequency and normalization characteristics. Finally, a graph plot can then be created from the monitor.

This process starts with the report definition.

1. Create a **Surface Average** report.
2. Rename the report to `Irradiation1`.
3. Select the **Irradiation1** node and set the following properties:

Property	Setting
Scalar Field Function	`Boundary Irradiation`
Parts	`bottom-solid1 [Interface 2]`

The reason for selecting the **bottom-solid1 [Interface 2]** boundary instead of the **bottom-solid1** boundary is that after initialization, the in-place interface conditions supersede the conditions of the original wall boundary. Therefore, any values you wish to extract from boundaries on region interfaces must always be from the in-place boundary interfaces.

The setup for the report is now complete, and a monitor and plot can be made from that report.

4. Create a monitor and a plot for the report.
5. Open the plot display.

   A new plot display that is called Irradiation1 Monitor Plot appears in the *Graphics* window.

   This display is automatically updated with boundary irradiation data when the run starts.

Create three more reports in the same way as described above.

6. Create a **Surface Average** report.
7. Rename the report to `HeatFlux1`.
8. Select the **HeatFlux1** node and set the following properties:

Property	Setting
Scalar Field Function	`Boundary Heat Flux`
Parts	`Fluid:top`

9. Create a **Maximum** report.
10. Rename the report to `MaxTempFluid1`.
11. Select the **MaxTempFluid1** node and set the following properties:

Property	Setting
Scalar Field Function	`Temperature`
Parts	`Fluid`

12. Create another **Maximum** report.
13. Rename the report to `MaxTempSolid1`.
14. Select the **MaxTempSolid1** node and set the following properties:

Property	Setting
Scalar Field Function	`Temperature`
Parts	`Solid`

15. Create monitors and plots for the reports:
    a) Select the **MaxTempSolid1** and **MaxTempFluid1** nodes.
    b) Create monitors and plots for them.
    c) In the *Create Plot From Reports* dialog, select **Single Plot**.
16. Open the **Reports Plot** plot.
    This completes the pre-processing setup for the case.
17. Save the simulation.

Initializing and Running the Simulation

To verify that the initial conditions for the regions are correct before you run the simulation, STAR-CCM+ allows you simply to perform solver initialization and then view the results using an appropriate scene.

Use the following steps:

1. Initialize the simulation.
2. To see the initialization result, look at the scalar scene display.



Based on your radiation model selections, STAR-CCM+ generates radiation patches and calculates view factors automatically. Patches are generated when the solution is initialized, and view factors are calculated during the first iteration of the simulation run. Hence, this process is reported in the *Output* window before the iteration information appears.

The simulation can now be run.

3. Run the simulation.
4. To watch the solution being updated while the simulation is running, look at the scalar scene display.
5. When the simulation has finished running, save it.

Visualizing the Results

To examine irradiation, temperature, and other results from the converged solution, look at the plots and the scene.

1. At the end of the calculation, to view the irradiation variation at the solid-fluid interface, activate the *Irradiation1 Monitor Plot* display.



2. To see the temperature result for the converged solution, look at the *Scalar Scene 1* display.



3. To make the display smoother, set *Contour Style* to `Smooth Filled`. The contours of the scalar display now appear smooth.



4. Look at the *Reports Plot* display to see the graph of the maximum temperature monitors for the fluid and the solid.



5. To zoom in on the relevant part of the graph, edit the **Plots** > **Reports Plot** > **Axes** node and set the following properties:

Node	Property	Setting
X Axis	*Minimum*	50
Y Axis	*Minimum*	470
	Maximum	570
Major Labels	*Spacing*	10

The graph now appears as shown below.



6. To determine the heat flux that is leaving the system at the top boundary, run the **HeatFlux1** report.

The *Output* window shows that heat flux.

```
Output - thermalInsulator

        Surface Average of Boundary Heat Flux on Volume Mesh

Part                                    Value (W/m^2)
------------------------------          -------------
Fluid: top                               8.362639e+02
                                        -------------
                           Total:        8.362639e+02
```

7. Save the simulation.

Summary

The Surface-to-Surface Radiation tutorial used two glass plates and air within a rectangular enclosure to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Splitting boundaries.
- Defining models for multiple continua.
- Defining material properties that are required for the selected models.
- Defining boundary conditions.
- Setting solver parameters.
- Setting up a monitoring report and plot.
- Initializing and running the solver until the residual magnitudes are sufficiently small.
- Analyzing results using the visualization, reports, and monitor plot facilities.

Multiband Surface-to-Surface Radiation: Solar Collector

This tutorial uses the simulation created for the Surface-to-Surface Radiation: Thermal Insulator.

This tutorial uses the geometry and initial problem setup from the Surface-to-Surface Radiation: Thermal Insulator. It solves the same problem, except that solar radiation effects are also included using a solar load and the multiband radiation model. The multiband radiation model is used to simulate the behavior of thin glass, which transmits most solar radiation while absorbing most infrared radiation. For the purposes of this tutorial, the glass is given a high transmissivity for wavelengths in the range 0.1 to 3.0 mm , and lower values in the range 3.0 mm to 1.0 m.

While the multiband radiation model is used here to represent radiative properties that are dependent on wavelength, modeling solar radiation does not require the multiband model. If property variations with wavelength are not being considered, solar loads may also be applied with the gray radiation spectrum model.

The influence of this additional heat transfer mechanism on the overall heat transfer to the collector is considered by comparing the two sets of results.

The analysis covers the following heat transfer modes:

- Conduction
- Multiband thermal radiation
- Solar load

Note: Radiative properties given in this tutorial are for demonstration only.

Contents

Prerequisites
Loading an Existing Simulation
Modifying the Model Selection
Setting Solar Loads Model Properties
Setting Multiband Thermal Radiation Properties
Setting Boundary Conditions and Values
Setting Solver Parameters
Reporting and Monitoring with Plots

Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Multiband Surface-to-Surface Radiation tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading an Existing Simulation

For this tutorial, you are provided with a simulation file that contains pre-defined objects.

Load the initial simulation:

1. Launch STAR-CCM+.

Note: To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file.

2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, locate the `thermalInsulator.sim` file that was created in the Thermal Insulator Tutorial.
5. Select `thermalInsulator.sim` then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `solarCollector.sim`.

Modifying the Model Selection

Set up the models for the solar collector case to include a solar load and multiband radiation.

Most of the models that are required in this case have already been selected (see Setting Up the Models in the Thermal Insulator Tutorial) and so do not need to be specified. It is, however, necessary to apply a solar load and set the properties for multiband radiation.

For the physics continuum, **Continua** > **Air**, change the choice of radiation models:

1. In the *Enabled Models* group box of the model selection dialog, deselect the **Gray Thermal Radiation** model.

   The *Radiation Spectrum* box appears on the left.
2. Select the following models in order:

Group Box	Model
Radiation Spectrum	**Multiband Thermal Radiation**
Solar Radiation	**Solar Loads**

3. Click **Close**
4. To review the models, open the **Air** > **Models** node.



Since solar radiation heat transfer is to be modeled in the glass solid as well as in the air, the Multiband Thermal Radiation and Solar Loads models must be activated in the solid continuum.

For the physics continuum, **Continua** > **Glass**, change the choice of radiation models:

5. In the *Enabled Models* group box of the model selection dialog, deselect the **Gray Thermal Radiation** model.

   The *Radiation Spectrum* box appears on the left.
6. Select the following models in order:

Group Box	Model
Radiation Spectrum	**Multiband Thermal Radiation**
Solar Radiation	**Solar Loads**

7. Click **Close**
8. To review the models, open the **Glass** > **Models** node.



9. Save the simulation.

Setting Solar Loads Model Properties

The Solar Loads Model requires definition of the sun direction vector in a specified coordinate system, the direct solar flux, and the diffuse solar flux.

The Sun Direction Vector is defined as having a horizontal component of 90 degrees and a vertical component of 0 degrees.

1. Select the **Continua** > **Air** > **Models** > **Solar Loads** > **Manual Specification** node and set *Azimuth* to `-90` with units of degrees.

The solar load model properties are the same for the glass continuum and may be copied and pasted from the air continuum node.

2. Copy from the **Continua** > **Air** > **Models** > **Solar Loads** node and paste onto the **Continua** > **Glass** > **Models** > **Solar Loads** node.

   If you prefer, examine the properties of the **Solar Loads** node to verify that the properties have copied correctly.
3. Save the simulation.

Setting Multiband Thermal Radiation Properties

Two spectral bands are created for this tutorial.

The first spectral band covers the range 0.1 to 3.0 μm, and the second covers the range 3.0 μm to 1.0 m.

1. Expand the **Continua** > **Air** > **Models** > **Multiband Thermal Radiation** node.
2. Right-click the **Spectral Bands** node and create a spectral band.

   A node appears within the **Spectral Bands** node: **Thermal 1**.
3. Create a second spectral band.
4. Edit the **Continua** > **Air** > **Models** > **Multiband Thermal Radiation** > **Spectral Bands** node and set the following properties:

Node	Property	Setting
Thermal 1	*Upper Wavelength*	`3.0E-6 m`
	Lower Wavelength	`1.0E-7 m`
Thermal 2	*Upper Wavelength*	`1.0 m`
	Lower Wavelength	`3.0E-6 m` This property has been set automatically based on the upper limit of the **Thermal 1** node.

Corresponding spectral bands must also be created for the **Glass** continuum. While it is possible to do this by repeating the same procedure, it can be more easily accomplished by copying and pasting.

5. Copy from the **Continua** > **Air** > **Models** > **Multiband Thermal Radiation** > **Spectral Bands** node and paste onto the **Continua** > **Glass** > **Models** > **Multiband Thermal Radiation** > **Spectral Bands** node.

   This operation adds two nodes, **Thermal 1** and **Thermal 2**, to the **Spectral Bands** node within the **Glass** continuum.
6. Save the simulation.

Setting Boundary Conditions and Values

Applying the Multiband Thermal Radiation model means that the optical properties at boundaries and interfaces must be revised.

For each boundary, it is now necessary to set the emissivity, reflectivity and transmissivity for each of the spectral bands that were created for the Multiband Thermal Radiation model. Property values that were set for the Thermal Insulator Tutorial are used for the **Thermal 2** spectral band, and new property values are assigned for the **Thermal 1** spectral band.

To make sure that almost all of the energy in the **Thermal 1** spectral band passes into the box, the top of the collector is assigned a reflectivity of 0 and a transmissivity of 0.98. The base of the collector is assigned a reflectivity of 0 and a transmissivity of 0 for both spectral bands.

Note: The non-radiative thermal conditions for the boundaries remain as before.

1. Select the **Regions** > **Fluid** > **Boundaries** > **collector** > **Physics Values** > **Surface Emissivity** > **Constant** node.

As there are multiple spectral bands present, the **Constant** node takes the form of an array profile value. The first value in the array corresponds to the **Thermal 1** spectral band and the second value corresponds to the **Thermal 2** spectral band. The array is shown in short notation, and can be fully visualized by opening the custom editor.

2. Edit the *Value* property of the **Constant** node.
3. In the *Constant - Value* dialog, set both *Thermal 1* and *Thermal 2* (the constant emissivity value) to `1.0`.
4. Click **OK**.
5. Edit the **collector** > **Physics Values** node.
6. Set the following properties:

Node	Property	Setting
Surface Reflectivity	*Method*	`Auto Calculate`
Surface Transmissivity		
▮▮ **Constant**	*Value*	`[0.0, 0.0]`

Set the optical properties for the **side-1**, **side-2**, and **side-3** wall boundaries:

7. Edit the **Regions** > **Fluid** > **Boundaries** > **side-1** > **Physics Values** node.
8. Set the following properties:

Node	Property	Setting
Surface Emissivity		
Constant	*Value*	`[0.95, 0.95]`
Surface Transmissivity		
Constant	*Value*	`[0.0, 0.0]`

The properties of the **side-1** node can be copied and pasted to the **side-2** and **side-3** nodes.

9. Copy from the **side-solid1** node and paste onto the **side-2** and **side-3** nodes.

Set the optical properties for the **top** boundary node.

10. Edit the **Regions** > **Fluid** > **Boundaries** > **top** > **Physics Values** node.
11. Set the following properties:

Node	Property	Setting
Surface Emissivity		
Constant	*Value*	`[0.02, 0.8]`
Surface Transmissivity		
Constant	*Value*	`[0.98, 0.15]`

As in the Thermal Insulator tutorial, it is also necessary to set the optical properties on either side of the solid-fluid interfaces. Set the fluid side properties first.

12. Edit the **Regions** > **Fluid** > **Boundaries** node.
13. Set the following properties:

Node	Property	Setting
bottom-solid1 [Interface 2]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`[0.1, 0.775]`

Node	Property	Setting
bottom-solid2 [Interface 4]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`[0.2, 0.725]`
top-solid1 [Interface 1]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`[0.15, 0.8]`
top-solid2 [Interface 3]		
Physics Values		
Surface Emissivity		
Constant	*Value*	`[0.175, 0.75]`

Set the same optical properties for the corresponding solid-side of each interface. An efficient way of doing this is to copy and paste the relevant nodes.

14. For each of the boundary names listed below, copy from the node in **Regions** > **Fluid** > **Boundaries**, and paste onto the boundary node with the matching name in **Regions** > **Solid** > **Boundaries**:
    - **bottom-solid1 [Interface 2]**
    - **bottom-solid2 [Interface 4]**
    - **top-solid1 [Interface 1]**
    - **top-solid2 [Interface 3]**

The final boundaries to set are the side wall boundaries of the **Solid** region.

15. Edit the **Regions** > **Solid** > **Boundaries** > **side-solid1** > **Physics Values** node.
16. Set the following properties:

Node	Property	Setting
Surface Emissivity		
Constant	*Value*	`[0.0, 0.0]`
Surface Transmissivity		
Constant	*Value*	`[0.0, 0.0]`

This setting results in a **Surface Reflectivity** of `1.0` for both spectral bands at the edges of the solid. These same settings are copied to the other solid boundary.

17. Copy from the **side-solid1** node and paste onto the **side-solid2** node.

Set the surface transmissivity at the interfaces.

18. Expand the **Interfaces** node.
19. In each interface node, select the **Physics Values > Surface Transmissivity > Constant** node and set *Value* as follows for each interface:

Interface Node	Surface Transmissivity Value
Interface 1	`[0.8, 0.15]`
Interface 2	`[0.8, 0.15]`
Interface 3	`[0.75, 0.175]`
Interface 4	`[0.75, 0.175]`

20. Save the simulation.

Setting Solver Parameters

To ensure solution convergence, reduce the under-relaxation factor for fluid segregated energy.

Use the following steps:

1. Select the **Solvers > Segregated Energy** node and change the *Fluid Under-Relaxation Factor* to `0.5`.
2. Save the simulation.

Reporting and Monitoring with Plots

The monitor plots created in the Thermal Insulator Tutorial are used to monitor the simulation progress.

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. The process involves setting up a report which defines the quantity of interest and the region parts that are monitored. A monitor is then defined based on the report, which controls the update frequency and normalization characteristics. A graph plot can also be created from the monitor.

Several reports and monitors have already been set up for the solution in the Reporting, Monitoring, and Plotting section of the Thermal Insulator Tutorial.

1. Open the **Irradiation1 Monitor Plot** and **Reports Plot** plots.

To monitor the irradiation of the collector for the first spectral band, create a report.

2. Create a **Surface Average** report.
3. Rename the report to `SpectralBand1Irradiation`.
4. Select the **SpectralBand1Irradiation** node and set the following properties:

Property	Setting
Scalar Field Function	`Boundary Irradiation of Thermal 1`
Parts	`collector`

The setup for the report is now complete and a monitor and plot can be made from that report.

5. Create a monitor and a plot for the report.
6. Open the plot display.

   A new plot display that is called SpectralBand1Irradiation Monitor Plot appears in the *Graphics* window.



This display is automatically updated with boundary solar irradiation data when the run starts.

This completes the pre-processing setup for the case.

7. Save the simulation.

Running the Simulation

The solution from the Thermal Insulator tutorial can be inspected by displaying the initial temperature profile.

Review the previous solution results.

1. Open *Scalar Scene 1*.

   The temperature profile from the previous solution is displayed in the *Graphics* window.



Before running this simulation, clear the previous solution:

2. Select **Solution** > **Clear Solution**.
3. In the *Clear Solution* dialog, click **OK**.

   The simulation can now be run.

Generate a new solution.

4. Run the simulation.

   View factors are recalculated during the first iteration of the simulation run. Hence, see this process reported in the Output window before the iteration information appears.
5. To watch the solution being updated while the simulation is running, look at the scalar scene, vector scene, and plot displays.
6. When the simulation has finished running, save it.

Visualizing the Results

Examine the plots and scenes.

1. At the end of the calculation, to view the irradiation variation at the solid-fluid interface, activate the *Irradiation1 Monitor Plot* display.



The boundary irradiation is higher than its Thermal Insulator Tutorial counterpart due to the additional solar irradiation.

2. To view the solar irradiation at the collector, look at the SpectralBand1Irradiation Monitor Plot display.



3. To see the temperature result for the converged solution, look at the *Scalar Scene 1* display.



The maximum temperature has increased by approximately 80 K compared with the previous solution.

4. Look at the *Reports Plot* display to see the graph of the maximum temperature monitors for the fluid and the solid.
5. To zoom in on the relevant part of the graph, edit the **Plots** > **Reports Plot** > **Axes** node and set the following properties:

Node	Property	Setting
X Axis	*Minimum*	50
	Maximum	450
Major Labels	*Spacing*	50
Y Axis	*Minimum*	550
	Maximum	650
Major Labels	*Spacing*	10

The graph now appears as shown below.



Compare with the Reports Plot from the previous solution. The maximum temperature of the fluid has increased from ~540 to ~620 as a result of the solar radiation effect. The maximum temperature of the solid has increased from ~490 to ~580.

The **Heat Flux 1** report may be run .

6. To determine the heat flux leaving the system at the top boundary, run the **HeatFlux1** report.
7. The *Output* window shows the heat flux that is leaving the system.



8. Save the simulation.

Summary

The Multiband Surface-to-Surface Radiation tutorial used the geometry and initial problem setup from the Thermal Insulator Tutorial to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Starting the code and loading an existing simulation.
- Introducing solar loads and multiband thermal radiation into the problem conditions.
- Defining boundary conditions for multiband thermal radiation.
- Setting up a monitoring report and plot.
- Running the solver until the residual magnitude is sufficiently small.
- Analyzing results using the visualization, reports and monitor plot facilities.

Thermal Comfort Wizard: Single Occupant in a Cabin

Use the Thermal Comfort Wizard to set up and run a thermal comfort analysis case for a single occupant in the cabin of a car.



The Thermal Comfort Management (TCM) model accounts for the properties of the human body. It allows you to observe the effects of flow scenarios on skin surface temperatures. The TCM model runs the TIM program (a German acronym for "Thermophysiological Occupant Model") which calculates the thermal state of the occupant, based on the surrounding conditions.

Contents

Prerequisites
Importing the Surface Mesh and Naming the Simulation
Visualizing the Imported Geometry and Surface Mesh
Generating the Volume Mesh
Setting Up the Physics Continuum
Preparing Boundaries for the Thermal Comfort Wizard

Activating the Thermal Comfort Wizard
Using the Thermal Comfort Wizard
Setting Stopping Criteria
Creating a Scalar Scene
Running the Solver
Visualizing the Results with Streamlines
Summary

Prerequisites

The instructions in the Thermal Comfort Wizard tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Surface Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

Note: It is recommended that STAR-CCM+ is launched from an empty working directory as a number of files are generated during the TCM run.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Surface Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/heatTransferAndRadiation/data`.
5. Select `singleOccTCM.dbs` then click **Open**.
6. In the *Import Surface Options* dialog, select **Create New Region** from the *Import Mode* group box and click **OK**.

   STAR-CCM+ provides feedback on the import process in the *Output* window.

7. Save the simulation as `personCabin.sim`.

Visualizing the Imported Geometry and Surface Mesh

Use the visualization features to view selected parts of the cabin and examine the mesh.

The tutorial geometry consists of a single occupant in the cabin of a vehicle. There are five inlet vents at the front of the cabin, and an outlet vent at the rear.

1. Look at the *Geometry Scene 1* display.
2. Click (**Make Scene Transparent**).
3. Use the mouse buttons to adjust the view so that it appears as follows:



4. Click (**Save-Restore-Select views**) and select **Store Current View**.

The boundaries defining each body segment have been named according to the requirements of the Thermal Comfort Wizard.

Examine the imported surface mesh.

5. Click (**Make Scene Transparent**).
6. Edit the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** > **Parts** node .
7. In the *Edit* dialog, expand the **Regions** > **Region 1** node and deselect the **Doors**, **Roof**, and **Windows** boundaries.



8. Click **OK**.
9. Right-click the **Outline 1** node and select **Toggle Visibility**.



10. Click (**Show All Meshes**).



Generating the Volume Mesh

Create a volume mesh for the car cabin.

The starting surface can be used to generate a volume mesh without the need of the surface meshing models.

To select the meshing models and specify some of the global mesh parameters:

1. For the mesh continuum, **Continua** > **Mesh 1**, select the following models in order:

Group Box	Model
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**

2. Click **Close**.
3. Edit the **Mesh 1** > **Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	*Value*	`40 mm`
Prism Layer Thickness	*Location*	`[-0.166693, 0.293814, 0.126133] m`
Relative Size	*Percentage of Base*	`10`

4. Click **Close**.

The mesh has now been defined.

5. To start the mesh generation process, click (**Generate Volume Mesh**) in the **Mesh** toolbar.

   The output from the meshers is shown in the *Output* window. When the meshing process is complete, the message `Volume Meshing Pipeline Completed` is displayed in the *Output* window.

6. Right-click the **Scenes** > **Geometry Scene 1** node and select **Apply Representation** > **Volume Mesh**.

   The scene appears as follows:



7. Save the simulation.

Setting Up the Physics Continuum

Models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution. In this example, the flow is turbulent with constant density.

To model the air flow through the cabin:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically) **Two-Layer All y+ Wall Treatment** (Selected automatically)
Optional Models	**Segregated Fluid Temperature** **Thermal Comfort**

2. Click **Close**.

   The color of the **Physics 1** node turns from gray to blue to indicate that air models have been activated and the continuum is valid.

3. To review the models, open the **Physics 1** > **Models** node.



Activating the thermal comfort model adds two new nodes to the object tree, **TCM Passengers** and **TCM Boundaries**, as well as the **Solvers > Thermal Comfort** node. The new nodes are placeholders for information from the Thermal Comfort Wizard.



Preparing Boundaries for the Thermal Comfort Wizard

To specify thermal properties of each boundary within the Thermal Comfort Wizard, first choose a thermal specification at the boundary level for each boundary in the model.

The thermal values for each boundary are set using the wizard.

Note: The wizard can run the simulation for a number of iterations before calling the TIM program. During the initial iterations, all physics values are read from the active simulation, and not from the wizard. Enter alternative physics values for the initial iterations here if they are required. Initial values are overwritten by the TIM program when the TCM model calls it.

To choose the adiabatic, convection and temperature thermal specifications for the boundaries in the model:

1. Rename the **Regions** > **Region 1** node to `Fluid`.
2. Expand the **Regions** > **Fluid** > **Boundaries** node.
3. Using the multi-select method, select all boundary nodes from **1 - Head** to **14 - RightFoot**, inclusive.
4. Edit the group of selected nodes.



5. Set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Temperature**

6. Click **Close**.

Use a similar method to set the thermal specification of a number of the cabin boundaries to `Convection`.

7. Using the multi-select method, edit the following boundary nodes:
   - **Cabin**
   - **Console**
   - **Doors**
   - **Roof**
   - **Windows**
8. Set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Convection**

9. Click **Close**.

Contents

Setting Up the Inlet and Outlet Boundaries

Setting Up the Inlet and Outlet Boundaries

The inlet flow values are also set using the Thermal Comfort Wizard. For them to be read correctly from the wizard into the simulation, select the correct boundary types before the TCM model is run.

Specify the inlet boundaries as mass flow inlets. Change the turbulence specification at the inlets to turbulence intensity with length scale, and make sure that the outlet boundary is set to the correct type. Use the multi-edit dialog to specify the properties of the inlets simultaneously.

1. Expand the **Regions** > **Fluid** > **Boundaries** node.
2. Using the multi-select method, edit all boundary nodes from **InletDashCentre** to **InletFeetRight**, inclusive.

In the *Multi-edit* dialog:

3. Set *Type* to **Mass Flow Inlet**.
4. Set the following properties:

Node	Property	Setting
Physics Conditions		
Turbulence Specification	*Method*	**Intensity + Length Scale**

5. Click **Close**.
6. Save the simulation.

Activating the Thermal Comfort Wizard

To activate the wizard, add it to a toolbar in the STAR-CCM+ workspace.

Use the following steps:

1. In the menu bar, select **Window > Toolbars > Customize...**
2. In the *Customize Toolbars* dialog, scroll down to the bottom of the list and select **Thermal Comfort Wizard**.



3. Drag and drop the **Thermal Comfort Wizard** icon from the dialog onto the end of one of the toolbars at the top of the workspace.



4. Click **Close**.

Using the Thermal Comfort Wizard

Use of the Thermal Comfort wizard begins with launching this feature.

To launch the Thermal Comfort wizard:

1. Click (**Thermal Comfort Wizard**).

Contents

Entering Passenger Properties
Specifying Advanced Settings
Setting Thermal Properties
Saving TCM Settings

Entering Passenger Properties

To begin setting up the thermal comfort simulation, specify key physical characteristics of the occupants of the car.

In the *Thermal Comfort Wizard*, each occupant is referred to as a "passenger". The properties of each passenger are entered in the *General Settings* tab of the wizard, such as their height, metabolic rate, and clothing resistance. In this tab, you also specify which boundary corresponds to each of the TCM body segments.

Number of Passengers and Person Name/Number:

1. As you have modeled only one passenger, you can retain the default settings for *Number of Passengers* and *Person*. Had the cabin contained several occupants, you would define each person with a different name and number. Notice that the first passenger is given the default name of `Driver`.



Height:

2. The TIM program uses the passenger height to calculate the body mass distributions for the vascular system and the surface area values. Set the *Height* to `1.75` m.



Clothing Resistance:

3. The clothing resistance defines the thermal conductivity. Set the *Clothing Resistance* to `0.02` K m^2/W.



Metabolic Rate:

The metabolic rate is the rate at which the person releases heat energy. It depends on the level of activity of each person. It is set by selecting one of the **Driver** or **Passenger** radio buttons, or by entering a **Custom** value. When the **Passenger** option is selected the metabolic rate is set to `58 W/m^2`, which is equivalent to a person who is seated and relaxed. When the **Driver** option is chosen the metabolic rate is set to `75.4 W/m^2` (`1.3` times the passenger value), indicating that the driver is more active. The **Custom** option allows you to specify an alternative value by entering a scaling factor relative to 58 W/m^2. It is usually in the range of `0.8` to `2` (where `0.8` is equivalent to reclining, and `2` corresponds to medium activity.)

Set a custom metabolic rate for the single occupant:

4. Under the *Metabolic Rate* property, select the **Custom** radio button and specify a value of `1.24`.



Assigning Body Segments:

Now specify which boundary corresponds to each passenger body segment. The wizard recognizes the body segment boundaries according to their names, so it is important to observe the Thermal Comfort Wizard naming conventions when preparing a model.

5. Make sure that **Fluid** is selected under the *Select Model Region* property.



6. Under the *Head* property, use the drop-down menu to select the **1 - Head** boundary.



Note: Only those boundaries with names corresponding to relevant body segment are listed in the drop-down menu.

7. The initial temperature for the **Head** boundary is `35 C`. make sure that the *Initial Temp* property is set to `35`. All the temperatures that you specify in the *Thermal Comfort Wizard* are in degrees Celsius.



8. The corresponding boundary and initial temperature for each body segment is given in the table below. Use the data in this table to select the appropriate boundary and set the initial temperature for the remaining body segments.

Boundary Segment	**Boundary**	**Initial Temperature (C)**
Head	1 - Head	35
Torso	2 - Torso	35
Upper Left Arm	3 - LeftUpperArm	33
Upper Right Arm	4 - RightUpperArm	33
Lower Left Arm	5 - LeftLowerArm	33
Lower Right Arm	6 - RightLowerArm	33

Boundary Segment	Boundary	Initial Temperature (C)
Left Hand	7 - LeftHand	34
Right Hand	8 - RightHand	34
Upper Left Leg	9 - LeftUpperLeg	32
Upper Right Leg	10 - RightUpperLeg	32
Lower Left Leg	11 - LeftLowerLeg	31
Lower Right Leg	12 - RightLowerLeg	31
Left Foot	13 - LeftFoot	34
Right Foot	14 - RightFoot	34

The panel appears as follows:



The data that you entered in the *Thermal Comfort Wizard* dialog is retained automatically, allowing you to move between each of the tabs and even close and reopen the wizard without losing any information. The data in the wizard is only cleared if the STAR-CCM+ client is shut down.

Specifying Advanced Settings

Refine the TCM settings such as the convergence tolerance and the frequency for generating output files.

Use the following steps:

1. Click the *Advanced Settings* tab at the top of the *Thermal Comfort Wizard* dialog.
2. Make the following settings:

Property	**Setting**
Temperature limiter (deg C)	`0.1`
Humidity Model	`Relative humidity will be held constant at (%)` `34`
Number of STAR iterations before initial call to TIM (steady-state)	`2`
Number of STAR iterations after subsequent calls to TIM	`1`
Convergence tolerance, defined as ...	`0.000001`
Update passenger_tim_??.dat file?	`Yes`

The panel appears as follows:



Setting Thermal Properties

Define the properties of any thermal boundaries in the model, and specify the mass flow rate and temperature at the inlets.

Use the following steps:

1. Click the *External Convection and Radiation* tab at the top of the *Thermal Comfort Wizard* dialog.

   The panel is divided into separate sections that correspond to the different types of thermal boundary, as shown in the diagram.



   The thermal specification that was defined earlier for each boundary corresponds to a particular section of this panel.

External Convection Settings:

To specify the thermal properties for the Convection boundaries, values are entered under the External Convection section.

A number of settings, common to all convection boundaries are specified in the External Convection Global Settings section. The settings include: vehicle velocity, reference velocity, and scaling exponent. The settings allow heat transfer coefficients for a reference velocity to be applied to a model operating at the vehicle velocity.

The external bulk temperature is also specified globally. The remaining external convection settings are boundary-specific.

For the tutorial, the boundaries are defined as having convection only.

To specify the global external convection settings:

2. In the *External Convection Global Settings* section of the dialog, set the *Vehicle Velocity* to `40`.



3. Set the *Reference Velocity* to `35`.



4. Set the *Scaling Exponent* to `0.8`.



5. Set the *External Bulk Temperature* to `20`.



To specify the reference heat transfer coefficient for each of the convection boundaries:

6. Make sure that the **Region** is set to `Fluid`. Use the drop-down menu to set the **Boundary** to `Windows`.



By default the Boundary Type is set to `External Convection`.

7. In the *External Convection Boundary Conditions* section on the right-hand-side of the panel, set the *Reference Heat Transfer Coefficient* to `7.9`.



Settings are saved automatically for each boundary.

Select the next convection boundary and set the heat transfer coefficient.

8. Using the drop-down menu, set the **Boundary** to `Roof`.



9. In the *External Convection Settings* section, set the *Reference Heat Transfer Coefficient* to `44.5`.



10. The properties of all the convection boundaries are shown in the table below. Use the technique described above and data from the table to set the properties of the remaining convection boundaries. You can check the properties that have been set for each boundary at any time by selecting the appropriate boundary from the drop-down list.

Boundary	**Reference Heat Transfer Coefficient (W/m^2 K)**
Windows	7.9
Roof	44.5
Cabin	44.5
Doors	44.5
Console	22.3

Inlet Boundary Settings

Set the inlet mass flow rate and temperature of the fluid conditions in the Inlet Boundary Settings section.

11. Set the **Boundary** to `InletFeetLeft`.



12. Select the **Inlet** radio button under *Boundary Type* to activate the input boxes for the Inlet Boundary Settings.



13. Enter a *Mass Flow Rate* of `0.003` in the `Inlet Boundary Settings` section.
14. Set the *Temperature* to `7`



15. Using this technique and values from the following table, set the properties of the remaining inlet boundaries.

Boundary	Mass Flow Rate (kg/s)	Temperature (C)
InletFeetLeft	0.003	7
InletDashCentre	0.006	14
InletDashLeft	0.007	13
InletDashRight	0.007	13
InletFeetRight	0.003	7

Saving TCM Settings

When you are finished working with the Thermal Comfort wizard, save the settings in a step that is distinct from the step of saving the simulation file.

To save TCM settings:

1. Click **Save TCM Settings**, and click **Close**.



The information from the wizard is added to the managers in the object tree.

Setting Stopping Criteria

Specify the maximum number of iterations.

1. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `300`.

Creating a Scalar Scene

To show the surface temperature of the different body segments of the person, set up a scalar scene.

Use the following steps:

1. Create a scalar scene.
2. Edit the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** node.
3. In the *Edit* dialog, expand the **Regions** > **Fluid** node.
4. Select the body boundaries **1 - Head** boundary to **14 - RightFoot**, inclusive.
5. Click **OK**.
6. Select the **Scalar 1** > **Scalar Field** node.
7. Set *Function* to `Temperature`.
8. Set *Units* to `C`.
9. Delete the **Displayers** > **Outline 1** node.
10. Create a new geometry displayer.
11. Edit the **Displayers** > **Geometry 1** > **Parts** node.
12. In the *Edit* dialog, expand the **Regions** > **Fluid** > **Boundaries** node.
13. Select the following boundaries:
    - **Cabin**

- **Console**
- **Doors**
- **Roof**
- **Seats**
- **Windows**

14. Click **OK**.
15. Select the **Geometry 1** node and set the following properties:
    - Deactivate *Outline*.
    - Activate *Surface*.
    - Set *Opacity* to `0.3`.

Contents

Adjusting the Color Bar

Adjusting the Color Bar

To optimize your view of the scalar data, adjust the color bar in the scene.

To adjust the size and position of the color bar in *Scalar Scene 1*:

1. Select the **Displayers** > **Scalar 1** > **Color Bar** node and set *Orientation* to `Vertical`.
2. In the *Graphics* window, click the left-hand edge of the color bar and drag it to decrease the width.



3. Drag the color bar into the top right corner of the scene.



4. Drag the bottom edge of the color bar to increase the length.



5. Click (**Save-Restore-Select views**) and select **Restore View > View 1**.

   The scene appears as follows:



6. Save the simulation.

Running the Solver

The simulation is now ready to run.

To obtain a solution:

1. Run the simulation.
2. To watch the solution being updated while the simulation is running, look at the scalar scene.

Note: The simulation is run for 2 iterations before the TIM program starts and updates the body surface temperatures.

3. When the simulation has finished running, save it.

Visualizing the Results with Streamlines

The solution for the surface temperature of each body segment is shown in the scalar scene. Also show the flow path of the air particles that enter the cabin through the inlet vents.

To show the air flow path, set up a streamlines displayer:

1. Make the *Scalar Scene 1* display active.
2. Create a streamline derived part.
3. In the *Create Streamline* inset panel, make the following settings:

Property	**Setting**
Input Parts	`[Fluid]`
Vector Field	`Velocity`
Seed Parts:	`InletDashLeft`

4. Click **OK**.
5. Click **Create** and **Close**.
6. Select the **Scenes > Scalar Scene 1 > Displayers > Section Stream 1 > Scalar Field** node.
7. Set *Function* to `Velocity > Magnitude`.
8. In the *Graphics* window, click in the center of the color bar and drag it to a more suitable position.

The scene appears as shown below.



9. Save the simulation.

Summary

The Thermal Comfort Wizard tutorial used an analysis of a single occupant in the cabin of a car to introduce a variety of STAR-CCM+ features.

The features that were introduced include:

- Generating the volume mesh.
- Setting up the physics continuum.
- Preparing boundaries for the *Thermal Comfort Wizard*.
- Activating the *Thermal Comfort Model*.
- Activating the *Thermal Comfort Wizard*.
- Using the *Thermal Comfort Wizard* to:
  - Enter passenger properties;

  - Specify advanced settings;
  - Set the thermal properties of other boundaries in the model;
  - Save TCM settings.
- Creating a scalar scene and adjusting the color bar.
- Running the TCM model.
- Visualizing streamline results.

Multiphase Flow

The tutorials in this set illustrate various STAR-CCM+ features for simulating multiphase fluid flow problems

Contents

VOF: Gravity-Driven Flow
VOF: Capillary Effects
VOF: Cavitation
VOF: Boiling
VOF: Melting-Solidification
Lagrangian: Particle-Laden Flow
Lagrangian: Solid Particle Erosion
Eulerian: Dispersed Phase Acceleration
Eulerian: Hibiki's Bubble Column
Eulerian: Mixture Settling
Eulerian: Bubble Formation in a Fluidized Bed
Eulerian: Particle-Laden Flow in a Rotating Body
Eulerian: Degassing Boundary
Eulerian: Wall Boiling
Eulerian: Conjugate Heat Transfer Wall Boiling
Fluid Film: Liquid Film Flow
Fluid Film: Binary Liquid Film Flow with Evaporation and Edge Stripping
Dispersed Multiphase: Airfoil Icing

VOF: Gravity-Driven Flow

This tutorial demonstrates how to set up a gravity-driven flow problem in STAR-CCM+. It simulates two-dimensional gravity-driven compressible flow through a channel connecting two chambers.

The geometry is shown below. Initially, the chamber on the left is filled with water; the one on the right and the connecting channel with air. All boundaries are solid walls except for the horizontal top left surface where a constant (atmospheric) static pressure is applied.



Under the action of gravity, water flows into the right chamber under assumed turbulent conditions. At the same time, water also flows in through the top left boundary to maintain a constant fluid level. The pressure in the right chamber increases due to the air compression, resulting in a reduction of the flow rate thorough the pressure boundary. After some time, all fluid elements are at rest and in a hydro-static equilibrium.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Converting to a Two-Dimensional Mesh
Visualizing the Mesh Interior
Selecting the Physics Models
Setting the Material Properties
Setting Initial Conditions and Reference Values
Setting Boundary Conditions

Setting Solver Parameters and Stopping Criteria
Visualizing and Initializing the Solution
Running the Simulation
Visualizing Results
Animating the Results
Summary

Prerequisites

The instructions in the VOF: Gravity-Driven Flow tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Importing a volume mesh	Simple Pipe Flow
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A one-cell-thick, three-dimensional, polyhedral cell mesh has been prepared for this analysis.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/gridfs.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.
5. Save the simulation as `freeSurface.sim`.

Converting to a Two-Dimensional Mesh

You can now convert the three-dimensional mesh to two dimensions.

There are special requirements in STAR-CCM+ for converting three-dimensional meshes into two-dimensional meshes. These requirements are:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The mesh that is imported for this tutorial is built with these requirements in mind. Were the grid not to conform to the above conditions, it would have been necessary to realign the region using the transformation and rotation facilities in STAR-CCM+.

To convert the mesh to two dimensions:

1. Select **Mesh > Convert to 2D**.

   The *Convert Regions To 2D* dialog appears.
2. Make sure that the **Delete 3D regions after conversion** option is activated.



3. Click **OK**.

   Once you click **OK**, the mesh conversion takes place and the new two-dimensional mesh is shown, viewed from the z-direction, in the *Geometry Scene 1* display. The mouse rotation option is suppressed for two-dimensional scenes.
4. Right-click the **Continua** > **Physics 1** node and select **Delete**.
5. In the confirmation dialog, click **Yes**.

Visualizing the Mesh Interior

After you have converted the volume mesh to a two-dimensional mesh, you can view the mesh structure.

To visualize the mesh:

1. In the simulation tree, select the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** node.
2. To display the surface mesh, activate the *Mesh* property.

The interior mesh structure is shown below.



Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this example, the flow is turbulent. Use the default K-Epsilon Turbulence model, and apply a gravitational force in the -y direction. As the problem also involves multiphase flow, two fluids (air and water) are required for the analysis. However, since these fluids occupy the same domain, only one continuum and one mesh region are required to set up the simulation. By default, a continuum that is called **Physics 1 2D** is created when the mesh is converted to two-dimensional.

To select the physics models:

1. Rename the **Physics 1 2D** node to **Chambers**.
2. For the physics continuum, **Continua** > **Chambers**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected) **Gradients** (Pre-selected)
Time	**Implicit Unsteady**

Group Box	Model
Material	**Eulerian Multiphase**
Eulerian Multiphase Model	**Volume of Fluid (VOF)**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Optional Models	**Gravity** **Segregated Fluid Isothermal**

3. Click **Close**.
4. To review the models, open the **Chambers** > **Models** node.



5. Save the simulation.

Setting the Material Properties

Define the materials corresponding to each of the mixture components under the **Eulerian Phases** node.

To define the phases and set the material properties:

1. In the **Chambers** continuum, select the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **H2O**.
3. For the **H2O** phase, select the following models:

Group Box	Model
Material	**Liquid**
Equation of State	**Constant Density**

4. Click **Close**.
5. Create a second phase and rename it to **Air**.
6. For the **Air** phase, select the following models:

Group Box	Model
Material	**Gas**
Equation of State	**Ideal Gas**

7. Click **Close**.

   The **Eulerian Multiphase** node contains the phase model objects as shown below:



The default selections for materials—water and air—are appropriate for this exercise, so there is no need to replace them with other materials from the STAR-CCM+ material database. The existing material

properties for air are suitable but the dynamic viscosity of water must be changed.

8. Select the **H2O** > **Models** > **Liquid** > **H2O** > **Material Properties** > **Dynamic Viscosity** > **Constant** node and change the dynamic viscosity **Value** to `0.001002 Pa-s`.



9. Save the simulation.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the physics continuum. You define a field function to use for specifying the initial condition for the spatial distribution of the two fluids in the continuum.

The usual practice for initially stationary fluids is to specify low, non-zero values for the turbulence parameters. If these quantities are set to zero, numerical problems can arise.

The initial condition for the spatial distribution of the two fluids in the continuum is that you only have water in the left-hand chamber and only air in the right-hand chamber and channel. Both fluids are also stationary. A convenient way of specifying such a distribution is by creating and using field functions.

The direction and magnitude of the gravity vector are set using the **Reference Values** node. In this case, a gravity force is applied in the negative y-direction.

To set the initial conditions and reference values:

Define the field function to use for specifying the initial condition for the spatial distribution of the two fluids in the continuum:

1. Right-click the **Tools** > **Field Functions** node and select **New** > **Scalar** from the pop-up menu.
2. Rename the **User Field Function 1** to `Initial Distribution`.
3. Select the **Initial Distribution** node and set the following properties:

Property	Setting
Function Name	`Initial Distribution`
Definition	`($$Position[0]<=-1)?1:0`

Set the initial conditions and reference values for the **Chambers** continuum.

4. Edit the **Initial Conditions** node and set the following properties:

Node	Property	Setting
Turbulence Intensity > Constant	**Value**	`0.01`
Turbulent Viscosity Ratio > Constant	**Value**	`100.0`
Volume Fraction	**Method**	**Composite**
Composite > H2O	**Method**	**Field Function**
H2O > Field Function	**Scalar Function**	**Initial Distribution**



It is not necessary to define a similar distribution for the Air node because any user-specified initial condition for the volume fraction of the last-defined phase in multiphase flows is always ignored. Instead, this condition is obtained by subtracting the sum of the volume fraction distributions of the other phases from 1.0.

5. Edit the **Reference Values** node and set the following properties:

Node	Property	Setting
Gravity	**Value**	`0,-9.81,0 m/s^2`



6. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values. The geometry that is used for this tutorial has six boundaries. Assign no-slip wall conditions to five of them, and assign the remaining boundary a pressure boundary condition.

To set the boundary conditions:

1. Rename the **Default_Fluid 2D** node to **Fluid**.
2. Edit the **Fluid** > **Boundaries** node and set the boundary types as follows:

Boundary	Type
Bottom	**Wall**
Left	**Wall**
Middle	**Wall**
Right	**Wall**
TopRight	**Wall**
TopLeft	**Pressure Outlet**

The default conditions are suitable for the wall boundaries, but the conditions for the pressure boundary must be reset.

3. Edit the **TopLeft** > **Physics Values** node and set the following properties:

Node	Property	Setting
Turbulence Intensity > Constant	**Value**	`0.01`
Turbulent Viscosity Ratio > Constant	**Value**	`100.0`
Volume Fraction > Constant	**H2O**	`1.0`
	Air	`0.0`



The **Volume Fraction** setting enforces the condition that only liquid water enters the solution domain through that boundary.

4. Save the simulation.

Setting Solver Parameters and Stopping Criteria

As you are solving an unsteady problem, it is necessary to specify the time-step size and the elapsed simulation time. Run this calculation for 5.0 s with a time-step size of 0.005 s.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following property:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.005 s`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated

Node	Property	Setting
	Maximum Physical Time	`5.0 s`
Maximum Steps	**Enabled**	Deactivate
Maximum Inner Iterations	**Maximum Inner Iterations**	`5`

3. Save the simulation.

Visualizing and Initializing the Solution

You create a scalar scene to view the results of the simulation. You can view the air and water distributions throughout the run.

To visualize the results of the simulation:

1. Right-click the **Scenes** node and select **New Scene > Scalar**.
   The *Scalar Scene 1* display appears.
2. Select the new **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Volume Fraction > Volume Fraction of H2O**.



3. Select the **Scalar 1** node and set *Contour Style* to **Smooth Filled**.
4. Select the **Attributes > Update** node and set *Trigger* to **Time-Step**.

5. Activate **Save To File**.
6. Specify the *Base Filename*, *Offscreen Width*, and *Offscreen Height* properties as desired.



7. Select the **Attributes** > **Update** > **Time-Step Frequency** node and set *Time-step Frequency* to `10`.



8. Click (**Initialize Solution**) in the top toolbar or use the **Solution** > **Initialize Solution** menu item.
9. Make the *Scalar Scene 1* display active to look at the initialization result.



At the start of the run, the left chamber is entirely filled with water whereas the right chamber and the connecting channel are entirely filled with air. A small region in which both fluids are apparently present is visible at the interface between the two fluids, but this effect is due to the coarseness of the mesh.

10. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers. You can also click the *Scalar Scene 1* tab at the top of the *Graphics* window to view the simulation progress.



During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 1000 time-steps are complete. The numbers that are displayed in the *Output* window represent the inner iterations, not the time-steps that the solver performs. As there are 5 inner iterations per time-step, you expect the run to perform up to 5,000 inner iterations.

When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Visualizing Results

View the results of the simulation.

The *Scalar Scene 1* display shows the water volume fraction profile at the end of the 5.0 s run.



To display the pressure profile:

1. Right-click the scalar bar in the display.
2. Select **Pressure** from the pop-up menu.

   The following screenshot shows the compression of the air in the right chamber.



3. Save the simulation.

Animating the Results

You can use the animation feature in STAR-CCM+ to animate the results.

In the Visualizing and Initializing the Solution section, you instructed STAR-CCM+ to save the water volume fraction profile every 10 time-steps. The resulting series of `png` files can be used to produce an animation using an external package, such as the `animate` tool included with most Linux and Unix systems. An example of such an animation is shown below.



You can also modify the simulation and use the animation feature in STAR-CCM+. See the following sections in the tutorial Solution Recording and Playback: Vortex Shedding:

- Setting Up the Solution History
- Creating a Recorded Solution View
- Creating an Animation from the Solution View

Summary

This tutorial demonstrated how to set up a gravity-driven flow problem in STAR-CCM+.

This tutorial introduced the following STAR-CCM+ features:

- Defining models for multiphase flow problems.
- Defining the material properties that are required for the selected models.
- Setting initial conditions and reference values.
- Creating and applying field functions.

- Displaying the results as an animation.

VOF: Capillary Effects

This tutorial demonstrates how to set up a capillary effects problem in STAR-CCM+. It simulates two-dimensional forced flow of liquid glycerine through a nozzle and into an air-filled chamber at atmospheric pressure.

The mesh that is used (shown below) is the same as in VOF: Gravity-Driven Flow except that, in this instance, the physical dimensions of the problem are scaled down by a factor of 1,000. Therefore, the nozzle width is about 1 mm. The boundary on the left of the problem geometry is an inlet with fluid velocity of 1 mm/s and the boundary on the right is at atmospheric pressure. The boundary at the bottom is a symmetry plane and all other boundaries are solid walls.

Initially, the left chamber is filled with liquid and the remainder of the solution domain is filled with air. For the given geometry and inlet velocity, the flow can be assumed to be laminar. Gravity acts in the positive x-direction and so helps to drive the flow through the nozzle. The shape of the free-surface that develops in the chamber downstream of the nozzle depends on the contact angle between liquid and wall. This contact angle is specified as 45º.



Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Converting to a Two-Dimensional Mesh
Scaling the Mesh
Selecting the Physics Models
Setting the Material Properties
Defining the Phase Interactions

Setting Initial Conditions and Reference Values
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Visualizing and Initializing the Solution
Running the Simulation
Visualizing Results
Effects of Changing the Contact Angle
Summary

Prerequisites

The instructions in the VOF: Capillary Effects tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A one-cell-thick three-dimensional polyhedral cell mesh has been predefined for this analysis.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/gridfs.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `capillaryEffects.sim`.

Converting to a Two-Dimensional Mesh

The mesh can now be converted from three dimensions to two dimensions.

There are special requirements in STAR-CCM+ for three-dimensional meshes that are to be converted to two-dimensional:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The mesh that is imported for this tutorial was built with these requirements in mind. Were the grid not to conform to the above conditions, it would have been necessary to realign the region using the transformation and rotation facilities in STAR-CCM+.

To convert the mesh to two dimensions:

1. Select **Mesh > Convert to 2D...**.
2. In the *Convert Regions to 2D* dialog, activate **Delete 3D regions after conversion**, and click **OK**.



The mesh conversion takes place and the two-dimensional mesh is created.

Note: The mouse rotation option is suppressed for two-dimensional scenes.

3. Right-click the **Physics 1** continuum node and select **Delete**.
4. Click **Yes** in the confirmation dialog.

Scaling the Mesh

The original mesh was not built to the correct scale and therefore requires scaling down by a factor of 1000.

To scale the mesh:

1. Select **Mesh** > **Scale Mesh** from the menu bar.
2. In the *Scale Mesh* dialog, make sure that the **Default_Fluid 2D** region is selected.



3. Enter a value of `0.001` for the scale factor.
4. Click **Apply** to scale the domain, and then click **Close**.

After clicking **Apply**, the mesh region reduces in size.

To restore the previous viewing distance for the scaled domain in the display:

5. Click the (**Reset View**) button in the visualization toolbar.

To verify that scaling has been applied, you can perform the mesh diagnostics check and review the output values.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this simulation, the flow is laminar, and a gravity force is applied in the +x direction. As the problem also involves multiphase flow, two fluids (air and glycerine) are required for the analysis. However, since these fluids occupy the same domain, only one continuum and one region are required to set up the simulation. By default, a continuum that is called **Physics 1 2D** is created when the mesh is converted to two-dimensional.

To select the physics models:

1. Rename the **Physics 1 2D** node to **Injector**.
2. For the physics continuum, **Continua** > **Injector**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase**
Eulerian Multiphase Model	**Volume of Fluid (VOF)**
Viscous Regime	**Laminar**
Optional Models	**Gravity**

3. Click **Close**.
4. To review the models, open the **Injector** > **Models** node.



5. Save the simulation.

Setting the Material Properties

Define the phases and set the appropriate material properties.

Define the materials corresponding to each of the mixture components under the **Eulerian Phases** node.

Note: In multiphase flow problems that include surface tension effects, define the densest fluid first because contact angle values at the wall are always specified relative to the first-defined fluid.

To define the phases and set the material properties:

1. In the **Injector** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **C3H8O3**.
3. For the **C3H8O3** phase, select the following models:

Group Box	Model
Material	**Liquid**
Equation of State	**Constant Density**

4. Click **Close**.

5. Create a second phase and rename it to **Air**.
6. For the **Air** phase, select the following models:

Group Box	Model
Material	**Gas**
Equation of State	**Constant Density**

7. Click **Close**.

   The **Eulerian Multiphase** node contains the phase model objects as shown below:



The default selection of air is appropriate for this exercise, but water must be replaced with glycerine from the STAR-CCM+ material database.

8. Right-click the **C3H8O3 > Models > Liquid > H2O** node and select **Replace with**.



9. In the *Replace Material* dialog, expand the **Material Databases** > **Standard** > **Liquids** node.
10. Activate **C3H8O3 (Glycerine)**.



11. Click **OK** to replace the material.

    The phase model objects appear as shown below.



12. Save the simulation.

Defining the Phase Interactions

Using the **Multiphase Interaction** model, you can define the interaction between the air and liquid phases.

In this tutorial, you are interested in the surface tension of glycerine. Therefore, use the Surface Tension Force and VOF multiphase interaction models:

To define the phase interactions:

1. In the **Injector** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. For the **Phase Interaction 1** phase interaction, select the following models in order:

Group Box	Model
Phase Interaction Topology	**VOF Phase Interaction**
Optional Models	**Surface Tension Force** **Multiphase Material** (selected automatically)

3. Click **Close**.

The value of the surface tension coefficient is changed for the phase interaction. Set the primary and secondary phases for the phase interaction.

4. Edit the **Phase Interaction 1 > Models** node and set the following properties:

Node	Property	Setting
Multiphase Material > Multiphase Material > Material Properties > Surface Tension > Constant	**Value**	`0.059688 N/m`
VOF Phase Interaction	**Primary Phase**	**C3H8O3**
	Secondary Phase	**Air**

5. Save the simulation.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the physics continuum.

The initial condition for the spatial distribution of the two fluids in the continuum is that you only have glycerine in the left-hand chamber and only air in the right-hand chamber and channel. A convenient way of specifying such a distribution is by creating and using a field function.

The direction and magnitude of the gravity vector is set using the Reference Values node. In this case, a gravity force is applied in the positive x-direction.

To set the initial conditions and reference values:

Define the field function to use for specifying the initial condition for the spatial distribution of the two fluids in the continuum:

1. Right-click the **Tools > Field Functions** node and select **New > Scalar** from the pop-up menu.
2. Rename the **User Field Function 1** to `Initial Distribution`.
3. Select the **Initial Distribution** node and set the following properties:

Property	Setting
Function Name	`Initial Distribution`
Definition	`($$Position[0]>=-0.001)? 0:1`

Set the initial conditions and reference values for the **Injector** continuum.

4. Edit the **Initial Conditions** node and set the following properties:

Node	Property	Setting
Volume Fraction	**Method**	**Composite**
Composite > C3H8O3	**Method**	**Field Function**
C3H8O3 > Field Function	**Scalar Function**	**Initial Distribution**


Servers
capillaryEffects
simulation
scene/plot
Continua
Injector
Models
Reference Values
Initial Conditions
Pressure
Constant
Velocity
Constant
Volume Fraction
Composite
C3H8O3
Constant
C3H8O3 - Properties
Properties
Method
Constant
Dimensions
Constant
Field Function
Table (iteration)
Table (r)
Table (time)
Table (x,y,z)
Table (xyz,time)
User Code



It is not necessary to specify a similar distribution for the Air node because any user-specified initial condition for the volume fraction of the last-defined phase in multiphase flows is always ignored. Instead, this condition is obtained by subtracting the sum of the volume fraction distributions of the other phases from 1.0.

5. Edit the **Reference Values** node and set the following properties:

Node	Property	Setting
Gravity	**Value**	`9.81,0,0 m/s^2`

6. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

To set the boundary conditions:

1. Rename the **Default_Fluid 2D** node to **Fluid**.

Set the boundary types.

2. Edit the **Fluid** > **Boundaries** node and set the boundary types as follows:

Boundary	Type
Bottom	**Symmetry Plane**
Left	**Velocity Inlet**
Middle	**Wall**
Right	**Pressure Outlet**
TopRight	**Wall**
TopLeft	**Wall**

The only condition that is specified for the wall boundaries is the contact angle of the glycerine-air interface.

3. Edit the **Middle** > **Phase Conditions** > **Phase Interaction 1** > **Physics Values** node and set the following property:

Node	Property	Setting
Contact Angle > **Constant**	**Value**	`45 deg`



4. Repeat the previous step for the other two wall boundaries: **TopLeft** and **TopRight**.

Set the boundary conditions for the inlet and pressure boundaries.

5. Edit the **Left** > **Physics Values** node and set the following properties:

Node	Property	Setting
Velocity Magnitude > Constant	**Value**	`0.001 m/s`
Volume Fraction > Constant	**C3H8O3**	`1.0`
	Air	`0.0`

   The **Volume Fraction** setting enforces the condition that only glycerine enters the solution domain through that boundary.

6. Edit the **Right** > **Physics Values** node and set the following properties:

Node	Property	Setting
Volume Fraction > Constant	**C3H8O3**	`0.0`
	Air	`1.0`

   This setting enforces the condition that only air enters the solution domain through that boundary.

7. Save the simulation.

Setting Solver Parameters and Stopping Criteria

As you are solving an unsteady problem, it is necessary to specify the time-step size and the elapsed simulation time.

Run this calculation for 2.0 s with a time-step size of 0.001 s, so 2,000 time-steps are required.

To set the solver parameters and the stopping criteria:

1. Edit the **Solvers** node and set the following property:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.001 s`



2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Steps	**Enabled**	Activated
	Maximum Steps	2000
Maximum Physical Time	**Enabled**	Deactivate



3. Save the simulation.

Visualizing and Initializing the Solution

Use a scalar scene to view the air and glycerine distributions throughout the run.

To set up the scalar scene:

1. Right-click the **Scenes** node and then select **New Scene > Scalar**.
2. In the *Graphics* window, right-click the scalar bar and select **Volume Fraction > C3H8O3**.



3. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node.
4. Change the *Contour Style* property setting to **Smooth Filled**.
5. Click (**Initialize Solution**) in the top toolbar or use the **Solution** > **Initialize Solution** menu option.
6. Activate the *Scalar Scene 1* display to look at the initialization results.



As expected, at the start of the run the left chamber is entirely filled with glycerine whereas the right chamber and the connecting channel are entirely filled with air. A small region in which both fluids are apparently present is visible at the interface between the two, but this effect is due to the coarseness of the mesh.

7. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.



You can click the *Scalar Scene 1* tab at the top of the *Graphics* window to view the simulation progress.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 2000 time-steps are complete. The numbers that are displayed in the *Output* window represent the inner iterations of the solver, not time-steps. As there are 5 inner iterations per time-step, you expect the run to perform 10,000 inner iterations.

When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Steps satisfied.
```

2. When the simulation has finished running, save it.

Visualizing Results

View the results of the simulation.

The *Scalar Scene 1* display shows the glycerine volume fraction profile at the end of the 2.0 s run. The observable contact angle is indeed close to 45°.



To view the results of the simulation:

1. Right-click the scalar bar in the display.
2. Select **Pressure** from the pop-up menu to display the pressure profile.



To display velocity vectors:

3. Right-click the **Scenes** node and select **New Scene** > **Vector**.

   The vector scene appears in the *Graphics* window.



The velocity vector plot shows high air velocities close to the free surface. These numerical inaccuracies are known as parasitic currents. These currents arise because the surface tension and pressure forces are much larger than all other terms in the momentum equations, and their balance on an irregular grid is difficult to achieve numerically due to the discontinuous variation of pressure across the free surface and a large discretization error that is associated with it. Parasitic currents become appreciable when the problem size is small and fluid velocity and viscosity are low. For flows where diffusion and convection forces are of a similar magnitude to surface tension forces, these problems are not so pronounced. Since artificial velocities are generated only within the air, their effect on the liquid flow (which is usually what you are trying to predict) is small.

4. Save the simulation.

Effects of Changing the Contact Angle

Consider the effects of changing the contact angle.

The contours of glycerine volume fraction shown below demonstrate the effect of changing the contact angle. In this case, the contact angle at all wall boundaries was changed from 45º to 135º. All other modeling options and material properties were kept the same and the analysis is run for a physical time of 2.0 s, as before.



A vector scene similar to the one shown below is also displayed.



Summary

This tutorial demonstrated how to set up a capillary effects problem in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Defining models for multiphase flow problems.
- Defining the material properties that are required for the selected models.
- Setting initial conditions and reference values.
- Creating and applying field functions.
- Defining boundary conditions.
- Setting solver parameters for an unsteady run.
- Initializing and running the solver for a given number of time-steps.

- Analyzing results using the built-in visualization facilities.
- Examining the effect of changing the contact angle value.

VOF: Cavitation

This tutorial demonstrates how to set up a cavitation problem in STAR-CCM+. It simulates a two-dimensional forced flow of water through a nozzle and into an air-filled chamber at atmospheric pressure.

The mesh that is used (shown below) is the same as in VOF: Gravity-Driven Flow except that, in this instance, the physical dimensions of the problem are scaled down by a factor of 1,000. Therefore, the nozzle width is about 1 mm. The boundary on the left of the problem geometry has a prescribed pressure of 5×10^7 Pa and the boundary on the right is at atmospheric pressure. The boundary at the bottom is a symmetry plane and all other boundaries are solid walls.

Initially, the left chamber is filled with water and the remainder of the solution domain is filled with air. Turbulence is modeled using the default K-Epsilon model. The sharp corner at the entrance to the narrow nozzle combined with the flow acceleration produces a low-pressure region in which cavitation takes place.



Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Converting to a Two-Dimensional Mesh
Scaling the Mesh
Selecting the Physics Models
Setting the Material Properties
Defining the Phase Interactions
Setting the Cavitation Model Parameters
Setting Initial Conditions

Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Visualizing and Initializing the Solution
Running the Simulation
Visualizing Results
Summary

Prerequisites

The instructions in the VOF: Cavitation tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A one-cell-thick three-dimensional polyhedral cell mesh has been predefined for this analysis.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/gridfs.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.
5. Save the simulation as `cavitation.sim`.

Converting to a Two-Dimensional Mesh

Convert the three-dimensional mesh to two dimensions.

There are special requirements in STAR-CCM+ for converting three-dimensional meshes to two-dimensional meshes:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The mesh that is imported for this tutorial is built with these requirements in mind. Had the grid not met them, it would have been necessary to realign the region using the mesh transformation and rotation capabilities in STAR-CCM+.

To convert the mesh region to two dimensions:

1. Select **Mesh > Convert to 2D...**.
2. In the *Convert Regions to 2D* dialog that appears, make sure that **Delete 3D regions after conversion** is activated, and click **OK**.



Once you click **OK**, the mesh conversion takes place and the new two-dimensional mesh is shown, viewed from the z-direction, in the *Geometry Scene 1* display.

Note: The mouse rotation option is suppressed for two-dimensional scenes.

Delete the three-dimensional physics continuum:

3. Right-click the **Physics 1** continuum node and select **Delete**.
4. Click **Yes** in the confirmation dialog.

Scaling the Mesh

The original mesh is not built to the correct scale and therefore requires scaling down by a factor of 1000.

To scale the mesh:

1. Select **Mesh > Scale Mesh** from the menu bar.
2. In the *Scale Mesh* dialog, make sure that the **Default_Fluid 2D** region is selected.
3. Enter a value of `0.001` for the scale factor.



4. Click **Apply** to scale the domain, and then click **Close**.
   The mesh region reduces in size.

To restore the previous viewing distance for the scaled domain in the display:

5. Click (**Reset View**) in the visualization toolbar.

To verify that scaling has in fact been applied, you can perform the mesh diagnostics check and review the output values.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this simulation, the flow is turbulent and the problem involves multiphase flow and cavitation. Three fluids (air, water, and water vapor) are required for the analysis. However, since these fluids occupy the same domain, only one continuum and one region are required to

set up the simulation. By default, a continuum that is called **Physics 1 2D** is created when the mesh is converted to two-dimensional.

To select the physics models:

1. Rename the **Physics 1 2D** continuum node to **Injector**.
2. For the physics continuum, **Continua** > **Injector**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected) **Gradients** (Pre-selected)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase**
Eulerian Multiphase Model	**Volume of Fluid (VOF)**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**

3. Click **Close**.
4. To review the models, open the **Injector** > **Models** node.



5. Save the simulation.

Setting the Material Properties

You define the phases and set the appropriate material properties.

Define the materials corresponding to each of the mixture components under the **Eulerian Phases** node.

To define the phases and set the material properties:

1. Right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **H2O**.
3. For the **H2O** phase, select the following models:

Group Box	Model
Material	**Liquid**
Equation of State	**Constant Density**

4. Click **Close**.
5. Create a second phase and rename it to **Air**.
6. For the **Air** phase, select the following models:

Group Box	Model
Material	**Gas**
Equation of State	**Constant Density**

7. Click **Close**.
8. Create a third phase and rename it to **H2O (G)**.
9. For the **H2O (G)** phase, select the following models:

Group Box	Model
Material	**Gas**
Equation of State	**Constant Density**

10. Click **Close**.

    The **Eulerian Multiphase** node contains the phase model objects as shown below:



If you were setting up a standard free-surface problem, only two phases would be required. However, because this problem involves cavitation, it is necessary to define a third phase for the product of cavitation, that is, water vapor. Hence, the air in the third phase is replaced with water vapor from the STAR-CCM+ material database.

11. Right-click the **H2O (G)** > **Models** > **Gas** > **Air** node and select **Replace with**.
12. In the *Replace Material* dialog, expand the **Material Databases** > **Standard** > **Gases** node.
13. Activate **H2O (Water)**.



14. Click **OK** to replace the material.

    The phase model objects appear as shown below.



The default material properties for the three-phases are suitable for this problem.

15. Save the simulation.

Defining the Phase Interactions

Using the **Multiphase Interaction** model, you can define the interaction between the gas and liquid phases.

To model cavitation, use the **VOF Multiphase Interaction** and **Cavitation** models.

To define the phase interactions:

1. In the **Injector** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. For the **Phase Interaction 1** phase interaction, select the following models in order:

Group Box	Model
Phase Interaction Topology	**VOF Phase Interaction**
Optional Models	**Multiphase Material**

3. Click **Close**.

Before you can select the Cavitation model, you set the primary and secondary phases for the phase interaction.

4. Edit the **Phase Interaction 1** > **Models** node and set the following properties:

Node	Property	Setting
VOF Phase Interaction	**Primary Phase**	**H2O**
	Secondary Phase	**H2O(G)**



You can now select the **Cavitation** model.

5. For the **Phase Interaction 1** phase interaction, select the following model:

Group Box	Model
Optional Models	**Cavitation**



6. Click **Close**.
7. Save the simulation.

Setting the Cavitation Model Parameters

Define the appropriate parameters for the Cavitation model. The saturation pressure is set on the relevant material property node for the liquid phase.

To set the Cavitation model parameters:

1. For the physics continuum, **Continua** > **Injector**, select the **Models** > **Eulerian Multiphase** > **Eulerian Phases** > **H2O** > **Models** > **Liquid** > **H2O** > **Material Properties Saturation Pressure** > **Constant** node and set *Value* to `2338.0` Pa.

2. Save the simulation.

Setting Initial Conditions

You set the initial conditions for the physics continuum. The initial condition for the spatial distribution of the two primary fluids in the continuum is as follows: you have only water in the left-hand chamber and only air in the right-hand chamber and channel. A convenient way of specifying such a distribution is by creating and using field functions.

To set the initial conditions:

Define the field function to use for specifying the initial condition for the spatial distribution of the two fluids in the continuum:

1. Right-click the **Tools > Field Functions** node and select **New > Scalar** from the pop-up menu.
2. Rename the **User Field Function 1** to `Initial Distribution (Air)`.
3. Select the **Initial Distribution (Air)** node and set the following properties:

Property	Setting
Function Name	`Initial Distribution (Air)`
Definition	`($$Position[0] >= -0.001) ? 1 : 0`

Set the initial conditions for the **Injector** continuum.

4. Edit the **Initial Conditions** node and set the following properties:

Node	Property	Setting
Volume Fraction	**Method**	**Composite**
Composite > Air	**Method**	**Field Function**
Air > Field Function	**Scalar Function**	**Initial Distribution (Air)**



5. Repeat the steps that are outlined above to specify the initial distribution of the H2O phase using a second field function named Initial Distribution (H2O) with the following definition:

   ```
   ($$Position[0] >= -0.001) ? 0 : 1
   ```

   Since these two field functions are the inverse of each other, the water vapor concentration is zero everywhere. This value is appropriate in this case as there is no cavitation in the stationary fluids.
6. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values. The geometry that is used for this tutorial has six boundaries, three of which have no-slip wall conditions that are assigned to them. Of the remaining three boundaries, define two as pressure outlet boundaries and define the last one as a symmetry plane.

To set the boundary conditions:

1. Rename the **Default_Fluid 2D** node to **Fluid**.
2. Edit the **Fluid** > **Boundaries** node and set the boundary types as follows:

Boundary	Type
Bottom	**Symmetry Plane**
Left	**Pressure Outlet**
Middle	**Wall**
Right	**Pressure Outlet**
TopRight	**Wall**
TopLeft	**Wall**

The only boundary conditions that are specified are at the two pressure boundaries.

3. Edit the **Left > Physics Values** node and set the following properties:

Node	Property	Setting
Pressure > Constant	**Value**	`5.0E7 Pa`
Volume Fraction > Constant	**H2O**	`1.0`
	Air	`0.0`
	H2O (G)	`0.0`





The **Volume Fraction** setting enforces the condition that only liquid water enters the solution domain through that boundary.

4. Edit the **Right > Physics Values** node and set the following properties:

Node	Property	Setting
Volume Fraction > Constant	**H2O**	`0.0`
	Air	`1.0`
	H2O (G)	`0.0`

This setting enforces the condition that only air enters the solution domain through that boundary.

5. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation. As you are solving an unsteady problem, it is necessary to specify the time-step size and the elapsed simulation time.

This calculation runs for 5×10^{-5} s with a time-step size of 10^{-7} s, so requires 500 time-steps.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following property:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`1e-7 s`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated

Node	Property	Setting
	Maximum Physical Time	`5e-5 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`5`

3. Save the simulation.

Visualizing and Initializing the Solution

View the distributions of the three-phases throughout the run by creating a scalar scene.

To visualize and initialize the solution:

1. Right-click the **Scenes** node and then select **New Scene > Scalar**.

   The *Scalar Scene 1* display appears.



2. In the *Graphics* window, right-click the scalar bar and select **Volume Fraction > H2O**.



3. Select the new **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node.
4. Set *Contour Style* to **Smooth Filled**.
5. Click (**Initialize Solution**) in the top toolbar or use the **Solution** > **Initialize Solution** menu item.
6. Make the *Scalar Scene 1* display active to check the initialization results.



As expected, at the start of the run the left chamber is entirely filled with water. Changing the scalar function to **Volume Fraction of Air** shows that the right chamber and the connecting channel are entirely filled with air. A small region in which both fluids are apparently present is visible at the interface between the two fluids, but this effect is due to the coarseness of the mesh.

7. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   You can click the *Scalar Scene 1* tab at the top of the *Graphics* window to view the simulation progress. Examine the distributions of each of the three-phases by changing the scalar function regularly.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 500 time-steps are complete. The numbers that are displayed in the *Output* window represent the inner iterations of the solver, not time-steps. As there are 5 inner iterations per time-step, you expect the run to perform 2,500 inner iterations.

   When the last iteration is complete, the following message is displayed in the *Output* window:

   ```
   Stopping criterion Maximum Steps satisfied.
   ```

2. When the simulation has finished running, save it.

Visualizing Results

You can view the results of the simulation.

Examine the final distributions of water and water vapor in the *Scalar Scene 1*, as shown below.





As expected, cavitation occurs at the nozzle entrance as a result of the low-pressure region that the sharp corner and the accelerating flow produce.

You can create a contour plot that approximately shows both the water and water vapor distributions by creating and plotting a field function.

To visualize the results:

Create the field function to plot.

1. Open the **Tools** node and then right-click the **Field Functions** node and select **New > Scalar**.
2. Rename the **User Field Function 1** node to `Volume Fraction (All Phases)`.
3. Select the **Volume Fraction (All Phases)** node and do the following:
   a) Change the *Function Name* to `Volume Fraction (All Phases)`.
   b) To open the custom editor, click the *Definition* box.
   c) In the editor window, enter the following syntax:
   `$VolumeFractionH2O+2*${VolumeFractionH2O (G)}`
   d) To exit from the editor, click **OK**.

4. In the *Scalar Scene 1* display, right-click the color bar and select **Volume Fraction (All Phases)**.

For this plotting technique to work, fix the legend scale to from 0 through 2.

5. Select the **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node and set *Max* to 2

   The contour plot appears as shown below.



To a reasonable approximation, the blue and turquoise areas are occupied predominantly by air, the green areas predominantly by liquid water and the yellow, orange, and red areas predominantly by water vapor.

Manipulating the results in this way provides a qualitative visual estimate of the spatial extent of each phase. However, the resulting plots are prone to inaccuracies in those places where air and water vapor occur together. Inaccuracies also occur where the liquid water volume fraction is only slightly greater than the water vapor volume fraction. In a case such as this, the latter leads to an over prediction of the size of the region of high water vapor concentration.

To display velocity vectors:

6. Right-click the **Scenes** node and select **New Scene** > **Vector**.

   The vector scene appears in the *Graphics* window.



7. Save the simulation.

Summary

This tutorial demonstrated how to set up a cavitation problem in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Defining a model for cavitating flow problems.
- Defining material properties that the model requires.
- Defining cavitation model parameters.
- Setting initial conditions.
- Creating and applying field functions.
- Defining boundary conditions.
- Setting solver parameters for an unsteady run.

- Initializing and running the solver for a given number of time-steps.
- Analyzing results.

VOF: Boiling

This tutorial demonstrates how to set up a boiling problem in STAR-CCM+. It simulates water boiling as it flows over a heated surface.

The purpose of this tutorial is to show the typical workflow in setting up and running a boiling simulation with phase change. One of the objectives is to show the impact of the boiling model parameters on the amount of heat transfer obtained. The mesh that is used is the same as in VOF: Gravity-Driven Flow except that, in this instance, you scale down the physical dimensions of the problem so that the nozzle width is about 50 mm.



Water enters the computational domain from the boundary on the left (with a prescribed velocity of 1 m/s and a temperature of 350 K). Water leaves from the boundary on the right (with a prescribed temperature of 370 K at atmospheric pressure). The boundary at the bottom is assumed to be a wall with a prescribed fixed temperature of 540 K. All other boundaries are assumed to be adiabatic solid walls.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Converting to a Two-Dimensional Mesh
Scaling the Mesh

Selecting the Physics Models
Setting the Material Properties
Defining the Phase Interactions
Defining the Boiling Model Parameters
Setting Initial Conditions
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Reporting, Monitoring, and Plotting
Visualizing the Solution
Running the Simulation
Visualizing Results
Modifying the Boiling Model Parameters
Running the Simulation
Visualizing Results
Summary
Bibliography

Prerequisites

The instructions in the VOF: Boiling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A one-cell-thick three-dimensional polyhedral cell mesh has been predefined for this analysis.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.

2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/gridfs.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `vofBoiling.sim`.

Converting to a Two-Dimensional Mesh

You convert the mesh from a three-dimensional mesh to two-dimensions.

There are special requirements in STAR-CCM+ to convert a three-dimensional mesh to two-dimensions. These requirements are:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The mesh that is imported for this tutorial is built with these requirements in mind. Had the grid not met them, it would have been necessary to realign the region using the mesh transformation and rotation capabilities in STAR-CCM+.

To convert the mesh to two-dimensions:

1. Right-click the **Scenes** node and then select **New Scene** > **Mesh**.
2. Select **Mesh** > **Convert to 2D**.
3. In the *Convert Regions To 2D* dialog that appears, do the following:
   a) Ensure that the **Default_Fluid** region is selected.
   b) Activate **Delete 3D regions after conversion**.
   c) Click **OK**.

The mesh conversion takes place and the new two-dimensional mesh is shown, viewed from the z-direction, in the *Mesh Scene 1* display. The mouse rotation option is suppressed for two-dimensional scenes.

Scaling the Mesh

The original mesh is not built to the correct scale and therefore requires scaling down by a factor of ten.

To scale the mesh:

1. Select **Mesh** > **Scale Mesh** from the menu bar.
2. In the *Scale Mesh* dialog:
   a) Ensure that the **Fluid 2D** region is selected.
   b) Enter a value of `0.1` for the scale factor.
   c) Click **Apply** to scale the domain and then click **Close**.



When the **Apply** button is clicked, the mesh region reduces in size.

3. To restore the previous viewing distance, click (**Reset View**) in the visualization toolbar.
4. Select **Mesh** > **Diagnostics** from the menu bar.
5. In the *Mesh Diagnostics* dialog, click **OK**.

The extent of the geometry is as given below:

```
x: [-4.0000e-01, 5.0000e-01] m

y: [-5.0586e-18, 4.0000e-01] m

z: [0.00000e+00, 0.0000e+00] m
```

The physics continuum that is associated with the deleted three-dimensional region is no longer required. To delete it:

6. Expand the **Continua** node. Right-click the **Physics 1** continuum node and select **Delete**.



7. Click **Yes** in the *Confirm Object Deletion* dialog.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this example, the flow is turbulent and the problem involves multiphase flow and boiling. Two fluids (water and water vapor) are required for the analysis. However, since these fluids occupy the same domain, only one continuum and one region are required to set up the simulation. By default, a continuum that is called **Physics 1 2D** is created when the mesh is converted to two-dimensional.

To select the physics models:

1. Rename the **Physics 1 2D** node to **Boiler**.
2. For the physics continuum, **Continua** > **Boiler**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase**
Eulerian Multiphase Model	**Volume of Fluid (VOF)**

Group Box	Model
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Optional Models	**Segregated Multiphase Temperature** **Gravity**

3. Click **Close**.
4. To review the models, open the **Boiler** > **Models** node.



5. Save the simulation.

Setting the Material Properties

Define the phases and set the appropriate material properties. Define the materials corresponding to each of the mixture components under the **Eulerian Phases** node.

To define the phases and set the material properties:

1. In the **Boiler** continuum, select the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **H2O**.
3. For the **H2O** phase, select the following models:

Group Box	Model
Material	**Liquid**
Equation of State	**Constant Density**

4. Click **Close**.
5. Create a second phase and rename it to **H2O (G)**.
6. For the **H2O (G)** phase, select the following models:

Group Box	Model
Material	**Gas**
Equation of State	**Constant Density**

7. Click **Close**.

The default gas must be replaced with water from the STAR-CCM+ material database.

8. Right-click the **H2O (G)** > **Models** > **Gas** > **Air** node and select **Replace with**.
9. In the *Replace Material* dialog, expand the **Material Databases** > **Standard** > **Gases** node.
10. Activate **H2O (Water)**.



11. Click **OK** to replace the material.

    The phase model objects appear as shown below.



    The default material properties for the two phases are suitable for this problem.

12. Save the simulation.

Defining the Phase Interactions

Define the interactions between the liquid and vapor phases. You specify the H2O phase as Primary Phase and the H2O (G) phase as Secondary Phase.

Using the **Multiphase Interaction** model, you can define the interaction between the liquid and vapor phases.

To define the phase interactions:

1. In the **Boiler** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. For the **Phase Interaction 1** phase interaction, select the following models in order:

Group Box	Model
Phase Interaction Topology	**VOF Phase Interaction**
Optional Models	**VOF Boiling** **Multiphase Material** (selected automatically)
Boiling Models	**Rohsenow Boiling**

3. Click **Close**.

   The selected models appear within the **Phase Interaction 1** > **Models** node, as shown below.



4. Save the simulation.

Defining the Boiling Model Parameters

Define the appropriate parameters for the Boiling model. You must specify which liquid phase is boiling and which gas phase is the product of boiling.

To define the Boiling model parameters:

1. Select the **Phase Interaction 1 > Models > VOF Phase Interaction** node and set **H2O** for the *Primary Phase* property and **H2O (G)** for the *Secondary Phase* property.



Wall heat flux due to boiling is a highly non-linear function at wall boundaries where the wall temperature is part of the numerical solution. For example, for given heat flux wall boundaries or solid-fluid interfaces. In this case, reduce the under-relaxation factor for heat flux due to boiling (*Under-Relaxation Factor* property under the **Rohsenow Boiling** node) to improve the convergence.



2. Save the simulation.

Setting Initial Conditions

Set the initial conditions for the physics continuum. You specify the initial volume fraction and temperature within the continuum. There is no initial water vapor in the continuum.

To set the initial conditions for the **Boiler** continuum:

1. Edit the **Initial Conditions** node and set the following properties:

Node	Property	Setting
Volume Fraction > Constant	**Value**	`[1.0, 0.0]`
Static Temperature > Constant	**Value**	`350 K`

2. Save the simulation.

Setting Boundary Conditions

You define the boundary types and specify the appropriate property values.

To set the boundary conditions:

1. Rename the**Regions** > **Default_Fluid 2D** node to **Fluid**.
2. Edit the **Fluid** > **Boundaries** node and set the boundary types as follows:

Boundary	Type
Bottom	**Wall**
Left	**Velocity Inlet**
Middle	**Wall**
Right	**Pressure Outlet**
TopRight	**Wall**
TopLeft	**Wall**

3. Edit the **Bottom** boundary and set the following properties:

Node	Property	Setting
Physics Conditions > Thermal Specification	**Condition**	**Temperature**
Physics Values > Static Temperature > Constant	**Value**	`540 K`



4. Edit the **Left** > **Physics Values** node and set the following properties:

Node	Property	Setting
Static Temperature > Constant	**Value**	350 K
Velocity Magnitude > Constant	**Value**	1 m/s
Volume Fraction > Constant	**H2O**	1.0
	H2O (G)	0.0

5. Edit the **Right** > **Physics Values** node and set the following properties:

Node	Property	Setting
Static Temperature > Constant	**Value**	370 K
Volume Fraction > Constant	**H2O**	1.0
	H2O (G)	0.0



6. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation. As you are solving an unsteady problem, it is necessary to specify the time-step size and the elapsed simulation time. This calculation runs for a maximum physical time of 3 s with a time-step size of 0.01 s.

Boiling flows are typically unsteady, with strong buoyancy effects, and large-scale flow instabilities. It is therefore recommended that unsteady solvers are used for this type of flows. You also have cases where the mean flow is dominant (as is the case in this tutorial) and where, eventually, a steady-state solution is possible. In these situations, try to simulate these flows using the steady-state model and an appropriate set of under-relaxation parameters. However, it is recommended that even for these cases, you use a time-marching solver with one iteration per time-step and a time-step size that secures a stable run.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.01 s`
Segregated Flow > Velocity	**Under-Relaxation Factor**	`0.8`

Node	Property	Setting
Segregated VOF	**Under-Relaxation Factor**	`0.1`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`3 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`1`

3. Save the simulation.

Reporting, Monitoring, and Plotting

You create a heat flux report, monitor, and plot for the bottom wall to assess convergence to a steady-state solution.

To create the report, monitor, and plot:

1. Right-click the **Reports** node and select **New Report > Heat Transfer.**



A new report is created named **Heat Transfer 1**.

To monitor heat flux on the bottom wall:

2. Rename the **Heat Transfer 1** node to `Heat Flux (Bottom Wall)`.

3. Make sure that the **Reports > Heat Flux (Bottom Wall)** node is selected and click to the right of *Parts* property.
4. Select **Bottom** and click **OK**.



5. Right-click the **Heat Flux (Bottom Wall)** node and select **Create Monitor and Plot from Report**.

   A new monitor and plot is created.



6. Select **Plots > Heat Flux (Bottom Wall) Monitor Plot** and set *x-Axis Monitor* to **Iteration**.

Visualizing the Solution

Create a scene to view the results of the simulation. You can view the distributions of the vapor volume fraction and fluid temperature throughout the run.

To view the results of the simulation:

1. Right-click the **Scenes** node and select **New Scene > Scalar**.

   The *Scalar Scene 1* display appears in the *Graphics* window.
2. Open the *scene/plot* window.
3. Rename the **Scalar Scene 1** node to `Vapor Volume Fraction`.
4. Select the **Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Volume Fraction > Volume Fraction of H2O (G)**.



5. Select the **Displayers** > **Scalar 1** node and set *Contour Style* to **Smooth Filled**.
6. Use the mouse buttons to zoom in and position the part to obtain a view similar to the view shown below.



7. Follow the above steps and create another scalar scene.
8. Rename the scene to `Fluid Temperature`.
9. Set *Function* to `Temperature`.
10. Right-click the **Plots** > **Heat Flux (Bottom Wall) Monitor Plot** and select **Open**.



11. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots. You can monitor the convergence by viewing the *Heat Flux (Bottom Wall) Monitor Plot*.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 300 time-steps are complete. When the last time-step is complete, the following message is displayed in the *Output* window:

   ```
   Stopping criterion Maximum Steps satisfied.
   ```

2. When the simulation has finished running, save it.

Visualizing Results

The scene tabs display the results of the simulation.

To view the results of the simulation:

1. Select the scene tabs in turn.

   A vapor volume fraction scene similar to the one shown below is displayed.



A fluid temperature scene similar to the one shown below is displayed.



A heat flux profile similar to the one shown below is displayed.



2. Save the simulation.

Modifying the Boiling Model Parameters

The Rohsenow empirical correlation is used to calculate the surface heat flux due to boiling. You specify the two parameters of this correlation, namely *C_qw* and *n_p*.

In STAR-CCM+, the Rohsenow empirical correlation is used to calculate the surface heat flux due to boiling. The two parameters of this correlation, namely *C_qw* and *n_p*, are specified. These empirical coefficients vary with the liquid-surface combination and the surface finish. Saiz Jabardo provides some typical parameters for various conditions.

Assume that the bottom wall is a polished copper and specify appropriate values for *C_qw* and *n_p* parameters.

To modify the Boiling model parameters:

1. Select the **Continua** > **Boiler** > **Models** > **Multiphase Interaction** > **Phase Interaction 1** > **Models** > **Rohsenow Boiling** node and set *C_qw* to `0.0128` and *n_p* to `1.7`.

   Leave the remaining parameters at their default values.



2. Save the simulation.

Running the Simulation

Clear the results of the previous simulation and run the simulation again.

To clear the results and run the simulation again:

Clear the previous solution.

1. In the menu, select **Solution** > **Clear Solution**.
2. In the *Clear Solution* dialog, click **OK**.



The simulation is now ready to be run.

3. Click (**Run**) in the *Solution* toolbar.
4. When the simulation has finished running, save it.

Visualizing Results

The scene tabs display the results of the simulation.

A vapor volume fraction scene similar to the one shown below is displayed.



A fluid temperature scene similar to the one shown below is displayed.



A heat flux profile similar to the one shown below is displayed.



From the above scenes, it is clear that the *C_qw* and *n_p* parameters have significant effects on the predicted levels of wall heat flux and the vapor volume fraction. In this example, you halved the heat flux on the bottom wall which resulted in the production of five times less water vapor. It is important therefore to make sure that correct values are used, to reflect the liquid-surface combination and the surface finish.

1. Save the simulation.

Summary

This tutorial demonstrated how to set up a boiling problem in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Defining a model for boiling flow problems.
- Defining and modifying boiling model parameters.
- Setting initial conditions.
- Setting solver parameters for an unsteady run.
- Analyzing wall boiling results.

Bibliography

The following papers were referenced in this tutorial.

Rohsenow, W.M. 1952. "A Method of Correlation Heat Transfer Data for Surface Boiling of Liquid", Trans. ASME, 74, 969.

Saiz Jabardo, J. M., Fockink da Silva, E., Ribatski, G. and de Barros, S. F. 2004. "Evaluation of the Rohsenow Correlation Through Experimental Pool Boiling of Halocarbon Refrigerants on Cylindrical Surfaces", J. of the Braz. Soc. of Mech. Sci. and Eng, April-June 2004, Vol. XXVI, No. 2.

VOF: Melting-Solidification

This tutorial demonstrates how to set up a melting and solidification analysis in STAR-CCM+. It simulates the freezing of water in a pipe.

The geometry is shown in the following figure:



Water flows along a pipe with freezing walls. This scenario can be modeled as an axisymmetric two-dimensional pipe. The wall temperature is a function of position, with the middle of the wall being at the lowest temperature according to:

$$T_{wall} = \min\left(273.1,\ 2010.0(x - 0.15)^2 + 253.0\right) \tag{5}$$



Table 1: Water and Ice Properties

Property	Value
Solidus temperature	273.0 K
Liquidus temperature	273.0 K
Density of water and ice	997.561 kg/m^3 (No expansion)
Viscosity of H2O	8.8871e-4 Pa s
Latent heat of fusion	100000.0 J/kg
Specific heat	4181.72 J /kg K
Thermal conductivity	Step function in T

The simulation is run using the melting-solidification model. The post-processing tools in STAR-CCM+ are used to visualize the results.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Selecting the Physics Models
Defining the Liquid Phase for the Melting-Solidification Model
Setting the Initial Conditions
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Creating a Mass Flow Monitor Plot
Setting Up Scenes to Visualize Results
Running the Simulation
Analyzing the Results
Summary

Prerequisites

The instructions in the VOF: Melting-Solidification tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+

Technique	Tutorial
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A mesh file has been prepared for this tutorial.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh...**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/freezingWaterDuct.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.



5. Save the simulation as `freezingWaterDuct.sim`.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this case, the transient Reynolds-Averaged Navier Stokes equations are solved for using the K-Epsilon turbulence model and the Segregated Flow solver. The Volume of Fluid (VOF) model is activated before the melting and solidification option becomes available for Eulerian phases.

To select the physics models:

1. Right-click **Continua** > **Physics 1** and click **Select models...**.
2. In the *Enabled Models* group box, deactivate **Two-Dimensional**.
3. Select the following models in order:

Group Box	Model
Space	**Axisymmetric**
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Volume of Fluid (VOF)** **Gradients** (Selected automatically) **Segregated Flow** (Selected automatically) **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Turbulent** **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically) **Two-Layer All y+ Wall Treatment** (Selected automatically)
Optional Models	**Segregated Multiphase Temperature**

4. Click **Close**.
5. To review the models, open the **Physics 1 > Models** node.



6. Save the simulation.

Defining the Liquid Phase for the Melting-Solidification Model

You define the liquid phase and apply the Melting-Solidification model.

You model the freezing water as a Eulerian phase in STAR-CCM+. The properties of water come from the material database, with some modifications to suit the problem. For this tutorial, the value of latent heat is reduced so that freezing takes place sooner than it would with the default value. For a real case, however, leave the value at its original setting.

To define the liquid phase for the Melting-Solidification model:

Create the H2O phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **H2O**.
3. For the **H2O** phase, select the following models in order:

Group Box	Model
Enabled Models	**Energy** (selected automatically) **Turbulent** (selected automatically)
Material	**Liquid**
Equation of State	**Constant Density**
Optional Models	**Melting-Solidification**

4. Click **Close**.

   The **Models** tree under **Eulerian Phases** appears as follows:



Set the solid volume fraction threshold above which flow is stopped in the fluid phase.

5. Edit the **Eulerian Phases** > **H2O** > **Models** node and set the following properties:

Node	Property	Setting
Melting-Solidification	**Flow Stop Relative Solid Fraction**	`0.999`
	Under-Relaxation Factor	`0.1`

Set the flow stop value in the **Melting-Solidification Flow Stop** node.

6. Edit the **Physics 1 > Models** node and set the following properties:

Node	Property	Setting
Melting-Solidification Flow Stop	**Flow Stop Solid Fraction**	`0.999`

A step function defines this property: H2O in liquid state flows without any resistance in cells where the solid volume fraction is 0.999 or less. STAR-CCM+ sets the fluid velocity to zero for cells with solid volume fraction above the flow stop value.

The thermal conductivity of H2O varies with phase and temperature. For this simulation, assume that the thermal conductivity is constant for each phase.

The thermal conductivity of ice is 2.33 W / m K while the conductivity of H2O is 0.620271 W/m K. Use a conditional expression to set the conductivity of the H2O phase. Set the conductivity to 2.33 W/m K when the temperature is less than or equal to 273 K, and 0.620271 W / m K for higher temperatures.

Create a step function to set the conductivity of H2O:

7. Right-click the **Tools > Field Functions** node and select **New > Scalar** from the pop-up menu.
8. Rename the **User Field Function 1** to `Variable Thermal Conductivity`.
9. Select the **Variable Thermal Conductivity** node and set the following properties:

Property	Setting
Function Name	`Variable Thermal Conductivity`
Definition	`(${Temperature} <= 273.) ? 2.33 : 0.620271`

10. Set **Dimensions** to **Power/Length-Temperature** by filling the *Dimensions* dialog with the following values:



The *Properties* window appears as shown.



Use this variable thermal conductivity for the H2O phase.

11. Edit the **Eulerian Phases > H2O > Models > Liquid > H2O > Material Properties** node and set the following properties:

Node	Property	Setting
Thermal Conductivity	**Method**	**Field Function**

Node	Property	Setting
Thermal Conductivity > Field Function	**Scalar Function**	**Variable Thermal Conductivity**
Latent Heat of Fusion > Constant	**Value**	`100000.0 J/kg`
Liquidus Temperature > Constant	**Value**	`273.0 K`
Solidus Temperature > Constant	**Value**	`273.0 K`

12. Save the simulation.

Setting the Initial Conditions

You set the initial conditions for the physics continuum.

A field function sets up the initial pressure field as a linear profile: this profile is close to the expected solution. Setting the initial pressure in this way reduces the time that is required to obtain a solution. The pressure profiles at the inlet and outlet are specified to define the linear pressure field correctly.

In this case, the inlet is not modeled as a velocity inlet, since the simulation diverges when the whole-cross section is frozen, and the flow stops. A combination of pressure outlets at both inlet and outlet allows the flow to be stopped completely in the duct without violating any boundary condition

When choosing initial conditions, it is important to select a field that is bounded by the boundary conditions to avoid divergence. To satisfy this criterion, specify a linear pressure field with maximum and minimum values that are set as the boundary conditions.

The initial temperature is set to the inlet temperature. The initial velocity field is set to be close to the expected average velocity during the initial stages of the flow, that is, before the freezing commences, for faster convergence.

To set the initial conditions:

1. Create the following field functions using the same procedure as in the previous section.

The desired gauge inlet pressure is taken as 1.0 Pa.

Node Title	**Target Inlet Pressure**
Dimensions	`Pressure`
Function Name	`TargetInletPressure`
Definition	`1.0`

As previously mentioned, the inlet is modeled as a pressure outlet: the dynamic pressure and turbulent stresses are added to make sure that the static inlet pressure is always 1.0 Pa.

Node Title	**Inlet Pressure**
Dimensions	`Pressure`
Function Name	`InletPressure`
Definition	`${TargetInletPressure} + 0.5*${Density}*pow(mag($$ {Velocity}), 2)`

The outlet pressure is set to 0.0 Pa.

Node Title	**Outlet Pressure**
Dimensions	`Pressure`
Function Name	`OutletPressure`
Definition	`0.0`

The initial pressure field is then taken as a linear profile from 1.0 Pa at X = 0 m to 0.0 Pa at X = 0.3 m at the lower end of the pipe.

Node Title	**Initial Pressure**
Dimensions	`Pressure`
Function Name	`InitialPressure`
Definition	`${TargetInletPressure} - $ ${Position}[0]/0.3*($ {TargetInletPressure} - $ {OutletPressure})`

2. Right-click **Tools** > **Field Functions** and select **Refresh** to order the created field functions alphabetically.



Set the initial conditions for the solved variables in the **Physics 1** continuum.

3. Edit the **Physics 1 > Initial Conditions** node and set the following properties:

Node	Property	Setting
Pressure	**Method**	**Field Function**
Field Function	**Scalar Function**	**Initial Pressure**
Static Temperature > Constant	**Value**	`273.1 K`
Velocity > Constant	**Value**	`[0.012, 0.0, 0.0] m/s`

Node	Property	Setting
Volume Fraction > Constant	**Value**	[1.0]



4. Save the simulation.

Setting Boundary Conditions

You define the boundary types and specify the appropriate property values. Set the boundary conditions for the inlet, outlet, and pipe wall.

To set the boundary conditions:

1. Expand the **Regions > Tube_2D > Boundaries** node to reveal the different boundaries.

Start with the Inlet boundary conditions. The inlet temperature is 273.1 K and the volume fraction is 1.0.

The dynamic pressure and contribution to stresses due to turbulence are removed from the specified pressure at a pressure outlet in case of reverse flow. The inlet pressure profile is therefore specified as: $1.0 + \frac{1}{2}\rho u^2$ as specified in the Inlet Pressure field function. This setting is to make sure that the static inlet pressure is always effectively 1.0 Pa when water flows into the pipe.

2. Edit the **Inlet** > **Physics Values** node and set the following properties:

Node	Property	Setting
Pressure	**Method**	**Field Function**
Pressure > Field Function	**Scalar Function**	**Inlet Pressure**
Static Temperature > Constant	**Value**	`273.1 K`
Volume Fraction > Constant	**Value**	`[1.0]`



The outlet pressure is set to atmospheric pressure, that is, 0 Pa gauge pressure. However, a field function is already defined for outlet pressure in the previous section and is reused here. The advantage of this method is: if you must change the pressure profile across the whole domain, you alter the outlet pressure value in only one place, rather than two.

3. Edit the **Outlet > Physics Values** node and set the following properties:

Node	Property	Setting
Pressure	**Method**	**Field Function**
Pressure > Field Function	**Scalar Function**	**Outlet Pressure**
Static Temperature > Constant	**Value**	`273.1 K`

Node	Property	Setting
Volume Fraction > Constant	**Value**	`[1.0]`

Specify the freezing wall boundary conditions.

4. Create the following field function using the same procedure as in the previous sections.

Node Title	`Wall Temperature`
Dimensions	`Temperature`
Function Name	`WallTemperature`
Definition	`min(273.1, 2010.0*pow($${Position}[0] - 0.15, 2) + 253.0)`

5. Edit the **Wall** boundary and set the following properties:

Node	Property	Setting
Physics Conditions > Thermal Specification	**Condition**	**Temperature**
Physics Values > Static Temperature	**Method**	**Field Function**
Field Function	**Scalar Function**	**Wall Temperature**



STAR-CCM+ applies a no-slip boundary condition to a wall by default, which is what is required in this simulation.

6. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation. You specify the time-step and the maximum number of iterations per time step.

The flow stop relative solid fraction property for melting and solidification is implemented as a step function. Therefore, cells containing H2O in liquid state require a few inner iterations to realize that velocity is zero in the newly frozen cells. This process is dynamic and many cells can be 'frozen' at subsequent iterations during one time-step. Therefore, many inner iterations are required to make sure that results are accurate.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.2 s`
Segregated Energy	**Fluid Under-Relaxation Factor**	`0.5`



2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	**Maximum Inner Iterations**	`40`
Maximum Steps	**Enabled**	Deactivate
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`20.0 s`



3. Save the simulation.

Creating a Mass Flow Monitor Plot

You create a monitor plot for normalized mass flow through the outlet.

To create a monitor plot for normalized mass flow through the outlet:

1. Right-click **Reports** and select **New Report** > **Mass Flow**.



2. Rename the **Mass Flow 1** report to **Mass Flow**.
3. Select the **Mass Flow** node and set the *Parts* to **Tube_2D > Outlet**.
4. Select **Outlet** from the dialog which appears when editing the *Parts* property value.



5. Right-click **Reports** > **Mass Flow** and select **Create Monitor and Plot from Report**.
6. Select **Monitors** > **Mass Flow Monitor**.
7. Set *Normalization Option* to **Manual**.
8. Make sure that the *Trigger* property is set to **Time Step**.



9. Expand the **Mass Flow Monitor** node and select **Manual Normalization**.

   The normalization factor is chosen as the mass outflow for the first iteration.
10. Set *Normalization Factor* to `3.74E-5`.

Edit the plot graphics settings.

11. Rename **Plots** > **Mass Flow Monitor Plot** to **Normalized Mass Flow Monitor Plot**.

12. Set the following properties:

Node	Property	Setting
Normalized Mass Flow Monitor Plot	*Title*	**Normalized Mass Flow Monitor Plot**
Axis > X Axis	*Minimum*	`0.0`
	Lock Minimum	Deactivated
	Maximum	`20.0`
	Lock Maximum	Activated
Title	*Title*	`Physical Time (s)`
Axis > Y Axis	*Minimum*	`0.0`
	Lock Minimum	Deactivated
	Maximum	`1.0`
	Lock Maximum	Activated
Title	*Title*	`Mass Flow Monitor`



13. Save the simulation.

Setting Up Scenes to Visualize Results

Create scalar scenes to visualize the results as the simulation progresses.

To create scalar scenes to visualize the results:

1. Right-click **Scenes** and select **New Scene** > **Scalar**.
2. Rename **Scalar Scene 1** to **Solid Volume Fraction**.

   The solid volume fraction function represents the state of the H2O phase in each cell. When solid volume fraction is 1.0, ice is present in the cell, and when solid volume fraction is 0.0, (liquid) water is present. Do not confuse solid volume fraction with volume fraction which represents the phase in the cell. Volume fraction is always 1.0 in this case, since there is only one phase in this simulation: H2O.

3. Select **Scenes > Solid Volume Fraction > Displayers > Scalar 1 > Scalar Field** and set *Function* to **Solid Volume Fraction > Solid Volume Fraction of H2O**.



The domain is long in X and narrow in Y: this aspect ratio makes it hard to visualize any scalar field.

Apply a simple transform to scale the Y axis for a better solution representation.

4. Right-click **Tools > Transforms** and select **New Graphics Transform > Simple Transform**.
5. Rename the **Simple Transform 1** node to **Stretch Y**.
6. Set *Scale* to `[1.0, 10.0, 1.0]`.



Apply the **Stretch Y** transform to the **Displayers** in the **Solid Volume Fraction** scene.

7. Select **Scenes** > **Solid Volume Fraction** > **Displayers** > **Outline 1** and set *Transform* to `Stretch Y`.
8. Select **Scenes** > **Solid Volume Fraction** > **Displayers** > **Scalar 1** and set *Transform* to `Stretch Y`.



Create more scenes to display pressure and velocity.

9. Click **Scenes > Solid Volume Fraction** and while holding the left mouse button, drag the **Solid Volume Fraction** node onto the **Scenes** node and rename the new node as **Pressure**.
10. Select **Scenes > Pressure > Displayers > Scalar 1 > Scalar Field** and set *Function* to **Pressure**.
11. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it. If left alone, the simulation continues until 10 seconds of physical time is complete.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

You can check the *Normalized Mass Flow Monitor* while the solution is running.

2. Double-click **Plots > Normalized Mass Flow Monitor Plot** to display the results in the *Graphics* window.



When the last time-step is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

3. When the simulation has finished running, save it.

Analyzing the Results

View and examine the results of the simulation.

The freezing starts in the middle of the top wall. The water in the whole pipe eventually solidifies as time progresses.

1. Click the *Solid Volume Fraction* tab in the *Graphics* window.



The velocity in the frozen region must be zero since to the *Flow Stop Relative Solid Fraction* property is set to a value less than one.

2. Click the *Pressure* tab in the *Graphics* window.



3. Save the simulation.

Summary

This tutorial demonstrated how to set up a melting-solidification problem in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Defining models for melting-solidification problems.
- Defining material properties that are required for the selected models.
- Setting solver parameters for a transient run.
- Plotting a normalized mass flow monitor.
- Running the solver to a specified maximum physical time.
- Displaying the results using the built-in visualization facilities.

Lagrangian: Particle-Laden Flow

This tutorial demonstrates how to set up a simple Lagrangian multiphase analysis in STAR-CCM+. It simulates a particle-laden flow of air through a partially blocked elbow duct.

Air at standard pressure (1 atm) enters the channel at a velocity of 10 m/s. The fluid, having passed a partial blockage and a 90 degree elbow, exits vertically through the outlet. All fluid properties are assumed to be constant.

The air stream is seeded with solid particles, uniformly distributed at the duct inlet. The volumetric loading of particles in the incoming air is 0.01%, corresponding to a volumetric particle flow rate of 6.4516×10^{-7} m^3 /s. The particle properties are as follows:

- Density 1.2×10^3 kg/m^3
- Diameter 4×10^{-5} m
- Initial velocity (10, 0, 0) m/s

Adiabatic and no-slip conditions are assumed at the duct walls. Turbulence is modeled using the default K-Epsilon model. All particles rebound perfectly from the walls and gravity effects are neglected.

The geometry that is used in the tutorial is shown below.



Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Imported Geometry
Selecting the Physics Models
Setting Continuous Phase Material Properties
Selecting the Lagrangian Phase Models
Setting the Lagrangian Phase Model Properties
Setting the Lagrangian Phase Boundary Conditions
Setting the Continuous Phase Boundary Conditions
Setting Up the Injector
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Initializing the Solution and Running the Simulation
Visualizing the Results
Displaying Particle Tracks
Using Tracks in Field Functions and Derived Parts

Summary

Prerequisites

The instructions in the Lagrangian: Particle-Laden Flow tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A three-dimensional hexahedral cell mesh has been prepared for this analysis. The grid has been saved in the .ccmg file format which contains all the required cell, vertex, and boundary information for the problem geometry.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/elbow.ccmg`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. One mesh region named **Default_Fluid** is created in the **Regions** node representing the grid domain. A geometry scene is created in the *Graphics* window.
5. Save the simulation as `obstructedElbow.sim`.

Visualizing the Imported Geometry

Examine the imported geometry.

Examine the *Geometry Scene 1* display in the *Graphics* window.

To visualize the imported geometry:

1. Use the mouse controls to adjust the view as shown below.



To make the interior features of the geometry visible, change the opacity of its surfaces.

2. Open the **Scenes** > **Geometry Scene 1** > **Displayers** node, then select the **Geometry 1** node.



3. Set the *Opacity* to `0.5`.



The obstruction in the elbow geometry is now clearly visible.

4. Save the simulation.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this simulation, the flow is turbulent and incompressible. The Segregated Flow model is used together with the default K-Epsilon

Turbulence model. The Lagrangian Multiphase model is used to model the dispersed phase.

Only one region for the continuous phase, air, is required for the analysis. By default, the continuum **Physics 1** is automatically created when the mesh is imported. Edit the continuum so that appropriate physical models are selected for the air fluid region.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Pre-selected)
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Gradient Metrics	**Gradients** (Selected automatically)
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**
Reynolds Averaged Turbulence	**K-Epsilon Turbulence**
Optional Models	**Lagrangian Multiphase**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Setting Continuous Phase Material Properties

You set the material properties of the continuous phase. The default material for the continuous phase is air and, for this tutorial, is considered incompressible. You update the default material properties for air so that it has a density of 1.205 kg/m^3 and a dynamic viscosity of 1.82E-5 Pa-s.

To set the continuous phase material properties:

1. Edit the **Physics 1 > Models > Gas > Air > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1.205 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`1.82E-5 Pa-s`

2. Save the simulation.

Selecting the Lagrangian Phase Models

You create the Lagrangian phase and select the appropriate phase models. These models represent the characteristics of the Lagrangian phase.

To create the Lagrangian phase and select the phase models:

1. In the **Physics 1** continuum, right-click the **Models** > **Lagrangian Multiphase** > **Lagrangian Phases** node and create a new phase.
2. For the **Phase 1** phase, select the following models:

Group Box	Model
Enabled Models	**Residence Time** (Selected automatically)
Particle Type	**Material Particles**
Material	**Solid**
Equation of State	**Constant Density**
Optional Particle Forces	**Drag Force**
Track Sampling	**Track File**
Optional Models	**Two-Way Coupling**

3. Click **Close**.

   The **Phase 1** phase node appears as shown below:



4. Save the simulation.

Setting the Lagrangian Phase Model Properties

You specify the appropriate properties of the Lagrangian phase models.

To set the properties for the Lagrangian Phase models:

1. Edit the **Lagrangian Phases > Phase 1 > Models** node and set the following properties:

Node	Property	Setting
Solid > Al > Material Properties > Density > Constant	**Value**	`1200 kg/m^3`
Drag Force > Drag Coefficient	**Method**	**Schiller-Naumann**
Track File	**Vectors**	**Parcel.Centroid Particle.Velocity**



2. Save the simulation.

Setting the Lagrangian Phase Boundary Conditions

You set the boundary interaction mode for Lagrangian Phase Boundary Conditions.

To set the Lagrangian phase boundary conditions:

1. Edit the **Lagrangian Multiphase > Lagrangian Phases > Phase 1 > Boundary Conditions > Wall** node and set the following properties:

Node	Property	Setting
Physics Conditions > Mode	**Active Mode**	**Rebound**
Physics Values > Normal Restitution Coefficient > Constant	**Value**	`1.0`

Node	Property	Setting
Tangential Restitution Coefficient > Constant	**Value**	`1.0`

Setting the **Value** to `1` represents a perfect elastic rebound.

2. Save the simulation.

Setting the Continuous Phase Boundary Conditions

Define the continuous phase boundary conditions. The default boundary conditions for the outlet and wall boundaries are suitable for this analysis. You specify the boundary conditions at the inlet, a fixed velocity of 10 m/s and appropriate turbulence parameters.

To set the continuous phase boundary conditions:

1. Edit the **Regions > Default_Fluid > Boundaries > inlet** node and set the following properties:

Node	Property	Setting
Physics Conditions > Turbulence Specification	**Method**	**Intensity + Length Scale**
Physics Values > Turbulent Length Scale > Constant	**Value**	`0.001 m`
Turbulence Intensity > Constant	**Value**	`0.005`
Velocity Magnitude > Constant	**Value**	`10 m/s`

2. Save the simulation.

Setting Up the Injector

You set up an injector to introduce the particles of the dispersed Lagrangian phase into the solution domain. The injector defines the initial state of the particles.

To set up the injector:

1. Right-click the **Injectors** node and select **New**.
2. Select the **Injectors** > **Injector 1** node and set *Type* to **Part Injector**.



3. Set *Inputs* to **Regions** > **Default Fluid** > **Boundaries** > **inlet**.



4. Click **OK**.
5. Set *Lagrangian Phase* to **Phase 1**.



The *Number of Points* property is 0 until the injection points are calculated during solution initialization. Specify the conditions and values of the injector to define the particles.

6. Open the **Injectors** > **Injector 1** node.



7. Edit the following nodes and set the appropriate properties:

Node	Property	Setting
Conditions > Flow Rate Specification	*Method*	**Volume Flow Rate**
Flow Rate Distribution	*Method*	**Per Injector** This signifies that the volume flow rate applies to the whole injector, rather than per injection point.
Velocity Specification	*Method*	**Magnitude + Direction**
Values > Volume Flow Rate > Constant	*Value*	`6.4516E-7` m^3/s

Node	Property	Setting
Particle Diameter > Constant	*Value*	`4.0E-5` m
Velocity Magnitude > Constant	*Value*	`10` m/s

8. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation. You create a new stopping criterion based on the continuity equation residual and set it to a minimum tolerance of 1E-4.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following property:

Node	Property	Setting
Lagrangian Multiphase > Steady	**Maximum Residence Time**	`0.1`

A new stopping criterion is created based on the continuity equation residual and with a minimum tolerance of 1E-4.

2. Right-click the **Stopping Criteria** node and select **Create New Criterion > From Monitor**.



3. Select **Continuity** from the Select Monitor dialog and click **OK**.



4. Edit the **Stopping Criteria** node and set the following property:

Node	Property	Setting
Continuity Criterion > Minimum Limit	**Minimum Value**	`1.0E-4`

5. Save the simulation.

Visualizing the Solution

View the results of the simulation. You can view the velocity vectors on a cross-section through the fluid region as the solution develops.

To visualize the solution:

1. Right-click the **Scenes** node, and select **New Scene > Vector**.
   The *Vector Scene 1* display appears.

Create a plane section part for display.

2. Right-click the **Derived Parts** node, and select **New Part > Section > Plane**.
   A new interactive in-place dialog, *Create Plane Section*, appears.
3. In the *Plane Parameters* box, set the normal to `(0,0,1)` and the origin to `(0,0,0.02)`.

4. In the *Display* box, select the radio button **Existing Displayer** to show the plane section in the current scene.

   The completed dialog is shown below.



5. Click **Create** to create the plane section.
6. Click **Close** in the *Create Plane Section* dialog.

   A node appears within the **Derived Parts** node, **plane section**.



The plane section part is displayed in the vector scene geometry. To display velocity vectors on the plane section, the part must also be added to the Vector 1 displayer.

7. Select the **Scenes** > **Vector Scene 1** > **Displayers** > **Vector 1** > **Parts** node and set *Parts* to **Derived Parts** > **plane section**.
8. Select the **Scenes** > **Vector Scene 1** > **Displayers** > **Outline 1** > **Parts** node and set *Parts* to **Regions** > **Derived Parts** > **plane section**.

The vector scene is now ready. No vectors appears in the scene until the run is started.



Create a scalar scene.

9. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.

   The *Scalar Scene 1* display appears.

To define the scalar scene:

10. Select the new **Scalar Scene 1** > **Displayers** > **Scalar 1** node and change the *Contour Style* to **Smooth Filled**.
11. Select the **Scalar 1** > **Parts** node and set *Parts* to **Regions** > **Default Fluid** > **Default_Boundary_Region**.
12. Right-click the scalar bar in the display. In the pop-up menu that appears, select **Incident Mass Flux** > **Phase1**.



13. Save the simulation.

Initializing the Solution and Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To initialize the solution and run the simulation:

1. Click (**Initialize Solution**) in the *Solution* toolbar or select **Solution** > **Initialize Solution** from the menu.
2. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. You can ignore the warning message about smoothing in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.
3. Select the *Vector Scene 1* tab to view the velocity vectors cross-section as the solution develops.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 1000 iterations are complete or the continuity residual criterion, 1.0E-4, is satisfied.



4. When the simulation has finished running, save it.

Visualizing the Results

You can view the results of the simulation. The *Vector Scene 1* displays the velocity vector cross-section for the converged solution. The *Scalar Scene 1* displays the incident mass flux profile for the converged solution.

To visualize the results:

1. Select the *Vector Scene 1* display to view the velocity vector cross-section for the converged solution.



2. Select the *Scalar Scene 1* display to view the incident mass flux profile for the converged solution.
3. Use the mouse controls to adjust the view as shown below.



The particle impacts are concentrated on the outer radius of the bend. In a realistic modeling exercise, a considerably larger number of particles tracks would be needed to increase confidence in the results.

Displaying Particle Tracks

You can display the particle tracks recorded by the Track File model.

To display the particle tracks:

1. Right-click the **Tools > Track Files** node and select **Track File**.



2. In the *Open* dialog, select the `obstructedElbow.trk` file and click **Open**.



Note: The simulation must be saved before the track file appears in the dialog.

A new node **obstructedElbow** appears under the **Track Files** node. Another new node, **Particle Tracks**, also appears in the simulation tree.

3. Open the **Particle Tracks** node to see the **Phase 1 tracks** node that was populated from the `obstructedElbow.trk` file.



Add particle tracks to the scalar scene.

4. Select the *Scalar Scene 1* display in the *Graphics* window.
5. Jump directly to the object tree for the **Scalar Scene 1** node by clicking the **scene/plot** button above the simulation tree.



Add feature lines first, and then add particle tracks.

6. Select **Scalar Scene 1** > **Displayers** > **Outline 1** and tick the *Feature Lines* property.



7. Select the **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** node and click the right half of the *Parts* property.

8. In the dialog that appears, use (**Column Selection**) to **Deselect All**, expand the **Particle Tracks** node and select **Phase 1 tracks** as shown below.



9. Click **OK**.
10. Continuing within the **Scalar 1** node, select the **Scalar Field** node.



11. Set the *Function* to **Track > Track: Particle Residence Time** using the drop-down menu.



12. In the *Scalar Scene 1* display, use the mouse controls to adjust the view as shown below.



The maximum time taken for all particle tracks to leave the domain is approximately 0.032 seconds. So, the solver parameter value that is specified for *Maximum Residence Time*, 0.1s, is more than sufficient.

13. With the **Scalar 1 > Scalar Field** node still selected in the partial tree for **Scalar Scene 1**, set *Function* to **Track: > Track: Particle Velocity > Magnitude**.

    The particle tracks are now colored according to the velocity of the particles.



14. Click **Simulation** to return to the STAR-CCM+ object tree.
15. Save the simulation.

Using Tracks in Field Functions and Derived Parts

In simulations where many particle tracks are present, it is likely that only a small number are required for plotting in a graphics scene. You can create a field function to filter the available tracks and apply the filter in a derived part.

A field function is created that returns the Track Parcel Index number for every eighth track. The function name for the Track Parcel Index is **TrackParcelId** (which can be found from the *Function Name* property of the **Track: Parcel Index** node in **Tools > Field Functions**).

To use tracks in field functions and derived parts:

1. Right-click the **Tools > Field Functions** node and select **New > Scalar** from the pop-up menu.
2. Rename the **User Field Function 1** node to `TrackFilter`.
3. Select the **Track Filter** node and set the following properties:

Property	Setting
Function Name	`Track Filter`
Dimensions	*Dimensionless*
Definition	`(fmod($TrackParcelId, 8) <= 0.0001) ? $TrackParcelId : -1`

The function entered here uses the fmod function and the conditional expression that is defined by the ternary operator "?:". The function essentially says that if an integer results from the division of **TrackParcelId** by 8, then the **TrackParcelId** is returned; otherwise the value -1 is returned.

Create a derived part based on a threshold to use this behavior.

4. Right-click **Derived Parts** and select **New Part > Threshold** from the pop-up menu.

   The *Create Threshold* dialog appears above the object tree.

5. In the *Input Parts* box, click the drop-down menu. In the dialog that appears, use (**Column Selection**) to **Deselect All**, expand the **Particle Tracks** node and select **Phase 1 tracks** item.



6. Click **OK**.

7. Click the drop-down menu in the *Scalar* box and choose item **Track Filter** as shown.



8. Set the **Extraction Mode** to **All Above** and the **Threshold Value** to `0.001`.
9. Select **No Displayer** in the *Display* box.

   The completed panel is shown below.



10. Click **Create** and then **Close**.

    A new derived part named **threshold** appears in the object tree.



11. Rename the **threshold** node to `Filtered Tracks`.

Modify *Scalar Scene 1* to display the filtered tracks:

12. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** node.



13. Click the right half of the *Parts* property.
14. In the dialog that appears, click the drop-down menu.
15. In the dialog that appears, use (**Column Selection**) to **Deselect All**, expand the **Derived Parts** node and select **Filtered tracks** as shown below.



16. Click **OK**.

Scalar Scene 1 now shows the filtered set of particle tracks.



Summary

This tutorial has demonstrated how to set up a simple Langrangian analysis in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Selecting models for a Lagrangian phase.
- Setting up an injector for Lagrangian Multiphase modeling.
- Defining material properties that are required for the selected models.
- Defining boundary conditions.
- Setting solver parameters.
- Initializing and running the solver until the residual magnitudes are sufficiently small.
- Analyzing results using the visualization facilities.
- Displaying particle tracks.
- Creating a field function and derived part based on particle tracks.

Lagrangian: Solid Particle Erosion

This tutorial demonstrates how to set up an erosion modeling analysis in STAR-CCM+. It simulates the erosion that is caused by solid particles in a liquid flow.

The geometry that is used in this tutorial was originally designed as a pressure reduction device for choke valves under erosive service, and is the subject of a US Patent. The device is shown below.



This device was tested under erosive service using a sand-liquid erosion test facility (as reported by Wallace), and the conditions that are used in this tutorial are based on the actual tests. In the erosion tests, the flow restrictor was mounted between lengths of 53.1 mm diameter piping. Static pressure measurements were made at two positions: 106 mm upstream of the restrictor and 318 mm downstream. Overall mass loss was obtained by weighing the flow restrictor periodically during the tests.

In this tutorial, water at room temperature is used for the liquid phase. Inlet conditions are as follows (corresponding to a flow rate of 28.57 liters/s in a 53.1 mm diameter pipe):

Table 2: Liquid Phase

Property	Value	Units
Fluid Density	997.561	kg/m^3
Fluid Viscosity	8.8871e-4	Pa-s
Inlet Velocity	12.901	m/s
Turbulence Intensity	0.029	
Turbulent Length Scale	0.0037	m

The particle loading by mass is 0.39%. The particle properties are as follows:

Table 3: Particle Phase

Property	Value	Units
Density	2650.0	kg/m^3
Diameter	2.75e-4	m
Initial axial velocity	12.901	m/s

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Imported Geometry
Selecting the Physics Models
Selecting the Lagrangian Phase Models
Setting the Material Properties
Setting the Lagrangian Phase Model Properties
Defining Field Functions for the Restitution Coefficients
Setting the Lagrangian Phase Boundary Conditions
Setting the Erosion Model
Setting Continuous Phase Boundary Conditions
Setting Up the Injector
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Reporting and Plotting the Pressure Drop
Initializing and Running the Simulation
Visualizing the Results
Estimating the Overall Erosion Rate

Comparison with Experimental Data
Implementing Alternative Functions for Erosion Ratio
Summary
Bibliography

Prerequisites

The instructions in the Lagrangian: Solid Particle Erosion tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A three-dimensional polyhedral cell mesh has been prepared for this analysis. The grid has been saved in the .ccm file format which contains all the required cell, vertex, and boundary information for the problem geometry.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/flow_restrictor.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. One mesh region named **Fluid** is created in the **Regions** node representing the grid domain. A geometry scene is created in the *Graphics* window.
5. Save the simulation as `erosion.sim`.

Visualizing the Imported Geometry

Examine the *Geometry Scene 1* display in the *Graphics* window.

To visualize the imported geometry:

1. Use the mouse controls to adjust the view as shown below.



2. To see the boundary regions that were imported with the mesh, open the **Regions** > **Fluid** > **Boundaries** node.



3. To observe the surfaces of the flow restrictor in the *Geometry Scene 1* display, select the **Restrictor** boundary node.



Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

Several models are required in this tutorial as there is not only a turbulent continuous phase to consider, but also the particle motion within the continuous phase. Two different strategies are employed within STAR-CCM+ to simulate these phases. The continuous fluid phase is modeled using an Eulerian formulation whereby fluid properties are obtained at fixed points throughout the fluid domain. The particulate phase is modeled using a Lagrangian method where the trajectories of representative particles are tracked throughout the continuous phase.

STAR-CCM+ permits two-way coupling between the continuous fluid phase and dispersed particulate phases. For this tutorial, however, the

coupling is not considered to reduce the amount of simulation time required.

Only one region for the continuous phase is required for this analysis. By default, the continuum **Physics 1** is automatically created when the mesh is imported. The continuum is edited so that appropriate physical models are selected for the water fluid region.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Pre-selected)
Time	**Steady**
Material	**Liquid**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**
Reynolds Averaged Turbulence	**K-Omega Turbulence**
Optional Models	**Lagrangian Multiphase**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Selecting the Lagrangian Phase Models

Create the Lagrangian phase and select the appropriate phase models.

The **Lagrangian Multiphase** model requires the selection of further models representing the characteristics of the Lagrangian phase.

To create the Lagrangian phase and select the phase models:

1. In the **Physics 1** continuum, right-click the **Models** > **Lagrangian Multiphase** > **Lagrangian Phases** node and create a new phase.
2. For the **Phase 1** phase, select the following models in order:

Group Box	Model
Enabled Models	**Residence Time** (Selected automatically)
Particle Type	**Material Particles**
Material	**Solid**
Equation of State	**Constant Density**
Optional Particle Forces	**Drag Force**

Group Box	Model
	Virtual Mass
Optional Models	**Turbulent Dispersion** **Erosion**
Track Sampling	**Track File**

3. Click **Close**.

   The **Phase 1** phase node appears as shown below:



4. Save the simulation.

Setting the Material Properties

Define the material properties of the continuous and Lagrangian phases.

This part of the tutorial involves working with material properties for two phases: continuous and Lagrangian. The default properties for water are acceptable for the continuous phase in this simulation. Sand particles are entrained in the liquid flow in this simulation and so the density of the Lagrangian phase must be set to the density of sand.

To set the material properties:

Specify the material properties for the Lagrangian phase.

1. Edit the **Physics 1** > **Models** > **Lagrangian Multiphase** > **Lagrangian Phases** > **Phase 1** > **Models** > **Solid** > **Al** > **Material Properties** node and set the following property:

Node	Property	Setting
Density > **Constant**	**Value**	`2650.0 kg/m^3`

2. Rename the **Al** node to **Sand**.
3. Save the simulation.

Setting the Lagrangian Phase Model Properties

Define the Lagrangian phase model properties. The Drag Force model uses the default method for the Drag Coefficient, which is the Schiller-Naumann correlation.

To set the properties for the Lagrangian phase models:

1. Edit the **Lagrangian Phases** > **Phase 1** > **Models** node and set the following properties:

Node	Property	Setting
Drag Force > **Drag Coefficient**	**Method**	**Schiller-Naumann**
Track File	**Vectors**	**Parcel.Centroid** **Particle.Velocity**



2. Save the simulation.

Defining Field Functions for the Restitution Coefficients

Create field functions to represent the restitution coefficients.

Restitution coefficients are used to predict the angle with which a particle rebounds off a solid wall. In this tutorial, the relations that were developed by Forder et al. are used:

$$\epsilon_T = 1 - 0.78a_p + 0.84a_p^2 - 0.21a_p^3 + 0.028a_p^4 - 0.022a_p^5$$

$$\epsilon_N = 0.988 - 0.78a_p + 0.19a_p^2 - 0.24a_p^3 + 0.027a_p^4$$

Represent these equations in STAR-CCM+ using user field functions. The quantities referred to in user field functions must come from the predefined system field functions or other user field functions. In this simulation you require a predefined system field function for the particle incidence angle.

To define field functions for the restitution coefficients:

1. Open the **Tools > Field Functions** node.



2. Scroll down the list until the **Particle Incidence Angle** function is visible.



3. Select **Particle Incidence Angle** and note the value of the *Function Name* property.



Having noted the required name of **ParticleIncidenceAngle**, it is now possible to create the function for the Normal Restitution Coefficient.

4. Right-click the **Tools > Field Functions** node and select **New > Scalar** from the pop-up menu.

   A new **User Field Function 1** node appears in the object tree.

5. Rename **User Field Function 1** to `ENormal`.
6. Select **ENormal** and set the following properties:

Property	Setting
Function Name	`NormalRestitutionCoefficient`
Dimensions	*Dimensionless*

Property	Setting
Definition	`0.988 - 0.78* $ParticleIncidenceAngle + 0.19*(pow($ParticleIncidenceAngle, 2)) - 0.024*(pow($ParticleIncidenceAngle, 3)) + 0.027*(pow($ParticleIncidenceAngle, 4))`

The same steps are now repeated for the tangential restitution coefficient.

7. Right-click the **Tools > Field Functions** node and select **New > Scalar** from the pop-up menu.

   A new **User Field Function 1** node appears in the object tree.
8. Rename **User Field Function 1** to `ETan`.
9. Select **ETan** and set the following properties:

Property	Setting
Function Name	`TangentialRestitutionCoefficient`
Dimensions	*Dimensionless*
Definition	`1 - 0.78* $ParticleIncidenceAngle + 0.84*(pow($ParticleIncidenceAngle, 2)) - 0.21*(pow($ParticleIncidenceAngle, 3)) + 0.028*(pow($ParticleIncidenceAngle, 4)) - 0.022*(pow($ParticleIncidenceAngle, 5))`

Setting the Lagrangian Phase Boundary Conditions

Define the Lagrangian phase boundary conditions.

The Lagrangian Phase Boundary Conditions are used to define the behavior of particles and droplets when they hit a solid boundary. Impact behavior can be defined for the following boundary types:

- Baffle

- Contact (interface)
- Wall

In this simulation, it is only necessary to define the boundary interaction models for the Wall boundary types, as there are no baffles or contact interfaces present. Parameters for the Erosion model are made in this section.

The field functions that you created in the previous section are now applied to the Lagrangian Phase Boundary Conditions.

To set the Lagrangian phase boundary conditions:

1. Edit the **Physics 1 > Models > Lagrangian Multiphase > Lagrangian Phases > Phase 1 > Boundary Conditions > Wall** node and set the following properties:

Node	Property	Setting
Physics Conditions > Mode	**Active Mode**	**Rebound**
Physics Values > Normal Restitution Coefficient	**Method**	**Field Function**
Field Function	**Scalar Function**	**ENormal**
Tangential Restitution Coefficient	**Method**	**Field Function**
Field Function	**Scalar Function**	**ETan**



2. Save the simulation.

Setting the Erosion Model

Select the appropriate erosion model.

When modeling erosion using CFD methods, it is necessary to select an erosion model that is based on a knowledge of the material undergoing erosion and the conditions under which the erosion is taking place. It is also important to appreciate that erosion modeling does not necessarily predict material loss with a consistent degree of accuracy and it is not recommended to estimate component lifetimes using CFD modeling alone. If estimating component lifetime is important, then additional experimental validation should be obtained to assess the effectiveness of the erosion model under appropriate conditions.

In this tutorial, use the Oka model initially, and use an additional user-defined model afterwards for comparison. The default coefficients that are given for the Oka model in STAR-CCM+ are for the sand erosion of carbon steel.

To set the erosion model:

1. In the **Continua** > **Physics 1** > **Models** > **Lagrangian Multiphase** > **Lagrangian Phases** > **Phase 1** > **Boundary Conditions** > **Wall** >

Physics Values node, select the **Erosion Ratio** node and set *Method* to **Oka**.



2. Select the **Erosion Ratio** > **Oka** node and review the model properties.



3. Save the simulation.

Setting Continuous Phase Boundary Conditions

Define the continuous phase boundary conditions.

The default continuous phase boundary conditions for the wall boundaries are sufficient for this analysis, and there are no user-definable parameters for the symmetry planes. However, the inlet and outlet boundaries must be modified.

To set the continuous phase boundary conditions:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Inflow** node and set the following properties:

Node	Property	Setting
Physics Conditions >	**Method**	**Intensity + Length Scale**

Node	Property	Setting
Turbulence Specification		
Velocity Specification	**Method**	**Components**
Physics Values > Turbulence Intensity > Constant	**Value**	`0.029`
Turbulent Length Scale > Constant	**Value**	`0.0037 m`
Velocity > Constant	**Value**	`[0.0, -12.901, 0.0] m/s`

2. Edit the **Regions > Fluid > Boundaries > Outflow** node and set the following properties:

Node	Property	Setting
Physics Conditions > Turbulence Specification	**Method**	**Intensity + Length Scale**
Physics Values > Turbulence Intensity > Constant	**Value**	`0.029`
Turbulent Length Scale > Constant	**Value**	`0.0037 m`

3. Save the simulation.

Setting Up the Injector

Define the injector that is used to introduce the particles into the solution domain.

The particles of the dispersed Lagrangian phase are introduced into the solution domain using an injector. The injector defines the initial state of the particles and their spatial distribution. In this tutorial, introduce particles near the inlet with a velocity matching that of the continuous phase.

A new derived part based on a presentation grid is created to define the injection points, as this improves the distribution of particles across the pipe section.

To set up the injector:

1. Right-click the **Derived Parts** node and select **New Part > Probe > Presentation Grid**.
2. In the *Create Presentation Grid* dialog, add the region Fluid to the **Input Parts**, click **Create**, and then **Close**.
3. Rename the **Derived Parts > presentation grid** node to `Injection Grid`.
4. Select the **Injection Grid** node and set the following values for the grid:

X Resolution	`100`
Y Resolution	`100`
Coordinate System	**Laboratory**
Point 1	`[0.11065, 0.22, -0.0265] m,m,m`
Point 2	`[0.137, 0.22, 0.0] m,m,m`
Normal	`[0.0, -1.0, 0.0] m,m,m`
Origin	`[0.0, 0.22, 0.0] m,m,m`

Create the injector:

5. Right-click the **Injectors** node and select **New**.
6. Select the **Injector 1** node and set *Type* to **Part Injector**.



7. Set *Inputs* to **Derived Parts > Injection Grid**.



8. Click **OK**.
9. Finally set *Lagrangian Phase* to **Phase 1** as shown.



The *Number of Points* property is 0 until the injection points are calculated during solution initialization. The conditions and values of the injector must now be set up to define the particles.

10. Edit the **Injectors** > **Injector 1** > **Conditions** node and set the following properties:

Node	Property	Setting
Flow Rate Specification	*Method*	**Mass Flow Rate**
Flow Rate Distribution	*Method*	**Per Injector**

Node	Property	Setting
Particle Size Specification	*Method*	**Particle Size**
Velocity Specification	*Method*	**Components**

The particular values for these quantities can now be set.

11. Edit the **Injectors** > **Injector 1** > **Values** node and set the following properties:

Node	Property	Setting
Velocity > **Constant**	*Value*	`[0.0, -12.901, 0.0] m/s`
Mass Flow Rate > **Constant**	*Value*	`0.027787 kg/s`
Parcel Streams	*Value*	`3`
Particle Diameter > **Constant**	*Value*	`2.75E-4 m`



12. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

In a Lagrangian multiphase simulation, it is important to set the maximum time for which a particle is tracked, as it is possible for

particles to become trapped within a fluid region (such as a stagnant zone) and remain there indefinitely.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following property:

Node	Property	Setting
Lagrangian Multiphase > Steady	**Maximum Residence Time**	`0.1`
	Verbosity	**High**

New stopping criteria are created based on the monitors that appeared when the segregated flow solver was selected. The goal is to achieve a residual of 1.0E-3 for the continuity and momentum quantities before stopping the solution.

2. Open the **Monitors** node.
3. Multi-select the following nodes:
   - **Continuity**
   - **X-momentum**
   - **Y-momentum**
   - **Z-momentum**



4. Right-click one of the selected nodes and select **Create Stopping Criterion from Monitor**.



5. Open the **Stopping Criteria** > **Continuity Criterion** node.

   A stopping criterion has been created for each of the monitors selected previously.



It is possible to change the properties for all the nodes in one operation, rather than modifying each node individually.

6. Multi-select the following nodes:

   - **Continuity Criterion**
   - **X-momentum Criterion**
   - **Y-momentum Criterion**
   - **Z-momentum Criterion**



7. In the *Properties* window, set **Logical Rule** to **And**.



8. Multi-select the following nodes:
   - **Continuity Criterion > Minimum Limit**
   - **X-momentum Criterion > Minimum Limit**
   - **Y-momentum Criterion > Minimum Limit**
   - **Z-momentum Criterion > Minimum Limit**



9. In the *Properties* window, set **Minimum Value** to `0.001`.
10. Save the simulation.

Visualizing the Solution

Create a scalar scene to view the results of the simulation.

Observe the pressure distribution on the symmetry planes as the solution develops.

To visualize the solution:

1. Right-click the **Scenes** node and select **New Scene > Scalar**.
2. Select the **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Parts** node and set *Parts* to **Regions > Fluid > Symmetry**.



3. Click **OK**.
4. Select the **Scalar 1 > Scalar Field** node and set *Function* to **Pressure**.
5. Use the mouse buttons to orient *Scalar Scene 1* as shown below.



6. Select the **Color Bar** node and set the following:

Setting	Value
Orientation	**Vertical** (Set this before setting *Width* and *Height*.)
Width	`0.1`
Height	`0.7`
Position	`[0.77, 0.11]`

Reporting and Plotting the Pressure Drop

Create reports of average pressure at the inlet and outlet boundaries, and use the reports in a field function that gives the pressure drop.

Experimental measurements are available for the pressure drop across the flow restrictor when the flowrate is 28.09 liters/s. These can be compared to the prediction of pressure drop for this simulation to give some estimate of the accuracy of results. To do this, it is necessary to create reports of average pressure at the inlet and outlet boundaries, and to use these reports in a user-defined field function that gives the

pressure drop. This field function can be plotted while running the simulation to estimate convergence.

The upstream static pressure was measured at a distance of two pipe diameters before the restrictor, and the downstream static pressure was measured at a distance of six pipe diameters after the restrictor. The inflow boundary corresponds to the upstream measurement, but a section plane must be created at the point of the downstream measurement. The section plane is created first:

1. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.

   The *Create Section* dialog appears in place above the object tree.
2. In the *Input Parts* box, select Fluid.
3. In the *Plane Parameters* box, set *Y origin* to `-0.239014 m`, the *X normal* to `0 m`, and the *Y normal* to `1 m`. The other parameters can be left as they are.
4. In the *Display* box, select **No Displayer** as shown below.



5. Click **Create**, and then **Close**.

   A new part that is called **plane section** is added to the object tree inside the **Derived Parts** node.

Create the reports that are used to extract the static pressure:

6. Right-click the **Reports** node and select **New Report > Surface Average**.



7. Rename the **Surface Average 1** node to `Upstream Pressure`.
8. Select the **Upstream Pressure** node and set *Scalar Field Function* to **Pressure**.
9. Click the right side of the *Parts* property. In the dialog that appears, expand the **Regions** and **Fluid** nodes and select the **Inflow** boundary as shown below.



10. Click **OK**.

A similar report is created for the downstream pressure:

11. Right-click the **Reports** node and select **New Report > Surface Average**
12. Rename the **Surface Average 1** node to `Downstream Pressure`.
13. Select the **Downstream Pressure** node and set the *Scalar Field Function* to **Pressure** and the *Parts* to **Plane Section**.

    The completed *Properties* window appears as shown below.



Two new field functions corresponding to these reports have been automatically added to the object tree.

14. Open the **Tools > Field Functions** node and scroll down to see the new nodes:



Examination of the properties for these nodes shows that the relevant field function names are **DownstreamPressureReport** and **UpstreamPressureReport** respectively. These function names are used to create an expression report for the pressure drop:

15. Right-click the **Reports** node and select **New Report > Expression**.
16. Rename the **Expression 1** node to `Pressure Drop`.
17. Select the **Pressure Drop** node.
18. Use (**Custom Editor**) to enter the *Definition* property as:

    `$UpstreamPressureReport - $DownstreamPressureReport`
19. Click **OK**.



20. Click (**Custom Editor**) for the *Dimensions* property. In the *Pressure Drop - Dimensions* dialog, set the **Pressure** dimension to `1` as shown below.



21. Click **OK**.
22. Finally, set the *Units* to `bar`.

    The *Properties* window for the **Pressure Drop** report appears as shown below.



A monitor and plot is created from this report.

23. Right-click the **Reports > Pressure Drop** node and select **Create Monitor and Plot from Report**.



The resulting monitor and plot can be seen in the object tree:



24. Right-click the **Plots > Pressure Drop Monitor Plot** and select **Open** from the pop-up menu. The empty plot is displayed.
25. Save the simulation.

Initializing and Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

Run the simulation in two stages. A steady state solution for the fluid phase in the absence of particles is obtained first, and then the particle tracks are obtained with a single additional iteration. This is a reasonable approach as two-way coupling has not been activated for the Lagrangian phase.

Fluid Phase

To deactivate the Lagrangian solver while the fluid phase solution is obtained:

1. Select the **Solvers > Lagrangian Multiphase** node and activate the **Solver Frozen** expert property.



To initialize the solution and run the simulation:

2. Click (**Initialize Solution**) in the *Solution* toolbar or select **Solution** > **Initialize Solution** from the menu.
3. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. You can ignore the warning message about smoothing in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.
4. Select the *Scalar Scene 1* tab to view the pressure distribution as the solution develops.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 1000 iterations are complete or the continuity residual criterion, 0.001, is satisfied for all appropriate quantities.

   The *Pressure Drop Monitor Plot* appears as shown below at the end of the simulation.



5. Click the *Scalar Scene 1* tab at the top of the *Graphics* window.

   A scene similar to the one shown below appears.



6. When the simulation has finished running, save it.

Lagrangian Phase

The Lagrangian Multiphase solver is reactivated to obtain the particle tracks and erosion predictions.

7. Select the **Solvers** > **Lagrangian Multiphase** node.
8. In the *Properties* window, deselect **Solver Frozen**.
9. Click (**Step**) in the *Solution* toolbar to perform a single iteration. It can take several minutes for the solver to compute all the particle tracks.
10. When the particle tracking is complete, save the simulation.

Visualizing the Results

Create a scalar scene to view the results of the simulation.

The main result of interest in this tutorial is the erosion rate across the restrictor surface. A scalar scene can be created showing this profile.

To visualize the results:

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
2. Open the **Scenes** > **Scalar Scene 2** node and select the **Displayers** > **Scalar 1** > **Parts** node.



3. Set *Parts* to **Regions** > **Fluid** > **Restrictor**.



4. Continuing within the same **Displayers > Scalar 1** node, select the **Scalar Field** node.
5. Set *Function* to **Erosion Rate**.
6. Select the **Scenes > Scalar Scene 2 > Displayers > Outline 1** node and clear the *Outline* property.

The view can be initialized using (**Save-Restore-Select Views**).

7. Click (**Save-Restore-Select Views**) and select **View > +Y > Up +Z** from the pop-up menus.
8. Select the **Color Bar** node and set the following:

Setting	Value
Orientation	**Vertical** (Set before setting *Width* and *Height*.)
Width	`0.1`
Height	`0.7`
Position	`[0.7, 0.1]`



9. Use (**Rubberband Zoom**) and mouse controls to zoom in on the eroded restrictor surface and position it as shown below.



The plot shows where the maximum erosion rates are expected.

Estimating the Overall Erosion Rate

Create a field function and report to estimate the overall erosion rate.

To obtain an estimate of the erosion rate in g/hr for the whole restrictor, it is necessary to create a field function that multiplies the computed erosion rate in each boundary face by the face area.

To estimate the overall erosion rate:

Create the field function.

1. Right-click the **Tools > Field Functions** node and select **New**.
2. Rename the **User Field Function 1** node to `Overall Erosion Rate`.
3. Edit the *Dimensions* property and set *Mass* to `1` and *Time* to `-1` as shown.



4. Click **OK**.
5. Set the *Function Name* to `OverallErosionRate`.
6. Edit the *Definition* property and enter the following definition in the dialog:

   ```
   $ErosionRate * mag($$Area)
   ```

7. Click **OK**.

   The completed *Properties* window appears as shown below.



It is possible to report the overall erosion rate in units other than kg/s by creating a user-defined unit. This is demonstrated by creating a unit for g/hr. The conversion factor is:

1 g/hr = 2.777778 x 10^{-7} kg/s

8. Right-click the **Tools > Units** node and select **New** from the pop-up menu.



9. Rename the **Units 1** node to `g/hr`.
10. Select the **g/hr** node and set *Description* to `g/hr`.
11. Edit the *Dimensions* property and set *Mass* to `1`.
12. Set *Time* to `-1`.
13. Click **OK**.
14. Enter the *Conversion* property as `2.777778E-7`.

    The completed *Properties* window is as shown below.



Create a report that uses the field function and unit:

15. Right-click the **Reports** node and select **New Report > Sum**.
16. Rename the **Sum 1** node to `Overall Erosion Rate`.
17. Set the *Scalar Field Function* to **Overall Erosion Rate**.
18. Click the right half of the *Parts* property, expand **Regions** and **Fluid** nodes and select **Restrictor**.
19. Set the *Units* to **g/hr** using the drop-down menu.

    The completed *Properties* window appears as shown below.



20. Right-click the **Overall Erosion Rate** report node and select **Run Report** from the pop-up menu.

    The *Output* window shows the result of running the report:



The predicted quantity is for a quarter annulus only. To get the full erosion rate for the real case, it is necessary to multiply by 4. This gives a predicted erosion rate of 0.63 g/hr for the whole component.

Comparison with Experimental Data

The simulation results are compared with experimental data.

Pressure Drop

For flow tests conducted at a flow rate of 28.09 liters/s, the pressure drop across the restrictor was found to be 5.44 bar. This value compares favorably with the predicted pressure drop of 5.99 bar at 28.57 liters/s obtained with the relatively coarse mesh that is used in this tutorial.

Erosion Rate

An erosion test for the restrictor device that is considered in this tutorial gave an average erosion rate of 0.58 g/hr over a 14 hour period. The flow rate and sand concentration that are used in this tutorial match those of the erosion test, and the predicted erosion rate using the Oka model is 0.63 g/hr.

Implementing Alternative Functions for Erosion Ratio

Erosion equations are often tuned for a particular material/sand combination that are based on experimental testing. In this section, implement the model that was developed by Zhang et al. at the Erosion/Corrosion Research Center in Tulsa as an example.

The model that is given by Zhang et al. is:

$$ER = C(BH)^{-0.59} F_s V_p^n F(\Theta)$$

$$F(\Theta) = 5.40\Theta - 10.11\Theta^2 + 10.93\Theta^3 - 6.33\Theta^4 + 1.42\Theta^5$$

where:

- ER is the Erosion Ratio.
- C and n are empirical constants that are given as 2.17×10^{-7} and 2.41 respectively.
- BH is the Brinell hardness of the wall material. For carbon steel, this is taken as 120 after Zhang et al.
- F_s is the particle shape coefficient, which is said to be 0.2 for fully rounded sand particles.

- V_p is the particle impact velocity.
- Θ is the particle impact angle in radians.

Implement the angle function in a field function, and reference this angle field function as a field function for erosion ratio:

1. Right-click the **Tools > Field Functions** node and select **New > Scalar**.
2. Rename the **User Field Function 1** node to `Zhang Angle Function`.
3. Select the **Zhang Angle Function** node and set the following properties:

Function Name	`ZhangAngleFunction`
Dimensions	*Dimensionless*
Definition	`(5.4 * $ParticleIncidenceAngle) - (10.11 * pow($ParticleIncidenceAngle, 2)) + (10.93 * pow($ParticleIncidenceAngle, 3)) - (6.33 * pow($ParticleIncidenceAngle, 4)) + (1.42 * pow($ParticleIncidenceAngle, 5))`

If you insert the numeric constants in the equation:

$$ER = C(BH)^{-0.59} F_s V_p^n F(\Theta)$$

the final equation to be implemented is:

$$ER = 2.17 \times 10^{-7} \cdot (120)^{-0.59} \cdot 0.2 \cdot V_p^{2.41} F(\Theta)$$

$$ER = 2.57497 \times 10^{-9} V_p^{2.41} F(\Theta)$$

4. Right-click the **Tools > Field Functions** node and select **New > Scalar**.
5. Rename the **User Field Function 1** node to `Zhang Erosion Ratio`.
6. Select the **Zhang Erosion Ratio** node and set the following properties:

Function Name	`ZhangErosionRatio`
Dimensions	*Dimensionless*

Definition	`2.57497E-9 * pow($ParticleVelocityMagnitude, 2.41) * $ZhangAngleFunction`

The field function must now be set in the Lagrangian model tree:

7. Select the **Continua > Physics 1 > Models > Lagrangian Multiphase > Lagrangian Phases > Phase 1 > Boundary Conditions > Wall > Physics Values > Erosion Ratio** node and set *Method* to **Field Function**.
8. Select the **Wall > Physics Values > Erosion Ratio > Field Function** node and set *Scalar Function* to **Zhang Erosion Ratio**.



With the new field function in place, it is only necessary to run a single iteration to obtain a further prediction of erosion.

9. Click (**Step**) in the **Solution** toolbar to perform a single iteration. It can take several minutes for the solver to compute all the particle tracks.
10. Save the simulation.
11. Run the **Overall Erosion Rate** report to obtain the mass loss rate from the new result:

```
Output - erosion

        Sum of Overall Erosion Rate

Part                                Value (g/hr)
-----------------------------       -------------
Restrictor                           1.097593e-01
                                    -------------
                        Total:       1.097593e-01
```

The erosion rate for the whole annulus is four times the predicted amount, which is 0.41 g/hr for the Zhang erosion model.

Summary

This tutorial demonstrated how to set up an erosion modeling analysis in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Selecting models for a Lagrangian phase.
- Setting up an injector for Lagrangian Multiphase modeling.
- Defining material properties that are required for the selected models.
- Defining boundary conditions.
- Setting solver parameters.
- Initializing and running the solver until the residual magnitudes are sufficiently small.
- Analyzing results using the visualization facilities.
- Analyzing results using field functions.

Bibliography

This tutorial references the following papers.

McCulloch, S., 2005. United States Patent No. US 6,886,598 B2, "Choke Restrictor Devices and Methods".

Wallace, M.S., 2001. "CFD-based erosion modeling of simple and complex geometries." PhD Thesis, University of Strathclyde.

Forder, A., Thew, M., Harrison, D., 1998. A numerical investigation of solid particle erosion experienced within oilfield control valves. Wear, Vol. 216, pp. 184-193.

Zhang, Y., Reuterfors, E.P., McLaury, B.S., Shirazi, S.A., Rybicki, E.R., 2007. "Comparison of computed and measured particle velocities and erosion in water and air flows." Wear, Vol. 263, pp. 330-338.

Eulerian: Dispersed Phase Acceleration

This tutorial demonstrates how to set up a dispersed phase acceleration analysis in STAR-CCM+. It simulates the acceleration of the dispersed phase due to the drag force of the continuous phase.

In this simulation, the continuous phase is air flowing at 5 m/s, and the dispersed phase is water particles flowing at 1 m/s. The geometry that is used in this simulation is shown below.



The tutorial introduces the phase interaction tool which allows you to define how the multiple phases interact. In this simulation, the drag force that the air flow exerts accelerates the particle phase.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Mesh
Converting to a Two-Dimensional Mesh
Splitting the Boundary by Angle
Renaming Continua, Regions, and Boundaries
Selecting the Physics Models
Creating Phases and Selecting Phase Models
Defining the Phase Interactions
Setting the Model Properties
Setting Initial Conditions and Reference Values
Setting Boundary Conditions
Setting Under-Relaxation Factors
Setting the Stopping Criteria
Creating a Scalar Scene
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Eulerian: Dispersed Phase Acceleration tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A three-dimensional, trimmed cell mesh has been prepared for this analysis.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/particlePipe.ccmg` then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.
5. Save the simulation as `particlePipe.sim`.

Visualizing the Mesh

Examine the mesh that has been provided for this tutorial.

A three-dimensional trimmed cell mesh has been prepared for this tutorial.

To visualize the mesh:

1. Click (**Show All Meshes**) in the **Vis** toolbar.

2. Use the mouse buttons to adjust the view so that it is similar to the view shown below.



Converting to a Two-Dimensional Mesh

Convert the three-dimensional mesh to two dimensions.

You wish to carry out a two-dimensional analysis, however the imported geometry is three-dimensional. The mesh is therefore converted to two-dimensions by using the **Convert to 2D** option. There are special requirements in STAR-CCM+ for three-dimensional meshes that are converted to two dimensions. These requirements are:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The mesh that is imported for this tutorial is built with these requirements in mind. Had the grid not met the above conditions, it would have been necessary to realign the region using the transformation and rotation facilities in STAR-CCM+.

To convert the mesh to two dimensions:

1. Select **Mesh** > **Convert to 2D**.



2. In the *Convert Regions to 2D* dialog, do the following:
   a) Make sure that the **Default_Fluid** region is selected.
   b) Activate **Delete 3D Regions After Conversion**.
   c) Click **OK**.



Once you click **OK**, the mesh conversion takes place and the new two-dimensional mesh is shown, viewed from the z-direction, in the *Geometry Scene 1* display. The mouse rotation option is suppressed for two-dimensional scenes.

The physics continuum that is associated with the deleted three-dimensional region is no longer required.

3. Expand the **Continua** node. Right-click the **Physics 1** continuum node and select **Delete**.



4. Click **Yes** in the confirmation dialog.

Splitting the Boundary by Angle

Split the boundary by angle to separate the four bounding walls.

To split the boundary by angle:

1. Expand the **Regions** > **Default_Fluid 2D** > **Boundaries** node.
2. Right-click the **Default_Boundary_Region** node and select **Split By Angle**.



3. In the *Split Boundaries by Angle* dialog that appears, do the following:
   a) Make sure that the **Default_Boundary_Region** is selected.

b) Leave the default value of **Angle (degrees)** unchanged.

c) Click **Preview** and make sure that splitting would result in 4 (in this case) separate boundaries.



4. Click **OK**.

Renaming Continua, Regions, and Boundaries

Change the default continuum, region, and boundary names to more suitable names.

To change the continuum, region, and boundary names:

1. Rename the continua **Physics 1 2D** to `Physics 1`.
2. Rename the region **Default_Fluid 2D** to `Fluid`.
3. Rename the boundaries within the **Fluid** node:

Original Name	New Name
Default_Boundary_Region	`Top`
Default_Boundary_Region 2	`Inlet`
Default_Boundary_Region 3	`Outlet`
Default_Boundary_Region 4	`Bottom`

4. Right-click the **Boundaries** node and select **Refresh** to rearrange the boundaries into alphabetical order.

   The object tree appears as follows:



Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this example, a simple laminar flow is set up in which water particles are accelerated in the flow of air. Use the Eulerian Multiphase option to allow you to define the two phases and the Multiphase Segregated Flow option to define how the two phases interact.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected) **Gradients** (Pre-selected)
Time	**Steady**
Material	**Eulerian Multiphase**
Eulerian Multiphase Model	**Multiphase Segregated Flow**
Viscous Regime	**Laminar**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

Create two new phases and specify them as air and water. Select the models that are associated with each phase. The air and particle phases are modeled as gas and liquid, respectively. They are both defined as having a constant density.

To create the phases and select the phase models:

Create the gas phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Air**.
3. For the **Air** phase, select the following models in order:

Group Box	Model
Material	**Gas**
Equation of State	**Constant Density**

4. Click **Close**.

Create the particle phase.

5. Create a second phase and rename it to **Particle**.
6. For the **Particle** phase, select the following models in order:

Group Box	Model
Material	**Liquid**
Equation of State	**Constant Density**

7. Click **Close**.

   The **Eulerian Phases** node appears as shown below.



Specify the material properties for the gas phase.

8. Edit the **Eulerian Phases** > **Air** > **Models** > **Gas** > **Air** > **Material Properties** node and set the following properties:

Node	Property	Setting
Density > **Constant**	**Value**	`1.0 kg/m^3`
Dynamic Viscosity > **Constant**	**Value**	`1.0E-5 Pa-s`

Specify the material properties for the particle phase.

9. Edit the **Eulerian Phases** > **Particle** > **Models** > **Liquid** > **H2O** > **Material Properties** node and set the following properties:

Node	Property	Setting
Density > **Constant**	**Value**	`2000 kg/m^3`
Dynamic Viscosity > **Constant**	**Value**	`1.0E-5 Pa-s`

10. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the air and particle phases. Specify the air phase as continuous, and the particle phase as dispersed.

In this tutorial, you are interested in the acceleration of the dispersed particles due to the drag force exerted by the continuous phase. You use the drag force multiphase interaction model with a constant drag force value.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to `Air-Particle`.
3. For the **Air-Particle** phase interaction, select the following models in order:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed**

Group Box	Model
Enabled Models	**Drag Force** (selected automatically) **Multiphase Material** (selected automatically) **Interaction Length Scale** (selected automatically) **Interaction Area Density** (selected automatically)

The *Air-Particle Model Selection* dialog is as shown below:



4. Click **Close**.
5. Edit the **Air-Particle > Models** node and set the following properties:

Node	Property	Setting
Drag Force > Drag Coefficient	**Method**	**Constant**
Drag Force > Drag Coefficient > Constant	**Value**	`0.44`
Eulerian Continuous-Dispersed	**Continuous Phase**	**Air**
	Dispersed Phase	**Particle**
Interaction Length Scale	**Interaction Length Scale**	`5e-4 m`



6. Save the simulation.

Setting the Model Properties

Define the model properties.

For most cases, the default model properties are appropriate. It is recommended that you do not change these properties unless you are an experienced user. In this simulation, use the default model properties.

To set the model properties:

1. In the **Physics 1** continuum, select the **Models** > **Multiphase Segregated Flow** node and ensure that the default property settings are used.

   The default **Multiphase Segregated Flow** properties are as shown below.



2. Save the simulation.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the physics continuum.

To set the initial conditions and reference values for the flow field:

1. Edit the **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and set the following properties:

Node	Property	Setting
Air > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[5.0, 0.0, 0.0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.99`
Particle > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[1.0, 0.0, 0.0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.01`



The sum of the volume fractions must be equal to one.

2. Edit the **Reference Values** node and set the following property:

Node	Property	Setting
Reference Pressure	**Value**	`0.0 Pa`

3. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

Set the boundary conditions for the **Inlet** and **Outlet** boundaries. First specify the velocity and volume fraction for each phase at the inlet. These values are set to the same values as the initial conditions to satisfy the continuity constraint.

To set the boundary conditions:

1. Edit the **Regions** > **Fluid** > **Boundaries** node and set the following property:

Node	Property	Setting
Inlet	**Type**	**Velocity Inlet**

2. Edit the **Regions** > **Fluid** > **Boundaries** > **Inlet** > **Phase Conditions** node and set the following properties:

Node	Property	Setting
Air > **Physics Values**		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`5.0 m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.99`
Particle > **Physics Values**		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`1.0 m/s`
Volume Fraction	**Method**	**Constant**

Node	Property	Setting
▮▮▮▮Constant	Value	`0.01`

3. Edit the **Regions** > **Fluid** > **Boundaries** node and set the following properties:

Node	Property	Setting
Outlet	Type	Flow-Split Outlet
▮▮Physics Values > Split Ratio	Split Ratio	`1.0`



The **Split Ratio** is 1.0 as all the flow exits through this boundary.

4. Save the simulation.

Setting Under-Relaxation Factors

Set the appropriate under-relaxation factors.

Specify a custom under-relaxation factor for the volume fraction solver.

To set the under-relaxation factor:

1. Select the **Solvers** > **Volume Fraction** node and make sure that the *Under-Relaxation Factor* is set to `0.5`.



Setting the Stopping Criteria

Set the appropriate stopping criteria for the simulation.

The simulation runs for 100 iterations.

To set the stopping criteria:

1. Edit the **Stopping Criteria** node and set the following property:

Node	Property	Setting
Maximum Steps	**Enabled**	Activated
	Maximum Steps	100



2. Save the simulation.

Creating a Scalar Scene

Create a scalar scene to view the results of the simulation.

A new scalar scene is created so that you can visualize the solution. This scene is used to display the velocity of the particle phase as it is accelerated in the flow of air.

1. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.

2. Click **Scene/Plot** above the object tree.



3. Select the **Displayers > Scalar 1 > Scalar Field** node.



The velocity components for the particle phase are denoted with the letters i, j and k, which correspond to the x, y and z components respectively. You are interested in the particle velocity in the x-direction.

4. Select i and set the *Function* to **Velocity > Velocity of Particle > Lab Reference Frame > Laboratory > i**.



5. Select the **Scalar 1** node.
6. Set the *Contour Style* to **Smooth Filled**.



7. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.



While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 100 iterations are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Steps satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

View the results of the simulation.

You can visualize the computed solution by viewing the scalar scene in the *Graphics* window.

To visualize the results:

1. Click the *Scalar Scene 1* tab at the top of the *Graphics* window.
2. Use the mouse buttons to zoom in on the scene.



Summary

This tutorial demonstrated how to set up a dispersed phase acceleration analysis in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the physics models.
- Creating phases and selecting phase models.
- Defining phase interaction.
- Setting model properties.
- Setting initial conditions and reference values.
- Setting boundary conditions.
- Setting under-relaxation factors.
- Setting stopping criteria.
- Creating a scalar scene.
- Running the simulation.
- Visualizing the results.

Eulerian: Hibiki's Bubble Column

This tutorial demonstrates how to model a Eulerian multiphase turbulent flow in STAR-CCM+.

This simulation models Hibiki's bubble flow column, which involves air and water phases being pumped through a vertical column. The geometry is shown in the diagram below.



The first part of the tutorial covers how to select the Eulerian multiphase and multiphase turbulence models, specify the properties of each phase and define the phase interaction.

In the second part of the tutorial, the S-Gamma model is activated to account for the size distribution of bubbles in the dispersed phase (air). The S-Gamma model predicts the Sauter mean diameter, which in turn is used to set the phase interaction length scale.

A validation section is also included where the results of the two simulations are compared with experimental data.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Mesh
Selecting the Physics Models
Creating Phases and Selecting Phase Models
Defining the Phase Interactions
Setting Initial Conditions and Reference Values
Setting Boundary Conditions
Setting Stopping Criteria and Solver Parameters
Setting Up the Plots
Running the Simulation
Exporting the Plot Data
Using the S-Gamma Model in the Simulation
Validating the Results
Summary
Bibliography

Prerequisites

The instructions in the Eulerian: Hibiki's Bubble Column tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.

2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/axisymmetricVerticalPipe.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `verticalPipe.sim`.

Visualizing the Mesh

Create a mesh scene and examine the mesh.

An axisymmetric hexahedral mesh of a vertical pipe has been prepared for this analysis, so further meshing is not required. The boundary names and types are pre-defined.

To visualize the mesh:

1. Create a mesh scene.

Examine the imported boundaries.

2. Expand the **Regions** > **Default Fluid** > **Boundaries** node and select the **Symmetry_Axis** node.



3. Make sure that the *Type* property is set to **Axis**.

4. Select the **Inlet** node and observe which part is highlighted in the scene. Make sure that the boundary *Type* is set to **Velocity Inlet**.
5. Select the **Outlet** boundary node and make sure that the *Type* is set to **Flow Split Outlet**.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

The multiphase turbulence models are available for selection once the multiphase flow and material models have been selected.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, in the *Enabled Models* group box, deactivate **Two-Dimensional**.
2. Select the following models in order:

Group Box	Model
Space	**Axisymmetric** **Gradients** (Selected automatically)
Time	**Steady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow** **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Turbulent**
Optional Models	**Gravity**

3. Click **Close**.

4. To review the models, open the **Physics 1** > **Models** node.



5. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

The two phases flowing through the section of vertical pipe are air and water. These phases are defined by creating new phases under the Eulerian physics model, and selecting the corresponding material properties. Create the liquid phase first:

To create the phases and select the phase models:

Create the liquid phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Water**.
3. For the **Water** phase, select the following models:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Liquid**

Group Box	Model
Enabled Models	**Reynolds-Averaged Navier-Stokes** (Selected automatically)

4. Deactivate **Auto-select recommended models**.
5. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

6. Activate **Auto-select recommended models**.
7. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**

8. Click **Close**.

Create the gas phase.

9. Create a second phase and rename it to **Air**.
10. For the **Air** phase, select the following models in order:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Gas**
Enabled Models	**Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence** **Realizable K-Epsilon Two-Layer** (Selected automatically) **Two- Layer All y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**

11. Click **Close**.

    The **Eulerian Phases** node appears as shown below.



Specify the material properties for the liquid phase.

12. Edit the **Eulerian Phases > Water > Models > Liquid > H2O > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1000 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`0.001 Pa-s`

Specify the material properties for the gas phase.

13. Edit the **Eulerian Phases > Air > Models > Gas > Air > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1 kg/m^3`

Node	Property	Setting
Dynamic Viscosity > Constant	**Value**	`1.8E-5 Pa-s`

14. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the water and air phases.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models > Multiphase Interaction > Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to **Water-Air**.
3. For the **Water-Air** phase interaction, deactivate **Auto-select recommended models**, and then select the following models in order:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed Drag Force** (selected automatically) **Multiphase Material** (selected automatically)
Interaction Length Scale	**Interaction Length Scale** (selected automatically)
Interaction Area Density	**Interaction Area Density** (selected automatically)
Optional Models	**Lift Force** **Turbulent Dispersion Force**

4. Click **Close**.

   The **Phase Interactions > Water-Air > Models** node appears as shown below.



Define the phase interaction properties.

5. Edit the **Water-Air > Models** node and set the following properties:

Node	Property	Setting
Drag Force > Drag Coefficient	**Method**	**Wang Curve Fit**
Eulerian Continuous-Dispersed	**Continuous Phase**	**Water**
	Dispersed Phase	**Air**
Interaction Length Scale	**Interaction Length Scale**	`0.003 m`
Lift Force > Lift Coefficient > Constant	**Value**	`-0.05`

The completed *Edit* dialog is shown below.



6. Save the simulation.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the physics continuum.

To set the initial conditions and reference values:

1. Edit the **Continua > Physics 1 > Models > Eulerian Multiphase > Eulerian Phases** node and set the following properties:

Node	Property	Setting
Water > Initial Conditions		
Turbulence Specification	**Method**	**Intensity + Length Scale**
Turbulent Length Scale	**Method**	**Constant**
Constant	**Value**	`0.005 m`
Velocity	**Method**	**Constant**
Constant	**Value**	`[1.23, 0.0, 0.0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.8`
Air > Initial Conditions		
Turbulence Specification	**Method**	**Intensity + Length Scale**
Turbulent Length Scale	**Method**	**Constant**
Constant	**Value**	`0.005 m`
Velocity	**Method**	**Constant**
Constant	**Value**	`[1.61, 0.0, 0.0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.2`

Set the reference values for the **Physics 1** continuum.

2. Edit the **Reference Values** node and set the following properties:

Node	Property	Setting
Reference Density	**Value**	`1000 kg/m^3`
Gravity	**Value**	`[-9.81, 0.0, 0.0] m/s^2`



3. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

To set the boundary conditions:

1. Go to **Regions > Default_Fluid > Boundaries > Inlet > Phase Conditions**. Select **Water > Physics Conditions > Turbulence Specification** and set *Method* to **Intensity + Length Scale**.
2. Edit the **Regions > Default_Fluid > Boundaries > Inlet > Phase Conditions** node and set the following properties:

Node	Property	Setting
Water > Physics Values		
Turbulent Length Scale	**Method**	**Constant**
Constant	**Value**	`0.005 m`

Node	Property	Setting
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`1.23 m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.8`
Air > Physics Values		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`1.61 m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.2`

3. Save the simulation.

Setting Stopping Criteria and Solver Parameters

Set the appropriate solver parameters and stopping criteria for the simulation.

Decrease the under-relaxation factors for the solvers to aid solution convergence.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Flow > Phase Coupled Velocity	**Under-Relaxation Factor**	`0.4`
Volume Fraction	**Under-Relaxation Factor**	`0.15`
K-Epsilon Turbulence	**Under-Relaxation Factor**	`0.5`



2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Steps	**Enabled**	Activated
	Maximum Steps	`1500`

3. Save the simulation.

Setting Up the Plots

Create a plot to view the simulation results.

A line probe is created and used to extract data across the radius of the bubble column. When creating a derived part, STAR-CCM+ automatically uses the active scene to display a preview of the part.

To set up the plots:

1. Right-click the **Derived Parts** node and select **New Part** > **Probe** > **Line**.



2. In the *Create Line Probe* panel, do the following:
   a) Make sure that **Default Fluid** is selected in the **Input Parts** box.
   b) Set the coordinates of **Point 1** to `[2.7, 0.0, 0.0] m`.
   c) Set the coordinates of **Point 2** to `[2.7, 0.025, 0.0] m`.
   d) In the **Display** box, select **No Displayer**.

   The panel appears as follows.



e) Click **Create** and then **Close**.

The x-axis in the plots indicates the radial position of the data points in the pipe. For convenience, use the normalized radial position (`r/R`), which can be specified using a field function:

3. Right-click the **Tools > Field Functions** node and select **New > Scalar**.
4. Rename the **User Field Function 1** node to `Normalized Radial Position`.
5. Set the *Function Name* to `Normalized Radial Position`.

6. Set the *Definition* to `$ $Position(@CoordinateSystem("Laboratory"))[1]/0.0254`

Note: Two dollar symbols (`$$`) are used when referencing a position field function.



The plot can now be created.

7. Click (**Create/Open Plots**) in the **Plot** toolbar and select **XY Plot**.
8. Rename the **Plots > XY Plot 1** node to `Volume Fraction of Air`.
9. Click **[]** next to the *Parts* property and select **line-probe** using the *Select Objects* dialog.
10. Select the **Volume Fraction of Air > X Type** node and set the *Type* to **Scalar**.



11. Select the **X Type** > **Scalar Function** node and set the *Field Function* property to **Normalized Radial Position**.



12. Select the **Y Types** > **Y Type 1** > **Scalar Function** node.



13. Set the *Field Function* property to **Volume Fraction** > **Volume Fraction of Air**.

The remaining plot shares many properties with the first plot, so you can copy and paste original plot to create the remaining plot and change some of the properties as necessary:

14. Right-click the **Plots** > **Volume Fraction of Air** node and select **Copy**, then right-click the **Plots** node and select **Paste**.
15. Rename the new plot `Axial Velocity of Air`.
16. Right-click the **Plots** node and select **Refresh**.



17. Select the **Axial Velocity of Air** > **Y Types** > **Y Type 1** > **Scalar Function** node and set the *Field Function* to **Velocity** > **Velocity of Air** > **Lab Reference Frame** > **Laboratory** > **i**.



18. To open the plots, right-click the **Plots** node and select **Open All Plots**. Use the tabs at the top of the *Graphics* window to display each plot.
19. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 1500 iterations are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Steps satisfied.
```

2. When the simulation has finished running, save it.

Exporting the Plot Data

Export the data from the plots so that it can be compared with the S-Gamma model and experimental data in the second part of the tutorial.

To export the plot data:

1. Right-click the **Volume Fraction of Air** node and select **Export**.
2. In the *Save* dialog that appears, save the file under the name `nonSGamma_voidFraction.csv`.
3. Repeat the steps for the other plot, naming the exported data for the **Axial Velocity of Air** plot as `nonSGamma_vAir.csv`.

Using the S-Gamma Model in the Simulation

This topic explains how the S-Gamma model is used in this simulation.

In dispersed two-phase flows, the strength of interphase forces, and the quantity of heat and mass transfer, are functions of the bubble size. Bubble sizes can change continuously due to coalescence and breakup, therefore being able to account for this size distribution leads to an improved prediction of these quantities. In STAR-CCM+, the S-Gamma method of moments has been implemented to account for particle size distributions in dispersed two-phase flows.

The S-Gamma model was developed to model the size distribution of dispersed bubbles with the consideration of breakup and coalescence effects:

- For breakup, the S-Gamma model considers the balance between disruptive forces (due to shear and turbulence) and restoring forces (due to surface tension) on the droplet.

- For coalescence, the S-Gamma model considers the probability of collisions of the droplets, the contact time of two colliding drops, and the drainage time of the liquid film between the drops.

Contents

Activating and Setting Up the S-Gamma Model
Modifying the Phase Interaction Settings
Setting S-Gamma Boundary Values
Clearing the Solution and Running the Simulation

Activating and Setting Up the S-Gamma Model

Activate the S-Gamma model and define the appropriate properties.

The S-Gamma model is activated in the *Phase Model Selection* dialog, where it can be selected from the *Optional Models* box.

To activate the S-Gamma model:

1. Save a copy of the simulation. Name the copy `verticalPipeSGamma.sim`.
2. Right-click the **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** > **Air** > **Models** node and choose **Select models...**.

In the *Air Model Selection* dialog:

3. Select the **S-Gamma** model from the *Optional Models* group box.
4. Click **Close**.

   The **Eulerian Phases** > **Air** node appears as shown below.



Define the S-Gamma properties.

5. Select the **Air** > **Models** > **S-Gamma** node and make sure that the *Number of S-Gamma Equations* property is set to **Two Equation**.

   This setting solves for both the transport of particle number density and the interfacial area density (zeroth and second moments).

6. Select the **Initial Conditions** > **Sauter Mean Diameter** > **Constant** node and set the *Value* to `0.003 m`.

Modifying the Phase Interaction Settings

Define the appropriate phase interaction settings.

Additional models that are related to the S-Gamma model are activated in the phase interaction node. The Constant Interaction Length Scale model is also replaced with the User-Defined Length Scale model.

To modify the phase interaction settings:

1. Select the **Physics 1** > **Models** > **Multiphase Interaction** > **Phase Interactions** and select the **Water-Air** > **Models** > **Interaction Length Scale** > **Interaction Length Scale** node, then set **Method** to **Field Function**.
2. Select **Interaction Length Scale** > **Field Function** and set **Scalar Function** to **Sauter Mean Diameter** > **Sauter Mean Diameter of Air**.

   The **Phase Interactions** > **Water-Air** > **Models** node appears as shown below.



3. Select **Phase Interactions** > **Water-Air** > **Models** and select the **S-Gamma Breakup** and **S-Gamma Coalescence** models from the *Optional Models* group box.
4. Click **Close**.

Define the phase interaction parameters that are related to the S-Gamma model:

5. Right-click the **Water-Air** > **Models** node and select **Edit**.

In the *Edit* dialog, do the following:

6. Expand the **Multiphase Material** > **Multiphase Material** > **Material Properties** > **Surface Tension** > **Constant** node and set the **Value** to `0.06 N/m`.
7. Expand the **S-Gamma Coalescence** node and make sure that the **Coalescence Factor** is set to `0.01`.
8. Expand the **S-Gamma Coalescence** > **Viscous Coalescence** node and set the **Viscous Coalescence Probability** to **Partially mobile (short collision time)**.

   The completed Edit dialog is shown below:


Models
Expand/Contract Tree
Expand/Contract Values
Nodes
Values
Models
Drag Force
Eulerian Continuous-Dispersed
Interaction Area Density
Interaction Length Scale
Interaction Length Scale
Method
Field Function
Dimensions
Length
Field Function
Scalar Function
SauterMeanDiameter of Air
Lift Force
Multiphase Material
Multiphase Material
Material Properties
Surface Tension
Dimensions
Force/Length
Method
Constant
Constant
Value
0.06 N/m
S-Gamma Breakup
S-Gamma Coalescence
Coalescence Factor
0.01
Hamaker Constant
5.0E-21
Viscous Coalescence
Viscous Collision Rate
Standard
Viscous Coalescence Probabil...
Partially mobile (short collision...
Inertial Coalescence
Turbulent Dispersion Force
Models
Model manager
Close
Help

9. Click **Close**.

Setting S-Gamma Boundary Values

Define the appropriate boundary conditions.

Define the Sauter mean diameter for the inlet boundary.

To set the S-Gamma bounday values:

1. Select the **Regions** > **Default Fluid** > **Boundaries** > **Inlet** > **Phase Conditions** > **Air** > **Physics Values** > **Sauter Mean Diameter** > **Constant** node and set the *Value* to `0.003 m`.



2. Save the simulation.

Clearing the Solution and Running the Simulation

Clear the results of the simulation and run the simulation again.

The stopping criteria and solution controls remain unchanged.

To clear the solution and run the simulation again:

1. Clear the solution from the previous run.
   Accept all default settings in the dialog that appears.
2. Click (**Run**) in the **Solution** toolbar.
3. Save the simulation when the solver finishes.

Validating the Results

Compare the results against experimental data.

The S-Gamma and non-S-Gamma solutions are validated against experimental data. The experimental data has been provided in two comma-separated values (.csv) files. You import this data in the simulation as tables.

To validate the results:

1. Right-click the **Tools** > **Tables** node and select **New Table** > **File Table**.
2. In the *Open* dialog, navigate to the `doc/startutorialsdata/multiphaseFlow/data` subdirectory of your STAR-CCM+ installation directory and multi-select `vAir.csv` and `voidFraction.csv`.
3. Click **Open**.
4. Using the same technique, import the .csv files that you exported in the first part of this tutorial:
   - `nonSGamma_vAir.csv`
   - `nonSGamma_voidFraction.csv`.

   The updated **Tables** node is shown below.



The tabular data is now added to the plots, starting with the **Volume Fraction of Air** plot:

5. Select the **Plots** > **Volume Fraction of Air** node.
6. Right-click the **Data Series** node and select **Add Data**.
7. In the Add Data Providers to Plot window, select **voidFraction**.
8. Select the **Data Series** > **voidFraction** node and do the following:
   a) Set the *Table* to **voidFraction** using the drop-down menu.

b) Set the *X Column* to **r/R**.

c) Set the *Y Column* to **Void Fraction of Air**.



9. Select the **voidFraction > Symbol Style** node and do the following:

   a) Set the *Color* to **Black**.

   b) Set the *Shape* to **Filled Circle**.

Add the non-S-Gamma data set to the plot.

10. Add a tabular data set as described previously and select the **Non-S-Gamma_voidFraction** data provider.



11. Select the **Data Series > Non-S-Gamma_voidFraction** node.

Do the following:

12. Set the *Table* to **nonSGamma_voidFraction** using the drop-down menu.
13. Set the *X Column* to **line-probe: Normalized Radial Position**.
14. Set the *Y Column* to **line-probe: Volume Fraction of Air**.
15. Select the **Non-S-Gamma_voidFraction > Line Style** node and do the following:
    a) Set the *Color* to a color of your choice. In this tutorial, **Blue** is used for the non-S-Gamma data set.
    b) Set the *Style* to **Solid**.
    c) Select the **Non-S-Gamma > Symbol Style** node and set the *Color* to **Blue**.

Rename the original data set for this simulation.

16. Select the **Volume Fraction of Air > Y Types > Y Type 1 > line-probe** node and set the *Legend Name* to `S-Gamma`.

Set the line type for the **S-Gamma** data set to **Solid**:

17. Select the **line-probe > Line Style** node and set the *Style* to **Solid**.

The *Volume Fraction of Air* plot appears as shown below.



Add the remaining data sets to the plot using the techniques described above.

18. Add the two data sets to the **Data Series** node in the **Axial Velocity of Air** plot. Rename one of the data sets to `vAir` and the other to `Non-S-Gamma_vAir`.
19. Right-click the **vAir** data set node and select **Edit**.
20. In the *Edit* dialog, set the properties as shown in the table below.

Node	Value
Table	vAir
X Column	r/R
Y Column	Axial Velocity of Air
Symbol Style > Color	Black
Symbol Style > Shape	Filled Circle

Do the same for the Non-S-Gamma data set.

21. Right-click the **Non-S-Gamma_vAir** data set node and select **Edit**.
22. In the *Edit* dialog, set the properties as shown in the table below.

Node	Value
Table	nonSGamma_vAir
X Column	Normalized Radial Position-line-probe
Y Column	Velocity of Air[i]-line-probe
Line Style > Color	Blue
Line Style > Style	Solid
Symbol Style > Color	Blue

Rename the original data set for this simulation:

23. Select the **Axial Velocity of Air > Y Types > Y Type 1 > line-probe** node and set the *Legend Name* to `S-Gamma`.

Set the line type for the **S-Gamma** data set to **Solid**:

24. Select the **line-probe > Line Style** node and set the *Style* to **Solid**.

The *Axial Velocity of Air* plot is shown below.



25. Save the simulation.

Contents

Validating the S-Gamma Model

Validating the S-Gamma Model

Compare the S-Gamma model results with experimental data.

The following section compares the plots of the Sauter mean diameter of the dispersed phase and the Interfacial Area Density to experimental data.

Create a field function that calculates the Interfacial Area Density. The Interfacial Area Density is defined by multiplying the **Second Size Distribution of Air** field function by π.

To validate the S-Gamma model:

1. Create a scalar field function and rename it to `Interfacial Area Density`, under the **Tools > Field Functions** node.



2. Select the **Interfacial Area Density** node and do the following:

   a) Click the area next to the *Dimension* property. In the dialog that appears, set the *Length* to `-1`.



   b) Click **OK**.

   c) Set the *Function Name* to `Interfacial Area Density`.

   d) Set the *Definition* to
   `$SecondSizeDistributionMomentAir*3.14`

Note: One dollar symbol ($) is used when referencing a scalar field function.

The completed *Properties* window is shown below.



Import the experimental data into STAR-CCM+.

3. Import the `interfacialAreaDensity.csv` and `meanAirDiameter.csv` files under the **Tools** > **Tables** node.

   These files can be found in the `doc/startutorialsdata/multiphaseFlow/data` subdirectory of your STAR-CCM+ installation directory.



The plots can now be created. The plot of the Interfacial Area Density is created first.

4. Create a **XY Plot** and rename it to `Interfacial Area Density`.

A new plot automatically appears in the *Graphics* window.

5. Right-click the **Interfacial Area Density** node and select **Edit**.
6. In the *Edit* dialog, set the properties as shown in the table below.

Node	Value
Parts	**line-probe**

Node	Value
X-Type > Type	**Scalar**
X-Type > Scalar Function > Field Function	**Normalized Radial Position**
Y-Types > Y-Type 1 > Scalar Function > Field Function	**Interfacial Area Density**
Y-Types > Y-Type 1 > line probe > Legend Name	**S-Gamma**
Y-Types > Y-Type 1 > line probe > Line Style > Color	**Red**
Y-Types > Y-Type 1 > line probe > Line Style > Style	**Solid**
Y-Types > Y-Type 1 > line probe > Symbol Style > Color	**Red**

7. Click **Close**.
8. Right-click the **Interfacial Area Density > Data Series** node and click **Add Data**.
9. In the Add Data Providers to Plot window, select **interfacialAreaDensity**.
10. Right-click the **Data Series > interfacialAreaDensity** node and select **Edit**.
11. In the *interfacialAreaDensity* dialog, set the properties as shown in the table below.

Node	Value
Table	**interfacialAreaDensity**
X Column	**r/R**
Y Column	**Interface_Area**
Symbol Style > Color	**Black**
Symbol Style > Shape	**Filled Circle**

12. Click **Close**.

Manually adjust the y-axis minimum value.

13. Select the **Interfacial Area Density > Axes > Y Axis** node, set the *Minimum* to `0.0`, and activate the checkbox next to *Lock Minimum*.

The *Interfacial Area Density* plot now has a minimum y value of 0, as shown in the following image:



Create and set up the plot of the Sauter Mean Diameter:

14. Create an **XY Plot** and rename it `Sauter Mean Diameter`.

A new plot automatically appears in the *Graphics* window.

15. Right-click the **Sauter Mean Diameter** node and select **Edit**.
16. In the *Edit* dialog, set the properties as shown in the table below.

Node	Value
Parts	**line-probe**
X-Type > Type	**Scalar**
X-Type > Scalar Function > Field Function	**Normalized Radial Position**
Y-Types > Y-Type 1 > Scalar Function > Field Function	**Sauter Mean Diameter > SauterMeanDiameter of Air**
Y-Types > Y-Type 1 > Scalar Function > Units	**mm**
Y-Types > Y-Type 1 > line probe > Legend Name	**S-Gamma**

Node	Value
Y-Types > Y-Type 1 > line probe > Line Style > Color	**Red**
Y-Types > Y-Type 1 > line probe > Line Style > Style	**Solid**
Y-Types > Y-Type 1 > line probe > Symbol Style > Color	**Red**

17. Click **Close**.
18. Right-click the **Sauter Mean Diameter > Data Series** node and click **Add Data**.
19. In the Add Data Providers to Plot window, select **meanAirDiameter**.
20. Right-click the **Data Series > meanAirDiameter** node and select **Edit**.
21. In the *Edit* dialog, set the properties as shown in the table below.

Node	Value
Table	**meanAirDiameter**
X Column	**r/R**
Y Column	**Mean_Diameter_of_Air**
Symbol Style > Color	**Black**
Symbol Style > Shape	**Filled Circle**

22. Click **Close**.

Manually adjust the y-axis minimum value.

23. Select the **Sauter Mean Diameter > Axes > Y Axis > Labels** node, set the *Minimum* to `0.0`, and activate the checkbox next to *Lock Minimum*.

    The *Sauter Mean Diameter* plot has a minimum y value of 0, as shown below:



Summary

This tutorial demonstrated how to model a Eulerian multiphase turbulent flow in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the Eulerian Multiphase and Multiphase Turbulence models
- Creating phases and selecting phase models
- Defining phase interactions
- Exporting results
- Activating the S-Gamma model
- Validating the results.

Bibliography

This tutorial references the following paper:

Hibiki, T., Ishii, M., Xiao, Z., 2001. "Axial interfacial area transport of vertical bubbly flows." International Journal of Heat and Mass Transfer, Vol. 44.

Eulerian: Mixture Settling

This tutorial demonstrates how to model mixture settling in STAR-CCM+. It simulates the settling of a two-phase mixture in a rectangular tank. The mixture inside the tank is assumed to be 90% water and 10% air.

A two-dimensional mesh is used to represent the tank. The geometry that used in this simulation is shown below:



The boundaries representing the tank are assumed to be no-slip walls. In this simulation, the implicit unsteady solver is used to study the separation of two phases with different densities within the tank. The buoyancy effect causes the lighter phase (air) to rise up and settle at the top of the heavier phase (water).

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Mesh
Converting to a Two-Dimensional Mesh
Reorienting the Mesh

Splitting the Boundary by Angle
Selecting the Physics Models
Creating Phases and Selecting Phase Models
Defining the Phase Interactions
Setting the Model Properties
Setting Initial Conditions
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Setting Up a Monitor
Creating a Scalar Scene
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Eulerian: Mixture Settling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A three-dimensional, trimmed cell mesh has been prepared for this analysis, which is subsequently converted to a two-dimensional mesh.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.

4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/tank.ccmg`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `settlingTank.sim`.

Visualizing the Mesh

Examine the imported mesh.

To visualize the imported mesh:

1. Right-click the **Scenes** node and then select `New Scene > Mesh`.
2. Use the mouse buttons to adjust the view so that it is similar to that shown below.



Converting to a Two-Dimensional Mesh

Convert the three-dimensional mesh to two dimensions.

To carry out a two-dimensional analysis, you need to convert the imported mesh to two-dimensions using the **Convert to 2D** option.

STAR-CCM+ has special requirements for converting three-dimensional meshes to two dimensions:

- The grid must be aligned with the X-Y plane.
- The grid must have a boundary plane at the Z = 0 location.

The imported mesh for this tutorial was built with these requirements in mind.

To convert the mesh to two dimensions:

1. Select **Mesh** > **Convert to 2D**.



2. In the *Convert Regions To 2D* dialog, do the following:
   a) Check that the **Default_Fluid** region is selected.
   b) Activate the **Delete 3D Regions After Conversion** option.



c) Click **OK**.

Once you click **OK**, the mesh conversion takes place and the new two-dimensional mesh is shown, viewed from the z-direction, in the *Mesh Scene 1* display. The mouse rotation option is suppressed for two-dimensional scenes.

Reorienting the Mesh

Change the orientation of the mesh to the appropriate direction.

You now use region rotation and translation to change the orientation of the imported mesh to make sure that the gravity vector acts in the correct direction.

To reorient the mesh:

1. Expand the **Regions** node, then right-click the **Default_Fluid 2D** node and select **Transform** > **Rotate...**.
2. In the *Rotate Regions* dialog, set the *Angle* to `90 deg`.



3. Click **Apply**, then **Close**.
4. Right-click the **Default_Fluid 2D** node for the second time and select **Transform > Translate**. In the *Translate Regions* dialog set the *X* component to `0.1 m`.



5. Click **Apply**, then **Close**.
6. Click (**Reset View**) in the visualization toolbar to reset the view as shown below.



You can perform the mesh diagnostics check by selecting **Mesh > Diagnostics** from the menu bar. The extent of the geometry is:

```
x: [-3.8858e-16, 1.0000e-01] m
y: [-3.8286e-17, 1.0000e-00] m
```

The physics continuum that was associated with the deleted three-dimensional region is no longer required and can be deleted.

7. Expand the **Continua** node. Right-click the **Physics 1** continuum node and select **Delete**.



8. Click **Yes** in the *Confirm Object Deletion* dialog.

Splitting the Boundary by Angle

Split the boundary by angle to separate the four bounding walls.

These boundaries are later defined as walls.

To split the boundary by angle:

1. Expand the **Regions** > **Default_Fluid 2D** > **Boundaries** node.
2. Right-click the **Default_Boundary_Region** node and select **Split By Angle**.



3. In the *Split Boundaries by Angle* dialog, click **Preview** and make sure that splitting would result in four separate boundaries.



4. Click **OK**.

The model currently has the default region and boundary names. Change these names to something more suitable to make setting up the simulation and understanding the model mechanics easier.

5. Change the region and boundary names as follows:

Node	Name
Default_Fluid 2D	`Fluid`
Default_Boundary_Region	`Right Wall`
Default_Boundary_Region 2	`Bottom Wall`
Default_Boundary_Region 3	`Top Wall`
Default_Boundary_Region 4	`Left Wall`

6. Right-click the **Boundaries** node and select **Refresh** to rearrange the boundaries into alphabetical order.

When you have renamed these boundaries, the object tree appears as follows:



Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this simulation, a simple laminar flow is set up to study the separation of two phases with different densities within a settling tank. Use the multiphase mixture option to allow you to define the two phases, and the segregated multiphase option to define how the two phases interact.

To select the physics models:

1. Rename the **Continua** > **Physics 1 2D** node to **Physics 1**.
2. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow**

Group Box	Model
	Multiphase Equation of State (Selected automatically)
Viscous Regime	**Laminar**
Optional Models	**Gravity**

3. Click **Close**.
4. To review the models, open the **Physics 1** > **Models** node.



5. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

Create two new phases and specify them as air and water. Select the models that are associated with each phase. The water and air phases are modeled as liquid and gas, respectively. They are both defined as having a constant density. Once the phase models have been selected, you can set the material properties for each phase.

To create the phases and select the phase models:

Create the liquid phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Water**.
3. For the **Water** phase, select the following models in order:

Group Box	Model
Enabled Models	**Laminar** (Selected automatically)
Material	**Liquid**
Equation of State	**Constant Density**

4. Click **Close**.

Create the gas phase.

5. Create a second phase and rename it to **Air**.
6. For the **Air** phase, select the following models in order:

Group Box	Model
Enabled Models	**Laminar** (Selected automatically)
Material	**Gas**
Equation of State	**Constant Density**

7. Click **Close**.

   The **Eulerian Phases** node appears as shown below.



Specify the material properties for the liquid phase.

8. Edit the **Eulerian Phases > Water > Models > Liquid > H2O > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1000 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`0.001 Pa-s`



Specify the material properties for the gas phase.

9. Edit the **Eulerian Phases > Air > Models > Gas > Air > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`1.0E-5 Pa-s`

10. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the water and air phases. The water phase is specified as continuous, while the air phase is specified as dispersed.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to **Water-Air**.
3. For the **Water-Air** phase interaction, select the following models:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed**
Enabled Models	**Drag Force** (selected automatically) **Multiphase Material** (selected automatically) **Interaction Length Scale** (selected automatically) **Interaction Area Density** (selected automatically)



4. Click **Close**.
5. Edit the **Water-Air** > **Models** node and set the following properties:

Node	Property	Setting
Eulerian Continuous-Dispersed	**Continuous Phase**	**Water**
	Dispersed Phase	**Air**

Node	Property	Setting
Drag Force > Drag Coefficient	**Drag Coefficients**	**Schiller-Naumann**
Drag Force > Drag Correction > Volume Fraction Exponent	**Volume Fraction Exponent (continuous)**	0



6. Save the simulation.

Setting the Model Properties

Define the model properties.

The default model properties are appropriate for most cases. However, given that the mesh is rather coarse in this simulation, you need to change the velocity convection method to first order.

To set the model properties:

1. In the **Physics 1** continuum, select the **Models** > **Multiphase Segregated Flow** node and set the following properties:

Property	Setting
Velocity Convection	**1st-order**
Volume Fraction Minimum	`1E-15`



2. Save the simulation.

Setting Initial Conditions

Set the initial conditions for the physics continuum.

Initialize the velocities within the tank and set the volume fraction for each of the phases. In this tutorial, you assume that the tank contains 90% water and 10% air. Note that the sum of volume fractions must be equal to one.

To set the initial conditions:

Set the initial conditions for the **Physics 1** continuum.

1. Edit the **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and set the following properties:

Node	Property	Setting
Water > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[0, -0.1, 0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.9`
Air > Initial Conditions		
Velocity		**Constant**
Constant	**Value**	`[0, 0.1, 0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.1`



2. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

By default, the conditions for boundaries representing the tank are set to no-slip walls. Check the conditions for each boundary and make any necessary changes.

To set the boundary conditions:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Bottom Wall** node and set the following properties:

Node	Property	Setting
Phase Conditions > Water > Physics Conditions > Shear Stress Specification	**Method**	**No-Slip**
Phase Conditions > Air > Physics Conditions > Shear Stress Specification	**Method**	**No-Slip**



2. Similarly, check the boundary setting for the other boundaries.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

As you are solving an unsteady problem, you need to specify the time-step size and the elapsed simulation time. This calculation runs for a maximum physical time of 18 seconds with a time-step size of 0.02 s.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.02 s`
Multiphase Segregated Flow >	**Under-Relaxation Factor**	`0.25`

Node	Property	Setting
Phase Coupled Velocity		
Multiphase Segregated Flow > Pressure	**Under-Relaxation Factor**	`0.1`
Volume Fraction	**Under-Relaxation Factor**	`0.1`



2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`18 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`25`
Maximum Steps	**Enabled**	Deactivated



3. Save the simulation.

Setting Up a Monitor

Create a monitor to help assess convergence of the simulation.

By default STAR-CCM+ produces a residuals plot that you can use to monitor the convergence of the solution. Additionally it is useful to create other monitors to help in assessing convergence of the solution. In this tutorial, you monitor the volume fractions at the center of the fluid domain.

To set up a monitor:

1. Right-click the **Derived Parts** node and select **New Part** > **Probe** > **Point**.

   The *Create Point* dialog appears in the *edit* window.
2. In the *Input Parts* box, make sure that **Fluid** is selected.
3. Set the *Point* coordinates to `[0.05, 0.5, 0]` m.
4. In the *Display* box, select **No Displayer**.



5. Click **Create** and then **Close**.

   A new derived part is created called **point**.



This point is used to monitor the water and air volume fractions in the fluid. Before you can create the monitor, create the report upon which these monitors are based. In this case, you create a Maximum report.

6. Right-click the **Reports** node and select **New Report** > **Maximum**.



A new report is created named **Maximum 1**, which is used to monitor the water volume fraction. Give the report a more suitable name.

7. Rename the **Maximum 1** node to `Water Volume Fraction`.
8. Select the **Reports** > **Water Volume Fraction** node and click to the right of the *Scalar Field Function* property.
9. In the *Water Volume Fraction - Scalar Field Function* window, select **Volume Fraction** > **Volume Fraction of Water**.



10. Click to the right of the *Parts* property. Expand the **Derived Parts** node, select **point** and click **OK**.



Create the report for the air volume fraction by copying and pasting the report for the water volume fraction.

11. Right-click the **Reports > Water Volume Fraction** node and select **Copy**.
12. Right-click the **Reports** node and select **Paste**.

    A new report is created, named **Copy of Water Volume Fraction**.

Give the report a more suitable name.

13. Rename the **Copy of Water Volume Fraction** node to `Air Volume Fraction`.
14. Select the **Reports > Air Volume Fraction** node and click to the right of the *Scalar Field Function* property.
15. In the *Air Volume Fraction - Scalar Field Function* window, select **Volume Fraction > Volume Fraction of Air**.



16. Multi-select the **Air Volume Fraction** and **Water Volume Fraction** nodes. Right-click the selected nodes and select **Create Monitor and Plot from Report**.
17. In the *Create Plot From Reports* dialog, click **Single Plot**.



A new monitor and plot are created.



Creating a Scalar Scene

Create a scalar scene so that you can visualize the solution.

This report is used to display the water volume fraction.

To create a scalar scene:

1. Right-click the **Scenes** node, and select **New Scene > Scalar**.
2. Click **Scene/Plot** above the object tree.
3. Select the **Displayers > Scalar 1 > Scalar Field** node and set the *Function* to **Volume Fraction > Volume Fraction of Water**.
4. Select the **Scalar 1** node and set the *Contour Style* to **Smooth Filled**.
5. Select the **Color Bar** node and do the following:
   a) Set the *Width* to `0.1`.
   b) Set the *Height* to `0.8`.
   c) Set the *Orientation* to **Vertical**.
   d) Set the *Position* to `[0.7, 0.1]`.



6. Click the **Simulation** button to return to the simulation object tree.
7. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

View the results of the simulation.

You can visualize the computed solution by viewing the scalar scene in the *Graphics* window.

To visualize the results:

1. Click the *Scalar Scene 1* tab at the top of the *Graphics* window.



As expected, water settles down to the bottom of the tank. Air, being lighter than water, moves to the top of the tank.

2. Right-click the **Plots** > **Reports Plot** and select **Open**.

   The Reports Plot is shown in the following screenshot.



The volume fractions of water and air at the mid point of the fluid domain settle to the expected values of 1.0 and 0.0, respectively. This result clearly indicates that the solution is stable and is converged.

Summary

This tutorial demonstrated how to model mixture settling in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the physics models.
- Creating phases and selecting phase models.
- Defining phase interaction.
- Setting model properties.
- Running the simulation.
- Visualizing the results.

Eulerian: Bubble Formation in a Fluidized Bed

This tutorial demonstrates how to model bubble formation in a fluidized bed in STAR-CCM+. It simulates the injection of gas into a fluidized bed and monitors the formation and growth of the bubble.

Fluidized beds are formed when a fluid is passed through a solid particulate substance to cause the mixture to flow like a fluid. The formation and propagation of bubbles within the bed affect its heat and mass transfer properties.

The shape of the bubble as well as its quantitative growth has been reported. The gas injection orifice is 15 mm wide. The superficial gas velocity at the orifice is 10 m/s. The rest of the inlet has a superficial gas velocity of 0.25 m/s; the minimum gas velocity to fluidize the bed. The geometry is shown below.



Index	Description
1	Air outlet
2	Inlet with air orifice
3	Bubble
4	Fluidized bed
5	Free surface

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting Up the Physics Models
Creating Phases and Selecting Phase Models

Defining the Phase Interactions
Setting Initial Conditions
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Setting Up a Monitor
Creating a Scalar Scene
Running the Simulation
Visualizing the Results
Summary
Bibliography

Prerequisites

The instructions in the Eulerian: Bubble Formation in a Fluidized Bed tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A two-dimensional, trimmed cell mesh has been prepared for this analysis. This mesh is relatively coarse, suitable only for a tutorial.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/fluidizedbed.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.

5. Save the simulation as `BubbleFormation.sim`.

Setting Up the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

The main phenomenon to solve for is the interaction of a granular Eulerian phase with a fluid Eulerian phase.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected) **Gradients** (Pre-selected)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow** **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Laminar**
Optional Models	**Gravity** **Granular Pressure** **Granular Temperature** (Selected automatically) **Granular Temperature Transport** (Selected automatically)

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



You now set the appropriate properties.

4. Select **Continua** > **Physics 1** > **Models** > **Granular Pressure** and set **Frictional Stress Formulation** to **Schaeffer**.
5. Select the **Granular Pressure** > **Schaeffer** node and set the **Angle of Internal Friction** to `25`.

Use second-order velocity convection.

6. Select the **Continua** > **Physics 1** > **Models** > **Multiphase Segregated Flow** node and set the **Volume Fraction Convection** property to **2nd-order**.
7. Make sure that the **Velocity Convection** is **2nd-order**.
8. Save the simulation.

Creating Phases and Selecting Phase Models

Create two new phases, specify them as air and particle, and set the appropriate material properties for each phase.

To create the phases and select the phase models:

Create the gas phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Air**.
3. For the **Air** phase, select the following models:

Group Box	Model
Enabled Models	**Laminar** (Selected automatically)
Material	**Gas**
Equation of State	**Constant Density**

4. Click **Close**.

   The **Air** phase node appears as shown below:



Create the particle phase.

5. Create a second phase and rename it to **Particle**.
6. For the **Particle** phase, select the following models:

Group Box	Model
Enabled Models	**Laminar** (Selected automatically)
Material	**Particle**
Enabled Models	**Constant Density** (Selected automatically) **Granular Temperature Transport** (Selected automatically)

Group Box	Model
	Granular Pressure (Selected automatically)

7. Click **Close**.

   The **Particle** phase node appears as shown below:



Specify the material properties for the gas phase.

8. Edit the **Eulerian Phases** > **Air** > **Models** > **Gas** > **Air** > **Material Properties** node and set the following properties:

Node	Property	Setting
Density > **Constant**	**Value**	`1.28 kg/m^3`
Dynamic Viscosity > **Constant**	**Value**	`1.7E-5 Pa-s`

Specify the material properties for the particle phase.

9. Edit the **Eulerian Phases** > **Particle** > **Models** > **Particle** > **Al** > **Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`2660.0 kg/m^3`
Particle Diameter > Constant	**Value**	`5.0E-4 m`

Set the maximum solid fraction.

10. Edit the **Eulerian Phases > Particle > Models > Granular Pressure** node and set the **Maximum Solid Fraction** to `0.65`.

Specify a reference value for the particle coefficient of restitution.

11. Edit the **Eulerian Phases > Particle > Reference Values > Coefficient of restitution** node and set the **Value** to `0.9`.
12. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the air and particle phases.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models > Multiphase Interaction > Phase Interactions** node and create a new phase interaction.
2. For the **Phase Interaction 1** phase interaction, select the following models:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed**
Enabled Models	**Drag Force** (selected automatically) **Multiphase Material** (selected automatically) **Interaction Length Scale** (selected automatically) **Interaction Area Density** (selected automatically)

3. Click **Close**.
4. Edit the **Phase Interaction 1 > Models** node and set the following properties:

Node	Property	Setting
Eulerian Continuous-Dispersed	**Continuous Phase**	**Air**
	Dispersed Phase	**Particle**
Drag Force	**Drag Force Coefficient Type**	**Linearized**
Drag Force > Linearized Drag Coefficient > Gidaspow	**Transition Volume Fraction**	`0.2`

5. Click **Close**.
6. Save the simulation.

Setting Initial Conditions

Set the initial conditions for the physics continuum.

Initially, the fluidized bed contains 40.2% particles and is 0.5 m deep. Before bubble formation, the particles are motionless and the air velocity sustains the fluidization of the bed. You can use field functions to set the initial volume fraction for each phase and the initial velocity of air.

To set the initial conditions:

Create the field functions for the initial volume fraction for each phase and the initial velocity of air.

1. Define the field function for the initial particle volume fraction as follows:
   a) Right-click the **Tools > Field Functions** node and select **New > Scalar**.
   b) Rename the **User Field Function 1** node to `InitialVFParticle`.
   c) Select the **InitialVFParticle** node and set the following properties:

Property	Setting
Function Name	`InitialVFParticle`
Definition	`($$Centroid[1] < 0.5) ? 0.402 : 1e-6.`

2. Use the same procedure to create field functions for the initial air volume fraction and the initial velocity of air, using the details below:

Name / Function Name	Definition
`InitialVFAir`	`($$Centroid[1] < 0.5) ? 0.598 : (0.999999)`
`InitialVelAir`	`0.25/$InitialVFAir`

Set the initial conditions for the **Physics 1** continuum.

3. Edit the **Continua > Physics 1 > Models > Eulerian Multiphase > Eulerian Phases** node and set the following properties:

Node	Property	Setting
Air > Initial Conditions		
Velocity	**Method**	**Composite**
Composite > Y Component	**Method**	**Field Function**
Field Function	**Scalar Function**	**InitialVelAir**
Volume Fraction	**Method**	**Field Function**
Field Function	**Scalar Function**	**InitialVFAir**
Particle > Initial Conditions		
Volume Fraction	**Method**	**Field Function**
Field Function	**Scalar Function**	**InitialVFParticle**



4. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

The boundaries representing the bed are no-slip walls, a velocity inlet, and a pressure outlet.

To set the boundary conditions:

Check the **LHS** and **RHS** boundary conditions.

1. Edit the **Regions** > **Fluidized_Bed** > **Boundaries** > **LHS** node and set the following property:

Node	Property	Setting
Phase Conditions > **Particle** > **Physics Conditions** > **Granular Temperature Specification**	**Option**	**Extrapolation**

2. Similarly, edit the **Regions** > **Fluidized_Bed** > **Boundaries** > **RHS** node.

To model the bubble formation, the inlet boundary requires field functions for the particle and gas volume fractions. Only gas enters through the inlet, but the granular flow model also requires an estimate of the kinetic energy of the particles.

3. Create the following scalar field functions:

Name / Function Name	Definition
`InletVFAir`	`1-$InletVFParticle`
`InletVFParticle`	`($$Centroid[0] > 0.2775) && ($$Centroid[0] < 0.2925) ? 0 : 0.402`
`InletVelAir`	`($$Centroid[0] > 0.2775) && ($$Centroid[0] < 0.2925) ? (10/$InletVFAir) : (0.25/$InletVFAir)`

Set the boundary conditions.

4. Edit the **Regions** > **Fluidized_Bed** > **Boundaries** > **Bottom** > **Phase Conditions** node and set the following properties:

Node	Property	Setting
Air > Physics Values		
Velocity Magnitude	**Method**	**Field Function**
Field Function	**Scalar Function**	**InletVelAir**
Volume Fraction	**Method**	**Field Function**
Field Function	**Scalar Function**	**InletVFAir**
Particle > Physics Values		
Volume Fraction	**Method**	**Field Function**
Field Function	**Scalar Function**	**InletVFParticle**

Allow the particles to slip along the wall boundaries.

5. Edit the **Regions** > **Fluidized_Bed** > **Boundaries** node and set the following properties:

Node	Property	Setting
LHS > Phase Conditions > Particle > Physics Conditions > Shear Stress Specification	**Method**	**Slip**
RHS > Phase Conditions > Particle > Physics Conditions > Shear Stress Specification	**Method**	**Slip**

Allow gas to escape through the **Top** boundary.

6. Edit the **Regions > Fluidized_Bed > Boundaries > Top** node and set the following property:

Node	Property	Setting
Phase Conditions > Particle	**Type**	**Phase Impermeable**

7. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

As you are solving an unsteady problem, it is necessary to specify the time-step size and the elapsed simulation time. This calculation runs for a maximum physical time of 0.2 seconds with a time-step size of 5.0E-4 s.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`5.0E-4 s`
Multiphase Segregated Flow > Phase Coupled Velocity	**Under-Relaxation Factor**	`0.1`

Node	Property	Setting
Multiphase Segregated Flow > Pressure	**Under-Relaxation Factor**	`0.1`
Volume Fraction	**Under-Relaxation Factor**	`0.1`



2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`0.2 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`20`
Maximum Steps	**Enabled**	Deactivated

3. Save the simulation.

Setting Up a Monitor

Create a monitor and plot of the bubble diameter to help assess convergence of the simulation.

By default, STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the solution. It is good practice to create more monitors on physical quantities to help in assessing convergence.

The bubble edge is defined where an 85% air concentration appears in the fluidized bed.

To monitor the bubble diameter:

1. Right-click the **Derived Parts** node and select **New Part > Isosurface**.

   The *Create Isosurface* dialog appears in the *edit* window.
2. In the *Input Parts* box, make sure that **Fluidized_Bed** is selected.
3. Set the **Scalar** function to **Volume Fraction > Volume Fraction of Particle**.
4. Set the **Extraction Mode** to **Single Value**.
5. Set the **Isovalue** to `0.15`.
6. In the *Display* group box, select **No Displayer**.



7. Click **Create** and then click **Close**.

   A new derived part, **iso**, is created.

However, the 85% concentration also occurs at the free surface at the top of the bed.

To remove the free surface at the top of the fluidized bed:

8. Right-click the **Derived Parts** node and select **New Part > Threshold**.

   The *Create Threshold* dialog appears in the *edit* window.
9. In the *Input Parts* box, select **iso**.
10. Set the **Scalar** to **Centroid > Laboratory > Y**.
11. Set the **Extraction Mode** to **All below**.
12. Set the **Threshold Value** to `0.4 m`.
13. In the *Display* group box, select **No Displayer**.
14. Click **Create** and then click **Close**.

To create the field function for the line integral:

15. Right-click **Tools > Field Functions** and select **New > Scalar**.
16. Rename the **User Field Function 1** node to `Unity`.
17. In the `Properties` window, set the **Function Name** property to `Unity`.
18. Enter a **Definition** of `1`.

Create a Line Integral report along the circumference of the bubble:

19. Right-click the **Reports** node and select **New Report > Line Integral**.
20. Rename the **Line Integral 1** node to `Circumference`.
21. Select the **Reports > Circumference** node and set the **Field Function** to **Unity**.
22. Set the **Parts** property to **Derived Parts > threshold**.

You can use the **Circumference** report to calculate the bubble diameter:

23. Right-click the **Reports** node and select **New Report > Expression**.
24. Rename the **Expression 1** node to `Bubble Diameter`.
25. Select the **Reports > Bubble Diameter** node and enter the definition:

    ```
    $CircumferenceReport /3.14
    ```

To create a monitor and plot of the bubble diameter at each time-step:

26. Right-click **Bubble Diameter** and set **Create Monitor and Plot from Report**.
27. Select the **Monitors > Bubble Diameter Monitor** node and set the **Trigger** property to **Time Step**.
28. Save the simulation.

Creating a Scalar Scene

Create a scalar scene to display the particle volume fraction in the tank.

To create a scalar scene:

1. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.
2. Rename the **Scalar Scene 1** node to `Bubble Shape`.
3. Click the **Scene/Plot** button above the object tree.
4. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set the **Function** property to **Volume Fraction** > **Volume Fraction of Particle**.
5. Select the **Scalar 1** node and set the **Contour Style** to **Smooth Filled**.
6. To return to the simulation object tree, click the **simulation** button.
7. Click the **Initialize Solution** button.

   The initial distribution of the phases appears in the *Graphics* window.



8. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   When the last iteration is complete, the following message is displayed in the *Output* window:

   ```
   Stopping criterion Maximum Physical Time satisfied.
   ```

2. When the simulation has finished running, save it.

Visualizing the Results

View the results of the simulation.

You can visualize the computed solution by viewing the scalar scene in the *Graphics* window.

To visualize the results of the simulation:

1. Click the *Bubble Shape* tab at the top of the *Graphics* window.

   A bubble has formed at the bottom of the fluidized bed.



Validation data is available for this case. Import this data and use it to validate the prediction of bubble diameter.

To compare the bubble diameter with experimental results:

2. Right-click the **Tools > Tables** and select **New Table > File Table**.
3. In the *Open* dialog navigate to the `doc/startutorialsdata/multiphaseFlow/data` subdirectory of your STAR-CCM+ installation, and select the file `kuipers_expt.csv`.
4. Click **Open**.
5. Right-click the **Plots > Bubble Diameter Monitor Plot > Data Series** node and select **Add Data**.
6. In the *Add Data Providers to Plot* dialog, expand the **Tables** node and select **kuipers_expt**.



7. Select the **kuipers_expt** node and set the **Legend Name** to `Experiment`, the **X Column** to `time` and the **Y Column** to `dia`.
8. Select the **Line Style** node and make sure that the **Style** property is set to **None**.
9. Select the **Symbol Style** node and do the following:
   a) Change the **Size** property to `12`.
   b) Change the **Color** property to **Blue**.
   c) Change the **Shape** property to **Filled Triangle**.



10. Select the **Plots** > **Bubble Diameter Monitor Plot** > **Axes** > **X Axis** node and set the **Maximum** to `0.2`.
11. Select the **Plots** > **Bubble Diameter Monitor Plot** > **Axes** > **X Axis** > **Title** node and set the **Title** to `Time (s)`.
12. Select the **Plots** > **Bubble Diameter Monitor Plot** > **Axes** > **Y Axis** > **Title** node and set the **Title** to `Diameter`.

    The experimental data now appears in the same plot as the predicted bubble diameter. A reports plot scene appears as shown below:



The results provide a reasonable agreement with the experiment.

13. Save the simulation.

Summary

This tutorial demonstrated how to model bubble formation in a fluidized bed in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up a simulation for granular flow using phase interaction models.
- Creating phases and selecting phase models.
- Defining phase interaction.
- Setting model properties.
- Creating a line integral report to monitor bubble diameter.
- Running the simulation.
- Visualizing the results.

Bibliography

This tutorial references the following papers:

Kuipers, J.A.M. 1990. "A two-fluid micro balance model of fluidized beds", PhD Thesis, Univ. of Twente, Twente, the Netherlands

Eulerian: Particle-Laden Flow in a Rotating Body

This tutorial demonstrates the separation of phases due to centrifugal forces in a rotating device using Eulerian Multiphase modeling in STAR-CCM+. Additionally, the tutorial also demonstrates solid pressure force model for modeling particle-particle interactions.

A simple two-dimensional geometry is used as shown below.



The initial mixture is assumed to be 90% water and 10% particles by volume. The particle density is set to be 1500 kg/m^3. The rotation (500 rpm) of the device has been simulated using both a rotating reference frame and rigid body motion. The simulation is run using an implicit unsteady solver for a physical time of 3 seconds.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Mesh
Visualizing the Boundaries
Visualizing the Interface

Selecting the Physics Models
Creating Phases and Selecting Phase Models
Defining the Phase Interactions
Setting Initial Conditions and Reference Values
Creating a Rotating Reference Frame
Setting Solver Parameters and Stopping Criteria
Setting Up a Monitor
Creating a Scalar Scene
Running the Simulation
Visualizing the Results
Creating a Rigid Body Motion
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Eulerian: Particle-Laden Flow in a Rotating Body tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A two-dimensional, trimmed cell mesh has been prepared for this tutorial.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.



4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/2d_rotating_geom.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `empWithMrf.sim`.

Visualizing the Mesh

Examine the imported mesh.

A two-dimensional cell mesh has been prepared for this tutorial. Display the imported mesh.

To visualize the mesh:

1. Right-click the **Scenes** node and then select **New Scene** > **Mesh**.
2. Use the mouse buttons to adjust the view so that it is similar to the view that is shown below.



Visualizing the Boundaries

Examine the mesh to view the boundaries.

Boundary definitions were also imported along with the volume mesh import.

To visualize the boundaries:

1. Open the **Regions** node and select the **mrf** node to check the location of the rotating body.



Notice that when you click the boundary, it is highlighted in the *Graphics* window.



2. Open the **Regions** > **mrf** node and select its sub nodes in turn to check their location.
3. Select the **station** node and its sub nodes to see their locations.

Alternatively, you can click various parts in the *Graphics* window to see which boundary they relate to.

4. Click anywhere outside the part in the *Graphics* window to clear the selected boundary.

Visualizing the Interface

View the interface between the wall boundaries.

All regions in a STAR-CCM+ simulation require an interface to be present between them that transfers the appropriate mass and energy quantities during the calculation. In this example, a so-called "In-place" interface has already been created between the wall boundaries enclosing the **mrf** and **station** regions.

To visualize the interface:

1. Open the **Interfaces** node and select the **In-Place 1** node.



A scene similar to the one shown below is displayed, showing the interface between the **mrf** and **station** regions.



Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

In this simulation, a turbulent flow is set up using the standard K-Epsilon turbulence model. Use the multiphase mixture option to allow you to define the two phases and the segregated multiphase option to define how the two phases interact.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Two Dimensional** (Pre-selected) **Gradients** (Pre-selected)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow** **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Turbulent**
Optional Models	**Gravity** **Solid Pressure**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

Create two new phases and specify them as water and particle. Select the models that are associated with each phase. The water and particle phases are modeled as liquid and particle, respectively. They are both defined as having a constant density.

To create the phases and select the phase models:

Create the liquid phase.

1. In the **Coolant** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Water**.
3. For the **Water** phase, select the following models:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Liquid**
Enabled Models	**Reynolds-Averaged Navier-Stokes** (Selected automatically)

4. Deactivate **Auto-select recommended models**.
5. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

6. Activate **Auto-select recommended models**.
7. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**

8. Click **Close**.

Create the particle phase.

9. Create a second phase and rename it to **Particle**.
10. For the **Particle** phase, select the following models:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Gas**
Enabled Models	**Constant Density** (Selected automatically) **Reynolds-Averaged Navier-Stokes** (Selected automatically)

11. Deactivate **Auto-select recommended models**.
12. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

13. Activate **Auto-select recommended models**.
14. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)

15. Click **Close**.

    The **Eulerian Phases** node appears as shown.



For the water phase, use the default material properties. For the particle phase, however, modify the density.

16. Edit the **Particle > Models > Particle > Al > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1500 kg/m^3`



17. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the water and particle phases. Specify the water phase as continuous, and the particle phase as dispersed.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to **Water-Particle**.

3. For the **Water-Particle** phase interaction, deactivate **Auto-select recommended models**, and then select the following models in order:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed**
Enabled Models	**Drag Force** (selected automatically) **Multiphase Material** (selected automatically) **Interaction Length Scale** (selected automatically) **Interaction Area Density** (selected automatically)
Optional Models	**Turbulent Dispersion Force**

4. Click **Close**.

Define the phase interaction properties.

5. Edit the **Water-Particle** > **Models** node and set the following properties:

Node	Property	Setting
Eulerian Continuous-Dispersed	**Continuous Phase**	**Water**
	Dispersed Phase	**Particle**
Drag Force	**Drag Force Coefficient Type**	**Linearized**
Drag Force > **Linearized Drag Coefficient**	**Method**	**Gidaspow**

6. Save the simulation.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the physics continuum.

Set the initial volume fraction for each of the phase. In this tutorial, you are assuming that the mixture is 90% water and 10% particle by volume. Note that the sum of volume fractions must be equal to one. By default,

the reference density is set to 1 kg/m^3. Change this value to correspond to the density of water.

To set the initial conditions and reference values:

Set the initial conditions and reference values for the **Physics 1** continuum.

1. Edit the **Physics 1 > Models > Eulerian Multiphase > Eulerian Phases** node and set the following properties:

Node	Property	Setting
Water > Initial Conditions		
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.9`
Particle > Initial Conditions		
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.1`



2. Edit the **Reference Values** node and set the following property:

Node	Property	Setting
Reference Density	**Value**	`997.561 kg/m^3`



3. Save the simulation.

Creating a Rotating Reference Frame

Create a moving reference frame for the rotating region.

Create a moving reference frame for the rotating region named **mrf**. Reference frames are created and managed within the **Reference Frames** node, and are then applied to regions as required.

To create a rotating reference frame:

1. Right-click the **Tools > Reference Frames** node and select **New > Rotating**.
2. Select the **Rotating** node and set the *Rotation Rate* to `500 rpm`.



Apply this motion to the **mrf** region.

3. Expand the **Regions** > **mrf** > **Physics Values** node and select **Motion Specification** and click to the right of the *Reference Frame* property and select **Rotating**.



In this section of the tutorial, you have left the *Motion* of the **mrf** region as **Stationary**.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

As you are solving an unsteady problem, it is necessary to specify the time-step size and the elapsed simulation time. This calculation runs for a maximum physical time of 3 seconds with a time-step size of 0.01 s.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.01 s`

Node	Property	Setting
Volume Fraction	**Under-Relaxation Factor**	`0.1`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`3 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`20`
Maximum Steps	**Enabled**	Deactivated

3. Save the simulation.

Setting Up a Monitor

Create a monitor to help assess convergence and stability of the solution.

By default STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the solution. In addition to this, it is useful to create other monitors to help in assessing convergence and stability of the solution. In this tutorial, you monitor the water and particle volume fractions at a point in the **station** region.

To set up a monitor:

1. Right-click the **Derived Parts** node and select **New Part** > **Probe** > **Point**

   The *Create Point* dialog appears in the *edit* window.
2. In the *Input Parts* box, select **station**.
3. Set the *Point* coordinates to `[0, 0.08, 0] m`
4. In the *Display* box, select **No Displayer**.



5. Click **Create**, then **Close**.

   A new derived part (called **point**) has been created.



This point is used to monitor the water and particle volume fractions in the stationary fluid region. Before you can create the monitor, first

create the report upon which these monitors are based. In this case, use a Maximum report.

6. Right-click the **Reports** node and select **New Report** > **Maximum**.



A new report is created named **Maximum 1**, which is used to monitor the water volume fraction.

Give the report a more suitable name.

7. Rename the **Maximum 1** node to `Water Volume Fraction`.
8. Having selected the **Reports** > **Water Volume Fraction** node, click to the right of *Scalar Field Function* property and select **Volume Fraction** > **Water**.



9. Click to the right of *Parts* property. Expand the **Derived Parts** node, select **Point**, and click **OK**.



Create the report for particle volume fraction by copying and pasting the report for the water volume fraction.

10. Right-click the **Reports** > **Water Volume Fraction** node and select **Copy**.
11. Right-click the **Reports** node and select **Paste**.

    A new report is created named **Copy of Water Volume Fraction**.

Give the report a more suitable name.

12. Rename the **Copy of Water Volume Fraction** node to `Particle Volume Fraction`.
13. Having selected the **Reports** > **Particle Volume Fraction** node, click to the right of *Scalar Field Function* property and select **Volume Fraction** > **Particle**.



Create monitors and a single plot from the reports.

14. Multi-select the **Water Volume Fraction** and **Particle Volume Fraction** report nodes. Right-click the selected nodes and select **Create Monitor and Plot from Report**.



15. In the *Create Plot From Reports* dialog, click **Single Plot**.



A new monitor and plot is created.



Creating a Scalar Scene

Create a scalar scene to view the results of the simulation.

A new scalar scene is created so that you can visualize the solution. This scene is used to display the particle volume fraction.

To create a scalar scene:

1. Right-click the **Scenes** node, and select **New Scene > Scalar**.
2. Click **Scene/Plot** above the object tree.
3. Select the **Displayers > Scalar 1 > Scalar Field** node and set the *Function* to **Volume Fraction > Particle**.



4. Select the **Scalar 1** node and set the *Contour Style* to **Smooth Filled**.
5. Select the **Color Bar** node and do the following:
   a) Set the *Width* to `0.2`.
   b) Set the *Height* to `0.9`.
   c) Set the *Orientation* to **Vertical**.
   d) Set the *Position* to `[0.7, 0.1]`.



6. Use the mouse buttons to adjust the view so that it is similar to the view that is shown below.



7. In the *Simulation* window, right-click the **Plots** > **Reports Plot** and select **Open**.
8. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

Visualize the computed solution by viewing the scalar scene.

You can visualize the computed solution by viewing the scalar scene in the *Graphics* window.

To visualize the results:

1. Click the *Scalar Scene 1* tab at the top of the *Graphics* window.

   A scene similar to the one shown below is displayed.



   Particles move away from the center because they are heavier than the water.

2. Click the *Reports Plot* tab at the top of the *Graphics* window.

   A reports plot scene similar to the one shown below is displayed.



The volume fractions of water and particle at the monitor point remain constant, indicating a stable and converged solution.

Other computed scalar quantities can be displayed by selecting the **Displayers** > **Scalar 1** > **Scalar Field**, and, setting the *Function* to the desired scalar.

In the next section, you convert the stationary, reference frame mode of rotation to a rigid body motion and thus demonstrating the required work through for this class of simulations.

Creating a Rigid Body Motion

Define a rigid body rotation motion.

In STAR-CCM+, motion is defined using motion objects, which are managed by the Motions manager. To define rigid body rotation motion, create a Rotation motion object and apply it to the region within which rotation takes place.

To create a rigid body motion:

1. Right-click the **Tools** > **Motions** node and select **New** > **Rotation**.
2. Select the **Rotation** node and do the following:
   a) Set the *Axis Direction* to [`0, 0, 1`].

b) Set the *Rotation Rate* to `500` rpm.



A reference frame node, called **ReferenceFrame for Rotation**, is added to the **Reference Frames** manager.



Apply this motion to the **mrf** region.

3. Expand the **Regions > mrf > Physics Values** node and select **Motion Specification** and set the *Motion* to **Rotation**.



Note: The *Reference Frame* property changes to **Lab Reference Frame** automatically.

Running the Simulation

Clear the results of the previous simulation and run the simulation again.

To clear the results and run the simulation again:

Clear the previous solution.

1. In the menu, select **Solution** > **Clear Solution**.
2. In the *Clear Solution* dialog, click **OK**.



The simulation is now ready to be run.

3. Click (**Run**) in the *Solution* toolbar.
4. When the simulation has finished running, save it.

Visualizing the Results

You can visualize the computed solution by viewing the scalar scene in the *Graphics* window.

To visualize the results:

1. Click the *Scalar Scene 1* tab at the top of the *Graphics* window.

   A scene similar to the one shown below is displayed.



2. Click the *Reports Plot* tab at the top of the *Graphics* window.

A reports plot scene similar to the one shown below is displayed.



Summary

This tutorial demonstrated the separation of phases due to centrifugal forces in a rotating device using Eulerian Multiphase modeling in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the physics models.
- Creating phases and selecting phase models.
- Defining phase interaction.
- Setting model properties.
- Setting reference frame mode of rotation.
- Setting rigid body rotation mode of motion.
- Running the simulation.
- Visualizing the results.

Eulerian: Degassing Boundary

This tutorial demonstrates how to model a Eulerian multiphase turbulent flow with a degassing boundary in STAR-CCM+. It simulates a simple aeration tank, where air enters the tank through two gas injectors at the bottom of the tank and escapes through the top surface.

The geometry is shown below.



The degassing surface (the top boundary) is modeled using a frictionless phase-permeable wall. This wall can share a region with a flow-split outlet or pressure outlet boundary.

STAR-CCM+ automatically calculates the mass flow rate through the phase-permeable boundary (the degassing rate). This rate can be plotted using the predefined Mass-Flow Rate field function.

This tutorial also covers how to select the Eulerian multiphase and multiphase turbulence models, specify the properties of each phase and define the phase interaction.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Selecting the Physics Models
Creating Phases and Selecting Phase Models
Defining the Phase Interactions

Deactivating Secondary Gradients
Setting Initial Conditions and Reference Values
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Eulerian: Degassing Boundary tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

A three-dimensional, trimmed cell mesh has been prepared for this analysis.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/aerationTank.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `aerationTank.sim`.

Create a mesh scene, and visualize the geometry:

6. Click (**Create/Open Scenes**) and select **Mesh**.
7. Use the mouse buttons to rotate and view the geometry.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Pre-selected) **Gradients** (Pre-selected)
Time	**Steady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow** **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Turbulent**
Optional Models	**Gravity** **Cell Quality Remediation**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

The two phases that are used in this case are air and water. Air bubbles are injected into a flowing stream of water. The phases are defined by creating new phases under the Eulerian Multiphase model.

To create the phases and select the phase models:

Create the liquid phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Water**.
3. For the **Water** phase, select the following models:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Liquid** **Reynolds-Averaged Navier-Stokes** (Selected automatically)

4. Deactivate **Auto-select recommended models**.
5. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

6. Activate **Auto-select recommended models**.
7. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**

8. Click **Close**.

Create the gas phase.

9. Create a second phase and rename it to **Air**.
10. For the **Air** phase, select the following models:

Group Box	Model
Material	**Gas** **Reynolds-Averaged Navier-Stokes** (Selected automatically)

11. Deactivate **Auto-select recommended models**.
12. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

13. Activate **Auto-select recommended models**.
14. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**

15. Click **Close**.

    The **Eulerian Phases** node appears as shown below.



16. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the water and air phases. The Wang Curve Fit model is suited for modeling dilute concentrations of air bubbles in water. The air bubbles being modeled have a diameter of 1 mm.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to **Water-Air**.

3. For the **Water-Air** phase interaction, select the following models in order:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed**
Enabled Models	**Drag Force** (selected automatically) **Multiphase Material** (selected automatically) **Interaction Length Scale** (selected automatically) **Interaction Area Density** (selected automatically)

4. Click **Close**.

   The **Water-Air > Models** node is shown below:



5. Edit the **Water-Air > Models** node and set the following properties:

Node	Property	Setting
Interaction Length Scale	**Interaction Length Scale**	`0.001 m`

Node	Property	Setting
Eulerian Continuous-Dispersed	Continuous Phase	Water
	Dispersed Phase	Air
Drag Force > Drag Coefficient	Method	Wang Curve Fit

The completed *Models* dialog is shown below.



6. Save the simulation.

Deactivating Secondary Gradients

Deactivate the secondary gradients option of the Multiphase Segregated Flow and Standard K-Epsilon models for each phase to help accelerate solution convergence.

In the following steps, you deactivate the secondary gradients option of the Multiphase Segregated Flow and Standard K-Epsilon models for each phase to help accelerate solution convergence.

To deactivate the secondary gradients:

Edit the Multiphase Segregated Flow model first.

1. Select the **Physics 1** > **Models** > **Multiphase Segregated Flow** node and set the *Secondary Gradients* to **Off**.

   The *Properties* window is shown below.



Edit the Multiphase Standard K-Epsilon model for each phase.

2. Select the **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** > **Water** > **Models** > **Standard K-Epsilon** node and set the *Secondary Gradients* to **Off**.
3. Repeat for the *Secondary Gradients* property of the **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** > **Air** > **Models** > **Standard K-Epsilon** node.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the **Physics 1** continuum.

To set the initial conditions and reference values:

1. Edit the **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and set the following properties:

Node	Property	Setting
Water > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[0.0, 0.0, 0.5] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`1.0`
Air > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[0.0, 0.0, 0.5] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.0`

Set gravity to act in the -y direction. It is common practice to use the density of the continuous phase as the reference density.

2. Edit the **Reference Values** node and set the following property:

Node	Property	Setting
Gravity	**Value**	`[0.0, -9.81, 0.0] m/s^2`
Reference Density	**Value**	`997.561 kg/m^3`



3. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

The boundary types have been defined before mesh generation.

You begin by setting the phase-permeable **Degassing Outlet** boundary so that only the air bubbles can escape as they rise to the top of the aeration tank. STAR-CCM+ ensures that the total gas inflow is equal to the total gas outflow and computes the internal gas volume fraction according to continuity.

To set the boundary conditions:

Allow gas to escape through the **Degassing Outlet** boundary.

1. Edit the **Regions** > **Aeration_Tank** > **Boundaries** > **Degassing Outlet** node and set the following property:

Node	Property	Setting
Phase Conditions > Air	**Type**	**Phase Permeable**



The top boundary of the aeration tank now acts as an outlet for the air phase. The water phase treats this boundary as a wall.

Make sure that this wall does not affect the flow of the water.

2. Edit the **Degassing Outlet** node and set the following property:

Node	Property	Setting
Phase Conditions > Water > Physics Conditions > Shear Stress Specification	**Method**	**Slip**

Define the **Gas Inlet** boundary. The velocity of each phase at the gas injector boundary is specified under the **Velocity Magnitude** node. For air bubbles of 1-mm diameter, the terminal rise velocity in water is about 0.25 m/s. You can use the terminal velocity at the inlets as bubbles in most bubbly flows reach their terminal velocity within a short distance from the injector.

Set the volume fraction of the gas injectors. This inlet is essentially a perforated plate where 10% of the area corresponds to the air nozzle holes.

3. Edit the **Gas Inlet > Phase Conditions** node and set the following properties:

Node	Property	Setting
Water > Physics Values		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`0.0 m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.9`
Air > Physics Values		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`0.25 m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.1`

Define the **Water Inlet** boundary. This inlet boundary introduces water into the aeration tank at a steady rate.

Set the velocity of the air phase to match the water velocity and set the air phase volume fraction to 0.0.

4. Edit the **Water Inlet > Phase Conditions** node and set the following properties:

Node	Property	Setting
Water > Physics Values		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`0.5 m/s`
Volume Fraction	**Method**	**Constant**

Node	Property	Setting
▮▮▮▮Constant	Value	1.0
Air > Physics Values		
▮▮Velocity Magnitude	Method	Constant
▮▮▮▮Constant	Value	0.5 m/s
▮▮Volume Fraction	Method	Constant
▮▮▮▮Constant	Value	0.0

5. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Multiphase Segregated Flow > Phase Coupled Velocity	Under-Relaxation Factor	0.3
Multiphase Segregated Flow > Pressure	Under-Relaxation Factor	0.3
Volume Fraction	Under-Relaxation Factor	0.2
K-Epsilon Turbulence	Under-Relaxation Factor	0.5
K-Epsilon Turbulent Viscosity	Under-Relaxation Factor	0.5

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Steps	Enabled	Activated

Node	Property	Setting
	Maximum Steps	700

3. Save the simulation.

Visualizing the Solution

Create two scalar scenes to display the results of the simulation.

To help you visualize the solution, create two scalar scenes. One displays the volume fraction of air through a central section plane, and the other displays the mass flow rate of the air phase on the degassing boundary. Create the section plane that cuts through the center of the geometry, including the two gas injectors. Then display the volume fraction of air onto this plane.

To visualize the solution:

1. Click (**Create/Open Scenes**) and select **Scalar**.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.
3. In the *Create Section* dialog that appears, select **Existing Displayer** > **Scalar 1** from the *Display* box.



4. Leave the remaining settings at their default values and click **Create** and then **Close**.

5. Click the **scene/plot** button to show the object tree for the scalar scene.
6. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set the *Function* to **Volume Fraction** > **Air**.



7. Select the **Displayers** > **Scalar 1**node and set the *Contour Style* to **Smooth Filled**.



8. Click (**Save-Restore-Select views**) and select **Views** > **+X** > **Up +Y**.

   This action switches the view to the side of the tank.
9. Select the **Scalar Scene 1** node at the top of the object tree and rename it to `Volume Fraction of Air`.

   The scene is shown below.



Create the second scalar scene that displays the mass flow rate of air through the degassing boundary.

10. Click (**Create/Open Scenes**) and select **Scalar**.
11. Right-click the **Displayers** > **Scalar 1** > **Parts** node and select **Edit**.
12. In the *Edit* dialog, select the **Regions** > **Aeration_Tank** > **Degassing Outlet** node, as shown below.



13. Click **OK**.
14. Select the **Scalar 1** > **Scalar Field** node and set the *Function* to **Mass Flow Rate** > **Air**.

    This action displays the mass flow of air across the degassing boundary. The mass per second through each boundary face is displayed, not the mass flow per unit area per second.
15. Select the **Displayers** > **Scalar 1** node and set the *Contour Style* to **Smooth Filled**.
16. Click (**Save-Restore-Select views**) and select **Views** > **+Y** > **Up - X**.

    This switches to the top view of the tank.
17. Select the **Scalar Scene 1** node at the top of the object tree and rename it to `Mass Flow Rate of Air`.
18. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 700 iterations are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

   ```
   Stopping criterion Maximum Steps satisfied.
   ```

2. When the simulation has finished running, save it.

Visualizing the Results

View the results of the simulation using the scalar scenes that you prepared earlier.

The solution can be visualized by displaying each of the scalar scenes you prepared earlier.

To visualize the results:

1. Select the *Volume Fraction of Air* tab at the top of the *Graphics* window.



The air bubbles are being carried downstream as they rise towards the degassing outlet, where they then escape through the phase-permeable wall. The water is unaffected by this boundary (due to the slip condition you applied) and is kept inside the tank.

2. Similarly, display the *Mass Flow Rate of Air* scene by selecting the corresponding tab.



This scene displays the mass flow rate of the escaping air bubbles through the phase-permeable wall. The faster moving bubbles are concentrated near the center of the bubble plume.

To make sure that no water is passing through the degassing outlet, you can do the following:

3. Right-click the color bar in the scene window and change the scalar field to display Mass Flow Rate of Water.

   No water flows through this boundary.

Summary

This tutorial demonstrated how to model a Eulerian multiphase turbulent flow with a degassing boundary in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the Eulerian Multiphase model.
- Creating phases and selecting phase models.
- Define phase interaction.

- Setting reference values and initial conditions.
- Setting up boundary conditions, including a phase-permeable wall boundary to allow the selected phase to exit through it.
- Running the simulation.
- Visualizing the results.

Eulerian: Wall Boiling

This tutorial validates the wall boiling model in STAR-CCM+ against experimental void and temperature data for high-pressure forced convective boiling.

The 1967 paper of Bartolomei provides 13 cases. The first case at 45 *bar* is of particular interest as it provides wall superheat data, and is the focus of this tutorial. Although the agreement between prediction and experiment is adequate rather than good, the same models can be used to obtain reasonable predictions for void in many of the further 29 cases spanning 30-150 *bar* given in the 1982 paper.

The geometry that is used is a Bartolomei boiling tube, where subcooled water flows upward through a 2 meter tall vertical pipe. The pressure at the outlet is 45 bar. Heat is applied uniformly at the pipe wall.

The basic geometry is a tube of 15.4-mm diameter and heated length of 2 m. To allow some flow development between the inlet and the heated section, a short unheated leader has been added. Similarly so that the outlet conditions in this developing flow do not influence the end of the heated section, a short trailer has also been added.

The conditions are defined as follows:

Length	L	2 m
Equivalent Diameter	D	0.0154 m
Heat Flux	q	570000 W/m^2
Mass Flux	G	900 kg/m^2
Inlet Temperature	T^{IN}	470.63 K
System Pressure	P^{SYS}	45.0 bar

Boiling starts when the wall temperature exceeds the water saturation temperature. This point is called the point of Inception of Boiling (IB) and small bubbles start to appear near the wall. When the degree of subcooling decreases, downstream of the point of IB, the volume fraction of the steam begins to grow. This point in the tube is called the point of Net Vapor Generation (NVG). Further downstream from the NVG point, the volume fraction of the steam continues to grow at an increasing rate and the bubbles start to detach from the wall and move into the liquid. As this detaching occurs, some of the bubbles condense, enhancing the heat transfer to the water. The mixture of steam and water exits at the top of the tube.

An axisymmetric, two-dimensional volume mesh of the Bartolomei boiling tube is provided, as shown below.



Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting Up User Field Functions
Selecting the Physics Models
Creating Phases and Selecting Phase Models
Defining the Phase Interactions
Defining the Phase Interaction Properties
Setting the Material Properties
Setting Reference Values
Setting Boundary Conditions
Setting Initial Conditions
Setting Solver Parameters and Stopping Criteria

Setting Up Monitoring Plots
Setting Up Analysis Plots
Visualizing the Solution
Running the Simulation
Final Mass Balance
Visualizing the Results
Comparison With Experiment
Turbulence Checks
Additional Information
Summary
Bibliography

Prerequisites

The instructions in the Eulerian: Wall Boiling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Input Files

The input files are in the `doc/startutorialsdata/multiphaseFlow/data` subdirectory of your STAR-CCM+ installation directory. It is recommended that you copy the following input files to your working directory before you start this tutorial:

- `wallboilingtutorialgrid.ccm`
- `wb_setupFieldFunctions.java`
- `wb_createAxialPlots.java`
- `wb_deleteAxialPlots.java`
- `wb_setupAnalysisPlots.java`
- `wb_setupMonitorPlots.java`
- `Bart1967_Figure3_AxisTemperature.csv`
- `Bart1967_Figure3_CalorimetricTemperature.csv`

- `Bart1967_Figure3_TrueVoid.csv`
- `Bart1967_Figure3_WallTemperature.csv`

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

The basic geometry is a tube of 15.4-mm diameter and heated length of 2 m. To allow some flow development between the inlet and the heated section, a short unheated leader has been added. Similarly so that the outlet conditions in this developing flow do not influence the end of the heated section, a short trailer has also been added. These additions are helpful and reasonable computational practices when you consider the drawings of the original rig.

The overall grid resolution is 115 cells in the axial direction and 20 cells in the radial direction between the axis and the tube wall. Both the grid and the grid cells have a high aspect ratio, and it is helpful later to set up a simple transform to view the mesh or the results.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/wallBoilingTutorialGrid.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `wallBoilingTutorial.sim`.

Setting Up User Field Functions

Create the field functions that are required for this tutorial.

This tutorial makes extensive use of field functions to define case parameters, to set up boundary conditions, and to analyze and present results.

For convenience, these field functions can be set up by running a macro before you set up any models.

To create the user field functions:

1. Select **File > Macro > Play Macro**.
2. In the *Open* dialog, navigate to your working directory and select `wb_setupFieldFunctions.java`.
3. Click the **Open** button to run the macro.

The field functions are tabulated below:

Case Parameter Field Functions

EquivalentDiameter	
Type	**Scalar**
Definition	`0.0154`
Ignore Boundary Values	**Deactivated**

HeatedLength	
Type	**Scalar**
Definition	`2.0`
Ignore Boundary Values	**Deactivated**

SystemPressure	
Type	**Scalar**
Definition	`4500000.0`
Ignore Boundary Values	**Deactivated**

SaturationTemperature	
Type	**Scalar**
Definition	`530.59`
Ignore Boundary Values	**Deactivated**

InletDensity	
Type	**Scalar**
Definition	`832.06`
Ignore Boundary Values	**Deactivated**

InletMassFlux	
Type	**Scalar**
Definition	`900.0`

InletMassFlux	
Ignore Boundary Values	**Deactivated**

InletTemperature	
Type	**Scalar**
Definition	`470.63`
Ignore Boundary Values	**Deactivated**

InletVelocity	
Type	**Vector**
Definition	`[ $InletMassFlux/ $InletDensity, 0, 0 ]`
Ignore Boundary Values	**Deactivated**

WallHeatFlux	
Type	**Scalar**
Definition	`570000.0`
Ignore Boundary Values	**Deactivated**

Model Field Functions

IterationsToFullPower	
Type	**Scalar**
Definition	`200.0`
Ignore Boundary Values	**Deactivated**

WallHeatFluxDistribution	
Type	**Scalar**
Definition	`($$Centroid[0] > 0 && $ $Centroid[0] < $HeatedLength) ? $WallHeatFlux : 0.0`
Ignore Boundary Values	**Deactivated**

WallHeatFluxIterationRamp	
Type	**Scalar**

WallHeatFluxIterationRamp	
Definition	`( $Iteration < $IterationsToFullPower ? ($WallHeatFluxDistribution * $Iteration )/ $IterationsToFullPower : $WallHeatFluxDistribution )`
Ignore Boundary Values	**Deactivated**

dMax	
Type	**Scalar**
Definition	`0.012`
Ignore Boundary Values	**Deactivated**

dMin	
Type	**Scalar**
Definition	`1.0E-6`
Ignore Boundary Values	**Deactivated**

ClippedSauterMeanDiameter	
Type	**Scalar**
Definition	`( $SauterMeanDiameterVapor < $dMin ? $dMin : ( $SauterMeanDiameterVapor > $dMax ? $dMax : $SauterMeanDiameterVapor ) )`
Ignore Boundary Values	**Deactivated**

Analysis Field Functions

InletThermodynamicQuality	
Type	**Scalar**
Definition	`$SpecificHeatLiquid * ($InletTemperature - $SaturationTemperature)/ ${LatentHeatLiquid Vapor Interface}`
Ignore Boundary Values	**Deactivated**

ThermodynamicQuality	
Type	**Scalar**
Definition	`$InletThermodynamicQuality+ 4 * $WallHeatFlux * $ $Centroid[0]/ ( $InletMassFlux * $ {LatentHeatLiquid Vapor Interface} * $EquivalentDiameter )`
Ignore Boundary Values	**Deactivated**

LiquidSubcooling	
Type	**Scalar**
Definition	`${TemperatureLiquid Vapor Interface} - $ {TemperatureLiquid}`
Ignore Boundary Values	**Activated**

WallSuperHeat	
Type	**Scalar**
Definition	`${TemperatureLiquid} - $ {TemperatureLiquid Vapor Interface}`
Ignore Boundary Values	**Deactivated**

QuenchingHeatFlux	
Type	**Scalar**
Definition	`(1-$ {WallDryoutAreaFractionLiquid Vapor Interface}) * $ {BubbleInducedQuenchingHeatTransferCoefficientLiquid Vapor Interface} * $ {BubbleInducedQuenchingTemperatureFactorLiquid Vapor Interface} * ( $WallSuperHeat + $LiquidSubcooling )`
Ignore Boundary Values	**Deactivated**

WallEvaporationRate	
Type	**Scalar**
Definition	`-${WallInterphaseMassFluxLiquid Vapor Interface}`
Ignore Boundary Values	**Deactivated**

BulkCondensationRate	
Type	**Scalar**
Definition	`( ${InterphaseMassTransferRateLiquid Vapor Interface} > 0 ? ${InterphaseMassTransferRateLiquid Vapor Interface} : 0 )`
Ignore Boundary Values	**Deactivated**

BulkEvaporationRate	
Type	**Scalar**
Definition	`( ${InterphaseMassTransferRateLiquid Vapor Interface} < 0 ? -${InterphaseMassTransferRateLiquid Vapor Interface} : 0 )`
Ignore Boundary Values	**Deactivated**

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

The multiphase turbulence models are available for selection once the multiphase flow and material models have been selected.

To select the physics models:

1. For the physics continuum, **Continua > Physics 1**, in the *Enabled Models* group box, deactivate **Two-Dimensional**.
2. Select the following models in order:

Group Box	Model
Space	**Axisymmetric**
Time	**Steady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow** **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Turbulent**
Optional Models	**Phase Coupled Fluid Energy** **Gravity**

3. Click **Close**.
4. To review the models, open the **Physics 1** > **Models** node.



5. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

The two phases, Liquid and Vapor, that are flowing through the section of vertical pipe represent water and steam. These phases are defined by creating phases under the Eulerian physics model, but material properties are set later once all other models have selected.

The main modeling choices that are made at this stage are:

- A K-Epsilon model with a two-layer wall treatment is selected for the continuous phase so that it can be used under a wide range of flow conditions.
- The Turbulence Response model is an appropriate model for the velocity fluctuations that are experienced in the dispersed phase for this bubbly flow.
- The S-gamma model is activated for the disperse phase so that it can track bubble size starting from bubble creation at the boiling wall, and including the effects of breakup, coalescence, and (most importantly) condensation.

To create the phases and select the phase models:

Create the liquid phase.

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Liquid**.
3. For the **Liquid** phase, select the following models:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Liquid** **Reynolds-Averaged Navier-Stokes** (Selected automatically)

4. Deactivate **Auto-select recommended models**.
5. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Realizable K-Epsilon Two-Layer**

6. Activate **Auto-select recommended models**.
7. Select the following models in order:

Group Box	Model
Enabled Models	**Two-Layer All y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**
Energy	**Segregated Fluid Temperature**

8. Click **Close**.

Create the gas phase.

9. Create a second phase and rename it to **Vapor**.
10. For the **Vapor** phase, select the following models:

Group Box	Model
Material	**Gas** **Reynolds-Averaged Navier-Stokes** (Selected automatically)

11. Deactivate **Auto-select recommended models**.
12. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**Turbulence Response**
Turbulence Response Models	**Issa Turbulence Response Model**

13. Activate **Auto-select recommended models**.
14. Select the following models in order:

Group Box	Model
Equation of State	**Constant Density**
Energy	**Segregated Fluid Temperature**
Optional Models	**S-Gamma**

15. Click **Close**.
16. Rename the **Vapor** > **Models** > **Gas** > **Air** node to `Steam`.

    The **Eulerian Phases** node appears as shown.



17. Edit the **Eulerian Phases** > **Vapor** > **Models** node and set the following properties:

Node	Property	Setting
Issa Turbulence Response Model	**Background Phase**	**Liquid**
S-Gamma	**Number of S-Gamma Equations**	**One Equation**

18. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the liquid and vapor phases.

The main modeling choice that is made here is the use of the combination of Hibiki-Ishii Nucleation Site Number Density and Kocamustafaogullari Bubble Departure Diameter which is applicable for water boiling under a wide range of pressures. The default models (Lemmert-Chawla and Tolubinsky) are much simpler, but can need testing and recalibration for use at different pressures.

To define the phase interactions:

1. In the **Physics 1** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to **Liquid Vapor Interface**.
3. For the **Liquid Vapor Interface** phase interaction, , deactivate **Auto-select recommended models**, and then select the following models in order:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed Drag Force** (selected automatically) **Multiphase Material** (selected automatically)
Interaction Length Scale	**Interaction Length Scale** (selected automatically)
Interaction Area Density	**Interaction Area Density** (selected automatically)
Optional Models	**Turbulent Dispersion Force** **Interphase Mass Transfer**
Interphase Mass Transfer Rate	**Boiling Mass Transfer Rate**
Optional Models	**Wall Boiling** **Bubble Induced Quenching Temperature** (selected automatically)
Nucleation Site Number Density	**Hibiki Ishii Nucleation Site Number Density**
Bubble Departure Diameter	**Kocamustafaogullari Bubble Departure Diameter**

Group Box	Model
Bubble Influence Wall Area Fraction	**Kurul Podowski Bubble Influence Wall Area Fraction**
Bubble Induced Quenching Heat Transfer Coefficient	**Del Valle Kenning Bubble Induced Quenching Heat Transfer Coefficient**
Bubble Departure Frequency	**Cole Bubble Departure Frequency**
Wall Dryout Area Fraction	**Basic Wall Dryout Area Fraction**

The length scale is predicted using the S-Gamma Model.

The **Symmetric Interaction Area Density** setting is because a wide range of volume fractions is expected.

4. Click **Close**.

The **Liquid Vapor Interface** > **Models** node is shown below:



Defining the Phase Interaction Properties

Define the interactions between the phases.

The phase interaction properties are now defined, starting with the definition of the continuous and dispersed phases.

To define the phase interaction properties:

1. Edit the **Liquid Vapor Interface** > **Models** node and set the following properties:

Node	Property	Setting
Eulerian Continuous-Dispersed	**Continuous Phase**	**Liquid**
	Dispersed Phase	**Vapor**

Once the Continuous and Dispersed phases have been defined, you select additional models:

2. For the **Liquid Vapor Interface** phase interaction, select the following models:

Group Box	Model
Optional Models	**S-Gamma Breakup**
	S-Gamma Coalescence

Define the remainder of the phase interaction properties.

3. Edit the **Liquid Vapor Interface** > **Models** node and set the following properties:

Node	Property	Setting
Boiling Mass Transfer Rate > Continuous Phase Nusselt Number	**Method**	**RanzMarshall**
Disperse Phase Nusselt Number > Constant	**Value**	`26.0`
Drag Force > Drag Coefficient	**Method**	**Tomiyama**
Drag Correction > Volume Fraction Exponent	**Volume Fraction Exponent (continuous)**	`0.0`

Node	Property	Setting
Turbulent Dispersion Force	**Use Logarithmic Form**	Activated
Interaction Length Scale > Interaction Length Scale	**Method**	**Field Function**
Field Function	**Scalar Function**	**ClippedSauterMean Diameter**



Selecting the **RanzMarshall** option makes STAR-CCM+ calculate the **Continuous Phase Nusselt Number** using the Ranz-Marshall correlation. The Ranz-Marshall correlation, in turn, is used to calculate the heat transfer coefficient for the heat transfer from the bulk water to the bubble interface.

The **Disperse Phase Nusselt Number** is used when calculating the heat transfer coefficient for the heat transfer from the bubble interface to the bulk steam. The exact value is not critical because the vapor is going to be at, or close to, saturation temperature. A value of 2 represents a slow response with little internal circulation. A higher value such as 26 can represent either a fast superficial response or the effect of internal mixing within the vapor.

The **Tomiyama** model is used for modeling bubbly flows. Drag reduction is not expected for small spherical bubbles.

The **Turbulent Dispersion Force** setting is important for wall boiling calculations in the presence of condensation because it treats the differencing of rapidly changing volume fraction gradients more accurately. Without this setting, the simulation can get unusual velocities at the edge of the spreading bubble cloud as well as unnecessary convergence difficulties.

The **Interaction Length Scale** is based on the bubble size predicted by the S-Gamma model. The **ClippedSauterMeanDiameter** field function is one of the user-defined field functions that was created earlier. This setting allows the Sauter Mean Diameter that is predicted with the S-Gamma model to be used as the interaction length scale for drag and interface transfers. It also filters the calculations through a maximum and a minimum diameter so that early unconverged S-Gamma solutions do not upset convergence.

4. Save the simulation.

Setting the Material Properties

Define the phases and set the appropriate material properties.

In this high-pressure calculation where wall dryout is not expected, the vapor remains at saturation temperature. Therefore it is valid to assume constant vapor properties and constant interface properties.

The liquid on the other hand can experience a significant change in density and in specific heat while it is heated from a subcooling of 60 K towards saturation temperature. Nevertheless a useful analysis can be made using constant properties for the liquid by respecting these two conditions:

- In order to obtain the correct heat balance, while also using the experimental inlet temperature, the following definition of liquid specific heat was chosen:

$$C_p = \frac{h_L^{SAT} - h_L^{IN}}{T^{SAT} - T^{IN}}$$

where *SAT* indicates saturation, *IN* indicates inlet, and *L* indicates liquid.

- The liquid density, viscosity, and conductivity are chosen at a temperature midway between inlet and saturation. The inlet velocity must be chosen such that inlet velocity times density matches the experimental mass flux.

In this tutorial, the process of setting the properties has been deferred until after all the models have been selected. Some properties appear only when demanded by particular models.

Specify the material properties for the liquid phase.

1. Edit the **Eulerian Multiphase** > **Eulerian Phases** > **Liquid** > **Models** > **Liquid** > **H20** > **Material Properties** node and set the following properties:

Node	Property	Setting
Density > **Constant**	**Value**	`$InletDensity`
Dynamic Viscosity > **Constant**	**Value**	`0.0001177 Pa-s`
Heat of Formation > **Constant**	**Value**	`1122143 J/kg`
Specific Heat > **Constant**	**Value**	`4648.3 J/kg-K`
Standard State Temperature > **Constant**	**Value**	`$SaturationTemperature`
Thermal Conductivity > **Constant**	**Value**	`0.6433 W/m-K`

Set the initial temperature of the liquid phase:

2. Select the **Liquid** > **Initial Conditions** > **Static Temperature** > **Constant** node and set the **Value** to `461.0 K`.

   The settings for the liquid phase are now complete.

Specify the material properties for the vapor phase.

In general the phase properties that are labeled as Heat of Formation and Standard State Temperature are used internally by STAR-CCM+ as a general Reference Enthalpy at a specified Reference Temperature. For

this boiling application, these properties are chosen to be Saturation Enthalpy and Saturation Temperature respectively from standard steam tables.

Latent Heat is computed automatically from the phase enthalpy definitions.

3. Edit the **Vapor > Models > Gas > Steam > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`22.7 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`1.78E-5 Pa-s`
Heat of Formation > Constant	**Value**	`2797997 J/kg`
Molecular Weight > Constant	**Value**	`18.0 kg/kmol`
Specific Heat > Constant	**Value**	`4227.9 J/kg-K`
Standard State Temperature > Constant	**Value**	`$SaturationTemperature`
Thermal Conductivity > Constant	**Value**	`0.0535 W/m-K`

4. Select the **Multiphase Interaction > Phase Interactions > Liquid Vapor Interface > Models > Boiling Mass Transfer Rate > Interface Temperature > Constant** node.
5. In the *Properties* window, set the **Value** to `$SaturationTemperature`.
6. Select the **Liquid Vapor Interface > Models > Multiphase Material > Multiphase Material > Material Properties > Surface Tension > Constant** node.
7. In the *Properties* window, set the **Value** to `0.024292 N/m`.

Setting Reference Values

Set the reference values for the physics continuum.

The reference gravity, density, and pressure for the simulation are set up. The main modeling choice here is to take the reference density for the buoyancy term from the liquid phase density. This choice allows the simulation to move smoothly from a zero initial field for pressure and vapor volume fraction, rather than requiring a hydrostatic initial pressure field.

To set the reference values:

Set the reference values for the **Physics 1** continuum.

1. Edit the **Reference Values** node and set the following properties:

Node	Property	Setting
Gravity	**Value**	`[-9.81, 0.0, 0.0] m/s^2`
Reference Density	**Value**	`$InletDensity`
Reference Pressure	**Value**	`$SystemPressure`



2. Save the simulation.

Setting Boundary Conditions

Define the boundary types and specify the appropriate property values.

Specify conditions on the wall and inlet boundaries.

To set the boundary conditions:

1. Edit the **Regions** > **Boiling_Tube** > **Boundaries** > **Wall** node and set the following properties:

Node	Property	Setting
Physics Conditions > Thermal Specification	**Condition**	**Heat Flux**
Physics Values > Heat Flux	**Method**	**Field Function**
Field Function	**Scalar Function**	**WallHeatFluxIterationRamp**

Set the phase conditions for the liquid and vapor phases at the inlet. The turbulence intensity and turbulent viscosity ratio are modified so that the settings match the Initial Conditions.

2. Edit the **In** > **Phase Conditions** node and set the following properties:

Node	Property	Setting
Liquid > Physics Values		
Static Temperature	**Method**	**Constant**
Constant	**Value**	`$InletTemperature`
Turbulence Intensity	**Method**	**Constant**
Constant	**Value**	`0.01`
Turbulent Viscosity Ratio	**Method**	**Constant**
Constant	**Value**	`200`
Velocity Magnitude	**Method**	**Field Function**
Field Function	**Scalar Function**	**InletVelocity > Laboratory > i**

Node	Property	Setting
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`1.0`

Node	Property	Setting
Vapor > Physics Values		
Sauter Mean Diameter	**Method**	**Constant**
Constant	**Value**	`0.001`
Static Temperature	**Method**	**Constant**
Constant	**Value**	`$SaturationTemperature`
Velocity Magnitude	**Method**	**Field Function**
Field Function	**Scalar Function**	**InletVelocity > Laboratory > i**
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.0`

Define the outlet boundary. The static temperature for the liquid and vapor phases at the outlet are set. These values are the same as for the inlet.

3. Edit the **Out > Phase Conditions** node and set the following properties:

Node	Property	Setting
Liquid > Physics Values		
Static Temperature	**Method**	**Constant**
Constant	**Value**	`$InletTemperature`
Turbulence Intensity	**Method**	**Constant**

Node	Property	Setting
▮▮▮▮Constant	Value	0.01
▮▮Turbulent Viscosity Ratio	Method	Constant
▮▮▮▮Constant	Value	10
▮▮Volume Fraction	Method	Constant
▮▮▮▮Constant	Value	1.0

Node	Property	Setting
Vapor > Physics Values		
▮▮Sauter Mean Diameter	Method	Constant
▮▮▮▮Constant	Value	0.001
▮▮Static Temperature	Method	Constant
▮▮▮▮Constant	Value	$SaturationTemperature
▮▮Volume Fraction	Method	Constant
▮▮▮▮Constant	Value	0.0



4. Save the simulation.

Setting Initial Conditions

Set the global initial conditions and the initial conditions for each phase.

To set the initial conditions:

1. Edit the **Physics 1 > Models > Eulerian Multiphase > Eulerian Phases** node and set the following properties for the liquid phase:

Node	Property	Setting
Liquid > Initial Conditions		
Static Temperature	**Method**	**Constant**
Constant	**Value**	`$InletTemperature`

Node	Property	Setting
▮▮Turbulence Intensity	Method	Constant
▮▮▮▮Constant	Value	`0.01`
▮▮Turbulence Specification	Method	Intensity + Viscosity Ratio
▮▮Turbulence Velocity Scale	Method	Constant
▮▮▮▮Constant	Value	`1.0 m/s`
▮▮Turbulent Viscosity Ratio	Method	Constant
▮▮▮▮Constant	Value	`200`
▮▮Velocity	Method	Field Function
▮▮▮▮Field Function	Vector Function	InletVelocity
▮▮Volume Fraction	Method	Constant
Constant	Value	`1.0`

2. Edit the same node and set the following properties for the vapor phase:

Node	Property	Setting
Vapor > Initial Conditions		
▮▮Sauter Mean Diameter	Method	Constant
▮▮▮▮Constant	Value	`0.01 m`
▮▮Size Distribution Specification	Method	Sauter Mean Diameter
▮▮Static Temperature	Method	Constant
▮▮▮▮Constant	Value	`$SaturationTemperature`
▮▮Velocity	Method	Field Function

Node	Property	Setting
▮▮▮▮Field Function	Vector Function	InletVelocity
▮▮Volume Fraction	Method	Constant
▮▮▮▮Constant	Value	0.0

The phase equations cannot be solved if volume fraction is identically zero. Therefore, for boundary conditions and for initial conditions, the code uses 10E-12 for vapor volume fraction rather than 0.0.



3. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

The under-relaxation factors for the solvers are decreased to help the solution converge. These relaxation factors have not been tuned for best performance in this one particular case. Rather, they have been chosen to also allow converged solutions to be obtained for many of the test conditions reported in Bartolemei (1982).

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Multiphase Segregated Flow > Phase Coupled Velocity	**Under-Relaxation Factor**	`0.5`
Multiphase Segregated Flow > Pressure	**Under-Relaxation Factor**	`0.1`
Volume Fraction	**Under-Relaxation Factor**	`0.1`
S-gamma	**Under-Relaxation Factor**	`0.1`
Segregated Energy	**Fluid Under-Relaxation Factor**	`0.5`
K-Epsilon Turbulence	**Under-Relaxation Factor**	`0.5`



The Sauter Mean Diameter is predicted from the ratio of S2 and volume fraction. Therefore the **Volume Fraction** and **S-gamma** solvers are given the same relaxation factors to ensure that these two fields are progressed together.

In addition to solver relaxation factors, there are further relaxation factors that are associated with some specific model terms.

2. Edit the **Continua** > **Physics 1** > **Models** > **Multiphase Interaction** > **Phase Interactions** > **Liquid Vapor Interface** > **Models** node and set the following properties:

Node	Property	Setting
Basic Wall Dryout Area Fraction	**Bubbly Layer Relaxation Factor**	`0.5`
Boiling Mass Transfer Rate	**Relaxation Factor**	`0.1`

Node	Property	Setting
Wall Boiling	**Relaxation Factor**	`0.1`

3. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Steps	**Enabled**	Activated
	Maximum Steps	`2500`

4. Save the simulation.

Setting Up Monitoring Plots

Create plots to monitor the progress of the simulation.

Since the simulation is starting with no boiling and with little vapor, the final residuals can be many orders of magnitude different from the initial levels.

To help determine when the simulation has converged, you can also monitor the progress of the following:

- maximum velocity
- maximum and minimum temperature
- maximum volume fraction
- mass balance

Any of these plots can be set up individually by applying the **Create Monitor and Plot from Report** action to selected groups of various reports operating on relevant fields or field functions.

For convenience, the complete set of monitoring plots for this tutorial can be set up by running a macro.

To set up the monitoring plots:

1. Select **File > Macro > Play Macro**.
2. In the *Open* dialog, navigate to your working directory and select `wb_setupMonitorPlots.java`.
3. Click **Open** to run the macro.

In this tutorial, the maximum and minimum plots are expected to confirm that the velocities and temperatures remain physically plausible throughout the convergence history. This result is due to the strategy of ramping up the power gradually as the flow field is established.

In more difficult or unfamiliar cases, maximum and minimum plots can be helpful to identify the first signs of when a simulation is starting to get out of hand and can need adjustments to relaxation factors.

The Mass Balance plot tracks three quantities, as follows:

- LiquidMassFlow, which computes the mass outflow of liquid less the mass inflow of liquid.
- VaporMassFlow, which computes the mass outflow of vapor less the mass inflow of vapor.
- NetEvaporationRateReport, which is constructed from the expression $WallEvaporationRateReport + $BulkEvaporationRateReport - $BulkCondensationRateReport.

  where:

  - BulkCondensationRateReport is the volume integral of the positive part of the interior IPMT rate from the vapor phase to the liquid phase.
  - BulkEvaporationRateReport is the volume integral of the negative part of the interior IPMT rate from the vapor phase to the liquid phase.
  - WallEvaporationRateReport is the surface integral of minus the wall IPMT rate from the vapor phase to the liquid phase.

The IPMT (Interphase Mass Transfer) rate is defined to be positive where mass is transferred from the second (dispersed) phase to the first (continuous) phase that is specified under the Eulerian Continuous-Dispersed phase interaction. For this tutorial, this definition means that a positive IPMT rate indicates condensation, and a negative value indicates evaporation.

On convergence, the mass balance is expected to be obeyed for each phase, at least to some finite precision:

- LiquidMassFlow + NetEvaporationRateReport = 0
- VaporMassFlow - NetEvaporationRateReport = 0

Setting Up Analysis Plots

Create plots to analyse the results of the simulation.

It is also instructive to observe quantities of interest as convergence is approached. You set up plots of Axis Subcooling and Wall Superheat which can be compared directly with data given by Bartolemei (1967). An important feature of the model is the automatic partitioning of the

applied wall heat flux into three components (liquid convection, evaporation, and vapor convection). This partitioning is presented as a Wall Heat Flux plot. A number of other submodel quantities, such as bubble size, are of interest to researchers in this field.

For convenience, the analysis plots for this tutorial can be set up by running a macro.

To set up the analysis plots:

1. Select **File** > **Macro** > **Play Macro**.
2. In the *Open* dialog, navigate to your working directory and select `wb_setupAnalysisPlots.java`.
3. Click **Open** to run the macro.

Many plots present results on the wall or on the axis against height. However, because of uncertainty in the exact location and conditions at the start of the heated section, a comparison with Bartolemei (1967) or Bartolemei (1982) is best done by using thermodynamic quality as a non-dimensional scale. Thermodynamic quality is easily related to height from the heat balance.

The definition of thermodynamic quality is:

$$x_{th} = \frac{h - h_L^{SAT}}{h_{LV}}$$

where h is the mixture enthalpy as calculated by this heat balance:

$$h = h_L^{IN} + \frac{4qz}{GD}$$

For convenience, the thermodynamic quality has been set up as a user field function.

The components of heat flux between the wall and the liquid, between the wall and the vapor, and between the wall and the evaporating interface, are available as standard field functions. These field functions are suffixed with the relevant phase or interaction name:

- Boundary Heat Flux of Liquid
- Boundary Heat Flux of Vapor
- Boundary Heat Flux of Liquid Vapor Interface

Boundary heat fluxes follow the sign convention in that they are positive when directed outwards. This sign convention is corrected for the presentation of the Wall Heat Flux plot.

Wall superheat is defined as the wall temperature excess over the saturation temperature. It is defined here as a user field function. The

wall temperature itself is obtained by evaluating the temperature of either phase at the wall.

Liquid subcooling is defined as the saturation temperature less the liquid temperature. To ensure that this value can also be plotted at the wall, the Ignore Boundary Values option is activated when the field function is created. This setting ensures that the liquid temperature in the cell next to the wall is used rather than the wall temperature.

Visualizing the Solution

Create a scalar scene to display the results of the simulation.

A scalar scene is created to display the volume fraction of vapor. This variable can be monitored while the solver is running.

To visualize the solution:

1. Create a scalar scene.

The view of the pipe in the scene is scaled using a simple transform so that the entire length of the pipe can be viewed in a smaller area. The transform also rotates the geometry so that it sits vertically.

2. Right-click the **Tools > Transforms** node and select **New Graphics Transform > Simple Transform**.



A new node, **Simple Transform 1**, is added to the **Transforms** node.

3. Rename the **Simple Transform 1** node to `Scene Transform`.
4. Select the **Scene Transform** node and do the following:
   a) Set the *Scale* to `[70.0, 1.0, 1.0]`.
   b) Set the *Rotation Angle* to `90 deg`.

   The completed *Properties* window is shown below.



5. Click **scene/plot**.
6. Select the **Scalar Scene 1 > Displayers > Scalar 1** node and do the following:
   a) Set the *Contour Style* to **Smooth Filled**.
   b) Set the *Transform* to **Scene Transform**.

Set the field function for the scalar scene:

7. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node and set the *Function* to **Volume Fraction > Volume Fraction of Vapor**.

The color bar is edited so that it is positioned beside the vertical pipe.

8. Select the **Displayers > Scalar 1 > Color Bar** node and do the following:
   a) Set the *Orientation* to **Vertical**.
   b) Set the *Height* to `0.8`.
   c) Set the *Width* to `0.15`.
   d) Set the *Position* to `[0.6, 0.1]`.

   The completed *Properties* window is shown below.



The **Scene Transform** is also applied to the **Outline 1** node.

9. Select the **Scalar Scene 1 > Displayers > Outline 1** node and set the *Transform* to **Scene Transform**.

   The scalar scene appears as shown below.



10. Click **simulation** to return to the STAR-CCM+ simulation object tree.
11. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 2500 iterations are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Steps satisfied.
```

As well as observing the progress of residuals during convergence, an interesting feature to follow is the development of the heat flux partitioning.

2. Expand the **Plots** node and then open the *Wall SuperHeat* and *Wall Heat Flux* plots.
3. In the *Graphics* window, click the tab of one of these plots and drag it to the left or right margin of the window.

   The two plots are then displayed side by side, as in the example below.



The wall superheat is initially negative (below boiling point), but rises as the wall power is ramped up, and becomes positive just before full power is achieved at 200 iterations. During this period, the whole of the heat flux is directed to the liquid only. Once boiling point is achieved, an evaporation heat flux takes up a share of the applied wall heat flux. As the calculation continues, the liquid

convection in the boiling region is replaced with a quenching heat flux due to the agitation or mixing of the liquid by departing bubbles. Both evaporation (latent heat) and quenching (agitation) are major enhancers of heat transfer during nucleate boiling.

At the inception of boiling, single phase heat transfer is the only contributor to the wall heat flux. As the wall superheat rises, more nucleations sites become active, and both the evaporation and quenching contributions rise accordingly. However, the quenching contribution goes through a maximum, because it is also proportional to liquid subcooling which is falling as the liquid temperature rises.

The overall balance between quenching and evaporation fluxes is a function of operating pressure, since quenching flux varies as departure size squared and evaporation flux varies as departure size cubed. The bubble departure size is smaller at higher pressures.

The wall superheat profile is established early thanks to the sensitivity of nucleation site density to a large power of wall superheat. Net vapor production appears a little later, at about 400 iterations, once the condensation capacity of the subcooled water exceeds the wall evaporation rate. At the same time, the flow begins to accelerate, as can be seen from the maximum velocity monitor plot.

True convergence of mass balance and of residuals can be seen to start after 600 iterations. The maximum/minimum plots show small changes after iteration 2000. The calculation is stopped at 2500 iterations to get a good final mass balance.

4. When the simulation has finished running, save it.

Final Mass Balance

The final mass balance is summarized.

Liquid Inflow	2.6681E-2
Wall Evaporation	1.9380E-3
Bulk Evaporation	3.6564E-4
Total Evaporation	2.3036E-3
Less Bulk Condensation	1.3281E-3
Net Evaporation	9.7553E-4

Net Evaporation as % of Liquid Inflow	3.66%

Vapor Outflow less Inflow	9.7534E-4
less Net Evaporation Rate	9.7553E-4
Error	-1.8670E-007
Error as % of Liquid Inflow	-0.001%

Liquid Outflow less Inflow	-9.6883E-4
plus Net Evaporation Rate	9.7553E-4
Error	6.7055E-6
Error as % of Liquid Inflow	0.03%

The net evaporation is 3.66% percent of the mass flow.

This figure also represents the true quality (mass fraction of steam) at the outlet.

Visualizing the Results

Examine the results of the simulation.

The scalar scene depicting the volume fraction of the vapor is shown below.



The first thing to look at is the vapor volume fraction distribution and to compare it with the liquid subcooling. This comparison can be done by switching the Scalar Field that is selected for the Scalar Scene.

The scalar scene depicting the liquid subcooling is shown below.



A small negative subcooling appears next to the wall, which in effect means that the liquid is superheated by a few degrees. The reason that it has not yet evaporated, is that the rate of evaporation is limited by the rate of heat transfer between the liquid and the liquid-vapor interface. Most of the vapor is confined to the region where the subcooling is not too high, below around 10 K. Elsewhere the vapor is quickly condensed as soon as it is produced on the wall. The vapor appears to spread quickly across the flow in this scaled scene, but in fact it is spread by the turbulent velocity fluctuations, modeled using the turbulent dispersion force.

Comparison With Experiment

Compare the results of the simulation with experimental data.

The main result of interest is the axial distribution of mean void. Since this result is a post processing calculation that is difficult to do without making a special assumption of a uniform grid, this calculation is made using a customized user macro:

1. Select **File** > **Macro** > **Play Macro**.
2. In the *Open* dialog, navigate to your working directory and select `wb_createAxialPlots.java`.
3. Click **Open** to run the macro.

   The macro creates two new plots entitled *Compare Axial Void* and *Compare Calorimetric Temperature*. At the same time, the macro loads experimental data from file.

Note: To refresh these plots, run the `wb_deleteAxialPlots.java` macro first.

The *Compare Axial Void* plot is shown below:



The prediction of axial void is good near the start of net vapor generation and passes through the first few experimental points up to an intermediate level of void. Beyond this intermediate level, it over

predicts the void by about 20% at the last experimental point. The agreement could be improved with model tuning. However, it is more satisfactory that this set of models and relaxation factors can provide solutions to many of the cases that are specified in Bartolemei (1982) without any further adjustments.

This wide range of applicability is largely due to calibrations of Kocamustafaogulari and of Hibiki-Ishii to a wide range of pressures up to 150 bar. At 150 bar the bubble departure size becomes very small (about 5 micron compared to 50 micron for this test at 45 bar). It is also due to the capability of the S-gamma model to track the bubble size in subcooled flow starting from the Kocamustaogulari model for departure diameter.

For more information, see the following sections:

- Hibiki Ishii Nucleation Site Number Density
- Kocamustafaogullari Bubble Departure Diameter

The Calorimetric Temperature that is provided by Bartolemei (1967) is not fully explained in the paper, and is a little difficult to interpret. The name as translated from Russian suggests a velocity-weighted temperature averaged across the diameter, and the macro has prepared such an average from the computed results.

The *Compare Calorimetric Temperature* plot is shown below:



Bartolemei (1967) does indeed provide a large number of points for Calorimetric Temperature as if these were experimental data. These points appear to lie on a straight line until a constant saturation temperature is reached exactly at zero thermodynamic quality.

The STAR-CCM+ results deviate from this line when vapor is generated. This result is as expected from a heat balance, because if vapor coexists in subcooled conditions (thermodynamic quality < 0), liquid must be below the temperature that is expected for a homogeneous mixture.

Now consider the liquid subcooling along the axis. The experimental data indicates that there was already a temperature gradient upstream of the heated section.

The source reference [Bartolemei (1967)] gives no details here, but you are able to proceed by assuming that the simulation can model a short length of thermal boundary layer development before any boiling begins. This tutorial also uses the same thermodynamic quality scale to compare experiment and simulation, so that after some distance any differences due to inlet conditions become irrelevant.

The *Axis Subcooling* plot is shown below:



Here, as expected, the liquid subcooling is still above zero as the zero thermodynamic quality line is crossed. Heat is going into vapor production near the wall instead of heating the liquid core. The over prediction that STAR-CCM+ makes of this core subcooling is consistent with its over prediction of outlet vapor.

Finally compare the predicted and experimental wall superheat. Bartolemei (1967) temperature measurements are used on the outside of the boiling tube, converted to inner face temperatures using Bartolemei's own estimate of temperature drop in the tube wall.

The *Wall Superheat* plot is shown below:



It is clear that there was heating ahead of the specified heated section in the experiment, while STAR-CCM+ takes a little distance to develop a thermal boundary layer. After this distance, the agreement is good at the start of boiling (zero superheat) and at fully developed boiling nearer the outlet where the superheat is of order 6K. In between, in the mixed convection/evaporation region, the current model underestimates superheat.

Again this model could be tuned to improve the superheat for this particular case (Bartolemei does not provide temperature data for any other conditions). However, it is good to see that this model could be useful for fixed-temperature applications as well as for fixed flux applications.

Turbulence Checks

Examine the results to check the turbulence profile.

If you use the Scalar Scene to examine the Turbulent Viscosity Ratio of Liquid, you see that there are two troughs in turbulence across the flow.

The Scalar Scene is shown below:



The first occurs at the inlet and is entirely due to the use of flat velocity and turbulence profiles at the inlet. The second is due to the acceleration of the flow due to the onset of vapor production. The point here is to check that the first artificial feature is not interfering with the prediction of onset of boiling nor with the onset of vapor production. This scenario is the case here since the artificial first feature is well downstream of the inception of wall boiling. Under different conditions, the right thing to do would be to replace the flat inlet profiles by fully developed profiles of velocity, turbulence, and perhaps temperature.

The second thing to check is that the predicted value of y+ at the first cell center next to the wall is well above 30.

The *Wall Yplus of Liquid* plot is shown below:



The computed y+ value confirms that a high-y+ wall treatment could have been used for this calculation. The two-layer treatment is retained so that this same model can be used under a wider range of conditions. The wall plot shows that y+ starts around 65 in the single phase region and then increases as the flow accelerates with increasing vapor content.

Additional Information

Examine the detailed results for specific models.

If you are interested, you can inspect the more detailed results for various submodels.

Perhaps the most important of these results is the bubble size prediction next to the wall shown in plot Near-Wall Bubble size. Bubbles depart with a size of about 50 micron under these conditions. Where the water is still heavily subcooled, the S-Gamma model predicts that bubble size next to the wall is immediately smaller due to condensation. Once the liquid subcooling is small enough (liquid temperature is approaching saturation) the local condensation rate can no longer balance wall

evaporation rate. At this point on the wall, net vapor production begins, and also bubble size starts to rise due to coalescence.

The Near-Wall Bubble Size plot is shown below:



Some other points of interest to boiling modeling are:

- Near wall volume fraction reaches levels of 0.8 near the wall in this calculation. This result might not be physical, and could be modified with model adjustments. However, due to the difficulties of making measurements across narrow tubes operating at very high pressure, the Bartolemei (1967) and Bartolemei (1982) data provides no information to support such adjustments to radial distribution.
- Bubble departure frequency is a simple model that essentially depends on departure size, so is also constant in this constant property calculation.
- Bubble Influence Wall Area Fraction shows how the heat flux between the wall and the liquid phase is partitioned between fully developed convective heat transfer and bubble-induced quenching heat flux. Initially this fraction is small as bubble nucleation sites are

far apart, but rises as the number of active sites increases with wall superheat.

- The nucleation site number density plot shows a rapid change in number of active sites with very small change in wall superheat. It is this behavior at the heart of the wall boiling model that both enhances heat transfer rapidly and keeps the wall superheat relatively uniform along a boiling wall.

Summary

This tutorial validated the wall boiling model in STAR-CCM+ against experimental void and temperature data for high-pressure forced convective boiling.

This tutorial has covered the following:

- Selection of suitable models for Eulerian simulation of forced convective boiling at high pressure.
- Definition of phase enthalpies that are based on steam table conventions.
- Selection of constant properties that give the same mass and energy balance as variable properties.
- Monitoring of min/max phase velocities, temperatures, and volume fractions as an indicator of relaxation needs in a sensitive calculation.
- Use of wall heat flux ramping during flow development to keep velocities and temperatures bounded.
- Verification of phase mass balances with interphase mass transfer.
- Methods for extracting results for axial void, liquid subcooling, wall superheat, and wall heat flux partitioning.
- Use of the thermodynamic quality scale to compare results against experiment.
- Extensive use of field functions and macros to scale and analyze the results.

Bibliography

This tutorial references the following papers.

Bartolomei, G. G. and Chanturiya, V. M. 1967 "Experimental Study of True Void Fraction When Boiling Subcooled Water in Vertical Tubes" in Thermal Engineering, 14(2).

Bartolomei, G. G. et al 1982 "An Experimental Investigation of True Volumetric Vapor Content with Subcooled Boiling in Tubes" in Thermal Engineering, 29(3).

Eulerian: Conjugate Heat Transfer Wall Boiling

This tutorial demonstrates how to set up a conjugate heat transfer wall boiling analysis in STAR-CCM+. It simulates a nuclear fuel pin that is cooled using subcooled boiling under flow conditions typical in the first meter of a 3.7-m long BWR6 reactor channel.

The channel operates at 71.7 bar with an average heat flux of 524900 W/m^2 at the cladding surface.

An axisymmetric, two-dimensional volume mesh of the boiling tube is provided, as shown below.



The fuel, cladding, and coolant are modeled as three finite regions. There is a gap between the fuel and the cladding. The properties of this

gap vary as the fuel reacts and daughter constituents fill the space between the fuel and the cladding. Therefore, the gap is treated as an empirical contact resistance, with its thickness added to the fuel.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Mesh
Selecting the Physics Models
Setting the Solid Material Properties
Creating Phases and Selecting Phase Models
Defining the Phase Interactions
Setting the Fuel-Cladding Interface Properties
Setting Region Sources and Boundary Conditions
Setting Initial Conditions and Reference Values
Setting the Stopping Criteria
Visualizing the Temperature and Vapor Fraction
Running the Simulation
Visualizing the Results
Validating the Results
Summary

Prerequisites

The instructions in the Eulerian: Conjugate Heat Transfer Wall Boiling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/WallBoilingCht.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Create a geometry scene.
6. Save the simulation as `wallBoilingCHT.sim`.

Visualizing the Mesh

Create a mesh scene to view the mesh.

An axisymmetric mesh of the fuel pin, cladding, and coolant has been prepared for this analysis, so further meshing is not required. The boundary names and types are as defined in the imported mesh. View the mesh to examine these boundaries.

To visualize the mesh:

1. Create a mesh scene.

The fuel pin is 1 m long in the X direction and 0.01 m in the Y direction. To visualize the scene with an improved aspect ratio, create a simple transform to stretch the domain in the Y direction:

2. Right-click **Tools > Transforms** and select **New Graphics Transform > Simple Transform**.
3. Rename the newly created **Simple Transform 1** node to `Resize Domain`.
4. In the *Properties* window, set the **Scale** property to `[1.0, 20.0, 1.0]`.
5. To apply this transform to the mesh scene, select **Scenes > Mesh Scene 1 > Displayers > Mesh 1** and set its **Transform** expert property to **Resize Domain**.



6. To examine the imported boundaries, expand the **Regions** > **Fuel** > **Boundaries** node and select the **Axis** node.



This boundary is highlighted in **Mesh Scene 1**.

7. In the *Properties* window, make sure that the **Type** is set to **Axis**.



8. Similarly, select the **Coolant** > **Boundaries** > **Inlet** node and observe which part is highlighted in the scene.



9. Save the simulation.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

Three continua are set up to model each of the three regions.

To set up the physics models for the coolant:

1. For the physics continuum, **Continua** > **Physics 1**, in the *Enabled Models* group box, deactivate **Two-Dimensional**.
2. Select the following models in order:

Group Box	Model
Space	**Axisymmetric** **Gradients** (Selected automatically)
Time	**Steady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Multiphase Segregated Flow** **Multiphase Equation of State** (Selected automatically)
Viscous Regime	**Turbulent**
Optional Models	**Gravity** **Phase Coupled Fluid Energy**

3. Click **Close**.
4. Rename the **Physics 1** node to **Coolant**.
5. To review the models, open the **Coolant** > **Models** node.



6. Save the simulation.

To assign the **Coolant** continuum to the **Coolant** region:

7. Select the **Regions** > **Coolant** node and set its **Physics Continuum** to **Coolant**.



To set up the physics models for the cladding and fuel:

8. For the physics continuum, **Continua** > **Physics 2**, in the *Enabled Models* group box, deactivate **Two-Dimensional**.
9. Select the following models in order:

Group Box	Model
Space	**Axisymmetric** **Gradients** (Selected automatically)
Time	**Steady**
Material	**Solid**
Optional Models	**Segregated Solid Energy**
Equation of State	**Constant Density**

10. Click **Close**.

The physics models for the cladding and fuel are identical. However, since different material properties are used to model each region, two separate continua are created for the cladding and fuel:

11. Right-click **Continua** > **Physics 2** and select **Copy**.
12. Right-click **Continua** and select **Paste**.
13. Rename **Copy of Physics 2** to **Cladding**.
14. Rename **Physics 2** to **Fuel**.
15. Assign these two continua to their respective regions.

    The icons of the **Fuel** and **Cladding** regions turn to gray to indicate that they are both solid regions.



16. Save the simulation.

Setting the Solid Material Properties

Define the properties for the solid materials.

To set the properties for the solid materials:

Specify the material properties for the fuel.

1. Rename the **Continua** > **Fuel** > **Models** > **Solid** > **Al** node to **UO2**.

2. Edit the **UO2 > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`10420.0 kg/m^3`
Specific Heat > Constant	**Value**	`327.7 J/kg-K`
Thermal Conductivity > Constant	**Value**	`2.586 W/m-K`



Specify the material properties for the cladding.

3. Rename the **Cladding > Models > Solid > Al** node to **Zircaloy**.
4. Edit the **Zircaloy > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`6579.0 kg/m^3`
Specific Heat > Constant	**Value**	`314.5 J/kg-K`
Thermal Conductivity > Constant	**Value**	`15.61 W/m-K`

5. Save the simulation.

Creating Phases and Selecting Phase Models

Create the phases and select the appropriate phase models.

The two phases flowing through the coolant section of vertical pipe are water and steam. These phases are defined by creating them under the Eulerian Phases model, and selecting the corresponding material properties.

To create the phases and select the phase models:

Create the liquid phase.

1. In the **Coolant** continuum, right-click the **Models > Eulerian Multiphase > Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Liquid**.
3. For the **Liquid** phase, select the following models:

Group Box	Model
Enabled Models	**Turbulent**
Material	**Liquid** **Reynolds-Averaged Navier-Stokes** (Selected automatically)

4. Deactivate **Auto-select recommended models**.
5. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

6. Activate **Auto-select recommended models**.
7. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**
Energy	**Segregated Fluid Temperature**

8. Click **Close**.
9. Rename the **Liquid** > **Models** > **Liquid** > **H2O** node to **H2O (l)**.

Create the gas phase.

10. Create a second phase and rename it to **Vapor**.
11. For the **Vapor** phase, select the following models:

Group Box	Model
Material	**Gas** **Reynolds-Averaged Navier-Stokes** (Selected automatically)

12. Deactivate **Auto-select recommended models**.
13. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**

14. Activate **Auto-select recommended models**.
15. Select the following models in order:

Group Box	Model
Enabled Models	**High y+ Wall Treatment** (Selected automatically)
Equation of State	**Constant Density**
Energy	**Segregated Fluid Temperature**

16. Click **Close**.
17. Rename the **Vapor** > **Models** > **Gas** > **Air** node to **H2O (g)**.

    The **Eulerian Phases** node appears as shown.



The material properties for each phase are now modified, starting with the water phase.

18. Edit the **Liquid > Models > Liquid > H20 (l) > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`745.65 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`9.253E-5 Pas`
Heat of Formation (Saturation Enthalpy) > Constant	**Value**	`1276100 J/kg`

Node	Property	Setting
Specific Heat > Constant	**Value**	`5310.0 J/kg-K`
Standard State Temperature (Saturation Temperature) > Constant	**Value**	`560.65 K`
Thermal Conductivity > Constant	**Value**	`0.574 W/m-K`

Next set the properties of the vapor phase:

19. Edit the **Vapor > Models > Gas > H2O (g) > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`37.52 kg/m^3`
Dynamic Viscosity > Constant	**Value**	`1.9034E-5 Pas`
Heat of Formation (Saturation Enthalpy) > Constant	**Value**	`2770300 J/kg`
Specific Heat > Constant	**Value**	`5440 J/kg-K`
Standard State Temperature (Saturation Temperature) > Constant	**Value**	`560.65 K`
Thermal Conductivity > Constant	**Value**	`0.06514 W/m-K`

20. Save the simulation.

Defining the Phase Interactions

Define the interactions between the phases.

Using the **Multiphase Interaction** model, you can define the interaction between the liquid and vapor phases.

To define the phase interaction:

1. In the **Physics 1** continuum, right-click the **Models** > **Multiphase Interaction** > **Phase Interactions** node and create a new phase interaction.
2. Rename the **Phase Interaction 1** node to **Liquid Vapor Interface**.
3. For the **Liquid Vapor Interface** phase interaction, select the following models in order:

Group Box	Model
Phase Interaction Topology	**Eulerian Continuous-Dispersed**
Enabled Models	**Drag Force** (selected automatically) **Multiphase Material** (selected automatically) **Models Interaction Length Scale** (selected automatically) **Interaction Area Density** (selected automatically)
Optional Models	**Turbulent Dispersion Force** **Interphase Mass Transfer**
Interphase Mass Transfer Rate	**Boiling Mass Transfer Rate**
Optional Models	**Wall Boiling** **Lemmert Chawla Nucleation Site Number Density** (selected automatically) **Kurul Podowski Bubble Influence Wall Area Fraction** (selected automatically) **Del Valle Kenning Bubble Induced Quenching Heat Transfer Coefficient** (selected automatically) **Tolubinsky Kostanchuk Bubble Departure Diameter** (selected automatically) **Cole Bubble Departure Frequency** (selected automatically)

Group Box	Model
	Basic Wall Dryout Area Fraction (selected automatically)

4. Click **Close**.

   The **Liquid Vapor Interface** > **Models** node is shown below:



Define the phase interaction properties, starting with the definition of the continuous and dispersed phases.

5. Edit the **Liquid Vapor Interface** > **Models** node and set the following properties:

Node	Property	Setting
Eulerian Continuous-Dispersed	**Continuous Phase**	**Liquid**
	Dispersed Phase	**Vapor**

Select a more appropriate interaction length scale model.

6. Select the **Liquid Vapor Interface** > **Models** > **Interaction Length Scale** > **Interaction Length Scale** node and set the **Method** property as **Kurul Podowski** .

Define the remainder of the phase interaction properties.

7. Edit the **Liquid Vapor Interface > Models** node and set the following properties:

Node	Property	Setting
Turbulent Dispersion Force	**Use Logarithmic Form**	Activated
Boiling Mass Transfer Rate > Continuous Phase Nusselt Number	**Method**	**RanzMarshall**
Boiling Mass Transfer Rate > Interface Temperature > Constant	**Value**	`560.65 K`
Multiphase Material > Multiphase Material > Material Properties > Surface Tension > Constant	**Value**	`0.01725 N/m`
Drag Force > Drag Coefficient	**Method**	**Tomiyama**

The logarithmic form of the **Turbulent Dispersion Force** is used to improve the convergence rate. When the volume fraction of one phase is low and varies rapidly, the logarithmic form is differenced more accurately than the standard form.

Selecting the **RanzMarshall** option makes STAR-CCM+ calculate the Nusselt number using the Ranz-Marshall correlation. The Nusselt number is used to calculate the heat transfer coefficient for the heat transfer from the bulk water to the bubble interface.

The interface temperature is the saturation temperature of the liquid phase.

The **Tomiyama** model is appropriate for modeling a small bubble in a low viscosity, air-water system.

8. Save the simulation.

Setting the Fuel-Cladding Interface Properties

Define the Fuel-Cladding interface contact resistance.

You set the **Fuel-Cladding** interface contact resistance which models the gap between the two regions.

To set the Fuel-Cladding Interface properties:

1. Select **Interfaces** > **Fuel-Cladding** > **Physics Values** > **Contact Resistance** > **Constant** and set its **Value** to `1.761E-4 m^2-K/W`.



2. Save the simulation.

Setting Region Sources and Boundary Conditions

Define the boundary types and specify the appropriate property values.

Conditions are specified in the fuel region, and on the inlet and outlet boundaries of the coolant.

To set the boundary conditions:

Specify the volume energy source in the fuel.

1. Edit the **Regions** > **Fuel** node and set the following properties:

Node	Property	Setting
Physics Conditions > Energy Source Option	**Method**	**Volumetric Heat Source**
Physics Values > Volumetric Heat Source > Constant	**Value**	`2.26412E8 W/m^3`



Set the **Inlet** boundary conditions.

2. Edit the **Regions > Coolant > Boundaries > Inlet > Phase Conditions** node and set the following properties:

Node	Property	Setting
Liquid > Physics Values		
Velocity Magnitude	**Method**	**Constant**

Node	Property	Setting
Constant	**Value**	2.0076
Volume Fraction	**Method**	**Constant**
Constant	**Value**	0.9999
Vapor > Physics Values		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	2.0076
Volume Fraction	**Method**	**Constant**
Constant	**Value**	1.0E-4

3. Edit the **Coolant > Boundaries > Inlet** node and set the following properties:

Node	Property	Setting
Phase Conditions > Liquid > Physics Values > Static Temperature > Constant	**Value**	551.55 K
Vapor > Physics Values > Static Temperature > Constant	**Value**	560.65 K
Phase Conditions > Liquid > Physics Values > Turbulence Intensity > Constant	**Value**	0.05
Turbulent Viscosity Ratio > Constant	**Value**	400
Phase Conditions > Vapor > Physics Values > Turbulence Intensity > Constant	**Value**	0.05

Node	Property	Setting
Turbulent Viscosity Ratio > Constant	**Value**	200

Set the **Outlet** boundary conditions. At the outlet boundary, the volume fraction for each phase is extrapolated from the interior of the domain.

4. Edit the **Coolant > Boundaries > Outlet** node and set the following properties:

Node	Property	Setting
Physics Conditions > Volume Fraction Specification	**Option**	**Extrapolation**
Phase Conditions > Liquid > Physics Values > Static Temperature > Constant	**Value**	560.65 K
Vapor > Physics Values > Static Temperature > Constant	**Value**	560.65 K
Phase Conditions > Liquid > Physics Values > Turbulence Intensity > Constant	**Value**	0.05
Turbulent Viscosity Ratio > Constant	**Value**	400
Phase Conditions > Vapor > Physics Values > Turbulence Intensity > Constant	**Value**	0.05
Turbulent Viscosity Ratio > Constant	**Value**	200

5. Save the simulation.

Setting Initial Conditions and Reference Values

Set the initial conditions and reference values for the physics continuum.

Set the reference gravity, density, and pressure for the coolant.

To set the initial conditions and reference values:

1. Edit the **Coolant > Models > Eulerian Multiphase > Eulerian Phases** node and set the following properties:

Node	Property	Setting
Liquid > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[2.0076, 0.0, 0.0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`0.9999`
Vapor > Initial Conditions		
Velocity	**Method**	**Constant**
Constant	**Value**	`[2.0076, 0.0, 0.0] m/s`
Volume Fraction	**Method**	**Constant**
Constant	**Value**	`1.0E-4`

Set the initial conditions for each phase.

2. Edit the **Models > Eulerian Multiphase > Eulerian Phases > Liquid > Initial Conditions** node and set the following properties:

Node	Property	Setting
Static Temperature > Constant	**Value**	`551.55 K`
Turbulence Intensity > Constant	**Value**	`0.05`
Turbulent Viscosity Ratio > Constant	**Value**	`400`



3. Edit the **Models** > **Eulerian Multiphase** > **Eulerian Phases** > **Vapor** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Static Temperature > Constant	**Value**	`560.65 K`
Turbulence Intensity > Constant	**Value**	`0.05`
Turbulent Viscosity Ratio > Constant	**Value**	`200`

Set the reference values for the **Coolant** continuum.

4. Edit the **Reference Values** node and set the following properties:

Node	Property	Setting
Gravity	**Value**	`[-9.81, 0.0, 0.0] m/s^2`
Reference Density	**Value**	`745.65 kg/m^3`
Reference Pressure	**Value**	`7170000.0 Pa`



5. Save the simulation.

Setting the Stopping Criteria

Set the appropriate stopping criteria for the simulation.

The simulation runs for 200 iterations.

To set the stopping criteria:

1. Edit the **Stopping Criteria** node and set the following property:

Node	Property	Setting
Maximum Steps	**Enabled**	Activated

Node	Property	Setting
	Maximum Steps	200

2. Save the simulation.

Visualizing the Temperature and Vapor Fraction

Create a scalar scene to view the results of the simulation.

You now create a scalar scene that displays the temperature of the fuel and cladding, and the coolant vapor volume fraction. Both field functions can be monitored while the solver is running.

To visualize the temperature and vapor fraction:

1. Create a scalar scene.
2. Rename **Scenes** > **Scalar Scene 1** to **Dual Scene**.
3. Display the temperature in the fuel and cladding regions:
   a) Select the **Dual Scene** > **Displayers** > **Scalar 1** > **Parts** node.
   b) In the *Properties* window, click ... (**Custom editor**) for the **Parts** property.
   c) In the *Parts* dialog, expand the **Regions** node.
   d) Select the **Cladding** and **Fuel** regions.



e) Click **OK**.

f) Select the **Scalar 1 > Scalar Field** node and set its **Function** property to **Temperature**.

g) Set the **Auto Range** property to **Off**.

h) Set the **Min** property to `550.0`.

i) Set the **Max** property to `1350.0`.



4. Create a new scalar displayer for the coolant:
   a) Right-click the **Dual Scene > Displayers** node and select **New Displayer > Scalar**.
   b) Select the **Scalar 2 > Parts** node.
   c) In the *Properties* window, click ... (**Custom editor**) for the **Parts** property.
   d) Select **Regions > Coolant** and click **OK**.
   e) Select **Scalar 2 > Scalar Field** and set its **Function** property to **Volume Fraction > Vapor**.
5. Apply the **Resize Domain** transform to **Outline 1**, **Scalar 1** and **Scalar 2** to improve the aspect ratio of the scene:

a) Select the **Dual Scene** > **Displayers** > **Outline 1** node and set its **Transform** property to **Resize Domain**.

b) Select the **Dual Scene** > **Displayers** > **Scalar 1** node and set its **Transform** property to **Resize Domain**.

c) Select the **Dual Scene** > **Displayers** > **Scalar 2** node and set its **Transform** property to **Resize Domain**.

6. Using the mouse buttons, move the legends and adjust the scene so that it appears as shown.



7. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 200 iterations are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Steps satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

Modify the scalar scene to better visualize the results of the simulation.

The scalar scene depicting the volume fraction of the vapor and the temperature of the cladding and fuel is shown below.



Make the vapor volume fraction representation smoother by changing the contour style of its scalar displayer.

To visualize the results:

1. Select **Scenes** > **Dual Scene** > **Displayers** > **Scalar 2**.
2. In the *Properties* window, set the **Contour Style** to (**Smooth Filled**).

   The *Dual Scene* updates as shown.



3. Save the simulation.

Validating the Results

Compare the results of the simulation with experimental data.

The computed temperature profile is validated by assuming that:

- Saturation conditions are obtained after a distance of 1 m.
- There are 6 K of superheat at the wall.
- An analytical solution for axisymmetric heat conduction can represent the heat transfer across the fuel and cladding.
- There is a contact resistance of 1.761E-4 (m^2K/W) at the cladding-fuel interface.
- The expected temperature at the center of the fuel is about 1000 C.

Import the expected radial temperature profile from a provided file:

1. Right-click **Tools** > **Tables** and select **New Table** > **File Table**.
2. In the *Open* dialog, navigate to the `doc/startutorialsdata/multiphaseFlow/data` folder of your STAR-CCM+ installation and select file `WallBoilingCht.csv`.
3. Click **Open** to start the import.

   The **WallBoilingCht** child node is created under **Tables**.

You create two derived parts that extract the radial temperature across the domain. Start with the fuel and cladding regions:

4. Right-click **Derived Parts** and choose **New Part** > **Probe** > **Line**.
5. Click the **[Cladding, Coolant, Fuel]** button in the *Input Parts* group-box.
6. Expand the **Regions** node and clear **Coolant**.
7. Click **OK**.
8. Enter the coordinates of *Point 1* and *Point 2* as shown:



9. Leave the *Coordinate System* as **Laboratory**.
10. Enter a *Resolution* of `5`.
11. Select **No Displayer** from the *Display* group box.

    The *Create Line Probe* window appears as shown.



12. Click **Create**, and then **Close**.
13. Rename **Derived Parts > line-probe** to **Fuel and Cladding**.



To create a derived part for the coolant region, copy the **Fuel and Cladding** line probe and edit its properties:

14. Right-click **Fuel and Cladding** and select **Copy**.
15. Right-click **Derived Parts** and select **Paste**.
16. Rename **Copy of Fuel and Cladding** to **Coolant**.
17. In the *Properties* window, set **Point 1** to `[0.97917, 0.0060, 0.0] m, m, m`.
18. Set **Point 2** to `[0.97917, 0.01, 0.0] m, m, m`.
19. Click ... (**Custom Editor**) for the **Parts** property.
20. In the *Coolant* dialog, right-click on a blank area and select **Deselect All**.
21. Activate **Coolant**.



22. Click **OK**.

    The *Coolant - Properties* window appears as shown.



Create an X-Y plot that displays the radial temperature profile:

23. Right-click **Plots** and select **New Plot > X-Y**.

24. Rename **XY Plot 1** to **Radial Temperature**.
25. In the *Properties* window, set its **Title** property to **Radial Temperature**.
26. Click ... (**Custom Editor**) for the **Parts** property.
27. In the *Radial Temperature* dialog, activate **Derived Parts**.

    Both **Coolant** and **Fuel and Cladding** parts are selected.



28. Click **OK**.



29. Select **Radial Temperature** > **X Type** and set its **Type** property to **Direction**.
30. Select **Radial Temperature** > **X Type** > **Vector Quantity** and set its **Value** property to `[0.0, 1.0, 0.0]`.
31. Select **Radial Temperature** > **Y Types** > **Y Type 1** > **Scalar Function** and set its **Field Function** property to **Temperature**.

Create a second Y Type that displays the temperature in the coolant:

32. Right-click **Radial Temperature** > **Y Types** and select **New**.
33. Select **Radial Temperature** > **Y Types** > **Y Type 2** > **Scalar Function** and set its **Field Function** property to **Temperature** > **Temperature of Liquid**.

The temperature profile in the coolant is not plotted. Remove it from the display:

34. Select **Y Types** > **Y Type 1** > **Coolant** and, in the *Properties* window, deactivate **Lock Legend Name** and then erase the **Legend Name**.



35. Select **Y Types** > **Y Type 1** > **Coolant** > **Symbol Style** and set its **Shape** to **None**.
36. Select **Y Types** > **Y Type 2** > **Fuel and Cladding** and, in the *Properties* window, deactivate **Lock Legend Name** and then erase the **Legend Name**.
37. Select **Y Types** > **Y Type 2** > **Fuel and Cladding** > **Symbol Style** and set its **Shape** to **None**.

Import the expected profile as a table and add it to the plot:

38. Right-click **Radial Temperature** > **Data Series** and select **Add Data**.
39. In the *Add Data Providers to Plot* dialog, select **WallBoilingCht** and then click **OK**.
40. Select the **Plots** > **Radial Temperature** > **Data Series** > **WallBoilingCht** node and change the **Legend Name** to `Expected Profile`.
41. Set the **X Column** to **r**.
42. Set the **Y Column** to **T**.

43. Select **WallBoilingCht** > **Symbol Style** and set its **Shape** to **None**.
44. Select the **WallBoilingCht** > **Line Style** node.
45. In the *Properties* window, set the **Style** to **Solid**.
46. Click (**Custom Editor**) for the **Color** property.
47. In the *Line Style - Colour* dialog, click the *Palette* tab.
48. Select **Red**.



49. Click **OK**.
50. To add the X-axis title, select **Radial Temperature** > **Axes** > **X Axis** > **Title** and, in the *Properties* window, set the **Title** to `Position [0.0, 1.0, 0.0] (m)`.
51. To add the Y-axis title, select **Radial Temperature** > **Axes** > **Y Axis** > **Title** and, in the *Properties* window, set the **Title** to `Temperature (K)`.

    The *Radial Temperature* plot appears as shown.



The computed temperature shows good agreement with the expected radial profile.

52. Save the simulation.

Summary

This tutorial demonstrates how to set up a conjugate heat transfer wall boiling analysis in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the Eulerian Multiphase and Multiphase Turbulence models.
- Creating phases and selecting phase models for wall boiling.
- Defining phase interactions.
- Visualizing two field functions on the same scene.
- Plotting a radial temperature profile using two line probes.

Fluid Film: Liquid Film Flow

This tutorial demonstrates how to use the fluid film modeling capability of STAR-CCM+. It simulates the progress of a rectangular film of water flowing down the wall of a three-dimensional duct.

Water flows down the wall (wetted wall in the following figure) due to gravity, forming a liquid film in the process.



The initial water is 1 mm thick and covers a wall length and width of 250 mm and 200 mm, respectively, and is placed at the top wall. The surrounding gas is initially at rest. The simulation is run using the implicit unsteady solver for a simulated time of 10 seconds, which is long enough to satisfy the similarity condition. For the first 1000 time steps, the time step size is set to 0.001 seconds. Thereafter, a time step size of 0.01 seconds is used.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Volume Mesh
Selecting the Physics Models

Setting Fluid Film Models and Initial Conditions
Setting Up Shell Region and Boundary Conditions
Setting Reference Values
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Fluid Film: Liquid Film Flow tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/duct.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
5. Save the simulation as `wallFilm.sim`.

Visualizing the Volume Mesh

Examine the imported volume mesh.

To visualize the volume mesh:

1. Right-click the **Scenes** node and select **New Scene** > **Mesh**.
2. Examine the *Mesh Scene 1* display in the *Graphics* window.
3. Click in the toolbar and select **Views** > **+X +Y +Z** > **Up +Y**.
4. Use the mouse buttons, zoom in and position the part so that the view is similar to the view shown below.



5. Click (**Save-Restore-Select views**) in the toolbar and select **Store Current View**.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

The main phenomena to solve for in this simulation is unsteady, laminar, incompressible, gaseous flow. Set up an implicit unsteady, laminar, incompressible, gaseous flow using the segregated flow and fluid film models.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	**Model**
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Laminar**
Optional Models	**Fluid Film** **Gravity**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Setting Fluid Film Models and Initial Conditions

Set the fluid film models and the initial conditions for the fluid film phases.

To set the fluid film models and initial conditions:

You first define the fluid film phases. In this tutorial, you define only one phase for water.

1. In the **Physics 1** continuum, right-click the **Models** > **Fluid Film** > **Fluid Film Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Water Film**.
3. For the **Water Film** phase, select the following models in order:

Group Box	Model
Space	**Shell Three-Dimensional Gradients** (selected automatically) **Flow Model** (selected automatically) **Laminar** (selected automatically)
Material	**Liquid**
Equation of State	**Constant Density** (selected automatically)

4. Click **Close**.

The expanded **Water Film** > **Models** node is shown below.



You specify the initial water thickness with a field function, as you only want film to appear initially at the top of the geometry. Assume that the initial water is 1 mm thick and covers a wall length and width of 250 mm and 200 mm, respectively.

5. Define the field function to use for specifying the initial water thickness:

   a) Right-click the **Tools > Field Functions** node and select **New > Scalar**.

   b) Rename the **User Field Function 1** node to `Initial Film Thickness`.

   c) Select the **Initial Film Thickness** node and set the following properties:

Property	Setting
Dimensions	**Length**
Function Name	`InitialFilmThickness`
Definition	`(($$Centroid[0]<0.25) && ($$Centroid[2]>0.025) && ($$Centroid[2]<0.225)) ? 0.001: 0`

Set the initial conditions for the fluid film.

6. Edit the **Water Film** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Fluid Film Thickness	**Method**	**Field Function**
Field Function	**Scalar Function**	**InitialFilmThickness**



7. Save the simulation.

Setting Up Shell Region and Boundary Conditions

Create a shell region and define the boundary conditions.

In this tutorial liquid film forms on the **Wetted Wall** boundary of the **Fluid** region. A shell region is required for all boundaries on which fluid film is to flow. This shell region is created from the Wetted Wall boundary.

To create a shell region and define the boundary conditions:

1. Right-click the **Regions** > **Fluid** > **Boundaries** > **Wetted Wall** node and select **Create Shell Region**.



A new shell region is added to the simulation tree. In addition, an interface is formed between the fluid region and the shell region, which is based on the geometry of the Wetted Wall boundary surface. The original boundary therefore becomes an interface to the film shell region.



The water film phase must be assigned to the **Wetted Wall Shell** region.

2. Select the **Regions** > **Wetted Wall Shell** node and set the **Continuum** to **Water Film**.



Specify the actual fluid film boundary conditions using the nodes within the Wetted Wall Shell region. These boundaries are formed from the edges of the original Wetted Wall boundary from which the shell region was created. By default all boundaries are set to no-slip wall, except the Interface to Fluid [In-place 1], which is set to the Fluid-Film Boundary type. So that the fluid film can leave the domain, you change the boundary type of the Bottom edges boundary. Set the type for this boundary to Flow-Split Outlet.

3. Select the **Regions** > **Wetted Wall Shell** > **Boundaries** > **Bottom edges** node and set **Type** to **Flow-Split Outlet**.



4. Save the simulation.

Setting Reference Values

Set the reference values for the **Physics 1** continuum.

The reference values of the **Physics 1** continuum must be defined to make sure that the acceleration due to gravity acts in the expected direction. In this tutorial, gravity acts in the positive X direction.

To set the reference values:

1. Edit the **Physics 1** > **Reference Values** node and set the following property:

Node	Property	Setting
Gravity	**Value**	`[9.81, 0.0, 0.0] m/s^2`



Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

By default the time step size is set to 1 s. Modify this setting and prescribe a variable time step size according to the following rule:

- When time < 1 s, time step = 0.001 s
- Otherwise, time step=0.01 s

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** > **Implicit Unsteady** node and set the following properties:

Property	Setting
Time-Step	`($Time<1) ? 0.001 : 0.01`
Temporal Discretization	**2nd-order**

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`10 s`

Node	Property	Setting
Maximum Inner Iterations	**Maximum Inner Iterations**	5
Maximum Steps	**Enabled**	Deactivated



3. Save the simulation.

Visualizing the Solution

Create reports, monitors, and plots to view the results of the solution.

The fluid film thickness can be monitored in a number of ways. In this section, create a report, monitor, and plot of the maximum film thickness, and an X-Y plot of film thickness along the vertical surface.

Contents

Creating a Maximum Report, Monitor, and Plot
Creating an X-Y Plot of Fluid Film Thickness
Creating a Scalar Scene

Creating a Maximum Report, Monitor, and Plot

Create a report, monitor, and plot to monitor the maximum film thickness along the wetted wall.

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. To monitor a quantity, you set up a report that defines the quantity of interest and the parts of the region to monitor. A monitor is then defined based on that report. This monitor can also be used to create an appropriate X-Y graph plot.

To create a maximum report, monitor, and plot:

Set up a report to monitor the maximum film thickness along the wetted wall.

1. Right-click the **Reports** node and select **New Report** > **Maximum**.



2. Rename the **Reports** > **Maximum 1** node to `Maximum Film Thickness`.
3. Select the **Maximum Film Thickness** node and do the following:
   a) Set the *Scalar Field Function* to **Fluid Film Thickness**.



b) Click to the right of the *Parts* property.

c) In the dialog that appears, expand the **Regions > Wetted Wall shell** node and select **Wetted Wall**.



d) Click **OK**.

Create a monitor and a plot for the report.

4. Right-click the **Maximum Film Thickness** node and select **Create Monitor and Plot from Report**.



5. Select the **Monitors > Maximum Film Thickness Monitor** node and check that the *Trigger* property is set to **Time Step**.



6. Right-click the **Plots > Maximum Film Thickness Monitor Plot** node and select **Open**.



Creating an X-Y Plot of Fluid Film Thickness

Create an X-Y plot of the fluid film thickness along a plane section through the center of the wetted wall.

To create an X-Y plot of the fluid film thickness:

Create the plane section.

1. Activate the *Mesh Scene 1* in the *Graphics* window.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.
3. In the *Create Section* window, do the following:
   a) In the *Input Parts* box, click the drop-down menu. In the dialog that appears, use (**Column Selection**) to clear selections, then expand the **Regions** > **Wetted Wall shell** node and select **Wetted Wall**.



   b) Click **OK**.

c) Set *Normal* to `[0.0, 0.0, 1.0]` m.

d) Under *Display*, select the **No Displayer** radio button.

e) Click **Create** and then **Close**.

A new derived part is created called **plane section**.

4. To view this part in the *Graphics* window, open the **Derived Parts** node and select **plane section**.



The location of the line running down the surface is highlighted in the *Graphics* window.

5. Rename the **Derived Parts > plane section** node to `Wall Profile`.

Create an X-Y plot for the fluid film thickness using the plane section.

6. Right-click the **Plots** node and select **New Plot > X-Y**.
7. Rename the **Plots > XY Plot 1** node to `Film Thickness Profile`.
8. With the **Film Thickness Profile** node selected, click to the right of the *Parts* property. In the dialog that appears, expand the **Derived Parts** node and select **Wall Profile**.



9. Click **OK**.
10. Select the **Film Thickness Profile > X Type** node and set the *Type* property it to **Scalar**.



11. Select the **Film Thickness Profile > X Type > Scalar Function** node and click to the right of the *Field Function* property and, in the *Scalar Function - Field Function* dialog, select **Centroid > Laboratory > X**.



12. Select the **Film Thickness Profile > Y Types > Y Type 1 > Scalar Function** node and click to the right of the *Field Function* property

and in the *Scalar Function - Field Function* dialog, select **Fluid Film Thickness**.



Creating a Scalar Scene

Create a scalar scene for fluid film thickness.

To create a scalar scene:

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
   A new node, **Scalar Scene 1**, is created under the **Scenes** node.
2. Rename the **Scenes** > **Scalar Scene 1** node to `Film Thickness`.
3. Above the object tree, click **scene/plot**.
4. Select the **Film Thickness** > **Displayers** > **Scalar 1** > **Parts** node, and click to the right of the *Parts* property.
5. In the dialog that appears, expand the **Regions** > **Wetted Wall shell** node and select **Wetted Wall**.



6. Click **OK**.
7. Select the **Film Thickness > Displayers > Scalar 1 > Scalar Field** node.
8. Click to the right of the *Function* property and, in the *Scalar Field - Function* window, select **Fluid Film Thickness**.



9. Click (**Save-Restore-Select views**) in the toolbar and select **Restore View > View 1**.



10. Select the **Film Thickness > Displayers > Scalar 1 > Color Bar** node and set *Width* to `0.1`, *Height* to `0.6`, *Orientation* to **Vertical** and *Position* to `[0.75, 0.33]`.



11. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

2. Activate the film thickness scene and observe the formation of fluid film as the solution progresses.



The example above shows the fluid film thickness at early stages of the run. The simulation continues until 10 seconds of physical time is complete.

When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

3. When the simulation has finished running, save it.

Visualizing the Results

Examine the results of the simulation.

The film thickness contour for the converged solution appears as shown below.



1. Activate the maximum film thickness monitor plot.

   The maximum film thickness plot for the converged solution appears as shown below.



Summary

This tutorial demonstrated how to use the fluid film modeling capability of STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Defining models for fluid film model.
- Setting fluid film models and initial conditions using field function.
- Setting up shell region and boundary conditions.
- Setting up variable time-step size for the implicit solver.
- Running the solver for a set physical time.

Fluid Film: Binary Liquid Film Flow with Evaporation and Edge Stripping

This tutorial demonstrates how to use the fluid film modeling capability of STAR-CCM+ to set up a multi-component film scenario with evaporation and edge stripping. It simulates the progress of a film mixture of water and ethylene glycol flowing down a partially heated wall.

The liquid mixture flows down the wetted wall due to gravity, forming a film. The film evaporates and edge stripping occurs at the 153.4º bend.



The initial liquid film is at 293.15 K and covers the whole wetted wall with a 0.343 mm thickness. The surrounding multi-component gas is

initially at rest and is composed of 99% dry air and 1% water vapor at 293.15 K. The case is run using the implicit unsteady solver for a simulated time of 5.0 seconds. The time step size is set to 0.01 seconds.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Volume Mesh
Selecting the Physics Models
Setting the Multi-Component Gas Properties and Initial Conditions
Setting Fluid Film Models and Initial Conditions
Setting the Fluid Film Evaporation and Condensation Model
Setting the Lagrangian Model
Setting the Gas Boundary Conditions
Setting Up Shell Region and Boundary Conditions
Setting Up the Injector
Setting Up the Film Stripping Model
Setting Reference Values
Setting Solver Parameters and Stopping Criteria
Setting Up Evaporation Rate Reports, Monitors, and Plots
Setting Up a Scalar Scene
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Fluid Film: Binary Liquid Film Flow with Evaporation and Edge Stripping tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

To import the mesh and name the simulation:

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/stripping.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.
5. Save the simulation as `strippingFilm.sim`.

Visualizing the Volume Mesh

Examine the imported volume mesh.

To visualize the volume mesh:

1. Create a mesh scene.
2. Position the part to obtain a similar view to the one shown below.



3. Click (**Save-Restore-Select views**) in the toolbar and select **Store Current View**.

Selecting the Physics Models

Physics models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution.

The main phenomena to solve for in this simulation is laminar, incompressible, multi-component gaseous flow. You use the segregated flow and fluid enthalpy models, and additionally select the multi-component fluid film, multiphase interaction, and Lagrangian multiphase models.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Time	**Implicit Unsteady**
Material	**Multi-Component Gas**
Reaction Regime	**Non-reacting**
Flow	**Segregated Flow** **Segregated Species** (selected automatically) **Gradients** (selected automatically)
Equation of State	**Constant Density**
Viscous Regime	**Laminar**
Optional Models	**Fluid Film** **Gravity** **Segregated Fluid Enthalpy** **Multiphase Interaction** **Lagrangian Multiphase**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Setting the Multi-Component Gas Properties and Initial Conditions

Define the gas components in the gas continuum and set their properties.

To set the multi-component gas properties and initial conditions:

1. Right-click **Multi-Component Gas > Gas Components** and choose **Select Mixture Components**.
2. In the *Select Mixture Components* dialog, expand the **Material Databases > Standard > Gases** node and select the following gases: **H2O (Water)**, **C2H4 (Ethylene)**, and **Air (Air)**.



3. Click **Apply**, and then **Close**.

Specify custom properties to model ethylene glycol vapor.

4. Rename the **Gas Components** > **C2H4** node to **C2H6O2**, the chemical symbol for ethylene glycol.
5. Select the **Gas Components** > **C2H6O2** > **Material Properties** > **Molecular Weight** > **Constant** node and change its *Value* to `62.07 kg/kmol`.



Set the initial conditions for the multi-component gas continuum. The gas is initially at 293.15 K.

6. Select **Continua > Physics 1 > Initial Conditions > Static Temperature > Constant** and set *Value* to `293.15 K`.

The gas initially consists of 99% dry air and 1% water vapor.

7. Select **Continua > Physics 1 > Initial Conditions > Species Mass Fraction > Constant** and click ... (**Custom Editor**) for the *Value* property.
8. In the *Constant - Value* dialog, set the values of *H2O* and *Air* to `0.01` and `0.99` respectively.



The sum of all mass fractions must be equal to `1.0`.

9. Click **OK**.

Setting Fluid Film Models and Initial Conditions

Set the fluid film models and the initial conditions for the fluid film phases.

You define the fluid film phases for H2O and C2H6O2 and set the initial conditions for the fluid film. The film is initially at 293.15 K and has a constant thickness of 0.343 mm. The mass fractions of water and ethylene glycol are initially equal.

To set the fluid film models and initial conditions:

Define the fluid film phases for H2O and C2H6O2.

1. In the **Physics 1** continuum, right-click the **Models** > **Fluid Film** > **Fluid Film Phases** node and create a new phase.
2. Rename the **Phase 1** node to **Fluid Film**.
3. For the **Fluid Film** phase, select the following models in order:

Group Box	Model
Space	**Shell Three-Dimensional Flow Model** (selected automatically) **Laminar** (selected automatically)

Group Box	Model
	Segregated Temperature (selected automatically)
Material	**Multi-Component Liquid** **Constant Density** (selected automatically) **Segregated Species** (selected automatically)

4. Click **Close**.

   The simulation tree appears as follows:



Select the liquid film components in the same way as for the multi-component gas continuum. Ethylene glycol is not listed in the liquid database; select ethanol and modify its name and properties to match the name and properties of ethylene glycol.

5. Right-click **Fluid Film > Models > Multi-Component Liquid > Liquid Components** and choose **Select Mixture Components**.
6. Expand the **Material Databases > Standard > Liquids** node and select **H2O (Water)** and **C2H6O (Ethanol)**.



7. Click **Apply**, and then **Close**.
8. Rename the **Liquid Components > C2H6O** node to **C2H6O2**.

Set the density and molecular weight values of ethylene glycol.

9. Edit the **Liquid Components > C2H6O2 > Material Properties** node and set the following properties:

Node	Property	Setting
Density > Constant	**Value**	`1113.2 kg/m^3`
Molecular Weight > Constant	**Value**	`62.07 kg/kmol`



Set the initial conditions for the fluid film.

10. Edit the **Fluid Film > Initial Conditions** node and set the following properties:

Node	Property	Setting
Fluid Film Thickness > Constant	**Value**	`3.43E-4 m`
Static Temperature > Constant	**Value**	`293.15 K`

Node	Property	Setting
Species Mass Fraction > Constant	**Value**	`[0.5, 0.5]`

11. Save the simulation.

Setting the Fluid Film Evaporation and Condensation Model

Set up the fluid film evaporation and condensation model.

The fluid film evaporation and condensation model is set as a Film-Eulerian phase interaction in STAR-CCM+.

To create this phase interaction:

1. Right-click **Continua > Physics 1 > Models > Multiphase Interaction > Phase Interactions** and select **New**.

   A new **Phase Interaction 1** node is added to the **Phase Interactions** folder.
2. Rename the **Multiphase Interaction > Phase Interactions > Phase Interaction 1** node to **Film Mixture - Gas**.

To set up the Film-Eulerian phase interaction:

3. Right-click **Multiphase Interaction > Phase Interactions > Film Mixture - Gas > Models** and choose **Select Models...**.

In the *Film Mixture - Gas Model Selection* dialog:

4. Select **Film-Eulerian Phase Interaction** from the *Phase Interaction Topology* box.
5. Click **Close**.
6. Select **Film Mixture - Gas > Models > Film-Eulerian Phase Interaction** and set the *Fluid Film Phase* to **Fluid Film**.

   The **Component Mapping** node appears in the simulation tree to indicate that the phase interaction is properly configured.



To select the evaporation and condensation model:

7. Select **Multiphase Interaction** > **Phase Interactions** > **Film Mixture - Gas** > **Models** and choose **Select Models...**.

In the *Film Mixture - Gas Model Selection* dialog:

8. Select **Evaporation/Condensation** and **Multiphase Material** from the *Optional Models* box.
9. Click **Close**.

Specify the connection between components in the gas phase and components in the fluid film phase for evaporation or condensation:

10. Select **Multiphase Interaction** > **Phase Interactions** > **Film Mixture - Gas** > **Models** > **Component Mapping** and click ... (**Custom Editor**) for the *Connectivity* property.
11. In the *Connectivity* dialog, map the fluid film components to their corresponding gas components, as shown below.



12. Click **OK**.

To model realistic saturation pressure behavior, use the **Antoine Equation** method. For this simulation, you can use default values for H2O:

13. Select **Continua > Physics 1 > Models > Fluid Film > Fluid Film Phases > Fluid Film > Models > Multi-Component Liquid > Liquid Components > H2O > Material Properties > Saturation Pressure** and set *Method* to **Antoine Equation**.

    The Saturation Pressure node turns to **Antoine Equation**.



To input the parameters of the *Antoine Equation* for C2H6O2:

14. Continuing with the **Liquid Components** node, select **C2H6O2 > Material Properties > Saturation Pressure** and set *Method* to **Antoine Equation**.

    The **Saturation Pressure** node turns to **Antoine Equation**.

15. Select the **Antoine Equation** node and set the properties to the values shown below:



Setting the Lagrangian Model

Set up a Lagrangian multiphase model to simulate stripping droplets.

To set the Lagrangian model:

1. Right-click **Continua** > **Physics 1** > **Models** > **Lagrangian Multiphase** > **Lagrangian Phases** and select **New**.
2. Rename the newly created **Phase 1** to **Droplets**.
3. For the Lagrangian phase, **Droplets**, select the following models in order:

Group Box	Model
Particle Type	**Material Particles** **Spherical Particles** (selected automatically) **Pressure Gradient Force** (selected automatically)
Material	**Multi-Component Liquid** **Species** (selected automatically)
Equation of State	**Constant Density**
Optional Particle Forces	**Drag Force**
Optional Models	**Two-Way Coupling** **Energy**

4. Click **Close**.

   The simulation tree appears as shown below.



To set up the droplet components:

5. Right-click **Multi-Component Liquid** > **Droplet Components** and choose **Select Mixture Components**.
6. In the *Select Mixture Components* dialog, expand the **Material Databases** > **Standard** > **Liquids** node and select **H2O (Water)** and **C2H6O (Ethanol)**.
7. Click **Apply**, and then **Close**.
8. Rename the **Droplets Components** > **C2H6O** node to **C2H6O2**.

To set the density and molecular weight values for ethylene glycol:

9. Select **Droplet Components** > **C2H6O2** > **Material Properties** > **Density** > **Constant** and set *Value* to `1113.2 kg/m^3`.
10. Select **Droplet Components** > **C2H6O2** > **Material Properties** > **Molecular Weight** > **Constant** and set *Value* to `62.07 kg/kmol`.

Setting the Gas Boundary Conditions

Define the boundary conditions and specify the appropriate property values.

The gas enters the domain at 1.0 m/s with a temperature of 293.15 K, and a composition of 99% dry air and 1% H2O.

To set the gas boundary conditions:

Set up the inlet.

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Top** node and set the following properties:

Node	Property	Setting
Physics Values > Species Mass Fraction > Constant	**H2O**	`0.01`
	Air	`0.99`
Physics Values > Static Temperature > Constant	**Value**	`293.15 K`



Leave the default velocity magnitude to 1.0 m/s.

Set up an outlet for the gas to leave at the bottom of the domain. The outlet gas composition is unknown and cannot be specified. The outlet gas composition is therefore matched to the composition inside the domain. Set the outlet temperature to 293.15 K.

2. Edit the **Regions** > **Fluid** > **Boundaries** > **Bottom** node and set the following properties:

Node	Property	Setting
Bottom	**Type**	**Pressure Outlet**

Node	Property	Setting
Physics Values > Species Mass Fraction	**Method**	**Composite**
Composite > H2O	**Method**	**Field Function**
Field Function	**Scalar Function**	**Mass Fraction > Mass Fraction of H2O**
Composite > C2H6O2	**Method**	**Field Function**
Field Function	**Scalar Function**	**Mass Fraction > Mass Fraction of C2H6O2**
Composite > Air	**Method**	**Field Function**
Field Function	**Scalar Function**	**Mass Fraction > Mass Fraction of Air**
Physics Values > Static Temperature > Constant	**Value**	`293.15 K`



Setting Up Shell Region and Boundary Conditions

Create a shell region and define the boundary conditions.

In this tutorial, the binary liquid film forms on the **Wetted Wall** boundary of the **Fluid** region. A shell region is required for all boundaries on which the fluid film is to flow. You create this shell region from the **Wetted Wall** boundary.

To create a shell region and define the boundary conditions:

1. Right-click the **Regions** > **Fluid** > **Boundaries** > **Wetted Wall** node and select **Create Shell Region**.



A new shell region is added to the simulation tree. In addition, an interface is formed between the fluid region and the shell region, based on the geometry of the **Wetted Wall** boundary surface.



Assign the fluid film phase to the **Wetted Wall Shell** region:

2. Select the **Regions** > **Wetted Wall Shell** node and set the **Continuum** to **Fluid Film**.

Specify the actual fluid film boundary conditions using the nodes within the **Wetted Wall Shell** region. These boundaries are formed from the edges of the original **Wetted Wall** boundary from which the shell region was created. By default all boundaries are set to no-slip wall, except the **Interface to Fluid [In-place 1]**, which is set to the **Fluid-Film Boundary** type.

Change the boundary type of the **Bottom edges**, **Symmetry edges**, and **Top edges** boundaries to their appropriate types.

3. Edit the **Regions** > **Wetted Wall Shell** > **Boundaries** node and set the boundary types as follows:

Boundary	Type
Bottom edges	**Flow-Split Outlet**
Symmetry edges	**Symmetry Plane**
Top edges	**Mass Flow Inlet**

The **Wetted Wall** is heated in one section only. You can define the heat flux in a field function.

4. Right-click the **Tools** > **Field Functions** node and select **New** > **Scalar** from the pop-up menu.
5. Rename the **User Field Function 1** to `Heat Flux Wetted Wall`.
6. Select the **Heat Flux Wetted Wall** node and set the following properties:

Property	Setting
Function Name	`heatFluxWettedWall`
Definition	`($$Position[2] <= 0.2)? 0 : 3000`

7. Set **Dimensions** as **Length** = `-2` and **Power** = `1`.



Specify the **Wetted Wall** heat flux.

8. Edit the **Regions** > **Wetted Wall Shell** > **Boundaries** > **Wetted Wall** node and set the following properties:

Node	Property	Setting
Physics Conditions > Thermal Specification	**Condition**	**Heat Flux**
Physics Values > Heat Flux	**Method**	**Field Function**
Field Function	**Scalar Function**	**Heat Flux Wetted Wall**

Set the boundary conditions for the inlet of the fluid film shell.

9. Edit the **Boundaries** > **Top edges** > **Physics Values** node and set the following properties:

Node	Property	Setting
Fluid Film Thickness > Constant	**Value**	`0.001 m`
Mass Flow Rate > Constant	**Value**	`0.015 kg/s`
Species Mass Fraction > Constant	**Value**	`[0.5, 0.5]`
Total Temperature > Constant	**Value**	`293.15 K`

10. Save the simulation.

Setting Up the Injector

Set up a Lagrangian injector to model the film surface from which droplets leave the film and are injected into the bulk flow.

A special type of injector is used to signify a film surface from which droplets leave the film and are injected into the bulk flow.

To set up the injector:

1. Right-click **Injectors** and select **New**.
2. Rename the **Injectors > Injector 1** node to **Binary Droplets**.
3. Set the injector *Type* to **Film Stripping**.
4. Set the *Lagrangian Phase* to **Droplets**.
5. Click (**Custom Editor**) for the *Inputs* property.
6. In the selection dialog, select **Regions > Fluid > Boundaries > Wetted Wall [In-place 1]** and click **OK**.

   The *Properties* window appears as follows:



7. Save the simulation.

Setting Up the Film Stripping Model

Define the film stripping model as a Film-Lagrangian phase interaction.

The edge film stripping model is set as a Film-Lagrangian phase interaction in STAR-CCM+.

To set up the film stripping model:

Create a phase interaction.

1. Right-click **Continua** > **Physics 1** > **Models** > **Multiphase Interaction** > **Phase Interactions** and select **New**.
2. Rename the newly created **Phase Interaction 1** node to **Film Mixture - Droplets**.

To set up the Film-Lagrangian phase interaction:

3. Right-click **Multiphase Interaction** > **Phase Interactions** > **Film Mixture - Droplets** > **Models** and choose **Select Models...**

In the *Film Mixture - Droplets Model Selection* dialog:

4. Select **Film-Lagrangian Phase Interaction** from the *Phase Interaction Topology* box.
5. Click **Close**.

Set the phases for the Film-Lagrangian interaction:

6. Select **Film Mixture - Droplets > Models > Film-Lagrangian Phase Interaction** and set the *Fluid Film Phase* to **Fluid Film** and the *Lagrangian Phase* to **Droplets**.

The **Component Mapping** node appears in the simulation tree to indicate that mapping is required between components of the film and droplet phases.



To set up the mapping between the phases:

7. Select **Multiphase Interaction > Phase Interactions > Film Mixture - Droplets > Models > Component Mapping** and click ... (**Custom Editor**) for the *Connectivity* property.
8. In the *Connectivity* dialog, map the fluid film components to their corresponding gas components, as shown below.



9. Click **OK**.

Select the edge stripping model and set its properties:

10. Right-click **Multiphase Interaction** > **Phase Interactions** > **Film Mixture - Droplets** > **Models** and choose **Select Models**.

In the *Film Mixture - Droplets Model Selection* dialog:

11. Select **Edge Stripping** from the *Optional Models* box.

    The **Multiphase Material** option is automatically selected.
12. Click **Close**.
13. Select the **Film Mixture - Droplets** > **Models** > **Edge Stripping** node.
14. To make sure that stripping occurs at the bend, set the *Minimum Corner Angle* to `10.0 deg`.



Having defined the **Film Mixture - Droplets** phase interaction, you can now set the boundary interaction mode for the fluid film boundary by editing the global Lagrangian phase boundary specification:

15. Select **Continua > Physics 1 > Models > Lagrangian Multiphase > Lagrangian Phases > Droplets > Boundary Conditions > Fluid-Film Boundary > Physics Conditions > Mode** and set the *Active Mode* to **Fluid Film**.



The **Fluid Film** mode is available only when there is a correct match between the fluid film continuum and the Lagrangian phase. This correct match occurs when components have been correctly mapped and the same models are used in both fluid film and Lagrangian phases. If you make any changes to the fluid film or the Lagrangian phase models that invalidate the **Film Mixture - Droplets** phase interaction, the *Active Mode* reverts to **Rebound**. The **Fluid Film** option is not available until you validate the phase interaction. You must set the *Active Mode* back to **Fluid Film**.

16. Save the simulation.

Setting Reference Values

Set the reference value for the physics continuum.

Define the reference values of the **Physics 1** continuum to make sure that acceleration due to gravity acts in the expected direction. In this tutorial, the gravity acts in the negative Z direction.

To set the reference value:

1. Edit the **Physics 1 > Reference Values** node and set the following property:

Node	Property	Setting
Gravity	**Value**	`[0.0, 0.0, -9.81] m/s^2`



2. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`0.01 s`

Node	Property	Setting
	Temporal Discretization	**2nd-order**

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`5 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`5`
Maximum Steps	**Enabled**	Deactivated

3. Save the simulation.

Setting Up Evaporation Rate Reports, Monitors, and Plots

Set up reports, monitors, and plots to monitor the film evaporation rate along the fluid film boundary.

To set up evaporation rate reports, monitors, and plots:

1. Right-click the **Reports** node and select **New Report > Surface Average**.
2. Rename the newly created **Surface Average 1** node to `H2O Evaporation Rate`.
3. Set the *Scalar Field Function* to **Film Evaporation Rate > Film Evaporation Rate of H2O**.
4. Click the ellipsis to the right of the *Parts* property.
5. In the dialog that appears, expand the **Regions > Fluid** node and select **Wetted Wall [In-place 1]**.



6. Click **OK**.

The *H2O Evaporation Rate - Properties* window looks as follows:

Copy this report to create a second one for the ethylene glycol evaporation rate:

7. Right-click **Reports > H2O Evaporation Rate** and select **Copy**.
8. Right-click **Reports** and select **Paste**.
9. Rename **Copy of H2O Evaporation Rate** to **C2H6O2 Evaporation Rate**.
10. Select **Reports > C2H6O2 Evaporation Rate** and, set the *Scalar Field Function* to **Film Evaporation Rate > Film Evaporation Rate of C2H6O2**.

Create a monitor and a plot for both reports.

11. Select both the **H2O Evaporation Rate** and **C2H6O2 Evaporation Rate** nodes, right-click either and select **Create Monitor and Plot from Report**.



12. Click **Single Plot** to display both evaporation rates on the same plot.

Set the monitors to update at each time step:

13. Select both the **Monitors > C2H6O2 Evaporation Rate Monitor** and the **H2O Evaporation Rate Monitor** nodes, and set *Trigger* to **Time Step**.



Format the monitor plot appearance:

14. Rename **Plots** > **Reports Plot** to **Evaporation Rates**.
15. Right-click the **Plots** > **Evaporation Rates** node and select **Open**.
16. Select **Plots** > **Evaporation Rates** > **Axes** > **X Axis** > **Title** and set the *Title* to `Time (s)`.
17. Select **Plots** > **Evaporation Rates** > **Axes** > **Y Axis** > **Title** and set the *Title* to `Surface Average of Film Evaporation Rate (kg/m^2-s)`.

Setting Up a Scalar Scene

Set up a scalar scene to display the fluid film thickness.

To set up a scalar scene:

Create a scalar scene for fluid film thickness.

1. Right-click the **Scenes** node and select **New Scene > Scalar**.

A new sub-node **Scalar Scene 1** is created under the **Scenes** node.

2. Rename the **Scenes > Scalar Scene 1** node to `Temperature`.
3. Click **scene/plot**.
4. Select the **Temperature > Displayers > Scalar 1 > Parts** node and then click the ellipsis to the right of the *Parts* property. In the dialog that appears, expand the **Regions > Wetted Wall shell** node and select **Wetted Wall**.
5. Expand the **Lagrangian Phases** node and select **Droplets** to visualize the edge film stripping.



6. Click **OK**.
7. Select the **Temperature > Displayers > Scalar 1 > Scalar Field** node and click to the right of the *Function* property and set it to **Temperature**.



8. Click (**Save-Restore-Select views**) in the toolbar and select **Restore View > View 1**.
9. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

2. Activate the temperature scene and observe the change in wall and droplets temperature as the solution progresses.



The example above shows the temperature in the early stages of the run. The simulation continues until 5 seconds of physical time is complete.

When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

3. When the simulation has finished running, save it.

Visualizing the Results

Examine the results of the simulation.

The temperature contour for the solution at the final time step is as shown below.



1. Activate the evaporation rates monitor plot.

   The evaporation rates plot for the converged solution is as shown below.



Summary

This tutorial demonstrated how to use the fluid film modeling capability of STAR-CCM+ to set up a multi-component film scenario with evaporation and edge stripping.

This tutorial has introduced the following STAR-CCM+ features:

- Defining models for fluid film model.
- Setting fluid film models.
- Setting Film-Eulerian and Film-Lagrangian phase interactions.
- Setting the fluid film evaporation and condensation model.
- Setting up shell region and boundary conditions.
- Setting up the edge stripping model.
- Running the solver for a set physical time.

Dispersed Multiphase: Airfoil Icing

When aircraft pass through air containing supercooled liquid water droplets, it is possible for the droplets to impact the aircraft surface and form layers of ice. Ice formation on an aircraft wing changes the aerodynamic properties of the wing, and leads to loss of lift, increase in drag, and a change in pressure distribution. STAR-CCM+ provides a technique for simulating the formation of ice on airfoils. This technique relies on the Dispersed Multiphase model to represent the liquid droplets in the air surrounding the airfoil.

The case simulated in this tutorial consists of a two-dimensional airfoil in an icing tunnel. The case is taken from a NASA study on ice formation and icing effects for modern airfoils.

In this simulation, a multiphase mixture of air and water droplets flows over an idealized two-dimensional airfoil, in a low-temperature, low-turbulence regime. The mixture flows with initial velocity of 129 m/s in the x-direction. The droplets are water particles with diameter of 2.1E-5 m. The droplets volume fraction, with respect to air, is of 3.42E-7.

The geometry that is used in the tutorial is displayed below.

In the DMP method, the background air affects the water droplets through drag and heat transfer. The water droplets, however, do not have any impact on the air. The liquid content of the droplet phase is transferred to the surface of the airfoil as a liquid film. For this reason, a shell region with fluid film properties is created on the surface of the airfoil. Droplets impact the shell region and form a liquid film.

Liquid film on the airfoil surface turns to ice. As more and more ice is formed, the shape of the airfoil must change. STAR-CCM+ converts the volume of ice in the film to a corresponding deformation in the wing surface. After deforming the wing surface using the mesh morpher, STAR-CCM+ removes the volume of ice from the film layer. In this way,

the shape of the shell region changes to reflect the buildup of ice on the airfoil surface.

Contents

Prerequisites
Importing the Volume Mesh
Selecting Physics Models
Specifying Initial Conditions
Creating a Shell Region
Selecting Fluid Film Models
Tracking the Ice Layer
Selecting Dispersed Phase Models
Defining Phase Interactions
Setting Up the Morpher
Defining Boundary Conditions
Setting Up Monitors
Preparing the Scenes
Setting Stopping Criteria
Setting the Solver Parameters
Running the Simulation
Visualizing the Results
Summary
DMP: Airfoil Icing Tutorial Bibliography

Prerequisites

The instructions in the Dispersed Multiphase - Airfoil Icing tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Volume Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the volume mesh for this tutorial.

1. Launch STAR-CCM+.
2. Start a new simulation.
3. Select **File > Import > Import Volume Mesh...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/multiphaseFlow/data`.
5. Select `dmpAirfoil.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window, and a mesh region is created under the **Regions** node.
6. Save the simulation as `dmpAirfoil.sim`.

   You can visualize the imported volume mesh with a method of your choice, such as creating a mesh scene. The mesh appears as shown below.



Selecting Physics Models

In the DMP approach, single-phase flow models are used to simulate the continuous gas phase, which in this tutorial is air. The continuous phase is not affected by the dispersed liquid phase, that is, the water droplets.

At this stage, select the single-phase models that represent the unsteady air flow as an ideal, turbulent gas, subject to gravity. Choose the Dispersed Multiphase model to solve for the water phase as droplets,

and the Fluid Film to solve for the water phase on the surface of the airfoil. You also require the Multiphase Interaction model, to define the influence of the gas phase on the droplet motion and energy, and the impingement of the droplets onto the water film phase.

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group box	Model
Enabled Models	**Three Dimensional** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Segregated Flow**
	Gradients (Selected automatically)
Equation of State	**Ideal Gas**
Energy	**Segregated Fluid Temperature**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Optional Models	**Fluid Film**
	Dispersed Multiphase
	Multiphase Interaction
	Gravity

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Specifying Initial Conditions

By default, gravity is set to act along the z-axis, whereas the geometry for this tutorial requires it to act along the y-axis. The simulation also requires initial values for fluid velocity and temperature.

To specify the initial conditions:

1. Edit the **Continua** > **Physics 1** node, and set the following properties:

Node	Property	Setting
Reference Values		

Node	Property	Setting
Gravity	*Value*	`[0,-9.81,0] m/s^2`
Initial Conditions		
Static Temperature > Constant	*Value*	`258.4 K`
Velocity > Constant	*Value*	`[129,0,0] m/s`

Creating a Shell Region

Define a shell region on the surface of the airfoil.

The DMP model permits mass within its liquid phase to transfer to a liquid film. In STAR-CCM+, liquid films exist only within a shell region. A shell region is a cell layer that is located on top of an existing boundary. Although the cell layer has zero-thickness in geometry, it contains a finite thickness relevant to the quantity being modeled. In the case of a liquid film, the thickness is the depth of the film layer.

In this tutorial, you define a shell region on the surface of the airfoil, so that the liquid film associated with the shell region can receive impinging droplets.

To create a shell region:

1. Right-click the **Regions** > **Region_1** > **Boundaries** > **Foil** node and select **Create Shell Region**.

   A shell region, **Foil shell**, is created within the **Regions** node. Also, an interface between the shell region and **Region 1** is created automatically within the **Interfaces** node.



2. Save the simulation.

Selecting Fluid Film Models

The set of models that solve for the liquid film phase is specified within the Fluid Film model node. Here, you first define a fluid film phase and select appropriate physics models. You then assign the film phase to the shell region on the surface of the airfoil.

To define a fluid film phase:

1. Expand the **Continua** > **Physics 1** > **Models** > **Fluid Film** node.
2. Right-click the **Fluid Film Phases** node and select **New**.

   A new node, **Phase 1**, is created within the **Fluid Film Phases** node.
3. Rename **Phase 1** to `Film Phase`.

Select the models for the film phase. Within this selection, activate the Melting-Solidification model that captures the solidification of water droplets into ice:

4. For the fluid film phase, **Film Phase**, select the following physics models:

Group box	Model
Fluid Film Segregated Energy	**Segregated Temperature***
Enabled Models	**Laminar** (selected automatically)
	Flow Model (selected automatically)
Space	**Shell Three Dimensional** (selected automatically)
Material	**Liquid**
Enabled Models	**Constant Density** (selected automatically)
Optional Models	**Melting-Solidification**

5. To review the models, expand the **Fluid Film Phases** > **Film Phase** > **Models** node.



As an initial condition for the fluid film phase, specify the temperature:

6. Expand the **Film Phase** > **Initial Conditions** node.
7. Select the **Static Temperature** > **Constant** node and set *Value* to `258.4 K`.

Tracking the Ice Layer

In an ice-accretion simulation, STAR-CCM+ calculates the thickness of solidified film (ice) at every time-step. At the start of the next time-step, the solid film thickness increment is passed to the mesh morpher so that it can deform the surface of the airfoil to match the film thickness. The solidified film mass is then removed from the simulation.

To allow the Melting-Solidification model to track the liquid-solid interface, you must set its Solidified Film Option to Remove Solidified Film.

To set the properties for removing the solidified film:

1. Expand the **Continua** > **Physics 1** > **Models** > **Fluid Film** > **Fluid Film Phases** > **Film Phase** > **Models** node.
2. Set the following properties:

Node	Property	Setting
Flow Model	*Maximum Film Thickness*	`1.0E-5 m`
Melting-Solidification	*Solidified Film Option*	**Remove Solidified Film**
	Solid Time-step Factor	`10`

The *Solid Time-step Factor* allows you to decrease the time that is required to run the simulation. A *Solid Time-step Factor* of 10, allows ice to form 10 times faster in the simulation.

The ice accumulation, visible as deformation of the wing shape, corresponds to an exposure time given by the product of the *Solid Time-step Factor* and the simulated physical time.

Finally, assign the fluid film phase models to the shell region, **Foil shell**:

3. Select the **Regions** > **Foil shell** node and set *Continuum* to **Film Phase**.
4. Save the simulation.

Selecting Dispersed Phase Models

A dispersed phase is defined within the Dispersed Multiphase model node as an Eulerian phase.

To create a dispersed phase:

1. Expand the **Continua** > **Physics 1** > **Models** > **Dispersed Multiphase** node.
2. Right-click the **Eulerian Phases** node, and select **New**.

   A new node, **Phase 1**, is added within the **Eulerian Phases** node.
3. Rename the **Phase 1** node to `DMP Phase`.

Select the physics models for the dispersed phase:

4. For the dispersed phase, **DMP Phase**, select the following physics models:

Group box	**Model**
Material	**Liquid**
Optional Models	**Flow Model**

5. To review the selected models, expand the **DMP Phase** > **Models** node.



For the dispersed phase, specify the droplets diameter and set their initial temperature and velocity to match the initial values that are set for the background fluid.

6. Expand the **DMP Phase** > **Models** > **Liquid** > **H2O** > **Material Properties** node.
7. Select the **Particle Diameter** > **Constant** node, and set *Value* to `2.1E-5 m`.
8. Edit the **DMP Phase** > **Initial Conditions** node, and set the following properties:

Node	Property	Setting
Static Temperature > **Constant**	*Value*	`258.4 K`
Velocity > **Constant**	*Value*	`[129, 0, 0] m/s`

9. Save the simulation.

Defining Phase Interactions

Interactions between phases are specified within the Multiphase Interaction Model node. In this case, the continuous gas phase interacts with the dispersed phase through drag force and heat transfer. The dispersed phase interacts with the fluid film phase through impingement.

To specify the interaction between the dispersed phase and the fluid film phase:

1. Expand the **Continua** > **Physics 1** > **Models** > **Multiphase Interaction** node.
2. Right-click the **Phase Interactions** node and select **New**.

   A new node, **Phase Interaction 1**, appears within the **Phase Interactions** node.
3. Rename **Phase Interaction 1** to `Film-DMP`.
4. For the phase interaction, **Film-DMP**, select the following physics models:

Group box	Model
Phase Interaction Topology	**Film-Dispersed Phase Interaction**
Optional Models	**Impingement**

5. Select the **Film-DMP > Models > Film-Dispersed Phase Interaction** node, and set the following properties:

Property	Setting
Fluid Film Phase	**Film Phase**
Dispersed Phase	**DMP Phase**

To specify the interaction between the dispersed phase and the continuous phase:

6. Repeat the steps to create a second phase interaction. Name this second interaction, `DMP-Continuum`.
7. For the phase interaction, **DMP-Continuum**, select the following physics models:

Group box	Model
Phase Interaction Topology	**Dispersed-Physics Continuum Interaction**
Optional Models	**Pressure Gradient Force**
	Interaction Length Scale
	Heat Transfer
	Drag Force

8. Select the **DMP-Continuum > Models > Dispersed-Physics Continuum Interaction** node, and set *Dispersed Phase* to **DMP Phase**.
9. Save the simulation.

Setting Up the Morpher

This case requires the morpher to apply mesh deformation for ice accretion. In STAR-CCM+, the morpher is defined as a motion.

To activate the morpher:

1. Right-click the **Tools > Motions** node and select **New > Morphing**.
2. Select the **Regions > Region_1 > Physics Values > Motion Specification** node, and set *Motion* to **Morphing**.

Assign the morphing motion to the airfoil surface:

3. Expand the **Regions > Region_1 > Boundaries > Foil [In-place 1] > Physics Conditions** node.
4. Select the **Morpher** node and set *Method* to **Internal**.

When the morpher method is set to Internal, the displacement fed to the morpher is computed by STAR-CCM+ internal models. This method is used by the Melting-Solidification model, that provides the solid film thickness increment to the morpher.

Motion Specification for the **Foil shell** region is automatically set to follow **Region_1**.

Defining Boundary Conditions

Specify the temperature, the volume fraction of the dispersed phase, and the Mach number at the inlet and outlet.

To define the boundary conditions:

1. Expand the **Regions** > **Region_1** > **Boundaries** > **Inlet** node.
2. Edit the **Physics Values** node, and set the following properties:

Node	Property	Setting
Mach Number > **Constant**	*Value*	`0.4`
Static Temperature > **Constant**	*Value*	`258.4 K`

3. Edit the **Phase Conditions** > **DMP Phase** > **Physics Values** node, and set the following properties:

Node	Property	Setting
Static Temperature > **Constant**	*Value*	`258.4 K`
Volume Fraction > **Constant**	*Value*	`3.42E-7`

4. Repeat the previous steps to set the boundary conditions for the **Outlet**, using the same settings that you entered for the **Inlet**.

Specify boundary conditions at the interface:

5. Edit the **Regions** > **Region_1** > **Boundaries** > **Foil [In-place 1]** node, and set the following properties:

Node	Property	Setting
Physics Conditions		
Wall Surface Specification	*Method*	**Rough**

Node	Property	Setting
Physics Values		
Roughness Height > Constant	*Value*	`0.1 mm`

6. Save the simulation.

Setting Up Monitors

Monitor key quantities like the film thickness, the maximum film thickness, and the film surface area, as functions of time.

To set up the reports and monitors:

1. Create a **Sum** report with the following settings:

Property	Setting
Scalar Field Function	**Area: Magnitude**
Parts	**Foil shell > Boundaries > Foil**

2. Rename the **Sum 1** node to `Film Surface`.
3. Create a **Surface integral** report with the following settings:

Property	Setting
Scalar Field Function	**Shell Thickness**
Parts	**Foil shell > Boundaries > Foil**

4. Rename the **Surface integral 1** node to `Film Volume`.
5. Create a **Maximum** report with the following settings:

Property	Setting
Scalar Field Function	**Shell Thickness**
Parts	**Foil shell > Boundaries > Foil**

6. Rename the **Maximum 1** node to `Maximum Film Thickness`.
7. Right-click the **Film Surface** node and select **Create Monitor and Plot from Report**.
8. Repeat this step for all the reports that you created.

Preparing the Scenes

Here you prepare scalar scenes to visualize the impingement rate, the volume fraction of the dispersed phase, and the total solid film thickness

increment. You also create a vector scene to visualize the morpher displacement.

To visualize the impingement rate:

1. Create a scalar scene and rename it to `Impingement Rate`.
2. Expand the **Scenes** > **Impingement Rate** > **Displayers** node.
3. Edit the **Outline 1** > **Parts** node, and remove **Region_1** from the selected Parts.
4. For the node, **Scalar 1**, set the following properties:

Node	Property	Setting
Parts	*Parts*	**Region_1 > Boundaries > Foil [In-place 1]**
Scalar Field	*Function*	**Film Impingement Rate > Film-Dispersed Impingement Rate of Film-DMP**

To visualize the volume fraction of the dispersed phase:

5. Create a scalar scene and rename it to `DMP Volume Fraction`.
6. Edit the **Scenes** > **DMP Volume Fraction** > **Displayers** node, and set the following properties:

Node	Property	Setting
Scalar 1 > Parts	*Parts*	**Region_1**
Scalar 1 > Scalar Field	*Function*	**Volume Fraction > Volume Fraction of DMP Phase**

The solid film thickness increment is computed at every time step. To visualize the total solid film thickness increment, accumulated during the total simulation time, you need to create a Field Sum monitor for the **Solid Film Thickness Increment** scalar function.

7. Within the **Monitors** node, create a new **Field Sum** monitor with the following settings:

Property	Setting
Field Function	**Solid Film Thickness Increment**

Property	Setting
Parts	**Foil shell > Boundaries > Interface to Region_1 [In-place 1]**

8. Rename the Field Sum 1 monitor to `Total Solid Film Thickness Increment`.
9. Create a scalar scene and rename it to `Total Solid Film Thickness Increment`.
10. Expand the **Scenes > Total Solid Film Thickness Increment > Displayers** node.
11. Edit the **Outline 1 > Parts** node, and remove **Region_1** from the selected Parts.
12. For the node, **Scalar 1**, set the following properties:

Node	Property	Setting
Parts	*Parts*	**Foil shell > Boundaries > Interface to Region_1 [In-place 1]**
Scalar Field	*Function*	**Sum of Solid Film Thickness Increment**

To visualize the morpher displacement:

13. Create a vector scene and rename it to `Morpher Displacement`.
14. Expand the **Scenes > Morpher Displacement > Displayers** node.
15. Select the **Outline 1** node and activate *Surface*.
16. Edit the **Outline 1 > Parts** node, deselect all Parts, and then select **Region_1 > Boundaries > Foil [In-place 1]**.
17. For the node, **Vector 1**, set the following properties:

Node	Property	Setting
Parts	*Parts*	**Region_1 > Boundaries > Foil [In-place 1]**
Vector Field	*Function*	**Morpher Displacement**

18. Save the simulation.

Setting Stopping Criteria

Specify the maximum number of iterations and the maximum physical time.

To set the stopping criteria for this simulation:

1. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`20`
Maximum Physical Time	*Maximum Physical Time*	`34.2 s`

Setting the Solver Parameters

Set appropriate parameters for the Dispersed Phase Flow solver and the Implicit Unsteady solver.

To set the solver parameters:

1. Edit the **Solvers > Dispersed Multiphase > Dispersed Phase Flow** node.
2. Set the following properties:

Node	Property	Setting
Velocity	*Under-Relaxation Factor*	`0.2`
Volume Fraction	*Maximum Volume Fraction*	`1.0E-4`
	Minimum Volume Fraction	`1.0E-8`

3. Select the **Solvers > Implicit Unsteady** node and set *Time-Step* to `0.1 s`.
4. Save the simulation.

Running the Simulation

The simulation is now fully set up and ready to be run.

To run the simulation:

1. Click (**Initialize Solution**).
2. Click (**Run**).
3. When the simulation has finished running, save the simulation.

Visualizing the Results

Examine the scenes that you prepared earlier.

The Impingement Rate scene shows the rate at which water droplets hit the surface of the airfoil:



The amount by which the morpher displaces the airfoil surface in each time step is shown in the Morpher Displacement scene:



The concentration of water droplets in air is shown in the DMP Volume Fraction scene:



The total amount of ice accumulated on the airfoil surface is shown in the scene for Total Solid Film Thickness Increment:



Summary

This tutorial demonstrated how to simulate ice formation on two-dimensional airfoils in STAR-CCM+. In particular, the simulation focuses on the ice build-up process, showing the effects on the shape of the wing. Since changes in the wing shape affect the lift and drag forces, it is straightforward to investigate the effects of the accumulated ice on aerodynamic performance.

This tutorial has introduced the following STAR-CCM+ features:

- Selecting the continuum physics models in the DMP approach
- Setting up a fluid shell region as particles impact boundary surface
- Using the DMP model to define a dispersed phase
- Defining interactions between phases
- Using the Melting-Solidification model in combination with the morpher

DMP: Airfoil Icing Tutorial Bibliography

Addy Jr, H. E. "Ice Accretions and Icing Effects for Modern Airfoils", NASA/TP – 2000-210031

Discrete Element Method

The tutorials in this set illustrate various STAR-CCM+ features for simulating Discrete Element Method problems

Contents

DEM Particles in a Conveyor
DEM Particle Settling

DEM Particles in a Conveyor

The purpose of this tutorial is to provide a guide to setting up a simulation using the Discrete Element Method.

The simulation consists of an injection chute, conveyor belt, and collection hopper. Grains of wheat are injected from points in the injection chute for 0.5 seconds. The conveyor then carries the grain up a slope, where it falls into the collection hopper. A visual representation of the transient solution is created, which shows the grain moving through the fluid region. The color of each piece of grain represents the velocity of that particle.



Contents

Prerequisites
Importing the Surface Mesh and Naming the Simulation
Visualizing the Imported Geometry and Surface Mesh
Generating the Volume Mesh
Selecting Physics Models
Defining Lagrangian Phases

Defining the DEM Particle Interaction
Setting Reference Values
Setting the Moving Conveyor Wall Condition
Creating an Injector
Setting the Solver Parameters and Stopping Criteria
Setting Up a Visualization Scene
Running the Simulation
Visualizing Results
Summary

Prerequisites

The instructions in this tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

To complete this tutorial, familiarize yourself with the following:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Surface Mesh and Naming the Simulation

Create a new simulation and import the mesh from `grainConveyor.x_b`.

Navigate to the `doc/startutorialsdata/discreteElementMethod/data` subdirectory of the STAR-CCM+ installation directory. Copy the `grainConveyor.x_b` file to the current working directory.

Start upSTAR-CCM+ in the usual manner and select the New Simulation option from the menu bar.

Continue by importing the fluid surface mesh and saving the simulation.

1. Select **File > Import > Import Surface Mesh** from the menu bar. In the *Open* dialog, navigate to the working directory and select the file `grainConveyor.x_b`. Click **Open**.
2. In the *Import Surface Options* dialog:
   a) Make sure that the **Create New Part** option is selected.

b) Leave the remaining settings as default.



c) Click **OK**.

STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.

3. Save the new simulation to disk under file name `tutorialDEM.sim`.

Visualizing the Imported Geometry and Surface Mesh

Examine the *Geometry Scene 1* display in the *Graphics* window.

Use the following steps:

1. Use the mouse buttons to rotate, zoom in, and position the part to obtain a similar view as shown below.



2. Save this view by clicking (**Save-Restore-Select views**) on the toolbar and choosing **Store Current View**.

Generating the Volume Mesh

The first step in generating a volume mesh is to create a fluid region from the imported geometry part.

Use the following steps:

1. Right-click the **Geometry** > **Parts** > **GrainSystem** node, and select **Assign Parts to Regions...**
2. In the *Assign Parts to Regions* dialog, make sure that the **GrainSystem** part is selected, and set the second dropdown to **Create a Boundary for Each Part Surface**.



3. Click **Apply**, and then **Close**.

A fluid boundary is created for each part surface, as shown below:



To create a meshing continuum with a base size of 0.1m:

4. Right-click the **Continua** manager and select **New** > **Mesh Continuum**.
5. Right-click the **Continua** > **Mesh 1** > **Models** node and select **Select Meshing Models**

In the *Meshing 1 Model Selection* dialog:

6. Select **Surface Remesher**
7. Select **Polyhedral Mesher**.
8. Click **Close**.
9. Expand the **Reference Values** manager, and select the **Base Size** node and set the value to `0.1 m`.



The volume mesh can now be generated.

10. Click (**Generate Volume Mesh**) in the **Mesh** toolbar to start the mesh generation process.

The output from the meshers is shown in the *Output* window. When the meshing process is complete, the message `Volume Meshing Pipeline Completed` is displayed in the *Output* window.

11. Save the simulation.

Selecting Physics Models

Physics models are required to define the behavior of the multiphase continuum.

To set the physics models:

1. For the physics continuum **Continua** > **Physics 1**, select the following models, in order:

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically)
	Gradients (Selected automatically)
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Ideal Gas**
	Coupled Energy (Selected automatically)
Viscous Regime	**Laminar**
Optional Models	**Gravity**
	Lagrangian Multiphase
	Discrete Element Model (DEM)
	Multiphase Interaction (Selected automatically)

2. Click **Close**.
3. Open the **Physics 1** > **Models** manager.

The selected models appear as shown below:



4. Save the simulation.

Defining Lagrangian Phases

A Lagrangian phase is used to define the solid particles in the simulation. The phase requires specific physics models and properties to characterize its behavior. In this case, the solid particles are grains of wheat, with the material properties defined manually.

To define the Lagrangian phase:

1. Right-click the **Lagrangian Multiphase** > **Lagrangian Phases** manager and select **New**.
2. Rename the phase as `Grain`.
3. Right-click the **Grain** > **Models** manager and select **Select models**.
4. In the *Grain Model Selection* dialog, select the following models, in order:

Group Box	Model
Particle Type	**DEM Particles**

Group Box	Model
Optional Particle Forces	**Pressure Gradient Force** (selected automatically)
Material	**Solid**
Equation of State	**Constant Density**
Particle Shape	**Composite Particles**

5. Click **Close**.

To define the material properties for the Grain particles:

6. Rename the **Rubber** node to `Wheat`.
7. Select the **Material Properties** > **Density** > **Constant** node and set the *Value* to `700 kg/m^3`.



8. Set the remaining properties as follows:

Property	Setting
Poisson Coefficient	`0.5`
Young's Modulus	`0.2 MPa`

To define the physical shape of the grain particles:

9. Right-click the **Model** > **Composite Particles** > **Shapes** node and select **New**.

10. Select the **Sphere 1** node and set the *Radius* to `0.8`.



11. Create two more shapes with the following properties:

Node	Property	Setting
Sphere 2	*Position*	`[1.0, 0.0, 0.0]`
	Radius	`1.0`
Sphere 3	*Position*	`[2.0, 0.0, 0.0]`
	Radius	`0.8`

12. Right-click the **Composite Particles** node and select **Show model**.

A representation of the grain particle is shown in the *Graphics* window.



13. Close the *Composite particle Grain* scene.

To define the physics conditions on the particle boundary:

14. Right-click the **Grain > Boundary Conditions > Wall > Physics Conditions > Boundary Properties > Solid > Al** node and select **Replace with**.



15. In the *Replace Material* dialog, expand the **Material Databases > Standard > Solids** manager, select **Rubber (Arlon Thermabond 99180T010)** and click **OK**.



16. Select the **Physics Conditions** > **Mode** node and check that the *Active Mode* property is set to **DEM mode**.



17. Save the simulation.

Defining the DEM Particle Interaction

The DEM particle interaction is used to define how the particles behave when they collide with one another and the region walls.

To define the grain to grain interaction:

1. Right-click the **Physics 1** > **Models** > **Multiphase Interaction** > **Phase Interactions** manager and select **New**.
2. Right-click the **Phase Interaction 1** > **Models** node and select the following models, in order:

Group Box	Model
Phase Interaction Topology	**DEM Phase Interaction**
Optional Models	**Hertz Mindlin**

3. Click **Close**.
4. Select the **Models** > **DEM Phase Interaction** manager and set both the *First phase* and *Second phase* properties to **Grain**.



5. Select the **Models** > **Rolling Resistance** > **Rolling Resistance Methods** > **Force Proportional** node and set the *Coefficient of rolling resistance* to `0.3`. The remaining properties can be left as default.



To define the Grain to Wall interaction, follow the steps above and create a particle interaction model, **Phase Interaction 2**, with the following settings:

6. Right-click the **Phase Interaction 2** > **Models** node and select the following models, in order:

Group Box	Model
Phase Interaction Topology	**DEM Phase Interaction**
Optional Models	**Hertz Mindlin**

7. Set the following properties:

Node	Property	Setting
DEM Phase Interaction	*First Phase*	**Grain**
	Second Phase	**Wall**
Hertz Mindlin	*Static Friction Coefficient*	0.8
	Normal Restitution Coefficient	0.5
	Tangential Restitution Coefficient	0.5

Node	Property	Setting
Rolling Resistance > Rolling Resistance Methods > Force Proportional	*Coefficient of rolling resistance*	`0.3`

This part of the simulation tree along with the *Hertz Mindlin - Properties* window appears as shown below:



8. Save the simulation.

Setting Reference Values

Define the reference values to make sure that the acceleration due to gravity acts in the expected direction.

To set the gravity reference value:

- Select the **Physics 1 > Reference Values > Gravity** node and set the *Value* to `[0.0, -9.81, 0.0] m/s^2`.



Setting the Moving Conveyor Wall Condition

Define the conveyor boundary with a translational velocity to represent the movement of the belt.

To set the moving wall condition:

1. Expand the **Regions** > **Region 1** > **Boundaries** > **GrainSystem.Conveyor** > **Physics Conditions** manager, and select the **Tangential Velocity Specification** node and set the *Method* to **Vector**.



2. Expand the **Physics Values** > **Relative Velocity** node and select the **Constant** node and set the *Value* to `[4.0, 0.0, 0.0] m/s`.



3. Save the simulation.

Creating an Injector

Create the local Cartesian coordinate system.

Use the following steps:

1. Right-click the **Tools > Coordinate System > Laboratory > Local Coordinate System** node and select **New > Cartesian**.



2. Using the in-place dialog, define the coordinate system with the following settings:

Property	**Setting**
Origin	`X: -0.5 m`
	`Y: 0.16 m`
	`Z: 0.0 m`
i Direction	`X: 1`
	`Y: 0`
	`Z: 0`
j Direction	`X: 0`
	`Y: 1`
	`Z: 0`

The dialog appears as follows.



3. Click **Create** then **Close**.

To create the presentation grid using the *Properties* window:

4. Right-click the **Derived Parts** manager and select **New Part > Probe > Presentation Grid**.



5. Using the in-place dialog, click **Create** and **Close**.

An arbitrary plane is created with the default settings.

To modify the settings:

6. Rename the **presentation grid** node to `Injector Grid`.
7. Use the following settings to define the injector grid:

Property	**Setting**
X Resolution	`10`
Y Resolution	`10`
Coordinate System	`Laboratory->Cartesian 1`
Point 1	`[0.2, -0.05, 0.0] m,m,m`
Point 2	`[0.0, -0.05, 0.2] m,m,m`
Origin	`[0.0, -0.05, 0.0] m,m,m`



Finally, the particle injector is created.

8. Right-click the **Injectors** node and select **New**.
9. Select the **Injector 1** node and set the *Type* to **Part Injector**.
10. To assign the injector to the previously created injector grid and click the ellipsis next to the *Inputs* property.



11. In the *Inputs* selection dialog, expand the **Derived Parts** node and tick the checkbox next to the **Injector Grid** node.



12. Click **OK** to close the dialog.
13. While still at the **Injector 1** node, set the *Lagrangian Phase* to **Grain**.



14. Expand the **Injector 1** node and set the following properties:

Node	Property	Setting
Conditions		
Flow Rate Distribution	**Method**	**Per Injector Point**
Flow Rate Specification	**Method**	**Particle Flow Rate**
Values		
Particle Diameter > Constant	**Value**	`0.01 m`

Node	Property	Setting
Particle Flow Rate > Constant	**Value**	`$Time < 0.5 ? 100 : 0`
Velocity > Constant	**Value**	`[0.0, -2.0, 0.0] m/s`

15. Save the simulation.

Setting the Solver Parameters and Stopping Criteria

As the simulation is transient, set the solver time-step and elapsed time.

To define the solver parameters:

1. To define the solver parameters, expand the **Solvers** node and set the following properties:

Node	Property	Value
Implicit Unsteady	**Time-Step**	`0.01`
DEM Solver	**Time Scale**	`0.1`

2. To define the stopping criteria, expand the **Stopping Criteria** node and set the following properties:

Property	Value
Maximum Inner Iterations	`1`
Maximum Physical Time	`2.5 s`
Maximum Steps > Enabled	off

3. To define the stopping criteria, select the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`1`
Maximum Physical Time	*Maximum Physical Time*	`2.5 s`
Maximum Steps	*Enabled*	Deactivated

4. Save the simulation.

Setting Up a Visualization Scene

Create a scalar scene to provide a visual representation of the grain particles within the fluid system.

Use the following steps:

1. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.
2. Click the **scene/plot** button above the object tree.
3. Select the **Outline 1** node.
   - Deactivate the *Outline* property.
   - Activate the *Surface* property.



4. Right-click the **Scalar 1** > **Parts** manager and select **Edit**.
5. In the *Edit* dialog, select **Derived Parts** > **Injector Grid** and **Lagrangian Phases** > **Grain**. Click **OK**.



6. Select the **Scalar Field** node and set the *Function* to **Particle > Particle Velocity > Magnitude**.



7. Select the **Scalar 1** node and set the *Point Scale Mode* to **Scalar**.



8. Select the **Point Size Scale Field** node and set the *Function* to **Particle > Particle Size**.



9. Select the **Attribute > Clip Plane > Plane 1** node and set the following properties:

Property	**Setting**
Normal	`[0.0, 0.0, -1.0]`
Enabled	`Ticked`
Origin	`[0.0, 0.0, 0.19] m, m, m`



10. Rearrange the color bar and display to provide the following view.



To export an image of the visualization at every time-step:

11. Select the **Attributes > Update** node and set the *Trigger* to **Time Step**.
12. Activate the checkbox for the *Save to File* property.



13. Expand the **Update** manager, select the **Update > Time-Step Frequency** node, and set the *Time-step Frequency* to 5.



14. Save the simulation.

Running the Simulation

The simulation is now ready to be run.

Use the following steps:

1. Run the simulation by clicking (**Run**) in the **Solution** toolbar.

The iterative progress can be monitored in the *Output* window. The scalar scene displays the injection and movement of the grain particles within the fluid region as the solution proceeds.

When the solution is complete, the following message is displayed in the *Output* window: `Stopping criterion Maximum Physical Time satisfied.`

2. Save the simulation.

Visualizing Results

When the solver is complete, the *Scalar Scene 1* in the *Visualization* window appears.

See below:



The scene shows the grain particles in the collection area, after being transported up the conveyor. A series of images have been created in the current working directory named **Scene_2imagexxxxx.png**, where **xxxxx** refers to the time-step number. These images show the particle movement through the region during the transient solution, and can be used to create an animation.

Summary

This tutorial provides a guide to setting up a simulation using the Discrete Element Method.

The tutorial has introduced the following STAR-CCM+ features:

- Selecting fluid physics models.
- Definition of Lagrangian phases.
- Definition of DEM particle interaction.
- Creation of a particle injector.
- Setting up visualization scenes.

DEM Particle Settling

This tutorial simulates a three-dimensional solid flow through a model blast furnace to predict the formation of the deadman region at the base of the furnace.

Based on the geometry and raceway size from Zhou *et al.*, a model of an iron making blast furnace is constructed as shown below.



The particles are consumed near the raceway in an actual furnace. In this case, particle loss is instead simulated by introducing a particle outlet at the raceway location. Changes in particle size and property are not taken into account.

The main aim of this tutorial is to present a comparison of the predicted deadman size due to solid descent (using particle residence time) with the laboratory scale experiments of Zhou *et al*.

The three stages of the simulation are:

1. Particle bed formation: In the first stage, the simulation is run for 11.5 s with the inlet open to allow the furnace to be filled with particles (glass beads) to create the initial bed.
2. Particle settlement: In this stage, the simulation is run for a further 2 s with the inlet closed to allow the bed to settle.

   For the first two stages the outlets are closed to allow the build-up and settlement of the glass beads in the furnace.
3. Particle discharge: In the final stage, the simulation is run for a further 36.5 s with the inlet and outlets (one on each side) open.

Particle properties in this simulation are as follows:

- Spherical, 10 mm diameter particles of density 2500 kg/m^3.
- Sliding friction coefficient of 0.4 for both particle-particle and wall-particle interactions.
- Particle mass flow rate of 18.33 kg/s.

Contents

Importing the Surface Mesh and Naming the Simulation
Visualizing the Imported Geometry
Generating the Volume Mesh
Selecting Physics Models
Defining Lagrangian Phases
Defining the DEM Particle Interaction
Setting Reference Values
Setting the Boundary Conditions
Creating an Injector
Setting the Solver Parameters and Stopping Criteria
Setting Up a Visualization Scene
Running the Simulation
Visualizing Results
Stage Two of the Simulation
Stage Three of the Simulation
Comparing the Results
Summary
Blast Furnace Tutorial Bibliography

Importing the Surface Mesh and Naming the Simulation

Start a new simulation and import `blastFurnaceSurfaceMesh.dbs`.

Use the following steps:

1. Navigate to the `doc/startutorialsdata/discreteElementMethod/data` subdirectory of the STAR-CCM+ installation directory. Copy the `blastFurnaceSurfaceMesh.dbs` file to the current working directory.
2. Start up STAR-CCM+ in the usual manner and select the New Simulation option from the menu bar.

A surface mesh for the blast furnace has been prepared for this analysis. Continue by importing the surface mesh and saving the simulation.

3. From the menu bar, select **File > Import > Import Surface Mesh...**
4. In the *Open* dialog, navigate to the working directory and select the file `blastFurnaceSurfaceMesh.dbs`.
5. To start the import, click **Open**.
6. In the *Import Surface Options* dialog:
   - Select the **Create new Region** option.
   - Leave the remaining settings as default.
7. Click **OK**.



Note: The **One boundary per cell type** option results in a different boundary being created for each cell type, property ID or solid depending on the input file type.

STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window. A geometry scene is created in the *Graphics* window.

8. Save the new simulation to disk under file name `blastFurnace.sim`.

Visualizing the Imported Geometry

Examine the *Geometry Scene 1* display in the *Graphics* window.

- Rotate, zoom, and position the part to obtain a similar view as shown below.



Generating the Volume Mesh

Use the surface remesher and polyhedral volume mesher to create a mesh with a base size of 0.01 m, as the fluid solution is not of particular importance for this tutorial.

Use the following steps:

1. Right-click the **Continua** > **Mesh 1** > **Models** node and choose **Select Meshing Models...**

In the *Meshing Model Selection* dialog:

2. Select the **Surface Remesher** from the *Surface Mesh* group box.
3. Select **Polyhedral Mesher** from the *Volume Mesh* group box.

The dialog appears as follows when complete.



4. Click **Close**.
5. Expand the **Reference Values** node, and select the **Base Size** node and set the value to `0.01 m`.



Generate the volume mesh:

6. To start the mesh generation process, click (**Generate Volume Mesh**) in the **Mesh** toolbar.

The output from the mesher is shown in the *Output* window. When the meshing process is complete, the message `Volume Meshing Pipeline Completed` is displayed in the *Output* window and the properties of the generated mesh are given.

The properties of the fluid volume mesh can be seen by selecting the **Representations > Volume Mesh > Finite Volume Regions > Region 1** node.

7. Save the simulation.

Selecting Physics Models

Physics models are required to define the behavior of the multiphase continua.

To set the physics models:

1. Expand the **Continua** > **Physics 1** node and select the following models, in order:

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically)
	Gradients (Selected automatically)
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Coupled Flow**
Equation of State	**Constant Density**
Viscous Regime	**Inviscid**
Optional Models	**Gravity**
	Lagrangian Multiphase
	Discrete Element Model (DEM)
	Multiphase Interaction

2. Click **Close**.

The color of the **Physics 1** node turns from gray to blue to indicate that fluid models have been activated and the continua is valid.

3. Open the **Physics 1** > **Models** node.

The selected models appear as shown below:



4. Save the simulation.

Defining Lagrangian Phases

A Lagrangian phase is used to define the solid particles in the simulation. The phase requires specific physics models and properties to characterize its behavior. In this case, the solid particles are glass beads, with the material properties defined manually.

To define the Lagrangian phase:

1. Right-click the **Lagrangian Multiphase** > **Lagrangian Phases** node and select **New**.
2. Rename the phase to `ORE`
3. Right-click the **ORE** > **Models** node and choose **Select models**



4. In the *ORE Model Selection* dialog, select the following models, in order:

Group Box	Model
Particle Type	**DEM Particles**
Particle Shape	**Spherical Particles**
Material	**Solid**
Equation of State	**Constant Density**
Optional Models	**Residence Time**

5. Click **Close**.

Now you can set an appropriate material for the Ore.

6. Right-click the **Rubber** node and select **Replace with...**



7. In the *Replace Material* dialog, expand the **Material Databases** > **Standard** > **Solids** node, select **Glass (Glass Plate)**, and click **OK**.



Now define the material properties for the glass beads:

8. Select the **Material Properties** > **Density** > **Constant** node and make sure that the *Value* property is set to `2500 kg/m^3`.



9. Select the **Material Properties** > **Poisson Coefficient** > **Constant** node and set *Value* to `0.3`
10. Select the **Material Properties** > **Young's Modulus** > **Constant** node and set *Value* to `10 MPa`

Define the physics conditions on the particle boundary using the surface material properties and interaction mode.

11. Select the **ORE Boundary Conditions** > **Wall** > **Physics Conditions** > **Boundary properties** > **Solid** > **Al** > **Material Properties Young's Modulus** > **Constant** node and set *Value* to `10 MPa`.



12. Select the **ORE Boundary Conditions** > **Wall** > **Physics Conditions** > **Mode** node and set *Active Mode* to **DEM mode**.



13. Save the simulation.

Defining the DEM Particle Interaction

The DEM particle interaction is used to define how the particles behave when they collide with one another and wall regions. The procedure to define the ORE to ORE interaction has been provided in detail, with a table summarizing the steps for ORE to wall interaction.

To define the ORE to ORE interaction:

1. Right-click the **Physics 1** > **Models** > **Multiphase Interaction** > **Phase Interactions** node and select **New**.
2. Right-click the **Phase Interaction 1** > **Models** node and select the following models, in order:

Group Box	Model
Phase Interaction Topology	**DEM Phase Interaction**
Optional Models	**Hertz Mindlin**



3. Right-click the **Phase Interaction 1 > Models** node and edit the node.
4. Click **Expand/Contract Tree** in the *Edit* dialog, and set the following nodes:

Node	Value
First phase	**ORE**
Second phase	**ORE**
Static Friction Coefficient	`0.4`
Normal Restitution Coefficient	`0.8`
Tangential Restitution Coefficient	`0.8`
Coefficient of rolling resistance	`0.0010`

The *Edit* dialog looks as shown below:



5. Rename the **Phase Interaction 1** node to `ORE-ORE`
6. Similarly, follow the above steps and create another phase interaction for the ORE to Wall interaction.

Group Box	Model
Phase Interaction Topology	**DEM Phase Interaction**
Optional Models	**Hertz Mindlin**

Node	Value
First phase	**ORE**
Second phase	**Wall**
Static Friction Coefficient	`0.4`
Normal Restitution Coefficient	`0.8`

Node	Value
Tangential Restitution Coefficient	`0.8`
Coefficient of rolling resistance	`0.0010`

The *Edit* dialog looks as shown below



7. Rename the **Phase Interaction 2** node to `ORE-Wall`

Setting Reference Values

The reference values must be defined to make sure that the acceleration due to gravity acts in the expected direction.

To set the gravity reference value:

1. Select the **Continua** > **Physics 1** > **Reference Values** > **Gravity** node and set *Value* to `[0.0, -9.81, 0.0] m/s^2`.



Setting the Boundary Conditions

By default, the conditions for boundaries representing the furnace are all set to no-slip walls. You now change the **pIn** boundary to **Flow-Split Outlet** with default **Split Ratio** of `1` to allow filling of the geometry with particles for the first 11.5 seconds of the simulation.

Use the following steps:

1. Navigate to the **Regions** > **Region 1** > **Boundaries** node.
2. Edit **pIn** and set the following properties:

Node	Property	Setting
pIn	**Type**	**Flow-Split Outlet**
Physics Values	**Split Ratio**	`1.0`

3. Save the simulation.

Creating an Injector

Particles are injected into the fluid region from the top of the blast furnace. The particles, 0.01 m in diameter, are injected at a mass flow rate of 18.33 kg/s with a velocity of 0.01 m/s in the negative Y-direction.

Create the particle injector:

1. Right-click the **Injectors** node and select **New**.
2. Select the **Injector 1** node and set *Type* to **Surface Injector**.



Now assign the injector to the **pIn** region of the blast furnace geometry and click the ellipsis next to the *Inputs* property.



3. In the *Injector 1* dialog, select **pIn** and click **OK**.



4. While still, set the *Lagrangian Phase* to **ORE**.



5. Expand the **Injector 1 > Values** node, and select the **Mass Flow Rate > Constant** node and set *Value* to `18.33 kg/s`.



6. Set the remaining properties as follows:

Property	Setting
Particle Diameter	0.01 m
Velocity	[0, -0.01, 0] m/s

7. Save the simulation.

Setting the Solver Parameters and Stopping Criteria

As the simulation is transient, the solver time-step and actual run time must be set. In this case, the time-step is 0.01 seconds, with a physical time of 11.5 seconds for the first stage of the simulation. Also as no fluid flow effect is taken into account, freeze the coupled implicit solver.

To define the solver parameters:

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	*Time-Step*	`0.01 s`
Coupled Implicit	*Solver Frozen*	Activated

To define the stopping criteria:

2. Expand the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`1`
Maximum Physical Time	*Maximum Physical Time*	`11.5 s`
Maximum Steps	*Enabled*	Deactivated

3. Save the simulation.

Setting Up a Visualization Scene

A scalar scene is created to provide a visual representation of the glass beads within the fluid system.

Use the following steps:

1. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.
2. Click the **scene/plot** button above the object tree.
3. Select the **Outline 1** node, deactivate *Outline*, and activate *Surface*.



4. Edit the **Scalar 1** > **Parts** node.
5. In the *Edit* dialog, expand the **Lagrangian Phases** node and select **ORE** and click **OK**.



6. Select the **Scalar Field** node and set *Function* to **Particle > Particle Residence Time**.



7. Select the **Scalar 1** node and set the *Point Scale Mode* to **Scalar**.



8. Select the **Point Size Scale Field** node and set *Function* to **Particle > Particle Size**.



9. Select the **Attribute** > **Clip Plane** > **Plane 6** node and set the following properties:
   - Enabled: `Activated`
   - Origin: `[0.0, 0.0, 0.02]` m, m, m



10. Rearrange the color bar and display to provide the following view.



Running the Simulation

The simulation is now ready to be run.

Use the following steps:

1. Run the simulation by clicking (**Run**) in the **Solution** toolbar.

The iterative progress can be monitored in the *Output* window. The scalar scene displays the injection and movement of the glass beads within the fluid region as the solution proceeds.

When the solution is complete, the following message is displayed in the *Output* window: `Stopping criterion Maximum Physical Time satisfied.`

2. Save the simulation.

Visualizing Results

When the solver is complete, the *Scalar Scene 1* in the *Graphics* window appears.

See below.



Stage Two of the Simulation

Having filled the furnace with beads in stage 1, it is now necessary to allow a period for the beads to settle before opening the outlets and allowing the beads to escape.

Change the **pIn** boundary to **Wall** to prevent the flow of particles into the furnace and to allow the glass beads to settle and come to rest within the furnace.

1. Expand the **Regions** > **Region 1** > **Boundaries** node, and select the **pIn** node and set *Type* to **Wall**.



2. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `13.5 s`.

Contents

Running the Simulation
Visualizing Results

Running the Simulation

The simulation is now ready to be run.

Use the following steps:

1. Run the simulation by clicking (**Run**) in the **Solution** toolbar.
2. Save the simulation.

Visualizing Results

When the solver is complete, the *Scalar Scene 1* in the *Graphics* window appears.

See below:



Stage Three of the Simulation

For this final stage of the simulation, open the outlets (one on each side) to allow the particles to discharge. Also, the inlet to allow new particles to enter the furnace to supplement the exiting ones already in the furnace.

Use the following steps:

1. Edit the **Regions** > **Region 1** > **Boundaries** node.
2. Set the following properties:

Node	Property	Setting
pIn	*Type*	**Flow-Split Outlet**

Node	Property	Setting
Physics Values > Split Ratio	*Split Ratio*	`1`
pOut-2	*Type*	**Velocity Inlet**
Physics Values > Velocity Magnitude > Constant	*Value*	`1e-4 m/s`
pOut-3	*Type*	**Velocity Inlet**
Physics Values > Velocity Magnitude > Constant	*Value*	`1e-4 m/s`



3. Select the **Stopping Criteria > Maximum Physical Time** node and set *Maximum Physical Time* to `50 s`.



Contents

Running the Simulation
Visualizing Results

Running the Simulation

The simulation is now ready to be run.

Use the following steps:

1. Run the simulation by clicking (**Run**) in the **Solution** toolbar.
2. Save the simulation.

In rare occasions, clumps of particles block the flow of particles at an outlet. This causes the scene to show a skewed distribution of particles. If blocking occurs, use the following procedure:

3. Stop the simulation by clicking (**Stop**) in the **Solution** toolbar.
4. Close the simulation without saving.
5. Reload the last saved simulation.
6. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `13.6 s`.
7. Select **Continua** > **Physics 1** > **Models** > **Multiphase Interactions** > **Phase Interactions** > **Phase Interaction 2** > **Models** > **Hertz-Mindlin** and set *Static Friction Coefficient* to `0`.

8. Run the simulation by clicking (**Run**) in the **Solution** toolbar.

   When the simulation stops, both outlets are unblocked.

Return to solving the case as before the blockage occurred:

9. Select the **Stopping Criteria** > **Maximum Physical Time** node and set the *Maximum Physical Time* to `50 s`.
10. Select **Continua** > **Physics 1** > **Models** > **Multiphase Interactions** > **Phase Interactions** > **ORE-ORE** > **Models** > **Hertz-Mindlin** and set the *Static Friction Coefficient* to `0.4`.
11. Run the simulation by clicking (**Run**) in the **Solution** toolbar.
12. Save the simulation.

Visualizing Results

When the solver is complete, the *Scalar Scene 1* in the *Graphics* window appears.

See below:



Particles that have residence time of more than 36.5 s constitute the deadman region. In this region, particles move rather slowly and hence have high residence times.

Comparing the Results

In this section, compare the height and shape of the deadman region with the experimental image of Zhou *et al*. To show the deadman region more clearly, you create the user-defined field function, *oldParticles*, which displays the particles with residence times of less than 36.5 s all in the same color.

To define the region in a field function:

1. Open the **Tools** node at the bottom of the main object tree. Right-click the **Field Functions** node and select **New > Scalar**.



2. Rename the **User Field Function 1** node to `Old Particles`
3. Select the **Old Particles** node and set the following properties:

Property	Setting
Function Name	`OldParticles`
Definition	`max(0,$ParticleResidenceTime - ($Time-13.5)) + ($Time-13.5)`

4. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1**> **Scalar Field** node and set the *Function* to **Old Particles**.



The *Scalar Scene 1* in the *Graphics* window appears as shown below:



Deep blue colored particles are newer particles, injected after time t=13.5 s. The remaining particles are the ones that had come into the domain before 13.5 s.

The following screenshot shows the comparison, with the experimental image superimposed on top of the Old Particles scene above and zoomed in using the mouse buttons:



The deadman size, height, and shape are well-predicted.

Summary

This tutorial simulated a three-dimensional solid flow through a model blast furnace.

This tutorial has introduced the following STAR-CCM+ features:

- Selecting fluid physics models.
- Definition of Lagrangian phases.
- Definition of DEM particle interaction.
- Creation of a particle injector.
- Setting solver and stopping criteria.
- Setting up visualization scenes.
- Running the simulation.
- Visualizing the results.

Blast Furnace Tutorial Bibliography

Zhou, Z., Zhu, H., Yu, A., Wright, B., Pinson, D. and Zulli, P. 2005, "Discrete Particle Simulation of Solid Flow in a Model Blast Furnace." *ISIJ International* , vol 45, no 12, pp 1828 – 1837.

Motion

The tutorials in this set illustrate various STAR-CCM+ features for simulating problems with moving geometries and meshes, dynamic fluid body interaction, and rigid body motion:

Contents

Moving Reference Frames: Rotating Fan
Rigid Body Motion: Rotating Fan
DFBI: Boat in Head Waves
DFBI: Lifeboat with Overset Mesh
Marine Resistance Prediction: KCS Hull with a Rudder
Moving Reference Frames: Marine Propeller in Open Water
Body Force Propeller Method: Marine Self-Propulsion
Morphing: Cylinder with Boundary Motion

Moving Reference Frames: Rotating Fan

This tutorial outlines the steps that are required to set up and run a rotating radial fan analysis. It uses the Moving Reference Frame model, a steady-state approach that involves two or more frames of reference that can be stationary or moving relative to each other.

The problem involves an idealized fan assembly whose axis and 12 equally spaced blades are rotating at 2,000 rpm while the external body remains stationary. Air enters through the inlet, that is shown in red below, at 5 m/s and exits through a pressure boundary, that is shown in orange.



In this case, two reference frames are used for the analysis — a spinning one representing the fan and blades and a stationary one representing the fan body and inlet section.

Contents

Importing the Mesh
Scaling the Mesh
Setting up Physics Models
Creating a Moving Reference Frame
Setting Initial Conditions

Creating Interfaces
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Reporting Monitoring and Plotting
Running the Simulation
Visualizing the Results
Summary

Importing the Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/motion/data`.
5. Select `fan.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window. Two mesh regions named **Fluid** and **Rotating** are created under the **Regions** node.
6. Save the simulation as `fan_mrf.sim`.

Scaling the Mesh

For the input file supplied, STAR-CCM+ interprets the length units in meters instead of millimeters. You can restore the correct dimensions by scaling the mesh.

The original grid was built using millimeters as the length unit, but no scaling factor was applied when the grid was saved in the .ccm file. You can refer to the output that is produced at import for information on the mesh regions. The mesh regions extend from roughly -113 to 240 units in the x-direction, -133 to 98 in the y-direction and -16 to 184 in the z-direction. Without any further action, STAR-CCM+ interprets these values as meters, which would make the fan rather large. Therefore, scale the region by a factor of 0.001 so that the dimensions are correctly specified:

1. Select **Mesh** > **Scale Mesh** from the menu bar.

2. In the *Scale Mesh* dialog, make sure that the **Fluid** and **Rotating** regions are selected.
3. Enter a value of `0.001` for the **Scale Factor**.



4. Click **Apply** to scale the domain and then click **Close**.

   When the Apply button is clicked, the mesh region is reduced in size.
5. To restore the previous viewing distance from the scaled domain in the display, click the **Reset View** button in the toolbar.

To verify that scaling has in fact been applied, you can perform the mesh diagnostics check and review the output values.

Setting up Physics Models

As previously noted, the mesh consists of two regions: Fluid and Rotating. However, only one continuum is needed for the simulation since the two regions contain the same material. The physical models for this continuum are defined next.

To set the physics models:

1. For the physics continuum, **Continua** > **Physics 1** , select the following models, in order:

Group Box	Models
Time	Steady
Material	Gas
Flow	Segregated Flow
Equation of State	Constant Density
Viscous Regime	Turbulent
Reynolds-Averaged Turbulence	K-Epsilon Turbulence

2. Click **Close**.
3. To review the models, open **Physics 1** > **Models** .
   The selected models appear as shown:



4. Save the simulation.

Creating a Moving Reference Frame

Create a moving reference frame for the rotating region.

Reference frames are created and managed within the **Tools** > **Reference Frames** node, and are applied to regions after they are created.

1. Right-click the **Tools** > **Reference Frames** node and select **New** > **Rotating**.



2. Select the **Rotating** node and set the **Rotation Rate** to `2000 rpm`.

   The other properties are left unchanged as the default definition for the axis of rotation is suitable for this problem.

Assign the reference frame to the **Rotating** region:

3. Select the **Regions** > **Rotating** > **Physics Values** > **Motion Specification** node.



4. In the *Properties* window, set **Reference Frame** to **Rotating**.
   The **Physics Values > Axis** node of the region is removed, as the axis specification is now taken from the Rotating reference frame.

Setting Initial Conditions

Change the turbulence parameters that are set as part of the initial conditions.

To set the initial conditions:

1. Edit the **Continua > Physics 1 > Initial Conditions** node and set the following properties:

Node	Property	Setting
Turbulence Specification	**Method**	**K + Epsilon**
Turbulent Dissipation Rate > Constant	**Value**	`150.0 m^2/ s^3`
Turbulent Kinetic Energy > Constant	**Value**	`1.5 J/kg`

2. Save the simulation.

Creating Interfaces

All regions in a STAR-CCM+ simulation require an interface to be present between them that transfers the appropriate mass and energy quantities during the calculation.

In this example, create two interfaces between the **Fluid** and **Rotating** regions. This process involves creating so-called “In-place” interfaces between the default wall boundaries enclosing each region.

The first interface is created between the inlet and rotating section:

1. Open the two **Boundaries** nodes inside the **Fluid** and **Rotating** nodes.
2. Multi-select the two nodes named **Interface1** from each boundary grouping.
   These two boundaries represent the interface between the inlet section of the fan and the rotating region and are now highlighted in pink in the Geometry Scene display.

3. Right-click either of the selected nodes and then select **Create Interface**.



The above step creates a node within each **Boundaries** node called **Interface1 [Interface 1]**.

The second interface is created between the fan body and the rotating region.

4. Multi-select the two nodes named **Interface2** from each boundary grouping.
   These two boundaries represent the interface between the rotating region and the rest of the stationary part of the fan and are now highlighted in pink in the *Geometry Scene* display.
5. Right-click either of the selected nodes and then select **Create Interface** to create another in-place interface called **Interface2 [Interface 2]**.
6. Save the simulation.

In addition to the in-place boundaries, an **Interfaces** node appears in the simulation tree. Within it, an interface node that is called **Interface 1** is created. This forms the connection between the **Fluid** and **Rotating** regions once the flow initialization step has been performed and supersedes the original wall boundaries forming the connection.

Setting Boundary Conditions

The fluid velocity at the inlet and the rotation conditions at the wall boundaries are specified.

STAR-CCM+'s default rotation setting for all wall boundaries is for them to remain stationary relative to the mesh. Therefore, in absolute terms, the default setting for the wall boundaries of the Fluid region is zero rotation while the default setting for the wall boundaries of the Rotating region is 2000 rpm. This setting is inappropriate for one wall boundary in each region:

- The **Axis** boundary in the **Fluid** region, which is rotating as it forms part of the axle rotating the blades.
- The **Walls** boundary in the **Rotating** region, which is stationary as the walls form part of the casing.

To set the boundary conditions for the **Fluid** region:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Inlet** node and set the following properties:

Node	Property	Setting
Physics Values		
Velocity Magnitude	**Method**	**Constant**
Constant	**Value**	`5.0 m/s`

2. Edit the **Regions** > **Fluid** > **Boundaries** > **Axis** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Tangential Velocity Specification	*Method*	**Rotation Rate**
Physics Values		
Wall Relative Rotation	*Method*	**Constant**
Constant	*Value*	`2000 rpm`

To set the boundary conditions for the **Rotating** region:

3. Edit the **Regions** > **Rotating** > **Boundaries** > **Walls** node and set the following properties.

Node	Property	Setting
Physics Conditions		
Tangential Velocity Specification	*Method*	**Fixed**
Reference Frame Specification	*Option*	**Lab Frame**

Setting Solver Parameters and Stopping Criteria

To make sure that a reasonable solution is obtained, alter some of the solver parameters.

To set the solver parameters and the stopping criteria:

1. Select the **Solvers** > **Segregated Flow** > **Velocity** node and change the **Under-Relaxation Factor** to `0.5`.
2. Similarly, change the **Pressure Under-Relaxation Factor** to `0.2`.
3. Select the **Pressure** > **AMG Linear Solver** node and change the **Convergence Tolerance** to `0.05`.

The simulation runs for 600 iterations. This number can be specified using a stopping criterion, as follows:

4. Select the **Stopping Criteria** > **Maximum Steps** node and change the **Maximum Steps** to `600`.
5. Save the simulation.

Visualizing the Solution

Visualize the solution while it is being updated by plotting velocity vectors on an x-y section bisecting the problem geometry.

1. Create the section:
   a) Right-click the Derived Parts node and select **New Part** > **Section** > **Plane**.
   b) In the *Create Section* dialog, set the following properties:

Property	Setting
Input Parts	**[Fluid, Rotating]**
Origin	`[0.0, 0.0, 0.0] m,m,m`
Normal	`[0.0, 0.0, 1.0] m,m,m`
Display	**New Geometry Displayer**

c) Click **Create** and then **Close**.

2. Create the vector scene:
   a) Right-click the **Scenes** node and select **New Scene > Vector**.
   b) Select the **Scenes > Vector Scene 1 > Vector 1 > Displayers > Parts** node and add plane section to the Parts list.
   c) Repeat this process so that the only part in the **Parts** list in the **Outline** displayer is the plane section.
   d) Click the camera icon in the Visualization toolbar and select **Projection Mode > Parallel**.
   e) Click the camera icon in the Visualization toolbar and select **Views > -Z > Up -Y**.
3. To show the entire section more clearly, use the middle mouse to zoom in a little.

Reporting Monitoring and Plotting

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. The process involves setting up a report which defines the quantity of interest and the parts of the region to monitor. A monitor is then defined based on the report which controls the update frequency and normalization characteristics. An X-Y graph plot can then be made from the monitor.

To track the area-averaged pressure at the inlet boundary, use a report monitor. This value is a useful measure of whether the solution has converged.

1. In the simulation tree, right-click the **Reports** node and select **New Report > Surface Average**.



2. Rename the **Reports > Surface Average 1**node to `Area Averaged Inlet Pressure`.
3. Select the **Area Averaged Inlet Pressure** node and add **Fluid: Inlet** to the **Parts** list and change the **Scalar Field Function** to **Pressure**.



4. To create the monitor and plot that are based on this report, right-click the **Area Averaged Inlet Pressure** node and select **Create Monitor and Plot from Report.**

   A new node named **Area Averaged Inlet Pressure Monitor Plot** appears under the **Plots** node. To preview the plot display, you can double-click this node, or right-click it and select **Open** from the pop-up menu. A new plot display that is called **Area Averaged Inlet Pressure Monitor Plot** appears in the *Graphics* window.
5. Save the simulation.

Running the Simulation

You can now run the simulation.

To start the analysis:

1. Click (Run) in the toolbar.
   The *Residuals* display is automatically created and shows the progress of the solver.

For this case, however, the *Area Averaged Inlet Pressure Monitor Plot* is a more useful plot to view during the run. The tabs at the top of the Graphics window make it possible to select any active display for viewing.

During the run, it is possible to stop the solution by clicking the (Stop) button in the toolbar. If you do halt the simulation, it can be continued by clicking the (Run) button. If left alone, the simulation continues until 600 iterations are completed.

To watch the solution as it is updated:

2. Click the tab for the *Vector Scene 1* display.



At the end of the calculation, the Area Averaged Inlet Pressure Monitor Plot looks as shown below



The pressure at the inlet has almost converged to a steady value. However, it is unlikely to do so completely as the solution to this problem is not truly steady-state.

3. Save the simulation.

Visualizing the Results

You can visualize the velocity vector results for the converged solution.

To visualize the results:

1. Click the tab for the *Vector Scene 1* display.



2. Save the simulation.

Summary

This tutorial demonstrated how to set up a moving reference simulation in STAR-CCM+.

This tutorial introduced the following STAR-CCM+ features:

- Defining a moving reference frame model.
- Setting initial conditions.
- Creating interfaces between different mesh regions.
- Specifying region-wise conditions for moving reference frames.
- Defining boundary conditions.
- Setting solver parameters.
- Setting up a monitoring report and plot.
- Running the solver for a set number of iterations.

- Analyzing the results.

Rigid Body Motion: Rotating Fan

This tutorial models the same radial fan problem as the Moving Reference Frames: Rotating Fan. However, instead of using the steady-state moving reference frame model, this tutorial uses the transient Rigid Body Motion model to solve the problem.

To avoid having to set up the entire case, this tutorial restarts from the `.sim` file that was created in the Moving Reference Frames Rotating Fan Overview tutorial.

Contents

Loading an Existing Simulation
Setting up the Rigid Body Motion Model
Setting Solver Parameters and Stopping Criteria
Monitoring the Solution
Initializing and Running the Simulation
Checking the Results
Summary

Loading an Existing Simulation

For this tutorial, you start from a previously saved simulation file.

Note: To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file.

1. Launch STAR-CCM+.
2. Select **File** > **Load Simulation**.
3. In the *Load Simulation* dialog, navigate to the `fan_mrf.sim` file that was created in the Moving Reference Frames: Rotating Fan tutorial.
4. Select `fan_mrf.sim` then click **Open**.
5. Click **OK**.
6. Save the simulation as `fan_rbm.sim`.

Setting up the Rigid Body Motion Model

The first step in setting up the Rigid Body Motion model is to switch the analysis type to unsteady.

For the physics continuum, **Physics 1**:

1. Clear the **Steady model** from the *Enabled Models* box.
2. Select **Implicit Unsteady** from the *Time* box.
3. Click **Close**.
4. To review the models, open the **Physics 1** > **Models** node.



Create a motion that uses the rigid body motion method and apply this new motion to the **Motion Specification** node of the rotating region.

5. Right-click the **Tools** > **Motions** node and select **New** > **Rotation**.



A new **Rotation** node is added to the **Motions** node.

6. Select the **Rotation** node and set the following properties:

Property	Setting
Axis Direction	[0.0, 0.0, 1.0]
Rotation Rate	2000 rpm

Apply this motion to the rotating region, and change the reference frame for the region to the stationary default frame.

7. Select the **Regions** > **Rotating** > **Physics Values** > **Motion Specification** node and set **Motion** to **Rotation**.

   Selecting this motion results in the **Reference Frame** property automatically reverting to the default **Lab Reference Frame**.

8. Save the simulation.

Setting Solver Parameters and Stopping Criteria

As the analysis is now unsteady, the time-step size is specified. For the specified rotational speed of 2000 rpm, it takes 0.03 s to complete one full revolution.

As a rule of thumb it is suggested to use 1 degree per time-step. This results in a time-step size of 0.03s/360 or 8.33E-5 s. For accuracy, a 2nd-order temporal discretization scheme is used. For demonstration purposes you solve three solution cycles, 0.09 s, only.

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicity Unsteady	**Time-Step**	`8.33333333e-5 s`
	Temporal Discretization	**2nd-order**
Segregated Flow > Velocity	**Under-Relaxation Factor**	`0.9`
Segregated Flow > Pressure	**Under-Relaxation Factor**	`0.9`

You can specify the maximum physical time to be 0.09 s. The maximum inner iterations value must be selected carefully: if the value is too low, the solution does not converge within a time-step; if it is too high the solution takes a long time to complete. Select a value of 7 and add a runtime monitor later.

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
	Maximum Physical Time	`0.09 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`7`
Maximum Steps	**Enabled**	Deactivated

3. Save the simulation.

Monitoring the Solution

Create a report, monitor, and plot to track the moment of force on the surface of the rotating blades during the simulation. Monitoring this value helps to make sure that the solution is converging within each time-step.

1. Right-click the **Reports** node and select **New Report > Moment**.

   A node that is called **Moment 1** is added to the **Reports** node and is selected automatically.
2. Select the **Moment 1** and set **Parts** to **Rotating > Blades**.



To create the monitor and plot based on this report:

3. Right-click the **Reports** > **Moment 1** node and select **Create Monitor and Plot from Report**.

   A new node named **Moment 1 Monitor** appears under the **Monitors** node.

4. Select the **Monitors** > **Moment 1 Monitor** node and make sure that the **Trigger** is set to **Iteration**. Monitoring the moment during time-steps provides the data for intra time-step convergence checks.



A new node named **Moment 1 Monitor Plot** also appears under the **Plots** node.

Although it is important to monitor convergence between time-steps, for most plots in unsteady cases it is more useful to examine the solution progress by reporting a quantity every time-step rather than every iteration. Modify the Area Averaged Inlet Pressure Monitor plot and the vector scene to update every time-step.

For the vector scene:

5. Select the **Scenes** > **Vector Scene 1** > **Attributes** > **Update** node and change the **Trigger** to **Time-Step**.

Next, set up the maximum wall pressure as the monitoring quantity:

6. Select the **Monitors** > **Area Averaged Inlet Pressure Monitor** node and change the **Trigger** to **Time-Step**.

To preview both of the plots:

7. Multi-select both plots. Right-click the selected nodes and select **Open**.



Make sure that the x-Axis Monitor is set to the correct type for each plot:

8. Select **Plots** > **Area Average Inlet Pressure Monitor Plot** and set **x-Axis Monitor** to **Physical Time**.
9. Select **Plots** > **Moment 1 Monitor Plot** and set **x-Axis Monitor** to **Iteration**.

   The selected plots appear in the *Graphics* window.
10. Save the simulation.

Initializing and Running the Simulation

The analysis starts from the same stationary initial flow field that was set up in the Moving Reference Frames: Rotating Fan tutorial, therefore the existing solution must be cleared first.

1. Select the **Solution** > **Clear Solution** > menu item and click **OK** on the displayed *Clear Solution* dialog.
2. Click (**Run**) in the toolbar.

   The *Residuals* display is created automatically and shows the progress of the solver.

During the run, make sure that the number of inner iterations that was selected is sufficient to achieve convergence.

3. Click the *Moment 1 Monitor Plot* tab at the top of the *Graphics* window.
4. Select the **Plots** > **Moment 1 Monitor Plot** > **Data Series** > **Moment 1 Monitor** > **Symbol Style** node and set **Shape** to **Filled Circle**.
5. In the Plot toolbar, select the (**Toggle Plot Zoom**) button to zoom into the graph over approximately the last 20 iterations.

   The plot appears similar to below (Note: this plot is for demonstration purposes only, results may differ based on the iterations selected).



From the above profile, it is clear the solution is indeed converging within each time-step. Convergence is evident by the leveling off of the profile in each time-step.

The simulation continues until all time-steps are complete.

6. After the run is finished, save the simulation.

Checking the Results

The vector plot is compared with the results for the Moving frame of Reference case.

1. Right-click the **Scene > Vector Scene 1** node and select **Open**.



These velocity vectors are similar to the vectors that were created during the Moving Reference Frames: Rotating Fan tutorial.

2. Click the *Area Averaged Inlet Pressure Monitor Plot* tab at the top of the *Graphics* window.

The large scale that is used for the y-axis makes it difficult to assess the inlet pressure variation over the solution cycles.

To alter the scale:

3. Edit the **Plots > Area Averaged Inlet Pressure Monitor Plot > Axes** node and set the following properties:

Node	Property	Setting
Y Axis	**Minimum**	`-240`
	Maximum	`40`
Y Axis > Major Labels	**Spacing**	`100.0`
	Precision	`-1`
X Axis	**Maximum**	`0.03`
	Lock Minimum	**Activated**
X Axis > Major Labels	**Spacing**	`0.005`

The plot now appears as shown:



Note: The pressure at the inlet is oscillating with a period of 0.0025 s, that is 1/12 of a revolution, as expected. Over time the amplitude of these oscillations become regular as the flow cycle settles down.

4. Save the simulation.

Summary

This tutorial demonstrated how to set up a rigid body motion simulation in STAR-CCM+.

This tutorial introduced the following STAR-CCM+ features:

- Defining a rigid body motion model.
- Setting solver parameters and stopping criteria for an unsteady analysis.

- Monitoring the solution progress using plots and vector scenes.
- Running the solver for a set time.
- Analyzing the results using STAR-CCM+'s visualization and XY-plot facilities.

DFBI: Boat in Head Waves

This tutorial demonstrates use of the Six Degrees Of Freedom solver with a free surface flow to model the motion of a boat hull facing oncoming waves.

The appropriate features in STAR-CCM+ provide initial conditions and wave generation.

Contents

Importing the Mesh and Naming the Simulation
Setting Up the Models
Setting the Material Properties
Selecting the 6-DOF Motion Model
Setting the Body Initial Coordinate System
Creating the 6-DOF Body
Defining the VOF Waves
Setting Initial Conditions
Setting the Boundary Conditions
Setting Up Solver Parameters and Stopping Criteria
Creating the Free Surface Isosurface
Creating a Custom Colormap
Visualizing the Free Surface
Creating DFBI Displayers
Exporting Images for Scene Animation
Reporting, Monitoring, and Plotting
Initializing and Running the Simulation
Visualizing the Results
Summary

Importing the Mesh and Naming the Simulation

You launch STAR-CCM+, import the supplied mesh, and save the simulation.

To import the mesh:

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the New Simulation option from the menu bar.

2. Import the three-dimensional hexahedral cell mesh, `boat.ccm` from the `doc/startutorialsdata/motion/data` subdirectory of your STAR-CCM+ installation.
3. Save the new simulation as `boat.sim`.

Setting Up the Models

Several physical models are activated to simulate the forces acting on the hull. The simulation models the behavior of two fluids (air and water) within the same continuum, and uses the Volume of Fluid model to do so.

As there are two fluids in different phases, the Eulerian Multiphase model is activated, and the effect of gravity acting on both is included using the Gravity model.

The effect of turbulence on the fluid is modeled using the default K-Epsilon turbulence model.

A default Physics Continuum was created when the mesh was imported. All necessary models are defined using this continuum.

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase**
Multiphase Model	**Volume of Fluid (VOF)**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Optional Models	**Gravity** **VOF Waves**

2. Click **Close**.

The selected models appear as shown below.



3. Save the simulation .

Setting the Material Properties

You set the material properties for the Eulerian phases.

In the boat window, open the **Continua** node. The color of the **Physics 1** node has turned from gray to blue to indicate that models have been activated for that continuum.

1. Open the **Physics 1** > **Models** node.

   The selected models now appear within that node.
2. Open the **Eulerian Multiphase** node.

Define the materials corresponding to each of the mixture phases under the **Eulerian Phases** node.

3. Right-click the **Eulerian Phases** node and select **New** to create a phase node that is called **Phase 1**
4. Rename the **Phase 1** node to `Water`

It is now necessary to set the models that determine the behavior of the *Water* phase.

5. Right-click the **Eulerian Phases** > **Water** > **Models** node and choose **Select models...**
6. Select **Liquid** from the *Material* group box.
7. Select **Constant Density** from the *Equation of State* group box.
8. Repeat the steps for a second phase except this time select **Gas** rather than **Liquid**, and specify **Constant Density**.
9. Rename the phase to `Air`

In both instances, you are using the default material appropriate to each phase (water for the liquid and air for the gas), so no further changes are necessary to material properties. The material selections can be seen by expanding the **Water** and **Air** nodes within the **Eulerian Phases** parent node.



10. Save the simulation.

Selecting the 6-DOF Motion Model

The 6-DOF motion solver is activated by creating the relevant DFBI motion object and assigning it to the fluid region. In this case, the DFBI rotation and translation motion is required.

To select the 6-DOF motion model:

1. Open the **Tools** > **Motions** node.

   The default **Stationary** motion is displayed.

2. Right-click the **Motions** node and choose **New** > **DFBI Rotation and Translation** as shown below.



3. A new **DFBI Rotation and Translation** node is added to the **Motion** manager.

Select this motion within the **Motion Specification** node for the fluid region.

4. Select the **Regions** > **Region 1** > **Physics Values** > **Motion Specification** node and set *Motion* to **DFBI Rotation and Translation**.



The action of setting the DFBI motion on the region exposes the **DFBI** manager node on the object tree, and adds the **6-DOF Solver** and **6-DOF Motion** nodes to the **Solvers** manager.



Setting the Body Initial Coordinate System

The 6-DOF Solver must know the initial orientation of the body relative to the laboratory coordinate system. Therefore, define a local coordinate system in which the x-axis is aligned with the forward direction of the body and the z-axis in the vertical direction.

To set the body initial coordinate system:

1. Make sure that *Geometry Scene 1* is still open.
2. Open the **Tools > Coordinate Systems** node and right-click the **Laboratory > Local Coordinate Systems** node. Choose **New > Cartesian** from the pop-up menu

   The *Create Cartesian Coordinate System* window is displayed in the Explorer Pane.
3. Set the *X* value of the *i Direction* to `-1` and the *Y* value of the *j Direction* to `-1` as shown below.



4. Click **Create** and then **Close**.

The node is shown as **Cartesian 1** in the object tree

5. Rename the **Cartesian 1** node to `Initial ship orientation`

Note: The location and orientation of the coordinate system can be checked visually in *Geometry Scene 1*. Select the **Initial ship orientation** node in the object tree and rotate the model in the scene as required.



Creating the 6-DOF Body

The next step is to define the surface of the floating body and its properties.

To define the surface and its properties, first create a 6-DOF body to which the boundary representing the boat hull is assigned:

1. Right-click the **DFBI > 6-DOF Bodies** node and select **New** from the pop-up menu.
2. Rename the **Body 1** node to `Boat`
3. Select the **Boat** node and click the custom editor button ... to the right of the *Body Surface* property. The *Select Objects* dialog appears.

4. Expand the **Regions** and **Region_1** nodes and tick the checkbox for the **hull** node. Click **OK** to close the dialog.



5. Set the *Body Mass* to `400.0 kg`.
6. So that the fluid solution can develop before starting the body motion, set the *Release Time* property to `0.5 s`.

In this simulation, the body is allowed to rotate about the y-axis only, to simulate motion in the head waves. This means that the moment of inertia about the y-axis must be defined.

7. Select the **Boat > Initial Values > Moment of Inertia node:**
   a) Set *Diagonal Components* to `[1000.0, 1000.0, 1000.0] kg-m^2`.
   b) Set *Coordinate System* to **Laboratory->Initial ship orientation**.
   c) Activate *Use Center of Mass*.



The previously created coordinate system defines the initial orientation of the body.

8. Select the **Boat > Initial Values > Orientation** node and set *Coordinate System* to **Laboratory->Initial ship orientation**.

Although the 6-DOF body has now been created with the Free Motion option, it is still necessary to specify the manner in which the body can move. In this case, the body must be free to translate in the Z direction (heave) and to rotate about the Y axis (pitch).

9. Select the **DFBI > 6-DOF Bodies > Boat > Free Motion** node and tick the *Z Motion* and *Y Rotation* properties.



10. Save the simulation.

Defining the VOF Waves

It is possible to define transient boundary conditions that simulate waves on the free surface using the VOF Wave physics model in STAR-CCM+.

The VOF Wave model generates a number of field functions that are based on wave parameters that you define. These field functions can be used as boundary and initial conditions in appropriate nodes on the object tree.

1. To create the VOF Wave, right-click the **Continua** > **Physics 1** > **Models** > **VOF Waves** > **Waves** node and select **New** > **FifthOrder**.
2. Rename the **FifthOrderVofWave 1** node to `Head Wave`.
3. Select the **Head Wave** node and set the *Current* and *Wind* properties to be `[2.5, 0.0, 0.0] m/s`
4. Set *WaveHeight* to `0.2 m`
5. Set *WaterDepth* to `10 m`
6. Set *Specification Type* to **Wave Length**.



7. Select the **Head Wave** > **WaveLengthSpecification** node and set *WaveLength* to `6.0 m`



Setting Initial Conditions

Set the initial conditions using field functions that are associated with the previously defined VOF Wave.

Initial conditions that are set are the free-surface wave profile, the velocity profile within the phases and the hydrostatic pressure.

To set the initial conditions:

1. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Volume Fraction** node and set *Method* to `Composite`.



The **Volume Fraction** node now contains a **Composite** node in response to this change of the *Method* property.

2. Select the **Volume Fraction** > **Composite** > **Water** node and set *Method* to `Field Function`.

   The **Water** node contains a **Field Function** node after this change has been made.



3. Select the **Water** > **Field Function** node and set *Scalar Function* to `Volume Fraction of Heavy Fluid of Head Wave`.

4. Repeat these steps with the **Volume Fraction** > **Air** node, but set *Scalar Function* to `Volume Fraction of Light Fluid of Head Wave`.
5. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Velocity** node and set *Method* to Field `Function`.

   The node changes to **Field Function** in response.



6. Select the **Field Function** node and set *Vector Function* to `Velocity of Head Wave`.



7. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Pressure** node and set *Method* to *Field Function*.
8. Select the **Field Function** node and set *Scalar Function* to `Hydrostatic Pressure of Head Wave`.

9. Save the simulation.

Setting the Boundary Conditions

There are four boundaries for this geometry:

1. No-slip wall boundary on the hull of the boat
2. Symmetry boundary on the Z-X plane at Y = 0.0
3. Pressure outlet on the downstream face
4. Velocity inlet on all remaining boundaries

The boundary types were correctly set during the initial mesh import, so it only remains to set the conditions corresponding to each type. Default conditions are suitable for the symmetry plane and the no-slip wall boundary on the boat hull. Wave conditions are applied to the pressure outlet and velocity inlet.

Contents

Inlet Boundary
Pressure Boundary

Inlet Boundary

The head wave enters the simulation through the inlet boundary. You use field functions to define this oncoming wave.

To set the inlet boundary:

1. Edit the **Regions** > **Region_1** > **Boundaries** > **inlet** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Velocity Specification	*Method*	**Components**
Physics Values		
Velocity	*Method*	**Field Function**
Field Function	*Vector Function*	**Velocity of Head Wave**
Volume Fraction	*Method*	**Composite**
Composite > Water	*Method*	**Field Function**

Node	Property	Setting
Field Function	*Scalar Function*	**Volume Fraction of Heavy Fluid of Head Wave**
Composite > Air	*Method*	**Field Function**
Field Function	*Scalar Function*	**Volume Fraction of Light Fluid of Head Wave**

Pressure Boundary

The head wave leaves the domain through the pressure boundary. You use field functions to define the wave at the boundary.

To set the pressure boundary:

1. Edit the **Regions > Region_1 > Boundaries > pressure** node and set the following properties:

Node	Property	Setting
Physics Values		
Pressure	*Method*	**Field Function**
Field Function	*Scalar Function*	**Hydrostatic Pressure of Head Wave**
Volume Fraction	*Method*	**Composite**
Composite > Water	*Method*	**Field Function**
Field Function	*Scalar Function*	**Volume Fraction of Heavy Fluid of Head Wave**
Composite > Air	*Method*	**Field Function**
Field Function	*Scalar Function*	**Volume Fraction of Light Fluid of Head Wave**

Setting Up Solver Parameters and Stopping Criteria

As this simulation is transient, set the time step, the maximum number of inner iterations that are permitted in each time step, and the overall physical time for which a solution is obtained.

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `0.01 s`
2. Select the **Stopping Criteria** > **Maximum Inner Iterations** node and make sure that the *Maximum Inner Iterations* is set to `5`
3. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Max Physical Time* to `2.0 s`

Creating the Free Surface Isosurface

It is instructive to visualize the motion of the boat hull in response to the oncoming wave over time.

This motion can be visualized by creating a scalar scene showing the boat hull with the free surface, and writing out this scene as a graphics file at each time step of the simulation. The resulting collection of files can be used to form an animation of the free surface and hull motion.

To plot the surface of the water in a scalar scene, first create an isosurface that is based on the volume fraction of water in a cell:

1. Right-click the **Derived Parts** node and select **New Part** > **Isosurface...**

   The *Create Isosurface* dialog opens over the object tree. In the *Create Isosurface* dialog:
2. Set *Scalar* to **Volume Fraction** > **Water**.
3. Leave the *Extraction Mode* set to `Single Value`.
4. Set *Isovalue* to `0.5`.
5. Select **Existing Displayer** > **Geometry 1** in the *Display* panel.

   The completed dialog is as shown below.



6. Click **Create**, and then **Close**
7. Rename the **Isosurface**node to `Water Surface`

Creating a Custom Colormap

To make the water surface appear in blue for visualization, you can create a custom colormap.

Set the color of the free surface artificially to a uniform blue. First create a custom **Colormap** within the **Tools** node as follows:

1. Right-click the **Tools** > **Colormaps** node and select **New**.
2. Rename the **user1** node to `Water`
3. Select the **Tools** > **Colormaps** > **Water** node and open the *Color Map* property for editing by clicking the editor icon.
4. Click and drag the blue control point all the way to the right so that the color bar is entirely blue.
5. Set *Color Space* to **RGB**.

   The completed dialog is shown below.



6. Click **OK**.

Visualizing the Free Surface

The **Water Surface** part is now used in creating the scalar scene.

1. Right-click the **Scenes** node, and select **New Scene** > **Scalar**.

2. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Outline 1** node and deactivate *Outline*.



3. Select the **Scalar 1** > **Parts** node and set **Parts** to **Derived Parts** > **Water Surface**.
4. Select the **Scalar Field** node and set *Function* to **Volume Fraction** > **Volume Fraction of Water**.
   a) Set *Clip* to **Off**.
   b) Set *Auto Range* to **Off**.
   c) Set *Min* to `0.5`.
   d) Set *Max* to `0.5`.

You can now return to the **Scenes** > **Scalar Scene 1** node to finish setting its properties.

5. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Color Bar** node and set *Color Map* to **Water**.
6. Deactivate *Shadow*.
7. Set *Levels* to `2`(as there are only two points on the **Water** color map).
8. Deactivate *Visible* because you are intending to display the free surface as a uniform color: the color bar is not necessary.

Some further changes are required to provide a complete representation of the boat hull.

9. Select the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node and set *Contour Style* to **Smooth Filled**.
10. Open the *Expert* properties and set *Transform* to **symmetry 1**.

Some adjustments are made to the scalar scene lighting to improve the appearance.

11. Right-click the **Scenes** > **Scalar Scene 1** > **Attributes** > **Lights** node and select the item **Head Light**. Only the first light remains active.



12. Open the **Lights** node and select the **Light 2** node.
13. Activate *Enabled* for the **Light 2** node.



Creating DFBI Displayers

Two displayers are added to the scene to display the center of mass of the boat and the forces and moments acting on it.

1. Right-click **Scenes** > **Scalar Scene 1** > **Displayers** and select **New Displayer** > **DFBI**

2. Select **DFBI 1** > **Parts** node and set *Parts* to **6-DOF Bodies** > **Boat**.
3. Select the **DFBI 1** node and set *Opacity* to `0.5`
4. Set *Transform* to **symmetry 1**.

The boat and an icon showing the location of the center of mass appear.

5. Right-click **Scenes** > **Scalar Scene 1** > **Displayers** and select **New Displayer** > **DFBI**
6. Select the **DFBI 2** > **Parts** node and set *Parts* to**6-DOF Bodies** > **Boat** > **External Forces and Moments** > **Fluid Force and Moment**.



7. Select **DFBI 2** > **Relative Size** node and set *Vector Glyph Length (%)* to `50`



8. Select **DFBI 2** node and set *Vector Length* to *Vector Magnitude*.

The force and moment are calculated during the run.

Exporting Images for Scene Animation

The final step in preparing for the free surface animation is to have STAR-CCM+ create a hardcopy of *Scalar Scene 1* each time-step.

1. Select the **Scenes** > **Scalar Scene 1** > **Attributes** > **Update** node and set *Trigger* to **Time-Step**.
2. Set the *Output Directory* to point to a suitable folder on your file system where multiple images can be saved.
3. Enter `wave_surface` for the *Base Filename* property. Image files that are created from this scene have `wave_surface` contained in the filename.
4. Activate *Save To File*.

   A red asterisk appears next to the **Scalar Scene 1** node in the object tree.
5. Select the **Update** > **Hardcopy** node and set the following properties:

Property	Setting
Output Width	`800`
Output Height	`600`

6. Save the simulation.

Reporting, Monitoring, and Plotting

Use report monitors in this simulation to plot the dynamic response of the motion of the boat to the incident waves over time.

1. Right-click the **Reports** node and select **New Report** > **6-DOF Body Translation**.

   A node that is called **Rigid Body Translation 1** appears within the **Reports** node.
2. Rename the **Rigid Body Translation 1** node to `Z Translation`
3. Select the **Z Translation** node and set *Direction* to `[0.0, 0.0, 1.0]`.



A second report is created to monitor the angle of rotation about the Y axis.

4. Right-click the **Reports** node in the object tree and select **New Report > 6-DOF Body Orientation**.
5. Rename the **Rigid Body Angle 1** node to `Y Rotation`
6. Select the **Y Rotation** node and set *Angle Report Option* to `Rotation Y Axis`.



Create monitors and plots that are based on these reports.

7. First select the **Reports > Y Rotation** node and then, while pressing down the `Ctrl` key, also select the **Z Translation** node.



8. Right-click the nodes and select item **Create Monitor and Plot from Report**

   The *Create Plot From Reports* dialog appears.
9. Select the **Multiple Plots (one per report)** option.

Two new nodes named **Z Translation Monitor Plot** and **Y Rotation Monitor Plot** appear under the **Plots** node.



Open the plots corresponding to these nodes, to preview.

The monitors that provide data for these plots are set to update every iteration by default. As this simulation is unsteady, it is better to update the monitors once per time-step instead at every iteration.

10. Open the **Monitors** node in the simulation tree and select both the **Y Rotation Monitor** and the **Z Translation Monitor** nodes using the `Ctrl Click` method.



11. In the *Properties* window, set *Trigger* to `Time Step`.
12. Save the simulation.

Initializing and Running the Simulation

To verify that the initial conditions applied to the regions are correct, STAR-CCM+ allows you to perform just solver initialization to start with and then view the results using an appropriate scene.

1. To initialize, click the (**Initialize Solution**) button in the toolbar or use the **Solution** > **Initialize Solution** menu item
2. Select the *Scalar Scene 1* window to see the initialization result. Use the mouse buttons to orient and size the model as shown below.



The force (red) and moment (blue) vectors for half of the boat are displayed during the run. To display the vectors for other half, enable the symmetry transform:

3. Select the **DFBI 2** node and set *Transform* to **symmetry 1**.

The force vector, a polar vector, is mirrored at the symmetry plane. The moment vector, an axial vector, maintains its magnitude but changes

direction due to reflection. To clear the forces or moments on the body untick the appropriate box in the **DFBI 2**node properties.

To reduce cluttering the scene, view the force and moment for one half of the boat only for this run.

4. Select the **DFBI 2** node:

   Set *Transform* to **Identity**.

   Activate **Forces** and **Moments**.

5. Click the (**Run**) button in the toolbar.

The *Y Rotation Monitor Plot* at the end of the simulation is shown below.



6. After the run is finished, save the simulation.

Visualizing the Results

The pressure distribution on the boat hull after 2.0 s is plotted in a scalar scene.

1. Right-click the **Scenes** node and select **New Scene > Scalar** from the pop-up menu.
2. Rename the **Scalar Scene 2** node as **Hull Pressure**
3. Select the **Scenes > Hull Pressure > Displayers > Outline 1** node and deactivate *Outline*.

   All visible outlines are removed from the *Hull Pressure* display.
4. Select the **Hull Pressure > Displayers > Scalar 1 > Parts** node and click the right-hand side of the *Parts* property.
5. In the dialog that appears, expand the **Regions** and **Region_1** nodes. Select the **hull** node and click **OK**.
6. Select the **Hull Pressure > Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Pressure**.
7. Select the **Hull Pressure > Displayers > Scalar 1** node and set *Contour Style* to **Smooth Filled**.
8. Expand the **Expert** properties and set *Transform* to **symmetry 1**.

Use the mouse controls to zoom and orient the scene as desired. A close-up of the bow is shown below.



9. Save the simulation.

Summary

This tutorial demonstrated use of the Six Degrees Of Freedom solver with a free surface flow to model the motion of a boat hull facing oncoming waves.

This tutorial introduced the following STAR-CCM+ features:

- Activating the 6-DOF model.
- Creating a VOF wave and using it for initial and boundary conditions.

- Obtaining a visualization of the free-surface flow using a derived part.
- Writing successive images of a scalar scene to disk during a transient simulation.

DFBI: Lifeboat with Overset Mesh

This tutorial demonstrates use of the Overset Mesh feature with free surface flow and DFBI to model the motion of a lifeboat falling into water. STAR-CCM+ automates the grid overlapping process.

The case is inspired by an experiment that is described in the paper "Simulation of Lifeboat Launching under Storm Conditions" by H.J. Morch et al. (2008). A lifeboat that is launched from a ramp undergoes linear and angular acceleration before hitting the water surface.



You model motion with three degrees of freedom: vertical and horizontal translation, and pitching rotation. This setup allows you to use symmetry conditions and include only half the geometry, thus reducing the computational cost of the simulation.

The full fall of the boat from rest on the ramp need not be modeled. Boat motion through the air can be pre-computed using simplified classical physics equations. The simulation is started with the boat initially positioned just above the water surface, with the following properties:

- Mass: 10,000 kg
- Moment of Inertia about axes passing through center of mass: 16000.0 kg m2
- Initial downward velocity: 22 m/s
- Initial angular velocity: 0 rad/s
- Initial inclination: 35

Contents

Importing the Mesh
Creating Regions and Setting Boundary Types

Selecting Physics Models
Coupling the Regions
Generating the Mesh
Transforming the Overset Region
Creating Eulerian Phases
Setting the DFBI Motion and 6-DOF Body
Defining the VOF Waves
Setting Initial Conditions
Setting the Boundary Conditions
Setting Up Solver Parameters and Stopping Criteria
Setting Scenes to Visualize the Solution
Initializing and Running the Simulation
Summary

Importing the Mesh

You import a CAD model consisting of two parts: a large block for the background region, and a smaller part from which the boat was subtracted to represent the overset region.

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment and create a New Simulation.
2. Import the CAD model, `lifeBoatOverset.x_b` from the `doc/startutorialsdata/motion/data` subdirectory of your STAR-CCM+ installation by selecting **File > Import > Import Surface Mesh**

The *Import Surface Options* dialog appears.

The default options are appropriate for the case.

3. Click **OK**.

Initially all parts of the imported geometry are shown as solid, gray surfaces with the boat hull hidden from sight. To see the hull shape in its imported orientation:

4. Click (**Save-Restore-Select Views**) and select **Views > -Y > Up +Z**.

Geometry Scene 1 displays the hull shape together with the outlines of the **Overset** and **Background** parts as shown below.



5. Save the new simulation as `lifeBoat.sim`.

Creating Regions and Setting Boundary Types

Create regions from the **Background** and **Overset** parts.

1. Multi-select the **Geometry > Parts > Background** and **Overset** nodes.
2. Right-click either selected node and select **Assign Parts to Regions**.
3. In the **Assign Parts to Regions** dialog:
   a) Set **Region Mode** to **Create a Region for Each Part**.
   b) Set **Boundary Mode** to **Create a Boundary for Each Part Surface**.
   c) Clear the **Create Interfaces From Contacts** checkbox.
   d) Click **Apply**, and then **Close**.

To be able to create an overset interface between the two regions, assign the `Overset Mesh` type to at least one boundary in the **Overset** region. This boundary type applies to all the boundary surfaces of the **Overset** region which are embedded inside the **Background** region, and do not form part of the boat hull. Boundaries of the **Overset** region which are coplanar with boundaries of the **Background** region (in this case, `Symmetry plane`) must be assigned the same boundary type.

Assign the boundary of each region to the correct type:

4. Edit the **Regions** > **Background** > **Boundaries** node and set the boundary type for each surface as follows:

Node	Boundary Type
Back	**Symmetry Plane**
Bottom	**Velocity Inlet**
Inlet	**Velocity Inlet**
Outlet	**Pressure Outlet**
Symmetry	**Symmetry Plane**
Top	**Velocity Inlet**

5. Edit the **Regions** > **Overset** > **Boundaries** node and set the boundary type for each surface as follows:

Node	Boundary Type
Boundaries	**Overset Mesh**
Lifeboat_Wall	**Wall**
Symmetry	**Symmetry Plane**

The simulation tree appears as follows:



6. Save the simulation .

Selecting Physics Models

A physics continuum must be present, and assigned to both regions, before the **Overset** region can be coupled to the **Background** region. You create the physics continuum and activate the several physical models that are required to simulate the forces acting on the hull.

The simulation models the behavior of two fluids (air and water) within the same continuum, and uses the Volume of Fluid model to do so. As there are two fluids in different phases, the Eulerian Multiphase model is activated, and the effect of gravity acting on both is included using the Gravity model. The flow is assumed to be laminar since you are only interested in simulating water entry and impact forces. Simulation of

real lifeboat cases over a longer distance would require a turbulence model.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional Gradients** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase**
Eulerian Multiphase Model	**Volume of Fluid (VOF)**
Viscous Regime	**Laminar**
Optional Models	**Gravity** **VOF Waves**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Coupling the Regions

You must create an interface between the background and the overset regions.

To create the Overset Interfaces:

1. Multi-select the **Regions** > **Background** and **Overset** nodes.
2. Right-click either selected node and select **Create Interface** > **Overset Mesh**.



A new **Interfaces** node appears in the simulation tree, with **Overset Mesh 1** as child node. This volume-type interface provides coupling of solutions on the domains using automatically generated sets of acceptor cells on one mesh and donor cells on the other mesh. Variable values at donor cells express variable values at acceptor cells through interpolation.

The interpolation method can be set as follows:

3. Select the **Interfaces** > **Overset Mesh 1** > **Physics Conditions** > **Interpolation Option** node and set *Option* to `Linear`.



4. Save the simulation.

Generating the Mesh

Proceed with the mesh creation for the case.

1. Right-click the **Continua** node and select **New > Mesh Continuum**.
2. Rename **Mesh 1** to **Background**

Select the mesh models for the **Background** mesh.

3. Right-click **Continua > Background > Models** and choose **Select Meshing Models...**

In the *Background Model Selection* dialog:

4. Select **Surface Remesher** from the *Surface Mesh* group box.
5. Select **Trimmer** from the *Volume Mesh* group box.
6. Click **Close**.

Create a second mesh continuum that is called **Overset** by copying the **Background** mesh.

7. Right-click the **Continua > Background** node and select **Copy**.
8. Right-click the **Continua** node and select **Paste**.
9. Rename **Copy of Background** to **Overset**

The simulation tree appears as follows, with all **Models** nodes expanded:



Each mesh is assigned to its corresponding region.

10. Select **Regions > Background** and make sure that *Mesh Continuum* is set to `Background`.
11. Select **Regions > Overset** and set *Mesh Continuum* to `Overset`.



When generating the mesh, pay particular attention to the following:

- To minimize errors in interpolating variables between the two meshes, use the same order of magnitude of grid density in the overlapping region of the **Overset** and **Background** meshes.

- To align the overset mesh with the flow around the lifeboat as it moves through water, generate the mesh with the lifeboat in horizontal position. Rotate the region to its initial position for simulation after the mesh creation.
- The grid must be finer around the boat hull and the water-air interface to capture relevant flow detail.

Set up volumetric controls for each mesh to help you achieve these goals. First create blocks inside which finer meshes are generated.

12. Right-click the **Geometry** > **Parts** node and select **New Shape Part** > **Block**.
13. Enter the following coordinates:



14. Click **Create**, and then **Close**.
15. Rename **Geometry** > **Parts** > **Block** to **Overlap**
16. Using the same procedure, create another block that is called **Water Surface** with the following coordinates:



17. Create another block that is called **Boat** with the following coordinates:



Set the base size for the **Background** mesh continuum and apply volumetric controls:

18. Select the **Continua** > **Background** > **Reference Values** > **Base Size** node and set *Value* to `5.0 m`
19. Right-click the **Continua** > **Background** > **Volumetric Controls** node and select **New**.
20. Rename **Volumetric Control 1** to **Overlap**
21. Select the **Volumetric Controls** > **Overlap** node and set *Parts* to **Parts** > **Overlap**.
22. To close the dialog, click **OK**.



23. Select the **Volumetric Controls > Overlap > Mesh Conditions > Surface Remesher** node and activate **Customize surface remesher**.
24. Select **Overlap > Mesh Values > Custom Size > Relative Size** and set *Percentage of Base* to `25.0`, resulting in an *Absolute Size* of `1.25 m`.

    This setting makes sure that the overset and background meshes are of similar order of magnitude where they overlap



25. Select the **Volumetric Controls** > **Overlap** > **Mesh Conditions** > **Trimmer** node and activate *Customize isotropic size*.

Create the volumetric control for the zone along the expected water surface:

26. Create a Volumetric Control and rename it **Water Surface**
27. Assign `[Water Surface]` to the *Parts* of this volumetric control.
28. Activate the **Custom Size** selection option for the Surface Remesher and set the **Relative Size** to a *Percentage of Base* of `100.0`
29. Select the **Water Surface** > **Mesh Conditions** > **Trimmer** node and activate *Customize anisotropic size*.
30. Select the **Water Surface** > **Mesh Values** > **Trimmer Anisotropic Size** node and activate *Custom X size*, *Custom Y size*, and *Custom Z size*.
31. Select the **Trimmer Anisotropic Size** > **Relative Z Size** node and set *Percentage of Base* to `50.0`



32. Set the **Percentage of Base** for the *Relative X size* and *Relative Y size* nodes to `100.0`

Align the mesh to the background region and change the mesh growth settings from the volumetric controls:

33. Select the **Background** > **Models** > **Trimmer** node and activate *Do mesh alignment*.
34. Select the **Background** > **Reference Values** > **Mesh Alignment Location** node and set *Location* to `[-125.0, 0.0 0.0] m`
35. Select the **Background** > **Reference Values** > **Template Growth Rate** node and set *Default Growth Rate* to `Slow`.
36. Set *Boundary Growth Rate* to `Slow`.

Set the base size for the **Overset** mesh continuum and add a volumetric control to refine the region around the lifeboat hull.

37. Select the **Continua** > **Overset** > **Reference Values** > **Base Size** node and set *Value* to `1.0 m`
38. Create a Volumetric Control for the **Overset** mesh continuum and rename it **Boat**
39. Assign `[Boat]` to the *Parts* of this volumetric control.
40. To activate the **Custom Size** selection option, select the **Volumetric Controls** > **Boat** > **Mesh Conditions** > **Surface Remesher** node and activate **Customize surface remesher**.
41. Select **Boat** > **Mesh Values** > **Custom Size** > **Relative Size**.
42. Set the **Relative Size** to a *Percentage of Base* of `50.0` to obtain an *Absolute Size* of `0.5 m`.

Generate the mesh:

43. Click the (**Generate Volume Mesh**) button in the toolbar or select the **Generate Volume Mesh** item in the **Mesh** menu.

The mesh can be visualized by creating a Mesh Scene:

44. Right-click **Scenes** and select **New Scene** > **Mesh**.
45. Click the (**Save-Restore-Select Views**) button and select **Views** > **-Y** > **Up +Z** to display the mesh.



The current mesh settings have been chosen for the tutorial to run in about 30 minutes on a 2.5 GHz processor. If you want to resolve wave propagation accurately, use a finer mesh along the water surface.

46. Save the simulation .

Transforming the Overset Region

Translate the lifeboat nearer to the water surface to decrease the simulation time and give it an initial tilt with a horizontal angle of 35.

1. Right-click **Regions** > **Overset** and select **Transform** > **Translate**.

   The *Translate Regions* dialog appears.
2. Set the **Translation Vector** to `[0 m, 0 m, -26 m]` as shown below, and click **Apply** and then **Close**.



Before rotating the lifeboat, specify the coordinate system in which the boat is rotated. Create a coordinate system whose origin is taken as the center of mass of the boat after the translation.

3. Right-click the **Tools > Coordinate Systems > Laboratory > Local Coordinate Systems** node and select **New > Cartesian**.

   The *Create Cartesian Coordinate System* panel appears.

4. Set Origin to `[-8.22, 0.0 -24.329] m, m, m`



5. Click **Create** and then **Close**.
6. Rename the **Cartesian 1** node to **Boat Rotation**

Rotate the **Overset** region:

7. Right-click **Regions** > **Overset** and select **Transform** > **Rotate**.

In the *Rotate Regions* dialog,

8. Set the **Angle** to `35 deg`
9. Set the **Axis Vector** to `[0 m, 1 m, 0 m]`
10. Set the **Coordinate System** to `Laboratory\Boat Rotation`.



11. Click **Apply** and then **Close**.

The mesh now appears as follows:



12. Save the simulation.

Creating Eulerian Phases

The materials corresponding to each of the mixture phases are defined as **Eulerian Phases**.

Create the **Eulerian Phases** required for the simulation and set their material properties:

1. Right-click the **Continua** > **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and select **New**.

   A new node, **Phase 1**, is created.
2. Rename the **Phase 1** node to `Water`

Define the models that determine the behavior of the *Water* phase.

3. Right-click the **Eulerian Phases** > **Water** > **Models** node and choose **Select models...**

In the *Phase Model Selection* dialog:

4. Select **Liquid** from the *Material* group box.
5. Select **Constant Density** from the *Equation of State* group box.

Repeat the above steps to create a second Eulerian phase.

6. Rename the second phase to **Air**
7. Choose **Gas** rather than **Liquid** for the material model. Select **Constant Density** as above.

In both instances, you are using the default material appropriate to each phase (water for the liquid and air for the gas), so no further changes are necessary to material properties. The model selections can be seen by expanding the **Water** and **Air** nodes within the **Eulerian Phases** parent node.



8. Save the simulation.

Setting the DFBI Motion and 6-DOF Body

The 6-DOF motion solver is activated when creating the relevant DFBI motion object and assigning it to the fluid region. For this tutorial, you require the DFBI rotation and translation motion.

1. Open the **Tools** > **Motions** node.

   The default **Stationary** motion is shown.
2. Right-click the **Motions** node and select **New** > **DFBI Rotation and Translation**.

A new **DFBI Rotation and Translation** node is added to the **Motion** manager.

Select this motion within the **Motion Specification** node for the Overset region.

3. Select the **Regions** > **Overset** > **Physics Values** > **Motion Specification** node and set *Motion* to **DFBI Rotation and Translation**.



Create the **Boat** body and set its properties.

4. Right-click the **DFBI > 6-DOF Bodies** node and select **New**.
5. Rename **Body 1** to **Boat**
6. Select the **DFBI > 6-DOF Bodies > Boat** node and click the ... (**Custom Editor**) button for *Body Surface*.

   The *Boat - Body Surface* dialog appears.
7. Select **Overset > Boundaries > Lifeboat_wall**.



8. To close the dialog, click **OK**.
9. Set the *Body Mass* to `10000.0 kg`.
10. Keep the default values for all other properties.



In this simulation, the body is allowed to rotate about the y-axis only, to simulate motion in the head waves. This means that the moment of inertia about the y-axis must be defined.

11. Select the **DFBI > 6-DOF Bodies > Boat > Initial Values > Moment of Inertia** node and do the following:

    a) Set *Diagonal Components* to `[16000.0, 16000.0, 16000.0] kg-m^2`

    b) Activate **Use Center Of Mass**.

    The moments of inertia are evaluated with respect to the center of mass.



12. Select the **Initial Values** > **Center of Mass** node and set *Position* to `[-8.22, 0.0, -24.329] m`.
13. Select the **Initial Values** > **Velocity** node and set *Value* to `[0.0, 0.0, -22.0] m/s`.

Although the 6-DOF body has now been created with the Free Motion option, it is still necessary to specify the manner in which the body can move. In this case, the body is free to translate in the X direction (surge), Z direction (heave), and to rotate about the Y axis (pitch).

14. Select the **DFBI** > **6-DOF Bodies** > **Boat** > **Free Motion** node and activate the *X Motion*, *Z Motion*, and *Y Rotation* properties.



15. Save the simulation.

Defining the VOF Waves

The VOF Wave model generates a number of field functions that are based on the wave parameters that you define. These field functions can be used as boundary and initial conditions in appropriate nodes on the object tree.

To create the VOF Wave:

1. Right-click the **Continua** > **Physics 1** > **Models** > **VOF Waves** > **Waves** node and select **New** > **Flat**.
2. Rename the **FlatVofWave 1** node to `Flat Wave`
3. Select the **Flat Wave** node and set *Point on Water Level* to `[0.0, 0.0, -30] m`



Note: The *Light Fluid Density* and *Heavy Fluid Density* properties must be the same as those specified for the Eulerian phases **Water** and **Air**.

Setting Initial Conditions

The initial conditions are set using field functions that are associated with the previously defined VOF Wave. Initial conditions that are set are the free-surface wave profile, the velocity profile within the phases and the hydrostatic pressure.

1. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Volume Fraction** node and set *Method* to `Composite`.

The **Volume Fraction** node now contains a **Composite** node in response to this change of the *Method* property.

2. Select the **Volume Fraction** > **Composite** > **Water** node and set *Method* to `Field Function`.
3. Select the **Water** > **Field Function** node and set *Scalar Function* to `Volume Fraction of Heavy Fluid of Flat Wave`
4. Repeat these steps with the **Volume Fraction** > **Air** node, but set *Scalar Function* to `Volume Fraction of Light Fluid of Flat Wave`.
5. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Velocity** node and set *Method* to `Field Function`.
6. Select the **Velocity** > **Field Function** node and set *Vector Function* to `Velocity of Flat Wave`.
7. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Pressure** node and set *Method* to `Field Function`.
8. Select the **Pressure** > **Field Function** node and set Scalar Function to `Hydrostatic Pressure of Flat Wave`.



9. Save the simulation.

Setting the Boundary Conditions

The boundary types have been correctly set previously, so it only remains to set the conditions corresponding to each type.

There are four boundaries for this geometry:

- No-slip wall boundary on the hull of the boat
- Symmetry boundary on the Z-X plane at Y = 0.0 and Y = 50.0
- Pressure outlet on the right face
- Velocity inlet on all remaining boundaries

Default conditions are suitable for the symmetry plane and the no-slip wall boundary on the boat hull. Wave conditions are applied to the pressure outlet and velocity inlet. Nothing is prescribed at the Overset Mesh boundaries.

Contents

Inlet Boundary
Outlet Boundary
Bottom Boundary
Top Boundary

Inlet Boundary

You specify the velocity and the volume fraction at the inlet boundary.

1. Open the **Regions** > **Background** > **Boundaries** > **Inlet** > **Physics Conditions** node.
2. Select the **Physics Conditions** > **Velocity Specification** node and set *Method* to `Components`.
3. Select the **Inlet** > **Physics Values** > **Velocity** node and set *Method* to `Field Function`.
4. Select the **Velocity** > **Field Function** node and set the *Vector Function* to `Velocity of Flat Wave`.
5. Continuing within the **Background** > **Boundaries** > **Inlet** node, select the **Physics Values** > **Volume Fraction** node.
6. Set *Method* to `Composite`.
7. Select the **Volume Fraction** > **Composite** > **Water** node and set *Method* to `Field Function`.
8. Select the **Composite** > **Water** > **Field Function** node and set *Scalar Function* to `Volume Fraction of Heavy Fluid of Flat Wave`.

9. Repeat the last two steps for the **Volume Fraction** > **Composite** > **Air** node, but set *Scalar Function* to `Volume Fraction of Light Fluid of Flat Wave`.



Outlet Boundary

At the oulet boundary, you set the pressure and the volume fraction to correspond to the outgoing wave by using field functions.

1. Open the **Regions** > **Background** > **Boundaries** > **Outlet** > **Physics Values** > **Pressure** node and set *Method* to `Field Function`.
2. Select the **Physics Values** > **Pressure** > **Field Function** node and set *Scalar Function* to `Hydrostatic Pressure of Flat Wave`.
3. Select the **Physics Values** > **Volume Fraction** node and set *Method* to `Composite`.
4. Select the **Volume Fraction** > **Composite** > **Water** node and set *Method* to **Field Function**.

5. Select the **Composite** > **Water** > **Field Function** node and set *Scalar Function* to `Volume Fraction of Heavy Fluid of Flat Wave`.
6. Repeat the last two steps for the **Volume Fraction** > **Composite** > **Air** node, but set *Scalar Function* to `Volume Fraction of Light Fluid of Flat Wave`.

Bottom Boundary

The following boundary conditions are appropriate for deep water conditions. A bottom boundary can also be a wall for shallow waters.

1. Open the **Regions** > **Background** > **Boundaries** > **Bottom** > **Physics Conditions** node.
2. Select the **Physics Conditions** > **Velocity Specification** node, and set *Method* to `Components`.
3. Select the **Bottom** > **Physics Values** > **Velocity** node and set *Method* to `Field Function`.
4. Select the **Velocity** > **Field Function** node and set *Vector Function* to `Velocity of Flat Wave`.
5. Continuing within the **Background** > **Boundaries** > **Bottom** node, select the **Physics Values** > **Volume Fraction** > **Constant** node and set *Value* to `[1.0, 0.0]`

Top Boundary

You specify the velocity and the volume fraction at the top boundary.

1. Open the **Regions** > **Background** > **Boundaries** > **Top** > **Physics Conditions** node.
2. Select the **Physics Conditions** > **Velocity Specification** node, and set *Method* to `Components`.
3. Select the **Top** > **Physics Values** > **Velocity** node and set *Method* to `Field Function`.
4. Select the **Velocity** > **Field Function** node and set *Vector Function* to `Velocity of Flat Wave`.
5. Continuing within the **Background** > **Boundaries** > **Top** node, select the **Physics Values** > **Volume Fraction** > **Constant** node and set *Value* to `[0.0, 1.0]`

Setting Up Solver Parameters and Stopping Criteria

As this simulation is transient, it is necessary to set the time step, the maximum number of inner iterations that are permitted in each time step, and the overall physical time for which a solution is obtained.

The time step value is selected as 0.0025 s so that, for any time step, the overset mesh moves by a distance smaller than the height of the smallest cell in the **Background** region. A finer mesh would require an even smaller time step.

1. Select the **Solvers** > **Implicit Unsteady** node and set the *Time-Step* to `0.0025 s`.

A second order time discretization is required for simulations requiring accurate wave propagation. If modeling wave propagation is not important, a first order time method can be used. Second order methods require smaller time steps.

2. Select the **Stopping Criteria** > **Maximum Inner Iterations** node and make sure that *Maximum Inner Iterations* is set to `5`.
3. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Max Physical Time* to `0.5 s`.
4. Select the **Stopping Criteria** > **Maximum Steps** node and deactivate *Enabled*.
5. Save the simulation.

Setting Scenes to Visualize the Solution

Create scenes to visualize the results as the simulation progresses.

To create a scalar scene:

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
2. Rename the **Scalar Scene 1** node to **Volume Fraction of Water**
3. Click the (**Save-Restore-Select Views**) button and select **Views** > **-Y** > **Up +Z**.
4. Select the **Scenes** > **Volume Fraction of Water** > **Displayers** > **Scalar 1** > **Parts** node and click the custom editor button to the right of the *Parts* property.

The *Parts* dialog appears.

5. Select the **Regions** > **Background** > **Symmetry**, **Regions** > **Overset** > **Lifeboat_Wall** and **Regions** > **Overset** > **Symmetry** nodes.



6. Click **OK** to close the dialog.
7. Expand **Scenes** > **Volume Fraction of Water** > **Displayers** > **Scalar 1**and select **Scalar Field** and set **Volume Fraction** > **Water** for the *Function* property.

To display overset cells, create an additional scene:

8. Click **Scenes** > **Volume Fraction of Water** and while holding the left mouse button, press `CTRL` and drag the **Volume Fraction of Water** node onto the **Scenes** node.
9. Rename the new node to **Overset Cells**
10. Expand **Scenes** > **Overset Cells** > **Displayers** > **Scalar 1** and select **Scalar Field** and set `Overset Cell Type` for the *Function* property.

You are now ready to run the simulation.

11. Save the simulation.

Initializing and Running the Simulation

To verify that the initial conditions applied to the regions are correct, STAR-CCM+ allows you to perform just solver initialization to start with and then view the results using an appropriate scene.

To initialize:

1. Click the (**Initialize Solution**) button in the toolbar or use the **Solution** > **Initialize Solution** menu item.
2. Select the *Volume Fraction of Water* display to see the initialization result.



3. Click the (**Run**) button in the toolbar.
4. The evolution of the volume of fraction of water can be visualized while the simulation is running.



5. After the run is finished, save the simulation.

Summary

This tutorial introduced the following STAR-CCM+ features:

- Using the Overset Grids model.
- Activating the 6-DOF model.
- Creating a VOF wave and using it for initial and boundary conditions.

Marine Resistance Prediction: KCS Hull with a Rudder

This tutorial demonstrates the workflow for setting up a resistance prediction simulation for a marine application.

The geometry that is used is the KRISO Container Ship (KCS). This geometry was used as part of a series of simulations that CD-adapco performed for the Gothenburg Workshop 2010 on Numerical Ship Hydrodynamics. The STAR-CCM+ simulation results were compared with experimental towing tank data.

The KCS ship hull with rudder is placed within a virtual towing tank as shown below:



The simulation is performed in model scale. The hull length is approximately 7.7 m (length between perpendiculars L_{pp}= 7.2786 m) with a scale factor of 31.6. The coordinate origin is at aft perpendicular

and still water surface. The still water surface is at 0.341772 m above keel.

Only half of the geometry is modeled due to symmetry conditions. The computational domain extends from:

- -18 m to 18 m in the flow direction
- -18 m to 9 m in the vertical direction
- 0 m to 18 m in the lateral direction

The following boundary conditions are applied:

- Symmetry plane at the hull midplane
- Symmetry plane at the side of the tank
- Hydrostatic pressure corresponding to the undisturbed water surface at the downstream boundary of the tank
- Inlets with prescribed velocity and volume fraction at the upstream, top, and bottom boundaries.

To obtain resistance data experimentally, the ship hull is towed through the towing tank. For the equivalent simulation, the ship remains in a fixed position, and an inlet velocity is prescribed at the upstream inlet boundary. This inlet velocity produces the relative movement. The ship is free to heave and pitch. The simulation starts from a flat water surface with the bottom of the hull being parallel to the free surface. The computation is performed for a Froude number of Fn = 0.2599.

You can work through the tutorial in three major steps:

- Mesh Generation
- Physical and Numerical Setup
- Visualization and Data Analysis

Contents

Mesh Generation
Physical and Numerical Setup
Visualization and Data Analysis
Summary
Bibliography

Mesh Generation

In order to achieve accurate resistance predictions, you are recommended to use trimmed hexahedral meshes with local refinements and prism layers along the walls of the wetted surface. In order to refine the computational mesh only where it is necessary, you create volumetric controls with particular shapes. A particular cell size is then assigned to these volumetric controls.

Contents

Importing the Geometry
Creating the Towing Tank
Assigning Parts to Regions
Generating the Volume Mesh

Importing the Geometry

The meshing process starts from a CAD geometry in Parasolid format.

To import the hull of the ship, including the rudder:

1. Create a directory for the tutorial called, `hullWithRudder`.
2. Start up STAR-CCM+ and create a Simulation by clicking **File > New Simulation...**
3. Select **File > Import > Import Surface Mesh...** Navigate to the `doc/startutorialsdata/motion/data` subdirectory of your STAR-CCM+ installation directory. Select the file `KCS_geometry.x_t` from the selection list and click **Open**.
4. In the *Import Surface Options* dialog, click **OK**.

   The imported geometry appears in the *Geometry Scene 1* window. The imported geometry resides under the **Geometry > Parts** node.
5. Open the **Parts > KCS_Geometry > Surfaces** node and inspect the imported surfaces.
6. Select **File > Save As** and save the simulation under the name `KCS_resistancePrediction.sim`.

Creating the Towing Tank

The imported geometry contains only the ship hull with the rudder. To perform a resistance prediction analysis, create a towing tank around the ship.

For the towing tank, create a block shaped geometry around the ship, and then subtract the ship hull from the block. The resulting volume is a contiguous and closed virtual towing tank that contains the ship.

To create a block shaped part forming the towing tank:

1. Right-click the **Geometry** > **Parts** node and select **New Shape Part** > **Block**.
2. In the *Create Block Part* dialog, enter the following Maximum and Minimum Coordinates for the block:
   a) Set **Corner 1** to [`-18 m, 0.0 m, -18.0 m`].
   b) Set **Corner 2** to [`18.0 m, 18.0 m, 9.0 m`].



3. Click **Create** and then **Close**.
4. Rename the newly created **Geometry** > **Parts** > **Block** node to `Tank`.
   The created block consists of one surface encompassing the entire tank.

In order to define different boundary conditions on each tank surface, split the tank surface:

5. Right-click the **Geometry** > **Parts** > **Tank** > **Surfaces** > **Block Surface** node and select **Split by Patch...**
6. In the *Split Part Surface by Patch* dialog, select **725** in the *Part Surface Patches* group box.
   The corresponding surface is highlighted in the *Graphics* window.
7. In the **Part Surface Name** text box, enter `Top`.
8. Click **Create**.

9. Split and rename the remaining block surfaces as described in the previous steps. Assign the following names to the different patches:

Patch	Part Surface Name
726	`Symmetry`
727	`Outlet`
728	`Side`
729	`Bottom`

Rename the final Patch 6 in the simulation tree:

10. Right-click the **Geometry** > **Parts** > **Tank** > **Surfaces** > **Block Surface** node and **Rename**... it `Inlet`.

To create the virtual towing tank, perform a Boolean subtract mesh operation:

11. Multi-select the **Tank** node and the **KCS_Geometry** node.
12. Right-click one of the highlighted items and select**Create Mesh Operation** > **Boolean** > **Subtract Parts.**
13. In the *Create Subtract Operation* dialog, click **OK**.

    The **Operations** > **Subtract** node is created.
14. Select the **Operations** > **Subtract** node and set *Target Part* to **Tank**.
15. Right-click the **Operations** > **Subtract** node and select **Execute**.
16. Rename the newly created **Parts** > **Subtract** part to `Virtual Towing Tank`.

To visualize the virtual towing tank:

17. Right-click in the empty space of the *Graphics* window and select **Hide All Parts**.
18. Select the **Geometry** > **Parts** > **Virtual Towing Tank** node and drag and drop it into the open geometry scene in the *Graphics* window.
19. In the pop-up menu, select **Part Displayers** > **Add to Geometry 1.**
20. Color the part surfaces.

    a) Select the **Scenes** > **Geometry Scene 1** > **Displayers** > **Geometry 1** node.

    b) In the *Properties* window, set *Color Mode* to **Distinguish Inputs**.

You can now inspect the combined geometry.

Assigning Parts to Regions

Before you create the volume mesh, transfer the virtual towing tank geometry to the regions level.

To assign the virtual towing tank part to the regions level:

1. Right-click the **Geometry > Parts > Virtual Towing Tank** node and select **Assign Parts to Regions...**
2. In the *Assign Parts to Regions* dialog:
   a) Select **Create a Region for Each Part.**
   b) Select **Create a Boundary for Each Part Surface.**
   c) Deactivate **Create Interfaces from Contacts.**
   d) Click **Apply** and then **Close**.



To view the transferred geometry in the simulation tree:

3. Expand the **Regions** node.
   You can see the newly created **Virtual Towing Tank** node.
4. Expand the **Virtual Towing Tank > Boundaries** node.
   You can see all of the transferred boundary surfaces.

Generating the Volume Mesh

Resistance analyses are generally performed using a trimmed volume mesh with prism layers on the wetted surface of the hull. Using trimmed cells means that the mesh is aligned with the undisturbed free surface.

To select the required mesh models:

1. Right-click the **Continua** node and select **New > Mesh Continuum**.
2. Right-click the **Mesh 1 > Models** node and click **Select Meshing Models...**.
3. In the **Mesh 1 Model Selection** dialog:
   a) Select **Surface Remesher** from the *Surface Mesh* group box.
   b) Select **Trimmer** from the *Volume Mesh* group box.
   c) Select **Prism Layer Mesher** from the *Optional Models* group box.
   d) Click **Close**.

To adjust the Prism Layer meshing model:

4. Select the **Mesh 1 > Models > Prism Layer Mesh** node.
5. In the *Properties* window:
   a) Set the **Minimum Thickness Percentage** to `1.0`.
   b) Set the **Boundary March Angle** to `75.0`.

To specify global mesh settings:

6. Expand the **Mesh 1 > Reference Values** node.



7. Set the meshing parameters according to the following table:

Node	Property	Setting
Base Size	**Value**	`0.1125 m`
Automatic Surface Repair > Minimum Proximity	**Minimum Proximity**	`0.001`
Maximum Cell Size > Relative Size	**Percentage of Base**	`1600`
Number of Prism Layers	**Number of Prism Layers**	`6`
Prism Layer Thickness	**Size type**	**Absolute**
Prism Layer Thickness > Absolute Size	**Value**	`0.02 m`
Surface Curvature > Basic Curvature	**# Pts./circle**	`36`
Surface Size > Relative Minimum Size	**Percentage of Base**	`6.25`

Node	Property	Setting
Surface Size > Relative Target Size	**Percentage of Base**	`50`
Template Growth Rate	**Default Growth Rate**	**Slow**
	Boundary Growth Rate	**Slow**

In addition to the globally applied mesh settings, you specify local mesh settings on the boundary surfaces.

To specify local mesh settings on boundaries:

8. Expand the **Regions > Virtual Towing Tank > Boundaries** node:
   a) Select **KCS_Geometry.Deck > Mesh Conditions > Custom Surface Size** and activate *Custom surface size.*
   b) Select **KCS_Geometry.Deck > Mesh Conditions > Customize Prism Mesh** and set *Customize Prism Mesh* to **Disable**.
   c) Select **KCS_Geometry.Deck > Mesh Values > Surface Size > Relative Minimum Size** and set *Percentage of Base* to `100`.
   d) Select **KCS_Geometry.Deck > Mesh Values > Surface Size > Relative Target Size** and set *Percentage of Base* to `200`.
9. Similarly, set the parameters for the other boundaries as listed in the following table:

Boundary Name	Mesh Conditions		Mesh Values	
	Custom Surface Size	**Customize Prism Mesh**	**Relative Minimum Size**	**Relative Target Size**
KCS_Geometry.Hull	Activate	Use Default Values	25	25
KCS_Geometry.Rudder	Activate	Use Default Values	6.25	12.5
KCS_Geometry.Rudder Edge	Deactivate	Disable		
KCS_Geometry.Shaft	Deactivate	Disable		
KCS_Geometry.Stern	Deactivate	Disable		

Boundary Name	Mesh Conditions		Mesh Values	
Tank.Bottom	Activate	Use Default Values	1600	1600
Tank.Inlet	Activate	Use Default Values	1600	1600
Tank.Outlet	Activate	Use Default Values	1600	1600
Tank.Side	Activate	Use Default Values	1600	1600
Tank.Symmetry	Activate	Use Default Values	1600	1600
Tank.Top	Activate	Use Default Values	1600	1600

In order to refine sharp corners or edges of your geometry, for example at the bow, or to improve the resolution of flow features, for example flow separation or the Kelvin wake, you can apply volumetric refinements to your volume mesh. For this application, it is particularly important to refine the free water surface area, that is, the interface area between water and air. By using trimmed cells, you can take advantage of anisotropic refinement in the normal direction of the free surface area. You define the geometrical shapes and dimensions of these volumetric controls in the form of part shapes.

To generate volumetric refinements:

10. Generate a part shape at the bow:
    a) Right-click the **Geometry** > **Parts** node and select **New Shape Part** > **Block**.
    b) In the *Create Block Part* dialog, enter the following Maximum and Minimum Coordinates for the block:
       - Set **Corner 1** to [`7.17 m, -0.087 m, -0.34 m`].
       - Set **Corner 2** to [`7.54 m, 0.087 m, 0.059 m`].
    c) Click **Create** and **Close**.
    d) Rename the **Block** part as `BOW`.
11. Generate a part shape at the stern:
    a) Right-click the **Geometry** > **Parts** node and select **New Shape Part** > **Block**.

b) In the *Create/Edit Block Part* dialog, enter the following Maximum and Minimum Coordinates for the block:

- Set **Corner 1** to [`-0.12 m, -0.09 m, -0.34 m`].
- Set **Corner 2** to [`0.28 m, 0.05 m, 0.008 m`].

c) Click **Create** and **Close**.

d) Rename the **Block** part as `STERN RUDDER - fine`.

12. Right-click the **Continua** > **Mesh 1** > **Volumetric Controls** node and select **New**.
13. Select the **Volumetric Control 1** node and rename it `BOW, STERN - fine`.
14. Assign the part shapes to the volumetric control.

    a) In the *Properties* window, click the custom editor ... next to *Parts*.

    b) In the *Parts* selection dialog, select **BOW** and **STERN RUDDER - fine**.

    c) Click **OK**.
15. Set the mesh conditions for this volumetric control.

    a) Select the **Volumetric Controls** > **BOW, STERN - fine** > **Mesh Conditions** > **Surface Remesher** node and activate *Customize surface remesher*.

    b) Select the **Volumetric Controls** > **BOW, STERN - fine** > **Mesh Values** > **Custom Size** > **Relative Size** node and, in the *Properties* window, set *Percentage of Base* to `12.5`.



In order to resolve flow features adequately around the hull, this tutorial specifies several part shapes.

To run the macro for creating part shapes and volumetric controls:

16. Select **File** > **Macro** > **Play Macro...**
17. In the *Open* dialog, navigate to the `doc/startutorialsdata/motion/data` subdirectory of your STAR-CCM+ installation directory. Select the file `partShapes.java` from the selection list and click **Open**.

    The macro is executed and creates the required part shapes and volumetric controls.

To view the created part shapes:

18. Expand the **Geometry** > **Parts** node.



To view the created volumetric controls:

19. Expand the **Continua > Mesh 1 > Volumetric Controls** node.



To inspect the mesh settings for the volumetric controls, expand the respective nodes.

Before you create the volume mesh, set the appropriate boundary types. By default, no prism layers are created for inflow, outflow, and symmetry boundaries.

To set the boundary types:

20. Select the **Regions** > **Boundaries** > **Tank.Bottom** node and set *Type* to **Velocity Inlet**.

Similarly, set the boundary types for the other boundaries as listed in the following table:

Boundary Name	Boundary Type
Tank.Inlet	Velocity Inlet
Tank.Outlet	Pressure Outlet
Tank.Side	Symmetry Plane
Tank.Symmetry	Symmetry Plane
Tank.Top	Velocity Inlet

All boundaries with the **KCS_Geometry** prefix remain as **Wall**.

To generate the volume mesh:

21. Select**Mesh > Generate Volume Mesh**.

    The meshing process starts. You can follow the meshing progress in the *Output* window.

To visualize the final volume mesh:

22. Create two planar sections:
    a) Right-click **Derived Parts** and select **New Part > Section > Plane...**
    b) In the *Create Section* dialog:
       - In the Input Parts group box, click **Select** and choose **Regions > Virtual Towing Tank** in the *Select Objects* dialog.
       - Click **OK**.
    c) In the *Plane Parameters* group box:
       - Set **origin** to [`0.0 m, 0.0 m, -0.118 m`].
       - Set **normal** to [`0.0 m, 0.0 m, 1.0 m`].
       - Click **Create**.
    d) For the second planar section, in the *Plane Parameters* group box:
       - Set **origin** to [`0.5 m, 0.0 m, 0.0 m`].
       - Set **normal** to [`1.0 m, 0.0 m, 0.0 m`].
       - Click **Create** and then **Close**.

    Two new **Derived Parts** nodes were created: **Plane Section** and **Plane Section 2**.
23. Click *scene/plot*.
24. Select the **Displayers > Geometry 1 > Parts** node and, in the *Properties* window, click (**Custom Editor**) next to *Parts*.
25. In the *Parts - Parts* dialog, deselect all and click **OK**.
26. Select the **Displayers > Outline 1 > Parts** node and, in the *Properties* window, click (**Custom Editor**) next to *Parts.*
27. In the *Parts - Parts* dialog, select **Regions > Virtual Towing Tank** and click **OK**.
28. Right-click the empty space in the *Graphics* window and select **Apply Representation > Volume Mesh**.
29. Click **Show All Meshes** .
30. Select the **Displayers > Outline 1** node and, in the *Properties* window, deactivate *Mesh*.

The inner trimmed volume mesh is displayed on the two planar sections.



Note the embedded refinement at an angle with respect to the track of the ship. This refinement is created to capture the Kelvin wake that is produced by the ship.

Physical and Numerical Setup

Flow resistance simulation is a complex transient multi-physics problem.

In STAR-CCM+, setting this simulation up involves:

- Initial conditions
- Boundary conditions
- VOF multiphase
- VOF Waves and wave damping
- Dynamic Fluid Body Interaction

- Transient numerical parameters

Contents

Selecting the Physical Models
Defining the Eulerian Phases
Defining the VOF Wave
Setting Initial Conditions
Damping Wave Reflections
Setting Boundary Conditions
Defining Dynamic Fluid Body Interaction (DFBI)
Setting Solver Parameters and Stopping Criteria

Selecting the Physical Models

Select the physical models that are required for this simulation.

1. Return to *simulation*.
2. Right-click the **Continua** > **Physics 1** > **Models** node and choose **Select Models**...
3. In the *Physics 1 Model Selection* dialog:
   a) Select **Implicit Unsteady** from the *Time* group box.
   b) Select **Eulerian Multiphase** from the *Material* group box.
   c) Select **Volume of Fluid (VOF)** from the *Multiphase Model* group box.
   d) Select **Turbulent** from the *Viscous Regime* group box.
   e) Select **K-Epsilon Turbulence** from the *Reynolds-Averaged Turbulence* group box.
   f) From the *Optional Models* group box, select **Gravity, Cell Quality Remediation,** and **VOF Waves**.



g) Click **Close**.

Defining the Eulerian Phases

Using the VOF multiphase model requires that you create Eulerian phases for water and air.

To define the Eulerian phases:

1. Right-click the **Continua** > **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and select **New**.
2. Rename **Phase 1** as `Water`.
3. Right-click the **Eulerian Phases** > **Water** > **Models** node and choose **Select Models...**
4. In the *Water Model Selection* dialog:
   a) Select **Liquid** from the *Material* group box.
   b) Select **Constant Density** from the *Equation of State* group box.
   c) Click **Close**.

5. Right-click the **Continua** > **Physics 1** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and select **New**.
6. Rename **Phase 1** as `Air`.
7. Right-click the **Eulerian Phases** > **Air** > **Models** node and choose **Select Models...**
8. In the *Air Model Selection* dialog:
   a) Select **Gas** from the *Material* group box.
   b) Select **Constant Density** from the *Equation of State* group box.
9. Click **Close**.



Defining the VOF Wave

The free surface water level changes over time during the simulation. STAR-CCM+ provides the VOF Waves model that lets you specify wave initial and boundary conditions. Here, the ship is towed through calm water. Therefore, define a flat VOF Wave.

To define a flat VOF Wave:

1. Right-click the **Continua** > **Physics 1** > **Models** > **VOF Waves** > **Waves** node and select **New** > **Flat.**
2. Select the **FlatVofWave 1** node.
3. In the *FlatVofWave 1 - Properties* window:

a) Set *Current* to [`-2.196, 0.0, 0.0`] m/s.

b) Set *Wind* to [`-2.196, 0.0, 0.0`] m/s.



Setting Initial Conditions

Set the initial conditions for pressure, velocity, and volume fraction.

To set the initial conditions:

1. Set the initial condition for Pressure:

a) Select the **Continua > Physics 1 > Initial Conditions > Pressure** node and set *Method* to **Field Function**.

b) Select the **Pressure > Field Function** node and set *Scalar Function* to **Hydrostatic Pressure of FlatVofWave 1**.



2. Set the initial condition for Velocity:

a) Select the **Continua > Physics 1 > Initial Conditions > Velocity** node and set *Method* to **Field Function**.

b) Select the **Velocity > Field Function** node and set *Vector Function* to **Velocity of FlatVofWave 1**.



3. Set the initial condition for Volume Fraction:
   a) Select the **Continua** > **Physics 1** > **Initial Conditions** > **Volume Fraction** node and set *Method* to **Composite**.
   b) Select the **Volume Fraction** > **Composite** > **Water** node and set *Method* to **Field Function**.
   c) Select the **Water** > **Field Function** node and set *Scalar Function* to **Volume Fraction of Heavy Fluid of FlatVofWave 1**.
   d) Select the **Volume Fraction** > **Composite** > **Air** node and set *Method* to **Field Function**.
   e) Select the **Air** > **Field Function** node and set *Scalar Function* to **Volume Fraction of Light Fluid of FlatVofWave 1**.



Damping Wave Reflections

Unless appropriate measures are applied, wave reflections can occur during a flow resistance simulation. These wave reflections must be

avoided, because they interact with the true wave field and can lead to invalid results.

There are two sources for wave reflections:

- Wave reflections from boundaries
- Wave reflections due to abrupt mesh transitions

The second cause was addressed by changing the *Template Default Growth Rate* from **Fast** (default) to **Slow** during mesh generation (see Generating the Volume Mesh).

To remedy wave reflections from boundaries, STAR-CCM+ provides the VOF Wave Damping functionality.

To activate VOF Wave damping at boundaries:

1. Select the **Regions** > **Virtual Towing Tank** > **Physics Conditions** > **VofWave Damping Option** node and activate **Damping**.
2. Select the **Regions** > **Virtual Towing Tank** > **Physics Values** > **VofWaveDamping Length** > **Constant** node and set *Value* to `9.0 m`.

Setting Boundary Conditions

Leave the boundary conditions for the boundaries labelled with the prefix **KCS_Geometry** at their default wall boundary condition settings.

To set the **Tank** boundary conditions:

1. Define the Tank.Bottom boundary:
   a) Select the **Regions** > **Virtual Towing Tank** > **Boundaries** > **Tank.Bottom** > **Physics Conditions** > **Velocity Specification** node and set *Method* to **Components**.
   b) Select the **Tank.Bottom** > **Physics Values** > **Velocity** node and set *Method* to **Field Function**.
   c) Select the **Tank.Bottom** > **Physics Values** > **Velocity** > **Field Function** node and set *Vector Function* to **Velocity of FlatVofWave 1**.
   d) Select the **Tank.Bottom** > **Physics Values** > **Volume Fraction** node and set Method to Composite.
   e) Select the **Tank.Bottom** > **Physics Values** > **Volume Fraction** > **Composite** > **Water** node and set *Method* to **Field Function**.
   f) Select the **Water** > **Field Function** node and set *Scalar Function* to **Volume Fraction of Heavy Fluid of FlatVofWave 1**.

g) Select the **Tank.Bottom > Physics Values > Volume Fraction > Composite > Air** node and set *Method* to **Field Function**.

h) Select the **Air > Field Function** node and set *Scalar Function* to **Volume Fraction of Light Fluid of FlatVofWave 1**.

2. Set the remaining boundary conditions according to the following tables:

Boundary Name	Physics Conditions	Method
Tank.Inlet	**Velocity Specification**	**Components**
	VofWaveDampingBoundaryOption	**yes**
Tank.Outlet	**VofWaveDampingBoundaryOption**	**yes**
Tank.Side	**VofWaveDampingBoundaryOption**	**yes**
Tank.Symmetry	**VofWaveDampingBoundaryOption**	**No**
Tank.Top	**Velocity Specification**	**Components**
	VofWaveDampingBoundaryOption	**No**

Boundary Name	Physics Values	Method	Function
Tank.Inlet	**Velocity**	**Field Function**	**Velocity of FlatVofWave 1**
	Volume Fraction	**Composite**	
		Water: Field Function	**Volume Fraction of Heavy Fluid of FlatVofWave 1**
		Air: Field Function	**Volume Fraction of Light Fluid of FlatVofWave 1**
Tank.Outlet	**Pressure**	**Field Function**	**Hydrostatic Pressure of FlatVofWave 1**

Boundary Name	Physics Values	Method	Function
	Volume Fraction	**Composite**	
		Water: Field Function	**Volume Fraction of Heavy Fluid of FlatVofWave 1**
		Air: Field Function	**Volume Fraction of Light Fluid of FlatVofWave 1**
Tank.Top	**Velocity**	**Field Function**	**Velocity of FlatVofWave 1**
	Volume Fraction	**Composite**	
		Water: Field Function	**Volume Fraction of Heavy Fluid of FlatVofWave 1**
		Air: Field Function	**Volume Fraction of Light Fluid of FlatVofWave 1**

Defining Dynamic Fluid Body Interaction (DFBI)

The Dynamic Fluid Body Interaction (DFBI) model simulates the motion of the ship according to the forces acting on it induced by the flow. For this simulation, the ship is allowed to move with two degrees of freedom to account for sinkage and trim.

To create the DFBI motion:

1. Right-click the **Tools > Motions** node and select **New > DFBI Rotation and Translation**.

Assign this motion to the Virtual Towing Tank region:

2. Select the **Regions > Virtual Towing Tank > Physics Values > Motion Specification** node and set *Motion* to **DFBI Rotation and Translation**.

   A new **DFBI** node appears in the simulation tree.



To define the 6-DOF body properties of the ship:

3. Right-click the **DFBI > 6-DOF Bodies** node and select **New**.
4. Rename Body 1 as `Ship`.
5. Select the **6-DOF Bodies > Ship** node.
6. In the *Ship - Properties* window:

   a) Click the custom editor ... next to *Body Surface* and select all boundaries with the prefix **KCS_Geometry**.

   b) Set *Body Mass* to `823.0451 kg`.

   c) Set *Release Time* to `1.0 s`.

   d) Set *Ramp Time* to `5.0 s`.



To define the initial conditions of the 6-DOF body:

7. Expand the **6-DOF Bodies > Ship > Initial Values node.**
   a) Select the **Center of Mass** node and set *Position* to `[3.53151899, 0.0, -0.1113924] m`.
   b) Select the **Moment of Inertia** node and set *Diagonal Components* to `[136.7, 2725.117, 2725.117] kg/m²`.
   c) Activate **Use Center of Mass**.
8. Create a new local initial coordinate system at the center of mass of the ship:
   a) Right-click **Tools > Coordinate Systems > Laboratory > Local Coordinate Systems** and select **New > Cartesian**.
   b) In the *Create Cartesian Coordinate System* dialog, set **Origin** to [`3.53151899 m, 0.0 m, -0.1113924 m`].
   c) Click **Create** and then **Close**.
   d) Rename **Cartesian 1** as `Initial COS`.
9. Select the **6-DOF Bodies > Ship > Initial Values node > Orientation** node and set *Coordinate System* to **Laboratory > Initial COS**.

The creation of the new 6-DOF body automatically creates a body local coordinate system, here called, **Ship-Csys**. This body local coordinate system is located at the center of mass of the body and moves with the body as the simulation progresses.



To define the free motion of the ship:

10. Select the **DFBI > 6-DOF Bodies > Ship > Free Motion** node.
11. In the *Free Motion - Properties* window:
    a) Activate *Z Motion* to allow motion in the Z direction (sinkage).
    b) Activate Y Rotation to allow rotation around the Y axis (trim).



Setting Solver Parameters and Stopping Criteria

A resistance prediction analysis is a transient simulation for which you must specify the time-step, the number of inner iterations per time-step, and the physical time. Additionally, you can set numerical parameters such as under-relaxation factors to enhance convergence per time-step.

To set the solver parameters:

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time Step* to `0.04 s`.
2. Select the **Solvers** > **6-DOF Solver** node and set *Maximum Number Of Iterations* to `5`.
3. Select the **Solvers** > **Segregated Flow** > **Velocity** node and set *Under-Relaxation Factor* to `0.7`.
4. Select the **Solvers** > **Segregated Flow** > **Pressure** node and set *Under-Relaxation Factor* to `0.4`.
5. Select the **Solvers** > **Segregated VOF** node and set *Under-Relaxation Factor* to `0.8`.

To set the stopping criteria:

6. Select the **Stopping Criteria** > **Maximum Inner Iterations** node and set *Maximum Inner Iteration* to `10`.

7. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `200.0 s`.
8. Select the **Stopping Criteria** > **Maximum Steps** node and deactivate *Enabled*.
9. Save the simulation.

Visualization and Data Analysis

Here, you create several plots and graphs for visualizing the results.

The following items are analyzed:

- Visualizing the Propagation of the Free Water Surface
- Visualizing the Wave Pattern
- Monitoring and Plotting Resistance Data
- Monitoring and Plotting Trim and Sinkage Data

Contents

Visualizing the Propagation of the Free Water Surface
Visualizing the Wave Pattern
Monitoring and Plotting Resistance Data
Monitoring and Plotting Trim and Sinkage Data
Running the Simulation
Visualizing the Results
Comparing with Experimental Results

Visualizing the Propagation of the Free Water Surface

You can plot the free water surface to get a good visual representation of the effect of hull resistance on the surrounding flow.

To visualize the propagation of the free water surface:

1. Create a contour plot of the volume fraction on the symmetry plane of the domain:
   a) Right-click the **Scenes** node and select **New Scene** > **Scalar**.
   b) Rename **Scalar Scene 1** as `FreeSurfaceHull`.
   c) Click **Scene/Plot**.
   d) Select the **Displayers** > **Scalar 1** node and set *Contour Style* to **Smooth Filled**.

e) Select the **Scalar 1** > **Parts** node and set *Parts* to **Regions** > **Virtual Towing Tank** > **Tank.Symmetry**.

f) Select the **Scalar 1** > **Scalar Field** node and set *Function* to **Volume Fraction** > **Volume Fraction of Air**.

2. Click anywhere in the displayed scene and then press `s` on your keyboard to obtain a side view of the ship and symmetry plane.

Visualizing the Wave Pattern

You can plot the elevation of the free water surface to get a good visual representation of the wake that the hull produces as it moves through the water.

To create a contour plot of the elevation of the water surface:

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
2. Rename **Scalar Scene 1** as `Waves`.
3. Create an iso-surface for the volume fraction of water at a value of 0.5:

   A value of 0.5 for the water volume fraction means that a computational cell is filled with 50% of water and, consequently, 50% of air. Therefore, this value gives the position of the water-air-interface corresponding to the free surface.

   a) Right-click the **Derived Parts** node and select **New Part** > **Isosurface...**

      The *Create Isosurface* dialog appears.

   b) In the *Input Parts* group box, select **Regions** > **Virtual Towing Tank**.

   c) In the *Scalar* group box, select **Volume Fraction** > **Volume Fraction of Water**.

   d) Set *Isovalue* to `0.5`.

   e) Select **No Displayer**.

   f) Click **Create** and **Close**.



4. Click **Scene/Plot**.
5. Select the **Scalar 1** node.
6. In the *Scalar 1 - Properties* window:
   a) Set *Contour Style* to **Smooth Filled + Lines**.
   b) Set *Transform* to **Tank.Symmetry 1**.
7. Select the **Scalar 1 > Parts** node and set *Parts* to **Derived Parts > Isosurface**. Click **OK**.
8. Select the **Scalar 1 > Scalar Field** node and set *Function* to **Position > Laboratory > Z.**
9. Select the **Displayers > Outline 1** node and de-activate *Outline*.

10. Select the **Outline 1** > **Parts** node.
11. In the *Parts - Properties* window:
    a) Click (**Custom Editor**) next to *Parts*.
    b) Expand **Regions** > **Virtual Towing Tank** > **Boundaries** and select all boundaries with the prefix **KCS**.
    c) Expand **Derived Parts** and select **Isosurface**.
    d) Click **OK**.



12. Click anywhere in the scene display area and then press `t` on your keyboard to obtain a top view of the ship and the wave pattern.
13. Return to *Simulation*.

Monitoring and Plotting Resistance Data

The resistance of a ship hull is defined as the force required to tow the ship in calm water at a constant velocity. Generally, a ship is designed to move through the water with as little force as possible. You can create a graph of the hull resistance over time to extract the value of this design parameter of your ship from the simulation. The drag force can be split into the two components: pressure drag and shear drag.

To analyze the drag of the ship hull:

1. Create a report for the shear drag over time:
   a) Right-click the **Reports** node and select **New Report > Force**.
   b) Rename **Force 1** as `DragShear`.
   c) In the *DragShear - Properties* window:
      - Set *Force Option* to **Shear**.
      - Set *Parts* to **Regions > Virtual Towing Tank > Boundaries** [all boundaries with the prefix **KCS_Geometry**] and click **OK**.
      - Set *Direction* to `[-1.0, 0.0, 0.0]`.
2. Create a report for the pressure drag over time:
   a) Copy and paste the **DragShear** node.
   b) Rename **Copy of DragShear** as `DragPressure`.
   c) In the *DragPressure - Properties* window, set *Force Option* to **Pressure**.
3. Create a report for the total drag over time:
   a) Copy and paste the **DragShear** node.
   b) Rename **Copy of DragShear** as `DragTotal`.
   c) In the *DragTotal - Properties* window, set *Force Option* to **Pressure + Shear**.



4. Create a monitor and plot for a combined display of shear and pressure drag:
   a) Multi-select **DragPressure** and **DragShear**.
   b) Right-click on any of the highlighted nodes and select **Create Monitor and Plot from Report**.
   c) In the pop-up window, click **Single Plot**.
   d) Rename **Plots** > **Reports Plot** as `ShearAndPressureDrag Plot`.
   e) Right-click the **Plots** > **ShearAndPressureDrag Plot** node and select **Open**.
5. Adjust the scale of the plot:
   a) Select the **Plots** > **ShearAndPressureDrag Plot** > **Axes** > **Y Axis** node and set *Maximum* to `40.0`.
6. Create a monitor and plot for the total drag:
   a) Right-click the **Reports** > **DragTotal** node and select **Create Monitor and Plot from Report**.
   b) Rename **Plots** > **DragTotal Monitor Plot** as `DragTotal Plot`.
   c) Right-click the **Plots** > **DragTotal Plot** node and select **Open.**
7. Adjust the scale of the plot:

a) Select the **Plots** > **DragTotal Plot** > **Axes** > **Y Axis** node and set *Maximum* to `80.0`.

Monitoring and Plotting Trim and Sinkage Data

A ship advancing through water experiences sinkage and trim due to the hydrodynamic forces acting on the ship hull. A towing tank ship model free to sink and trim can experience an increase in resistance. You can create a graph of the trim and sinkage over time to extract the value of this design parameter of your ship from the simulation.

To analyze trim and sinkage of the ship hull:

1. Create a report for the trim:
   a) Right-click the **Reports** node and select **New Report** > **6-DOF Body Orientation**.
   b) Rename **Rigid Body Angle 1** to `Trim`.
   c) In the *Trim - Properties* window:
      - Set Angle Report Option to Rotation Y Axis.
      - Activate *Constrained Angle Range*.
2. Create a report for the sinkage:
   a) Right-click the **Reports** node and select **New Report** > **6-DOF Body Translation**.
   b) Rename **Rigid Body Translation 1** to `Sinkage`.
   c) In the *Sinkage - Properties* window, set *Direction* to `[0.0, 0.0, 1.0]`.



3. Create a monitor and plot for trim and sinkage:
   a) Multi-select the **Reports** > **Sinkage** and **Trim** nodes.
   b) Right-click on any of the highlighted nodes and select **Create Monitor and Plot from Report**.
   c) In the pop-up window, click **Single Plot**.
   d) Rename **Plots** > **Reports Plot** as `Motion Plot`.
   e) Right-click the **Plots** > **Motion Plot** node and select **Open**.
4. Save the simulation.

The setup is now complete and the simulation is ready to run.

Running the Simulation

You are advised to run this tutorial on a multi-core machine with four cores or more. Depending on the specifications of your computer, this tutorial takes one to two days to complete.

To run the simulation:

1. Click the (**Run**) button on the STAR-CCM+ toolbar.

While the simulation is running, you can click the tabs at the top of the *Graphics* window to view each of the scenes. The Residuals display is created automatically and shows the progress of the solvers.

Visualizing the Results

You can follow the progress of the solution using each of the scenes that you created previously.

The following screenshot shows a close-up of the free water surface around the ship hull on the symmetry plane at the end of the simulation:



The following screenshot shows the wave pattern around the ship hull at the end of the simulation:



The following screenshot shows a plot of shear and pressure drag over time:



Here, you can see strong oscillations at the beginning of the simulation. These oscillations level out as the simulation progresses to reach a quasi steady-state solution after approximately 130 s.

The following screenshot shows a plot of total drag over time:



The following screenshot shows a plot of trim and sinkage over time:



Comparing with Experimental Results

For the Gothenburg Workshop 2010, experimental results were provided for comparison with simulations.

Here, the resistance coefficient is investigated:

$$c_d = \frac{F_d}{\frac{\rho}{2} v^2 A}$$

where:

Symbol	Description	Value
c_d	Resistance coefficient	
F_d	Total drag force (N)	41.817
v	Ship velocity (m/s)	2.196
A	Scaled area of hull and rudder (m^2)	9.55275

The total drag force in the table above is obtained from the simulation.

1. Export the total drag over time plot:
   a) Right-click the **Plots** > **DragTotal Plot** node and select **Export...**.

b) Set **File Name** to `TotalDrag.csv` and click **Save**.

2. Import the table into a spreadsheet program of your choice and perform a numerical averaging (arithmetic mean) of the total drag values over time from 160 – 200 s.
3. Calculate the resistance coefficient according to the equation above.

   Note that the simulation models only half of the ship. To match the experimental data, either the drag or the resistance coefficient must be multiplied by 2.

Resistance Coefficient (Predicted)	Resistance Coefficient (Experimental)	Deviation
0.003640	0.003711	1.92%

This result is obtained on a grid that is rather coarse for this type of simulation. It is expected that the deviation decreases with increasing mesh refinement.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Importing a ship hull geometry.
- Creating a towing tank around the ship.
- Generating a locally refined volume mesh using volumetric controls.
- Setting up a VOF multiphase simulation.
- Setting up a flat VOF Wave.
- Setting up boundary-wise wave damping.
- Setting up a DFBI rotation and translation motion.
- Setting up and running a transient simulation.
- Creating specific post-processing items for flow resistance prediction analysis.

Bibliography

http://www.simman2008.dk/KCS/kcs_geometry.htm. Downloaded July 2013.

Moving Reference Frames: Marine Propeller in Open Water

This tutorial demonstrates the workflow for setting up an open water simulation for a marine propeller.

The geometry that is used in this simulation is the Potsdam Propeller Test Case (PPTC) VP1304. The results that you obtain in STAR-CCM+ are compared with experimental data that is provided by SVA. This data was presented at the workshop on propeller performance during the International Symposium on Marine Propulsors held in Hamburg in June 2011.

The PPTC VP1304 propeller is placed in a virtual water tank as shown below:



The PPTC VP1304 propeller is a controllable pitch propeller that includes a 0.3 mm gap between the hub and the propeller blade at the leading and trailing edges. This gap is maintained in the simulation.

The rotation of the propeller is modelled using MRF.

The simulation is run for a range of advance coefficient values J (ranging from 0.6 to 1.4). The operating conditions for the simulation are listed as follows:

Density of Water (at 17.5°C)	998.67 kg/m^3
Dynamic Viscosity of Water (at 17.5°C)	0.001068 Pa s
Propeller Diameter	0.25 m
Number of Revolutions	15 rps

You are provided with the propeller performance data plot. The thrust coefficient K_T, the torque coefficient K_Q, and the open water efficiency η_O are calculated for each value of J and are plotted in the propeller performance plot for comparison.

Contents

Prerequisites
Loading the Simulation File
Reviewing the Geometry
Selecting the Mesh Models
Generating the Mesh
Setting Up Boundary Conditions and Values
Creating the Reports for the Performance Data
Monitoring the Propeller Performance Plot
Running the Simulation
Plotting the Propeller Performance Data
Summary
Bibliography

Prerequisites

The instructions in the Open Water tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading the Simulation File

For this tutorial, you are provided with a simulation file that contains pre-defined simulation objects. Due to the large number of cells, and the time it takes the solver to run, it is recommended that you run this tutorial in parallel on at least 10 cores.

Load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/motion/data`
5. Select `openWaterPropeller_start.sim`, then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `openWaterPropeller_run.sim`.

Reviewing the Geometry

For open water simulations, it is recommended that your geometry model exhibits certain characteristics.

The geometry that is provided in the starting simulation for this tutorial demonstrates several best practices for this type of simulation:

- There is a feature curve along the leading and trailing edge of the propeller to demarcate the pressure and suction sides of the propeller. During the meshing process, this feature curve allows you to capture the shape of the blade with a relatively coarse mesh.



- Small details in the geometry are closed off. In this simulation, the clearance gap between the propeller hub and the adjacent shaft has been bridged with an additional surface, as indicated in the following screenshot.



The motion of the propeller is modelled using MRF. The starting simulation file contains two regions, one for the rotating propeller, and the other for the static fluid domain. The static region is extended using the extruder mesh model.



Selecting the Mesh Models

For this type of simulation, the quality of the result depends to a large extent on how successful the mesh is at capturing features of the flow around the propeller.

In this simulation, you use the trimmer mesh model for the rotating region and its surroundings. The static region surrounding the shaft is meshed using the extruder mesher as this minimizes the computational costs.

An important criteria to consider when choosing mesh sizes is the wall y + value. In general, you are advised to keep the wall y+ values outside the buffer layer, that is values from 5 to 30. In this tutorial, you aim for wall y+ values greater than 30.

1. Create a mesh continuum under the **Continua** node and select the following models in order:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Trimmer**
Optional Models	**Prism Layer Mesher**
Optional Models	**Extruder**

2. Click **Close**.

Activate per-region meshing to allow you to apply the mesh continuum to multiple regions:

3. Select the **Continua** > **Mesh 1** node and activate *Per-Region Meshing*.

By default, mesh alignment is deactivated. Activate this feature to specify a point where the mesh should align:

4. Select the **Continua** > **Mesh 1** > **Models** > **Trimmer** node and activate *Do mesh alignment*. The default value is acceptable for this tutorial.

Contents

Specifying the Reference Values
Refining the Propeller Blades
Controlling the Prism Layer Size in the Clearance Gap
Specifying Prism Layer Size on the Interface
Specifying the Far Field Surface Size
Refining the Area Around the Propeller
Setting the Extruder Options

Specifying the Reference Values

Specify the reference mesh values for this simulation:

1. Edit the **Continua** > **Mesh 1** > **Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	*Value*	9 mm
Number of Prism Layers	*Number of Prism Layers*	3
Prism Layer Thickness > **Relative Size**	*Percentage of Base*	20

Node	Property	Setting
Surface Growth Rate	*Surface Growth Rate*	`1.3`
Surface Size > Relative Minimum Size	*Percentage of Base*	`0.5`
Template Growth Rate	*Default Growth Rate*	**Medium**

2. Click **Close**.

Refining the Propeller Blades

Set up custom mesh conditions to refine the mesh on the propeller blades.

1. Edit the **Regions > Rotating Region > Boundaries > Blades > Mesh Conditions** node.
2. Set the following properties:

Node	Property	Setting
Custom Boundary Growth Rate	*Custom Boundary Growth Rate*	**Activated**
Custom Surface Curvature	*Custom curvature*	**Specify Custom Values**
Custom Surface Size	*Custom Surface Size*	**Activated**

3. Edit the **Blades > Mesh Values** node and set the following properties:

Node	Property	Setting
Boundary Growth Rate	*Boundary Growth Rate*	**Fast**
Surface Curvature		
▮▮ **Basic Curvature**	# *Pts/circle*	`72`
Surface Size		
▮▮ **Relative Minimum Size**	*Percentage of Base*	`5`
▮▮ **Relative Target Size**	*Percentage of Base*	`25`

Use the feature curve that splits the pressure side and suction side of the propeller to refine the mesh on the blade edges:

4. Select the **Regions** > **Rotating Region** > **Feature Curves** > **Blade Edges** > **Mesh Conditions** > **Custom Surface Size** node and activate *Custom surface size*.
5. Edit the **Blade Edges** > **Mesh Values** > **Surface Size** node and set the following properties:

Node	Property	Setting
Relative Minimum Size	*Percentage of Base*	2
Relative Target Size	*Percentage of Base*	5

Controlling the Prism Layer Size in the Clearance Gap

The propeller in this simulation is a controllable pitch propeller, which means that there is a small clearance gap between each blade and the hub. In order to accurately resolve the boundary layer in this gap, modify the prism layer settings.

1. Edit the **Continua** > **Mesh 1** > **Models** > **Prism Layer Mesher** node and set the following properties:

Node	Property	Value
Prism Layer Mesher	*Minimum Thickness Percentage*	1
	Layer Reduction Percentage	5
	Boundary March Angle	25

Specifying Prism Layer Size on the Interface

The interface between the rotating region and the surrounding static region is cylindrical in shape. It is good practice to have cells that are normal to each other on either side of an interface.

To achieve this result, grow a single prism layer cell on from either side of the interface:

1. Multi-select the **Regions** > **Rotating Region** > **Boundaries** > **Interface** and the **Regions** > **Static Region** > **Boundaries** > **Interface** nodes.
2. Right-click one of the selected nodes and select **Edit...**
3. Set the following properties:

Node	Property	Setting
Mesh Conditions > Customize Prism Mesh	*Customize Prism Mesh*	**Specify Custom Values**
Mesh Values		
Number of Prism Layers	*Number of Prism Layers*	1
Prism Layer Thickness > Relative Size	*Percentage of Base*	50

4. Click **Close**.

Specifying the Far Field Surface Size

As flow behavior far from the propeller region is not particularly important for these simulations, you can specify a coarser mesh on the far field, inlet, and outlet boundaries.

1. Expand the **Regions > Static Region > Boundaries** node.
2. Multi-select the **Far Field**, **Inlet**, and **Outlet** nodes.
3. Right-click one of the nodes, choose **Edit...**, and set the following properties:

Node	Property	Setting
Mesh Conditions > Custom Surface Size	*Custom Surface Size*	**Activated**
Mesh Values > Surface Size > Relative Target Size	*Percentage of Base*	1600

4. Click **Close**.

Refining the Area Around the Propeller

The area around the propeller requires a finer mesh to make sure that the flow characteristics are adequately captured.

Create a volumetric control region around the propeller:

1. Right-click the **Geometry > Parts** node and choose **New Shape Part > Cylinder**.
2. Set the following parameters:

Property	Setting
Start Coordinate	`[0.2, 0.0, 0.0] m`
End Coordinate	`[-0.25, 0.0, 0.0] m`
Cylinder Radius	`0.18 m`

3. Create a volumetric control under the **Continua** > **Mesh 1** > **Volumetric Controls** node.
4. Select **Volumetric Control 1** and set *Parts* to **Cylinder**.
5. Edit the **Volumetric Control 1** node and set the following properties:

Node	Property	Setting
Mesh Conditions		
Trimmer	*Customize isotropic size*	**Activated**
Mesh Values		
Custom Size > **Relative Size**	*Percentage of Base*	`50`

Setting the Extruder Options

Use the extruder mesh to extend the domain in the inflow and outflow directions. As the flow in these directions is generally normal to the cell faces, it is acceptable to generate stretched cells. In this tutorial, the largest cells have an aspect ratio of approximately 300.

First, extend the domain in the inflow direction:

1. Edit the **Regions** > **Static Region** > **Boundaries** > **Inlet** node and set the following properties:.

Node	Property	Setting
Mesh Conditions		
Extrusion Type	*Extrusion Type*	**Constant Rate**
Mesh Values		
Extrusion Parameters	*Direction Magnitude*	`[1.0, 0.0, 0.0] m, m, m`
	Number of Layers	`10`
	Stretching	`2.0`

2. Edit the **Regions** > **Static Region** > **Boundaries** > **Outlet** node and set the following properties:

Node	Property	Setting
Mesh Conditions		
Extrusion Type	*Extrusion Type*	**Constant Rate**
Mesh Values		
Extrusion Parameters	*Direction Magnitude*	`[-3.0, 0.0, 0.0] m, m, m`
	Number of Layers	`40`
	Stretching	`2.0`

3. Save the simulation.

Generating the Mesh

The simulation is ready for meshing. The mesh generation process takes roughly 10 minutes to complete, depending on the hardware you are using.

1. Generate the volume mesh.
2. Open **Mesh Scene 1**.
3. Right-click a blank area in the scene and select **Apply Representation** > **Volume Mesh**. You may need to refresh the scene for the section plane to display correctly.



4. Save the simulation.

Setting Up Boundary Conditions and Values

In this tutorial, most of the boundary conditions are already defined for you. All that remains for you to do is to define the velocity of the flow, V_A, at the inlet and as an initial condition in the fluid domain.

This velocity is driven by the advance coefficient, J:

$$J = \frac{V_A}{n \cdot D}$$

where:

- V_A is the advance velocity
- n is the rate of revolution
- D is the propeller diameter

Contents

Defining the Advance Coefficient
Specifying Boundary Conditions

Defining the Advance Coefficient

In this steady state simulation, you specify a range of J values using two field functions: J and i. The field function, i, specifies the iteration on which the value of J changes.

In this tutorial, the starting value of J is 0.6. The value of J increases by 0.2 every 501 iterations. This number of iterations is sufficient for the solution to converge for each value of J.

First, define the field function **i**:

1. Create a custom scalar field function and rename it to `i`.
2. Select the **i** node and set the following properties:

Property	Setting
Definition	`floor ($Iteration/501)`
Function Name	`i`

Create and define the field function for **J**:

3. Create a custom scalar field function and rename it to `J`.
4. Select the **J** node and set the following settings:

Property	Setting
Definition	`0.6 + 0.2 * $i`
Function Name	`J`

Proceed to create field functions that supply the parameters for the advance velocity:

5. Create three custom scalar field functions and name them `n`, `D`, and `VA` respectively.
6. Set the following properties for the field functions:

Field Function	Property	Setting
n	*Definition*	`15`
	Function Name	`n`
D	*Definition*	`0.25`

Field Function	Property	Setting
	Function Name	`D`
VA	*Definition*	`$J * $n * $D`
	Function Name	`VA`

Specifying Boundary Conditions

Here, you apply the advance velocity to the inlet boundary.

To set the inlet boundary:

1. Edit the **Regions** > **Static Region** > **Boundaries** > **Inlet** node and set the following properties:

Node	Property	Setting
Physics Values		
Velocity Magnitude	Method	Field Function
Field Function	Scalar Function	VA

Use the advance velocity as the initial velocity for the fluid domain:

2. Expand the **Continua** > **Physics 1** > **Initial Conditions** node.
3. Select the **Velocity** > **Constant** node and set *Value* to `[-1 * $VA, 0.0, 0.0]`.

Creating the Reports for the Performance Data

The propeller performance data consists of the advance coefficient, the thrust coefficient, the torque coefficient, and the open water efficiency.

For this case, a custom field function that specifies the value of density is provided for you.

Thrust Coefficient

The thrust coefficient, K_T, is defined as:

$$K_T = \frac{T}{\rho \cdot n^2 \cdot D^4} \tag{6}$$

where:

- T is the thrust that is generated by the propeller

- ρ is the density of water
- n is the rate of revolution of the propeller
- D is the propeller diameter

In order to calculate this value, you must first calculate the thrust that the propeller generates. Create the report for the thrust:

1. Right-click **Reports** and select **New Report** > **Force**.
2. Rename this report to `T`.
3. Select the **T** report and set the following properties:

Property	Setting
Direction	`[1.0, 0.0, 0.0]`
Parts	**[Rotating Region: Blades, Rotating Region: Hub]**

4. Right-click **Reports** and select **New Report** > **Expression**.
5. Rename this report to `KT`.
6. Select the **KT** report and set *Definition* to `$TReport/($rho*pow($n,2)*pow($D,4))`.

Torque Coefficient

The torque coefficient, K_Q, is defined as:

$$K_Q = \frac{Q}{\rho \cdot n^2 \cdot D^5} \tag{7}$$

where:

- Q is the torque that is generated by the propeller
- ρ is the density of water
- n is the rate of revolution of the propeller
- D is the propeller diameter

In order to calculate this value, you must first calculate the torque that the propeller generates. Create the report for the torque:

1. Right-click **Reports** and select **New Report** > **Moment**.
2. Rename this report to `Q`.
3. Select the **Q** report and set the following properties:

Property	Setting
Direction	`[-1.0, 0.0, 0.0]`

Property	Setting
Parts	**[Rotating Region: Blades, Rotating Region: Hub]**

4. Right-click **Reports** and select **New Report > Expression**.
5. Rename this report to `KQ`.
6. Select the **KQ** report and set *Definition* to `$QReport/($rho*pow($n,2)*pow($D,5))`.

Open Water Efficiency

The open water efficiency, η_O, is defined as:

$$\eta_O = \frac{J}{2\pi} \cdot \frac{K_T}{K_Q} \tag{8}$$

where:

- J is the advance coefficient
- K_T is the thrust coefficient
- K_Q is the torque coefficient

Create the report for the open water efficiency:

1. Right-click **Reports** and select **New Report > Expression**.
2. Rename this report to `hO`.
3. Select the **hO** report and set *Definition* to `($KTReport/$KQReport)*($JReport/(2*3.1428571429))`.

Monitoring the Propeller Performance Plot

The propellor performance plot contains values of J and 10KQ. Here you create reports that provide these values.

1. Create an expression report and rename it to `J`.
2. Select the **J** report and set *Definition* to `$J`. This report is monitoring the value of J from the corresponding custom field function.
3. Create an expression report and rename it to `10KQ`.
4. Select the **10KQ** report and set *Definition* to `$KQReport*10`.

The propeller performance plot includes the values of J, KT, 10KQ, and ηO. In the following steps you create monitors from these reports and modify the update frequency. The value of J changes every 501

iterations. The converged values of these reports are saved every 500 iterations.

5. Multi-select the **J**, **10KQ**, **KT**, and **hO** reports.
6. Right-click one of the nodes and choose **Create Monitor from Report**.
7. Expand the **Monitors** node and multi-select the **10KQ Monitor**, **KT Monitor**, **hO Monitor**, and **J Monitor** nodes.
8. Right-click one of the nodes and choose **Edit...**.
9. Select the **Iteration Frequency** node and set *Iteration Frequency* to `500`.
10. Click **Close**.
11. Right-click **Monitors** > **J Monitor** and select **Create Plot from Monitor**.
12. Save the simulation.

Running the Simulation

Set the stopping criteria and run the simulation.

1. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `2500`.
2. Run the simulation.

   The simulation can take a few hours to complete, depending on the number of cores you are running.
3. Save the simulation.

Plotting the Propeller Performance Data

When complete, the simulation contains all the data that is required to plot the propeller performance data and compare it with the experimental results. The tutorial starting simulation already has a basic plot for you to use.

Begin by exporting the data to a `.csv` file:

1. Expand the **Monitors** node.
2. Multi-select the **10KQ Monitor**, **KT Monitor**, **hO Monitor**, and **J Monitor** nodes.
3. Right-click on one of the selections and select **Export...**

4. In the *Save* dialog, save the file as `propellerPerformanceData.csv` in the working directory.

   The data is exported into a single file that contains five columns of data.
5. Right-click **Tools** > **Tables** and select **New Table** > **File Table**.
6. In the *Open* dialog, select `propellerPerformanceData.csv`.

Update the plot:

7. Expand the **Plots** > **Propeller Performance Data** node.
8. Right-click **Data Series** and select **Add Data**.
9. In the *Add Data Providers to Plot* dialog, select **propellerPerformanceData** and click **OK**.
10. Edit the **propellerPerformanceData** node and set the following properties:

Node	Property	Setting
propellerPerformanceData	*Legend Name*	`KT`
	X Column	**J Monitor: Expression**
	Y Column	**KT Monitor: Expression**
Line Style	*Color*	**Black**
	Style	**Solid**
	Width	`2`
Symbol Style	*Size*	`10`
	Color	**Black**
	Shape	**Filled Triangle**

11. Create two more data sets from the **propellerPerformanceData** data, one for the `10KQ` data, and the other for the `hO` data.
12. Edit the **propellerPerformanceData 2** node and set the following properties:

Node	Property	Setting
propellerPerformanceData 2	*Legend Name*	`10KQ`
	X Column	**J Monitor: Expression**

Node	Property	Setting
	Y Column	**10KQ Monitor: Expression**
Line Style	*Color*	**Black**
	Style	**Solid**
	Width	`2`
Symbol Style	*Size*	`10`
	Color	**Black**
	Shape	**Filled Circle**

13. Edit the **propellerPerformanceData 3** node and set the following properties:

Node	Property	Setting
propellerPerformanceData 3	*Legend Name*	`hO`
	X Column	**J Monitor: Expression**
	Y Column	**hO Monitor: Expression**
Line Style	*Color*	**Black**
	Style	**Solid**
	Width	`2`
Symbol Style	*Size*	`10`
	Color	**Black**
	Shape	**Filled Square**

14. Open the **Propeller Performance Data** plot.

The following plot displays the STAR-CCM+ calculations as black solid lines, and the experimental measurements as red long dashed lines.



Summary

This tutorial demonstrated the workflow for setting up an open water simulation for a marine propeller.

This tutorial has introduced the following STAR-CCM+ features:

- Importing a propeller geometry and checking that it contains recommended features for marine simulations.
- Extending the fluid domain using the extruder mesher.
- Setting up a steady state simulation to run with varying advance coefficent values.
- Setting up and exporting propeller performance data to a .csv file.
- Plotting the propeller performance data that is calculated in STAR-CCM+ against experimental data.

Bibliography

http://www.sva-potsdam.de/pptc_data.html#geometry. Downloaded October 2013.

Body Force Propeller Method: Marine Self-Propulsion

One of the challenges in marine engineering is to predict the speed with which a ship hull moves through water in response to the thrust supplied by a spinning propeller. STAR-CCM+ provides a methodology by which you can predict this speed.

This tutorial simulates the motion of a ship through calm water in response to the thrust produced by its propeller. The propeller is placed centrally at the stern of the ship and rotates at 2300 rpm. Propeller effects are important in these type of simulations, as they can significantly affect the performance of the hull due to the swirl and pressure gradients.

To model self-propulsion, the ship is placed inside a virtual towing tank in full scale. In this tutorial, you apply the Body Force Propeller method of the Virtual Disk model to represent the effect of the propeller on the ship. The Body Force Propeller method allows you to model the propeller and hull interaction without the added cost of modeling the propeller geometry with a fine mesh and smaller time scales. For more information, see Body Force Propeller Method in the User Guide.



Specific models used in this tutorial are:

- Volume of Fluid (VOF) for the water and air. The simulation starts from a flat water surface.
- Dynamic Fluid Body Interaction (DFBI) for calculating the ship motion in response to fluid forces acting on it. The ship is not allowed to move forwards or backwards, but is free to heave and pitch.

In this simulation, the ship moves at a steady velocity with no acceleration. To achieve this condition in STAR-CCM+, the thrust induced by the propeller must balance the drag forces on the ship hull. However, you do not know in advance the inlet water velocity that is required to generate drag forces that balance the thrust. This tutorial describes a process for discovering the inlet velocity that results in an overall force balance of zero. Obtaining this balancing inlet velocity is equivalent to obtaining the hull speed as it represents the speed of the ship traveling through calm water relative to a fixed reference frame.

The model uses the full geometry of the ship and does not employ symmetry conditions.

The following boundary conditions are applied:

- inlets with prescribed velocity and volume fraction at the upstream, top, bottom, and side boundaries
- hydrostatic pressure corresponding to the undisturbed water surface at the downstream boundary of the tank
- walls on the ship hull surface

Contents

Prerequisites
Loading the Simulation File
Selecting the Physics Models
Setting Up Dynamic Fluid Body Interaction (DFBI)
Modeling the Propeller by Using the Virtual Disk Model
Applying the Thrust Force to the Ship
Creating Mesh Refinement for the Virtual Disk
Generating the Volume Mesh
Defining the Water and Air Phases
Defining the VOF Wave
Setting Initial Conditions
Damping Wave Reflections
Setting Up Boundary Conditions
Setting Solver Parameters and Stopping Criteria

Monitoring and Plotting the Net Force in X Direction
Visualizing the Pressure Distribution
Running the Simulation
Visualizing the Results
Obtaining the Balancing Inlet Velocity
Summary

Prerequisites

The instructions in this tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction
Using visualization tools, scenes, and plots	Introduction
Meshing and visualization of virtual towing tank applications	Marine Resistance Prediction: KCS Hull with a Rudder

Loading the Simulation File

For this tutorial, you are provided with a simulation file that contains pre-defined objects and settings.

The pre-defined objects and settings include:

- geometry
- mesh operations
- boundary types

Due to the large number of cells, and the time it takes the solver to run, launch this tutorial in parallel on at least 4 cores.

To load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Process Options* box, select an appropriate parallel option and set the number of *Compute Processes* to an amount not less than `4`.
4. In the *Load Simulation* dialog, click **Browse...**

5. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/motion/data`.
6. Select `marineSelfPropulsion_start.sim` then click **Open**.
7. In the *Load Simulation* dialog, click **OK**.
8. To visualize the imported geometry, open **Geometry Scene 1**.



9. Save the simulation as `marineSelfPropulsion_run.sim`.

Selecting the Physics Models

This marine self-propulsion simulation is a transient complex multi-physics problem for which you select a variety of physics models.

This simulation involves two fluid phases - water and air - that are modeled using the VOF multiphase model. The flow field is transient and turbulent. To model the effects of the propeller, you select the virtual disk model. To account for the hydrostatic pressure distribution under water, you activate gravity. You also use VOF waves that specify initial and boundary conditions for VOF phases as the simulation progresses.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Gradients** (selected automatically) **Three Dimensional** (selected automatically)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (selected automatically)
Eulerian Multiphase Model	**Volume of Fluid (VOF)** **Multiphase Equation of State** (selected automatically) **Segregated Flow** (selected automatically)
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Omega Turbulence** **SST (Menter) K-Omega** (selected automatically)
Optional Models	**Virtual Disk** **VOF Waves** **Gravity**

2. To review the models, open the **Physics 1** > **Models** node.



3. Save the simulation.

Setting Up Dynamic Fluid Body Interaction (DFBI)

The Dynamic Fluid Body Interaction (DFBI) model simulates the motion of the ship according to the forces acting on it. For this simulation, the ship is allowed to move with two degrees of freedom to account for sinkage and trim.

To set up the dynamic fluid body interaction:

1. Right-click the **Tools > Motions** node and select **New > DFBI Rotation and Translation**.
2. Select the **Regions > Domain > Physics Values > Motion Specification** node and set **Motion** to **DFBI Rotation and Translation**.

   The DFBI Rotation and Translation motion is assigned to the Domain region and a new **DFBI** node appears in the simulation tree..
3. Right-click the **DFBI > 6-DOF Bodies** node and select **New**.
4. Rename **Body 1** to `Ship`.
5. Edit the **6-DOF Bodies** node and set the following properties:

Node	Property	Setting
Ship	**Body Surface**	**Domain: Deck, Domain: Hull, Domain: Stern**
	Body Mass	`13775 kg`
	Ramp Time	`3.0 s`
Free Motion	**Z Motion**	Activated
	Y Rotation	Activated
Initial Values		
Center of Mass	**Position**	`[6.807, 0.0, 0.9] m,m,m`
Moment of Inertia	**Diagonal Components**	`[1.0, 181000.0, 1.0] kg-m^2`
	Use Center of Mass	Activated

6. Save the simulation.

Modeling the Propeller by Using the Virtual Disk Model

For this simulation, you model the effects of the propeller on the ship using the body force propeller method. This method belongs to the virtual disk model.

Using the virtual disk for a propeller requires you to define the following items:

- dimensions and location of the disk
- orientation of the disk axis. This orientation defines the direction of the thrust force.
- propeller performance data, where the thrust coefficient K_T and the torque coefficient K_Q are provided as a function of advance ratio J
- velocity plane on which to obtain an average velocity and density for the inflow surface of the virtual disk

To set up the virtual disk model:

1. Right-click the **Continua** > **Physics 1** > **Models** > **Virtual Disk** > **Virtual Disks** node and select **New**.
2. Edit the **Virtual Disk 1** node and set **Method** to **Body Force Propeller Method**.

Before you can specify the propeller curve, you must import the propeller table data file (here in .csv format) into STAR-CCM+.

3. Right-click **Tools > Tables** and select **New Table > File Table**.
4. Navigate to [INSTALL_DIR]/doc/startutorialsdata/motion/data/BSeries_Z5_BAR1p05_PD0p723.csv and click **Open**.

You can now specify the propeller curve and set up the virtual disk. You set the *Velocity Plane Radius* to a value that is 10% larger than the virtual disk radius. The **Velocity Plane Offset** is 10% of the virtual disk diameter.

5. Edit the **Virtual Disks > Virtual Disk 1** node and set the following properties:

Node	Property	Setting
Propeller Curve	**Method**	**Table (J)**
Table (J)	**Table**	**BSeries_Z5_BAR1p05_PD0p723**
	Table: J	**J**
	Table: Kt	**KT**
	Table: Kq	**KQ**
	Table: Eta	**ETA**
Disk Geometry	**Inner Radius**	`0.036 m`
	Outer Radius	`0.18 m`
	Origin	`[-6.307, 0.0, -1.25] m,m,m`
Normal And Coordinate System	**Coordinate System**	**Laboratory -> Ship-CSys**
	Disk Normal	`[1.0, 0.0, 0.0]`
Inflow Velocity Plane	**Velocity Plane Radius**	`0.198 m`
	Velocity Plane Offset	`0.018 m`
RotationRate (n)	**Rotation Rate**	`2300.0 rpm`

Note: You must associate the **Coordinate System** property that defines the direction of thrust with the local coordinate system of the ship body, here identified as **Ship-CSys**. This association establishes the connection between the thrust produced by the virtual disk and the DFBI motion of

the ship. The ship coordinate system moves with the ship according to the fluid forces calculated by the DFBI model.

6. Expand the **Tools** > **Coordinate Systems** > **Laboratory** > **Local Coordinate Systems** and drag and drop the **Virtual Disk 1-CSys 1** node into the currently active *Geometry 1* scene.

   The virtual disk local coordinate system is displayed in the scene. The thrust produced by the virtual disk acts in the Z direction of this coordinate system.

7. To visualize the virtual disk, select the **Virtual Disk 1** node within **Physics Continuum 1** and zoom in on the virtual disk.



8. Save the simulation.

Applying the Thrust Force to the Ship

To apply the thrust force produced by the virtual disk to the ship, you must create an external force for the 6-DOF ship body.

To apply the virtual disk thrust to the ship:

1. Right-click the **DFBI > 6-DOF Bodies > Ship > External Forces and Moments** node and select **New > Virtual Disk Force**.
2. Select the **External Forces and Moments > VirtualDiskForce 1** node and set **Virtual Disk** to **Virtual Disk 1**.

   The force produced by **Virtual Disk 1** is now associated with the **Ship** body.



3. Save the simulation.

Creating Mesh Refinement for the Virtual Disk

To obtain a good resolution of the flow field around the virtual disk, you are advised to create a cylindrical volumetric refinement. You create the refinement using part shapes and mesh custom controls.

To create the mesh refinement for the virtual disk:

1. Right-click the **Geometry > Parts** node and select **New Shape Part > Cylinder**.
2. In the *Create/Edit Cylinder Part* dialog set the following parameters:

	Start Coordinate	End Coordinate	Cylinder Radius	Coordinate System
X	-5.4145 m	-8.192 m	0.357	**Laboratory \Ship-CSys**
Y	0.0 m	0.0 m		
Z	-1.25 m	-1.25		



You create the cylindrical shape part with respect to the local body coordinate system of the ship.

3. Click **Create**, then **Close**.
4. Rename the **Parts** > **Cylinder** node as `Disk Refinement`.

Continue by assigning custom mesh controls to the **Disk Refinement** part shape:

5. Right-click the **Geometry** > **Operations** > **Automated Mesh** > **Custom Controls** node and select **New** > **Volumetric Control**.
6. Rename the **Volumetric Control** node as `Disk Refinement`.
7. Edit the **Disk Refinement** node and set the following properties:

Node	Property	Setting
Disk Refinement	**Parts**	**Disk Refinement**
Controls > Trimmer	**Customize Isotropic Size**	Activated
Values > Custom Size	**Size Type**	**Relative to base**
Relative Size	**Percentage of Base**	`10.0`

8. Save the simulation.

Generating the Volume Mesh

Suitable for this type of application, you generate a trimmed hexahedral mesh with local refinements and prism layers along the walls of the wetted surface.

In the initial simulation file, the required mesh models and corresponding settings are already set up. The wake pattern of the boat is resolved through cone-shaped volumetric refinements.

To generate the volume mesh:

1. Select **Mesh** > **Generate Volume Mesh**

   The mesh generation process several seconds to complete.
2. Open **Mesh Scene 1**.
3. Right-click a blank area in the scene and select **Apply Representation** > **Volume Mesh**

   You may need to refresh the scene for the section plane to display correctly.



4. Save the simulation.

Defining the Water and Air Phases

Using the VOF multiphase model requires that you create Eulerian phases for water and air.

To define the water and air phases:

1. In the **Physics 1** continuum, right-click the **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and create a new phase.
2. Rename the **Phase 1** node to `Water`.
3. For the new **Water** phase, select the following models:

Group Box	Model
Material	**Liquid**
Equation of State	**Constant Density**

4. Create a second phase and rename it to `Air`.
5. For the **Air** phase, select the following models:

Group Box	Model
Material	**Gas**
Equation of State	**Constant Density**

You leave the material properties for water and air at their default values.

The **Eulerian Multiphase** node contains the following models:



6. Save the simulation.

Defining the VOF Wave

The free surface water level changes over time during the simulation. STAR-CCM+ provides the VOF Waves model that provides wave initial and boundary conditions. Here, you define a flat VOF wave as the ship is moving through calm water.

To define a flat VOF Wave:

1. In the **Physics 1** continuum, right-click the **VOF Waves > Waves** node and select **New > Flat**.
2. Select the **FlatVofWave 1** node and set the following properties:

Property	Setting
Current	`[-4.57, 0.0, 0.0] m/s`
Wind	`[-4.57, 0.0, 0.0] m/s`

The inlet velocity that you specify here results from the procedure for obtaining the inlet velocity to balance the forces in X direction.

3. Save the simulation.

Setting Initial Conditions

You set the initial conditions for pressure, velocity, and volume fraction.

To set the initial conditions:

1. In the **Physics 1** continuum, edit the **Initial Conditions** node and set the following properties:

Node	Property	Setting
Pressure	**Method**	**Field Function**
Field Function	**Scalar Function**	**Hydrostatic Pressure of FlatVofWave 1**
Velocity	**Method**	**Field Function**
Field Function	**Vector Function**	**Velocity of FlatVofWave 1**
Volume Fraction	**Method**	**Composite**
Composite > Water	**Method**	**Field Function**
Field Function	**Scalar Function**	**Volume Fraction of Heavy Fluid of FlatVofWave 1**
Composite > Air	**Method**	**Field Function**
Field Function	**Scalar Function**	**Volume Fraction of Light Fluid of FlatVofWave 1**

2. Save the simulation.

Damping Wave Reflections

Unless appropriate measures are applied, waves can be reflected from the surrounding boundaries due to the nature of the numerical methods employed. These non-physical wave reflections must be avoided as they interact with the true wave field and can lead to invalid results.

STAR-CCM+ provides the VOF Wave damping model to minimize the risk of wave reflection.

To activate VOF Wave damping at boundaries:

1. Edit the **Regions** > **Domain** node and set the following properties:

Node	Property	Setting
Physics Conditions		
VofWave Damping Option	**VofWave Damping Option**	**Damping**
Physics Values		
VofWave Damping Length	**Method**	**Constant**
Constant	**Value**	`15.0 m`

2. Save the simulation.

VOF wave damping using the specified damping length is later activated for each boundary.

Setting Up Boundary Conditions

To specify wave damping, velocities, and variable water surface level at the surrounding boundaries of the ship, you apply field functions from the flat VOF wave.

To set up the boundary conditions:

1. Edit the **Regions** > **Domain** > **Boundaries** > **Aft** node and set the following properties:

Node	Property	Setting
Physics Conditions		
VofWaveDampingBoundaryOption	**Method**	**yes**
Physics Values		
Pressure	**Method**	**Field Function**
Field Function	**Scalar Function**	**Hydrostatic Pressure of FlatVofWave 1**
Volume Fraction	**Method**	**Composite**
Composite > Water	**Method**	**Field Function**

Node	Property	Setting
▮▮▮▮▮▮Field Function	Scalar Function	Volume Fraction of Heavy Fluid of FlatVofWave 1
▮▮▮▮Composite > Air	Method	Field Function
▮▮▮▮▮▮Field Function	Scalar Function	Volume Fraction of Light Fluid of FlatVofWave 1

2. Multi-edit the **Forward**, **Port**, and **Starboard** nodes and set the following properties:

Node	Property	Setting
Physics Conditions		
▮▮Velocity Specification	Method	Components
▮▮ VofWaveDampingBoundaryOption	Method	yes
Physics Values		
▮▮Velocity	Method	Field Function
▮▮▮▮Field Function	Vector Function	Velocity of FlatVofWave 1
▮▮Volume Fraction	Method	Composite
▮▮▮▮Composite > Water	Method	Field Function
▮▮▮▮▮▮Field Function	Scalar Function	Volume Fraction of Heavy Fluid of FlatVofWave 1
▮▮▮▮Composite > Air	Method	Field Function
▮▮▮▮▮▮Field Function	Scalar Function	Volume Fraction of Light Fluid of FlatVofWave 1

3. Multi-edit the **Bottom** and **Top** nodes and set the following properties:

Node	Property	Setting
Physics Conditions		

Node	Property	Setting
Velocity Specification	**Method**	**Components**
Physics Values		
Velocity	**Method**	**Field Function**
Field Function	**Vector Function**	**Velocity of FlatVofWave 1**
Volume Fraction	**Method**	**Composite**
Composite > Water	**Method**	**Field Function**
Field Function	**Scalar Function**	**Volume Fraction of Heavy Fluid of FlatVofWave 1**
Composite > Air	**Method**	**Field Function**
Field Function	**Scalar Function**	**Volume Fraction of Light Fluid of FlatVofWave 1**

Leave all wall boundaries at their default settings.

4. Save the simulation.

Setting Solver Parameters and Stopping Criteria

This self-propulsion simulation is a transient analysis for which you must specify a set of numerical parameters. The choice of parameters has a strong influence on the performance of STAR-CCM+ for these type of analyses..

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following parameter:

Node	Property	Setting
Implicit Unsteady	*Time-Step*	`0.02 s`
6-DOF Solver	*Maximum Number of Iterations*	`7`

2. Edit the **Stopping Criteria** node and set the following parameter:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`
Maximum Physical Time	*Maximum Physical Time*	`80.0 s`
Maximum Steps	*Enabled*	**Deactivated**

3. Save the simulation.

Monitoring and Plotting the Net Force in X Direction

You create several reports for the different forces such as drag, trim, and the virtual disk force components. The final expression report for the net force in X direction uses these reports.

To analyze the net force in the X direction acting on the ship hull:

1. Create a report for the total drag over time:
   a) Right-click the **Reports** node and select **New Report** > **Force**.
   b) Rename the **Force 1** node to `Drag`.
   c) Select the **Drag** report and set the following properties:

Property	Setting
Direction	`[-1.0, 0.0, 0.0]`
Parts	**(Regions)[Domain:Deck, Domain:Hull, Domain:Stern]**

2. Create a report for the trim:
   a) Right-click the **Reports** node and select **New Report** > **Rigid Body Angle**.
   b) Rename this report to `Trim Down`.
   c) Select the **Trim Down** report and set *Angle Report Option* to **Rotation Y Axis**.
3. Create a report for the thrust created by the virtual disk:
   a) Right-click the **Reports** node and select **New Report** > **Virtual Disk Force**.
   b) Rename this report to `Longitudinal Thrust`.
   c) Select the **Longitudinal Thrust** report and set the following properties:

Property	Setting
Virtual Disk	**Virtual Disk 1**
Force Component Option	**Z-Axis Force**

4. Create a report for the X component of the thrust created by the virtual disk:
   a) Right-click the **Reports** node and select **New Report > Expression**.
   b) Rename this report to `X Component of Thrust`.
   c) Select the **X Component of Thrust** report and set *Definition* to `${LongitudinalThrustReport}*cos(abs($TrimDownReport))`.
5. Create a report for the net force in the X direction:
   a) Right-click the **Reports** node and select **New Report > Expression**.
   b) Rename this report to `X Component of Net Force`.
   c) Select the **X Component of Net Force** report and set *Definition* to `${XComponentofThrustReport} - $DragReport`.
6. Create monitors and plots from all reports:
   a) Multi-select all reports that you just created.
   b) Right-click on any one of the high-lighted items and select **Create Monitor and Plot from Report**.
   c) In the *Create Plot From Reports...* dialog, select **Multiple Plots (one per report)**.
7. Right-click the **Plots > X Component of Net Force Monitor Plot** node and select **Open**.
8. Save the simulation.

If you want to obtain an averaged value of the net force in X direction, right-click the **X Component of Net Force Monitor Plot** node and **Export...** the force values over time to a file in .csv format. You can then import this file into a spreadsheet software of your choice and average the force values. You can use this value to obtain the balancing inlet velocity.

Visualizing the Pressure Distribution

You can obtain a good visualization of the pressure distribution around the ship hull by subtracting the hydrostatic pressure component.

To visualize the pressure distribution:

1. Create a custom field function for pressure:
   a) Right-click the **Tools > Field Functions** node and create a new scalar field function.
   b) Rename this field function to `Pdynamic`.
   c) Select the **Pdynamic** node and set the following properties:

Property	Setting
Function Name	`Pdynamic`
Definition	`$Pressure + $Density * 9.81 * $$Centroid[2]`

2. Create a scalar scene.
3. Edit the **Scenes > Scalar Scene 1 > Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1 > Parts	*Parts*	**(Regions) [Domain: Deck, Domain: Hull, Domain: Stern, (Derived Parts) Centerplane**
Scalar Field	*Function*	**Pdynamic**
Scalar 1	*Contour Style*	**Smooth Filled**

4. Select the **Scalar Scene 1 > Attributes > Update** node and set *Trigger* to **Time Step**.
5. Save the simulation.

Running the Simulation

You are advised to run this tutorial on a multi-core machine with four cores or more. Depending on the specifications of your computer, this tutorial takes several hours or up to a day to complete.

To run the simulation:

1. Click the (**Run**) button on the STAR-CCM+ toolbar.
   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view each of the scenes. The *Residuals* display is created automatically and shows the progress of the solvers.

Visualizing the Results

You can follow the progress of the solution using each of the scenes that you created previously.

The following screenshot shows a close-up of the pressure distribution without the hydrostatic pressure component:



Obtaining the Balancing Inlet Velocity

You specify the inlet velocity such that the forces acting on the ship hull are balanced in the forward direction.

In this simulation, the propeller rotates at a fixed rate. This rotation rate produces a certain thrust and the ship moves forward. The movement of the ship through water induces a resistance due to the viscous and pressure forces acting on the ship hull—the drag force.

For this simulation setup, the ship is not moving freely in the forward (X) direction, but remains in a fixed position. Instead, you prescribe an inlet velocity that produces a relative motion of the ship.

The inlet velocity is not known in advance. You can determine it by running, for example, three simulations with three different inlet velocities. For each simulation, you calculate the average net force acting

on the ship hull in the X direction. By plotting the net force against inlet velocity for each of the three simulations, you can find the velocity for which the net force becomes zero. An example is shown in the following plot:



For this simulation, the inlet velocity at which thrust and drag are balanced is 4.57 m/s. You set this value for the wave velocity in Defining the VOF Wave.

To set up the reports, monitors, and plots to obtain the averaged X net force, see Monitoring and Plotting the Net Force in X Direction.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Setting up a marine self-propulsion simulation.
- Modeling a marine propeller by using the Virtual Disk model.
- Setting up the Virtual Disk model in conjunction with the DFBI model.
- Obtaining the unknown inlet velocity to balance the thrust and drag forces on the ship hull.
- Setting up a VOF multiphase simulation.
- Setting up a flat VOF Wave.

- Setting up boundary-wise wave damping.
- Setting up and running a transient simulation.

Morphing: Cylinder with Boundary Motion

STAR-CCM+ contains a Morphing Motion Model that enables you to define motion on boundary surfaces using several methods.

In this tutorial, two of the available methods are used to impose motion on the outer wall of a simple cylinder. A grid velocity is applied in the first part of the tutorial to cause contraction of the cylinder, and a boundary displacement is applied in the second part to cause expansion. These deformations are illustrated below.



Contents

Importing the Geometry and Naming the Simulation
Preparing the Region and Boundaries
Generating the Volume Mesh
Setting up the Physics
Setting up the Motion
Creating the Local Coordinate System
Creating a Field Function for the Grid Velocity
Setting Boundary Conditions and Values
Setting Initial Conditions
Setting Solver Parameters
Setting Stopping Criteria
Creating Additional Monitors
Visualizing the Solution
Setting Solver Verbosity

Running the Simulation
Visualizing the Results
Expanding the Cylinder
Summary

Importing the Geometry and Naming the Simulation

The first step in this tutorial is to import the cylinder geometry from the Parasolid file supplied.

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment and select the New Simulation option from the menu bar. Click **OK** to accept the default options.
2. Select **File** > **Import** > **Import Surface Mesh** from the menu bar. In the *Open* dialog, navigate to the `doc/startutorialsdata/motion/data` subdirectory of your STAR-CCM+ installation directory and select file `cylinder.x_t`.
3. Click the **Open** button to start the import. The *Import Surface Options* dialog appears.
4. Select **Create new Region** from the **Import Mode** box



5. Click **OK** to import the cylinder.

   STAR-CCM+ provides feedback on the import process in the **Output** window. A new geometry scene is created in the **Graphics** window showing the imported cylinder.
6. Save the new simulation to disk under file name `contraction.sim`.

Preparing the Region and Boundaries

There is only one boundary for the entire surface at this stage, but as you need to have separate surfaces for the inlet and outlet boundaries, it is necessary to split the boundary before creating the volume mesh.

Rename the single region:

1. Rename the **Regions** > **Body 1** node to `Fluid`.

2. Right-click the **Regions** > **Fluid** > **Boundaries** > **Boundary 1** node and select **Split By Angle**.



The *Split Boundaries by Angle* dialog appears.

3. In this dialog do the following:

   a) Make sure that the **Boundary 1** boundary is selected in the *Boundaries* property.

   b) Also make sure that the *Representation* property is set to **Import**.

   c) Accept the default *Angle*, click **Preview** and make sure that splitting would result in three boundaries.

   d) Click **OK**.

   Information on the splitting action is shown in the *Output* window.

Three boundaries are shown in the object tree.

For this case generate a trimmed cell volume mesh aligned in the axial direction, and has two layers of prism cells next to the cylinder wall. You do not wish to have prism cells on the inlet and outlet boundaries, however, and so to avoid this you must define the boundary types before you create the mesh. (The volume meshing tools do not generate prism layers on non-wall surfaces by default).

4. Rename the **Regions** > **Fluid** > **Boundaries** > **Boundary 1** node to `Outlet`.

5. With the **Outlet** boundary still selected, set the *Type* to **Pressure Outlet**.



6. Rename **Boundary 1 2** to `Wall`.
7. Rename **Boundary 1 3** to `Inlet`.
8. Set the *Type* to **Velocity Inlet**.

   The boundary tree is as shown below.



Generating the Volume Mesh

The volume mesh is generated by selecting the required meshing models in the Mesh Continuum, setting the appropriate sizing parameters, and running the mesh generator.

1. Open the **Continua** > **Mesh 1** node that was automatically created when you imported the geometry.



2. Right-click the **Models** node and choose **Select Meshing Models**.
3. In the *Meshing Model Selection* dialog:
   a) Select **Surface Remesher** from the **Surface Mesh** box.
   b) Select **Trimmer** from the **Volume Mesh** box.
   c) Select **Prism Layer Mesher** from the **Optional Models** box.

   The completed dialog appears as shown below.



4. Click **Close**.

Reference sizes must now be set at the global level to make sure cell sizes are appropriate for the simulation.

5. Open the **Continua** > **Mesh 1** > **Reference Values** node and select the **Base Size** node.



6. Set *Value* to `0.1 m`.



Several of the other meshing parameters are set relative to the base size by default, and so can be left unchanged. Change the Prism Layer Thickness, as the default setting results in a fairly abrupt change in cell size between the prism layer and the core mesh.

7. Select the**Continua > Mesh 1 > Reference Values > Prism Layer Thickness > Relative Size** node and set *Percentage of Base* to `65.0`.



The volume mesh can now be created.

8. Click (**Generate Volume Mesh**) in the toolbar.

Setting up the Physics

A default Physics Continuum, **Physics 1**, was created when the volume mesh was generated. This continuum is used to define the physical models required for the simulation.

1. Expand the **Continua** > **Physics 1** node.
2. Right-click the **Models** node and choose item **Select models**.
   The *Physics Model Selection* dialog appears.

Continue by selecting models from the various groups as follows:

3. Select **Implicit Unsteady** from the **Time** box.
4. Select **Liquid** from the **Material** box.
5. Select **Segregated Flow** from the **Flow** box.
6. Select **Constant Density** from the **Equation of State** box.
7. Select **Turbulent** from the **Viscous Regime** box.
8. Select **K-Epsilon Turbulence** from the **Reynolds-Averaged Turbulence** box.
   The *Physics Model Selection* dialog appears as shown below.



9. Click **Close**.

Setting up the Motion

Motion is defined in STAR-CCM+ using motion objects, which are managed by the Motions manager.

To define morphing motion, a Morpher motion object must be created and applied to the region within which morphing takes place.

1. Right-click the **Tools** > **Motions** node and select **New** > **Morphing**.



A new motion node, called **Morphing**, is added to the **Motions** manager.



This motion is now applied to the **Fluid** region.

2. Expand the **Regions** > **Fluid** > **Physics Values** node and select **Motion Specification**.



3. In the **Properties** window, click to the right of the *Motion* property and choose **Morphing**.

   This change adds the mesh morpher solver to the **Solvers** manager.

Creating the Local Coordinate System

Define the cylinder wall motion by a grid velocity in the radial direction. It is necessary to create a local cylindrical coordinate system in which to define the radial velocity as one does not exist by default.

1. Right-click the**Tools > Coordinate Systems > Laboratory > Local Coordinate Systems** node and select **New > Cylindrical**.

   The *Create Cylindrical Coordinate System* dialog appears in the **Edit** window.

   The default parameters are acceptable for this analysis, as the radial direction of the new system is in line with the x axis of the global cartesian system, and the axial direction of the new system is coincident with the z axis of the global system.



2. Click **Create** and then **Close**.

Creating a Field Function for the Grid Velocity

A radial velocity is to be set at the cylinder wall using a field function based on the axial position.

The variation of the radial velocity with axial distance is shown in the figure below:



The equation for this profile is:

$$v_r = \begin{cases} -\frac{1}{8}x_z & \text{for } x_z < 2 \\ -\frac{1}{8}\left(4 - x_z\right) & \text{for } x_z \geq 2 \end{cases}$$

where v_r is the radial velocity in m/s and x_z is the z (or axial) coordinate of the face centroid in m.

The field function for this profile is implemented as follows:

1. Right-click the **Tools** > **Field Functions** node and select **New** > **Vector**.
2. Rename the **User Field Function 1** node to `Cylinder Radial Velocity`.
3. Select the **Cylinder Radial Velocity** node and set the following properties:

Property	**Setting**
Function Name	`CylinderRadialVelocity`
Dimensions	**Velocity**
Definition	`[(($$Position("Cylindrical 1")[2] < 2) ? -0.125*$`

Property	Setting
	`$Position("Cylindrical 1")[2] : -0.125*(4-$ $Position("Cylindrical 1")[2])), 0, 0]`

Note: the function `$$Position("Cylindrical 1")` returns a cartesian vector of the position coordinates in local coordinate system **Cylindrical 1**; the subscript operator `[]` is appended to access a single component of the position vector (with an index of `[0]`, `[1]`, or `[2]`).

Setting Boundary Conditions and Values

The boundary types were set before the volume meshing stage to make sure prism layers were obtained only on the wall.

Set physical values where necessary, and also specify conditions to control morphing of the mesh. Note: Morpher-specific conditions are required even on boundaries that are not explicitly deformed.

1. Expand the **Regions** > **Fluid** > **Boundaries** > **Inlet** > **Physics Conditions** node. Select the **Morpher** node and set the *Method* to **In Plane**.



The **In Plane** condition permits vertices on the inlet boundary to move in the plane of the boundary surface itself, but not orthogonal to it. Selecting this option adds a **Morpher Plane Options** node to the **Physics Values** sub-tree of the **Inlet** boundary.

2. Select the **Inlet** > **Physics Values** > **Morpher Plane Options** node.



3. In the **Properties** window, set the *Damping Factor* to `0.01 m`.
4. Continue by selecting the **Physics Values** > **Velocity Magnitude** > **Constant** node.
5. Enter a *Value* of `10.0 m/s`.

This completes the settings for the **Inlet** boundary. The **Outlet** boundary has the same Morpher settings as the **Inlet** boundary.

6. Select the **Regions** > **Fluid** > **Boundaries** > **Outlet** > **Physics Conditions** > **Morpher** node.
7. Set the *Method* to **In Plane**.
8. Select the **Outlet** > **Physics Values** > **Morpher Plane Options** node.
9. Set the *Damping Factor* to `0.01 m`.

The **Wall** boundary is to be deformed during the simulation. Its conditions are set as follows:

10. Select the **Regions** > **Fluid** > **Boundaries** > **Wall** > **Physics Conditions** > **Morpher** node.



11. Set the *Method* to **Grid Velocity**.
12. Select the **Wall** > **Physics Values** > **Grid Velocity** node.
13. In the **Properties** window, set the *Coordinate System* to **Laboratory-** > **Cylindrical 1** and the *Method* to **Field Function**. The final properties are as shown below.



14. Select the **Grid Velocity** > **Field Function** node and set *Vector Function* to **Cylinder Radial Velocity**.

Setting Initial Conditions

For this tutorial, the only initial condition to set is the flow axial velocity.

Set the axial velocity to 10.0 m/s:

1. Expand the **Continua** > **Physics 1** > **Initial Conditions** > **Velocity** node and select the **Constant** node.



2. In the **Properties** window, set the *Value* to `[0, 0, 10.0] m/s`.

Setting Solver Parameters

This is a simple transient simulation and so an appropriate time step must be set.

As there are no complex flow phenomena a reasonably large time step can be used. Make sure that the mesh motion proceeds at a pace that permits adequate convergence of the fluid solution in each time step.

1. Select the **Solvers** > **Implicit Unsteady** node.
2. In the **Properties** window, set the *Time-Step* to `0.02 s`.

Setting Stopping Criteria

Stopping Criteria is created based on the Monitors for Continuity and Momentum to make sure that the solution does not iterate unnecessarily within a single time-step.

Set the Maximum Physical Time to 1.0 s.

1. Right-click the **Stopping Criteria** node and select **Create New Criterion** > **From Monitor...**.



2. In the *Select Monitor* dialog, select the **Continuity** monitor and click **OK**.
3. Repeat this step to create criteria for **X-momentum**, **Y-momentum**, and **Z-momentum**.
4. Select the **Stopping Criteria > Continuity Criterion > Minimum Limit** node, and then use `Ctrl Click` to select the **Minimum Limit** nodes for the **X-momentum Criterion**, the **Y-momentum Criterion** and the **Z-momentum Criterion**.



5. Set the *Minimum Value* to `0.001`.
6. Select the **Stopping Criteria > Continuity Criterion** node, and then use `Ctrl Click` to select the **X-momentum Criterion**, **Y-momentum Criterion**, and **Z-momentum Criterion** nodes.



7. In the **Properties** window change the *Logical Rule* to **And**.



Having a *Logical Rule* of **And** means that the value of all four monitors (**Continuity**, **X-momentum**, **Y-momentum** and **Z-momentum**) must fall to 0.001 before the simulation assumes that an adequate level of convergence has been achieved and move on to the next time step. It is possible that one of the other Stopping Criteria could activate before reaching this level of convergence, in which case the simulation would move on to the next time step regardless.

8. Select the **Stopping Criteria** > **Maximum Inner Iterations** node and set the *Maximum Inner Iterations* to `10`.
9. Select the **Stopping Criteria** > **Maximum Physical Time** node and make sure that the *Max Physical Time* property is set to `1.0 s`.

Creating Additional Monitors

It is generally useful to create one or two additional monitors to help assess convergence of the simulation.

1. Right-click the **Reports** node and select **New Report** > **Mass Flow Averaged**.



2. Rename the **Mass Flow Averaged 1** node to `Inlet Pressure`.
3. In the **Properties** window, do the following:
   a) Set *Scalar Field Function* to **Pressure**.
   b) Set *Parts* to **Inlet**.

Use the **Inlet Pressure** report to create a monitor and associated plot:

4. Right-click the **Reports** > **Inlet Pressure** node and select **Create Monitor and Plot from Report**.
5. Expand the **Reports**, **Monitors** and **Plots** nodes to see the related items.



6. Select the **Monitors** > **Inlet Pressure Monitor** node.
7. Right-click the **Plots** > **Inlet Pressure Monitor Plot** node and select **Open**.
8. Save the simulation.

Visualizing the Solution

A scalar scene is created to show the deformation of the cylinder wall as the simulation proceeds.

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
2. Rename the **Scalar Scene 1** node to `Cylinder Wall`.
3. Select the **Cylinder Wall** node and use the **scene/plot** button at the top of the simulation tree to switch to the tree for this scene.
4. Select the **Cylinder Wall** > **Displayers** > **Outline 1** node.
5. In the **Properties** window, tick the *Surface* property, change the *Opacity* value to `0.1`, and set the *Representation* to **Import**.



6. Continuing within the **Cylinder Wall** tree, select the **Displayers** > **Scalar 1** > **Parts** node and set *Parts* to **Wall**.
7. Select the **Scalar 1** > **Scalar Field** node and set *Function* to **Centroid** > **Laboratory-** > **Cylindrical 1** > **r**.
8. Click the **simulation** button to switch back to the simulation tree.

Setting Solver Verbosity

The Mesh Morpher solver has an option to display additional information in the **Output** window during runtime. This can be useful in gauging the proportion of CPU time devoted to the Morphing process.

1. Select the **Solvers** > **Mesh Morpher** node.
2. In the **Properties** window, set the *Verbosity* to **High**.
3. Save the simulation.

Running the Simulation

The setup is complete and you can now run the simulation.

To initialize the solution and run the analysis:

1. Click the (**Initialize Solution**) button in the **Solution** toolbar or use the **Solution** > **Initialize Solution** menu item.
2. Click the (**Run**) button in the **Solution** toolbar to run the analysis.

The *Residuals* display is automatically created and shows the solver's progress. If necessary, click the **Residuals** tab to bring the *Residuals* plot into view.

At the end of the simulation:

3. Click the **Cylinder Wall** tab at the top of the **Graphics** window.
   A cylinder wall scene similar to the one shown below appears.



The cylinder would continue to contract if the simulation were run for longer, as no time limiting factor was built into the field function used to specify the grid velocity.

Modify the scales of the **Inlet Pressure Monitor Plot** to view the solution between 0.2 s and 1.0 s.

4. Select the **Plots** > **Inlet Pressure Monitor Plot** > **Axes** > **X Axis** node and set *Minimum* to `0.2` and *Maximum* to `1.0`.
5. Select the **Axes** > **Y Axis** node and set *Minimum* to `18000.0` and *Maximum* to `26000.0`.

The resulting **Inlet Pressure Monitor Plot** is shown below.



Visualizing the Results

A new scene is created to show streamlines of velocity magnitude running through the geometry.

1. Right-click the **Scenes** node and select **New Scene > Geometry**.
2. Rename the **Scenes > Geometry Scene 2** node to `Streamlines`.

A derived part is created for the streamlines which can be added to the new scene.

3. Right-click the **Derived Parts** node and select **New Part > Streamline**.
4. In the *Create Streamline* panel, click **Select...** in the **Seed Parts** box. In the dialog that appears, select **Inlet** and click **OK**.

   The completed panel is as follows:



5. Click **Create** and then **Close**.

6. Select the **Scenes** > **Streamlines** node then click the **Scene/Plot** button to show only the relevant portion of the tree.
7. Select the **Streamlines** > **Displayers** > **Section Stream 1** > **Scalar Field** node and set *Function* to **Velocity: Magnitude**.
8. Select the **Displayers** > **Section Stream 1** node and set Mode to **Tubes** and Width to `0.03`.
9. Finally, select the **Displayers** > **Geometry 1** node.
10. Set the *Opacity* to `0.1`.

    A streamlines scene similar to the one shown below appears.



11. Switch back to the simulation tree using the **simulation** button.
12. Save the simulation.

Expanding the Cylinder

The same exercise can be repeated to model an expansion of the cylinder instead of a contraction. A displacement is specified at the wall boundary instead of a grid velocity.

1. Select **File** > **Save As** and save the simulation with a filename of `expansion.sim`.

2. Right-click **Tools > Field Functions** and select **New > Vector**.
3. Rename the **User Field Function 1** to `Cylinder Wall Displacement`.
4. Select the **Cylinder Wall Displacement** node and set the following properties:

Property	**Setting**
Function Name	`CylinderWallDisplacement`
Dimensions	**Length**
Definition	`[(($$Position("Cylindrical 1")[2] < 2) ? 0.00625*$$Position("Cylindrical 1")[2] : 0.00625*(4-$$Position("Cylindrical 1")[2])), 0, 0]`

The morpher conditions on the cylinder wall must be modified to use the new field function.

5. Select the **Regions > Fluid > Boundaries > Wall > Physics Conditions > Morpher** node.
6. Set the *Method* to **Displacement**.
7. Select the **Physics Values > Morpher: LinearDisplacement** node and set the following properties:
   a) Set *Coordinate System* to **Laboratory- > Cylindrical 1**.
   b) Set *Method* to **Field Function**.
8. Continue by selecting the **Morpher: LinearDisplacement > Field Function** node.
9. Set the *Vector Function* to **Cylinder Wall Displacement**.

This completes the changes necessary for the expansion simulation.

10. Select **Solution > Clear Solution** from the menu bar and click **OK**.
11. Proceed to initialize and run the simulation as before using either the **Solution > Run** option or the (**Run**) button on the toolbar.

At the end of this second simulation:

12. Click the **Cylinder Wall** tab at the top of the **Graphics** window.
13. Use the mouse buttons to zoom out and position the part to obtain a view similar to that shown below.



Summary

This tutorial covered the following features:

- Import of a simple CAD geometry.
- Splitting boundaries and renaming.
- Generating a trimmed cell mesh with prism layers on the wall only.
- Selecting the Morpher motion and applying it to a region.
- Creating a field function to provide a grid velocity vector for the Morpher motion.
- Applying appropriate boundary conditions for the Morpher.
- Creating a scene to display the centroid on the deformed surface of the geometry.
- Creating a streamline derived part and using it to visualize the velocity field.
- Creating a field function to provide a displacement vector for the Morpher motion model.
- Using the Morpher motion model to simulate expansion as well as contraction.

Combustion

The tutorials in this set illustrate various STAR-CCM+ features for simulating problems with combustion

Contents

3-Step Eddy Breakup: Non-Premixed Propane
PPDF Equilibrium: Non-Premixed Propane
PPDF Flamelets: Non-Premixed Hydrogen
Complex Chemistry: Premixed Hydrogen
Surface Chemistry: Methane Deposition
Complex Chemistry: Reduced CPU Time with ISAT
Droplet Combustion: Non-Adiabatic Toluene
Coherent Flame Model: Premixed Propane
Adiabatic PCFM: Partially-Premixed Methane
Progress Variable Model: Sandia Flame D with Participating Media Radiation
Fire and Smoke Wizard: Steckler Room

3-Step Eddy Breakup: Non-Premixed Propane

This tutorial models propane combustion in air using a 3-step Eddy Break-up model:

$$C_3H_8 + 1.5O_2 \rightarrow 3CO + 4H_2$$

$$CO + 0.5O_2 \rightarrow CO_2$$

$$H_2 + 0.5O_2 \rightarrow H_2O$$

The physical properties of the air components (23.3% O2 and 76.7% N2, by mass) and the rest of the reaction components (C3H8, CO, H2, CO2, H2O) are defined as follows:

	O2	N2	C3H8	CO	H2	CO2	H2O
Molecular weights	32.0	28.008	44.1	28.01	2.018	44.01	18.02
Density	Ideal Gas						
Molecular viscosity	1.716×10^{-5}Pas						
Specific Heat	Determined using thermodynamic polynomial functions						
Thermal conductivity	Determined using the Lewis number						

Air enters the combustion chamber through the three air inlets and propane gas enters through the fuel inlet. Both air and fuel are at a pressure of 1 bar and a temperature of 293 K at the inlets.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting up the Models
Setting Material Properties
Defining Reactions

Setting Initial Conditions and Reference Values
Creating Interfaces
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Reporting, Monitoring, and Plotting
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the 3-Step Eddy Breakup: Non-Premixed Propane tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `combustor.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `ebuCombustion.sim`.

Setting up the Models

Models define the primary variables of the simulation, including pressure, temperature, velocity, and what mathematical formulation is used to generate the solution. In this example, the flow involves a turbulent, compressible, multi-component gas whose components are reacting chemically. The Segregated Flow model is used together with the standard K-Epsilon turbulence model and the eddy break-up reaction model.

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically)
Time	**Steady**
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Non-Premixed Combustion**
Non-Premixed Combustion models	**Eddy Break-up**
Flow	**Segregated Flow**
Equation of State	**Ideal Gas (with combustion)** (Selected automatically)
Viscous Regime	**Turbulent**

2. Deactivate **Auto-select recommended models**.
3. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

4. Click **Close**.

   The Standard K-Epsilon model is preferable to the default Realizable Two-layer K-Epsilon model because the near-wall mesh is too coarse to resolve the boundary layer properly. Two-layer models are not recommended for meshes without refined near-wall layers.

5. To review the models, open the **Physics 1** > **Models** node:



To keep run times low, the mesh used in this tutorial is rather coarse. As a result, a solution can be obtained much faster if first-order convection term differencing is used for the momentum, species, energy, and turbulence equations. To change the differencing schemes:

6. Select the **Segregated Flow** node and set *Convection* to **1st-order**.



7. Repeat this process for the **Segregated Species**, **Segregated Fluid Enthalpy**, and **Standard K-Epsilon** nodes.
8. Save the simulation.

Setting Material Properties

Define the components of the mixture and its material properties.

The options included within the **Material Properties** node define the material properties of the gas mixture as a whole. Specify these methods before setting up the individual mixture components.

1. Expand the **Physics 1 > Models > Multi-Component Gas > Material Properties** node.
2. Set the following Properties:

Node	Property	Setting
Dynamic Viscosity	*Method*	**Constant**
Thermal Conductivity	*Method*	**Lewis Number**

You are now ready to define the seven components of the mixture: C3H8, O2, CO, H2, CO2, H2O, and N2. All these species are included in the STAR-CCM+ materials database so detailed specification of their properties is not required.

3. Right-click the **Gas Components** node and choose **Select Mixture Components**.
4. In the *Select Mixture Components* dialog, expand the **Material Databases**, **Standard**, and **Gases** nodes to reveal the available gaseous components.
5. Select the required components: **C3H8 (Propane)**, **CO (Carbon Monoxide)**, **CO2 (Carbon Dioxide)**, **H2 (Hydrogen)**, **H2O (Water)**, **N2 (Nitrogen)**, and **O2 (Oxygen)**.



6. When complete, click **Apply** then **Close**.
7. Right-click the **Gas Components** node and select **Reorder Mixture Components**.
8. In the *Reorder Mixture Components* dialog, drag and drop components until the components are in the following order:
   - C3H8 (Propane)
   - O2 (Oxygen)
   - CO (Carbon Monoxide)
   - H2 (Hydrogen)
   - CO2 (Carbon Dioxide)
   - H2O (Water)
   - N2 (Nitrogen)



9. Click **OK**.

   This order is important for specifying initial and boundary conditions correctly and also to make sure that N2 is the last species (since the species transport equations are solved for n-1 species).

   New nodes corresponding to each of these components appear within the **Gas Components** node.



Once you have specified all components, it is necessary to define their specific heats using thermodynamic polynomials that are retrieved from the materials database.

10. Select the **C3H8 > Material Properties > Specific Heat** node.

    Method must be set to **Thermodynamic Polynomial Data**.



11. Select each of the remaining six gas components in turn and make sure that **Material Properties > Specific Heat** *Method* is set to **Thermodynamic Polynomial Data**.
12. Save the simulation.

Defining Reactions

1. Select the **Models > Eddy Break-up** node.
2. Set *Reaction Control* to **Standard EBU**.
3. Right-click the **Reacting > Reactions** node and select **New Reaction**.

   A **Reaction 1** node appears under the **Reactions** node. Three more nodes, named **Reactants**, **Products**, and **Properties**, are also added under the **Reaction 1** node.

Define the first reaction:

4. Right-click the **Reactants** node and select **Add Reactant > Gas Mixture > C3H8**.



A node named **C3H8** appears underneath the **Reactants** node.

5. Right-click the **Reactants** node and select **Add Reactant > Gas Mixture > O2**

   A node, **O2,** appears.

6. Now add CO and H2 as products by right-clicking the **Products** node and selecting **Add Product > Gas Mixture**.

   This section of the simulation tree appears as shown in the following screenshot.



With all the necessary species added to **Reaction 1**, specify the stoichiometry of the reaction:

7. Select the **C3H8** node and enter `1.0` for the *Stoich. Coeff*.



8. Repeat this procedure for O2, CO and H2, assigning stoichiometric coefficients of `1.5`, `3.0` and `4.0`, respectively.

   The specification of **Reaction 1** is now complete.

9. Follow the same procedure to define the remaining two reactions of the chemical reaction scheme:

$$CO + 0.5O_2 \rightarrow CO_2$$

$$H_2 + 0.5O_2 \rightarrow H_2O$$

10. Save the simulation.

Setting Initial Conditions and Reference Values

Increase the minimum allowable temperature to match the lower temperature limit for the polynomials.

The lower temperature limit for the specific heat polynomials that are imported from the materials database is 200 K. Although this temperature is much lower than you would expect to find in the converged solution, temperatures below this value could appear early in the run. For this reason, it is necessary to increase the minimum allowable temperature to match the lower temperature limit for the polynomials.

1. Select the **Reference Values > Minimum Allowable Temperature** node and set *Value* to `200` K.

The initial condition of the combustor is a stationary flow field consisting entirely of air at 293 K.

2. Select the **Initial Conditions > Species Mass Fraction > Constant** node and set *Value* to:

Property	Setting
C3H8	0.0
O2	0.233
CO	0.0
H2	0.0
CO2	0.0
H2O	0.0
N2	0.767

3. Click **OK.**
4. Set the following properties:

Node	Property	Setting
Static Temperature > Constant	*Value*	`293` K
Turbulence Specification	*Method*	**Intensity + Length Scale**
Turbulence Intensity > Constant	*Value*	`0.05`
Turbulent Length Scale > Constant	*Value*	`0.2`

5. Save the simulation.

Creating Interfaces

All regions and boundaries already have suitable names so you can proceed to create the periodic interface linking the two plane, rectangular, cyclic boundaries.

1. Expand the **Regions** > **Default_Fluid** > **Boundaries** node.
2. Multi-select the **Cyclic1** and **Cyclic2** nodes.
3. Right-click one of the highlighted nodes and select **Create Interface**.

   Two periodic boundary nodes appear under the **Boundaries** node and a node named **Interface 1** appears under the **Interfaces** node.
4. To change this interface to periodic, select the **Interfaces** > **Interface 1** node and set *Topology* to `Periodic`.

Setting Boundary Conditions and Values

All wall boundaries, including the baffle, are adiabatic, no-slip walls. Since **no-slip wall** is the default wall boundary type, no changes are required here. The default settings are also suitable for the outlet, so the four inlets are the only boundary conditions that you specify.

1. Expand the **Air_Inlet1** > **Physics Conditions** node and set the following properties:

Node	Property	Setting
Turbulence Specification	*Method*	**Intensity + Length Scale**
Velocity Specification	*Method*	**Components.**

The air and fuel are made to swirl on entry to the combustor by specifying inlet velocity vectors in a cylindrical coordinate system. To create the latter:

2. Right-click the **Tools** > **Coordinate Systems** > **Laboratory** > **Local Coordinate Systems** node and select **New** > **Cylindrical**.
3. In the *Create Cylindrical Coordinate System* in-place dialog, to accept the default properties, click **Create** then **Close**.

   A cylindrical coordinate system, **Cylindrical 1**, is created.

4. Expand the **Air_Inlet1 > Physics Values** node and set the following properties:

Node	Property	Setting
Velocity	*Coordinate System*	**Laboratory - > Cylindrical 1**
Species Mass Fraction > Constant	O2 *Value*	0.233
	N2 *Value*	0.767
Static Temperature	*Constant*	`293` K
Turbulence Intensity	*Constant*	`0.1`
Velocity	*Constant*	`-11, -24, 42` m/s

The required boundary conditions for **Air_Inlet1** are identical to the **Air_Inlet2** boundary conditions and similar to the **Air_Inlet3** and **Fuel_Inlet** boundary conditions. Therefore, it is convenient to copy the **Air_Inlet1** conditions to the other three inlet boundaries and change settings where necessary.

5. Right-click the **Air_Inlet1** node and select **Copy**.
6. Right-click the nodes of each of the other three inlet boundaries, and select **Paste** each time.
7. Expand the **Air_Inlet3 > Physics Values** node.
8. Select the **Velocity > Constant** node and set *Value* to `0,0,15` m/s.
9. Expand the **Fuel_Inlet > Physics Values** node.
10. Select the **Species Mass Fraction > Constant** node and set *Value* to:

Property	Setting
C3H8	1.0
O2	0.0
CO	0.0
H2	0.0
CO2	0.0
H20	0.0
N2	0.0

11. Click **OK**.
12. Select the **Velocity > Constant** node and set *Value* to `-28,-60,100` m/s.

    Specification of the boundary conditions is now complete.

13. Save the simulation.

Setting Solver Parameters and Stopping Criteria

The default under-relaxation factors for the flow and turbulence equations are suitable. However, reduce the under-relaxation fractions for the species and energy equations to make sure that the solution converges.

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Species	*Under-Relaxation Factor*	`0.8`
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.8`

   It is important that the under-relaxation factors for the species and energy equations are the same to make sure that the two solutions remain synchronized. Other species and energy modeling settings, such as the choice of differencing scheme, are also kept the same.

The simulation is run for 500 iterations, which is sufficient to achieve a steady-state solution. This number can be specified using a stopping criterion.

2. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `500`.

   The solution does not run for more than 500 iterations, unless this stopping criterion is changed or disabled.
3. Save the simulation.

Visualizing the Solution

View the velocity vector field as the solution develops on a plane bisecting the mesh and passing through the axis of symmetry. A scene showing the temperature contours is created for the same plane.

1. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.
2. In the *Create Section* in-place dialog set the following properties:

Property	Setting
Input Parts	**Default_Fluid**
Plane Parameters	origin `0,0,0`
	vector `0,1,0`
Display	**No Displayer**

3. Click **Create** and **Close**.

To create the temperature contour scene:

4. Create a Scalar Scene.
5. Expand the **Scalar Scene 1 > Displayers** node.
6. Set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Derived Parts > plane section**

7. Click **OK**.
8. In the *Graphics* window, right-click the scalar bar and select **Temperature**.



9. Rotate the scalar scene until the view is roughly perpendicular to the plane section (which is colored beige in the geometry scene), and the inlet boundaries are on the left.



10. Save the simulation.

Reporting, Monitoring, and Plotting

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. This procedure requires setting up a report defining the quantity of interest and the parts of the monitored region. A monitor is then defined based on that report. The former also helps to create an appropriate X-Y graph plot.

Report monitors are used in this case to compare the mass of carbon atoms flowing into the solution domain with the mass of such atoms flowing out of it. Once the two monitors reach the same value, it is reasonable to assume that the solution has converged also. To set up these reports, create a field function that calculates the mass flow rate of carbon. In view of the species that is involved in the 3-step reaction mechanism, this field function can be defined as follows:

$$M(C) = M\left(\frac{A_r(C)}{M_r(CO_2)}[CO_2] + \frac{A_r(C)}{M_r(CO)}[CO] + \frac{A_r(C)}{M_r(C_3H_8)}[C_3H_8]\right)$$

where$M(C)$ is the mass flow rate of carbon, M is the total mass flow rate, and $M_r(X)$, $A_r(X)$, and $[X]$ are the molecular weight, the atomic weight, and the mass fraction of species X, respectively. The corresponding field function is therefore:

$$M(C) = \frac{3}{77} M\left(7\left[CO_2\right] + 11\left[CO\right] + 21\left[C_3H_8\right]\right)$$

To create the field function:

1. Right-click the **Tools** > **Field Functions** node, and select **New** > **Scalar**.
2. Rename the **User Field Function 1** node to `Carbon`.
3. Select the **Carbon** node and set the following properties:

Property	Setting
Function name	`Carbon`
Dimensions	**Dimensionless**
Definition	`(3/77)* $FaceFlux*($MassFractionCO2*7+$MassFractionCO*11+ $MassFractionC3H8*21)`

The value of the Carbon field function is positive over outflow boundaries and negative over inflow boundaries. Therefore, to make a side-by-side comparison of carbon inflow and outflow, create another field function for the inflow boundaries that is the negative of the previously created one.

4. As before, right-click the **Field Functions** node and select **New** > **Scalar**.
5. Following the same procedure, change both the node name and function name of the field function to `MinusCarbon` and enter its definition as `0-$Carbon`.

You can now create the reports, monitors, and plots.

6. Right-click the **Reports** node and select **New Report** > **Sum** to create a report node named **Sum 1**
7. Repeat the previous process to create a report node named **Sum 2**.
8. Rename the **Sum 1** node to `Carbon In`.
9. Rename the **Sum 2** node to `Carbon Out`.
10. Select the **Carbon In** node and add the four inlet boundaries to the *Parts* list.
11. Select the **MinusCarbon** field function for the *Scalar Field Function* property.
12. Similarly, edit the properties of the **Carbon Out** node so that the outlet boundary is the only item in the Parts list and the *Scalar Field Function* property is set to **Carbon**.

13. Right-click the **Carbon In** node and select **Create Monitor and Plot from Report**.
14. Right-click the **Carbon Out** node and select **Create Monitor from Report**.

    No plot is created for the Carbon Out Monitor because it is added to the Carbon In Monitor Plot so that the convergence of the two can be observed more easily.
15. Right-click the **Plots** > **Carbon In Monitor Plot** > **Data Series** node and select **Add Data**.
16. In the *Add Data Providers to Plot* dialog, select **Carbon Out Monitor** and click **OK**.
17. Rename the **Carbon In Monitor Plot** node to `Carbon Balance`. Make the same entry for the *Title* property of this node.

To display the empty plot in the **Graphics** window:

18. Right-click the **Plots** > **Carbon Balance** node and select **Open**.
19. The analysis is now ready to be run.
20. Save the simulation.

Running the Simulation

1. To run the simulation, click (**Run**).

Note: If this button is not displayed, use the **Solution** > **Run** menu item. Alternatively, activate the **Solution** toolbar by selecting **Window** > **Toolbars** > **Solution** and then clicking the toolbar button.

The Residuals display is created automatically and shows the progress of the solver. If necessary, click the **Residuals** tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.

The tabs at the top of the **Graphics** window make it possible to select any of the active displays for viewing. During the run, it is possible to stop the solution by clicking (**Stop**) on the toolbar. If you do halt the simulation, to continue the simulation later, click (**Run**). If left alone, the simulation continues until 500 iterations are complete.

2. After the run has finished, click the **Carbon Balance** tab to make sure that the solution has indeed converged.



3. Save the simulation.

Visualizing the Results

1. To view the temperature profile for the converged solution, select the *Scalar Scene 1* display.



2. In the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction > C3H8**.



3. In the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction > CO2**.



4. In the *Scalar Scene 1* display, right-click the scalar bar and select **Velocity > Magnitude**.



Summary

This tutorial introduced the following STAR-CCM+ features:

- Defining models for eddy break-up combustion.
- Defining material properties that are required for multi-component gases.
- Defining chemical reactions.
- Setting initial conditions and reference values.
- Creating interfaces.
- Defining boundary conditions.
- Running the solver for a set number of iterations.
- Analyzing the results.

PPDF Equilibrium: Non-Premixed Propane

In this tutorial, a propane combustion problem is simulated using the Presumed Probability Density Function (PPDF) reaction model for non-premixed flames.

To avoid having to set up the entire case, the simulation file you created in the 3-Step Eddy Break-Up tutorial is reloaded and modified.

The PPDF model assumes that the combustion is mixing-limited and thus represents the state-space variables such as density, temperature and species concentration as a function of mixture fraction and enthalpy. The mixture fraction variable tracks the mixing between the fuel and oxidizer streams. A Presumed Probability Density Function (PPDF) in the shape of a β function represents the local subgrid scale mixing between these two streams. β is evaluated by the mean and variance of the mixture fraction. The mean state space variables are computed by integrating the instantaneous values, the chemical equilibrium, over the PPDF function.

Adiabatic conditions are assumed in this tutorial: as a result, enthalpy and the other state space variables can be computed as a function of mixture fraction alone (see the Adiabatic PPDF Equilibrium Model section in the User Guide). For non-adiabatic boundary conditions and/or gas-phase radiation, use the Non-Adiabatic PPDF Equilibrium model.

Contents

Prerequisites
Loading an Existing Simulation
Setting up the Models
Checking Mixture Components
Generating the PPDF Table
Setting Boundary Conditions and Values
Setting Stopping Criteria
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the PPDF Equilibrium: Non-Premixed Propane tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading an Existing Simulation

For this tutorial, you are provided with an existing simulation file that contains pre-defined objects.

To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to the directory that contains `ebuCombustion.sim`.
5. Select `ebuCombustion.sim` then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `adiabaticPPDF.sim`.

Setting up the Models

Models define the physical variables and phenomena in the simulation. In this simulation, the flow involves a turbulent, compressible, multi-component gas whose components are reacting chemically. The Segregated Flow model is used together with the standard K-Epsilon turbulence model and the PPDF reaction model.

For the physics continuum, **Continua > Physics 1**:

1. Clear models from the *Enabled Models* box so that only the following models remain selected:
   - **Non-Premixed Combustion**
   - **Reacting**
   - **Multi-Component Gas**
   - **Steady**
   - **Gradients**
   - **Three-Dimensional**
2. Select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Enabled Models	**Non-Premixed Combustion** **Reacting** **Multi-Component Gas** **Steady** **Gradients** **Three-Dimensional**
Non-Premixed Combustion models	**Adiabatic PPDF**
PPDF Reaction	**PPDF Equilibrium**
Enabled Models	**PPDF Ideal Gas** (Selected automatically) **Segregated Flow** (Selected automatically) **Turbulent** (Selected automatically) **Reynolds-Averaged Navier-Stokes** (Selected automatically)

3. Deactivate **Auto-select recommended models**.
4. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

5. Click **Close**.
6. To review the models, open the **Physics 1 > Models** node.



7. Select the **Segregated Flow** node and set *Convection* to **1st-order**.



8. Repeat this process for the **Adiabatic PPDF** and **Standard K-Epsilon** nodes.
9. Save the simulation.

Checking Mixture Components

Check the **Gas Components** for the **Multi-Component Gas** model.

The PPDF table for this tutorial includes the same seven gaseous species of the 3-Step Eddy Break-Up tutorial. This setup provides a solution of reasonable accuracy. However, for a more realistic solution you can include additional intermediate species in the table.

To check the gas components corresponding to the species:

1. Open the **Continua** > **Physics 1** > **Models** > **Multi-Component Gas** > **Gas Components** node.



You do not need to change the default component and gas mixture properties.

Generating the PPDF Table

PPDF tables contain information about chemical species and reactions, that the PPDF Equilibrium model needs.

First, change the number of heat loss ratio points in the table:

1. Expand the **Continua > Physics 1 > Models > PPDF Equilibrium** node.
2. Make sure that the following properties are set:

Node	Property	Setting
PPDF Equilibrium Table	*Relative pressure of the mixture*	`0.0 Pa`
Non PPDF Variables > Heat Loss Ratio	*Table Dimension*	`1`

Define the fuel and oxidizer streams:

3. Expand the **Physics 1 > Fluid Stream Manager** node and set the following properties:

Node	Property	Setting
Fuel Stream	*Temperature of the stream*	`293` K

4. Right-click the **Fuel Stream > Components** node and select **Add Fluid Stream Component > C3H8**.
5. In the **Physics 1 > Fluid Stream Manager** node continue to set the following properties:

Node	Property	Setting
Fuel Stream > Components > C3H8	*Mass Fraction*	`1.0`
Oxidizer Stream	*Temperature of the stream*	`293` K

6. Right-click the **Oxidizer Stream > Components** node and select **Add Fluid Stream Component > O2**.
7. Right-click the **Oxidizer Stream > Components** node and select **Add Fluid Stream Component > N2**.

   Two nodes, **O2** and **N2**, appear under the **Oxidizer Stream > Components** node.



8. Expand the **Oxidizer Stream** > **Components** node and set the following properties:

Node	Property	Setting
O2	*Mass Fraction*	`0.233`
N2	*Mass Fraction*	`0.767`

The PPDF Equilibrium table set-up is now complete.

To construct the table:

9. Right-click the **PPDF Equilibrium Table** node and select **Construct**. When the table generation is complete, the message `Generating PPDF Table: Finished` appears under the simulation tree window.



10. Save the simulation.

Setting Boundary Conditions and Values

The boundary types are already defined, however, appropriate property values must be set.

Set the mixture fraction at the fuel inlet region to one.

1. Expand the **Fuel_Inlet** > **Physics Values** node.
2. Edit the **Mixture Fraction Profile** > **Constant** node and set *Value* to [`1.0`].

Setting Stopping Criteria

Set the appropriate stopping criteria for the simulation.

Adjust the maximum number of iterations so that the calculation runs for 600 iterations, which is sufficient for a steady-state solution. This number can be specified using a stopping criterion.

1. Select the **Stopping Criteria** > **Maximum Steps** node.
2. Set *Maximum Steps* to `600`.

The solution does not run beyond 600 iterations, unless this stopping criterion is changed or disabled.

3. Save the simulation.

Running the Simulation

Run the simulation.

To monitor the solution:

1. Right-click the **Scenes** > **Scalar Scene 1** node and select **Open**.

Clear the previous solution:

2. In the menu bar, select **Solution** > **Clear Solution**.
3. In the *Clear Solution* dialog, click **OK**.

To run the simulation:

4. Click (**Run**).
5. When the solution is complete, save the simulation.

Visualizing the Results

This section of the tutorial covers post-processing.

The temperature scene appears as shown in the following screenshot:



1. In the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction > C3H8**.



2. In the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction > CO2**.



The predicted maximum temperature in the model is lower than what would normally be expected from this type of combustor. It is instructive to consider the likely accuracy of the temperature distribution. A good quality measure to use for this purpose is the ratio:

$$\frac{\sqrt{\dot{f}^2}}{\dot{f}} = \frac{\sqrt{g_f}}{\dot{f}}$$

where $\dot{f}$ is the mean mixture fraction and g_f its variance.

Create a field function for this ratio and display its distribution within the combustor geometry:

3. Right-click the **Tools > Field Functions** node and select **New > Scalar**.
4. Rename the **User Field Function 1** node to `Quality Ratio`.
5. Select the **Quality Ratio** node and set the following properties:

Property	Setting
Function Name	`QualityRatio`

Property	Setting
Definition	`($MixtureFraction0 > 1e-4) ? sqrt($MixtureFractionVariance0) / $MixtureFraction0 : 0`

The value of this function in regions of low mixture fraction is set to zero. The function is a useful measure only within the combustion region.

6. Expand the **Scalar Scene 1 > Displayers > Scalar 1** node.
7. Set *Contour Style* to **Filled**.
8. In the *Scalar Scene 1* display, right-click the scalar bar and select **Quality Ratio**.

   The resulting scene is shown in the following screenshot.



The quality ratio values are rather high, implying that the predicted maximum temperature is considerably lower than the adiabatic flame temperature. In a realistic modeling exercise, a considerably

finer mesh is required to increase confidence in the calculated temperatures.

Summary

This tutorial demonstrated how simulate a propane combustion problem using the Presumed Probability Density Function (PPDF) reaction model for non-premixed flames.

The steps covered were:

- Defining an adiabatic PPDF combustion model.
- Generating a PPDF table.
- Setting stopping criteria.
- Running the solver for a set number of iterations.
- Analyzing the results using STAR-CCM+'s visualization facilities.

PPDF Flamelets: Non-Premixed Hydrogen

This tutorial demonstrates how to set up a hydrogen combustion simulation using the **PPDF Laminar Flamelets** reaction model for non-premixed flames.

The model assumes adiabatic conditions (no heat loss) and accounts for non-equilibrium and finite-rate chemistry effects. To avoid having to set up the entire case, the simulation file from the 3-Step Eddy Break-Up tutorial is reloaded and modified.

For adiabatic PPDF, STAR-CCM+ does not directly utilize the physical properties of the fuel (hydrogen in this case) as no additional transport equations requiring these properties are solved. Temperature, density and species mass fractions are evaluated using the β function formulation (see the Adiabatic Equilibrium Model section in the User Guide).

Contents

Prerequisites
Loading an Existing Simulation
Modifying Monitors, Plots, and Reports
Setting up the Models
Constructing the PPDF Flamelet Table
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the PPDF Flamelets: Non-Premixed Hydrogen tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading an Existing Simulation

For this tutorial, start with a simulation file that contains pre-defined objects.

Note: To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally, or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file.

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to the directory that contains the `ebuCombustion.sim` file.
5. Select `ebuCombustion.sim` then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `flameletsPPDF.sim`.

Modifying Monitors, Plots, and Reports

Remove the monitors, plots, and field functions that are not required for this tutorial.

The user field functions **Carbon** and **MinusCarbon** were created using Propane as an input fuel in the 3-Step Eddy Break-Up tutorial. Remove them and the reports, monitors, and plots that depend on them. In STAR-CCM+ dependent objects must be removed first.

1. Delete the following plots, monitors, and field functions nodes:
   - **Plots > Carbon Balance**.
   - **Monitors > Carbon In Monitor**.
   - **Monitors > Carbon Out Monitor**.
   - **Tools > Field Functions > MinusCarbon**.

- **Tools > Field Functions > Carbon**.

Setting up the Models

Models define the primary variables of the simulation. In this example, the flow involves a turbulent, compressible, multi-component gas whose components are reacting chemically.

The Segregated Flow model is used together with the standard K-Epsilon turbulence model and the PPDF reaction model.

1. For the physics continuum, **Continua > Physics 1**:
   a) In the *Physics 1 Model Selection* dialog *Enabled Models* box, clear all of the models except **Steady**, **Gradients** and **Three-Dimensional**.
   b) Select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Enabled Models	**Steady** (Selected automatically) **Gradients** (Selected automatically) **Three-Dimensional** (Selected automatically)
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Non-Premixed Combustion**
Non-Premixed Combustion models	**Adiabatic PPDF**
PPDF Reaction	**PPDF Flamelet (Adiabiatic)**
Enabled Models	**PPDF Ideal Gas** (Selected automatically) **Segregated Flow** (Selected automatically) **Turbulent** (Selected automatically) **Reynolds-Averaged Navier-Stokes** (Selected automatically)

2. Deactivate **Auto-select recommended models**.
3. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

4. Click **Close**.
5. To review the models, open the **Physics 1** > **Models** node.



6. Select the **Segregated Flow** node and set *Convection* to **1st-order**.



7. Repeat this process for the **Adiabatic PPDF** node.

8. Make sure that the **Standard K-Epsilon** node has *Convection* set to **2nd-order**.
9. Save the simulation.

Constructing the PPDF Flamelet Table

PPDF tables contain information about chemical species and reactions, that the PPDF Flamelet model needs.

STAR-CCM+ is not currently capable of producing PPDF Flamelet tables. However, it contains facilities for importing pre-generated PPDF tables or constructing them from flamelet library data, the DARS package produces. In this tutorial, the latter method is used.

Before constructing the table, copy the directory containing the flamelet library files into your working directory:

1. Navigate to the `doc/startutorialsdata/combustion/data` directory of your STAR-CCM+ installation.
2. Copy the directory named `ppdf` into the directory that contains your simulation file.
3. In STAR-CCM+, expand the **Continua > Physics 1 > Models** node.
4. Right-click the **PPDF Flamelet (Adiabiatic) > PPDF Flamelet Table** node and select **Construct**.

   The parameters control the accuracy and size of the look-up table. Their default values are suitable for this tutorial so no changes are necessary.
5. In the *Open* dialog, select the `ppdf` directory and click **Open**.

   The PPDF look-up table is constructed from the flamelet data files in directory `ppdf`. Once complete, the following message appears in the *Output* window:

   ```
   PpdfFlameletTable::generate() end
   ```

The species that are stored in the flamelet libraries is now listed under the **Multi-Component Gas > Gas Components** node.



No further specification of species or reactions is required since all necessary information is contained within the table.

6. Save the simulation.

Setting Boundary Conditions and Values

Specify the appropriate boundary property values.

The only change that is now required is to set the mixture fraction at the fuel inlet region to one.

1. Expand the **Fuel_Inlet** > **Physics Values** node.
2. Select the **Mixture Fraction Profile** > **Constant** node and set *Value* to [`1.0`].



Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

To ensure rapid convergence, reduce the under-relaxation factor for the pressure equation:

1. Select the **Solvers** > **Segregated Flow** > **Pressure** node and set *Under-Relaxation Factor* to `0.15`.

Change the default stopping criterion:

2. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `600`.

   The solution stops after 600 iterations, unless this stopping criterion is changed or disabled.
3. Save the simulation.

Running the Simulation

Run the simulation.

To monitor the solution:

1. Right-click the **Scenes** > **Scalar Scene 1** node and select **Open**.

Clear the previous solution:

2. In the menu bar, select **Solution** > **Clear Solution**.
3. In the *Clear Solution* dialog, click **OK**.

To run the simulation:

4. Click (**Run**).

When the solution is complete:

5. Save the simulation.

Visualizing the Results

This section of the tutorial covers post-processing.

The temperature scene now appears as shown in the following screenshot:



1. In the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction** > **H2**.



2. In the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction > H2O**.



Summary

This tutorial demonstrated how to set up a hydrogen combustion simulation using the **PPDF Laminar Flamelets** reaction model for non-premixed flames.

The steps covered were:

- Defining an adiabatic PPDF combustion model.
- Constructing a PPDF Flamelet table.
- Setting stopping criteria.
- Running the solver for a set number of iterations.
- Analyzing the results using STAR-CCM+'s visualization facilities.

Complex Chemistry: Premixed Hydrogen

This tutorial models hydrogen combustion in air using a complex chemistry model of 9 species and 19 reversible reactions.

The following equations detail these reactions:

$$H_2 + O_2 \rightarrow 2OH$$

$$H_2 + OH \rightarrow H_2O + H$$

$$H + O_2 \rightarrow OH + O$$

$$O + H_2 \rightarrow OH + H$$

$$H + O_2 + M_1 \rightarrow HO_2 + M_1$$

$$H + 2O_2 \rightarrow HO_2 + O_2$$

$$H + O_2 + N_2 \rightarrow HO_2 + N_2$$

$$OH + HO_2 \rightarrow H_2O + O_2$$

$$H + HO_2 \rightarrow 2OH$$

$$O + HO_2 \rightarrow O_2 + OH$$

$$2OH \rightarrow O + H_2O$$

$$H_2 + M_2 \rightarrow 2H + M_2$$

$$O_2 + M \rightarrow 2O + M$$

$$H + OH + M_3 \rightarrow H_2O + M_3$$

$$H + HO_2 \rightarrow H_2 + O_2$$

$$2HO_2 \rightarrow H_2O_2 + O_2$$

$$H_2O_2 + M \rightarrow 2OH + M$$

$$H_2O_2 + H \rightarrow HO_2 + H_2$$

$$H_2O_2 + OH \rightarrow H_2O + HO_2$$

	H2	O2	H	O	OH	HO2	H2O	H2O2	N2
Molecular weights	2.01	32.0	1.005	16.0	17.005	35.005	18.01	24.01	28.0

For each species:

Density	Ideal gas
Molecular viscosity	1.76 x 10-5 Pa.s
Specific heat	Determined via thermodynamic polynomial functions
Thermal conductivity	Determined via the Lewis number

A premixed mixture of hydrogen and air enters the pipe through an inlet at a pressure of 1 bar and a temperature of 1000K.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting up the Models
Importing the Complex Chemistry Definition
Setting Material Properties
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Plotting Simulation Data
Reporting, Monitoring, and Plotting
Running the Simulation
Examining Plots
Summary

Prerequisites

The instructions in the Complex Chemistry: Premixed Hydrogen tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `flameTube.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `flameTube.sim`.

Setting up the Models

Models define the primary variables of the simulation. In this simulation, the flow involves an inviscid, compressible, multi-component gas whose components are reacting chemically.

In this tutorial, the **Segregated Flow** model is used together with the **Homogeneous Reactor** Model.

1. For the physics continuum, **Continua > Physics 1**, select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Enabled Models	**Two Dimensional** (Selected automatically)
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Premixed Combustion**
Premixed Combustion Models	**Homogeneous Reactor**
Complex Chemistry	**Operator Splitting**
Enabled Models	**Dars-CFD Reaction** (Selected automatically)

Group Box	Model
	Dars-CFD Ideal Gas (Selected automatically)
Flow	**Segregated Flow**
	Segregated Species
	Segregated Fluid Enthalpy
Time	**Steady**
Viscous Regime	**Inviscid**
Enabled Models	**Laminar Flame Concept** (Selected automatically)

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



Although the **Dars-CFD Ideal Gas** model was selected automatically, the standard Ideal Gas equation is appropriate for this tutorial.

4. Select the **Dars-CFD Ideal Gas** node and set *Ideal gas equation* to STAR-CCM+.
5. Save the simulation.

Importing the Complex Chemistry Definition

This step involves importing the complex chemistry definition from external files.

1. Expand the **Continua > Physics 1 > Models** node.
2. Select the **Dars-CFD Reaction** node and deactivate *Binary Problem Description*.
3. Right-click the **Dars-CFD Reaction** node and select **Import Complex Chemistry Definition (Dars-CFD format)**.
4. In the *Open* dialog that appears, navigate to the `doc/startutorialsdata/combustion/data` subdirectory of your STAR-CCM+ installation directory.
5. Select `h2o2` and click **Open**.

   In the **Dars-CFD Reaction Model - Properties** window, the *Case Path* property contains the name and path of the folder that contains the complex chemistry definition. The *Is DarsCFD library loaded* property is also activated.



Setting Material Properties

Species are automatically imported with the complex chemistry definition. For each species, verify that the specific heat method uses the **Thermodynamic Polynomial** method.

1. Expand the **Continua > Physics 1 > Models > Multi-Component Gas** node.
2. Expand each of the nine **Gas Components > [Species]** nodes, and verify that **Material Properties > Specific Heat** is set to **Read-Only Thermodynamic Polynomial Data**.



Since reactions are automatically defined with the complex chemistry definition, you are now ready to set initial conditions.

3. Save the simulation.

Setting Initial Conditions

The initial temperature of the domain is set to 1000K. You do not expect to see low or high temperatures for this run, so the temperature range between the minimum and maximum allowable temperatures is sufficient.

1. Expand the **Continua** > **Physics 1** > **Initial Conditions** node.
2. Select the **Static Temperature** > **Constant** node and set *Value* to `1000` K.
3. Select the **Species Mass Fraction** > **Constant** node and set *Value* to the following:

H2	`0.025`
O2	`0.2`
H	`0.0`
O	`0.0`
OH	`0.0`
HO2	`0.0`
H2O	`0.0`

H2O2	`0.0`
N2	`0.775`

4. Click **OK**.
5. Select the **Velocity** > **Constant** node and set *Value* to `1,0,0`.
6. Save the simulation.

Setting Boundary Conditions and Values

Define the boundary types and specify the appropriate property values.

The boundary **symmetry-5** is set to `Symmetry Plane` for this problem. When the mesh was originally created, this boundary was designated to be a wall, so a change of boundary type is required. Also, specify boundary conditions for the inlet and outlet. For the outlet, the backflow conditions are provided, which are identical to the initial conditions specified above. For the inlet, a slightly faster velocity field is specified.

1. Select the **Regions** > **A** > **Boundaries** > **symmetry-5** node and make sure that *Type* is set to **Symmetry Plane**.
2. Expand the **Boundaries** > **Outlet** > **Physics Values** node.
3. Select the **Species Mass Fraction** > **Constant** node and set *Value* to:

H2	`0.025`
O2	`0.2`
H	`0.0`
O	`0.0`
OH	`0.0`
HO2	`0.0`
H20	`0.0`
H2O2	`0.0`
N2	`0.775`

4. Select the **Boundaries** > **Inlet** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Velocity Specification	Method	`Magnitude + Direction`
Physics Values		

Node	Property	Setting
Velocity Magnitude > Constant	*Value*	`1.2`
Static Temperature > Constant	*Value*	`1000`
Species Mass Fraction > Constant	*Value*	Enter the same values as for the outlet

The specification of the boundary conditions is now complete.

5. Save the simulation.

Setting Solver Parameters and Stopping Criteria

The default under-relaxation factors for the flow and turbulence equations are suitable. However, reduce the under-relaxation factors for the species and energy equations to make sure that the solution converges.

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Species	*Under-Relaxation Factor*	`0.5`
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.5`

The under-relaxation factors for the species and energy equations must be the same to make sure that the two solutions remain synchronized. Other species and energy modeling settings, such as the choice of differencing scheme, are also kept the same.

Run the simulation for 300 iterations to achieve a steady-state solution. This number can be specified using a stopping criterion.

2. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `300`.
3. Save the simulation.

Plotting Simulation Data

Create three plots to visualize the solution: temperature, pressure, and mass fractions.

1. Create an X-Y plot.
2. Rename the plot `Temperature Plot`.
3. Select the **Temperature Plot** node and set the following properties:

Node	Property	Setting
Temperature Plot	*Parts*	**A**
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Temperature**
Axes > Y Axis > Title	*Title*	`Temperature (K)`

4. Repeat the steps above and create an X-Y plot called `Pressure Plot`.
5. Select the **Pressure Plot** node and set the following properties:

Node	Property	Setting
Pressure Plot	*Parts*	**A**
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Pressure**
Axes > Y Axis > Title	*Title*	`Pressure (Pa)`

6. Repeat the steps a third time and create an X-Y plot called `Mass Fractions Plot`.
7. Select the **Mass Fractions Plot** node and set the following properties:

Node	Property	Setting
Mass Fractions Plot	*Parts*	**A**
Y Types		
Y Type 1 > Scalar Function	*Field Function*	**Mass Fraction > Mass Fraction of H**

Create additional objects to plot the other mass fractions.

8. Right-click the **Y Types** node and select **New**.

9. Select the **Y Type 2** > **Scalar Function** node and set *Field Function* to the next Mass Fraction scalar, H2.
10. Repeat these steps until you have added Y types for all of the mass fraction scalars. (There must be a total of 9).
11. Save the simulation.

Reporting, Monitoring, and Plotting

Create reports to monitor and plot a comparison of the CPU time to perform complex chemistry calculations with the CPU time for the entire solver calculation.

Complex chemistry calculations are time-consuming and that they take a significant portion of the total solver CPU time. STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. In this tutorial, compare the CPU time to perform complex chemistry calculations with the CPU time for the entire solver calculations. However, complex chemistry calculations also offer the most detailed description of the underlying fluid chemistry. Use complex chemistry to predict physical values of interest accurately.

1. Create two reports, one **Total Solver CPU Time** and the other **Maximum**.
2. Select the **Reports** > **Maximum** node and set the following properties:

Property	Setting
Scalar Field Function	`Complex Chemistry Cpu Time`
Parts	**Region** > **A**

3. Multi-select both reports.
4. Right-click one of the highlighted reports and select **Create Monitor and Plot from Report**.
5. In the dialog that appears, choose **Single Plot**.
6. Rename the **Plots** > **Reports Plot** node to `CPU Plot`.
7. Save the simulation.

Running the Simulation

Run the simulation.

1. To run the simulation, click (**Run**).

The Residuals display is created automatically and shows the progress of the solver. This example looks different from the residuals that are generated in other simulations, since the plot depends on the models that are selected.

The simulation continues until 300 iterations are complete.

2. Save the simulation.

Examining Plots

Examine the results in the plots that you set up.

1. Click the *Temperature* tab.



2. Click the *Pressure* tab.



3. Click the *Mass Fractions* tab.



Summary

This tutorial introduced the following STAR-CCM+ features:

- Defining models for complex chemistry combustion with Dars-CFD.
- Importing the complex chemistry input for Dars-CFD.
- Defining material properties for multi-component gases.
- Setting up plots for total CPU time.

Surface Chemistry: Methane Deposition

This tutorial models methane on platinum deposition using complex chemistry models.

The complex chemistry models are described in the files `chem.inp` and `surf.inp`. In this simulation, a premixed combination of methane and air flows over a platinum plate at a pressure of 1 bar and a temperature of 600 K.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting up the Models
Importing the Complex Chemistry Definition
Setting Material Properties
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Plotting Simulation Data
Reporting and Monitoring with Plots
Running the Simulation
Examining Plots
Appendix: Using a Fortran Compiler for a Binary Description
Summary

Prerequisites

The instructions in the Surface Chemistry: Methane Deposition tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `methaneOnPt.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `methaneOnPt.sim`.

Setting up the Models

Models define the primary variables of the simulation. In this example, the flow involves a laminar, compressible, multi-component gas whose components are reacting chemically with the platinum surface.

In this tutorial, the Segregated Flow model is used together with the Homogeneous Reactor with Surface Chemistry Model.

To select the physics models:

1. For the physics continuum, **Continua > Physics 1**, select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically)
	Gradients (Selected automatically)
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Premixed Combustion**
Premixed Combustion	**Homogeneous Reactor With Surface Chemistry**

Group Box	Model
Enabled Models	**Surface Chemistry Model** (Selected automatically)
Surface Chemistry Reaction	**Operator Splitting (with Surface Reactions)**
Enabled Models	**Dars-CFD Reaction** (Selected automatically)
	Dars-CFD Ideal Gas (Selected automatically)
	Laminar Flame Concept (Selected automatically)
Flow	**Segregated Flow**
Enabled Models	**Segregated Species** (Selected automatically)
	Segregated Fluid Enthalpy (Selected automatically)
Time	**Steady**
Viscous Regime	**Laminar**

2. Click **Close**.
3. To review the models, open the **Physics 1 > Models** node.



4. Save the simulation.

Although the **Dars-CFD Ideal Gas** model was selected automatically, the standard Ideal Gas equation is appropriate for this tutorial.

5. To apply this setting, select the **Dars-CFD Ideal Gas** node and set *Ideal gas equation* to STAR-CCM+.



The decoupled operator splitting algorithm must be enabled and, in steady state simulations, be provided with a starting iteration.

The cold flow simulation is run for the first 50 iterations to achieve a stable steady-state field.

6. Expand the **Continua** > **Physics 1** > **Models** node.
7. Select the **Operator Splitting (with Surface Reactions)** node and set the following properties:

Property	Setting
Source Enabled Trigger	`Iteration`
Begin	`51`

8. Save the simulation.

Importing the Complex Chemistry Definition

Import the complex chemistry definition from Chemkin-formatted files.

There are two options at this stage: importing the chemistry definition in a binary format, which requires the use of a Fortran compiler, or importing the definition in a text format. For this tutorial, the text format is used. For industrial cases, however, you most likely want to use the binary format for significant performance improvements. Instructions on importing a binary format are given in a separate appendix.

1. Expand the **Continua** > **Physics 1** > **Models** node.
2. Select the **Dars-CFD Reaction** node and deactivate *Binary Problem Description*.



3. Copy the following files from `doc/startutorialsdata/combustion/data` to the directory where `methaneOnPt.sim` is located:
   - `chem.inp`
   - `surf.inp`
   - `surf-therm.dat`
   - `therm.dat`
   - `tran.dat`

The two chemkin files `chem.inp` and `surf.inp` define the surface/gas chemistry mechanism. Thermodynamic data comes from `therm.dat` for species in the gas phase, and from `surf-therm.dat` for species in the surface phase. The file `tran.dat` has molecular data.

The next step is to bring the Chemkin formatted files into the simulation.

4. Right-click the **Models** > **Dars-CFD Reaction** node and select **Import Complex Chemistry Definition (Chemkin format)**.

5. In the dialog that appears, set the following properties:

Property	Setting
Project name	`MethaneOnPlatinum`
Chemkin files	`chem.inp` and `surf.inp`
Thermodynamic Properties Files	`therm.dat` and `surf-therm.dat`
Molecular Properties Files	`tran.dat`



6. Click **OK**.

   STAR-CCM+ starts importing the complex chemistry definitions directly from the specified Chemkin-formatted files. The following actions take place:

   - STAR-CCM+ converts Chemkin files to Dars-CFD ASCII input files.
   - The ReduceKin library is called to convert the ASCII input files into case-specific F90 files.
   - The general Dars-CFD library is loaded for use with the case-specific ASCII input files.

See the *Output* window for more details. The following screenshots show various sections from the *Output* window.

```
Output - methaneOnPt
ReduceKin version 1.3.21
  --------
  In ReadSpeciesDataOneMec
   Mechanism name = gasphaseMech
   -------
   In ReadSpeciesDataOnePhase
    Phase Name = NoName
    Species Names = H2              O2              OH              H2O
    All information about phase NoName read.
  --------
  In ReadSpeciesDataOneMec
   Mechanism name = CATALYST
```

```
Output - methaneOnPt
Reading of kinetics for mechanism gasMech done
  -------
  In ReadKineticsOneMech
  Mechanism = CATALYST
  Mechanism header read
  Number of reactions=   25
Reaction name = r1b:H2+2PT(S)=>2H(S)
Reaction name = r2b:2H(S)=>H2+2PT(S)
Reaction name = r3b:O2+2PT(S)=>2O(S)#
Reaction name = r4b:O2+2PT(S)=>2O(S)##
Reaction name = r5b:2O(S)=>O2+2PT(S)
Reaction name = r6b:H2O+PT(S)=>H2O(S)
```

At the end of the import process, STAR-CCM+ creates a number of directories in your current working directory. The main one, which is named `DarsInputFiles`, contains all the dars-cfd input files created by STAR-CCM+. These include the following files:

- ChemicalSources.f90.
- ChemistryInformation.f90.
- ChemistryStateFunctions.f90.
- InputRedKinMec.txt.
- InputRedKinTherm.txt.
- ReactionCalculations.f90.

7. In your current working directory, open the `DarsInputFiles` directory and make sure that the required files have been created successfully.
8. Select the **Dars-CFD Reaction**node and activate *Is DarsCFD library loaded*.



9. Save the simulation.

Setting Material Properties

Verify that the specific heat method for each gas phase species uses the Thermodynamic Polynomial.

Gas and surface species are automatically imported with the complex chemistry definition. When the complex chemistry definition is given in Chemkin format, gas and surface phase reactions are also imported.

1. Expand the **Continua** > **Physics 1** > **Models** > **Multi-Component Gas** node.
2. Open the **Gas Components** node.



The **Gas Components** node contains the species for this problem.

3. Open the node of each species, and within it, the **Material Properties > Specific Heat** node. This exposes the blue method nodes, which are **Read-Only Thermodynamic Polynomial Data**.



Since reactions are automatically defined with the complex chemistry definition, the next step is to set mixture properties.

4. Open the **Multi-Component Gas > Material Properties** node. It contains the mixture properties for this problem.



5. Open the **Dynamic Viscosity**, **Molecular Diffusivity** and **Thermal Conductivity** nodes.

   If any of the subnodes do not match those that are shown for the **Multi-Component Gas > Material Properties** node, select the mixture property node and specify the correct *Method*.

6. Save the simulation.

Setting Initial Conditions

Specify appropriate values for the initial conditions.

The initial temperature of the domain is set to 600 K. Temperatures for this run are not expected to be low or high, so the temperature range between the minimum and maximum allowable temperatures is sufficient.

1. Expand the **Continua > Physics 1 > Initial Conditions** node and set the following properties:

Node	Property	Setting
Static Temperature > Constant	*Value*	600 K

Node	Property	Setting
Species Mass Fraction > Constant	*Value*	O2: `0.228`
		CH4: `0.228`
		N2: `0.752`

Note: For the **Species Mass Fraction > Constant** *Value* settings that are not specified, maintain the remaining default values.



2. Save the simulation.

Setting Boundary Conditions and Values

Specify boundary conditions for the inlet and outlet boundaries.

The boundaries **symmetry plane**, **symmetry plane-2**, and **symmetry plane-3** are set to `Symmetry Plane` for this problem. The boundary **Default Boundary Region** is set to a reactive wall and boundary conditions for surface species are specified. Also, boundary conditions for the inlet and outlet need to be specified. For the outlet, the backflow

conditions are provided, which are identical to the initial conditions specified above. For the inlet, a slightly faster velocity field is specified.

1. Expand the **Regions** > **Default_Fluid** > **Boundaries** > **pressure** node.
2. Set the following properties:

Node	Property	Setting
Physics Values		
Species Mass Fraction > Constant	*Value*	O2: `0.228`
		CH4: `0.02`
		N2: `0.752`

Note: For the **Species Mass Fraction > Constant** *Value* settings that are not specified, maintain the remaining default values.

3. Select the **inlet** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Velocity Specification	*Method*	`Magnitude + Direction`
Physics Values		
Velocity Magnitude > Constant	*Value*	`5.0` m/s
Static Temperature > Constant	*Value*	`600` K
Species Mass Fraction > Constant	*Value*	O2: `0.228`
		CH4: `0.02`
		N2: `0.752`

Note: For the **Species Mass Fraction > Constant** *Value* settings that are not specified, maintain the remaining default values.

4. Expand the **Boundaries** > **Default_Boundary_Region** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Surface Mechanism Option	*Mechanism*	**CATALYST**
Thermal Specification	*Condition*	**Temperature**
Physics Values		
Surface Bulk Fraction > Constant	*Value*	PT(B): `1.0`
Surface Site Fraction > Constant	*Value*	PT(S): `1.0`
Static Temperature > Constant	*Value*	`1000` K

Specification of the boundary conditions is now complete.

5. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Species	*Under-Relaxation Factor*	`0.5`
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.5`

It is important that the under-relaxation factors for the species and energy equations are the same to make sure that the two solutions remain synchronized. Also keep other species and energy modeling settings, such as the choice of differencing scheme, the same.

Run the simulation for 1100 iterations, which is sufficient to achieve a steady-state solution:

2. Select the **Stopping Criteria** > **Maximum Steps** node.

3. Set *Maximum Steps* to `1100`.

   The solution runs for 1100 iterations, unless this stopping criterion is changed or disabled.
4. Save the simulation.

Plotting Simulation Data

Create plots to help you visualize the solution.

Three plots are created: temperature, pressure, and mass fractions. These plots use a type of derived part that is called a line probe.

1. To create a line-probe, right-click the **Derived Parts** node and select **New Part > Probe > Line**.
2. In the *Create Line Probes* in-place dialog, set the following properties:

Property	Setting
Point 1	X: `2e-4`
	Y: `0.0`
	Z: `0.0`
Point 2	X: `2e-4`
	Y: `0.0`
	Z: `0.05`
Display	**No Displayer**

The completed dialog is shown below:



3. Click **Create** then **Close**.

Create the temperature plot:

4. Right-click the **Plots** node and select **New Plot > X-Y**.
5. Rename the **XY Plot 1** node to `Temperature Plot`. Make the same entry for the *Title* property of this node.
6. Expand the **Temperature Plot** node and set the following properties:

Node	Property	Setting
Temperature Plot	*Parts*	**Derived Parts > line-probe**

Node	Property	Setting
X Type > Vector Quantity	*Value*	[`0.0,0.0,1.0`]
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Temperature**

Create a second plot for the CH4 mass fraction and set the properties.

7. Create an X-Y plot and rename the **XY Plot 1** node to `CH4 Plot`. Make the same entry for the *Title* property of this node.



8. Expand the **CH4 Plot** node and set the following properties:

Node	Property	Setting
CH4 Plot	*Parts*	**Derived Parts > line-probe**
X Type > Vector Quantity	*Value*	[`0.0,0.0,1.0`]
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Mass Fraction > Mass Fraction of CH4**

Create a third plot for the Site Surface Fraction of Platinum.

9. Create an X-Y plot and rename the **XY Plot 1** node to `Pt(S) Plot`. Make the same entry for the *Title* property of this node.
10. Expand the **Pt(S) Plot** node and set the following properties:

Node	Property	Setting
Pt(S) Plot	*Parts*	**Default_Fluid: Default_Boundary_ Region**
X Type > Vector Quantity	*Value*	[`0.0,0.0,1.0`]
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Surface Site > Site Fraction > Site Fraction of Pt(S)**

11. Save the simulation.

Reporting and Monitoring with Plots

In this tutorial, the CPU time needed to perform complex chemistry calculations is compared with the CPU time for the entire solver calculations.

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops. This requires setting up a report defining the quantity of interest and the parts of the region to be monitored. A monitor is then defined based on that report. The former also helps to create an appropriate X-Y graph plot. From the plots, it can be observed that complex chemistry calculations are time-consuming and that they take a big portion of the total solver CPU time. However, complex chemistry calculations also offer the most detailed description of the underlying fluid chemistry, and for some problems these calculations are needed to predict physical values of interest accurately.

1. Create two reports, one **Total Solver CPU Time** and one **DarsCFD Quantity Maximum**.
2. Select the **DarsCFD Quantity Maximum** node and set *DarsCFD Variable* to **Cpu Time**.
3. Multi-select both reports then right-click them and select **Create Monitor and Plot from Report**.



4. In the dialog that appears, click **Single Plot**.
5. Rename the **Plots** > **Reports Plot** node to `CPU Plot`.

   The analysis is now ready to be run.
6. Save the simulation.

Running the Simulation

Run the simulation.

1. To run the simulation, click (**Run**).

   The Residuals display is created automatically and shows the progress of the solver. This example looks different from the residuals that are generated in other simulations, since the plot depends on the models that are selected.

   The simulation continues until 1100 iterations are complete.
2. Save the simulation.

Examining Plots

Examine the results in the plots.

- Temperature



- Mass Fraction of CH4



- Site Fraction of Pt(S)



Appendix: Using a Fortran Compiler for a Binary Description

Create a case-specific Fortran library in a folder that is located within your working directory.

The complex chemistry calculations run faster if a case-specific DARS-CFD library is provided at the point of importing the complex chemistry definition. To achieve this, you need to have a compatible Fortran 90 compiler available on your system, to which the path is correctly set before you launch STAR-CCM+.

The steps that are required to generate the binary library are the same as described in section, Importing the Complex Chemistry Definition, with the exception of activating the *Binary Problem Description* property of the **Dars-CFD Reaction** node. When you complete the selection of Chemkin files in the *Select chemkin files for importing* dialog, STAR-CCM+ calls the Fortran 90 compiler after it has converted the ASCII input files into case-specific F90 files. At the end of the process, a case-specific Fortran library has been created in a folder that is located within the working directory. The location of this folder is given in the *Case Path* property of the **Dars-CFD Reaction** node.

Summary

This tutorial demonstrated how to model methane on platinum deposition using complex chemistry models.

The steps covered were:

- Importing the mesh and saving the simulation.
- Defining models for complex chemistry combustion with Dars-CFD.
- Importing the Chemkin-formatted complex chemistry input for Dars-CFD.
- Defining material properties that are required for multi-component gases.
- Setting initial conditions.
- Defining boundary conditions.
- Setting solver parameters and stopping criteria.
- Creating XY-plots for examining the results.
- Setting up monitoring reports and plots for total CPU time, and for the CPU time that is needed to perform complex chemistry calculations.
- Running the solver for a set number of iterations.
- Analyzing the results using STAR-CCM+'s plotting facilities.

Complex Chemistry: Reduced CPU Time with ISAT

This tutorial demonstrates how In-Situ Adaptive Tabulation (ISAT) can be used to reduce the overall computational time to perform complex chemistry calculations.

Complex chemistry models are suitable for simulations where a detailed prediction of non-fuel and oxidizer species is required. However, since some reactions are fast and other ones are slow, any ODE (Ordinary Differential Equation) solver needs to take a very large number of small time sub-steps to solve them. Such a process requires considerable computational resources.

ISAT is a tabulation tool that, instead of calculating the solution to the computationally expensive function (such as the ODE that arises from the reaction system), it tries to approximate that function based on the previously calculated solutions that it has stored in its internal table. It is performed in-situ, which means that it starts with an empty table and adds the points as the calculation progresses.

For each point (in this case, each initial condition for the ODE) that it stores, the ISAT algorithm also calculates the region of accuracy. If the next query point is within this region, it means that the approximated value will satisfy certain tolerance requirements with the respect to the actual solution. Since the region of accuracy is also only approximate, the ISAT algorithm tries to re-evaluate these regions of accuracy around certain points with each new query.

The Sandia piloted methane-air jet flame geometry was developed by Sydney University and is used here to show the required steps.



The two-dimensional axisymmetric geometry consists of the main jet with nozzle diameter of 7.2 mm, burning a premixture of 25% methane and 75% dry air by volume. The pilot inlet surrounds the nozzle. The inlet has a diameter of 18.2 mm and burns a premixture of C2H2, H2, air, CO2, and N2.

The chemistry model uses a reduced 16-species methane combustion mechanism that has been precompiled and distributed with the STAR-CCM+ installation.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting up the Models
Importing the Complex Chemistry Definition
Turning on ISAT Approximation
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Monitoring Simulation Data

Running the Simulation
Examining Plots
Summary
ISAT Sandia Flame Tutorial Bibliography

Prerequisites

The instructions in the Complex Chemistry: Reduced CPU Time with ISAT tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

See the user guide for more information on setting up complex chemistry simulations.

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `sandia.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `sandiaFlameSteady.sim`.

Setting up the Models

Models define the primary variables of the simulation. In this example, the flow involves a turbulent, compressible, multi-component gas whose components are reacting chemically.

The Segregated Flow model is used together with the Homogeneous Reactor model.

1. For the physics continuum, **Continua** > **Physics 1**:
   a) In the *Enabled Models* box, deactivate **Two Dimensional**.
   b) Select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Space	**Axisymmetric**
Enabled Models	**Gradients** (Selected automatically)
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Non-Premixed Combustion**
Non-Premixed Combustion	**Homogeneous Reactor**
Complex Chemistry	**Operator Splitting**
Enabled Models	**Dars-CFD Reaction** (Selected automatically)
	Dars-CFD Ideal Gas (Selected automatically)
Turbulence Chemistry Interactions	**Laminar Flame Concept**
Flow	**Segregated Flow**
Enabled Models	**Segregated Species** (Selected automatically)
	Segregated Fluid Enthalpy (Selected automatically)
Time	**Steady**
Viscous Regime	**Turbulent**

Group Box	Model
Enabled Models	**Reynolds-Averaged Navier Stokes** (Selected automatically)

2. Deactivate **Auto-select recommended models**.
3. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-EpsilonTurbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

4. Click **Close**.
5. To review the models, open the **Physics 1** > **Models** node.



6. Save the simulation.

Importing the Complex Chemistry Definition

Import the pre-compiled complex chemistry definition of some of the most commonly used mechanisms.

A complex chemistry definition is a binary-formatted description of the following:

- Species: elemental composition, molecular weight, transport properties, and thermodynamic properties.
- Reactions and their properties that are used in calculating reaction rates.

The STAR-CCM+ repository contains a pre-compiled definition of some of the most commonly used mechanisms:

- methane combustion
- a CO/H2 flame
- hydrogen combustion
- ethylene
- n-heptane
- n-dodecane
- PRF

By default the specific heat material property uses read-only Thermodynamic Polynomial data as its method.

Import the complex chemistry definition from the STAR-CCM+ repository.

1. Right-click the **Continua > Physics 1 > Models > Dars-CFD Reaction** node and select **Import Case From STAR-CCM+ Repository > Libraries with numerical Jacobian > Methane combustion**.

STAR-CCM+ provides feedback on the import process in the *Output* window.

When the import process is complete, look at the *Properties* window. The name and path of the folder containing the complex chemistry definition appears in the *Case Path* property of the Dars-CFD Reaction Model node.

2. Check that the *Is DarsCFD library loaded* property is activated.



Turning on ISAT Approximation

Activate ISAT approximation and specify a constant pressure specification.

1. Expand the **Continua** > **Physics 1** > **Models** > **Dars-CFD Reaction** node.
2. Select the **Dars-CFD Properties** > **Approximation Option** node and set *Method* to `In Situ Adaptive Tabulation`.



The **Approximation Option > In Situ Adaptive Tabulation** node appears in the simulation tree.



3. Since this is a constant-pressure case, select the **In Situ Adaptive Tabulation** node and change *Pressure Specification* from **Variable** to **Constant**.

   If you leave this property set to **Variable** for a constant-pressure case, you get a larger table and slower calculation. If you use a setting of **Constant** for a variable-pressure case, reactions get calculated only with the initial pressure value, yielding incorrect results.

4. Leave the other properties at the defaults.

   The default settings are sufficient to cover a generic case with complex chemistry. These other properties do the following:

   - The *Restart* property facilitates resuming work on an existing simulation with ISAT. When activated, it loads the ISAT table from a file that was saved in the previous session.
   - The *Absolute Tolerance* property sets limits on the Ellipsoid of Accuracy for the ISAT approximation.
   - The *Table Size (MB)* property sets a limit on the size of the ISAT table. This limit restricts computational expense, but it also reduces the number of query points, and with it, accuracy of the results.

   In the current version of STAR-CCM+, the *Time Step Specification* property is preset to one of the following:

   - In unsteady simulations, this property is constant.

- In steady simulations, STAR-CCM+ uses a special iterative-tables technique to accommodate variation of the time-step. This variation is not suitable for ISAT tabulation in most steady cases.

The expert properties are mainly for administering technical details of a simulation that is run in parallel. These settings can also be left at the defaults.

Setting Initial Conditions

Adjust the conditions for the start of the calculation: temperature, species mass fraction, the turbulence specification method, and the axial velocity of the gas.

1. Expand the **Continua** > **Physics 1** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Species Mass Fraction > **Constant**	*Value*	`0.233` for **O2**
		`0.767` for **N2**
Static Temperature > **Constant**	*Value*	`293`K
Turbulence Specification	*Method*	**K + Epsilon**
Turbulent Dissipation Rate > **Constant**	*Value*	`50000` m2/s3
Turbulent Kinetic Energy > **Constant**	*Value*	`30` J/kg
Velocity > **Constant**	*Value*	[`0.9, 0, 0`] m/s

Setting Boundary Conditions and Values

In this section you set the boundary conditions for the inlets, outer boundaries, and outlet boundaries. For the outer and outlet boundaries, you set the backflow conditions to the extrapolation method, which extrapolates the flow direction from the interior of the domain. This case requires extrapolation because the flow is not necessarily normal to the boundary.

Contents

Inlet-1 Boundary
Inlet-2 Boundary
Inlet-3 Boundary
Outer Boundary
Outlet Boundary

Inlet-1 Boundary

Set the **inlet-1** boundary physics values.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-1** > **Physics Values** node and set the following properties:

Node	Property	Setting
Species Mass Fraction > **Constant**	*Value*	`0.1563` for **CH4**
		`0.1965` for **O2**
		`0.6472` for **N2**
Static Temperature > **Constant**	*Value*	`293` K
Turbulence Intensity > **Constant**	*Value*	`0.1`
Turbulent Viscosity Ratio > **Constant**	*Value*	`100`
Velocity Magnitude > **Constant**	*Value*	`49.6` m/s

3. Click **Close**.

Inlet-2 Boundary

Set the **inlet-2** boundary physics values.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-2** > **Physics Values** node and set the following properties:

Node	Property	Setting
Species Mass Fraction > **Constant**	*Value*	`0.095` for **H2O**
		`0.117` for **CO2**

Node	Property	Setting
		`0.054` for **O2**
		`0.734` for **N2**
Static Temperature > Constant	*Value*	`1880` K
Turbulence Intensity > Constant	*Value*	`0.1`
Turbulent Viscosity Ratio > Constant	*Value*	`100`
Velocity Magnitude > Constant	*Value*	`11.4` m/s

3. Click **Close**.

Inlet-3 Boundary

Set the **inlet-3** boundary physics values.

1. Expand the **Regions > fluid > Boundaries** node.
2. Edit the **inlet-3 > Physics Values** node and set the following properties:

Node	Property	Setting
Species Mass Fraction > Constant	*Value*	`0.233` for **O2**
		`0.767` for **N2**
Static Temperature > Constant	*Value*	`293` K
Turbulence Intensity > Constant	*Value*	`0.1`
Turbulent Viscosity Ratio > Constant	*Value*	`100`
Velocity Magnitude > Constant	*Value*	`0.9` m/s

3. Click **Close**.

Outer Boundary

Change the **Backflow Direction Specification** method to **Extrapolated**.

1. Expand the **Regions** > **fluid** > **Boundaries** > **outer** > **Physics Conditions** node.
2. Select the **Physics Conditions** > **Backflow Direction Specification** node and set *Method* to **Extrapolated**.

Set the outer boundary physics values.

3. Expand the **Regions** > **fluid** > **Boundaries** node.
4. Edit the **outer** > **Physics Values** node and set the following properties:

Node	Property	Setting
Species Mass Fraction > **Constant**	*Value*	`0.233` for **O2**
		`0.767` for **N2**
Static Temperature > **Constant**	*Value*	`293` K
Turbulence Intensity > **Constant**	*Value*	`0.01`
Turbulent Viscosity Ratio > **Constant**	*Value*	`10`

5. Click **Close**.

Outlet Boundary

Change the **Backflow Direction Specification** method to **Extrapolated**.

1. Expand the **Regions** > **fluid** > **Boundaries** > **outlet** > **Physics Conditions** node.
2. Select the **Physics Conditions** > **Backflow Direction Specification** node and set *Method* to **Extrapolated**.

Set the outlet boundary physics values.

3. Expand the **Regions** > **fluid** > **Boundaries** node.
4. Edit the **outlet** > **Physics Values** node and set the following properties:

Node	Property	Setting
Species Mass Fraction > **Constant**	*Value*	`0.233` for **O2**
		`0.767` for **N2**
Static Temperature > **Constant**	*Value*	`293` K

Node	Property	Setting
Turbulence Intensity > Constant	*Value*	`0.01`
Turbulent Viscosity Ratio > Constant	*Value*	`10`

5. Click **Close**.
6. Save the simulation.

Setting Solver Parameters and Stopping Criteria

The default under-relaxation factors for the flow equation are suitable. To ensure solution convergence, reduce the under-relaxation factors for species, energy, and turbulence.

1. To set up the energy and turbulence equations, expand the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Species	*Under-Relaxation Factor*	`0.8`
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.8`
K-Epsilon Turbulence	*Under-Relaxation Factor*	`0.6`

   It is important that the under-relaxation factors for the species and energy equations are the same to ensure that the two solutions remain synchronized.

Set the maximum number of steps to 2200, which is sufficient to achieve a steady-state solution.

2. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `2200`.

Monitoring Simulation Data

Create reports to monitor the total complex chemistry CPU time and the volume average of temperature on the fluid region.

Create a plot for the axial temperature distribution.

Contents

Create Reports for Monitors and Plots
Create a Temperature Plot

Create Reports for Monitors and Plots

Create a number of reports for monitoring ISAT performance.

1. Right-click the **Reports** node and select **New Report > DarsCFD Quantity Maximum**.

   This type of report takes a Dars-CFD variable that you specify and calculates the maximum.

2. With the **Dars-CFD Quantity-Maximum** node selected, set the *DarsCFD Variable* property to `Cpu Time`.

   The name of the report node changes automatically to **Complex Chemistry Cumulative Cpu Time-Maximum**.



3. Create three more reports using the same **New Report > DarsCFD Quantity Maximum** action. Set their *DarsCFD Variable* property to the following quantities, all of which are important for evaluating the performance of ISAT:

   a) **Adds** -- points stored to the table.

b) **Grows** -- extensions of the region of accuracy around certain points based on new data points (queries Y and their ODE solutions, F(Y)).

c) **Retrieves** -- for query point Y0, a returned approximation of F(Y0) using F(Y_in_table).

4. Multi-select all of the **Reports** nodes.
5. Right-click one of the highlighted nodes and select **Create Monitor and Plot from Report**.



6. In the *Create Plot From Reports* dialog, click **Multiple Plots (one per report)**.

Create a Temperature Plot

Create an X-Y plot to show the axial temperature distribution.

1. Right-click the **Plots** node and select **New Plot > X-Y**.
   An **XY Plot 1** node is created under the **Plots** node.
2. Rename the **XY Plot 1** node to `TAxis`.
3. Select the **TAxis** node and set the following properties:

Node	Property	Setting
TAxis	*Parts*	**Regions > fluid > Boundaries > axis**
X Type > Vextor Quantity	*Value*	[`1.0, 0.0, 0.0`]

Node	Property	Setting
Y Types > Y Type 1 > Scalar Function	*Field Function*	**Temperature**

4. Save the simulation.

Running the Simulation

The simulation can now be run.

1. To run the simulation, click (**Run**).

You can monitor the iteration progress in the *Output* window. The *Residuals* display is created automatically and shows the progress of the solver. If necessary, click the *Residuals* tab to bring the *Residuals* plot into view.

2. Save the simulation.

Examining Plots

Examine the results in the plots that you set up.

Click the tab of each of the following plots to view them:

1. *TAxis*



2. *Complex Chemistry Cumulative Cpu Time-Maximum Monitor Plot*



3. *DarsCFD-ISAT Cumulative Adds-Maximum Monitor Plot*



4. *DarsCFD-ISAT Cumulative Grows-Maximum Monitor Plot*



You can see that when the solution reaches approximately 200 iterations, the numbers of adds and grows become constant (the rate of increase drops down to almost zero). This result means that ISAT did not have to calculate any more points, but could retrieve them all from the table. At this point in the process, your simulation becomes much less computationally expensive.

Look at the residuals. At 200 iterations, the simulation did not yet reach the steady state. However, with more computing power available from this point on, you can reach the steady state much faster.

Given that the number of grows is of the order of magnitude larger than the number of adds, ISAT is constantly reevaluating its regions of accuracy to ensure that it is efficiently and accurately retrieving the points from the table.

Unlike the adds and grows, the number of retrieves continues to increase for the entire simulation run. However, retrieves are computationally inexpensive.



Further Practice

To examine ISAT further, run the base simulation without it. Use the results of this other simulation to measure the effectiveness of ISAT.

Summary

This tutorial demonstrated how to use In-Situ Adaptive Tabulation (ISAT) to reduce the overall computational time to perform complex chemistry calculations.

The steps covered were:

- Using complex chemistry systems from the STAR-CCM+ repository.
- Using ISAT for speeding up complex chemistry calculations.

ISAT Sandia Flame Tutorial Bibliography

Masri, A. R., Dibble, R. W., and Barlow, R. S. 1996 “The structure of turbulent nonpremixed flames revealed by Raman-Rayleigh-LIF measurements”, Prog. Energy Combust. Sci. 22(4) pp. 307-362 .

Droplet Combustion: Non-Adiabatic Toluene

In this tutorial, a droplet combustion case is set up using the Non-Adiabatic PPDF Equilibrium reaction model for non-premixed flames. The fuel is toluene (C7H8) and is injected into the combustor in the form of droplets that then evaporate into combustible vapor.

Start by reloading and modifying the simulation file that was created in the 3-Step Eddy Break-Up tutorial.

Contents

Prerequisites
Loading an Existing Simulation
Modifying Mixture Components
Modifying Monitors, Plots, and Reports
Modifying the Models
Setting the Convection Differencing Scheme
Setting Initial Conditions
Setting Fuel and Oxidizer Stream Properties
Generating the PPDF Table
Selecting Lagrangian Phase Models
Setting Lagrangian Phase Model Properties
Checking Lagrangian Phase Boundary Conditions
Setting up the Injector
Setting Boundary Conditions and Values
Setting Solver Variables and Stopping Criteria
Running the Simulation
Visualizing the Results
Displaying Particle Tracks
Summary

Prerequisites

The instructions in the Droplet Combustion: Non-Adiabatic Toluene tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Loading an Existing Simulation

For this tutorial, you start with a simulation file.

Note: To start this tutorial, you can open one of two versions of the simulation file: the one that you created locally (if applicable), or a previously solved simulation file that is available online. The following instructions apply to a locally created simulation file.

Load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to the`ebuCombustion.sim` file in the directory that you created locally.
5. Select the `ebuCombustion.sim` file then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `nonAdiabaticPPDF.sim`.

Modifying Mixture Components

Change the fuel from propane (C3H8) to toluene (C7H8). Also enable solid carbon (C(S)).

To do this, first delete the dependent reactions:

1. Expand the **Continua > Physics 1 > Models > Reacting > Reactions** node.
2. Delete **Reaction 1**, **Reaction 2**, and **Reaction 3**.
3. Expand the **Continua > Physics 1 > Models > Multi-Component Gas > Gas Components** node.
4. Delete **C3H8**.
5. Right-click the **Gas Components** node and select **Select Mixture Components**.

6. In the *Select Mixture Components* dialog, expand the **Material Databases**, **Standard**, and **Gases** nodes to reveal the components in the materials database.
7. Select the components C(S) and C7H8.
8. When complete, click **Apply** then **Close**.
9. Right-click the **Gas Components** node and select **Reorder Mixture Components**.



10. In the *Reorder Mixture Components* dialog, drag and drop components until you get the components in the following order:
    - C(S)
    - CO
    - H2O
    - H2
    - C7H8
    - O2
    - CO2
    - N2
11. Click **OK**.

    This order is important for specifying initial and boundary conditions correctly and also to make sure that N2 is the last species (since the species transport equations are solved for n-1 species).

The nodes of the components appear as shown below.



The default component and mixture properties can be used. For reference, check one of these:

12. Expand the **Gas Components > C7H8 > Material Properties > Specific Heat** node and note that it is set to **Thermodynamic Polynomial Data**.
13. Save the simulation.

Modifying Monitors, Plots, and Reports

Remove the monitors, plots, and reports that were created in the 3-Step Eddy Break-Up tutorial.

The user field functions **Carbon** and **MinusCarbon** were created using Propane as an input fuel in the 3-Step Eddy Break-Up tutorial. Remove them and the reports, monitors, and plots that depend on them. In STAR-CCM+ dependent objects must be removed first.

Start by removing the **Carbon Balance** plot, then the monitors, and finally the field functions.

1. Delete the following nodes:

- **Plots > Carbon Balance**
- **Monitors > Carbon In Monitor** and **Carbon Out Monitor**.
- **Reports > Carbon In** and **Carbon Out**.
- **Tools > Field Functions > MinusCarbon** and **Carbon**.

Modifying the Models

Select the appropriate models.

First clear the models that are not required:

1. For the physics continuum, **Continua > Physics 1**:
   a) In the *Enabled Models* box, deselect the following models:
      - **High y+ Wall Treatment**
      - **Standard K-Epsilon**
      - **K-Epsilon Turbulence**
      - **Reynolds-Averaged Navier-Stokes**
      - **Turbulent**
      - **Segregated Fluid Enthalpy**
      - **Ideal Gas (with combustion)**
      - **Segregated Species**
      - **Segregated Flow**
      - **Eddy Break-up**
   b) Select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Non-Premixed Combustion models	**Non-Adiabatic PPDF**
PPDF Reaction	**PPDF Equilibrium**
Ideal Gas Equation of State for PPDF	**PPDF Ideal Gas**
Enabled Models	**Segregated Flow** (Selected automatically) **Turbulent** (Selected automatically)

Group Box	Model
	Segregated Fluid Enthalpy (Selected automatically) **Reynolds-Averaged Navier-Stokes** (Selected automatically)

2. Deactivate **Auto-select recommended models**.
3. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**
Optional Models	**Lagrangian Multiphase**

4. Click **Close**.
5. To review the models, open the **Physics 1** > **Models** node.



6. Save the simulation.

Setting the Convection Differencing Scheme

Change the convection term differencing scheme to first-order for the non-adiabatic PPDF, momentum, enthalpy, and turbulence equations.

1. Expand the **Continua** > **Physics 1** > **Models** node.
2. Select the **Non-Adiabatic PPDF** node and set *Convection* to **1st-order**.
3. Repeat this process for the **Segregated Flow**, **Segregated Fluid Enthalpy**, and **Standard K-Epsilon** nodes.
4. Save the simulation.

Setting Initial Conditions

Adjust the conditions for the start of the calculation.

Set the initial temperature to 600 K.

1. Expand the **Continua** > **Physics 1** > **Initial Conditions** node.
2. Select the **Static Temperature** > **Constant** node and set *Value* to `600` K.

Setting Fuel and Oxidizer Stream Properties

Specify the property settings for the fuel and oxidizer streams.

To set the fuel stream properties:

1. Expand the **Continua** > **Physics 1** > **Fluid Stream Manager** node.
2. Select the **Fuel Stream** node and set *Temperature of the stream* to `350` K.
3. Right-click the **Fuel Stream** > **Components** node and select **Add Fluid Stream Component** > **C7H8**.
4. Set *Mass Fraction* to `1.0`.

To set the oxidizer stream properties:

5. Select the **Fluid Stream Manager** > **Oxidizer Stream** node and set *Temperature of the stream* to `600` K.
6. Right-click the **Oxidizer Stream** > **Components** node and select **Add Fluid Stream Component** > **N2**.
7. Set *Mass Fraction* to `0.767`.

8. Similarly, add O2 as the second component of the oxidizer stream and set the *Mass Fraction* of O2 to `0.233`.

Generating the PPDF Table

PPDF tables contain information about chemical species and reactions, that the PPDF Equilibrium model needs.

Before generating the PPDF table, specify the PPDF equilibrium table properties:

1. Expand the **Continua > Physics 1 > Models** node.
2. Edit the **PPDF Equilibrium** node and set the following properties:

Node	Property	Setting
PPDF Equilibrium Table > Non PPDF Variables > Heat Loss Ratio	*Table Dimension*	`11`
PPDF Equilibrium Table	*Type of the Equilibrium Algorithm*	`Dars Equilibrium`
PPDF Equilibrium Table > PPDF Variables > Mixture Fraction	*Table Dimension*	`71`
PPDF Equilibrium Table > PPDF Variables > Mixture Fraction Variance	*Table Dimension*	`41`

This completes the PPDF Equilibrium table set-up. Generate the table:

3. Right-click the **Continua > Physics 1 > Models > PPDF Equilibrium > PPDF Equilibrium Table** node and select **Construct**.

The message `Generating PPDF Table: Finished` appears under the *Simulation* window once the table generation is complete.

4. Save the simulation.

Selecting Lagrangian Phase Models

Select models representing the characteristics of the Lagrangian phase.

First create the phase.

1. Expand the **Continua** > **Physics 1** > **Models** > **Lagrangian Multiphase** node.
2. Right-click the **Lagrangian Phases** node and select **New**.
3. Rename the **Phase 1** node to `Fuel`.
4. Right-click the **Fuel** > **Models** node and select **Select models**.
5. In the *Fuel Model Selection* dialog, select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Enabled Models	**Residence Time** (Selected automatically)
Particle Type	**Material Particles**
	Spherical Particles (Selected automatically)
	Pressure Gradient Force (Selected automatically)
Material	**Liquid**
Equation of State	**Constant Density**
Mass Transfer	**Quasi-Steady Evaporation**
	Energy (Selected automatically)
Optional Particle Forces	**Drag Force**
Optional Models	**Turbulent Dispersion**
	Two-Way Coupling
Track Sampling	**Track File**

When the **Material Particles** model is selected, a **Fuel** > **Boundary Conditions** node appears in the object tree automatically.

6. Click **Close**.
7. Expand the **Fuel** > **Models** node.

   The selected models now appear within that node.



Setting Lagrangian Phase Model Properties

Set the correct fuel properties by replacing H2O with toluene (C7H8).

1. Expand the **Continua > Physics 1 > Models > Lagrangian Multiphase** node.
2. Right-click the **Fuel > Models > Liquid > H2O** node and select **Replace with**.

The *Replace Material* dialog appears.

3. Expand the **Material Databases > Standard > Liquids** node and select **C7H8 (Toluene)**



4. Click **OK**.

Set the properties for the Lagrangian Phase models.

5. Select the **Fuel** > **Models** > **Quasi-Steady Evaporation** node and set *Vapor Component* to **C7H8**.
6. Expand the **Models** > **Liquid** > **C7H8** > **Material Properties** node and set the following properties:

Node	Property	Setting
Saturation Pressure	*Method*	**Wagner Equation**
Latent Heat of Vaporization	*Method*	**Clausius-Clapeyron**

The default settings for the remaining properties are acceptable for use in this case.

Checking Lagrangian Phase Boundary Conditions

The wall boundary interaction mode for the Lagrangian Phase is set to Rebound by default.

To check this setting:

1. Expand the **Continua** > **Physics 1** > **Models** > **Lagrangian Multiphase** > **Lagrangian Phases** node.
2. Select the **Fuel** > **Boundary Conditions** > **Wall** > **Physics Conditions** > **Mode** node and make sure that *Active Mode* is set to **Rebound.**



3. Save the simulation.

Setting up the Injector

As droplets are to be injected along a line, it is necessary to create a line part.

The liquid droplets of the dispersed Lagrangian phase are introduced into the combustor via an injector. The injector is located along a short line close to the combustor head. The injector defines the initial state of the liquid droplets.

Create a line part:

1. Right-click the **Scenes > Geometry Scene 1** node and select **Open**.
2. Right-click the **Derived Parts** node and select **New Part > Probe > Line**.
3. In the *Create Line Probe* panel, set the following:

Property	Setting
Input Parts	**Default Fluid**
Point 1	X: `0.00453154` m
	Y: `-0.00211309` m
	Z: `0.02` m
Point 2	X: `0.00453154` m
	Y: `0.00211309` m
	Z: `0.02` m
Resolution	`20`
Display	**No Displayer**

4. Click **Create** then **Close**.

To view this part in the *Graphics* window:

5. Select the **Derived Parts > Line Probe** node.

The location of the point is highlighted in the *Graphics* window:



Create the injector:

6. Right-click the **Injectors** node and select **New**.
7. Rename the **Injector 1** node to `Fuel`.
8. Edit the **Injectors > Fuel** node and set the following properties:

Property	Setting
Type	**Part Injector**
Inputs	**Derived Parts > Line Probe**
Lagrangian Phase	**Fuel**

9. In the *Edit* dialog, expand the **Conditions** and **Values** nodes.
10. Set the following properties:

Table 4: Conditions

Node	Property	Setting
Conditions		
Flow Rate Distribution	*Method*	**Per Injector Point**

Node	Property	Setting
Velocity Specification	*Method*	**Magnitude + Direction**
Values		
Injection Direction > Constant	*Value*	`[0.447214, 0.0, 0.894427]`
Mass Flow Rate > Constant	*Value*	`4.152E-4` kg/s
Particle Diameter > Constant	*Value*	`2.5E-4` m
Particle Temperature > Constant	*Value*	`350` K
Velocity Magnitude > Constant	*Value*	`4.4721` m/s

11. Click **Close**.

Setting Boundary Conditions and Values

As the fuel is injected into the combustor as droplets, the boundary that was defined previously as **Fuel_Inlet** is no longer required.

Redefine this boundary as a wall and combine it with the boundary named **In_Wall**.

1. Select the **Regions** > **Default_Fluid** > **Boundaries** > **Fuel_Inlet** node and set *Type* to `Wall`.
2. Multi-select the **In_Wall** and **Fuel_Inlet** nodes.
3. Right-click the highlighted nodes and select **Combine**.

Note: The combined boundary retains the name **In_Wall**.

Use the multi-edit function to specify the temperature for the three air inlets.

4. Multi-select the **Air_Inlet1**, **Air_Inlet2**, and **Air_Inlet3** nodes.
5. Right-click the highlighted nodes and select **Edit**.
6. In the *Multi-edit* dialog, expand the **Physics Values** > **Static Temperature** > **Constant** node and set *Value* to `600` K.

The completed dialog appears as shown:



7. Click **Close**.

Setting Solver Variables and Stopping Criteria

Set the appropriate solver parameters and stopping criteria.

Change the fluid under-relaxation factor.

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.8`
Lagrangian Multiphase > Two-Way Coupling	*Linearize Sources*	deactivated

The maximum number of iterations is set to 500 which is suitable.

2. Save the simulation.

Running the Simulation

Clear the previous solution and run the simulation.

Monitor the solution:

1. Right-click the **Scenes > Scalar Scene 1** node and select **Open**.

Clear the previous solution:

2. In the menu bar, select **Solution > Clear Solution**.
3. In the *Clear Solution* dialog, click **OK**.

Run the simulation:

4. To run the simulation, click (**Run**).
   The Residuals display is created automatically and shows the progress of the solver. The simulation stops after 500 iterations.
5. Save the simulation.

Visualizing the Results

View temperature and mass fraction contours for the converged solution.

1. To view the temperature contour for the converged solution, in the *Graphics* window, click the *Scalar Scene 1* tab.



2. To view the mass fraction of C7H8 contour for the converged solution, in the *Scalar Scene 1* display, right-click the scalar bar and select **Mass Fraction > C7H8**.



Displaying Particle Tracks

This section of the tutorial covers post processing.

To display the particle tracks that were recorded in the track file:

1. Right-click the **Tools > Track Files** node and select **Track File**.
2. In the *Open* dialog, select the `nonAdiabaticPPDF.trk` file and click **Open**.
   Two nodes appear in the simulation tree:
   - **Track Files > nonAdiabaticPPDF**.
   - **Particle Tracks**.



3. Select the **Particle Tracks > Fuel tracks** node.



Note: The **Fuel tracks** node was populated from the `nonAdiabaticPPDF.trk` file.

The particle tracks are displayed in the *Scalar Scene 1* display in the *Graphics* window.



Display the particle residence time.

4. Select the **Scalar Scene 1** node and click the **scene/plot** tab.

Activate the feature lines first and add the particle tracks.

5. Expand the **Displayers** node and set the following properties:

Node	Property	Setting
Outline 1	*Feature lines*	activated
Scalar 1 > Parts	*Parts*	**Particle Tracks > Fuel tracks**
Scalar 1 > Scalar Field	*Function*	**Track > Particle Residence Time**

The particle residence time appears as shown below.



You can see that the maximum time taken for all particle tracks to leave the domain is approximately 0.115 seconds. The default solver parameter value for *Maximum Residence Time* of 100 s is more than sufficient.

6. Save the simulation.

Summary

This tutorial demonstrated how to set up a droplet combustion case using the Non-Adiabatic PPDF Equilibrium reaction model for non-premixed flames.

The steps covered were:

- Reloading and modifying an existing simulation.
- Defining a non-adiabatic PPDF combustion model.
- Generating a PPDF table.
- Selecting models for a Lagrangian phase.
- Setting up an injector for Lagrangian Multiphase modeling.
- Defining boundary conditions.
- Creating scalar displays for examining the results.

- Running the solver for a set number of iterations.
- Analyzing the results using STAR-CCM+'s visualization facilities.
- Displaying particle tracks.

Coherent Flame Model: Premixed Propane

This tutorial simulates a lean premixed propane combustion in a V-shaped combustor using the Coherent Flame Model (CFM).

The combustor geometry with the imposed boundary conditions is shown below.



A two-dimensional geometry is used to represent the combustor. The physical properties of the fuel components (3.69% C3H8, 22.44% O2, and 73.87% N2 by mass) and the rest of the reaction components (CO2 and H2O) are defined as follows:

	O2	N2	C3H8	CO2	H20
Dynamic viscosity	Constant				
Molecular weight	Determined from elemental composition				
Specific heat	Determined via thermodynamic polynomal functions				
Thermal conductivity	Determined via the Lewis Number				

The fuel mixture, at a temperature of 600 K and pressure of 1 atmosphere, enters the combustor through the inlet with velocities of 29 m/s and 0.9 m/s, in the X and Y directions, respectively.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Visualizing the Boundaries

Setting up the Models
Setting Material Properties
Setting the Coherent Flame Model Properties
Setting Initial Conditions
Setting the Ignitor
Setting Unburnt Gas Composition
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Reporting with Monitors and Plots
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Coherent Flame Model: Premixed Propane tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `vCombustor.ccm` then click **Open**.

STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.

6. Save the simulation as `cfmCombustion.sim`.

Visualizing the Boundaries

Boundary definitions were also imported along with the volume mesh import.

To view where the boundary definitions are placed:

1. Expand the **Regions** > **Fluid** > **Boundaries** node.
2. Select the **inlet** node.

   The boundary is highlighted in the **Graphics** window.
3. Select the other nodes in turn to see their location.

Note: Alternatively, you can click various parts in the **Graphics** window to see which boundary they relate to.

4. To clear the selected boundary, click anywhere outside the part in the **Graphics** window.

Setting up the Models

Models define the primary variables of the simulation. In this example, the flow involves a turbulent, compressible, multi-component gas whose components are reacting chemically.

The Segregated Flow model is used together with the realizable K-Epsilon turbulence model and the coherent flame model.

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Note: Make sure that **Auto-select recommended models** is activated.

Group Box	Model
Enabled Models	**Two Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Time	**Steady**
Material	**Multi-Component Gas**

Group Box	Model
Reaction Regime	**Reacting**
Reactive Flow	**Premixed Combustion**
Premixed Combustion models	**Coherent Flame Model**
Enabled Models	**Ideal Gas (with combustion)** (Selected automatically) **Turbulent** (Selected automatically) **CFM reaction model** (Selected automatically) **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Flow	**Segregated Flow**
Enabled Models	**Segregated Fluid Enthalpy** (Selected automatically)

2. Deactivate **Auto-select recommended models**.
3. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

4. Click **Close**.
5. To review the models, open the **Physics 1** > **Models** node.



6. Save the simulation.

Setting Material Properties

Specify the components of the multi-component gas mixture and define the thermal conductivity.

Use the Lewis number to specify mixture thermal conductivity.

1. Expand the **Multi-Component Gas > Material Properties** node and set the following properties:

Node	Property	Setting
Thermal Conductivity	*Method*	**Lewis Number**
Lewis Number	*Lewis Number*	`1.0`

Define the five components of the mixture: C3H8, O2, CO2, H2O, and N2. All these species are included in the STAR-CCM+ materials database so detailed specification of their properties is not required.

2. Right-click the **Multi-Component Gas > Gas Components** node and choose **Select Mixture Components**.
3. In the *Select Mixture Components* dialog, expand the **Material Databases**, **Standard** and **Gases** nodes to reveal the available gaseous components.
4. Select the required components: **C3H8 (Propane)**, **CO2 (Carbon Dioxide)**, **H2O (Water)**, **N2 (Nitrogen)**, and **O2 (Oxygen)**.
5. When complete, click **Apply** then **Close**.

   Nodes that correspond to each of these components appear within the **Gas Components** node.



Note: The order of mixture components may be different to the one shown above. However, reordering of the mixture components is not necessary because species are determined from algebraic relationships rather than through their transport equations.

6. Open the **C3H8 > Material Properties > Specific Heat** node.

   Make sure that **Method** is set to **Thermodynamic Polynomial Data**.
7. Make sure that the remaining four gas components have **Method** set to **Thermodynamic Polynomial Data**.

Setting the Coherent Flame Model Properties

Specify the model parameters for the CFM.

1. Edit the **Models > Coherent Flame Model** node and set the following properties:

Node	Property	Setting
CFM Constant, Alpha	*Value*	`2.1`
Gulder Laminar Flame Speed	**Fuel Name**	**Propane**

2. Click **Close**.
3. Save the simulation.

Setting Initial Conditions

Set the initial temperature, turbulence quantities, and velocity values for the computational domain.

1. Edit the **Initial Conditions** node and set the following properties:

Node	Property	Setting
Fuel Mass Fraction	*Value*	`0.0369`
Static Temperature	*Value*	`600 K`
Turbulence Specification	*Method*	**K + Epsilon**
Turbulent Dissipation Rate	*Value*	`20` m^2/s^3
Turbulent Kinetic Energy	*Value*	`1` J/kg
Velocity	*Value*	`[1, 0, 0]` m/s

2. Click **Close**.

Setting the Ignitor

Specify a Flame Area Density (FAD) ignitor to initiate combustion.

When turned on, the FAD ignitor sets the value of FAD of the cells in the volume to the value specified (this value must non-zero) to start the

combustion process. In a spark-ignition engine, for example, this process emulates a spark to ignite the fuel and air mixture.

To define a FAD ignitor, first define a volume shape:

1. Right-click the **Tools** > **Volume Shape** node and select **New Shape** > **Block**.
2. In the *Create Block Volume Shape* dialog, set the following coordinates:

Corner 1	[`0.1, 0.0108, -0.005`] m
Corner 2	[`0.13, 0.029, 0.005`] m

3. Click **Create** then **Close**.

   A **Volume Shapes** > **Block 1** node appears in the simulation tree.
4. Rename the **Block 1** node to `IgnitorBrick`.

Create the ignitor.

5. Right-click the **Continua** > **Physics 1** > **Ignitors** node and select **New**.

   An **Ignitors** > **FlameAreaDensityIgnitor 1** node appears in the object tree.

Set the cells in the IgnitorBrick volume to have a flame density (ignitor variable) value of 10 from iteration 100 to 101.

6. Select the **FlameAreaDensityIgnitor 1** node and set the following properties:

Property	Setting
Activator	**Pulse**
Ignitor Variable Value	`10`
Shapes	**IgnitorBrick**

7. Select the **Ignitors** > **FlameAreaDensityIgnitor 1** > **Pulse** node and set the following properties:

Property	Setting
End	`101`
Begin	`100`

Setting Unburnt Gas Composition

Specify the temperature and composition of the unburnt gas mixture.

Set the unburnt gas temperature to 600 K and composition to 3.69% C3H8, 22.44% O2, and 73.87% N2, by mass.

1. Select the **Continua > Physics 1 > Fluid Stream Manager > Unburnt Gas** node and set the *Temperature of the stream* to `600` K.
2. Expand the **Continua > Physics 1 > Fluid Stream Manager** node.
3. Right-click the **Unburnt Gas > Components** node and select **Add Fluid Stream Component > C3H8**.
4. Repeat the previous step to add O2 and N2 components.
5. Expand the **Components** node and set the following properties:

Node	Property	Setting
C3H8	*Mass Fraction*	`0.0369`
O2	*Mass Fraction*	`0.2244`
N2	*Mass Fraction*	`0.7387`

6. Save the simulation.

Setting Boundary Conditions

All wall boundaries are adiabatic, no-slip walls with zero-flux for the scalars. Since these conditions are the defaults, no changes are required here. However, you need to specify the inlet and outlet boundary conditions.

Specify the inlet conditions:

1. Edit the **inlet** node and set the following properties.

Node	Property	Setting
Physics Conditions		
Turbulence Specification	*Method*	**K + Epsilon**
Velocity Specification	*Method*	**Components**
Physics Values		
Fuel Mass Fraction	*Value*	`0.0369`
Static Temperature	*Value*	`600` K

Node	Property	Setting
Turbulent Dissipation Rate	*Value*	`203` m^2/s^3
Turbulent Kinetic Energy	*Value*	`4.8` J/kg
Velocity	*Value*	`[29, 0.9, 0]` m/s

2. Click **Close**.

Specify the outlet conditions:

3. Expand the **outlet** > **Physics Values** node.
4. Select the **Static Temperature** > **Constant** node and set *Value* to `300` K.

Setting Solver Parameters and Stopping Criteria

The default under-relaxation factors for the flow and turbulence equations are suitable. However, to aid convergence, reduce the under-relaxation factors for the energy equation and CFM Combustion model.

1. Reduce the under-relaxation factors for the energy equation and **CFM reaction** model to the following values:

Node	Property	Setting
Solvers		
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.7`
CFM Combustion	*Under-Relaxation Factor*	`0.7`

Set the simulation to run for 1500 iterations, which is sufficient to achieve a steady-state solution.

2. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `1500`.

   The solution does not run for more than 1500 iterations, unless this stopping criterion is changed or disabled.
3. Save the simulation.

Visualizing the Solution

To visualize the solution, set up scalar scenes for temperature, propane mass fraction, velocity magnitude, and flame area density.

To create the temperature contour scene:

1. Create a scalar scene.
2. Rename the **Scenes** > **Scalar Scene 1** node to `Temperature`.
3. Click the **scene/plot** tab.
4. Expand the **Temperature** > **Displayers** > **Scalar 1** node.
5. Set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Scalar Field	*Function*	**Temperature**
Color Bar	*Width*	`0.7`
	Height	`0.15`
	Position	`[0.2, 0.05]`

6. Use the mouse buttons to zoom in and position the part as shown below:



Create a scalar scene for the propane mass fraction.

7. Click the *simulation* tab.
8. Right-click the **Scenes** > **Temperature** node and select **Copy**, then right-click the **Scenes** node and select **Paste**.
9. A **Copy of Temperature** node appears in the object tree.
10. Rename the **Scenes** > **Copy of Temperature** node to `Propane Mass Fraction`.
11. Expand the **Scenes** > **Propane Mass Fraction** > **Displayers** node.
12. Select the **Scalar 1** > **Scalar Field** node and set *Function* to **Mass Fraction** > **C3H8**.

Create scalar scenes for velocity magnitude and flame area density .

13. Using the copy-and-paste technique outlined above, create two more scalar scenes.
14. Rename the newly created scalar scenes `Velocity Magnitude` and `Flame Area Density`.
15. Set the following properties:

Node	Property	Setting
Scenes		
Velocity Magnitude	*Function*	**Velocity > Magnitude**
Flame Area Density	*Function*	**Flame Area Density**

16. Right-click the **Scenes** node and select **Open All Scenes**.



Reporting with Monitors and Plots

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops.

This feature requires setting up a report defining the quantity of interest and the parts of the region to be monitored. A monitor is then defined based on that report. The former also helps to create an appropriate X-Y graph plot.

Set up a report to monitor the bulk outlet temperature.

1. Right-click the **Reports** node and select **New Report > Mass Flow Averaged**.
2. Rename the **Reports > Mass Flow Averaged 1** node to `Bulk Outlet Temperature`.
3. Select the **Bulk Outlet Temperature** node and set the following properties:

Property	Setting
Scalar Field Function	**Temperature**
Parts	**Fluid: outlet**

Create a monitor and a plot for the report.

4. Right-click the **Bulk Outlet Temperature** node and select **Create Monitor and Plot from Report**.
5. Right-click the **Plots > Bulk Outlet Temperature Monitor Plot** node and select **Open**.

   The analysis is now ready to be run.
6. Save the simulation.

Running the Simulation

Run the simulation.

1. To run the simulation, click (**Run**).

   The Residuals display is created automatically and shows the progress of the solver. If necessary, click the **Residuals** tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.

The tabs at the top of the **Graphics** window make it possible to select any of the active displays for viewing. During the run, it is possible to stop the solution by clicking the (**Stop**) button on the toolbar. If you do halt the simulation, it can be continued later by clicking the (**Run**) button. If left alone, the simulation continues until 1500 iterations are complete.

2. After the run has finished, click the **Bulk Outlet Temperature Monitor Plot** tab to make sure that the solution has indeed converged.



3. Save the simulation.

Visualizing the Results

View the results of the converged solution.

1. In the *Graphics* window, click the *Temperature* tab .



2. In the *Graphics* window, click the *Propane Mass Fraction* tab .



3. In the *Graphics* window, click the *Velocity Magnitude* tab.



4. In the *Graphics* window, click the *Flame Area Density* tab.



Summary

This tutorial demonstrated how to simulate a lean premixed propane combustion in a V-shaped combustor using the Coherent Flame Model (CFM).

The steps covered were:

- Importing the mesh and saving the simulation.
- Visualizing the geometry and boundaries.
- Defining models for premixed CFM combustion.
- Defining material properties that are required for multi-component gases.
- Setting initial conditions.
- Defining boundary conditions.
- Checking solver parameters and setting stopping criteria.
- Creating scenes to visualize the results.
- Setting up monitoring report and plot.
- Running the solver for a set number of iterations.

Adiabatic PCFM: Partially-Premixed Methane

This tutorial demonstrates how to use the Partially-Premixed Coherent Flame Model (PCFM).

To show the workflow, a methane combustion case is set up using the adiabatic PCFM with PPDF laminar flamelet reaction model for non-premixed flames. For the latter, STAR-CCM+ does not directly utilize the physical properties of the fuel (methane in this case) as no additional transport equations requiring these properties are solved. Temperature, density and species mass fractions are evaluated using the β function formulation (see the Adiabatic PPDF Model section in the User Guide).

The Sandia piloted methane-air jet flame geometry that was developed by Sydney University is used here to show the required steps.



The two-dimensional axisymmetric geometry consists of the main jet with nozzle diameter of 7.2 mm, burning a premixture of 25% methane and 75% dry air by volume. This is surrounded by the pilot inlet with

diameter of 18.2 mm, burning a premixture of C2H2, H2, air, CO2, and N2.

The PPDF laminar flamelet reaction model uses a reduced 16-species methane combustion mechanism that has been prepared and distributed with the STAR-CCM+ installation.

You set up this tutorial in two stages:

1. A cold flow (non-reacting) stage, where you set up most of the geometry, boundary conditions and visualization scenes.
2. A combustion (reacting) stage, where you specify the combustion-specific model selection, boundary condition modification and additional visualization scenes and plots.

To set up other combustion models, you can use the resulting simulation file from stage 1 and save under a different file name.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting Up the Models
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Solver Parameters and Stopping Criteria
Visualizing the Solution
Running the Base Simulation
Visualize the Results
Setting Up the Combustion Simulation
Summary
Adiabatic PCFM Tutorial Bibliography

Prerequisites

The instructions in the Adiabatic PCFM: Partially-Premixed Methane tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+

Technique	Tutorial
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `sandia.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `sandiaFlameBase.sim`.

Setting Up the Models

Models define the primary variables of the simulation, including pressure, temperature, velocity and other mathematical formulations that are required to generate the solution. For the first stage of this tutorial, you set up a steady, turbulent, compressible, gaseous, non-reacting flow using the segregated flow model.

1. For the physics continuum, **Continua** > **Physics 1**:
   a) From the *Enabled Models* box, deselect **Two Dimensional**.
   b) Select the following models in order:

Note: Make sure that **Auto-select recommended models** is activated.

Group Box	Model
Space	**Axisymmetric**
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**

Group Box	Model
Equation of State	**Ideal Gas**
Viscous Regime	**Turbulent**
	Reynolds-Averaged Navier Stokes (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Energy	**Segregated Fluid Temperature**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Setting Initial Conditions

Specify appropriate values for the turbulence and velocity initial conditions.

1. Edit the **Continua** > **Physics 1** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Turbulence Specification	*Method*	**K + Epsilon**
Turbulent Dissipation Rate	*Value*	50000.0 m^2/s^3
Turbulent Kinetic Energy	*Value*	30 J/kg
Velocity	*Value*	[0.9, 0.0, 0.0] m/s

2. Click **Close**.

Setting Boundary Conditions and Values

In this section you set the boundary conditions for the inlets, outer and outlet boundaries. For the outer and outlet boundaries, you set the backflow conditions to the extrapolation method.

Contents

inlet-1 boundary
inlet-2 boundary
inlet-3 boundary
outer boundary
outlet boundary

inlet-1 boundary

Specify the physics values for the **inlet-1** boundary.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-1** node and set the following properties:

Node	Property	Setting
Static Temperature	*Value*	300.0 K
Turbulence Intensity	*Value*	0.1
Turbulent Viscosity Ratio	*Value*	100
Velocity Magnitude	*Value*	49.6 m/s

3. Click **Close**.

inlet-2 boundary

Specify the physics values for the **inlet-2** boundary.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-2** node and set the following properties:

Node	Property	Value
Static Temperature	*Value*	`300` K
Turbulence Intensity	*Value*	`0.1`
Turbulence Viscosity Ratio	*Value*	`100`
Velocity Magnitude	*Value*	`11.4` m/s

3. Click **Close**.

inlet-3 boundary

Specify the physics values for the **inlet-3** boundary.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-3** node and set the following properties:

Node	Property	Setting
Static Temperature	*Value*	`300` K
Turbulence Intensity	*Value*	`0.1`
Turbulent Viscosity Ratio	*Value*	`100`
Velocity Magnitude	*Value*	`0.9` m/s

3. Click **Close**.

outer boundary

Specify the physics conditions and values for the **outer** boundary.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **outer** node and set the following properties:

Node	Property	Setting
Backflow Direction Specification	*Method*	**Extrapolated**
Static Temperature	*Value*	`300` K (set by default)
Turbulence Intensity	*Value*	`0.01` (set by default)
Turbulent Viscosity Ratio	*Value*	`10` (set by default)

3. Click **Close**.

outlet boundary

Specify the physics conditions and values for the **outlet** boundary.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **outlet** node and set the following properties:

Node	Property	Setting
Backflow Direction Specification	*Method*	**Extrapolated**
Static Temperature	*Value*	`300` K (set by default)
Turbulence Intensity	*Value*	`0.01` (set by default)
Turbulent Viscosity Ratio	*Value*	`10` (set by default)

3. Click **Close**.
   Specification of the boundary conditions is now complete.
4. Save the simulation.

Setting Solver Parameters and Stopping Criteria

The default under-relaxation factors for the flow, energy, and turbulence equations are suitable.

By default, the maximum number of steps is set to 1000. This number is sufficient to run the cold flow stage of this tutorial.

Visualizing the Solution

This section of the tutorial covers post-processing.

1. Create a scalar scene.
2. Rename the **Scalar Scene 1** node to `Velocity Magnitude`.
3. Click the **scene/plot** tab.
4. Expand the **Velocity Magnitude** > **Displayers** > **Scalar 1** node.
5. Set the following properties:

Node	Property	Setting
Scalar Field	*Function*	**Velocity > Magnitude**
Color Bar	*Width*	`0.7`
	Height	`0.15`
	Position	`[0.2, 0.05]`

6. Zoom in and position the part to obtain a similar view as shown below:



7. Save the simulation.

Running the Base Simulation

Run the simulation.

1. In the toolbar, click (**Run**).

   The Residuals display is created automatically and shows the progress of the solver. If necessary, click the **Residuals** tab to bring the Residuals plot into view. An example of a residual plot is shown in a separate part of the User Guide. This example looks different from your residuals, since the plot depends on the models selected.

   If left alone, the simulation continues until 1000 iterations are complete.

   When the last iteration has been completed, the message `Stopping Criterion Maximum Steps satisfied` is displayed in the *Output* window.
2. Save the simulation.

Visualize the Results

View the velocity magnitude contour for the converged solution.

1. To view the velocity magnitude contour for the converged solution, select the *Velocity Magnitude* display.



This completes the first stage of the simulation. In the next section you set up the combustion model, modify some of the boundary condition values and create additional visualization scenes and a monitoring plot.

Setting Up the Combustion Simulation

In this section, you set up the PCFM and adiabatic PPDF laminar flamelet reaction model for non-premixed flames. For the latter, you use a pre-prepared 16-species methane combustion mechanism.

If you have exited the previous simulation file:

1. Start up STAR-CCM+ in a manner that is appropriate to your working environment.
2. In the toolbar, click (**Load Simulation**).
3. In the *Load Simulation* dialog:
   a) Click **Browse**.
   b) In the *Open* dialog, navigate to your local `sandiaFlame` directory, select `sandiaFlameBase.sim` and click **Open**.
   c) In the *Load Simulation* dialog, click **OK**.

4. Save a copy of the simulation file under a different file name, `adiabaticPCFM.sim`, for the combustion calculation to follow.

Contents

Setting Up the Models
Setting the Coherent Flame Model Properties
Constructing the PPDF Flamelet Table
Checking Mixture Components
Checking Fluid Stream Composition
Setting Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Reporting with Monitors and Plots
Running the Simulation
Visualize the Results

Setting Up the Models

Models define the physical variables and phenomena in the simulation. In this simulation, the flow is a reacting, partially premixed, multi-component gas.

1. For the physics continuum, **Continua** > **Physics 1**:
   a) From the *Enabled Models* box, deselect all models, except **Steady**, **Gradients**, and **Axisymmetric**.
   b) Select the following models in order:

Group Box	Model
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Partially-Premixed Combustion**
Partially-Premixed Combustion models	**Adiabatic Partially Premixed CFM**
	Adiabatic PPDF (Selected automatically)
	Coherent Flame Model (Selected automatically)
	CFM partially premixed reaction model (Selected automatically)

Group Box	Model
	PPDF Flamelet
	PCFM Ideal Gas (Selected automatically)
	Segregated Flow (Selected automatically)
	Turbulent (Selected automatically)
	Reynolds-Averaged Navier-Stokes (Selected automatically)

2. Deactivate **Auto-select recommended models**.
3. Select the following models in order:

Group Box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
K-Epsilon Turbulence Models	**Standard K-Epsilon**
K-Epsilon High y+ Wall Treatment	**High y+ Wall Treatment**

4. Click **Close**.
5. To review the models, open the **Physics 1** > **Models** node.



Setting the Coherent Flame Model Properties

Specify the model parameters for the CFM.

1. Edit the **Models** > **Coherent Flame Model** node and set the following properties:

Node	Property	Setting
CFM Constant, Alpha	*Value*	`2.1`
Gulder Laminar Flame Speed	*Fuel Name*	**Methane**

2. Click **Close**.

Constructing the PPDF Flamelet Table

PPDF tables contain information about chemical species and reactions, that the PPDF Flamelet model needs.

Before you can construct the table, the directory containing the flamelet library files must be copied into your working directory.

STAR-CCM+ is not currently capable of producing internally-generated PPDF Flamelet tables. However, it does contain facilities for importing pre-generated PPDF tables or constructing them from flamelet library data that are produced by the DARS package. In this tutorial, the latter method is used.

1. Copy the flamelet library files into your working directory:
   a) Navigate to the `doc/startutorialsdata/combustion/data` directory of your STAR-CCM+ installation.
   b) Copy the directory named `ppdfMethane16` into the directory that contains your simulation file.
2. In STAR-CCM+, expand the **Continua** > **Physics 1** > **Models** node.
3. Right-click the **PPDF Flamelet** > **PPDF Flamelet Table** node and select **Construct**.
   The parameters control the accuracy and size of the look-up table. Their default values are suitable for this tutorial so no changes are necessary.
4. In the *Open* dialog, select the `ppdfMethane16` directory and click **Open**.
   The PPDF look-up table is constructed from the flamelet data files in directory `ppdfMethane16`. Once complete, the following message appears in the **Output** window:
   `PpdfFlameletTable::generate() end`

You can display an XY-plot of variables within the PPDF look-up table. For example:

5. Right-click the **PPDF Flamelet Table** > **PPDF XY Plots** node and select **New Plot** > **DRHODP-SPECIESOH** > **Density**.
   The **PPDF XY Plots** > **Density-1** node is created. A corresponding plot is also created in the *Graphics* window.
6. Right-click the **Density-1** node and select **Add New Table View for Plotting**.
7. In the *Select points for plotting* dialog, accept the default settings and click **OK**.



In the **Graphics** window, the *Density-1* plot looks as shown below.



Checking Mixture Components

Once you generate the PPDF table successfully, the components that are stored in the flamelet libraries are listed under the **Multi-Component Gas** node. There are 16 components listed.

1. To check the gas components, open the **Multi-Component Gas > Gas Components** node.



No further specification of species or reactions is required since all necessary information is contained within the table.

Checking Fluid Stream Composition

Once you have successfully generated the PPDF table, the fluid stream node is populated with correct components and mass fractions.

To check the fuel and oxidizer stream compositions:

1. Select the **Continua** > **Physics 1** > **Fluid Stream Manager** > **Fuel Stream** node.
2. Make sure that:
   a) *Temperature of the stream* is set to `300.0 K`.
   b) Each of the **Component** nodes has the following mass fractions:

Component	MassFraction
CH4	0.1563
N2	0.6472
O2	0.1965

3. Similarly, make sure that the **Oxidizer Stream** temperature is set to `300 K`, and has the following mass fractions (to three significant figures):

Component	Mass Fraction
N2	0.767
O2	0.233

4. Save the simulation.

Setting Boundary Conditions

You have already set up the flow boundary conditions. Here you only need to specify the inlet boundary conditions for the combustion-related variables.

1. Expand the **Regions** > **Fluid** > **Boundaries** node.
2. Edit the **inlet-1**, **inlet-2**, and **inlet-3** nodes, and set the following properties:

Node	Property	Setting
inlet-1		
Flame Area Density	*Value*	`0.0`
Fuel Mass Fraction	*Value*	`0.1563`
Mixture Fraction Profile	*Value*	`1.0`
inlet-2		
Flame Area Density	*Value*	`100.0`
Fuel Mass Fraction	*Value*	`0.0`
Mixture Fraction Profile	*Value*	`0.27`
inlet-3		
Flame Area Density	*Value*	`0.0`
Fuel Mass Fraction	*Value*	`0.0`

Node	Property	Setting
Mixture Fraction Profile	*Value*	`0.0`

Setting Solver Parameters and Stopping Criteria

The default under-relaxation factors for the flow, turbulence equations are suitable. However, to aid convergence, reduce the under-relaxation factors for combustion models.

1. Expand the **Solvers** node.
2. Set the following properties:

Node	Property	Setting
Segregated Flow > Velocity	*Under-Relaxation Factor*	`0.7`
Segregated Flow > Pressure	*Under-Relaxation Factor*	`0.3`
K-Epsilon Turbulence	*Under-Relaxation Factor*	`0.8`
CFM Combustion	*Under-Relaxation Factor*	`0.7`
PPDF Combustion	*Under-Relaxation Factor*	`0.7`

Run the simulation for 1600 iterations, which is sufficient to achieve a steady-state solution.

3. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `1600`.

Reporting with Monitors and Plots

STAR-CCM+ can dynamically monitor virtually any quantity while the solution develops.

This requires setting up a report which defines the quantity of interest and the parts of the region to be monitored. A monitor is then defined based on that report. The former also helps to create an appropriate X-Y graph plot.

Set up a report to monitor the mass flow averaged temperature at the outlet.

1. Create a **Mass Flow Averaged** report and rename it `Mass Flow Averaged Temperature`.
2. Select the **Mass Flow Averaged Temperature** node and set the following properties:

Property	Setting
Scalar Field Function	**Temperature**
Parts	**Fluid > outlet**

Create a monitor and a plot for the report.

3. Right-click the **Mass Flow Averaged Temperature** node and select **Create Monitor and Plot from Report**.
4. Right-click the **Plots > Mass Flow Averaged Temperature Monitor Plot** node and select **Open**.

   The analysis is now ready to be run.
5. Save the simulation.

Running the Simulation

Clear the previous solution and re-run the simulation.

Clear the previous solution:

1. Select **Solution > Clear Solution**.
2. In the *Clear Solution* dialog, click **OK**.

Re-run the simulation:

3. To run the simulation, click (**Run**).

The iteration progress can be monitored in the *Output* window. Monitor the convergence by viewing the Residuals using the visualization tabs in the *Graphics* window.

4. After the run has finished, click the **Mass Flow Averaged Temperature Monitor Plot** tab to make sure that the solution has indeed converged.



5. Save the simulation.

Visualize the Results

This section of the tutorial covers post-processing.

Activate the *Velocity Magnitude* scene in the *Graphics* window.

1. Select the **Velocity Magnitude > Displayers > Scalar 1** node and set *Contour Style* to **Smooth Filled**.



The velocity magnitude contour for the converged solution is now as shown below.



2. Create another scalar scene for temperature by copying and pasting the velocity magnitude scene that was created previously.

a) To switch back to the main object tree, click the *simulation* window .

b) Right-click the **Scenes** > **Velocity Magnitude** node and select **Copy**, then right-click the **Scenes** node and select **Paste**.

A node called **Scenes** > **Copy of Velocity Magnitude** is added to the object tree.

3. Rename the **Copy of Velocity Magnitude** node to `Temperature`.
4. Click the **scene/plot** tab.
5. Select the **Temperature** > **Displayers** > **Scalar 1** > **Scalar Field** node and set the following properties:

Property	Setting
Function	**Temperature**
Min	`300`
Max	`2200`

The temperature contour for the converged solution appears as shown below:



6. Create two scalar scenes for the CH4 and CO2 mass fractions.

7. Rename the scalar scene nodes `CH4 Mass Fraction` and `CO2 Mass Fraction`.
8. Set the following properties:

Node	Property	Setting
CH4 Mass Fraction	*Function*	**Mass Fraction > CH4**
CO2 Mass Fraction	*Function*	**Mass Fraction > CO2**
	Min	`0.0`
	Max	`0.13`

These settings are based on copying and pasting the velocity magnitude node for both mass fractions.

The CH4 mass fraction contour for the converged solution is now as shown below:



The CO2 mass fraction contour for the converged solution is now as shown below:



Summary

This tutorial demonstrated how to use the Partially-Premixed Coherent Flame Model (PCFM).

The steps covered were:

- Defining models for partially-premixed CFM combustion and PPDF laminar flamelet for non-premixed flames.
- Displaying variables within the PPDF look-up table.
- Setting initial conditions.
- Defining boundary conditions.
- Creating scenes to visualize the results.

Adiabatic PCFM Tutorial Bibliography

Masri, A. R., Dibble, R. W., and Barlow, R. S., Prog. Energy Combust. Sci. 22:307-362 (1996).

Progress Variable Model: Sandia Flame D with Participating Media Radiation

This tutorial demonstrates how to model participating media radiation with a gray spectrum model.

The Sandia piloted methane-air jet flame geometry was developed by Sydney University, and is used here to show the required steps.



The two-dimensional axisymmetric geometry consists of the main jet with nozzle diameter of 7.2 mm, burning a premixture of 25% methane and 75% dry air by volume. The pilot inlet surrounds the nozzle. The inlet has a diameter of 18.2 mm, and burns a premixture of C2H2, H2, air, CO2, and N2.

The Progress Variable Model (PVM) is used to reduce the computational cost of this tutorial. In the PVM combustion model, the effects of the detailed chemistry of combustion are stored in a precomputed table. This PVM table is supplied with this tutorial, and is used for density calculations.

In the second part of this tutorial, the effect of participating media radiation is taken into account, using the Discrete Ordinate Method (DOM). The Weighted Sum of Gray Gases (WSGG) method is used to calculate the absorption coefficient of the gas mixture.

You set up this case in two stages:

1. The case is run without radiation, and the peak temperature in the domain is noted.
2. The case is run with radiation. The peak temperature of the domain is compared with the first case. The literature predicts a drop of 15 K to 20 K with radiation modeling.

Contents

Prerequisites
Importing the Mesh and Naming the Simulation
Setting Up the Models
Importing the PVM Table
Setting Initial Conditions
Setting Boundary Conditions and Values
Setting Solver Parameters
Monitoring the Peak Temperature
Creating a Temperature Scene
Running the Simulation Without Radiation
Visualizing the Results
Exporting Axial Temperature for Comparison with Radiation Case
Setting Up the Radiation Model
Setting the Surface Emissivity at the Boundaries
Monitoring the Radiative Heat Source
Creating Radiative Energy Source and Absorption Coefficient Scenes
Running the Simulation with Radiation
Analyzing the Results
Summary
PVM Sandia Flame Tutorial Bibliography

Prerequisites

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Also, see the User Guide for more details about importing a volume mesh.

Importing the Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/combustion/data`.
5. Select `sandia.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `pvmSandiaNoRadiation.sim`.

Setting Up the Models

In this example, the flow is turbulent (Re = 22400), compressible, and composed of a multi-component gas whose components participate in chemical reactions.

Since the flow speeds that are involved are small compared with the speed of sound, the Segregated Flow model is appropriate for this tutorial. The Progress Variable Model is appropriate for chemical reactions.

To select the models:

1. For the physics continuum, **Continua** > **Physics 1**,
   a) In the *Enabled Models* group box, deselect **Two Dimensional**.
   b) Select the following models in order:

Note: **Auto-select recommended models** must be activated.

Group Box	Model
Space	**Axisymmetric**
Material	**Multi-Component Gas**
Reaction Regime	**Reacting**
Reactive Flow	**Non-Premixed Combustion**
Non-Premixed Combustion models	**Progress Variable Model**
	PVM Reaction (Selected automatically)
	PVM Ideal Gas (Selected automatically)
Flow	**Segregated Flow**
	Gradients (Selected automatically) **Turbulent** (Selected automatically) **Segregated Fluid Enthalpy** (Selected automatically) **Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
	Realizable K-Epsilon Two-Layer (Selected automatically)
	Two-Layer All y+ Wall Treatment (Selected automatically)
Time	**Steady**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Importing the PVM Table

The PVM table contains pre-computed values of molecular weights. The PVM Ideal Gas model uses these values to compute densities.

Import the PVM table:

1. Expand the **Continua** > **Physics 1** > **Models** node.
2. Right-click the **PVM Reaction** > **PVM Table** node and select **Import...**.
3. In the *Open* dialog, navigate to the `doc/startutorialsdata/combustion/data` subdirectory of your STAR-CCM+ installation directory, select `PVMSandia_16.tbl` (the PVM table file), and click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window.
4. Expand the **Continua** > **Physics 1** > **Models** > **Multi-Component Gas** > **Gas Components** node.



The gas components have been selected automatically when you imported the PVM table.

5. Save the simulation.

Setting Initial Conditions

Set the initial turbulence dissipation rate and the initial turbulent kinetic energy, then initialize the velocity field with the lowest inlet value.

1. Expand the **Continua > Physics 1 > Initial Conditions** node.
2. Set the following properties:

Node	Property	Setting
Turbulence Specification	*Method*	**K + Epsilon**
Turbulent Dissipation Rate > Constant	*Value*	50000 m^2/s^3
Turbulent Kinetic Energy > Constant	*Value*	30 J/kg
Velocity > Constant	*Value*	[0.9, 0, 0] m/s

Setting Boundary Conditions and Values

In this section, you set the boundary conditions for the inlets, outer, and outlet boundaries.

For the outer boundary, you set the backflow conditions to the extrapolation method because the flow is parallel to this pressure outlet.

Contents

Inlet-1 Boundary
Inlet-2 Boundary
Inlet-3 Boundary
Outer Boundary
Outlet Boundary

Inlet-1 Boundary

Set the Inlet-1 boundary physics values.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-1** node and set the following properties:

Table 5: Physics Values

Node	Property	Setting
Mixture Fraction Profile > Constant	*Value*	[1.0]
Turbulence Intensity > Constant	*Value*	0.1
Turbulent Viscosity Ratio > Constant	*Value*	100
Velocity Magnitude > Constant	*Value*	49.6 m/s

3. Click **Close**.

Inlet-2 Boundary

Set the Inlet-2 boundary Physics values.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-2** node and set the following properties:

Node	Property	Setting
Mixture Fraction > Constant	*Value*	0.27
Progress Variable > Constant	*Value*	1.0
Static Temperature > Constant	*Value*	1880 K
Turbulence Intensity > Constant	*Value*	0.1
Turbulence Viscosity Ratio > Constant	*Value*	100
Velocity Magnitude > Constant	*Value*	11.4 m/s

3. Click **Close**.

Inlet-3 Boundary

Set the Inlet-3 boundary physics values.

1. Expand the **Regions** > **fluid** > **Boundaries** node.
2. Edit the **inlet-3** node and set the following properties:

Node	Property	Setting
Turbulence Intensity > Constant	*Value*	0.1
Turbulent Viscosity Ratio > Constant	*Value*	100
Velocity Magnitude > Constant	*Value*	0.9 m/s

3. Click **Close**.

Outer Boundary

Set the physics conditions for the outer boundary.

For this tutorial, the flow is parallel to the pressure boundary. Therefore, the reverse flow direction is extrapolated from the interior of the domain.

1. Expand the **Regions** > **fluid** > **Boundaries** > **outer** node.
2. Select the **Physics Conditions** > **Backflow Direction Specification** node and set *Method* to **Extrapolated**.

Outlet Boundary

Set the physics conditions and values for the outlet boundary.

The default physics conditions and values for the outlet boundary are suitable for this simulation.

1. Maintain the default settings for the **Outlet** boundary.
2. Save the simulation.

Setting Solver Parameters

Increase the under-relaxation factor for the velocity solver.

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Flow > **Velocity**	*Under-Relaxation Factor*	`0.85`
Segregated Energy	*Fluid Under-Relaxation Factor*	`0.9`
PVM Combustion	*Under-Relaxation Factor*	`0.9`

Monitoring the Peak Temperature

Create a report to monitor the maximum value of temperature in the domain.

Create a report:

1. Create a maximum report and rename it `Peak Temperature`.
2. Select the **Reports** > **Peak Temperature** node and set the following properties:

Property	Value
Scalar Field Function	**Temperature**
Parts	**fluid**

3. Click **OK**.

Create the monitor and plot:

4. Right-click the **Reports** > **Peak Temperature** node and select **Create Monitor and Plot from Report**.
5. Right-click the **Plots** > **Peak Temperature Monitor Plot** node and select **Open**.
6. Save the simulation.

Creating a Temperature Scene

Create a temperature scene to monitor the jet development as the solution progresses.

1. Create a scalar scene.
2. Rename **Scalar Scene 1** to `Temperature Scene`.
3. Open **Scenes** > **Temperature Scene** and click **scene/plot** at the top of the simulation tree.
4. Expand the **Temperature Scene** > **Displayers** > **Scalar 1** node.
5. Select the **Scalar Field** node and set *Function* to **Temperature**.

Running the Simulation Without Radiation

Run the simulation without radiation to provide a set of results that can be used to measure the effectiveness of radiation.

1. Click (**Run**).
   The iteration progress can be monitored in the *Output* window. Monitor the convergence by viewing the Residuals using the visualization tabs in the *Graphics* window. When the last iteration is complete, the message `Stopping Criterion Maximum Steps satisfied` is displayed in the *Output* window.

Visualizing the Results

View the temperature contour for the converged solution, without radiation.

View the temperature contour:

1. Select the *Temperature Scene* display.



Note the maximum temperature in the domain for this case:

2. In the model tree, right-click the **Reports** > **Peak Temperature** node, and select **Run Report**.

   The maximum temperature in the domain is displayed in the *Output* window:

```
        Maximum of Temperature on Volume Mesh

Part                                Value (K)
------------------------------      -------------
fluid                                1.983717e+03
                                    -------------
                        Total:       1.983717e+03
```

Save the converged solution for the simulation without radiation:

3. Save the simulation.

Exporting Axial Temperature for Comparison with Radiation Case

In order to compare the results from this case with the radiation run, export the temperature along the axis.

To view the axial temperature profile:

1. Create an X-Y plot.
2. Rename **XY Plot 1** to `Axial Temperature`.
3. Expand the **Plots** > **Axial Temperature** node and set the following properties:

Node	Property	Setting
Axial Temperature	*Parts*	**fluid** > **Boundaries** > **axis**
Y Types		
Y Type 1 > **Scalar Function**	*Field Function*	**Temperature**
Y Type 1 > **fluid: axis**	*Legend Name*	`Without Radiation`

Export the axial temperature profile:

4. Right-click the **Plots** > **Axial Temperature** node and select **Export...**.
5. Save the export file as `WithoutRadiation.csv` in the working directory.
6. Save the simulation.

Setting Up the Radiation Model

Activate the radiation model to determine the effect that it has on the simulation.

First clear the solution and save the case to a new file name:

1. From the menu bar, select **Solution** > **Clear Solution...**.
2. In the *Clear Solution* dialog, click **OK**.
3. Save the simulation as `pvmSandiaRadiation.sim`.

Select the radiation model:

4. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Optional Models	**Radiation**
	Participating Media Radiation (DOM) (Selected automatically)
Radiation Spectrum (Participating)	**Gray Thermal Radiation**

5. Click **Close**.
6. To review the models, open the **Physics 1** > **Models** node.



CO2 and H2O gases dominate emission and absorption among the combustion gases. These species require that the spectral variation of radiative properties be taken into account. The Weighted Sum of Grey Gases (WSGG) model is a compromise between the simplistic gray gas model, and a complete spectral property model.

Select the WSGG model to model the absorption coefficient of the gas mixture in the domain:

7. Expand the **Continua > Physics 1 > Models > Multi-Component Gas** node.
8. Select the **Material Properties > Absorption Coefficient** node and set *Method* to **Weighted Sum of Grey Gases**.



9. Save the simulation.

Setting the Surface Emissivity at the Boundaries

The Sandia flame D is a jet flame and is not confined by any real wall surfaces. Therefore, surface emissivity values are set on boundaries.

Such open boundaries for flames are modeled as black walls at ambient temperature. With this model, any radiation that falls on the wall is fully absorbed without raising the wall temperature (that you specify as static temperature). All the radiative heat escapes through these open boundaries. Also, the emission from these boundaries equals the ambient emission entering the domain. This method is standard for radiation modeling.

Set surface emissivity at each boundary.

1. Expand the **Regions** > **fluid** > **Boundaries** > **inlet-1** > **Physics Values** node.
2. Select the **Surface Emissivity** > **Constant** node and set *Value* to `1.0`.
3. Similarly, set the surface emissivity of the **inlet-2**, **inlet-3**, **outer** and **outlet** boundaries to `1.0`.

Monitoring the Radiative Heat Source

The volume radiative heat source in the domain should be equal to the total radiative heat flux over the boundaries.

Set up two reports to determine these values. Start with the volume radiative source:

1. Create a Volume Intergral report and rename it `Total Radiative Source`.
2. Select the **Reports** > **Total Radiative Source** node and set the following properties:

Property	Setting
Scalar Field Function	**Radiative Energy Source**
Parts	**Regions** > **fluid**

Create a report to calculate the total boundary flux through the boundaries:

3. Create a Surface Integral report and rename it `Total Boundary Heat Flux`.
4. Select the **Reports** > **Total Boundary Heat Flux** node and set the following properties:

Property	Setting
Scalar Field Function	**Boundary Radiation Heat Flux**
Parts	**fluid > Boundaries**

To select the boundaries without selecting the region, click the button to the left of the **Boundaries** node:



Creating Radiative Energy Source and Absorption Coefficient Scenes

Create two scenes to monitor the radiative energy source and absorption coefficient.

1. Create a scalar scene.
2. Rename **Scalar Scene 1** to `Radiative Energy Source Scene`.
3. Open **Scenes > Radiative Energy Source Scene** and click **scene/plot** at the top of the simulation tree.
4. Expand the **Radiative Energy Source Scene > Displayers > Scalar 1**.
5. Select the **Scalar Field** node and set *Function* to **Radiative Energy Source**.
6. Repeat these steps to create a scalar scene for absorption coefficient.

7. Rename the scene to `Absorption Coefficient Scene` and set *Scalar Field* to **Absorption Coefficient**.

Running the Simulation with Radiation

Run the simulation with the radiation model activated.

1. Click (**Run**).
   When the last iteration has been completed, the message `Stopping Criterion Maximum Steps satisfied` is displayed in the *Output* window.
2. Save the simulation.

Analyzing the Results

Analyze the results for the converged solution, with radiation.

1. To view the temperature contour for the converged solution, select the *Temperature Scene* display.



2. Click the *Radiative Energy Source Scene* display.
   The scene appears as shown below:



3. Click the *Absorption Coefficient Scene* display.
   The absorption coefficient scene appears as shown below:



Calculate the peak temperature in the domain with radiation:

4. Right-click the **Reports** > **Peak Temperature** node, and select **Run Report**.

   The Output window displays the peak temperature:

```
        Maximum of Temperature on Volume Mesh

Part                                Value (K)
------------------------------      -------------
fluid                                1.965602e+03
                                    -------------
                         Total:      1.965602e+03
```

The peak temperature with radiation is 18 K below the predicted maximum temperature without radiation. This temperature difference is in good agreement with the 15 – 20 K temperature drop that is predicted by the literature.

5. Run the **Total Boundary Heat Flux** and **Total Radiative Source** reports:

```
        Surface Integral of Boundary Radiation Heat Flux on Volume Mesh

Part                                Value (W)
------------------------------      -------------
fluid: axis                          0.000000e+00
fluid: inlet-1                       8.072476e-03
fluid: inlet-2                       5.800014e-02
fluid: inlet-3                       6.598733e+00
fluid: outer                         5.322614e+01
fluid: outlet                        9.833222e+00
                                    -------------
                        Total:       6.972417e+01

        Volume Integral of Radiative Energy Source on Volume Mesh

Part                                Value (W)
------------------------------      -------------
fluid                               -6.972416e+01
                                    -------------
                        Total:      -6.972416e+01
```

The volume integrated radiation energy source matches with the surface integrated boundary radiative heat flux summed over all boundaries.

Compare the axial temperature profile against the results of the case without radiation, and against the experimental data:

6. Expand the **Plots** > **Axial Temperature** > **Y Types** node.
7. Select the **Y Type 1** > **fluid:axis** node and set the following properties:

Node	Property	Setting
fluid:axis	*Legend Name*	`With Radiation`
Symbol Style	*Color*	`Blue [r=0,g=0,b=255]`



8. Right-click the **Tools > Tables** node and select **New Table > File Table**.
9. Copy `doc/startutorialsdata/combustion/data/SandiaD.csv` to the working directory.
10. Multi-select `SandiaD.csv` and `WithoutRadiation.csv`.
11. Click **Open**.

Add the tables to the *Axial Temperature* plot:

12. Right-click the **Plots > Axial Temperature > Data Series** node and select **Add Data**.
13. In the *Add Data Providers to Plot* dialog, select **Without Radiation** and click **OK**.
14. Select the **Data Series > WithoutRadiation** node and set the following properties:

Property	Setting
Table	**Without Radiation**
X Column	**fluid axis: Direction [1,0,0]**
Y Column	**fluid axis: Temperature**

Similarly, create another data set to plot the experimental data:

15. Add the data set **SandiaD**.
16. Select the **SandiaD** node and set the following properties:

Property	Setting
Table	**SandiaD**
X Column	**x/m**
Y Column	**T/K**

The *Axial Temperature* plot appears as show below:



Summary

This tutorial demonstrated how to model participating media radiation with a gray spectrum model.

The steps covered were:

- Using the Progress Variable Model (PVM).
- Importing a PVM table.

- Modeling participating media radiation.
- Modeling absorption with the Weighted Sum of Grey Gases model.
- Importing and exporting plot data.

PVM Sandia Flame Tutorial Bibliography

Masri, A. R., Dibble, R. W., and Barlow, R. S., Prog. Energy Combust. Sci. 22:307-362 (1996).

Barlow, R., and Frank, J. "Piloted CH4/Air Flames C,D,E, and F—Release 2.1" Sandia National Laboratories, Livermore, CA, 2007.

Fire and Smoke Wizard: Steckler Room

This tutorial demonstrates the Fire and Smoke Wizard. The geometry and fire settings closely resemble those Steckler reported, including radiation and wall heat absorption. This setup allows a preliminary comparison with experimental data.

Before you begin, take note of the following important Conventions in the **Fire and Smoke Wizard**:

- Gravity is in negative z-direction (z pointing up).
- Reference Density is 1.2 Kg/m^3.
- Reference Altitude is 0, 0, 0 m.

The problem geometry is shown below.



In this scene, the location of the line-probe corresponds to the experimental measurement stack.

Contents

Prerequisites
Importing the Geometry
Splitting the Boundary
Generating the Volume Mesh

Setting up Fire Properties
Reviewing Model Settings
Setting up Solver Parameters and Stopping Criteria
Visualizing the Solution
Running the Simulation
Analyzing Results
Settings for Steckler Room Fire Validation
Comparison with Measurements
Adjusting Stopping Criteria and Continuing
Analyzing Validation Results
Summary

Prerequisites

The instructions in the Fire and Smoke Wizard: Steckler Room tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Geometry

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied geometry.

Create a directory for the tutorial called `stecklerFire`. Navigate to the `doc/startutorialsdata/combustion/data` subdirectory of your STAR-CCM+ installation directory. Copy the file containing the problem geometry database (`tutorial.dbs`), along with the other files `DoorUvelExpt.csv` and `DoorTempExpt.csv`, into this directory.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh**.
4. In the *Open* dialog, navigate to the `stecklerFire` directory.

5. Select `tutorial.dbs` then click **Open**.
6. In the *Import Surface Options* dialog, in the **Import Mode** box, select **Create new Region**.
7. To accept the remaining default *Import Surface Options*, click **OK**.
8. Save the simulation as `roomfire.sim`.

The *Geometry Scene 1* display in the *Graphics* window can be examined by adjusting the viewing angle and zooming as appropriate:

9. Open the **Scenes** > **Geometry Scene 1** > **Displayers** node.
10. Select the **Geometry 1** node and activate the **Mesh** property.
11. Using the mouse buttons, zoom in and place the parts as shown below:



12. In the *Vis* toolbar, click (**Save-Restore-Select views**) and select **Store Current View**.
13. Save the simulation.

Splitting the Boundary

The model is imported as a single region, a surface mesh enclosing two volumes–the fire volume, and the volume surrounding the fire. Split this single region into two separate volumes that represent the fire and

surrounding volumes (room, door space and exterior) by selecting the surface topology option.

1. Right-click the **Regions** > **Region 1** node and select **Split by Surface Topology**.
2. In the *Split Region Options* dialog, click **OK** to accept the default settings.

In the simulation tree, the original **Region 1** node is now split into two nodes, **Region 1** and **Region 1 2**.

3. Rename the **Region 1** node to `Room`.
4. Rename the **Region 1 2** node to `Fire`.

Display the **Fire** and **Room** regions that have been created:

5. Click the **Fire** and **Room** nodes in turn.

As each region is selected, it is highlighted in *Geometry Scene 1*, as shown for the fire region in the following screenshot.



The model has a single shell entity that is defined for all room surfaces. Split this shell further in the following steps.

6. Expand the **Room** > **Boundaries** node.

7. Right-click the **Surface shells** node and select **Split By Feature Curves**.
8. In the *Split Boundaries by Feature Curves* dialog set the following properties:

Property	Setting
Boundaries	**Room: Surface Shells**
Feature Curves	**Room: Edge lines**

9. Click **OK**.

   Several boundaries are created to allow separate boundary types to be prescribed on each surface if necessary. In this case, only a pressure ambient boundary is required.

Surface shells 16 represents the top surface of the exterior volume.

10. Select the **Surface shells 16** node and set *Type* to **Pressure Outlet**.



11. Save the simulation.

Generating the Volume Mesh

Set up and generate the volume mesh.

1. For the meshing continuum, **Continua** > **Mesh 1**, select the following models in order:
   - **Surface Remesher**
   - **Polyhedral Mesher**
   - **Prism Layer Mesher**
2. Select the **Mesh 1** > **Reference Values** > **Base Size** node and set *Value* to `0.1 m`.
3. To generate the volume mesh, click **(Generate Volume Mesh)**. When the volume mesh is generated, the message `Volume Meshing Pipeline Completed` is displayed in the *Output* window. The number of Cells, Faces, and Vertices are also reported.

Setting up Fire Properties

Set up the properties of the fire using the **Fire and Smoke Wizard**.

To access the wizard:

1. In the menu bar, select **Window** > **Toolbars** > **Customize**.
2. In the *Customize Toolbars* dialog, scroll down to locate the **Fire and Smoke Wizard**.



3. Drag and drop the **Fire and Smoke Wizard** icon onto the end of one of the toolbars near the top of the workspace.
4. In the *Customize Toolbars* dialog, click **Close**.
5. To open the *Fire and Smoke Wizard* dialog, click (**Fire and Smoke Wizard**).
6. Click the *Fire Dynamics* tab and set the following properties:

Property	**Setting**
Turn On Fire Regions	**Fire**
Select Pressure Boundary	**Room: Surface shells 16**
Max Simulation Time	`200 sec`
alpha =	`3`
p =	`2`
Limit Q	`0.0629 MW`

7. Click **Set Graph Parameters**.
8. In *Graph Parameters Dialog*, set *End X Range* to `50` and click **OK**.



When these steps are complete, the dialog appears as follows:



9. Click **Apply**.

In this process, a fire source, height-varying pressure boundary setting, and scalar source are all set-up automatically.

10. Save the simulation.

Reviewing Model Settings

To make sure that the model setup is accurate, you can inspect the selected models and edit them as necessary.

This tutorial describes an unsteady problem in which:

- A three-dimensional grid is employed.
- The fluid is a slightly compressible gas (air) that obeys the ideal gas equation of state.
- The flow is turbulent and non-isothermal.
- The $k-\varepsilon$ low Re model is used for representing turbulence effects.
- The segregated flow solver is used.
- Gravity effects are taken into account.

All of the above conditions are automatically set up as properties of the fire. To make sure that the model setup is accurate, you can inspect the selected models and edit them as necessary.

1. Open the **Continua > Physics 1 > Models** node.



In this case, the model setup is correct and no change is required.

Setting up Solver Parameters and Stopping Criteria

Before running the analysis, specify the maximum physical time over which the simulation is run.

Set the time step and the temporal (and spatial) discretization schemes. Also reduce the maximum number of iterations that are carried out for each time step.

1. Expand the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	*Time-Step*	`1.0 s`

Node	Property	Setting
	Temporal Discretization	**2nd-order**
Stopping Criteria > Maximum Physical Time	*Maximum Physical Time*	`200.0 s`
Maximum Inner Iterations	*Maximum Inner Iterations*	`5`

Visualizing the Solution

To view the progress of the simulation as it runs, set up scalar scenes using a section plane.

First, create a section through the fire for use in a temperature scene:

1. Right-click the **Derived Parts** node and select the option **New Part > Section > Plane**.
2. Set the following properties:

Property	Setting
Input Parts	**Room** and **Fire**
Plane Parameters > Origin	X: 0 Y: 0 Z: 0
Plane Parameters > Normal	X: 0 Y: 1 Z: 0

3. Click **Create** then **Close**.
4. Open the **Derived Parts** node.

   A node, **plane section**, has been added to the **Derived Parts** node. The position of the plane section is shown in *Mesh Scene 1*:



Create a scalar scene to display the temperature on the plane section.

5. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
6. Open the **Scalar Scene 1** > **Displayers** > **Outline 1** node. Right-click the **Parts** node and select **Edit**.

In the *Edit* dialog:

7. To select all objects, click once on the **Column Selection** button in the upper left portion of the dialog.



8. Expand the nodes and make sure that all the objects are selected.



9. Click **OK**.
10. Expand the **Displayers** > **Scalar 1** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**

Node	Property	Setting
Parts	*Parts*	**Derived Parts > plane section**
Scalar Field	*Function*	**Temperature**

11. Use the mouse buttons to display *Scalar Scene 1* as shown below:



During the Fire properties set up: reports, monitors, and plots for the fire heat release rate are all set up automatically.

12. Open the **Reports**, **Monitors**, and **Plots** nodes.



13. Save the simulation.

Running the Simulation

Run the simulation.

1. To start the solver, click **(Run)**.

   The analysis now starts and the results, in the form of a graph of residuals vs. iteration number, is displayed automatically in the *Graphics* window. An example of a residual plot is shown in a separate part of the User Guide.

   The solution stops after 200 time-steps.

Analyzing Results

Analyze the contours of temperature for the converged solution.

The contours of temperature on the plane section that is displayed in *Scalar Scene 1* are dynamically updated with each time-step.

Adjust the temperature range that is displayed in the scalar scene to help you visualize the results.

1. Expand the **Scenes** > **Scalar Scene 1** > **Displayers** > **Scalar 1** node.
2. Select the **Scalar Field** node and set the following properties:

Property	Value
Min	300
Max	400
Clip	**Off**

A scene similar to the one shown below is displayed.



Examine the Heat Release Rate in the Fire region:

3. Expand the **Plots** node.
4. Right-click the **Fire heat release rate plot** node and select **Open**.
5. Select the **Fire heat release rate plot** > **Axes** > **X Axis** node and set *Maximum* to 50.

   A scene similar to the one shown below is now displayed.



6. Save the simulation.

Settings for Steckler Room Fire Validation

Set heat transfer coefficients for the surfaces that represent the walls and ceiling.

In the Steckler room fire experiment, the walls and ceiling were composed of conduction brick of thickness 0.1m and conductivity of 0.69W/m-K. Set an effective heat transfer coefficient of 6.9 W/m^2-K for each of the five shell surfaces representing the walls and ceiling of the room.

1. Expand the **Regions** > **Room** > **Boundaries** > **Surface shells 9** node.
2. Select the **Physics Conditions** > **Thermal Specification** node and set *Condition* to **Convection**.
3. Expand the **Physics Values** node and set the following properties:

Node	Property	Setting
Heat Transfer Coefficient > Constant	*Value*	`6.9` (default unit W/m^2-K)
Ambient Temperature > Constant	*Value*	`300.0 K`

4. Repeat the above steps for **Surface shells 10**, **11**, **14**, and **18**.

Switch on the radiation model and set the gray-gas absorptivity to a field function based on temperature:

5. Open the **Continua > Physics 1** node.
6. Right-click the **Models** node and select **Select Models**.
7. In the *Physics Model Selection* dialog, select the following models in order:

Group Box	Model
Optional Models	**Radiation**
Radiation	**Participating Media Radiation (DOM)**
Radiation Spectrum (Participating)	**Gray Thermal Radiation**

The *Physics Model Selection* dialog has the following physics models enabled:



8. Click **Close**
9. Select the **Models** > **Gas** > **Air** > **Material Properties** > **Absorption Coefficient** node and set the following properties:

Node	Property	Setting
Absorption Coefficient	*Method*	**Field Function**
Field Function	*Scalar Function*	**AbsorTempFunction**

10. Save the simulation.

Comparison with Measurements

To compare the predictions with the experimental data, first load the files containing the vertical temperatures and horizontal velocities at the door.

1. Expand the **Tools** node.
2. Right-click the **Tables** node and select **New Table** > **File Table**.

3. In the *Open* dialog, navigate to your `stecklerFire` directory and select the `DoorTempExpt.csv` and `DoorUvelExpt.csv` files.
4. To read the data, click **Open**.

   The `DoorTempExpt.csv` and `DoorUvelExpt.csv` files are now listed under the **Tables** node.



Create a line probe against which to plot the predicted profiles:

5. Click the *Geometry Scene 1* tab at the top of the *Graphics* window to activate the scene.
6. Right-click the **Derived Parts** node and select **New Part > Probe > Line**.
7. Set the following properties for the line probe:

Property	Setting
Input Parts	**Room**
Point 1	X: `1.4` m Y: `0` m Z: `0` m
Point 2	X: `1.4` m Y: `0` m Z: `1.83` m
Display	**New Geometry Displayer**
Resolution	`20`

8. Click **Create** then **Close**.

   A **Derived Parts > line-probe** node appears in the simulation tree.

Create a plot for temperature comparison with measurement using the following settings:

9. Right-click the **Plots** node and select **New Plot > X-Y**.
10. Rename the **XY Plot 1** node to `Door Temperature`.
11. Edit the **Plots > Door Temperature** node and set the following properties:

Node	Property	Setting
Door Temperature	*Title*	`Door Temperature`
	Parts	**Derived Parts > line-probe**
X Type > Vector Quantity	*Value*	[`0, 0, 1`] m
Y Types > Y Type 1	*Smooth Values*	activated
Scalar Function	*Scalar Function*	**Temperature**
line-probe	*Sort Plot Data*	activated
Line Style	*Style*	**Solid**
Axes		
X-Axis > Title	*Title*	`Position (m)`
Legend	*x-Position*	`0.8`
	y-Position	`0.1`

12. Right-click the **Plots > Door Temperature > Data Series** node and select **Add Data**.
13. In the *Add Data Providers to Plot* dialog, select **DoorTempExpt** and click **OK**.
14. Select the **Door Temperature > Y-Axis > Title** node and set *Title* to `Temperature (K)`.
15. Edit the **DoorTempExpt** node and set the following properties:

Node	Property	Setting
DoorTempExpt	*Legend Name*	`Experimental Temp.`
	Lock Legend Name	activated
	Table	DoorTempExpt

Node	Property	Setting
	X Column	Y
	Y Column	Temperature [K]
	Sort Plot Data	activated
Symbol Style	*Color*	**Blue**
	Shape	**Empty Square**

The usual orientation for the position vector is the vertical axis.

16. Select the **Axes** node.
17. Edit the **Axis Orientation** and select the x-up, y-right option.



18. Click **OK**.

The temperature comparison graph setup is now complete. Create a similar plot for the door velocity.

19. Copy and paste the **Door Temperature** plot and rename as `Door Velocity`.
20. Edit the **Door Velocity** node and set the following properties:

Node	Property	Setting
Door Velocity	*Title*	`Door Velocity`
Y Types		
Y Type 1 > Scalar Function	*Scalar Function*	**Velocity[i]**

Node	Property	Setting
Axes		
Y-Axis > Title	*Title*	`Velocity (m/s)`

21. Expand the **Door Velocity** > **Data Series** > **DoorTempExpt** node and set the following properties:

Property	Setting
Legend Name	`Experimental Velocity`
Table	**DoorUvelExpt**
X Column	**Y**
Y Column	**V [m/s]**

The velocity comparison graph setup is now complete.

Adjusting Stopping Criteria and Continuing

Increase the maximum physical time of the simulation.

To allow you to solve for a further 200 time-steps, increase the maximum physical time to 400 seconds.

1. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `400.0 s`.
2. Save the simulation as `roomFireValidation.sim`.
3. To start the solver, click **(Run)**.
   The solution stops after 400 time-steps.

Analyzing Validation Results

This section of the tutorial covers the analysis of the validation results.

The following scene shows the solution after 400 time-steps, plotted in the temperature range of 300 K to 400 K.



1. Open the *Door Temperature* plot which shows comparison at the defined line-probe location.



The temperature profile from the simulation is in good agreement with the profile based on experimental data.

2. Open the *Door Velocity* plot which shows comparison at the defined line-probe.



The simulated velocity profile shows good agreement with the experimental data up to a height of 1.5 m. After this point, the experimental data indicates that the maximum velocity in the boundary layer is higher than the simulation predicts. One cause for this difference could be an insufficient density of cells in the upper region of the doorway. Good simulation practice would be to reduce the volume mesh **Base Size** to 0.05 m, remesh the domain, and rerun the simulation using the finer mesh. Using a finer mesh size gives closer agreement between the simulated and experimental results.

Summary

This tutorial demonstrated how to use the Fire and Smoke Wizard.

The steps covered were:

- Invoking the Fire and Smoke Wizard.
- Defining the simulation models.

- Setting up radiation properties.
- Analyzing the results using the built-in plotting and visualization facilities, including the XY Plot, comparing the experimental data with the predictions.

Solid Stress

The tutorials in this set illustrate various STAR-CCM+ features for computing deformation, strain, and stresses in solid regions. They also show how such computations can be coupled to the fluid behavior in an analysis of fluid-structure interaction.

Contents

Solid Stress: Exhaust Manifold
Fluid Structure Interaction: Vibrating Pipe

Solid Stress: Exhaust Manifold

This tutorial demonstrates how to use the Solid Stress solver in STAR-CCM+, including the application of rigid body constraints.

The part used in the simulation is an exhaust manifold whose geometry is shown below.



In this tutorial a surface mesh is imported and used to create a polyhedral volume mesh. The mesh extrusion tool is used to add a length of pipe to the outlet to improve resolution of the highly turbulent flow. The simulation incorporates both conjugate heat transfer (CHT) and solid stress analysis.

In this example you are seeking the static, steady-state solutions in the fluid and solid. In this situation it is both convenient and accurate to initially solve for the fluid flow and the conjugate heat transfer and then later solve for the solid stress while holding the thermal field fixed. This is achieved by freezing and unfreezing the appropriate solvers.

The same result could also be achieved by running the fluid flow and CHT solvers simultaneously with the stress solver, but the stress solver may require more iterations to converge the case.

The post-processing tools in STAR-CCM+ are used to visualize the results and aid in analyzing the stress through the solid region.

Contents

Prerequisites

Importing the Surface Mesh
Visualizing the Imported Geometry
Splitting the Region
Renaming the Boundaries
Setting Up the Mesh
Selecting Physics Models
Changing the Solid Material and Specifying Material Properties
Setting Global Initial Conditions
Assigning Physics Continua
Defining Boundary Conditions
Running the Polyhedral Volume Mesher
Applying Rigid Body Constraints
Setting Up Monitors
Setting Solver Variables
Freezing Solvers
Setting Stopping Criteria
Running the Simulation
Freezing and Unfreezing Solvers
Running the Simulation
Assessing Convergence
Visualizing the Results
Summary

Prerequisites

The instructions in the Solid Stress - Exhaust Manifold tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Surface Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Surface Mesh...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/solidStress/data`.
5. Select `manifold.dbs`, then click **Open**.
6. In the *Import Surface Options* dialog, select **Create New Region**.
7. To accept the remaining defaults, click **OK**.
8. Save the simulation as `manifold.sim`.

Visualizing the Imported Geometry

Display the imported geometry and store a convenient view.

1. Use the mouse buttons to obtain a similar view as shown below.



2. Click (**Save-Restore-Select views**) and select **Store Current View**.
   The view has been stored as **View 1**.
3. Save the simulation.

Splitting the Region

Split the imported region into two separate regions, representing the manifold and the fluid inside the manifold.

During preparation of the surface mesh a set of boundaries were defined. Following the import process they were grouped under a single region. To allow different properties to be applied to fluid and solid regions, the region needs to be split into two separate regions:

1. Expand the **Regions** node.
2. Right-click the **Region 1** node and select **Split by Surface Topology**.
3. In the *Split Region Options* dialog, make sure that *Combine Regions With Shared Boundaries* is not selected.
4. Click **OK**.

   The region is split into two separate regions, **Region 1** and **Region 1 2**.
5. Rename **Region 1** to `Solid` and **Region 1 2** to `Gas`.

Renaming the Boundaries

The model currently has the default boundary names. Change these to more suitable names to aid in setting up the simulation and to ease understanding of the model mechanics.

1. Expand the **Regions** > **Solid** > **Boundaries** node.
2. Rename the boundaries as follows:

Original Name	New Name
Boundary 3	`solid-wall`
Boundary 4	`flanges`
Boundary 5	`boltheads`
Boundary 6	`outside-wall`
Boundary 7	`exhaust-connection`
Boundary 8	`shank`

3. Expand the **Regions** > **Gas** > **Boundaries** node.
4. Rename the boundaries as follows:

Original Name	New Name
Boundary 1	`inlets`

Original Name	New Name
Boundary 2	`outlet`
Copy of Boundary 3	`gas-wall`

Notice that the **Copy of Boundary 3 [Boundary 3]** node changes automatically. This is because it comprises part of an internal interface between the solid and fluid regions in the manifold.

The interface is defined by a boundary in each of the regions, that is **gas-wall** in the fluid region and **solid-wall** in the solid region. The interface boundaries also display the interface name which is shown in square brackets. Change the contents of the square brackets by renaming the interface:

5. Rename the **Interfaces** > **Boundary 3** node to `interface`.

Rearrange the regions and boundaries nodes into alphabetical order:

6. Right-click the **Regions** node and select **Refresh**.
7. Right-click the **Gas** > **Boundaries** and **Solid** > **Boundaries** and select **Refresh**.
8. Save the simulation.

Setting Up the Mesh

Set up the mesh for this simulation. This includes selecting the meshing models, setting global mesh reference values, and setting the meshing model parameters.

Due to the geometry of the manifold, the gas flowing through the part is highly turbulent at the outlet. As this could limit the likelihood of the solution converging, use the mesh extruder feature in STAR-CCM+ to add a length of pipe at the manifold outlet. This aids good flow characteristics and helps the solution to converge. Although it is unlikely that such a length of pipe represents an exhaust pipe in reality, it is appropriate for a case such as this where you are mainly concerned with stress concentrations in the solid region.

Use the polyhedral mesher and prism layer mesher to create a volume mesh. The part already has a good surface mesh, so the surface remesher tool is not necessary (although this would be necessary if a different mesh size were required).

Contents

Selecting Meshing Models

Setting the Global Mesh Reference Values
Specifying Prism Layer Parameters
Setting the Extrusion Parameters

Selecting Meshing Models

Select meshing models for the imported geometry.

Use the polyhedral mesher and prism layer mesher to create a volume mesh. Also use the mesh extruder model to add a length of pipe at the manifold outlet.

1. For the mesh continuum, **Mesh 1**, select the following meshing models:

Group box	Model
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**
Optional Models	**Extruder**

2. Save the simulation.

Setting the Global Mesh Reference Values

Define the global reference values for the volume mesh.

Set the mesh base size to a custom value to better suit the intricacy of the manifold geometry.

1. Expand the **Continua** > **Mesh 1** > **Reference Values** node.
2. Select the **Base Size** node and set *Value* to `0.005 m`.
3. Save the simulation.

Specifying Prism Layer Parameters

Set the prism layers to grow only in the Gas region, and to occur at the interface.

Prism layers are desirable in the fluid region of the model only, where they are used to resolve boundary layers.

To set the prism layers to grow from the interface only in the Gas region:

1. Expand the **Regions** > **Solid** > **Mesh Conditions** node.

2. Select the **Customize Prism Mesh** node and set *Customize Prism Mesh* to **Disable**.
3. Expand the **Interfaces** > **interface** > **Mesh Conditions** node.
4. Select the **Interface Prism Layer Option** node and activate *Grow Prisms from Interface*.



5. Save the simulation.

Setting the Extrusion Parameters

Set up the mesh extruder to add a straight section of pipe at the manifold outlet.

1. Expand the **Regions** > **Gas** > **Boundaries** > **outlet** > **Mesh Conditions** node.
2. Select the **Extrusion Type** node and set *Extrusion Type* to **Constant Rate Normal**.
3. Expand the **Regions** > **Gas** > **Boundaries** > **outlet** > **Mesh Values** node.
4. Edit the **Normal Extrusion Parameters** node and set the following properties:

Property	Setting
Magnitude	`0.25 m`
Number of Layers	`100`
Stretching	`10`
Specify a New Region	**None**

5. Save the simulation.

Selecting Physics Models

As this model contains a fluid and a solid region, two physics continua are required. The physics models define the primary variables of the simulation, including pressure, temperature, velocity, and which mathematical formulation is used to generate the solution.

In this example, the flow is turbulent, so use the K-Omega turbulence model. Use the option of cell quality remediation to limit the effect of bad cells on the solution.

To select the physics models:

1. Create two new physics continua.
2. Rename the **Continua > Physics 1** node to `Gas`, and the **Continua > Physics 2** node to `Solid`.
3. For the physics continuum, **Continua > Gas**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Ideal Gas**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Omega Turbulence**
Optional Models	**Cell Quality Remediation**
Energy	**Segregated Fluid Temperature**

4. Click **Close**.

   The **Continua > Gas > Models** node appears as shown below:



In this example, use the **Solid Stress Iterative Solver** to resolve the stress throughout the solid part. Proceed by selecting the physics models for the Solid continuum.

5. For the physics continuum, **Continua** > **Solid**, select the following models in order:

Group box	**Model**
Space	**Three-Dimensional**
Time	**Steady**
Material	**Solid**
Optional Models	**Finite Volume Solid Stress**
	Segregated Solid Energy
Equation of State	**Constant Density**
Optional Models	**Cell Quality Remediation**

6. Click **Close**.

The **Continua** > **Solid** > **Models** node appears as shown below:



7. Save the simulation.

Changing the Solid Material and Specifying Material Properties

Change the default solid material to ductile iron, as this is more appropriate to the manifold case.

Select a different material model, and alter the material properties.

1. Expand the **Continua** > **Solid** > **Models** > **Solid** node.
2. Right-click the **Al** node and select **Replace with...**



3. In the *Replace Material* dialog, open the **Material Database > Standard > Solids** node, and select **Ductile Iron (Ductile Iron)**.



4. Click **OK**.

   The **Al** node changes to **Ductile Iron**.

5. Expand the **Continua** > **Solid** > **Models** > **Solid** > **Ductile Iron** node.
6. Edit the **Material Properties** node and set the following properties:

Node	Property	Setting
Poisson Coefficient > **Constant**	*Value*	`0.283`
Thermal Expansion Coefficient > **Constant**	*Value*	`1.38E-5 /K`
Young's Modulus > **Constant**	*Value*	`138000 MPa`
Zero Thermal Strain Reference Temperature > **Constant**	*Value*	`485 K`

7. Save the simulation.

Setting Global Initial Conditions

Specify the global initial values.

As initial conditions, specify the fluid temperature and velocity:

1. Edit the **Continua** > **Gas** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Static Temperature > **Constant**	*Value*	`450 K`
Velocity > **Constant**	*Value*	`[0, 0, 10] m/s`

2. Save the simulation.

Assigning Physics Continua

Assign physics continua to regions.

To assign the physics continua to regions:

1. Select the **Regions** > **Solid** node and set *Physics Continuum* to **Solid**.



The default settings for the **Regions** > **Gas** node are appropriate.

The settings for the **Gas** and **Solid** regions appear as shown below:

Node	Property	Setting
Regions > **Gas**	*Physics Continuum*	**Gas**
	Type	**Fluid Region**
Regions > **Solid**	*Physics Continuum*	**Solid**
	Type	**Solid Region**

2. Save the simulation.

Defining Boundary Conditions

Specify the type of each boundary and set the specific initial conditions that are required for the simulation.

To specify the boundary type:

1. Edit the **Regions** > **Gas** > **Boundaries** node.
2. Set the **Type** property for the **inlets** and **outlet** nodes:

Boundary	Type
inlets	**Velocity Inlet**
outlet	**Pressure Outlet**

The **inlets** and **outlet** node icons have changed to denote a change in type. It is useful to become familiar with such icons, as you are able to see the type of each boundary at a glance.



Set the velocity and temperature of the gas at the manifold inlets:

3. Edit the **Regions** > **Gas** > **Boundaries** > **inlets** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Velocity Specification	*Method*	**Components**
Physics Values		
Velocity > **Constant**	*Value*	`[0, 0, 10] m/s`
Static Temperature > **Constant**	*Value*	`500 K`

4. Save the simulation.

Running the Polyhedral Volume Mesher

The mesh has now been fully defined, and as the part already has a surface mesh you are able to proceed straight to meshing the volume.

If you had run the prism layer mesher before specifying the **inlets** node as **Velocity Inlet** and the **outlet** node as **Pressure Outlet**, the mesher would have treated them in the same way as a wall boundary, and would have created prism layers on those boundaries.

1. Generate the volume mesh.

   The output from the meshers is shown in the *Output* window and the ongoing status is shown in the status bar.

The volume mesh is shown below:



2. Save the simulation.

Applying Rigid Body Constraints

Apply rigid body constraints to prevent any rigid body motion.

When using the solid stress solver, it is important to make sure that the part is fully constrained to prevent any rigid body motion. This is a necessity as, depending on the geometry of the part and the loads applied, moments might be set up and forces applied which would theoretically cause the part to move or rotate in space. By constraining the part in such a way as to inhibit any potential motion you can resolve the resulting deformations, strains and stress concentrations.

In this example constrain the part near the inlets and outlet to mimic a real life scenario. It is good practice to consider which constraints are used and where they are applied when setting up the initial CAD model. Similarly, the limitations of FVA (Finite Volume Analysis) as compared to FEA (Finite Element Analysis) must be considered, and the model designed accordingly.

STAR-CCM+ offers a variety of rigid body constraints which may be applied at boundaries. In this example, use the Fixed and Normal Displacement constraints.

1. Expand the **Regions** > **Solid** > **Boundaries** node.

Use the same settings for the boltheads and exhaust-connection nodes:

2. Edit the **boltheads** node, and set the following properties:

Node	Property	Setting
Physics Conditions		
Stress/ Displacement	*Method*	**Fixed**
Thermal Specification	*Condition*	**Convection**
Physics Values		
Ambient Temperature > Constant	*Value*	`360 K`
Heat Transfer Coefficient > Constant	*Value*	`6.667 W/m^2-K`

3. Repeat the previous step for the **exhaust-connection** node.

   The Fixed constraint has been applied to the **boltheads** and **exhaust-connection** to prevent motion in all degrees of freedom.

4. Edit the **flanges** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Stress/ Displacement	*Method*	**Normal Displacement**
Thermal Specification	*Condition*	**Convection**
Physics Values		
Normal Displacement > Constant	*Value*	`0 m`
Ambient Temperature > Constant	*Value*	`360 K`

Node	Property	Setting
Heat Transfer Coefficient > Constant	*Value*	`333.33 W/m^2-K`

The part is now fully constrained: a zero normal displacement constraint has been applied to the **flanges** node to restrict movement normal to the plane of the flanges.

By default, the stress or displacement method applied to boundaries is Free (that is, no constraint is applied). As a constraint is applied the boundary node icon changes; use this feature to make sure that the remaining solid boundaries are set to the default **Stress/Displacement** method.

5. Expand the **Regions** > **Solid** > **Boundaries** node.
   The nodes appear as shown below:



6. Edit the **outside-wall** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Convection**
Physics Values		

Node	Property	Setting
Ambient Temperature > Constant	*Value*	`360 K`
Heat Transfer Coefficient > Constant	*Value*	`6.667 W/m^2-K`

The **shank** and **solid-wall** boundaries are set to Adiabatic . This is the default Thermal Specification.

The stress/displacement methods and thermal specifications have now been set for all the solid boundaries.

7. Save the simulation.

Setting Up Monitors

Create monitors for temperature and solid stress displacement.

By default STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the solution. In addition to this it is useful to create additional monitors to help in assessing convergence of the solution. In cases where rigid body constraints have been applied, avoid placing monitors near to the location of the constraints as the solution in these locations may be unreliable.

Create a plane section through the part to monitor the temperature and solid stress displacement while the solvers are running.

1. Within the **Derived Parts** node, create a new **Plane Section** with the following settings:

Property	Setting
Input Parts	**Solid, Gas** (default selection)
origin	`[ -0.1, 0, 0.1 ] m`
normal	`[ 0.4, 0, 1 ] m`
Display	**No Displayer**

To set up a temperature monitor and plot:

2. Right-click the **Reports** node and select **New report > Surface Average**.
3. Rename the **Reports > Surface Average 1** node to `Temperature`.

4. Edit the **Reports** > **Temperature** node and set the following properties:

Property	Setting
Scalar Field Function	**Temperature**
Parts	**[plane section]**

5. Right-click the **Temperature** report and select **Create Monitor and Plot from Report**.

   A new monitor and plot have been created, named **Temperature Monitor** and **Temperature Monitor Plot**, respectively.



To set up a solid stress monitor and plot:

6. Copy and paste the **Reports** > **Temperature** node onto the **Reports** node.
7. Rename the **Reports** > **Copy of Temperature** node to `Solid Stress Displacement`.
8. Edit the **Reports** > **Solid Stress Displacement** node and set the following properties:

Property	Setting
Scalar Field Function	**Solid Stress Displacement > Magnitude**
Units	**um**

9. Right-click the **Solid Stress Displacement** report and select **Create Monitor and Plot from Report**.

   The necessary monitors have now been created. To display all plots, right-click the **Plots** node and select **Open All Plots**.

10. Save the simulation.

Setting Solver Variables

Modify the Segregated Energy and FVA Stress Solver parameters.

STAR-CCM+ allows users to enter custom values for a number of the variables used by the solvers. The default values are usually appropriate to most cases, however in this example use a custom under-relaxation factor for the solid segregated energy solver, and turn on Successive Over-Relaxation in the solid stress iterative solver. Use successive over-relaxation to accelerate solution convergence when the solid stress iterative solver is employed.

1. Select the **Solvers > Segregated Energy** node and set *Solid Under-Relaxation Factor* to `1`.
2. Select the **Solvers > FVA Stress Solver** node and activate *Successive Over-Relaxation*.



3. Save the simulation.

Freezing Solvers

Freeze the solid stress solver to solve for fluid flow and conjugate heat transfer.

The simulation is now ready to be run, and if desirable, you could set it to run at this stage. However, due to the nature of the case used in this example it is advantageous to solve for fluid flow and conjugate heat transfer before solving for solid stress. To do this, freeze the solid stress iterative solver and then run the simulation until convergence is achieved. When you have solved for conjugate heat transfer and fluid flow, freeze all solvers apart from the solid stress iterative solver and then solve the solid stress part of the case.

To freeze the solid stress iterative solver:

1. Select the **Solvers > FVA Stress Solver** node and activate *Solver Frozen*.



2. Save the simulation.

Setting Stopping Criteria

Set the maximum number of iterations for the first run.

The default number of iterations is 1000 iterations, but as you are freezing and unfreezing solvers in this example, reduce this to 300 iterations for the first run. Add an additional set of iterations to solve for solid stress.

1. Edit the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `300`.
2. Save the simulation.

Running the Simulation

The simulation is now ready to be run.

1. Run the simulation.

   The iteration progress can be monitored in the *Output* window. Monitor the convergence by viewing the **Temperature Monitor Plot** and the **Residuals** using the visualization tabs in the *Graphics*

window. When the last iteration (that is, no. 300) has been completed, the message **Stopping criterion Maximum Steps satisfied** is displayed in the *Output* window.

2. When the stopping criterion has been satisfied, save the simulation.

Freezing and Unfreezing Solvers

You now have a solution for the conjugate heat transfer case and can solve for solid stress.

You can see from the Temperature Monitor Plot that the solution has converged:



You now have a solution for the conjugate heat transfer case and can solve for solid stress. To accept the solution for conjugate heat transfer, and to prevent further calculations from being carried out, freeze the fluid flow and conjugate heat transfer solvers. To solve for solid stress, unfreeze the solid stress iterative solver.

The fluid flow and conjugate heat transfer solvers can be frozen simultaneously.

1. Expand the **Solvers** node.

2. Multi-select the **Wall Distance**, **Segregated Flow**, **Segregated Energy**, **K-Omega Turbulence** and **K-Omega Turbulent Viscosity** nodes, and activate *Solver Frozen*.



3. To unfreeze the solid stress iterative solver, edit the **Solvers > FVA Stress Solver** node and deactivate *Solver Frozen*.

   The solvers have been frozen and unfrozen as required. You must add an additional set of iterations to solve the solid stress case.
4. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `450`.
5. Save the simulation.

Running the Simulation

The simulation is now ready to be run for a second time.

To run the simulation:

1. Click (**Run**).
2. When the stopping criterion has been satisfied, save the simulation.

Assessing Convergence

Creating additional monitors can provide a useful tool in cases such as this where complex geometry and flow characteristics result in fluctuations in the residuals. By selecting the Solid Stress Displacement Monitor Plot visualization tab in the Graphics window, you can use the surface average solid stress displacement monitor to assess convergence.

To be sure that the solution has converged to an acceptable level, you can view the plot in the Graphics window, or read the values yielded for a specific monitor after each iteration.

In the Solid Stress Displacement Monitor Plot at 300 iterations, when the solid stress iterative solver is unfrozen, the solid stress displacement

leaps to almost 70 micrometers. Plots created in STAR-CCM+ are automatically scaled, however in this case, this is not useful.

Alter the scale on the X and Y axes to provide a better view of the monitor value.

1. Edit the **Plots** > **Solid Stress Displacement Plot** > **Axes** node and set the following properties:

Node	Property	Setting
X Axis	*Minimum*	300
	Maximum	450
Y Axis	*Minimum*	63
	Maximum	65
Major Labels	*Spacing*	0.5

2. Click the **Solid Stress Displacement Monitor Plot** visualization tab in the *Graphics* window to make the plot active.



The calculated values for the **Solid Stress Displacement monitor** can also be viewed in the *Output* window.

3. Use the scroll bars to locate values for **Solid Stress Displacement** in the *Output* window.
4. Save the simulation.

   Once the accepted level of convergence has been reached the post-processing stage can begin.

Visualizing the Results

To visualize the stress throughout the solid region of the part create a scalar scene. Use the threshold function to analyze areas of high stress which may result in failure due to shear stress and tension.

Create a scalar scene to display maximum principal stress and equivalent stress throughout the solid region.

1. Create a new scalar scene and rename it to `Solid Stress`.
2. Edit the **Scenes > Solid Stress > Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1 > Parts	*Parts*	**boltheads**
		exhaust-connection
		flanges
		outside-wall
		shank
		solid-wall
		solid-wall [interface]
Scalar 1 > Scalar Field	*Function*	**Solid Stress Max. Principal**

   The solid boundaries are now visible in the *Graphics* window.

The **Solid Stress Max. Principal** scalar function is useful for identifying potential failure sites due to tension.

3. Use the mouse buttons to rotate the part to observe areas of high stress.
   To restore the view that you saved at the beginning of this tutorial,

   click (**Save-Restore-Select Views**) and select **Restore View > View 1**.



You may notice that there are stress concentrations around the boltheads. As this is the boundary at which one of the rigid body constraints was applied, results in this area may not give an accurate representation of what would happen in reality. For this reason, it is more useful to examine the stress further away from the points at which the constraints were applied.

4. Save the simulation.

Contents

Creating a Scalar Threshold Part

Creating a Scalar Threshold Part

Create a threshold part in a geometry scene which can be used to show areas of stress above a certain level.

1. Create a new geometry scene and rename it to `Stress Threshold`.
2. Within the **Derived Parts** node, create a new **Threshold** with the following settings:

Property	Setting
Input Parts	**Solid**
Scalar	**Solid Stress Equivalent**
Extraction Mode	**All Between**
Min	`100 MPa`
Max	`250 MPa`
Display	**New Scalar Displayer**

3. To close the *Create Threshold* dialog, click **Create**, then **Close**.
4. Expand the **Scenes** > **Stress Threshold** > **Displayers** > **Geometry 1** node.
5. Edit the **Parts** node and, in the *Parts* selection dialog, deselect the **Regions** > **Gas** and **Regions** > **Solid** > **solid-wall [Interface]** nodes.
6. Edit the **Scenes** > **Stress Threshold** > **Displayers** node and set the following properties:

Node	Property	Setting
Geometry 1	*Opacity*	0.3
Threshold Scalar 1 > **Scalar Field**	*Function*	**Solid Stress Equivalent**

The *Stress Threshold* scene appears as shown below:



7. Save the simulation.

Summary

This tutorial demonstrated how to use the FVA Solid Stress solver in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Setting extrusion parameters.
- Changing the solid material and specifying material properties.
- Assigning physics continua to regions.
- Applying rigid body constraints.
- Assessing convergence.
- Visualizing the stress distribution throughout the solid.
- Visualizing areas where the stress is above a specified threshold.

Fluid Structure Interaction: Vibrating Pipe

This tutorial demonstrates how to set up fluid-structure interaction (FSI) problems in STAR-CCM+. Water flow in a slender pipe under a 350 g load is used in the simulation to describe the required steps.

The geometry with some of the boundary types is shown below.



The geometry consists of a simple half-pipe model, 0.02 m in diameter and 1.20 m in length, where flow enters through a velocity inlet and exits through a pressure outlet. The inner wall of the pipe acts as a solid boundary to the fluid flow but is dynamically coupled to the fluid domain only at the middle section of the overall geometry. In this section, the pipe is represented as a solid domain 0.0005 m thick and 0.40 m long. Both ends of this section are fixed and its outer surface is stress-free. The rest of the pipe is treated as a rigid boundary.

This tutorial is comprised of two parts:

1. In the first part of the simulation the case is run under a body load of 350 g applied to the pipe to obtain a deformed steady-state solution.
2. In the second part of the simulation the body load is removed and the pipe is allowed to vibrate in real time.

The mass of the pipe is relatively small compared to the mass of the fluid it contains. So the physics between the interaction of the fluid and the structure in the second transient part is inherently strongly coupled. This tutorial demonstrates the robustness of implicit coupling techniques for these types of problems.

Contents

Prerequisites
Importing the Volume Mesh
Visualizing the Volume Mesh

Creating Interfaces
Selecting Fluid Physics Models
Selecting Solid Physics Models
Selecting Motion Models
Assigning Motion Models
Specifying Fluid Material Properties
Specifying Solid Material Properties
Defining Body Load
Setting Global Initial Conditions
Defining Fluid Boundary Conditions
Defining Solid Boundary Conditions
Specifying Explicit Fluid Solid Coupling
Setting Solver Variables
Setting Stopping Criteria
Setting up Probe, Report, Monitor, and Plot
Running the Simulation
Viewing Velocity Vectors
Viewing Y-Displacement Contours
Viewing the FSI Displacement Time Monitor Plot
Setting Up Transient Solution
Running the Simulation
Summary

Prerequisites

The instructions in the Fluid Structure Interaction - Vibrating Pipe tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Volume Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/solidStress/data`.
5. Select `vibratingPipe.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `vibratingPipeH2O.sim`.

Visualizing the Volume Mesh

Create a mesh scene and adjust the view using the scene properties.

1. Create a mesh scene.
2. Edit the **Scenes > Mesh Scene 1 > Attributes > View** node and set the following properties:

Property	Setting
Position	`[-0.06, 0.03, 0.87] m`
Focal Point	`[-0.004, -0.002, 0.73] m`
View Up	`[0.04, 1.0, -0.2] m`

3. Click **Apply** then **Close**.
4. Click (**Save-Restore-Select views**) and select **Store Current View**.

   The mesh is shown below:



5. Save the simulation.

Creating Interfaces

Create internal interfaces between the fluid and solid pipe regions.

The interface is defined by existing boundaries in each of the regions, that is PipeContact in the fluid region and FluidContact in the solid region. When created, the interface makes copies of the existing boundaries, which are distinguished from existing boundaries by having a different icon, and by displaying the interface name alongside the boundary name.

1. Expand the **Regions** > **Fluid** > **Boundaries** and **Regions** > **Pipe** > **Boundaries** nodes.
2. Multi-select **PipeContact** and **FluidContact**.



3. Right-click either **PipeContact** or **FluidContact** and select **Create Interface**.



This creates **PipeContact [Interface 1]** and **FluidContact [Interface 1]** entries, as well as the **Interfaces** node in the simulation tree.



4. Save the simulation.

Selecting Fluid Physics Models

The problem is solved as a transient, three-dimensional, laminar, compressible flow case. Water with constant velocity of 0.01 m/s is used. A user-defined density for water is prescribed using a field function.

To select the physics models:

1. Rename the **Continua** > **Physics 1** node to `Fluid`.
2. For the physics continuum, **Continua** > **Fluid**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically)
	Gradients (Selected automatically)
Time	**Implicit Unsteady**
Material	**Liquid**
	H2O (Selected automatically)
Flow	**Segregated Flow**
Equation of State	**User Defined EOS**
Viscous Regime	**Laminar**
Optional Models	**Segregated Fluid Temperature**

3. Click **Close**.
4. To review the models, expand the **Continua** > **Fluid** > **Models** node.



5. Save the simulation.

Selecting Solid Physics Models

Select physics models for the solid continuum.

To select the physics models:

1. Rename the **Continua** > **Physics 2** node to `Solid`.
2. For the physics continuum, **Continua** > **Solid**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Solid**
Optional Models	**Finite Volume Solid Stress**
	Linear isotropic Elastic(Selected automatically)
	Gradients (Selected automatically)
	Segregated Solid Energy
	Gravity
Equation of State	**Constant Density**

3. Click **Close**.
4. To review the models, expand the **Continua** > **Solid** > **Models** node.



5. Save the simulation.

Selecting Motion Models

Select the appropriate motion models to simulate the deformation of the fluid and solid regions under the external load on the pipe.

There are three broad categories for defining motion in STAR-CCM+:

1. Actual displacement of mesh vertices in real time. Use these methods with a transient analysis.
2. An unsteady motion problem is converted to a steady-state problem by imposing a moving frame of reference on a static mesh. Rigid rotation and /or translation can be simulated in this way.
3. Imposed modal displacement of the flutter boundary that is applied to the frequency domain solver. This displacement is a corollary of the real-time displacement and applies only to harmonic balance flutter.

Models from the first category are the most appropriate for this tutorial. Allow the fluid mesh representing the inner region of the pipe to be deformed under the external load on the pipe (morphing). Allow the solid mesh representing the pipe to be displaced in real time (solid displacement).

1. Right-click the **Tools** > **Motions** node and select **New** > **Morphing**.
2. Repeat the above step and select **New** > **Solid Displacement**.

   This portion of the object tree appears as shown below:



3. Save the simulation.

Assigning Motion Models

Assign the motion models to the appropriate region.

1. Expand the **Regions** > **Fluid** > **Physics Values** node.
2. Select the **Motion Specification** node and set *Motion* to **Morphing**.



3. Expand the **Regions** > **Pipe** > **Physics Values** node.
4. Select the **Motion Specification** node and set *Motion* to **Solid Displacement**.
5. Save the simulation.

Specifying Fluid Material Properties

Use field functions to specify the speed of sound, density, and density pressure derivative properties. Use a combination of predefined system and user field functions.

In this tutorial, the speed of sound in H_2O is assumed to be constant at 200. To define a constant field function that specifies the speed of sound:

1. Right-click the **Tools** > **Field Functions** node and select **New** > **Scalar**.
2. Rename the **User Field Function 1** node to `H2O Speed Of Sound`.
3. Select the **H2O Speed Of Sound** node and set the following properties:

Property	Setting
Function Name	`H2OSpeedOfSound`
Definition	`200`

The density expression is a function of the speed of sound. It is: 1000+Pressure/(Speed of Sound)^2. To define the field function for density:

4. Create another scalar field function.
5. Rename the **User Field Function 1** node to `H2ODensity`.
6. Select the **H2ODensity** node and set the following properties:

Property	Setting
Function Name	`H2ODensity`
Definition	`1000+$Pressure/ pow($H2OSpeedOfSound,2)`

The density pressure derivative expression is a function of the speed of sound. It is: 1/(Speed of Sound)^2. To define the field function for the density pressure derivative:

7. Create another scalar field function.
8. Rename the **User Field Function 1** node to `H2O Density Pressure Derivative`.
9. Select the **H2O Density Pressure Derivative** node and set the following properties:

Property	Setting
Function Name	`H2ODensityPressureDerivative`
Definition	`1/pow($H2OSpeedOfSound,2)`

Specify the fluid material properties:

10. Select the **Continua** > **Fluid** > **Models** > **User Defined EOS** node and activate *Compressible*.



11. Expand the **Continua** > **Fluid** > **Models** > **Liquid** > **H2O** node.
12. Edit the **Material Properties** node and set the following properties:

Node	Property	Setting
Speed of Sound	*Method*	**Field Function**

Node	Property	Setting
Field Function	*Scalar Function*	**H2O Speed Of Sound**
Density	*Method*	**Field Function**
Field Function	*Scalar Function*	**H2O Density**
Density pressure derivative	*Method*	**Field Function**
Field Function	*Scalar Function*	**H2O Density Pressure Derivative**

13. Save the simulation.

Specifying Solid Material Properties

Modify the default properties of Aluminum.

1. Expand the **Continua** > **Solid** > **Models** > **Solid** > **Al** node.
2. Edit the **Material Properties** node and set the following properties:

Node	Property	Setting
Density > **Constant**	*Value*	`4500.0 kg/m^3`
Poisson Coefficient > **Constant**	*Value*	`0.34`
Young's Modulus > **Constant**	*Value*	`102.7 GPa`

3. Save the simulation.

Defining Body Load

The body load is applied to the pipe by increasing the gravity value.

Increase the gravity value to 350 g, equivalent to 3433.5 m/s^2 . This value is selected as it causes a suitable level of displacement in the pipe.

1. Edit the **Continua** > **Solid** > **Reference values** > **Gravity** node and set *Value* to `[0,-3433.5,0] m/s^2`.



2. Save the simulation.

Setting Global Initial Conditions

Specify the global initial values.

Specify the fluid initial velocity:

1. Edit the **Continua** > **Fluid** > **Initial Conditions** > **Velocity** > **Constant** node and set *Value* to `[0, 0, 0.001] m/s`.
2. Save the simulation.

Defining Fluid Boundary Conditions

Define the specific conditions at each boundary.

1. Expand the **Regions** > **Fluid** > **Boundaries** node.
2. Select the **Inlet** > **Physics Values** > **Velocity Magnitude** > **Constant** node and set *Value* to `0.001 m/s`.
3. Edit the **Symmetry** > **Physics Values** > **Morpher Plane Options** node and set the following properties:

Property	Setting
Damping Factor	`0.002 m`
Ignore Factor	`5.0e-4`

4. Save the simulation.

Defining Solid Boundary Conditions

The pipe is fixed at both ends. Specify the rigid body boundary condition for the pipe region.

1. Edit the **Regions** > **Pipe** > **Boundaries** node and set the following properties:

Node	Property	Setting
Left		
Physics Conditions > Stress/ Displacement	*Method*	**Fixed**
Right		
Physics Conditions > Stress/ Displacement	*Method*	**Fixed**

2. Save the simulation.

Specifying Explicit Fluid Solid Coupling

Select the morpher method and visualize the coupling settings.

Since in this first part of the simulation the intent is to reach a steady-state solution, the physical coupling does not need to be strong. Therefore using explicit coupling is sufficient, providing a suitable level of computational speed and stability.

1. Expand the **Regions > Fluid > Boundaries > PipeContact > Physics Conditions** node.
2. Select the **Morpher** node and set *Method* to **Solid Stress**.

The default settings for the coupling method are appropriate. To review the settings:

3. Select the **Interfaces > Interface 1 > Physics Conditions > FSI Coupling Algorithm** node.

   Method is set to the default **Explicit**.



4. Save the simulation.

Setting Solver Variables

Set the solver properties.

In this first part of the simulation, use the Implicit Unsteady solver to reach a steady-state solution. Since the result is independent of the time step, the actual time step is not too important. Use the default value of 1.0 s with the default first order Temporal Discretization. The first order scheme with this relatively large time step has sufficient numerical damping to drive the solution to the steady-state quickly. In addition, there is no reason to consider the flux that is related to the motion of the pipe since, in the steady-state, the pipe velocity is zero. Customize the Mesh Morpher solver parameters to neglect the grid flux terms in this first part of this simulation.

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	*Time-Step*	`1.0 s`
Mesh Morpher	*Recompute Interfaces*	**Deactivated**
	Ignore gridflux terms	**Activated**
FVA Stress Solver > AMG Linear Solver > F Cycle	*Post-Sweeps*	`1`

Note: If *Recompute Interfaces* is not deactivated, the displacement solution on the interface is over-written at the end of each iteration.

2. Save the simulation.

Setting Stopping Criteria

Set the standard stopping criteria and create an additional stopping criterion from the maximum Y-force values.

To create a stopping criterion from the maximum Y-force values:

1. Right-click the **Stopping Criteria** node and select **Create New Criterion > Create from Monitor...**
2. Select **Y-Force** from the *Select Monitor* dialog.
3. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`
Maximum Physical Time	*Maximum Physical Time*	`100 s`
Y-Force Criterion > Minimum Limit	*Minimum Value*	`5e-4`

4. Save the simulation.

Setting up Probe, Report, Monitor, and Plot

Create a monitor point in the middle of the pipe to monitor the solid stress displacement while the solvers are running.

1. Within the **Derived Parts** node, create a new point probe with the following settings:

Property	Setting
Input Parts	**Pipe**
Point	`[-0.01025, 0, 0.2] m`
Display	**No Displayer**



2. In the Edit window, click **Create** and then **Close**.



Create a report for the probe, then a monitor and plot.

3. Right-click the **Reports** node and select **New Report > Maximum**.
4. Rename the **Maximum 1** node to `FSI Displacement Time`.
5. Edit the **Reports > FSI Displacement Time** node, and set the following properties:

Property	Setting
Scalar Field Function	**Solid Stress Displacement > Laboratory > j**
Parts	**Derived Parts > point**
Units	**mm**

6. Right-click the **FSI Displacement Time** report and select **Create Monitor and Plot from Report**.

   A new monitor named **FSI Displacement Time Monitor** and its associated plot named **FSI Displacement Time Monitor Plot** have been created in the **Monitors** and **Plots** nodes, respectively.
7. To give the plot axes more suitable titles, edit the **Plots > FSI Displacement Time Monitor Plot > Axes** node, and set the following properties:

Node	Property	Setting
X Axis > Title	*Title*	`Time (s)`
Y Axis > Title	*Title*	`Y-Displacement (mm)`

Open the **FSI Displacement Time Monitor Plot**, to visualize it in the *Graphics* window.

8. Save the simulation.

Running the Simulation

Initialize the solution and run the simulation.

The simulation is now ready to be run:

1. Click (**Initialize solution**).
2. Click (**Run**).
3. When the stopping criterion has been satisfied, save the simulation.

Viewing Velocity Vectors

Create a vector scene to display the velocity vectors in the pipe.

1. Create a vector scene.
2. Edit the **Scenes > Vector Scene 1 > Displayers** node and set the following properties:

Node	Property	Setting
Vector 1 > Parts	*Parts*	**Regions**
Vector 1 > Color Bar	*Width*	`0.9`
	Height	`0.15`
	Position	`[0.04,0.04]`

3. Restore View 1.



4. Save the simulation.

Viewing Y-Displacement Contours

Create a scalar scene to display the Y-Displacement contour.

To display the Y-Displacement contour:

1. Create a scalar scene.

   The **Scalar Scene 1** node appears within the **Scenes** node with a corresponding display in the *Graphics* window.
2. Edit the **Scenes > Scalar Scene 1 > Displayers** node and set the following properties:

Node	Property	Setting
Outline 1		
Parts	*Parts*	**Pipe**
Scalar 1	*Contour Style*	**Smooth Blended**
Parts	*Parts*	**Pipe**

Node	Property	Setting
Scalar Field	*Function*	**Solid Stress Displacement > Laboratory > j**
	Units	**mm**
Color Bar	*Width*	`0.85`
	Height	`0.15`
	Position	`[0.07,0.2]`

3. Click (**Save-Restore-Select views**) and select **Views > -X > Up +Y**.
4. Use the mouse buttons to zoom-in and position the display as shown in the following screenshot.



5. Save the simulation.

Viewing the FSI Displacement Time Monitor Plot

Examine the FSI Displacement Time Monitor Plot.

A scene similar to the one shown below is displayed.

The deformed steady-state analysis run is now complete. Set up the second part of the tutorial where the body load is removed and the pipe is allowed to vibrate in real time.

Setting Up Transient Solution

The transient solution is found by removing the body load from the pipe, allowing it to vibrate.

To remove the body load:

1. Remove the gravity model for the **Solid** continuum.
   The updated **Continua** > **Solid** > **Models** tree appears as below:



This part of the simulation requires a strong fluid-solid coupling. Change the FSI coupling method to implicit:

2. Expand the **Interfaces** > **Interface 1** > **Physics Conditions** node.
3. Select the **FSI Coupling Algorithm** node and set *Method* to **Implicit**.

In this transient part of the simulation, the second order Temporal discretization provides a more accurate solution with little damping. A time step size of about 1/400th of the natural period of the pipe/water system is sufficiently small to give high accuracy and stability.

Modify the solver and stopping criteria for this part of the simulation:

4. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	*Time Step*	`2.4e-5`
	Temporal Discretization	**2nd-Order**
Mesh Morpher	*Ignore gridflux terms*	**Deactivated**

5. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`50`
Maximum Physical Time	*Maximum Physical Time*	`0.0036 s`

The transient solution set-up is now complete.

6. Save the simulation.

Contents

Visualizing the Pipe Deformation
Writing Animation Files for External Use (Optional)

Visualizing the Pipe Deformation

The deformation of the pipe can be displayed by using the warp feature of the STAR-CCM+ software. Create a scalar scene to display the warped solid stress displacement magnitude.

To create a vector warp:

1. Create a scalar scene.
2. Right-click the **Derived Parts** node and select **New Part > Warp > Vector Warp**.
3. In the *Create Warp Surface* window:
   a) Select as *Input Parts* all the boundaries within the **Pipe** region.
   b) Set the remaining properties as follows:

Property	Setting
Vector	**Solid Stress Displacement**
Scale Factor	200
Display	**Existing Displayer > Scalar 1**

4. Click **Create**, then **Close**.
   A derived part, **vector warp**, is created within the **Derived Parts** node.



5. Edit the **Scenes > Scalar Scene 2 > Displayers** node and set the following properties:

Node	Property	Setting
Outline 1		
Parts	*Parts*	**Pipe**
Scalar 1	*Contour Style*	**Smooth Blended**
Scalar Field	*Function*	**Solid Stress Displacement > Magnitude**
	Units	**mm**
Color Bar	*Width*	`0.9`
	Position	`[0.04,0.15]`

6. Click (**Save-Restore-Select views**) and select **Views > +X > Up +Y**.
7. Save the simulation.

Writing Animation Files for External Use (Optional)

Write out a scene as a graphics file at specified time-steps, to form animations using external software packages.

Write out the solid stress displacement magnitude scene every six time-steps, resulting in writing out a total of 25 files. These files can then be used to form an animation of the pipe deformation.

1. Edit the **Scalar Scene 2 > Attributes** node and set the following properties:

Node	Property	Setting
Update	*Save To File*	**Activated**
	Trigger	**Time Step**
Time-Step Frequency	*Time-step Frequency*	`6`
Hardcopy	*Output Width*	`800`
	Output Height	`600`

   The files with `.png` extension are saved in the current working directory and have a base name of image.
2. Save the simulation as `vibratingPipeH2OTransient`.

Running the Simulation

Reset the iteration number and run the simulation.

Before running the transient simulation, reset the iteration number to start from one:

1. From the **Solution** menu, select **Clear Solution...**
2. In the *Clear Solution* dialog, clear *Fields* and *Reset Mesh*, then click **OK**.



3. Run the simulation.

   The iteration progress can be monitored in the *Output* window. Monitor the convergence by viewing the **Residuals** using the visualization tabs in the *Graphics* window.

   When the last time step corresponding to the physical time of `0.0036 s` has been completed, the message Stopping Criterion Maximum Physical Time satisfied is displayed in the *Output* window.
4. When the stopping criterion has been satisfied, save the simulation.

Contents

Warped Displacement Magnitude Plot
FSI Displacement Time Monitor Plot

Warped Displacement Magnitude Plot

1. Use the mouse buttons to zoom-in and position the display as shown in the following screenshot:



FSI Displacement Time Monitor Plot

Modify the Y-axis scale:

1. Edit the **Plots > FSI Displacement Time Monitor Plot > Axes** node and set the following properties:

Node	Property	Setting
Y Axis	*Minimum*	`-0.25`
	Maximum	`0.25`
Major Labels	*Spacing*	`0.05`

Setting the *Minimum* and *Maximum* properties automatically activates the *Lock Minimum* and *Lock Maximum* properties.

A scene similar to the one shown below is displayed.



Summary

This tutorial demonstrated how to set up fluid-structure interaction (FSI) problems in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Creating interfaces.
- Selecting fluid and solid physics models.
- Selecting and assigning motion models.
- Specifying fluid and solid material properties.
- Defining body load.
- Specifying explicit and implicit fluid-solid coupling.
- Applying rigid body constraints.
- Setting up probes.
- Setting stopping criteria.
- Visualizing scalar stress.
- Visualizing stress displacement time plot.

- Creating a scalar warp derived part.

Aeroacoustics

The tutorials in this set illustrate various STAR-CCM+ features for solving aeroacoustic simulations.

Contents

Broadband Models: Noise from a Cylinder (Preparation)
DES and FW-H On-The-Fly: Noise from a Cylinder (Unsteady Analysis)
Ffowcs Williams-Hawkings: Sound Propagation
Signal Post-Processing: FFT and Wavenumber
Acoustic Wave Modeling: Noise from a Cylinder

Broadband Models: Noise from a Cylinder (Preparation)

This tutorial demonstrates how to set up aeroacoustics problems in STAR-CCM+. The flow around a cylinder is used in the simulation to describe the required steps.

The geometry is shown below.



The geometry consists of a cylinder, 20 mm in diameter, placed at the center of a three-dimensional circular computational domain, 0.6 m in diameter. The free-stream air velocity and temperature are set to be 50 m/s and 300K, respectively, corresponding to a Mach number of 0.144 (50.0/347.2) and a Reynolds number of 6.38E4 (1.184*50*0.02/1.85563E-5).

According to the recommended procedure, you first run in steady state to:

- Use the broadband models to understand where the aeroacoustics sources are located and to indicate where the mesh needs refining.
- Use the Mesh Frequency Cutoff estimator to prescribe the absolute cell size in the source region to satisfy frequency resolution requirements in transient flow.

Then, you activate an LES-based turbulence model in transient mode to capture the aeroacoustics sources directly and to:

- Use the post-processing tools to perform spectral analysis.
- Export the CFD fields for far-field noise propagation (suitable for non-compact sources).

Contents

Loading the Base Simulation
Visualizing the Geometry
Generating the Volume Mesh
Selecting Fluid Physics Models
Setting Global Initial Conditions
Defining Boundary Conditions
Setting Up a Plane Section
Setting Up Scalar Scenes
Monitoring Force Coefficients
Setting Stopping Criteria
Running the Simulation
Viewing the Results
Applying the Broadband Noise Sources
Results Analysis and Consequences
Refining the Volume Mesh
Setting New Stopping Criteria
Running the Simulation
Viewing Results After Mesh Refinement
Summary

Loading the Base Simulation

A base simulation file has been prepared for this analysis. Continue by copying the file from the STAR-CCM+ installation directory, starting the simulation, loading the base simulation file and saving it under a different file name.

Note: Use the double precision version of STAR-CCM+ for all aeroacoustic and hydroacoustic analyses.

1. Create a directory for the tutorial called `aeroacoustics`.
2. In this directory, copy the file containing the problem geometry (`cylinderBase.sim`) from the `doc/startutorialsdata/`

`aeroacoustics/data` subdirectory of your STAR-CCM+ installation directory. Also copy the two macro files, `compareRmsDragLiftCoefficient.java` and `compareRmsDragLiftCoefficientRefinedMesh.java`.

3. Start up a double precision version of STAR-CCM+ in a manner that is appropriate to your working environment.
4. Click (**Load Simulation**) from the tool bar and open `cylinderBase.sim`.
5. Save the new simulation to disk under file name `cylinder.sim`.

Visualizing the Geometry

Display the geometry scene.

1. Right-click the **Scenes** > **Geometry Scene 1** node and select **Open**.
2. Click (**Make Scene Transparent**) from the tool bar.

A similar view to the one shown below appears.



Store the current view.

3. Click (**Save-Restore-Select views**) in the toolbar and select **Store Current View**.
4. Examine the *Geometry Scene 1* display in the *Graphics* window by adjusting the viewing angle and zooming as appropriate. Click and select **Restore View** > **View 1**.

Generating the Volume Mesh

The meshing models and mesh reference values for this simulation have already been set up. A volumetric control has been created for mesh refinement, to capture the separation behind the cylinder, and to capture the vortex shedding.

The mesh is ready to be generated.

1. Click (**Generate Volume Mesh**) in the **Mesh Generation** toolbar.

To visualize the volume mesh:

2. Right-click the **Scenes** node and select **New Scene** > **Mesh**.

The mesh appears as shown below.



Selecting Fluid Physics Models

1. For the physics continuum,**Continua** > **Physics 1**, select the following models in order:

The models define the primary variables of the simulation, including pressure, temperature and velocity. Specify the mathematical formulation that is used to generate the solution.

Group Box	**Model**
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Ideal Gas**
Energy	**Segregated Fluid Temperature**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Omega Turbulence**

2. Click **Close**.

The selected models appear as shown below:



Setting Global Initial Conditions

To enhance convergence, initialize the flow field with a constant velocity of 50 m/s.

To set the initial conditions:

1. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Velocity** > **Constant** node and set *Value* to `[ 50, 0.0, 0.0]` m/s.

   All other values remain at default.
2. Save the simulation.

Defining Boundary Conditions

It is only necessary to set the Mach number of the freestream boundary, as all other settings have acceptable default values.

- Select the **Regions** > **Body 1** > **Boundaries** > **Freestream** > **Physics Values** > **Mach Number** > **Constant** node and set the *Value* to `0.144`.

Setting Up a Plane Section

Create a plane section through the part with which to display the scalar and vector quantities while the solvers are running.

1. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane...**

   The *Create Section* dialog appears in the *cylinder* window.
2. Specify exact values to define the plane section:

Property	Value
Input Parts	**Body 1**
Plane Parameters	
▮▮*origin*	`[0, 0, 0.02]` m
▮▮*normal*	`[0, 0, 1]` m
Display	**No Displayer**

The *Create Section* dialog now appears as shown below.



3. Click **Create**, then **Close**.

   A new derived part, **plane section**, is created under the **Derived Parts** node.

4. Save the simulation.

Setting Up Scalar Scenes

Two scalar scenes are created.

- Velocity Magnitude Contour
- Wall Y+ Contour

Contents

Velocity Magnitude Contour
Wall Y+ Contour

Velocity Magnitude Contour

Create a scalar scene to display the velocity magnitude contour on the plane section.

Use the following steps:

1. Right-click the **Scenes** node and select **New Scene** > **Scalar**.
2. Rename the **Scalar Scene 1** node to `Velocity`.
3. Select the **Displayers** > **Outline 1** node and activate *Surface*.
4. Select the **Displayers** > **Outline 1** > **Parts** node and set *Parts* to **Regions** > **Body 1** > **Boundaries** > **Cylinderwall**.
5. Select the **Displayers** > **Scalar 1** node and set the *Contour Style* to `Smooth Filled`.
6. Select the **Displayers** > **Scalar 1** > **Parts** node and set *Parts* to **Derived Parts** > **plane section**.
7. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and click to the right of *Function* and select **Velocity** > **Magnitude** from the list.



8. Set *Max* to `75`
9. Set *Clip* to `Off`.
10. Use the mouse buttons, as described previously, to adjust the viewing angle, zoom, and position the scene as shown below.



11. Store the current view. Click the (**Save-Restore-Select views**) button and select **Store Current View**.

Wall Y+ Contour

Create a scalar scene to display the Wall Y+ scalar field.

1. Create another scalar scene and rename it to **Wall Y+**.
2. Select the **Displayers** > **Outline 1** > **Parts** node and set *Parts* to **Regions** > **Body 1** > **Boundaries** > **Cylinderwall**.
3. Select the **Displayers** > **Scalar 1** > **Parts** node and set *Parts* to **Regions** > **Body 1** > **Cylinderwall**.
4. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set *Function* to **Wall Y+**.
5. Click the (**Save-Restore-Select views**) button, and select **Projection mode** > **Perspective**.
6. Use the mouse buttons, to zoom and position the scene as shown below.



Monitoring Force Coefficients

By default STAR-CCM+ provides a residuals plot which can be used to monitor the convergence of the solution. In addition to this, it is useful

to create additional reports and monitors to help in assessing convergence.

Create two reports for monitoring force coefficients while the solvers are running:

1. Right-click the **Reports** node and select **New Report > Force Coefficient**.
2. Rename the **Force Coefficient 1** node to `Drag Coefficient`.
3. In the *Properties* window, set the following properties:

Property	Setting
Reference Density	`1.18 kg/m^3`
Reference Velocity	`50 m/s`
Reference Area	`8E-4 m^s`
Parts	**Regions > Body 1 > Boundaries > Cylinderwall**

This report can be copied and used as a basis for the lift coefficient report:

4. Copy and paste the **Drag Coefficient** node.
5. Rename the **Copy of Drag Coefficient** node to `Lift Coefficient`.
6. In the *Properties* window, set the *Direction* to `[0, 1, 0]`



Create a monitor and plot from the reports:

7. Multi-select the **Drag Coefficient** and **Lift Coefficient** nodes. Right-click the **Drag Coefficient** node and select **Create Monitor and Plot from Report.**

8. Click **Single Plot** to display the coefficients for both probes in a single graph.

Two new monitors have been created, with their associated plot.

9. Rename the **Plots** > **Reports Plot** to **Force Coefficients Plot**
10. Save the simulation.

Setting Stopping Criteria

Set the **Maximum Steps** as stopping criterion for this simulation.

Set the stopping criteria.

1. Edit the **Stopping Criteria** node and set the **Maximum Steps** to `200`.
2. Save the simulation.

Running the Simulation

Before running the simulation, reorder the mesh to optimize it for the solver.

1. Right-click **Regions** > **Body 1** and choose **Reorder Mesh**.
2. In the *Reorder Regions* dialog, activate **Use directional reordering**, then click **Apply**, then **Close**.

The simulation is now ready to be run.

3. Click (**Run**) in the *Solution* toolbar.
   The solution runs for 200 iterations. Observe that the residuals and other plots adopt an oscillatory pattern about a mean. This pattern is evidence that the steady state approach does not satisfy the real physics of the problem; an unsteady approach is required. However, the purpose of this steady state solution is to qualify the mesh before committing resources to the full unsteady simulation.
4. Save the simulation.

Viewing the Results

View the Velocity Magnitude Contour and Wall Y+ Contour scenes.

1. Click the *Velocity Magnitude Contour* tab.
   A scene similar to the one shown below appears.



The above plot illustrates the representative wake field.

2. Click the *Wall Y+ Contour* tab.

   A scene similar to the one shown below appears:



Wall Y+ values are everywhere below 1.

3. Close the *Wall Y+* scene in the *Graphics* window.

Contents

Extracting Drag and Lift Coefficient RMS Values

Extracting Drag and Lift Coefficient RMS Values

Examine the root-mean-squares of drag and list coefficients.

1. Select **File** > **Macro** > **Play Macro**
2. In the *Open* dialog, navigate to your working directory and select `compareRmsDragLiftCoefficient.java`.
3. Click **Open** to run the macro.

The values appear in the *Output* window.

Applying the Broadband Noise Sources

The Broadband Noise Source models are intended to be used as tools for qualifying a volume mesh before a transient acoustic analysis. They are applied in the post-processing stage of a preliminary steady analysis.

Contents

Curle Surface Acoustic Power dB Contour
Proudman Acoustic Power Contour
Mesh Frequency Cutoff Contour

Curle Surface Acoustic Power dB Contour

Activate the Curle model first and create a scalar scene to display the Curle surface acoustic power contour on the cylinder wall.

The broadband noise source models are activated from the model selection dialog in the **Continua** > **Physics 1** node.

1. Open the **Continua** > **Physics 1** node and select the following models, in order:

Group Box	Model
Optional Models	**Aeroacoustics**
Aeroacoustics Models	**Broadband Noise Sources**

Group Box	Model
Broadband Noise Source Models	**Curle**

The selected models appear as shown below:



2. Copy and paste the **Scenes** > **Wall Y+** node.
3. Rename the new node to `Curle AC Power`.
4. Right-click the **Curle AC Power** node and choose **Open**.
5. Click the *scene/plot* tab.
6. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set *Function* to **Curle Surface Acoustic Power dB**.
7. Select the **Displayers** > **Scalar 1** node and set the *Contour Style* to **Smooth Blended**.
8. Click the (**Step**) button in the **Solution** toolbar.

A scene similar to the one shown below appears:



The Curle acoustic power model, which assumes an isotropic turbulent flow field, computes the sound that dipole sources generates. This model therefore represents the fluctuating surface pressure with which the solid boundary acts on the flow. The plot displays the noise level that the turbulent boundary layer emits over a surface at low Mach number.

9. Close the *Curle AC Power* scene in the *Graphics* window.

Proudman Acoustic Power Contour

Activate the Proudman model instead of the Curle model and create a scalar scene to display the Proudman surface acoustic power on the symmetry plane. This model indicates acoustic quadrupole sources.

1. Open the **Continua** > **Physics 1** node and select the following models, in order:

Group Box	Model
Enabled Models	**Curle**
Broadband Noise Source Models	**Proudman**

2. Copy and paste the **Scenes** > **Velocity** node.
3. Rename the new node to `Proudman AC Power`.
4. Right-click the **Proudman AC Power** node and choose **Open**.
5. Select the **Displayers** > **Scalar 1** > **Scalar Field** node and set *Function* to **Proudman Acoustic Power dB**.

A scene similar to the one shown below appears:



The Proudman acoustic power (in dB) estimates the local contribution of the quadrupole sources as generated by isotropic turbulence. It displays the acoustic power per unit volume as emitted by the turbulence structures, in the flow-field surrounding the cylinder.

6. Close the *Proudman AC Power* scene in the *Graphics* window.

Mesh Frequency Cutoff Contour

The mesh frequency cutoff field function becomes available when the Curle, Proudman, and Goldstein broadband noise source model are activated. Create a scalar scene to display the mesh frequency cutoff contour on the symmetry plane.

1. Copy and paste the **Scenes > Proudman AC Power** node.
2. Rename the new node to `Mesh Frequency Cutoff`.
3. Right-click the **Mesh Frequency Cutoff** node and choose **Open**.
4. Select the **Displayers > Scalar 1 > Scalar Field** node and set **Function** to **Mesh Frequency Cutoff Hz**.

   A scene similar to the one shown below appears:



This plot indicates that the mesh is able to resolve up to 2000 Hz, approximately, in the main volume source regions the Proudman correlation indicates. The mesh frequency cut-off estimate becomes less accurate in the boundary-extruded layer where the cells are non-isotropic.

5. Save the simulation.

Results Analysis and Consequences

The cylinder and cone refinements seem sufficient in terms of capturing the aeroacoustics sources behind the cylinder (as the Proudman source plot illustrates). Typically, however, the peak of the human hearing range is between 1-3 kHz. Therefore, refine the mesh further to increase the resolvable frequency.

Refining the Volume Mesh

Continuing from the `cylinder.sim` file, you now create a cylindrical part, 0.02 m in radius, which defines the outer extent of the region to be refined.

1. Click the *Geometry Scene 1* tab in the *graphics* window.
2. Right-click the **Geometry** > **Parts** node and select **New Shape Part** > **Cylinder**.
3. In the *Create/Edit Cylinder Part* panel, set the following values:

Property	Setting
Start Coordinate	`[0, 0, 0] m`
End Coordinate	`[0, 0, 0.041] m`
Cylinder Radius	`[0.02] m`

4. Select the **Geometry** > **Parts** > **Cylinder** node, to display the cylinder part that you created in *Geometry Scene 1*.



You now create a volumetric control and select the shape that defines the source.

5. Expand the **Geometry** > **Operations** > **Automated Mesh** node.
6. Right-click the **Custom Controls** node and select **New** > **Volumetric Control**.

   A new node, **Volumetric Control 2**, is created in the simulation tree.
7. Select the **Volumetric Control 2** node and set *Parts* to **Cylinder**.
8. Expand the **Volumetric Controls 2** > **Controls** node and select **Trimmer**.
9. In the *Properties* window, enable *Customize isotropic size*.



A new node,**Values**, is added to the simulation tree.

10. Select the **Volumetric Controls 2** > **Values** > **Custom Size** node and set *Size Type* to **Absolute**.
11. Select the **Custom Size** > **Absolute Size** node and set *Value* to `8.0E-4 m`.



12. Click (**Generate Volume Mesh**) in the **Mesh Generation** toolbar.

The output from the meshers is shown in the *Output* window and the ongoing status is shown in the status bar. The solution obtained on the coarse mesh is interpolated onto the new mesh.

13. Save the simulation.

Setting New Stopping Criteria

Set a new value for **Maximum Steps**.

Use the following steps:

1. Open the **Stopping Criteria** node and select **Maximum Steps**.
2. In the *Properties* window, set the *Maximum Steps* to `400`.

Running the Simulation

The simulation is now ready to be run.

1. Run the simulation by clicking (**Run**) in the **Solution** toolbar.
2. Save the simulation on completion.

Viewing Results After Mesh Refinement

Visualize the updated *Mesh Frequency Cutoff* scene in the *Graphics* window.

1. Select the **Scenes** > **Mesh Frequency Cutoff** > **Displayers** > **Scalar 1** > **Scalar Field** node and set the *Min* to `0` and the *Max* to `5000`.

   A scene similar to the one shown below appears:



From the above plot, it is clear that the fine mesh adequately resolves higher frequencies close to the cylinder.

The steady-state analysis runs are now complete. Set up the final transient run in the next part of the tutorial.

2. Save the simulation.

Contents

Extracting Drag and Lift Coefficient RMS Values

Extracting Drag and Lift Coefficient RMS Values

Examine the drag and list coefficients again, after the mesh has been refined.

1. Select **File** > **Macro** > **Play Macro**
2. In the *Open* dialog, navigate to your working directory and select `compareRmsDragLiftCoefficientRefinedMesh.java`.
3. Click **Open** to run the macro.

The values appear in the *Output* window.

Summary

This tutorial demonstrated how to set up aeroacoustics problems in STAR-CCM+, using the flow around a cylinder.

The tutorial introduced the following STAR-CCM+ features:

- Loading an existing base simulation.
- Visualizing the geometry.
- Generating the volume mesh.
- Selecting physics models.
- Setting up initial and boundary conditions.
- Setting up stopping criteria.
- Creating plane sectione and scenes to visualize the results.
- Applying the broadband noise sources.
- Running the solver for a set number of iterations.
- Refining the volume mesh and rerunning the solver.

DES and FW-H On-The-Fly: Noise from a Cylinder (Unsteady Analysis)

In the previous section of this tutorial, you completed the steady-state analysis runs for qualifying the volume mesh. In this section, continuing from the previous simulation file, you set up the final transient run for an unsteady analysis of the near field and an FW-H analysis of the far field.

Two point receivers are used for the far field analysis, each 1 m from the cylinder.

Note: Use the double precision version ofSTAR-CCM+ for all aeroacoustic and hydroacoustic analyses.

If you have closed the previous simulation file:

- Click (**Load Simulation**) from the tool bar and load `cylinder.sim`.

Contents

Saving File for Transient Simulation
Selecting the Transient Calculation Physics Models
Setting the Solver Properties and Stopping Criteria
Visualizing the Mean Velocity
Monitoring Pressure
Monitoring Force Coefficients
Saving Transient Surface Data
Preparing the On-The-Fly FW-H Analysis
Visualizing the Solution
Running the Simulation
Visualizing the Results
Spectral Analysis
Results Analysis
Summary
DES and FW-H On-The-Fly Tutorial Bibliography

Saving File for Transient Simulation

The steady-state analysis runs are now complete. Set up the final transient run.

- Save a copy of the simulation file under the name `cylinderUnsteady.sim` for the transient calculation to follow.

Selecting the Transient Calculation Physics Models

Modify the physics model selections to activate the unsteady solver.

1. For the physics continuum, **Continua** > **Physics 1**, deselect the following models from the *Enabled Models* list:
   - **Proudman**
   - **Broadband Noise Sources**
   - **Aeroacoustics**
   - **All y+ Wall Treatment**
   - **SST (Menter) K-Omega**
   - **K-Omega Turbulence**
   - **Reynolds-Averaged Navier-Stokes**
   - **Turbulent**
   - **Steady**
2. Deactivate the **Auto-select recommended models** option.
3. Select the following models, in order:

Group Box	Model
Time	**Implicit Unsteady**
Viscous Regime	**Turbulent**
Turbulence	**Detached Eddy Simulation**
Detached Eddy Simulation	**SST (Menter) K-Omega Detached Eddy**
K-Omega Wall Treatment	**All y+ Wall Treatment**

The physics models appear as below:



4. Open the **Continua > Physics 1 > Models** node, select **Segregated Flow** and set *Convection* to **Hybrid**.
5. Open the **Continua > Physics 1 > Models** node, select **SST (Menter) K-Omega Detached Eddy** and set the *Formulation Option* property to **DDES**.

Setting the Solver Properties and Stopping Criteria

Choosing the correct time-step is an important step in setting up an aeroacoustic simulation.

You typically evaluate the following criteria, and choose the smallest value:

- Convective Courant Number: Implicit solvers are usually stable at maximum values in the range 10-100 locally, but the overall mean value should be around 1. For this case, a time-step of 2.5E-5 s gives a Convective Courant Number whose mean value is approximately 1.
- Maximum resolvable frequency: For example, to obtain 10 points in a wave at 3000 Hz, (around the point at which human hearing is most sensitive), the time-step size should be 10 x 3000 Hz = 30000 Hz, or 3.33E-5 s.

- Local Strouhal shedding: At the Reynolds number observed in this tutorial, 63800, a cylinder sheds vortices at a Strouhal number of approximately 0.22. This number corresponds to a shedding frequency of 500 Hz. Capturing this shedding frequency requires a time-step size of 2.0E-4 s.

The smallest time-step size is the one required by the Convective Courant Number criterion, 2.5E-5 s.

To set the time-step size:

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `2.5E-5 s`.
2. Set *Temporal-Discretization* to `2nd-order`.

Now set suitable under-relaxation factors:

3. Select the **Solvers** > **Segregated Flow** > **Velocity** node and set *Under-Relaxation Factor* to `0.9`.
4. Select the **Segregated Flow** > **Pressure** node and set *Under-Relaxation Factor* to `0.7`.

The case is run until 0.12 s and with 5 iterations per time-step:

5. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`5`
Maximum Physical Time	*Maximum Physical Time*	`0.12` s
Maximum Steps	*Enabled*	Deactivated

As the simulation is unsteady, update the trigger on the monitors for the lift and drag coefficients:

6. Within the **Monitors** node, multi-select the **Drag Coefficient Monitor** and **Lift Coefficient Monitor** nodes, and set *Trigger* to **Time Step**.
7. Save the simulation.

Visualizing the Mean Velocity

Create a monitor for displaying the mean velocity magnitude in a scalar scene.

1. Right-click the **Monitors** node and select **New Monitor > Field Mean**.

   Field Mean 1 is now added to the simulation tree.
2. Select the **Field Mean 1** node and set the following properties:

Property	Setting
Field Function	**Velocity Magnitude**
Parts	**Regions > Body 1**

3. Select the **Field Mean 1 > Time-Step Frequency** node and set *Start Time-step* to `800`.

Create a scalar scene to display the mean velocity field:

4. Copy and paste the **Scenes > Velocity** node and paste it into the **Scenes** node.
5. Rename the **Copy of Velocity** node to `Mean Velocity Magnitude`.
6. Right-click the **Mean Velocity Magnitude** node and choose **Open**.
7. Select the **Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Mean of Velocity: Magnitude**.



Monitoring Pressure

Create three probes that are used to monitor pressure while the solvers are running.

Use the following steps:

1. Right-click the **Scenes > Geometry Scene 1** node and select **Open**.
2. In the simulation tree, right-click the **Derived Parts** node and select **New Part > Probe > Point**.
3. In the *Create Point* panel, set the following properties:

Group Box	Setting
Input Parts	**Body 1**
Point	`[0.0101, 0, 1.0E-4]` m, m, m

Group Box	Setting
Display	**No Displayer**

A new derived part is created, called **point**.

4. Rename the **Derived Parts** > **point** node to `Probe-back`.

Create two more probes by copying and pasting:

5. Copy the **Probe-back** entry and paste it into the **Derived Parts** node.
6. Rename the **Derived Parts** > **Copy of Probe-back** node to `Probe-bottom`.
7. In the *Properties* window, set *Point* to [`0,-0.0101,1.0E-4`] m, m, m.
8. Copy the **Probe-back** node for a second time and paste it into the **Derived Parts** node.
9. Rename the **Derived Parts** > **Copy of Probe-back** node to `Probe-top`.
10. In the *Properties* window, set *Point* to [`0,0.0101,1.0E-4`] m, m, m.

Display the probes in the *Graphics* window:

11. Multi-select the **Probe-back**, **Probe-bottom**, and **Probe-top** nodes.
12. Click **(Make Scene Transparent)** and, using the mouse buttons, zoom and position the scene as shown below.



Create three reports, one for each probe:

13. Right-click the **Reports** node and select **New Report > Maximum**.
14. Rename the **Maximum 1** node to `Max Pressure-back`.
15. Select the **Max Pressure-back** node and set the following properties:

Property	Value
Scalar Field Function	**Pressure**
Parts	**Derived Parts > Probe-back**

16. Copy the **Max Pressure-back** node and paste it into the **Reports** node.
17. Rename the **Copy of Max Pressure-back**node to `Max Pressure-bottom`.
18. Select the **Max Pressure-bottom** node and set *Parts* to **Derived Parts > Probe-bottom**.
19. Create another copy of **Max Pressure-back.**
20. Rename the new node to `Max Pressure-top`.
21. Select the **Max Pressure-top** node and set *Parts* to **Derived Parts > Probe-top**.

Create a monitor and plot from the reports:

22. Multi-select the **Max Pressure-back**, **Max Pressure-bottom**, and **Max Pressure-top** nodes. Right-click the **Max Pressure-back** node and select **Create Monitor and Plot from Report**
23. Click **Single Plot** to display all three values in a single graph.

STAR-CCM+ creates three new monitors and the associated plot.

24. Rename the **Plots > Reports Plot** node to `Pressure Probes Plot`.

Change the auto save criteria to be based on time-step rather than iterations. Also specify the simulation to be saved every six hundred time steps.

25. Select **File > Auto Save...** .
26. In the *Auto Save* dialog, set *Max Autosaved Files* to `20`.
27. Under *Update*, activate the **Enabled** checkbox and select `Time Step` for the *Trigger* property.
28. Under *Time-Step Frequency*, set the *Time-step Frequency* to `600`.



29. Click **Close**.
30. Save the simulation.

Monitoring Force Coefficients

By default STAR-CCM+ provides a residuals plot which can be used to monitor the convergence of the solution. In addition to this, it is useful to set up additional reports and monitors to help in assessing convergence.

Two monitor reports were created in the Broadband Models tutorial, for Drag Coefficient and Life Coefficient. Adapt these and the Force Coefficient Plot:

1. Select **Monitors** > **Drag Coefficient Monitor** and set the *Trigger* property to **Time Step**.
2. Select **Monitors** > **Lift Coefficient Monitor** and set the *Trigger* property to **Time Step**.
3. Select **Plots** > **Force Coefficient Plot** and set the *X-Axis Monitor* property to **Physical Time**.
4. Change the *Title* property to `Force Coefficients Plot`.
5. Save the simulation.

Saving Transient Surface Data

Assuming that you have sufficient disk space (8.4 Gb), set up STAR-CCM+ to export volume or surface data at regular intervals throughout the transient simulation.

In this tutorial, you export pressure data on selected surfaces after each time-step, so that you can perform surface spectral analysis when the simulation completes.

1. Select **File** > **Auto Export...**

   The *Auto Export* dialog appears.
2. Set the following values:

Nodes	Values
File Format	**Surface fft files (*.trn)**
Base Name	**cylinderPressureData**
Append to File	Activated
Export Solution Data Only	Activated
Scalar Functions	**Pressure**
Boundaries	**Body 1: Cylinderwall**
Derived Parts	**plane section**
Update	
▮▮*Enabled*	Activated
▮▮*Trigger*	**Time Step**
Time-Step Frequency	
▮▮*Time-step Frequency*	`1`
▮▮*Start Time-step*	`50`

The *Auto Export* dialog looks as shown below.



3. Click **Close**.
4. Save the simulation.

Note: The Auto Export function can also be used to export data for use in third-party aeroacoustics propagation codes. When this is required, choose the **CD-adapco CCM Files (*.ccm)** option for the *File Format*, as some third-party software (Actran, VA One, and Virtual.Lab.Acoustics) are able to read the `.ccm` format directly. Follow the advice from these packages as to the data required. For volume data, select appropriate regions using the *Regions* property; for surface data, select appropriate boundaries or derived parts in the manner that is shown above.

Preparing the On-The-Fly FW-H Analysis

The two key steps in preparing STAR-CCM+ for a Ffowcs Williams-Hawkings analysis are to define the FW-H Surfaces and then define the FW-H Receivers. FW-H Receivers are points at specific distances from the noise sources.

To choose the FW-H models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models, in order:

Group Box	Model
Optional Models	**Aeroacoustics**
Aeroacoustics Models	**Ffowcs Williams-Hawkings Unsteady**
FW-H Unsteady Models	**On-The-Fly FW-H**



2. Select the **FW-H Receivers** node and set the following properties:

Property	Setting
Far-Field Density	`1.18 kg/m^3`
Far-Field Speed of Sound	`347.2 m/s`

Contents

Setting the FW-H Surfaces
Setting the FW-H Receivers
Checking FW-H Solver Properties

Setting the FW-H Surfaces

Two additional nodes were added to the simulation tree when the Ffowcs Williams-Hawkings model was selected. Set the FW-H Surfaces.

1. Right-click the **FW-H Surfaces** node and select **New > Impermeable Surface**.



2. Select the **Impermeable Surface 1** node and set *Boundaries* to **Body 1:Cylinderwall**.

Setting the FW-H Receivers

Set the point-receivers.

1. Right-click the **FW-H Receivers** node and select **New Point Receiver**.



2. Select **Point Receiver** and set *FW-H Surface* to **Impermeable Surface 1**.

3. Set *Position* to `[0.0, 1.01, 0.0] m,m,m`.
4. Rename the node to `OTF_FWH_Plus90deg100r`.
5. Similarly, create a second point receiver node.
6. Select the newly created node and set *FW-H Surface* to **Impermeable Surface 1**.
7. Set *Position* to `[0.0, -1.01, 0.0] m,m,m`.
8. Rename the node to `OTF_FWH_Minus90deg100r`.



Checking FW-H Solver Properties

The **FW-H Solver** node has a *Start Time* property that defines the physical time when point-receivers start computing the sound pressure.

To set the start time:

- Select the **FW-H Unsteady Solver** node and set *Start Time* to `0.015 s`.

This setting is based on observed statistically steady flow of the 3D simulation, in which the mean level remains constant.



Visualizing the Solution

Add scenes for the noise emitted from the cylinder surface and received at each receiver.

The noise has a loading component for the unsteady motion of the force distribution. The body is stationary which means the thickness component is zero. The total surface noise is simply the summation of the two components.

To create the first total noise contour scene:

1. Create a scalar scene.
2. Rename the Scene to `TotalSurfaceTerm_OTF_FWH_Plus90deg100r`.
3. Select the **Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Scalar Field	*Function*	**FW-H Noise Source > OTF_FWH_Plus90deg100r > Total Surface Term**

Node	Property	Setting
Parts	*Parts*	**Regions > Body 1 > Boundaries > Cylinderwall**
Outline 1 > Parts	*Parts*	**Regions > Body 1 > Boundaries > Cylinderwall**

4. In the main toolbar, click **Reset View**.

Only the cylinder is shown.

Add a second scalar scene for the other receiver similar to the first.

5. Copy the **Scenes > TotalSurfaceTerm_OTF_FWH_Plus90deg100r** node and paste it to the **Scenes** node.
6. Rename the **Copy of TotalSurfaceTerm_OTF_FWH_Plus90deg100r** to `TotalSurfaceTerm_OTF_FWH_Minus90deg100r`.
7. Select the **Scalar 1 > Scalar Field** node and set *Function* to **FW-H Noise Source > OTF_FWH_Minus90deg100r > Total Surface Term**.
8. Save the simulation.

Running the Simulation

The simulation is now ready to be run. It is best that you run this simulation in parallel using several CPUs, as otherwise the run time is considerable.

1. Run the simulation by clicking the (**Run**) button in the **Solution** toolbar.

The solution is saved every six hundred time-steps. Save the final solution file (cylinderUnsteady@04800.sim) under the original file name.

2. Select **File > Save As**
3. In the *Save* dialog, select cylinderUnsteady.sim for the File Name.
4. Click **Save**.
5. Click **OK** to overwrite the previous file.

Contents

Force Coefficients Time Plot
Maximum Pressure Time Plot
Mean Velocity Magnitude

Force Coefficients Time Plot

Modify the axis range of the force coefficients plot.

Use the following steps:

1. Right-click the **Plots** > **Force Coefficients Plot** node and select **Open**.
2. Select the **Plots** > **Force Coefficients Plot** > **Axes** > **X Axis** node and set *Minimum* to `0.0` and *Maximum* to `0.12`.
3. Select the **X Axis** > **Major Labels** node and set *Label Spacing* to `0.02`.
4. Select the **X Axis** > **Title** node and set *Title* to `Time (s)`.
5. Select the **Y Axis** > **Title** node and set *Title* to `Force Coefficient`.

A scene similar to the one shown below is displayed.

As expected, the mean lift is around zero and the mean drag is around 1 [452].

Maximum Pressure Time Plot

Create a plot of maximum pressure versus time.

1. Right-click the **Plots** > **Pressure Probes Plot** node and select **Open**.
2. Open the **Plots** > **Pressure Probes Plot** > **Axes** > **X Axis** node and set *Minimum* to `0.0` and *Maximum* to `0.12`.
3. Select the **X Axis** > **Major Labels** and set *Label Spacing* to `0.02`.
4. Select the **X Axis** > **Title** node and set *Title* to `Time (s)`.
5. Select the **Y Axis** > **Title** node and set *Title* to `Maximum Pressure (Pa)`.

The Graphics window contains a scene similar to the following one.



The mean pressure level remains constant while the top and bottom pressure probes identify anti-symmetric fluctuations.

Mean Velocity Magnitude

View the mean velocity magnitude in a scalar scene.

Examine the *Mean Velocity Magnitude* scene.



The mean velocity magnitude is symmetric around the cylinder. The difference in the wake is due to the refinement level.

Visualizing the Results

You can now visualize the acoustic pressure on the cylinder, based on input from the two point-receivers.

1. Open the node **Scenes** > **TotalSurfaceTerm_OTF_FWH_Plus90deg100r**. A scene similar to the one shown below is displayed.



2. Open the node **Scenes** > **TotalSurfaceTerm_OTF_FWH_Minus90deg100r**. A scene similar to the one shown below is displayed.



Spectral Analysis

Using Fast Fourier Transforms, you can determine sound pressure as a function of frequency at each of the point receivers.

Set up the spectral analysis.

1. Open the **Tools** node, right-click **Data Set Functions** and select **New > Point Time Fourier Transform (G(p))**.



A new node, **G(p) 1**, is created in the simulation tree.

2. Set the following properties:

Property	**Value**
Start Time	`0.02 s`
Cut-off Time	`0.12 s`
Amplitude Function	**Sound Pressure Level**
Analysis Blocks	`3`
Overlap Factor	`0.5`
Window Function	**Hann**

3. Right-click the **G(p) 1 > Monitor** node and select **New derived data from monitor**.
4. Rename the **Monitor** node to `derived monitor-back`.
5. In the *Properties* window, set *Input data 1* to **Max Pressure-back Monitor**.



6. Right-click the **G(p) 1 > Monitor** node and select **New derived data from monitor** for a second time.
7. Rename the **Monitor** entry to `derived monitor-bottom`.
8. In the *Properties* window, set *Input data 1* to **Max Pressure-bottom Monitor**.
9. Right-click the **G(p) 1 > Monitor** node and select **New derived data from monitor** for a third time.
10. Rename the **Monitor** entry to `derived monitor-top`.
11. In the *Properties* window, set *Input data 1* to **Max Pressure-top Monitor**.

    This portion of the object tree is shown in the following screenshot:



12. For each of the derived monitor nodes, activate the *Update Active* property.
13. Rename the **G(p) 1** node to `FFT-SPL`.

You now add the derived monitor nodes to a monitor plot.

14. Right-click the **Plots** node and select **New Plot > Monitor**.
15. Rename the **Monitor Plot 1 node** to `Monitor Plot - SPL`.
16. Select the **Plots > Monitor Plot - SPL > Data Series** node, right-click, and select **Add Data**.

    This opens the *Add Data Providers to Plot* dialog.
17. Expand the Derived Data node and select the following:
    - derived monitor-back
    - derived monitor-bottom
    - derived monitor-top



18. Select the **Plots > Monitor Plot - SPL > Axes > X Axis** node.
19. In the *Properties* window, activate the *Logarithmic* property. Set the following properties:

Property	Value
Logarithmic	Activated
Minimum	`10`
Maximum	`20000`

20. Select the **X Axis > Title** node and set *Title* to `Frequency (Hz)`.

21. Select the **Y Axis** > **Title** node and set *Title* to `Sound Pressure Level (dB)`.

A scene similar to the one shown below is displayed.



From the above monitor plots, it is seen that there are two peaks: one at 500 Hz for the top and bottom probes (lift mode) and one at 1000 Hz for the back probe (drag mode).

22. Save the simulation.

Results Analysis

Compare the peaks in the spectra with the Mesh Frequency Cutoff plot from the final steady-state simulation.

Use the following steps:

1. Open the simulation file for the steady-state calculation after 400 iterations (that is, the file `cylinder.sim` containing the refined mesh) alongside the existing `cylinderUnsteady` simulation.
2. In the same *Simulation* window, open the **Scenes** node.

3. Right-click **Mesh Frequency Cutoff** and select **Open.**
4. Select the **Mesh Frequency Cutoff > Displayers > Scalar 1 > Scalar Field** node and set *Max* to `10000`.



Use the mouse controls to obtain a plot similar to the following:



Compare the predicted spectra with the Mesh Frequency Cutoff plot from the steady-state simulation. You can see that there is a close correlation between the steady-state estimated cut-off frequency and the high-end frequency resolution from the DES calculation. All spectra show a change of gradient at ~6000 Hz. The change in gradient represents a change in the energy cascade to the higher frequencies, resulting artificially from inadequate resolution of the turbulent structures. These gradient kinks correspond with the fine-mesh steady-

state Mesh Frequency Cutoff values local to those points. The spectral peak at ~500 Hz is the predicted vortex shedding frequency.

Many researchers (for example, Strouhal, Oertel, and Williamson, among others) have studied the sound that is generated by the flow around a circular cylinder.

Strouhal found that the frequency f of the sound that is radiated from a cylinder of diameter D is related to the free-stream velocity U_∞ of a uniform flow. This relation is known as the Strouhal number $St = fD / U_\infty$. The Strouhal number is 0.2 to 0.22 for the Reynolds number range $300 \leq \mathrm{Re} \leq 10^4$(Inoue). For this case of a cylinder 20 mm in diameter and a free-stream air velocity of 50 m/s, the Reynolds number is 6.38×10^4 and the Strouhal number is 0.2. The frequency f of the sound is 500 Hz. This predicted frequency is identical to the shedding frequency of vortices behind the cylinder.

The SPL vs. frequency plot can numerically predict the frequency peak of 500 Hz generated by vortex shedding from the cylinder surface into its wake. This suggests that DES predicts the vortex shedding well and that the Strouhal number is close to the experimental value.

5. Save and close the simulation.

Summary

This tutorial continued from the previous one, doing an unsteady analysis of the near field and an FW-H analysis of the far field.

The tutorial introduced the following STAR-CCM+ features:

- Setting up a transient calculation.
- Modifying solver and stopping criteria parameters.
- Setting up monitoring report and plot.
- Saving transient surface data.
- Setting up spectral analysis.

DES and FW-H On-The-Fly Tutorial Bibliography

Strouhal, V. 1878, “Ueber eine besondere art der tonerregung.” *Annu. Phys. Chem. (Wied. Annu. Phys.)* Vol. 5, pp. 216-251.

Oertel, Jr. H., 1990, “Wakes behind blunt bodies” *Annu. Rev. Fluid Mech.* Vol 22, pp. 539-564.

Williamson, C. H., K., 1996, "Vortex dynamics in the cylinder wake" *Annu. Rev. Fluid Mech.* Vol 28, pp. 477-539.

Inoue O., Hatakeyama, N., 2002 "Sound generation by a two-dimensional circular cylinder in a uniform flow" *J. Fluid Mech.* Vol. 471, pp. 285-314.

Travin, A., Shur, M., 2000 "Detached-Eddy Simulations Past a Circular Cylinder" *Combustion*, Vol 63(1), pp. 293-313

Ffowcs Williams-Hawkings: Sound Propagation

This tutorial compares the Post FW-H model with the FW-H On-The-Fly model.

In the previous tutorial, DES and FW-H On-The-Fly: Noise from a Cylinder (Unsteady Analysis), you completed the DES and FW-H On-the-Fly transient runs. In this tutorial, you first import the transient surface data file that was generated during the transient runs. You then use the Post FW-H option of the FW-H model to calculate sound pressure levels at locations corresponding to the far field monitor points for which transient data was collected before. By comparing the Post FW-H model predictions with the results of the FW-H On-The-Fly analysis, you can verify that the results obtained from the `.trn` file with the Post FW-H model are identical to those obtained while actually running the unsteady simulation.

Contents

Importing the Transient Surface Data
Preparing for Analysis with the Post FW-H Model
Plotting Sound Pressure Levels from FW-H Data
Summary
Using the Transient Surface Data Tutorial Bibliography

Importing the Transient Surface Data

If the transient surface data file has been saved as instructed previously, it can be loaded in the last saved simulation, cylinderUnsteady@04800.sim. This section describes the steps for doing so.

Use the following steps:

1. Start up a double-precision version of STAR-CCM+ in a manner that is appropriate to your working environment and load cylinderUnsteady@04800.sim.
2. Select **File > Import > Import CAE Model.**
3. Select `CylinderPressureData.trn` and click **Open**.
4. In the *Import Options* dialog, deactivate **Open geometry scene after import** and click **OK**.

This action adds the **Imported Models** node to the object tree.



5. Expand the **Imported Models** node and select the **trn:cylinderPressureData** node.



Note the properties. These are used to examine the pressure data at different points in time.

6. Save the simulation as 3Dcylinder_FWH.sim.

Preparing for Analysis with the Post FW-H Model

In this section, you select the Post FW-H model, set up two Post FW-H receivers, then run the Post FW-H solver. Receivers are points at specific distances from the noise sources where sound pressure is calculated. The receivers take their data from the imported CylinderPressureData.trn file.

Contents

Choosing the Post FW-H Model
Creating Two Derived Point Parts
Setting the Post FW-H Receivers
Running the Post FW-H Solver

Choosing the Post FW-H Model

Add the Post FW-H model to the physics models.

To select the Post FW-H model:

1. Right-click the **Continua** > **Physics 1** > **Models** node and choose **Select Models.**

In the *Physics 1 Model Selection* dialog:

2. Deselect the **On-The-Fly FW-H** model.
3. Select **Post FW-H** from the *FW-H Unsteady Models* box.



4. Click **Close**.
5. Select the **Post FW-H Receivers** node and set the following properties:

Property	Setting
Far-Field Density	`1.18 kg/m^3`
Far-Field Sound Speed	`347.2 m/s`

Creating Two Derived Point Parts

Create two point parts at the positions of the point receivers.

Use the following steps:

1. Right-click the **Scenes > Geometry Scene 1** node and select **Open**.
2. In the *Simulation* window, right-click the **Derived Parts** node and select **New Part > Probe > Point**.
3. In the *Create Point* panel, do the following:

   In the *Input Parts* box, select **Body 1** and **Imported models > trn:cylinderressureData > Imported Surfaces > Body 1-Cylinder wall**.

   Set the *Point* to `[0, 1.01, 0]` m, m, m.

   Under *Display*, select **No Displayer**.



Click **Create** and then **Close**.

4. Rename the **Derived Parts** > **point** node to `Plus90deg100r`.

Create another probe by copying and pasting the first one:

5. Right-click the **Plus90deg100r** entry and select **Copy**.
6. Right-click the **Derived Parts** node and select **Paste**.
7. Rename the **Derived Parts** > **Copy of Plus90deg100r** node to `Minus90deg100r`.
8. In the *Properties* window, set *Point* to [`0,-1.01,0`] m, m, m.

Setting the Post FW-H Receivers

Two point-receivers are created for this analysis, both at a distance of 1.01 m from the cylinder surface. These point receivers take the point probes created previously as an input.

Set the point-receivers:

1. Right-click the **Post FW-H Receivers** node and choose **New** to create a new Point Receiver.



2. Select the **Post Point Receiver 1** node and set *Imported FW-H Surfaces* to **Body 1-Cylinderwall**.
3. For the same node, set *Point Part* to **Minus90deg100r**.
4. Rename the node to `Post_FWH_Minus90deg100r`.



5. Using copy and paste, create a second point receiver node from the first one.
6. For this second node, set *Point Part* to **Plus90deg100r**.
7. Rename the node to `Post_FWH_Plus90deg100r`.

Running the Post FW-H Solver

You are now ready to run the Post FW-H solver.

- Right-click on the **Post FW-H Solver** node and select **Execute Solver**.



Plotting Sound Pressure Levels from FW-H Data

While generating the .trn file for the DES and FW-H On-The-Fly tutorial, STAR-CCM+ collected sound pressure data from two point receivers for the FW-H On-The-Fly model. In the following sections, you prepare plots of the sound pressure levels using data from both the On-The-Fly FW-H model and the Post FW-H model.

By plotting both sets of data on the same plot, you can verify that the Post FW-H model gives identical results to the On-The-Fly FW-H model. To use this data, the following steps are necessary:

1. Export the data from the point-receivers.
2. Import this data as a table.
3. Create a Fast Fourier Transform function that performs a spectral analysis using the sound pressure level data from the table created.
4. Plot this data with appropriate scales.

Contents

Exporting Sound Pressure Data from the On-The-Fly FW-H Receivers
Exporting Sound Pressure Data from the Post FW-H Receivers

Importing Sound Pressure Data
Applying Fast Fourier Transform to Table Data
Creating a Plot of Sound Pressure Level

Exporting Sound Pressure Data from the On-The-Fly FW-H Receivers

Export the sound pressure data to external files.

The sound pressure data from the unsteady simulation is contained within the FW-H point-receivers.

Use the following steps:

1. Right-click the **FW-H Receivers > OTF_FWH_Minus90deg100r** node and choose **Export Surface Total**.



   The *Save* dialog appears.
2. For **File Name**, enter `OTF_FWH_Minus90deg100r`. Click **Save**.
3. Similarly, right-click the **FW-H Receivers > OTF_FWH_Plus90deg100r** node, choose **Export Surface Total**, and set **File Name** to `OTF_FWH_Plus90deg100r`. Click **Save**.

Two files are created in the working directory.

Exporting Sound Pressure Data from the Post FW-H Receivers

Export the sound pressure data to external files.

The sound pressure data generated by the Post FW-H model is contained within the Post FW-H point receivers.

1. Right-click the **Post FW-H Receivers > Post_FWH_Minus90deg100r** node and choose **Export...**



   The *Save* dialog appears.

2. Enter `Post_FWH_Minus90deg100r` in the **File Name** entry box of the dialog and click **Save** to export the data.
3. Similarly, right-click the **Post FW-H Receivers > Post_FWH_Plus90deg100r** node, choose **Export...**, and enter `Post_FWH_Plus90deg100r` in the **File Name** entry box of the *Save* dialog. Click **Save** to export the data.

Two files are created in the working directory.

Importing Sound Pressure Data

Import the sound pressure data as tables.

To import the sound pressure data from the four files:

1. Right-click the **Tools > Tables** node and select **New Table > File Table**.
2. In the *Open* dialog, select the **OTF_FWH_Minus90deg100r.csv** file and click **Open**.



3. Repeat these steps for the following data files:
   - **OTF_FWH_Plus90deg100r.csv**
   - **Post_FWH_Minus90deg100r.csv**
   - **Post_FWH_Plus90deg100r.csv**

Four additional nodes are created beneath the **Tables** manager.



Applying Fast Fourier Transform to Table Data

Use a fast Fourier transform data set to analyze the sound pressure level at different frequencies.

Use the following steps:

1. Right-click the **Tools > Data Set Functions** node and choose **New > Point Time Fourier Transform (G(p))**.



   A node, **G(p) 1**, is added to the **Data Set Functions** manager node.

2. Right-click the **G(p) 1 > Tabular** node and select **New derived data from table**.
3. Rename the node to `OTF_FWH_Minus90deg100r`.
4. For this node, set the following properties:

Property	Setting
Input data	**OTF_FWH_Minus90deg100r**
X Column 1	**Time**
Y Column 1	**Point Receiver (0 -1.01 0)**



5. Repeat this process to create three more Tabular nodes, for a total of four. The four nodes have identical properties, except for names and *Input data*:
6. Select the **G(p) 1** node and set the following properties:

Property	**Setting**
Update Interval	1
Start Time	0.02 s
Cut-off Time	0.12 s
Amplitude Function	**Sound Pressure Level**
Analysis Blocks	3
Overlap Factor	0.5
Window Function	**Hann**

7. Select the *Reference Sound Pressure* node below the **G(p) 1** node and set *Value* to `2.0E-5 Pa`.
8. Select each of the four nodes under the **Tabular** node and activate *Update Active* in each one.

Creating a Plot of Sound Pressure Level

The spectral analysis plot takes its data from the fast Fourier transform created previously.

Use the following steps:

1. Right-click the **Plots** node and choose **New Plot** > **Monitor**.



2. Rename the **Monitor Plot 1** node to `FFT`.
3. Open the **FFT** node, right-click on the **Data Series** node, and select Add Data.

   The *Add Data Provides to Play* dialog opens.
4. In this dialog, select and expand Derived Data, then select the following items:
   - OTF_FWH_Minus90deg100r
   - OTF_FWH_Plus90deg100r
   - Post_FWH_Minus90deg100r
   - Post_FWH_Plus90deg100r



Change the X-axis:

5. Select the **FFT** > **Axes** > **X-Axis** node and activate *Logarithmic*.
6. Set *Minimum* to `100` and *Maximum* to `10000`.



7. Select the **FFT** node and set the *Title* property to `On-The-Fly FW-H and Post FW-H`.

The final plot appears similar to the one shown below.



The Post FW-H model allows you to obtain noise prediction at receiver locations not defined in the original unsteady simulation. As the computational cost of running a Post FW-H analysis is much lower than the cost of running an unsteady analysis, you can continue to explore the solution space for little extra effort.

Summary

This tutorial compared the results of the previous tutorial with those of the Post FW-H solver.

This tutorial introduced the following STAR-CCM+ features:

- Post-processing transient surface data
- Applying fast Fourier transforms
- Plotting sound pressure levels

Using the Transient Surface Data Tutorial Bibliography

Strouhal, V. 1878, "Ü ber eine besondere art der tonerregung." *Annu. Phys. Chem. (Wied. Annu. Phys.)* Vol. 5, pp. 216-251.

Oertel, Jr. H., 1990, "Wakes behind blunt bodies" *Annu. Rev. Fluid Mech.* Vol 22, pp. 539-564.

Williamson, C. H., K., 1996, "Vortex dynamics in the cylinder wake" *Annu. Rev. Fluid Mech.* Vol 28, pp. 477-539.

Inoue O., Hatakeyama, N., 2002 "Sound generation by a two-dimensional circular cylinder in a uniform flow" *J. Fluid Mech.* Vol. 471, pp. 285-314.

Travin, A., Shur, M., 2000 "Detached-Eddy Simulations Past a Circular Cylinder" *Combustion*, Vol 63(1), pp. 293-313

Signal Post-Processing: FFT and Wavenumber

Several further analyses are possible using the transient surface data file from the earlier tutorial, DES and FW-H On-The-Fly: Noise from a Cylinder (Unsteady Analysis). These analyses are covered in this tutorial.

Contents

Importing the Transient Surface Data
Plotting Pressure at Different Time-Steps
Plotting the Sound Pressure Level at Different Frequencies
Determining the Group Velocity of the Disturbances
Summary

Importing the Transient Surface Data

Create a new simulation, based on the transient surface data.

Use the following steps:

1. Start up a double-precision version of STAR-CCM+ in a manner that is appropriate to your working environment.
2. Select the New Simulation option from the menu bar.
3. Select **File > Import > Import CAE Model**.
4. Select `cylinderPressureData.trn` and click **Open**.
5. In the *Import Options* dialog, make sure that **Open geometry scene after import** is deactivated and click **OK**.

   This action adds the **Imported Models** node to the object tree.



6. Expand the **Imported Models** node and select the **trn:cylinderPressureData** node.



   Note the properties. These are used to examine the pressure data at different points in time.

7. Save the simulation as `cylinderPressureData.sim`.

Plotting Pressure at Different Time-Steps

A scalar scene showing pressure is created; the **trn:cylinderPressureData** node is used to change the time for which the data is shown.

1. Create a scalar scene.
2. Select the **Scalar Scene 1 > Displayers > Scalar 1** node and set *Contour Style* to **Smooth Filled**.
3. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Parts** node and click to the right of the *Parts* property. In the object selection dialog that opens, select all the nodes under the **Imported Models > trn:cylinderPressureData** node.
4. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node.
5. Click to the right of the *Function* property and choose **Time History:Pressure**.

The scalar scene is updated to show the pressure field at the point in time specified by the *Transient Display Time* property of the **Imported**

Models > trn:cylinderPressureData node. In this instance, the initial value is set to 0.12 s.



6. Select the **Imported Models > trn:cylinderPressureData** node and click to the right of *Transient Display Time*. Use the slider bar (or its accompanying edit box) to set the value to `0.06`.



The scalar scene updates to show the pressure field at 0.06 s.

Plotting the Sound Pressure Level at Different Frequencies

The entire transient surface data file can be used in a Fast Fourier Transform to plot the sound pressure level across the stored surfaces at selected frequencies.

1. Expand the **Tools > Field Functions** node and select the **Fft:Pressure** node.
2. In the *Properties* window, set the the following properties:

Property	Setting
Amplitude Function	**Sound Pressure Level**
Analysis Blocks	`3`
Overlap Factor	`0.5`
Fft Input Start Time	`0.02 s`
Fft Input Cut-off Time	`0.12 s`

3. Select the **Fft:Pressure > Window Function** node and set the *Window Type* to **Hann**.



4. Select the **Fft:Pressure > Frequency Value** node and set the Value to `500 Hz`.

Plot the sound pressure level using *Scalar Scene 1*.

5. Make sure *Scalar Scene 1* is open and visible in the *Graphics* window.
6. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node and set the *Function* to **Fft:Pressure**.

   STAR-CCM+ calculates the Fast Fourier Transform for the specified frequency. When the calculation completes, the scene updates.
7. Manipulate the scene using the mouse controls to obtain something similar to the following:



The sound pressure level at different frequencies can be viewed by modifying the **Frequency Value** of the **Fft:Pressure** field function. The SPL (dB) values at the top and bottom of the cylinder exactly match the point spectra dB levels (plotted previously) at 500 Hz.

8. Close the scene.

Determining the Group Velocity of the Disturbances

A wavenumber-frequency diagram (or fk plot) is used to determine the group velocity at which disturbances propagate along the domain.

To determine the group velocity for this simulation, first compute cross spectral densities or auto spectral densities using the data from the `.trn` file. Determine the range for which the data is meaningful for analysis and use this range to compute a line spatial Fourier transform. By plotting the spectra from this Fourier transform, you can compute the group velocity of the disturbances. You require an `aero-vibro` feature license to complete this analysis.

This section describes the steps for doing so:

1. Creating a line probe along which transient data is extracted.
2. Creating a line time history to record data from the line probe.
3. Calculating the spectral density along the line probe.

4. Calculating the line spatial Fourier transform and using it to determine the group velocity.

Contents

Creating a Line Probe
Calculating the Spectral Density
Calculating the Line Spatial Fourier Transform

Creating a Line Probe

Create a line probe to sample the pressure time history.

1. Create a geometry scene.
2. Right-click the **Derived Parts** node and select **New Part > Probe > Line...**.

   The *Create Line Probe* window appears.
3. Click the **Select** button under *Input Parts* and activate **Imported Models > plane section** in the dialog box.
4. Click **OK**.
5. In the *Create Line Probe* window, set the following properties:

Property	Setting
Point 1	[`0.0229, 0.001, 0.02`] m, m, m
Point 2	[`0.2229, 0.001, 0.02`] m, m, m
Resolution	`100`
Display	**No Displayer**

The *Create Line Probe* window appears as shown.



6. Click **Create**, and then click **Close**.

Calculating the Spectral Density

To calculate the spectral density, you must create a line time history, a line time Fourier transform, and a derived data set. You can then visualize the spectral density.

To create a line time history to record data from the line probe:

1. Right-click the **Tools > Data Set Functions** node and select the **New > Line Time History (h[l])** node.
2. Right-click the **h(l) 1 > Imported** node and select **New derived data from imported surface and derived part**. The node **Multi-Point Imported** opens below.
3. Click on **Multi-Point Imported** and set the following properties:

Property	Setting
Field Function 1	**Time History:Pressure**
Data Surfaces 1	**plane section**
Part	**line-probe**
Update Active	Activated



To create a line time Fourier transform:

4. Right-click the **Tools > Data Set Functions** node and select **New > Line Time Fourier Transform (G[h[l]])**.
5. Rename the newly created **G(l) 1** node to `Gxx(l)`.
6. Select the **Tools > Data Set Functions > Gxx(l)** node and make sure the *Amplitude Function* is set to **Power Spectral Density**.

To create the derived data set node into which you import the data:

7. Right-click the **Gxx(l) > Imported** node and select **New derived data from derived data**.
8. Rename the newly created **Multi-Point Imported 2** node to `Power Spectra`.



9. Select the **Gxx(l) > Imported > Power Spectra** node.
10. Set *Input time history* to **Multi-Point Imported**.
11. Activate *Update Active*. All other properties are disabled for editing and updated automatically after you activate *Update Active*.



To visualize the spectral density:

With the *Update Active* property activated, the Fourier transform is automatically calculated once you assign the last property. This process takes a few minutes, depending on the processor speed of your computer.

When the Fourier transform has been computed, visualize the spectral density as a function of distance. Use a logarithmic scale for the spectral density axis.

12. Select the **Gxx(l) > Imported > Power Spectra > Output control** node.



13. Create a monitor plot and rename it to `Power Spectra`.
14. Select the **Plots** > **Power Spectra** > **Data Series** node, right-click, and select Add Data.

    The *Add Data Providers to Plot* dialog opens.
15. In the dialog, expand Derived Data and select Power Spectra.



The plot of Power Spectral Density against Distance appears in the *Graphics* window.



The expected profile is obtained. The average correlation is maximum at zero where the separation is minimum.

Calculating the Line Spatial Fourier Transform

Use the values in this range to create a line spatial Fourier transform. The velocity of propagation is calculated from this transform.

1. Right-click the **Tools > Data Set Functions**node and select **New > Line Spatial Fourier Transform (F[G[h[l]]])**.
2. Rename the newly created **F(G(l)) 1** node to `Fxx(l)`.
3. Select the **Data Set Functions > Fxx(l)** node and set *Start Signal* to `0.0 m`.
4. Set *End Signal* to `0.2 m`.



5. Right-click **Fxx(l) > Derived** and select **New derived data from derived data**.
6. Rename the newly created **Line Derived** node to `PSD`.
7. Select **Fxx(l) > Derived > PSD** and set *Input Derived Data* to **[Power Spectra]**.



With the *Update Active* property activated, the line spatial Fourier transform is automatically calculated. This process takes a few minutes, depending on the processor speed of your computer.

When the Fourier transform has been computed, visualize the spectra:

8. Select **PSD** > **Output control** and activate *Log data*.

Create a new scalar scene and rename it to `PSD`.

9. Select **PSD** > **Displayers** > **Outline 1** and deactivate *Outline*.
10. Select **PSD** > **Displayers** > **Scalar 1** > **Parts** and select **Derived Data** > **PSD**.
11. Click **OK**.

The power spectral density plot appears in the *Graphics* window. Use the mouse to adjust the view.

12. To color the plot, select **PSD** > **Displayers** > **Scalar 1** > **Scalar Field** and set *Function* to **Derived Data** > **PSD**.



13. To visualize the Power Spectral Density as a two-dimensional contour, click (**Save-Restore-Select views**) in the toolbar and select **Projection Mode > Parallel**.
14. To align with the X-Y plane, press the <T> key.



The slope of the distinctive line in the wavenumber-frequency plot is used to calculate the group velocity of the disturbances traveling along the domain. For this case, the group velocity is estimated as

$$2\pi\frac{10{,}000}{1560}\ \mathrm{m/s} = 40.3\ \mathrm{m/s}$$

which is, as expected, roughly 80% of the inlet velocity.

Summary

This tutorial covered post-processing for spectral and wavenumber analysis.

The tutorial introduced the following STAR-CCM+ features:

- Setting up a transient calculation.
- Modifying solver and stopping criteria parameters.
- Setting up monitoring report and plot.
- Saving transient surface data.
- Setting up spectral analysis.
- Post-processing transient surface data.

Acoustic Wave Modeling: Noise from a Cylinder

This tutorial demonstrates how to set up an acoustic wave simulation in STAR-CCM+. The flow around a cylinder is used in the simulation to describe the required steps.

The Acoustic Wave model uses the theory of Acoustic Perturbation Equations (APE) to predict the noise generated by the fluid flow (which is typically incompressible flow). The Acoustic Wave solver is one-way coupled with the flow solver through this APE-based noise source term, and models the acoustic pressure field inside a specified zone of interest. Using the Acoustic Wave model in a hybrid CFD/CAA simulation produces better results (with fewer spurious effects) than a pure compressible flow simulation.

The process for performing an acoustic wave simulation is as follows:

1. Perform a steady turbulent simulation to get an initial flow solution.
2. Perform an unsteady simulation to get a fully developed flow solution.
3. From the unsteady solution, compute the mean flow pressure.
4. Activate the Acoustic Wave model in the unsteady simulation.
5. Specify the noise source (the zone of interest in the flow) and the acoustic damping function to use in the acoustic wave simulation.
6. Run the acoustic wave simulation to obtain the acoustic pressure values throughout the flow region.

You can perform post-processing of the acoustic data to compare the simulation results against other data. STAR-CCM+ provides built-in data set functions, such as Fourier Transforms, that you can use to apply signal processing to the acoustic data.

The geometry that is used in this tutorial is shown below.



The geometry consists of a cylinder, 20 mm in diameter, placed at the center of a three-dimensional circular computational domain, 0.6 m in diameter. The air velocity and temperature are set to be 50 m/s and 300K, respectively.

Contents

Prerequisites
Loading the Base Simulation
Maintaining the Appropriate Acoustic CFL in the Wave Equation
Generating the Volume Mesh
Obtaining a Steady Solution
Obtaining an Unsteady Solution
Obtaining the Acoustic Wave Solution

Summary

Prerequisites

The instructions in the Acoustic Wave Solver tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Input Files

The input files that you require for this tutorial are in the `doc/startutorialsdata/aeroacoustics/data` subdirectory of your STAR-CCM+ installation directory. You are advised to copy the following input files to your working directory before you start this tutorial:

- `cylinderBase.sim` — contains the problem geometry.
- `CreateUserFieldFunctions.java` — sets up the field functions that you require to specify the acoustic noise source and damping regions.

Loading the Base Simulation

A base simulation file has been prepared for this analysis. Continue by starting the simulation, loading the base simulation file, and saving it under a new file name.

To load the base simulation:

1. Start up a double precision version of STAR-CCM+ in a manner that is appropriate to your working environment.
2. Load the base simulation file:

   a) Click (**Load Simulation**) from the tool bar.

   b) In the *Load Simulation* dialog, under **Process Options**, select **Parallel on Local Host**.

You are strongly advised to run acoustic wave simulations in parallel. This tutorial case has a high number of cells and uses 10 inner iterations, so can be slow to run.

c) In the **Compute Processes** box, set an appropriate number for the machine on which you are working.

To make the run time reasonable, 10 processors are recommended.

d) Navigate to your working directory and open `cylinderBase.sim`.

e) Click **OK**.

3. Save the simulation to disk under file name `acoustic_wave.sim`.

Maintaining the Appropriate Acoustic CFL in the Wave Equation

To maintain the accuracy of acoustic predictions, apply criteria for points per wavelength (PPW) and the Acoustic CFL. This section outlines the requirements and calculations for this case.

The mesh size in the acoustic zone affects two criteria that are important for accurate acoustic predictions: the Acoustic CFL and the points per wavelength (PPW). The appropriate calculations for this tutorial are shown below.

Points Per Wavelength (PPW)

The recommended PPW for an acoustic calculation is PPW ≥ 30-40, maintained throughout the undamped zone. You can calculate the expected wavelength for this simulation from the vortex shedding frequency across the cylinder. This frequency comes from the Strouhal number, St:

$$St = \frac{fD}{v} \tag{9}$$

Assuming a Strouhal number of 0.235 and rearranging the equation gives the expected frequency:

$$f = St\frac{v}{D} = 0.235 \times 50 / 0.02 = 587 \text{ Hz}$$

The wavelength in turn is calculated as c / f where c is the speed of sound (340 m/s). As $\text{PPW} = (c / f) / \Delta x$, the required mesh size for PPW of 40 is:

$$\Delta x = \frac{(340 / 587)}{40} = 0.0145 \text{ m}$$

Acoustic CFL

For the recommended PPW of 40, aim to have Acoustic CFL ≤ 4 in the zone of interest. Acoustic CFL is defined as:

$$\text{Acoustic CFL} = c\frac{\Delta t}{\Delta x} \tag{10}$$

Rearranging and inserting a time-step Δt of 2.5e-5 s gives a candidate mesh size of:

$$\Delta x = 340 \times (2.5 \times 10^{-5}) / 4 = 0.002125\text{m}$$

Overall Mesh Size

Taking the lower computed amount of 0.002125 m and rounding down gives a final target mesh size of 0.002 m in the acoustic zone.

Generating the Volume Mesh

The meshing models and mesh reference values for this simulation have already been set up. A volumetric control has been created for mesh refinement, to capture the separation behind the cylinder, and to capture the vortex shedding.

The mesh size is 0.008 m in the coarse zone and 0.002 m in the fine zone.

To generate the volume mesh:

1. In the **Mesh Generation** toolbar, click (**Generate Volume Mesh**).

To visualize the volume mesh:

2. Right-click the **Scenes** node and select **New Scene** > **Mesh**.

   The mesh appears as shown below.



Obtaining a Steady Solution

Run a steady turbulent simulation to get an initial flow solution. This solution is used as a starting point for the unsteady simulation.

The free-stream air velocity and temperature are set to be 50 m/s and 300K, respectively.

To obtain an initial flow solution, you do the following:

1. Set up the steady physics continuum.
2. Set the appropriate initial conditions.
3. Set the appropriate boundary conditions.

4. Set the appropriate solver parameters and stopping criteria.
5. Run the steady simulation.

Contents

Selecting the Physics Models (Steady)
Setting the Initial Conditions (Steady)
Setting the Boundary Conditions (Steady)
Setting the Solver Parameters and Stopping Criteria (Steady)
Running the Simulation (Steady)

Selecting the Physics Models (Steady)

Set up the physics models for the initial steady simulation.

The result of the steady simulation provides a starting point for the next phase: an unsteady simulation from which you compute the mean flow pressure of the fully developed flow. To ensure good noise source predictions, the physics continuum requires an incompressible flow and a high-resolution turbulence model.

To select the physics models:

1. Rename the physics continuum, **Continua** > **Physics 1**, as **Continua** > **Physics Steady**.
2. For this physics continuum, select the following models in order:

Group Box	Model
Space	**Three Dimensional** (Selected automatically)
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow Gradients** (Selected automatically)
Equation of State	**Constant Density**
Viscous Regime	**Turbulent Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Omega Turbulence SST (Menter) K-Omega** (Selected automatically)

Group Box	Model
	All y+ Wall Treatment (Selected automatically)
Optional Models	**Segregated Fluid Temperature**

3. Click **Close**.
4. To review the models, open the **Physics Steady** > **Models** node.



5. Save the simulation.

Setting the Initial Conditions (Steady)

To enhance convergence, initialize the flow field with a constant velocity of 50 m/s.

To set the initial conditions:

1. Select the **Continua** > **Physics Steady** > **Initial Conditions** > **Velocity** > **Constant** node and set **Value** to `[50.0, 0.0, 0.0]` m/s.

   All other values remain at default.

Setting the Boundary Conditions (Steady)

Set the flow velocity at the outer boundary; all other settings have acceptable default values.

To set the boundary condition:

1. Select the **Regions** > **Body 1** > **Boundaries** > **Freestream** node and rename it to **Outer**.
2. Select this node and set *Type* to **Velocity Inlet**.
3. Select the **Outer** > **Physics Values** > **Velocity Magnitude** > **Constant** node and set **Value** to 50 m/s.

Setting the Solver Parameters and Stopping Criteria (Steady)

Specify the Segregated Flow solver under-relaxation factors and set the appropriate stopping criteria for the steady simulation.

The under-relaxation factors for the Velocity and Pressure solvers are modified to enhance convergence.

To set the solver parameters and stopping criteria for the steady simulation:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Flow		
▮▮**Velocity**	**Under-Relaxation Factor**	`0.9`
▮▮**Pressure**	**Under-Relaxation Factor**	`0.1`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Steps	**Enabled**	Activated
	Maximum Steps	`500`

3. Save the simulation.

Running the Simulation (Steady)

Preparation of the steady simulation is now complete, and you can run the simulation.

To run the steady simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 500 iterations are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

   ```
   Stopping criterion Maximum Steps satisfied.
   ```

2. When the simulation has finished running, save it.

Obtaining an Unsteady Solution

Starting from the solution of the steady simulation, perform an unsteady simulation until the flow reaches a fully developed state.

The Acoustic Wave solver requires the mean of flow pressure in order to compute the instantaneous flow pressure fluctuations that are used in computing the noise source. You pre-compute the mean of flow pressure before you start the Acoustic Wave solver.

To compute the mean flow pressure, perform an unsteady simulation until the flow reaches a fully developed state. Use implicit time marching with sufficient inner iterations to ensure convergence.

To obtain a fully developed flow solution and compute the mean flow pressure:

1. Set up the unsteady physics continuum.
2. Set the appropriate solver parameters and stopping criteria.
3. Create a Mean of Pressure monitor.
4. Create a scalar scene to visualize the mean flow pressure for the fully developed flow.
5. Run the unsteady simulation.

Contents

Selecting the Physics Models (Unsteady)
Setting the Solver Parameters and Stopping Criteria (Unsteady)
Visualizing the Mean Flow Pressure Field
Running the Simulation (Unsteady)

Selecting the Physics Models (Unsteady)

Create a physics continuum for the unsteady simulation and select the appropriate turbulence models.

To select the physics models:

1. Make a copy of the physics continuum, **Continua** > **Physics Steady**, and rename the copy as **Continua** > **Physics Unsteady**.
2. For the physics continuum **Continua** > **Physics Unsteady**, deactivate the following models:

Group Box	Model
Enabled Models	**All y+ Wall Treatment** **SST (Menter) K-Omega** **K-Omega Turbulence** **Reynolds-Averaged Navier-Stokes** **Steady** **Turbulent**

3. Deactivate the **Auto-select recommended models** option.
4. Select the following models in order:

Group Box	Model
Time	**Implicit Unsteady**
Viscous Regime	**Turbulent**
Turbulence	**Detached Eddy Simulation**
Detached Eddy Simulation	**SST (Menter) K-Omega Detached Eddy**
K-Omega Wall Treatment	**All y+ Wall Treatment**

5. Click **Close**.
6. To review the models, open the **Physics Unsteady** > **Models** node.



7. Edit the **Regions** > **Body 1** node and set the following property:

Property	Value
Physics Continuum	**Physics Unsteady**

8. Save the simulation.

Setting the Solver Parameters and Stopping Criteria (Unsteady)

Modify the Segregated Flow solver under-relaxation factors to values appropriate for an unsteady simulation. Also set the necessary stopping criteria.

To ensure convergence, use 10 inner iterations per time step.

To set the solver parameters and stopping criteria for the unsteady solution:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Segregated Flow		
Velocity	*Under-Relaxation Factor*	0.9

Node	Property	Setting
Pressure	*Under-Relaxation Factor*	`0.7`
Implicit Unsteady	*Time-Step*	`2.5e-5 s`
	Temporal Discretization	**2nd-order**

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Enabled*	Activated
	Maximum Inner Iterations	`10`
Maximum Physical Time	*Enabled*	Activated
	Maximum Physical Time	`0.01 s`
Maximum Steps	*Enabled*	Deactivated

3. Save the simulation.

Visualizing the Mean Flow Pressure Field

Create a scalar scene to visualize the mean flow pressure field from the fully developed flow solution.

First create a plane section:

1. Right-click the **Derived Parts** node and select **New Part > Section > Plane**.
2. Edit the **plane section** node and set the following properties:

Property	Setting
Origin	`[0.0, 0.0, 0.02] m, m, m`
Normal	`[0.0, 0.0, 1.0] m, m, m`
Parts	**Regions > Body 1**

Next, create a Field Mean Pressure monitor:

3. Right-click the **Monitors** node and select **New Monitor > Field Mean**.
4. Edit the **Field Mean 1** node and set the following properties:

Property	Setting
Parts	**Regions > Body 1**
Field Function	**Pressure**

Finally, create a scalar scene to display the Mean of Pressure:

5. Right-click the **Scenes** node and select **New Scene > Scalar**.
6. Select the **Scenes > Scalar Scene 1** node and rename it as **Mean of Pressure**.
7. Edit the **Scenes > Mean of Pressure > Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Derived Parts > plane section**
Scalar Field	*Function*	**Mean of Pressure**

8. Save the simulation.

Running the Simulation (Unsteady)

Preparation of the unsteady simulation is now complete, and you can run the simulation.

To run the unsteady simulation:

1. Click (**Run**) in the *Solution* toolbar.

   You can click the *Mean of Pressure* tab at the top of the *Graphics* window to view the development of the mean flow pressure field as the simulation progresses.



The mean flow pressure reference is for all cells in the flow region.

The simulation continues until 400 time-steps (which represents 0.01 seconds of physical time) are complete. The numbers that are displayed in the *Output* window represent the inner iterations of the solver, not time-steps. As there are 10 inner iterations per time-step, you expect the run to perform 4000 inner iterations.

When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Obtaining the Acoustic Wave Solution

Starting from the solution of the unsteady simulation (a fully developed flow solution), activate the Acoustic Wave model, specify the noise source region, and then run the simulation to compute the acoustic pressure.

You use the mean flow pressure that was computed in the unsteady simulation. This value is available as a field function.

This simulation models acoustic waves that are generated by turbulent flow. You specify the noise source with a field function that defines a volume of interest within the flow region. To suppress acoustic reflections from outside the noise source volume, you specify an acoustic damping function. Physical damping is required in this case, since the mesh becomes coarse away from region of interest.

To obtain the acoustic wave solution, you do the following:

1. Activate the Acoustic Wave model.
2. Specify how the Mean Flow Pressure is obtained.
3. Create the appropriate field functions to define the noise source filtering and acoustic damping zones.
4. Set up the noise source to use in the Acoustic Wave model.
5. Create a scalar scene to visualize the acoustic pressure throughout the flow region.
6. Run the simulation with both the flow and the Acoustic Wave models activated.

Contents

Setting up the Acoustic Wave Model
Defining the Noise Source Region and Damping Zone
Setting the Solver Parameters and Stopping Criteria (Acoustic Wave)
Visualizing the Acoustic Pressure
Running the Simulation (Acoustic Wave)
Visualizing the Results

Setting up the Acoustic Wave Model

To model the acoustic waves that are generated by the fluid flow in the physics continuum, you select the Acoustic Wave model and specify the appropriate acoustic perturbation-based noise source.

You specify the following:

- The method by which the acoustic waves are generated in the flow region.

  This simulation models the acoustic waves that are generated by the fluid flow in the physics continuum.

- The method by which the mean pressure field is obtained.

  This tutorial uses the Mean of Pressure field function that was calculated previously from the fully developed unsteady flow.

To set up the Acoustic Wave model:

1. For the physics continuum **Continua** > **Physics Unsteady**, select the following models in order:

Group Box	**Model**
Optional Models	**Aeroacoustics**
Aeroacoustics Models	**Acoustic Wave**

2. Click **Close**.
3. To review the models, open the **Physics Unsteady** > **Models** node.



4. Edit the **Continua** > **Physics Unsteady** > **Models** > **Acoustic Wave** node and set the following properties:

Node	Property	Setting
Acoustic Wave	*Noise Source*	**Acoustic Perturbation Based Source**
	Mean Flow Pressure Reference	**Specified**
Mean Flow Pressure Reference	*Method*	**Field Function**
Field Function	*Scalar Function*	**Mean of Pressure**

5. Save the simulation.

Defining the Noise Source Region and Damping Zone

When you use the Acoustic Wave model, you specify the noise source and the acoustic damping by setting the appropriate physics conditions and physics values of the flow region.

The Acoustic Wave model requires you to define field functions for two coefficients:

- The *Noise Source Weighting Coefficient*, which defines the space within which noise is generated.
- The *Acoustic Damping Coefficient*, which STAR-CCM+ uses to suppress acoustic reflections outside the noise source enclosure space.

In this tutorial, the `NoiseZone` and `DampingZone` field functions define these quantities, respectively. Both of these field functions use a Hanning window for the transition between the values of 0 and 1.

For convenience, in this tutorial, you run a macro to set up the appropriate field functions. The macro, `CreateUserFieldFunctions.java`, is available as an input file in the `doc/startutorialsdata/aeroacoustics/data` subdirectory of the STAR-CCM+ installation directory.

To set up the user field functions:

1. Select **File** > **Macro** > **Play Macro**.
2. In the *Open* dialog, navigate to your working directory and select `CreateUserFieldFunctions.java`.

3. Click **Open** to run the macro.

To specify the noise source region and the acoustic damping function:

4. Edit the **Regions > Body 1 > Physics Conditions** node and set the following properties:

Node	Property	Setting
Acoustic Damping Option	*Acoustic Damping Option*	**Specified**
Noise Source Weighting Option	*Noise Source Region Option*	**Specified**

5. Edit the **Body 1 > Physics Values** node and set the following properties:

Node	Property	Setting
Noise Source Weighting Coefficient	*Method*	**Field Function**
Field Function	*Scalar Function*	**NoiseZone**
Acoustic Damping Coefficient	*Method*	**Field Function**
Field Function	*Scalar Function*	**DampingZone**

6. Save the simulation.

Visualizing the Noise Source Region and Damping Zone

The noise source region, as defined by the `NoiseZone` field function, is shown below:



The noise source is the region (in red) around the cylinder, where the airflow (moving from left to right) generates acoustic waves by vortex shedding from the cylinder. Note the fine mesh throughout the noise source zone.

The acoustic damping zone, as defined by the `DampingZone` field function, is shown below:



The region of interest (in blue) is undamped and everywhere else is fully damped (in red). Note the fine mesh throughout the undamped zone and the damping that is applied at the mesh transitions.

Field Function Definitions

The field functions that are used in this tutorial are tabulated below:

Function Name	Type	Definition
CombinedShape	**Scalar**	`(${JoinZones} > 1) ? 1 : ${JoinZones}`
DampingZone	**Scalar**	`1-${CombinedShape}`
dnSlope	**Scalar**	`${ydn_up}+($$Position[0]*tan((`

Function Name	Type	Definition
		${ydn_dn}-${ydn_up})/(${xdn}-${xup})))
HanningCircle	**Scalar**	1-(0.5*(1-cos($pi*($Radius-$Rm)/$thk)))
HanningLong	**Scalar**	($$Position[1] < ${upSlope}+${thkLong} && $$Position[1] > ${upSlope}) ? 1-(0.5*(1-cos($pi*($$Position[1]+${thkLong}-${upSlope})/${thkLong}))) : 1-(0.5*(1-cos($pi*($$Position[1]+${thkLong}-${dnSlope})/${thkLong})))
HanningSrc	**Scalar**	1-(0.5*(1-cos($pi*($Radius-${Rsrc})/${thkSrc})))
JoinZones	**Scalar**	sqrt((${ZoneLong}*${ZoneLong}) + (${ZoneCircle}*${ZoneCircle}))
NoiseZone	**Scalar**	(${Radius} < ${Rsrc}) ? 1 : ((${Radius} > ${Rsrc} && ${Radius} < ${Rsrc}+${thkSrc}) ? ${HanningSrc} : 0)
pi	**Scalar**	4*atan(1)
Radius	**Scalar**	sqrt(($$Position[0]*$$Position[0])+($

Function Name	Type	Definition
		`$Position[1]*$ $Position[1]))`
Rm	**Scalar**	`0.05`
Rp	**Scalar**	`$Rm+$thk`
Rsrc	**Scalar**	`0.025`
thk	**Scalar**	`0.02`
thkLong	**Scalar**	`0.03`
thkSrc	**Scalar**	`0.02`
upSlope	**Scalar**	`${yup_up}+($ $Position[0]*tan(( ${yup_dn}-$ {yup_up})/(${xdn}- ${xup})))`
xdn	**Scalar**	`0.3`
xup	**Scalar**	`0.0065`
ydn_dn	**Scalar**	`-0.1`
ydn_up	**Scalar**	`-0.02`
yup_dn	**Scalar**	`0.1`
yup_up	**Scalar**	`0.02`
ZoneCircle	**Scalar**	`(${Radius} < $ {Rm}) ? 1 : (($ {Radius} > ${Rm} && ${Radius} < $ {Rp}) ? $ {HanningCircle} : 0)`
ZoneLong	**Scalar**	`($$Position[0] > $ {xup}) ? ( (abs($ $Position[1]) > $ {upSlope}+$ {thkLong}) ? 0 : ((abs($ $Position[1]) < $ {upSlope}) ? 1 : 1-$`

Function Name	Type	Definition
		`{HanningLong} ) ) : 0`

Setting the Solver Parameters and Stopping Criteria (Acoustic Wave)

Check that the solver parameters are correct and set the stopping criteria to give the simulation run another 0.01 seconds of physical time.

The Segregated Flow solver and Implicit Unsteady solver parameters are unchanged from the previous unsteady simulation. The default settings for the Acoustic Wave solver are acceptable for this simulation.

To set the solver parameters and stopping criteria:

1. Select the **Solvers** node and check that the following properties are set:

Node	Property	Setting
Acoustic Wave	*Start-time*	`0.0 s`
	Newmark Delta	`0.5`
Segregated Flow		
▮▮**Velocity**	*Under-Relaxation Factor*	`0.9`
▮▮**Pressure**	*Under-Relaxation Factor*	`0.7`
Implicit Unsteady	*Time-Step*	`2.5e-5 s`
	Temporal Discretization	**2nd-order**

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	*Enabled*	Activated
	Maximum Physical Time	`0.02 s`

3. Save the simulation.

Visualizing the Acoustic Pressure

Create a scalar scene to display the acoustic pressure throughout the flow region.

To visualize the acoustic pressure:

1. Right-click the **Scenes** node and select **New Scene > Scalar**.
2. Select the **Scalar Scene 1** node and rename it as **Acoustic Pressure**.
3. Edit the **Acoustic Pressure > Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Derived Parts > plane section**
Scalar Field	*Function*	**Acoustic Pressure**

4. Save the simulation.

Running the Simulation (Acoustic Wave)

Preparation of the acoustic wave simulation is now complete, and the simulation can be run.

To run the acoustic wave simulation:

1. Click (**Run**) in the *Solution* toolbar.

   You can click the *Acoustic Pressure* tab at the top of the *Graphics* window to view the development of the acoustic pressure field as the simulation progresses.



The simulation continues until 0.02 seconds of physical time is complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

When the simulation is complete, examine the acoustic pressure prediction.

The flow over the cylinder is illustrated by the velocity magnitude scene that is shown below:



The acoustic prediction (below) shows a field that is typically observed in laminar vortex shedding from a cylinder.



Summary

This tutorial demonstrated how to set up an acoustic wave simulation in STAR-CCM+, using the flow around a cylinder.

This tutorial demonstrated how a flow-generated noise source can be identified and propagated through the zone of interest using a hybrid CFD/CAA modeling technique. In this case, noise that is generated by flow over the cylinder is captured and the noise source information is passed onto the acoustic wave solver.

This tutorial has covered the following:

- Obtaining the mean flow pressure from an unsteady flow solution.
- Selecting the Acoustic Wave model.

- Specifying the noise source zone to enclose all of the important noise sources that are generated in the flow.
- Minimizing spurious acoustics: specifying the acoustic damping zone to ensure damping in the fine-to-coarse mesh transition regions away from the zone of interest.
- Setting the solver parameters and stopping criteria.
- Displaying the acoustic pressure in a scalar scene.

Electromagnetism

The following tutorial illustrates features for solving problems that involve an electric field.

Contents

Ohmic Heating: Domestic Fuse

Ohmic Heating: Domestic Fuse

This tutorial demonstrates some of the capabilities of the Ohmic Heating model available in STAR-CCM+.

A simple domestic fuse geometry is used in the simulation to demonstrate the required steps. The geometry is shown below.



A fuse is a safety device to protect electrical equipment against unexpected high currents. Because of the Joule effect, the current going through the thin wire inside the fuse increases its temperature. If the current hits the critical current, the wire temperature exceeds its melting point, and breaks.

The fuse assembly in this tutorial is 30 mm long and is composed of metal, ceramic, and silicium. The wire inside the ceramic enclosure is 27.5 µm in diameter and has a critical current of 7 A with ±0.7% error.

A CAD model has been prepared in which various parts are imprinted correctly on adjacent parts. When this model is imported into STAR-CCM+ (as geometry parts), coincident faces are automatically detected and used to generate part contact data. The scenario being simulated is rather simple: to increase the potential difference between the two

terminals of the wire, thus inducing a current increase until it reaches the critical point.

The post-processing tools in STAR-CCM+ are used to visualize the results.

Contents

Prerequisites
Importing the Surface Mesh and Naming the Simulation
Examining the Imported Geometry
Increasing the Tessellation Density of the Fuse Wire
Assigning Geometry Parts to Regions
Setting up the Mesh
Adding a Volumetric Control
Generating the Mesh
Specifying Physics Models and Material Properties
Assigning Physics Continua
Setting Global Initial Conditions
Defining Boundary Conditions
Setting up Reports, Monitors, and Plots
Setting up a Plane Section
Setting Up a Scalar Scene
Setting Solver Parameters
Setting Stopping Criteria
Running the Simulation
Summary

Prerequisites

The instructions in the Ohmic Heating: Domestic Fuse tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	

Importing the Surface Mesh and Naming the Simulation

A CAD file containing the fuse assembly has been prepared for this analysis.

Begin by copying the file from the STAR-CCM+ installation directory, starting the simulation, importing the CAD model and saving the simulation.

1. Create a directory for the tutorial called `ohmicHeating`.
2. In this directory, copy the files containing the problem geometry (`fusePartsAssembly.x_b`) and the potential levels data (`potentialLevels.csv`) from the `doc/startutorialsdata/electromagnetism/data` subdirectory of your STAR-CCM+ installation directory.
3. Start up STAR-CCM+ in a manner that is appropriate to your working environment and create a New Simulation.
4. Click (**Import Surface Mesh**) from the tool bar.
5. In the *Open* dialog, select file `fusePartsAssembly.x_b`.



6. Click **Open** to start the import.
7. In the *Import Surface Options* dialog that appears, do the following:
   a) Select the **Create New Part** from the **Import Mode** box.
   b) Make sure that **Create Part Contacts from Coincident Entities** option is selected.
   c) Click **OK** to retain the default options.

STAR-CCM+ provides feedback on the import process in the **Output** window. In particular, note the message `Created 8 new Part Contacts from 10 pairs of coincident faces`. A geometry scene is created in the **Graphics** window.

8. Finally, save the simulation to disk under file name `fuseWire.sim`.

Examining the Imported Geometry

The imported CAD parts, which represent the different parts of the fuse and its holder, are contained in the **Geometry** > **Parts** manager node.

The various parts were previously imprinted on each other to provide pairs of matching faces on the surfaces that are in contact. STAR-CCM+ detects the matching faces on import, and creates part/part contact data that are used later to define interfaces between regions.

These contacting surfaces can be viewed successively as follows:

1. Open the **Geometry** > **Parts** node and multi-select all of the parts. Right-click one of the selected nodes and select **Browse Part/Part Contacts**.

   The *Part/Part Contact Browser Options* dialog appears:



2. In the *Part/Part Contact Browser Options* dialog, make sure that all the parts are selected and click **OK**.

   The *Part/Part Contact Browser* dialog opens in the **edit** window and a scene in the **Graphics** window. .



3. Select **Wire** from the **Source Parts** menu. Use the navigation buttons under the **Target Parts** menu to step through the **Clip+ve**, **Clip-ve**, and **Silicium** parts with which the wire is in contact.
4. Click **Close**.

The surfaces and feature curves of geometry parts are stored as individual entities on which operations can be performed if necessary. No further operations are required in this tutorial, but it is instructive to examine the surfaces for the electric potential.

5. Open the **Parts** > **Clip+ve** > **Surfaces** and **Parts** > **Clip-ve** > **Surfaces** nodes. Multi-select the **terminal+** and **terminal-** nodes.



These surfaces are highlighted in the **Graphics** window.



6. Use the mouse buttons to zoom in and position the part so that the view appears as shown.

7. Click (Save-Restore-Select views) in the *Vis* toolbar and select **Store Current View**.

Increasing the Tessellation Density of the Fuse Wire

To ensure that the thin fuse wire has a well-defined surface, increase the tessellation density of the imported surface to represent the surface with a higher density of triangles.

1. Right-click the **Geometry** > **Parts** > **Wire** node and select **Re-tessellate**.
2. In the *Tessellation Options* dialog, set **Tessellation Density** to **Fine**.



3. Click **OK**.

Assigning Geometry Parts to Regions

Assign geometry parts to regions before surface or volume meshing takes place. STAR-CCM+ allows you to create regions and boundaries that are based on the parts and their surfaces.

1. Multi-select all of the parts. Right-click the selected nodes and select **Assign Parts to Regions**.
2. In the *Assign Parts to Regions* dialog:
   a) Set the Region mode to **Create a Region for Each Part**.
   b) Set the Boundary mode to **Create a Boundary for Each Part Surface**.
   c) The Feature Curve mode can be left at **Create One Feature Curve for All Part Curves**. The **Create Interfaces From Contacts?** option remains selected.
   d) Click **Apply** and then **Close**.

The necessary regions and boundaries are created, along with interfaces corresponding to the part/part contact information reviewed earlier.

3. Expand the **Regions** and **Interfaces** nodes.



There are eight interfaces corresponding to the eight sets of part/part contact data.

Setting up the Mesh

Use the surface remesher with the polyhedral mesher to create the volume mesh.

To create the mesh:

1. Right-click **Continua** and create a new mesh continuum.
2. For the **Continua** > **Mesh 1** continuum, select the following mesh models:

Group Box	Model
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**

3. Click **Close**.
4. Edit the **Continua** > **Mesh 1** > **Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	*Value*	`1 mm`
Surface Curvature > Basic Curvature	*# Pts/Circle*	`72.0`
Tet/Poly Density	*Growth Factor*	`1.5`



To have adequate resolution on the surface of the fine wire, set a custom surface size on the interface between the wire and the surrounding solid.

5. Expand the **Interfaces** > **Silicium/Wire** > **Mesh Conditions** node and select the **Custom Surface Size** node. In the **Properties** window, activate the *Custom surface size*.



6. Select the **Mesh Values > Surface Size Relative Minimum Size** node and set the *Percentage of Base* to `2.5`.



7. Now select the **Relative Target Size** node and set the *Percentage of Base* to `5.0`



Adding a Volumetric Control

Add a volumetric control to provide adequate resolution of the volume mesh within the wire. When adding volumetric controls, it is first necessary to create the volume shape within which custom mesh controls applies, and thereafter to select this shape within a volumetric control itself.

To add a volumetric control:

1. Create the Volume Shape:

   a) Expand the **Tools** node, right-click the **Volume Shapes** node and choose **New Shape > Cylinder**.

   Assuming that you have the **Geometry Scene 1** window still visible, the *Create Cylinder Volume Shape* dialog appears underneath the **edit** tab.

   b) In the **Start Coordinate** box set the coordinates to `[-0.01, 0.0065, 0.0] m`.

c) In the **End Coordinate** box set the coordinates to `[0.01, 0.0065, 0.0] m`.

d) Set the *Cylinder Radius* to `0.000028 m`.
The dialog appears as shown below.



e) Click **Create** and then **Close**.
A **Cylinder 1** node is added to the **Volume Shapes** manager.

2. Add the Volumetric Control:

a) Expand the **Continua** > **Mesh 1** node. Right-click the **Volumetric Controls** node and select **New**.

b) Rename the **Volumetric Control 1** node to `Cylinder`.

c) In the **Properties** window for this node, click the ... **(Custom Editor)** button to the right of the *Shapes* property.

d) In the *Cylinder* dialog that appears, click the **Select Cylinder 1** button.



e) Click **OK**.

In the simulation tree, the **Cylinder** node changes color to show that a volume shape has been selected.

f) Select the **Mesh Conditions Polyhedral Mesher** node and activate the *Customize polyhedral mesher*.

g) Select the **Surface Remesher** node and activate the *Customize surface remesher*.

A **Custom Size** node is added to the **Mesh Values** folder.

h) Select the **Mesh Values** > **Custom Size** > **Relative Size** node and set *Percentage of Base* size to `1.5`.



Generating the Mesh

You can now create the volume mesh for the simulation.

To generate the volume mesh:

1. Click (**Generate Volume Mesh**) in the **Mesh Generation** toolbar.

   The mesh generation process takes some time depending on the speed of your computer. The output from the surface and volume meshers are shown in the *Output* window and the ongoing status is shown in the status bar.

2. Save the simulation .

Specifying Physics Models and Material Properties

You can now specify the physical models and the material properties for this simulation.

Three physics continua are required for this simulation:

- Copper (Ohmic Heating): For the Clip+ve, Clip-ve, and Wire regions through which electric current passes.
- Ceramic: For the Ceramic region, which is the outer casing of the central portion of the fuse.
- Silicium: Fills the annulus between the wire and the ceramic casing.

Contents

Selecting Copper (Ohmic Heating) Physics Models
Specifying Copper (Ohmic Heating) Material Properties
Selecting Ceramic Physics Models
Specifying Ceramic Material Properties
Selecting Silicium Physics Models
Specifying Silicium Material Properties

Selecting Copper (Ohmic Heating) Physics Models

Define the models for the Copper (Ohmic Heating) continuum.

By default STAR-CCM+ has created a **Physics 1** continuum.

To select the physics models:

1. Rename the **Continua** > **Physics 1** node to `Copper (Ohmic Heating)`.
2. For the physics continuum, **Continua** > **Copper (Ohmic Heating)**, select the following models in order:

Group Box	Model
Time	**Steady**
Enabled Models	**Three-Dimensional** (Selected automatically) **Gradients** (Selected automatically)
Material	**Solid**
Equation of State	**Constant Density**

Group Box	Model
Optional Physics Models	**Segregated Solid Energy** **Electromagnetism** **Ohmic Heating**
Electromagnetism	**Electrodynamic Potential**

3. Click **Close**.

Specifying Copper (Ohmic Heating) Material Properties

The properties of pure aluminum are set on a solid continuum by default. Replace this material with Copper.

To specify the material properties:

1. Right-click the **Models > Solid Al** node and select **Replace with**.
2. In the *Replace Material* dialog, expand the **Material Databases > Standard > Solids** nodes and select **Cu (Copper)**.



3. Click **OK**.

Define a field function to relate the electrical conductivity of copper to its temperature.

4. Select the **Cu** > **Material Properties** > **Electrical Conductivity** node and set the **Method** to **Field Function**.



The electrical conductivity expression depends on the temperature as follows:

$$\text{Electrical Conductivity} = 26.08 \cdot 10^9 \cdot T^{-27/25}$$

5. Right click the **Tools** > **Field Functions** node and select **New** > **Scalar**.
6. Rename the **User Field Function 1** node to `Copper Electrical Conductivity`.
7. Select the **Copper Electrical Conductivity** node and set the following properties:

Property	**Setting**
Dimensions	**ElectricConductance**
Function Name	`CopperElectricalConductivity`
Definition	`26.08e9 * pow($Temperature,(-27/25))`

8. Return to the **Electrical Conductivity** node, select **Field Function**, and set *Scalar Function* to **Copper Electrical Conductivity**.



Selecting Ceramic Physics Models

Define the models for the Ceramic continuum.

To select the physics models:

1. Create a physics continuum and rename the **Physics 1** node to `Ceramic`.
2. For the physics continuum, **Continua** > **Ceramic**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Steady**
Material	**Solid**
Optional Physics Models	**Segregated Solid Energy**
Equation of State	**Constant Density**

3. Click **Close**.

Specifying Ceramic Material Properties

You now specify the material properties for the ceramic.

To specify the material properties:

1. Rename the **Ceramic** > **Models** > **Solid** > **Al** node to `Porcelain`.
2. Select the **Porcelain** > **Material Properties** > **Density** > **Constant** node and set the **Value** to `2500 kg/m^3`.



3. Select the **Specific Heat** > **Constant** node and set the **Value** to `1080 J/kg-K`.
4. Select the **Thermal Conductivity** > **Constant** node and set the **Value** to `0.928 W/m-K`.

Selecting Silicium Physics Models

Define the models for the Silicium continuum by copying and pasting the Ceramic continuum.

To select the Silicium physics models:

1. Right-click the **Continua** > **Ceramic** node and select **Copy**.



2. Right-click the **Continua** node and select **Paste**.



An entry **Copy of Ceramic** is added to the object tree.

3. Rename the **Continua** > **Copy of Ceramic** node to `Silicium`.

Specifying Silicium Material Properties

You now specify the material properties for silicium.

To specify the material properties:

1. Rename the **Silicium** > **Models** > **Solid** > **Porcelain** node to `SiO2`.
2. Select the **SiO2** > **Material Properties** > **Density** > **Constant** node and set the **Value** to `2200 kg/m^3`



3. Select the **Specific Heat** > **Constant** node and set the **Value** to `740 J/kg-K`.
4. Select the **Thermal Conductivity** > **Constant** node and set the **Value** to `1.4 W/m-K`.

Assigning Physics Continua

For a multi-region simulation, you must assign a physics continuum to each region. Through this assignment, the different physical models and material properties are assigned to the respective region.

To assign physics continua to regions:

1. Select the **Regions** > **Ceramic** node and set the *Physics Continuum* to **Ceramic**.



2. Select the **Regions** > **Silicium** node and set the *Physics Continuum* to **Silicium**.



The other three regions are associated with the Copper (Ohmic Heating) continuum by default.

3. Save the simulation .

Setting Global Initial Conditions

Specify the initial electrical potential value for the Copper (Ohmic Heating) model.

To set the initial conditions:

1. Select the **Continua** > **Copper (Ohmic Heating)** > **Initial Conditions** > **Electric Potential** > **Constant** node and set the *Value* to `0.2 V`.



Defining Boundary Conditions

Set the boundary conditions for the Ceramic, Clip+ve, and Clip-ve regions.

To define the boundary conditions:

1. Select the **Regions** > **Ceramic** > **Boundaries** > **outerceramic** > **Physics Conditions** > **Thermal Specification** node and set the *Condition* to **Temperature**.
2. Select the **Physics Values** > **Static Temperature** > **Constant** node and set the *Value* to `310` K.



Set the boundary conditions for the Clip+ve region to have a stepped electrical potential, read from a table of Iteration vs. Potential Levels. The table is supplied as a CSV file, `potentialLevels.csv`.

3. Right-click the **Tools** > **Tables** node and select **New Table** > **File Table**.
4. In the *Open* dialog, select `potentialLevels.csv` and click **Open**.

   potentialLevels is added to the object tree under the **Tables** node.
5. Right-click the **Tables** > **potentialLevels** node and select **Tabulate**.

   The content of the file is displayed in the *Tabular Data* dialog as shown below.



6. Click **Close**.
7. Select the **Regions** > **Clip+ve** > **Boundaries** > **terminal+** > **Physics Conditions** > **Electric Potential Specification** node and set the *Method* to **Electric Potential**.



8. Select the **Physics Values** > **Electric Potential** node and set the *Method* to **Table (iteration)**.



9. Select the **Electric Potential > Table (Iteration)** node. In the **Properties** window, do the following:

   a) Set the *Table* to **potentialLevels**.

   b) Set the *Table: Data* to **PotentialLevel**.

   c) Set the *Interpolation* to **Step**



10. Select the **Regions** > **Clip-ve** > **Boundaries** > **terminal-** > **Physics Conditions** > **Electric Potential Specification** node and set the *Method* to **Electric Potential**.
11. Select the **Physics Values** > **Electric Potential** > **Constant** node and set the *Value* to `0.02 V`.



12. Save the simulation.

Setting up Reports, Monitors, and Plots

By default STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the solution. It is useful to create more monitors to help in assessing convergence.

Create two sets of reports, monitors and plots for the current and maximum temperature in the wire.

For the current-in-wire report, generate a plane section normal to the X direction using the wire as the input part. The plane section, is placed half way along the wire. It is used to create a surface integral of the current density in the Y direction (A/m^2).

1. Click the **Geometry Scene 1** tab in the **Graphics** window.
2. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.



In the *Create Section* dialog that appears in the **edit** window, do the following:

3. In the **Input Parts** box, select **Wire**.
4. Set the **origin** to `[0, 0.00242566, 0] m` and the **normal** to `[1, 0, 0] m` for X, Y, and Z.
5. In the **Display** box, select **No Displayer**.

   The *Create Section* now appears as follows:



6. Click **Create**, then click **Close**.
   A **plane section** derived part is created.

7. Rename the **Derived Parts > plane section** node to `Wire Cross-section`.
8. Right-click the **Reports** node and select **New Report > Surface Integral**.



9. Rename the **Surface Integral 1** node to `Current in Wire`.

Reports can be created for all derived parts and boundaries.

10. Select the **Reports > Current in Wire** node.
11. In the **Properties** window, set *Scalar Field Function* to **Electric Current Density > Magnitude**.
12. Set *Parts* to **Derived Parts > Wire Cross-section**.
13. Click **OK**.

Create a report for the maximum temperature in the wire.

14. Right-click the **Reports** node and select **New Report > Maximum**.
15. Rename the **Maximum 1** node to `Maximum Temperature in Wire`.

Select the **Wire** region and also set the scalar field to **Temperature**.

16. Make sure that the **Reports** > **Maximum Temperature in Wire** node is selected. In the **Properties** window, do the following:
    a) Set *Scalar Field Function* to **Temperature**.
    b) Set *Parts* to **Regions** > **Wire**.

Monitors and plots can be created from reports, and used to judge convergence while the simulation is running.

17. Multi-select the **Current in Wire** and the **Maximum Temperature in Wire** report nodes.
18. Right-click and select **Create Monitor and Plot from Report**.



19. Click **Multiple Plots (one per report)** to display the plots in separate graphs.



Two monitors and plots are created.



20. Right-click the **Plots** > **Current in Wire Monitor Plot** node and select **Open**.
21. Similarly, open the **Maximum Temperature in Wire Monitor Plot**.

Setting up a Plane Section

You can use a plane section to visualize the temperature within the parts.

Generate a plane section normal to the Z direction using the ceramic, silicium, and wire as the input parts. The plane section is placed at the center of the fuse assembly. Use the plane in a scalar scene to display the temperature within the parts.

1. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane**.
2. In the *Create Section* in-place dialog:
   a) Set **Input Parts** to **Ceramic**, **Silicium**, and **Wire** only.
   b) Set **origin** to `[0, 0.0024209, 0] m` for X, Y, and Z.
   c) Set **normal** to `[0, 0, 1] m` for X, Y, and Z.
   d) In the **Display** box, select **No Displayer**.
3. Click **Create**, then click **Close**.
4. Rename the **plane section** node to `Fuse Cross-section`.
5. Save the simulation.

Setting Up a Scalar Scene

To display the electric potential for the two clip regions and the temperature for the fuse cross-section derived part, you create a single scalar scene.

To create a scalar scene with two scalar displayers:

1. Create a scalar scene.
2. Expand the **Scalar Scene 1** > **Displayers** > **Scalar 1** node and set the following properties:

Node	Property	Value
Parts	*Parts*	**Clip+ve** **Clip-ve**
Color Bar	*Levels*	2
	Number of Labels	3

3. In the **Scalar Scene 1 Graphics** window, right-click **Select Function** and select **Electric Potential**.

Add another scalar displayer for the temperature of the fuse cross-section derived part.

4. Right-click the **Scalar Scene 1** > **Displayers** node and select **New Displayer** > **Scalar**.
5. Expand the **Displayers** > **Scalar 2** node and set the following properties:

Node	Property	Value
Parts	*Parts*	**Fuse Cross-section**
Color Bar	*Color Map*	`cool-warm`

6. In the **Scalar Scene 1 Graphics** window, right-click **Select Function** and select **Temperature**.
7. Use the mouse buttons to zoom in and position the part as shown:



8. Save the simulation.

Setting Solver Parameters

You change the solver settings to help the solution converge quickly and maintain stability.

To set the solver parameters:

1. Select the **Solvers** > **Segregated Energy** node and set the *Solid Under-Relaxation Factor* to `1`.



2. Select the **Solvers** > **Electric Potential** node and set the *Under-Relaxation Factor* to `1`.

Setting Stopping Criteria

You now specify for how long the simulation runs.

Set the simulation to run for 500 iterations.

1. Select the **Stopping Criteria** > **Maximum Steps node** and set the *Maximum Steps* to `500`.



2. Save the simulation.

Running the Simulation

The simulation is now ready to be run.

To run the simulation:

1. Initialize the solution by clicking (**Initialize Solution**) in the **Solution** toolbar.

   A scalar scene appears as shown:



   The temperature and electric potential within the fuse assembly are initialized to the predefined values.

2. Run the simulation by clicking (**Run**) in the **Solution** toolbar.

   The iteration progress can be monitored in the *Output* window. You can monitor the convergence by viewing the Residuals using the visualization tabs in the *Graphics* window.

   When the last iteration has been completed, the message `Stopping Criterion Maximum Steps satisfied` is displayed in the *Output* window.

3. Save the simulation.

Contents

Temperature and Electric Potential Plot

Report Plots

Temperature and Electric Potential Plot

You can now visualize the final temperature and electric potential distribution.

To visualize the temperature and electric potential plot:

1. Click the **Scalar Scene 1** tab at the top of the **Graphics** window.
   A scene appears as shown:



Report Plots

You visualize the current and the maximum temperature in the wire in the form of x-y plots.

To visualize the plots:

1. Click the **Current In Wire Monitor Plot** tab at the top of the **Graphics** window.
   A scene appears as shown



2. Click the **Maximum Temperature Monitor Plot** tab at the top of the **Graphics** window.

   A scene similar to the one shown appears



Summary

This tutorial demonstrated how to set up a simulation involving Ohmic Heating in STAR-CCM+.

This tutorial has introduced the following STAR-CCM+ features:

- Importing a surface mesh, creating a part and naming the simulation.
- Visualizing and examining the imported geometry.
- Generating a volume mesh using a surface remesher and polyhedral mesher.
- Setting the physics models for an Ohmic heating analysis.
- Implementation of user field functions.
- Setting up report, monitors, and plots.
- Post-processing electric potential and temperature.

Battery

The tutorials in this set illustrate various STAR-CCM+ features for setting up a battery model:

Contents

Battery Simulation Module: Cell Thermal Analysis
Cylindrical Cells: Cell Thermal Analysis
Li-Ion Battery Cell Model: Cell Electrochemistry Analysis

Battery Simulation Module: Cell Thermal Analysis

This tutorial introduces the battery modeling capabilities of STAR-CCM +. It simulates the battery heat generation that occurs during the first 90 seconds of a discharge cycle.

The tutorial takes you through the steps that are required to run a battery analysis in STAR-CCM+ using files that are created in Battery Design Studio® (BDS).

Additional licenses are required to use the battery modeling capabilities:

- For battery modeling, STAR-CCM+ requires a `batterysim` license in addition to the other license features.
- BDS, which is supplied for use with the STAR-CCM+ battery modeling capability, requires either a `batterysim` or a `batterydesignstudio` license.

The tutorial is divided into three parts:

1. Preparing files in BDS
2. Running a STAR-CCM+ simulation with a Single Cell
3. Running a STAR-CCM+ simulation with 3 Cells in Series

The battery heat generation that occurs during the first 90 seconds of a discharge cycle is simulated. A simple scenario with a single cell battery is run first. The simulation is then rerun after modifying the battery, adding two more cells in series.



Both simulations use a voltage plot to monitor the battery discharge during the cycle. A scalar scene displays how battery heat generation is distributed on the battery cells.

Contents

Prerequisites
Preparing Files in Battery Design Studio®
Running a STAR-CCM+ Simulation with a Single Cell
Running a STAR-CCM+ Simulation with 3 Cells in Series
Summary

Prerequisites

Before attempting this tutorial, make sure that you are familiar with the following techniques.

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Preparing Files in Battery Design Studio®

Prepare battery data in BDS and import it into STAR-CCM+.

Battery Design Studio (BDS) is a tool for the design, simulation, and analysis of batteries. Battery data from BDS is imported into STAR-CCM + and can be used to set up complex scenarios, such as simulating cooling by including a fluid region around the battery cells.

To use a battery cell in a STAR-CCM+ simulation, the following files must be generated in BDS:

- `.tbm`—a text battery model file, containing all the information about the battery. This information includes construction, sizing, materials, and physics models.
- `.prg`—a program file, containing details of the battery load cycle

With STAR-CCM+, it is possible to import battery data from BDS and set up complex scenarios that include a fluid region around the battery cells to simulate cooling.

In the following topics, `tbm` and `prg` files are generated for a lithium-ion battery cell in BDS.

If you do not wish to work through the steps involving BDS, you can jump directly to Running a STAR-CCM+ simulation with a Single Cell.

Contents

Creating a Battery Cell
Creating a Program File

Creating a Battery Cell

Create a battery cell in BDS.

Use the Li-Ion stack cell that is available in BDS as a starting point.

To create a battery cell in BDS:

1. Launch BDS.

2. Create a workspace by selecting **File > New Workspace**.



The *New Workspace* dialog appears.

3. Enter `BasicCell` as the *Name* in the textbox, and select an appropriate location.



4. Click **OK**.
5. If the location does not exist, confirm the directory creation.

Components of the battery are created using the tree on the left of the BDS interface. Create a battery cell.

6. Right-click the **Batteries** folder and select **Add New Battery > LiIon Stack > LiIonStack**.



A new node, **LiIonStack**, is added below the **Batteries** node folder, and a corresponding `.tbm` file is created in the workspace directory.

The cell can be renamed immediately after it is created.

7. Rename the battery cell `TutorialCell`.



8. To see a diagram of the battery cell, select **Main** in the *TutorialCell.tbm* feature tree.



The position of the tabs on each electrode is specified using an offset value. To reposition the tabs, change the offset values.

9. Select the **+ Electrode** node in the *TutorialCell.tbm* window.
10. Set *Tab offset (Otab)* to `3.500`.



11. Select the **- Electrode** node in the *TutorialCell.tbm* window.
12. Set *Tab offset (Otab)* to `50.500`.



Select a battery model. This model defines how the battery behaves in response to different electrical loads. In this tutorial, the NTG 3D model is used.

13. Scroll down to the **Models** node in the *TutorialCell.tbm* window, and select the NTG 3D model from the **IET** folder.



The addition of a * to the *TutorialCell.tbm* window title indicates that you have changed the cell definition. Once the asterisk appears, **Build** the cell. This process rewrites the `.tbm` file in the workspace directory and prompts BDS to check the cell definition for any invalid changes.

14. Right-click the **Batteries** > **TutorialCell** node in the BDS feature tree and select **Build**.



15. If prompted to save, accept.
    A message box appears informing you that the build was successful.
16. Click **OK** to close the message box.

The cell definition is complete.

Creating a Program File

Create a program file in BDS.

The next step is to define a procedure for cycling the battery. In this case, specify a constant current discharge of 10 A, for 1 hour. The capacity of the cell is 10 Ah, so this procedure is equivalent to a 1 C discharge. You can find out the cell capacity by selecting the **Report** node in the *TutorialCell.tbm* window, and reading the value from the **Cell Properties** table.



To create a procedure:

1. Right-click the **Procedures** node and select **Add new procedure > IET Cycler**.



The **IET Cycler** is a simulation model that allows you to define the Current (I), Voltage (E), and Temperature (T). A **Cycler** node is added under the **Procedures** folder, and a corresponding program file (.prg) is created in the workspace directory. A new window that is called *Cycler.prg* appears in the work area.

2. Rename the program file `Tutorial1CDischarge`.



The name of the *Cycler.prg* window changes to *Tutorial1CDischarge.prg*.

3. Click **<add>** under the *Type* column for the **Step** in the *Tutorial1CDischarge.prg* window.



4. Click the cell twice for the drop-down menu to appear and select **Discharge**.



The constant current discharge is specified as a **Control Condition**.

5. Click in the first cell of the *Type* column for the **Control Condition** and select **Current**.



6. Enter `10A` as the *Value*.



The discharge finishes after 1 hour, which is specified as an **End Condition**.

7. Click in the *Type* column for the **End Condition** and select **Time** from the drop-down list.



8. Set the *Value* of **Time** to `3600s`.



Note: The value of `Time` must be entered in seconds.

9. Click **<add>** under the *Type* column for the **Step** and select **End**.



The definition of the battery cell and program is now complete.

10. Save the workspace by selecting **File > Save Workspace** from the menu.



11. Click **OK**.

Running a STAR-CCM+ Simulation with a Single Cell

Import the files from BDS into STAR-CCM+ and run a simple single cell simulation.

In this section of the tutorial, the files that you created in BDS are imported into STAR-CCM+, and a simple single cell simulation is run.

To create the simulation:

If you skipped the section involving BDS, create a working directory and copy the files that are required for this tutorial.

1. Create a working directory called `BasicCell`.
2. Copy the following files from the `doc/startutorialdata/battery/data` subdirectory of your STAR-CCM+ installation directory into the `BasicCell` directory:
   - `TutorialCell.tbm`
   - `Tutorial1CDischarge.prg`

Create the simulation.

3. Launch STAR-CCM+.
4. Start a simulation.
5. Save the simulation as `batteryTutorial.sim` in the `BasicCell` directory.

Contents

Selecting the Physics Models
Importing the Battery Cell
Generating a Battery Module
Setting Up the Battery Mesh
Loading the Program Data
Creating a Circuit
Preparing a Scalar Scene
Setting Up a Voltage Monitor
Setting Solver Parameters and Stopping Criteria
Running the Simulation
Visualizing the Results

Selecting the Physics Models

Physics models define the primary variables of the simulation and what mathematical formulation is used to generate the solution.

To select the physics models:

1. Create a physics continuum and rename the **Continua** > **Physics 1** node to **Battery Physics**.
2. For the physics continuum, **Continua** > **Battery Physics**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Implicit Unsteady**
Material	**Solid**
Optional Models	**Battery**
Energy	**Segregated Solid Energy**
Equation of State	**Constant Density**

3. Click **Close**.
4. To review the models, open the **Battery Physics** > **Models** node.



The **Batteries** node appears in the STAR-CCM+ feature tree. The battery pre-processing is carried out primarily under this node.

5. Save the simulation.

Importing the Battery Cell

The battery file created in BDS is imported into STAR-CCM+ as a battery cell.

The battery (.tbm) file that you created in BDS is imported into STAR-CCM+ as a battery cell.

To import the battery cell:

1. Expand the **Batteries** node.
2. Right-click **Batteries** > **Battery Cells** and select **Create from Tbm...**.



3. In the *Import Battery Data From File* dialog, navigate to the BDS workspace directory, `BasicCell`, and select the file `TutorialCell.tbm`.
4. Click **Open**.
5. In the *Import Battery Options* dialog, ensure that all objects are selected, then click **OK**.



During the import process, STAR-CCM+ builds an idealized model of the cell using the 3D-CAD module. The model is displayed in the *3D-CAD View 1* scene in the *Graphics* window.

A battery module comprises one or more battery cells. The cells can be arranged in series and parallel. You specify the arrangement and number of cells. As a module typically contains multiple instances of the same cell, STAR-CCM+ automatically copies the imported cell for more instances.

In this section of the tutorial, the module contains a single cell.

6. Right-click **Batteries** > **Battery Modules** and select **New**.

   A new **Battery Module** node is added to the tree. In the *Battery Module - Properties* window, the *N Series* and *N Parallel* properties define the number of cells in the module. By default these values are both set to `1`.

The next step is to set up the module cells.

7. Right-click **Battery Module** > **Battery Module Cells** and select **Setup Battery Module Cells**.



The panel that appears allows you to:

- Define the offset of battery cells in relation to each other.
- Flip or rotate cells.
- Assign different cells to each position in the battery module.

In this case, the default settings are suitable.

8. Click **Apply** to assign the imported cell to the battery module.



9. Click **Close**.

Generating a Battery Module

Generate the battery module in STAR-CCM+.

You are now ready to generate the battery module. This process involves generating parts, and then generating regions. For modules that contain several cells, this process also duplicates the battery cells according to the battery module setup.

To generate the battery module:

1. Right-click the **Battery Modules** > **Battery Module** node and select **Generate Battery Parts.**.



2. Right-click the **Battery Modules** > **Battery Module** node and select **Generate Regions...**.



An *Additional (non-Battery) Parts to Generate Regions* dialog appears. This dialog is empty because you do not have any non-battery parts in this simulation.



3. Click **OK**.
4. Expand the **Geometry** > **Parts** node, and the **Regions** node, to see the parts and regions that were created during the process.



Interfaces between the stack, tab roots, and tab stems have automatically been created during the generate regions action.

5. Select the **Interfaces > Battery Module: Battery Module Cell 0,0: +Ve Tab Root/Battery Module: Battery Module Cell 0,0: Stack** node.



The selected interface is highlighted in *3D-CAD View 1*.



6. Save the simulation.

Setting Up the Battery Mesh

Define the battery mesh and set the appropriate properties.

The battery mesher generates a thin mesh on the tab stems, tab roots, and stack for thermal mesh and also generates a separate Cartesian electrical mesh on the stack. The mesh parameters for the tabs and thermal mesh are specified in a mesh continuum, while the number of cells in the electrical mesh are specified on the battery cell. Define the electrical mesh first.

To set up the mesh continuum:

1. Right-click the **Batteries** > **Battery Modules** > **Battery Module** node.
2. Select **Setup Battery Meshing**.



A new **Mesh 1** node appears under the **Continua** node. This new battery mesh continuum already contains the **Surface Remesher** and **Thin Mesher** models selected.

3. Rename the new Mesh 1 node to `Battery Mesh`..

Confirm that the **Battery Mesh** continuum contains the correct meshing models.

4. Extend the **Continua** > **Battery Mesh** node.
5. Ensure that the **Surface Remesher** and **Thin Mesher** models are selected.
6. Click **Close**.

Modify the **Battery Mesh** continuum **Reference Values**.

7. Expand the **Battery Mesh** > **Reference Values** node.
8. Select the **Reference Values** > **Base Size** node and set the **Value** to `20 mm`.

The properties of the **Battery Cell** node include the number of cells in the X, Y, and Z directions for the electrical stack. Initially, these

parameters are defined based on the values in the `.tbm` file, but they can be changed to modify the stack electrical mesh resolution. Further information on setting up the mesh is provided in the STAR-CCM+ User Guide. You cannot view the electrical mesh.

To change the number of cells in the electrical mesh:

9. Select the **Batteries** > **Battery Cells** > **Battery Cell** node.
10. Set the *Electrical Mesh Dimensions* to `[5,5,1]`.



11. Click (**Generate Volume Mesh**) to build the volume mesh.
12. Create a Mesh Scene.

    Mesh Scene 1 is displayed in the *Graphics* window.

    The volume mesh is shown below.



Loading the Program Data

Load the discharge procedure that you created in BDS into STAR-CCM+.

Load the discharge procedure that you created in BDS into STAR-CCM+ using the **Program Files** node in the STAR-CCM+ tree.

To load the program data:

1. Right-click **Batteries** > **Program Files** and select **Load**.



2. In the *Open* dialog, navigate to the `BasicCell` directory and select the `Tutorial1CDischarge.prg` file.
3. Click **Open**.

   The imported file is displayed under the **Program Files** node.



Creating a Circuit

Create a circuit to simulate the response of the battery.

The program file that you loaded in the previous step represents an electrical load on the battery. To simulate the response of the battery, the electrical load must be connected to the battery in a circuit.

Connect the electrical load using the **Circuit Solver** tool.

1. Expand the **Circuit** > **Circuit Solver** node.

The **Circuit Solver** uses Elements and Connections. Each element is either a battery module or a program, and a connection joins elements together in the circuit.

First, you define the battery module and program as elements of the circuit.

2. Right-click the **Batteries** > **Battery Modules** > **Battery Module** node and select **Create Circuit Element**.



3. Right-click the **Program Files** > **Tutorial1CDischarge.prg** node and select **Create Circuit Element**.



Next, the positive and negative terminals of the elements are connected.

4. Expand the **Circuit** > **Circuit Solver** > **Circuit Elements** node.
5. Use `[Ctrl-Click]` to select the **Battery Module** > **(+) Battery Module** node and the **Tutorial1CDischarge.prg** > **(-) Tutorial1CDischarge.prg** node.
6. Right-click one of the nodes and select **Create Connection**.



7. Repeat for the remaining terminals: **(-) Battery Module** and **(+) Tutorial1CDischarge.prg**.

   Each connection is shown under the **Circuit Connections** node.



Preparing a Scalar Scene

Create a scalar scene to display the results of the simulation.

A scalar scene is created to display the volumetric heat of the battery during the simulation.

1. Create a scalar scene.
2. Click **scene/plot**.
3. Right-click **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** and select **Edit**.

   The *Edit* dialog appears.

4. Click (**Column Selection**) until all parts are deselected.
5. Expand the **Regions** node and the **Battery Module: Battery Cell: Stack** node.
6. Select **Battery Module: Battery Cell 0,0: Stack. Stack Front**.



7. Click **OK**.
8. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node.
9. In the *Scalar Field - Properties* window, set *Function* to **Battery > Battery Volumetric Heat**.
10. Click the **simulation** tab to return to the STAR-CCM+ simulation object tree.

Setting Up a Voltage Monitor

Define a voltage monitor to show the voltage of the battery over the course of the simulation.

To set up a voltage monitor:

1. Right-click **Reports**.
2. Select **New Report > Volume Average**.

   A new **Volume Average 1** node appears under the **Reports** node.
3. Rename the **Reports > Volume Average 1** node to `Voltage`.
4. Select the **Reports > Voltage** node.
5. In the Voltage *Properties* window, set **Scalar Field Function** to **Battery > Battery Volts**.
6. Click ... (**Custom Editor**) to the right of the *Parts* property.

   The *Voltage - Parts* dialog appears.
7. Expand the **Regions** node.
8. Click **Battery Module: Battery Cell: Stack**.



9. Click **OK** to close the dialog.
10. Right-click the **Reports > Voltage** node and select **Create Monitor and Plot from Report**.

    A new **Voltage Monitor** node appears under the **Monitors** node and a new **Voltage Monitor Plot** node appears under the **Plots** node.



The **Voltage Monitor** node lets you monitor the voltage of the battery with time.

11. Select the **Monitors** > **Voltage Monitor** node.
12. In the Voltage Monitor - Properties window, make sure that *Trigger* is set to **Time Step**.
13. To open the *Voltage Monitor Plot*, double-click the **Plots** > **Voltage Monitor Plot** node.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

As you are modeling a single cell without a fluid domain, you require few iterations per time-step.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`1.0 s`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	**Enabled**	Activated
Maximum Physical Time	**Maximum Physical Time**	`90 s`
Maximum Inner Iterations	**Maximum Inner Iterations**	`7`





3. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 90 seconds of physical time are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

View the results of the simulation.

To visualize the results:

1. Select *Scalar Scene 1*.



The scalar scene is shown with the mesh turned on. It shows you that the highest concentration of battery heat generation occurs near the tab roots area in the battery cell, as expected.

2. Select the *Voltage Monitor Plot*.



The voltage plot displays the steady discharge of the battery during the first 90 seconds.

Running a STAR-CCM+ Simulation with 3 Cells in Series

Use the battery cell from the previous simulation to model a battery with three cells in series.

Model a battery with three cells in series, also called a string. A tab connector is imported into STAR-CCM+ to model the physical connections between the cells. Complete Part 2 of the tutorial (Running a STAR-CCM+ Simulation with a Single Cell) before attempting this part.

1. Copy the tab connector geometry file `tabConnector.x_b` from the `doc/startutorialdata/battery/data` subdirectory of your STAR-CCM+ installation directory into the `BasicCell` directory.

Clear the solution from the single cell simulation.

2. Select **Solution** > **Clear Solution** in the menu bar.
3. In the *Clear Solution* dialog, click **OK**.
4. Save the simulation.

Contents

Setting Up the Battery Module
Setting Up Tab Connections
Generating the Battery Module
Generating the Volume Mesh
Running the Simulation
Visualizing the Results

Setting Up the Battery Module

Set up the battery module with three cells in series.

To set up the battery module:

Increase the number of cells in the battery module from 1 to 3.

1. Select the **Batteries** > **Battery Modules** > **Battery Module** node.
2. In the *Battery Module - Properties* window, set *N Series* to 3.



3. Expand the **Battery Module** > **Battery Module Cells** node.



Three cells are shown. The cells are named according to their position in the battery module matrix.

To define the position and orientation of the cells in the module:

4. Right-click the **Battery Module Cells** node and select **Setup Battery Module Cells**.



The **edit** tab opens which contains the **Setup Battery Module Cells** panel.

In the *Setup Battery Module Cells* panel, there are two places to specify cell spacing values:

- The *Parallel String Spacing* group box allows you to define the relative position of multiple strings that are connected in parallel.
- The *Cell Spacing* group box allows you to define the relative position of cells that are connected in series within a particular string.

The cells are arranged in a stack formation, with an offset from one cell to the next of 0.01m. Use the default *Offset Vector* in the **Cell Spacing** box to define the position of the 3 cells. The **Parallel String Spacing** is not required.

5. In the **Cell Spacing** group box, make sure that the *Offset Vector* is set to `[0, 0, 0.01]m`.
6. Click **Set**.

At the bottom of the panel, notice that each cell in the table now has a different offset position. In addition to defining the cell spacing, flip the middle cell in the series. This positions the positive tab of the middle cell next to the negative tab of the first cell, which is necessary for the tab connections in this case.

7. Activate **Flip?** for the middle cell.



8. Click **Apply** and then click **Close**.

Setting Up Tab Connections

Set up tab connectors to define the connection between the positive and negative tabs of the battery cells.

A physical connection between the positive and negative tabs of the battery cells can be modeled using tab connectors. For a regular formation of cells, a single tab connector can be used, and STAR-CCM+ uses copies of the connector to join all the tabs in the battery module. A simple tab connector has been prepared for this tutorial.

To set up tab connectors:

1. Select **File** > **Import** > **Import Surface Mesh** from the menu bar.
2. In the *Open* dialog, select the `tabConnector.x_b` file.
3. Click **Open**.
4. In the *Import Surface Options* dialog, do the following:
   a) Set **Import Mode** to **Create new Part**.
   b) Deactivate **Open geometry scene after import** under **Import Options**.
5. Click **OK**.

The imported geometry consists of 3 parts: a positive post, a negative post, and a strap. These parts are added to the **Geometry** > **Parts** node in the STAR-CCM+ tree.



Use the imported tab connector parts in the battery module definition.

6. Right-click the **Batteries** > **Battery Modules** > **Battery Module** > **Tab Connections** node and select **Setup Tab Connections**.



The *Setup Tabs Connections* in-place dialog appears.

7. Set **Connections Setup Method** to **Copy Single Tab Connector**.

The positive tab of the first cell in a string of cells that are connected in series generally becomes the positive terminal for the battery. Similarly, the negative tab on the last cell in a string of cells that are connected in series becomes the negative terminal for the battery. This means that the positive tab on the first cell and the negative tab on the last cell in the series remain open when the tab connectors are added.

The tab connector parts are assigned using the *First Post Parts*, *Strap Parts*, and *Second Post Parts* options.

8. Click to the right of the *First Post Parts* property.
9. In the *Set Connection Geometry Parts* dialog, select **Tab Connector: -ve**.



10. Click **OK**.
11. Set the *Strap Parts* to **Tab Connector: Strap**.
12. Set the *Second Post Parts* to **Tab Connector: +ve**.

    The completed dialog is shown below:



13. Click **Apply**, then click **Close**.

Generating the Battery Module

Generate the battery module.

The three-cell battery is connected to the same electrical load, so no changes to the program file or circuit solver are necessary.

To generate the battery module:

1. To display the 3D-CAD View in the *Graphics* window, select the **3D-CAD View 1** tab.
2. Right-click the **Battery Modules** > **Battery Module** node and select **Generate Battery Parts**.

The battery module contains three cells. Parts have been created for each battery cell in the module.

3. Right-click the **Battery Modules** > **Battery Module** node and select **Generate Connector Parts**.



Parts are created for each tab connection. Before generating regions, imprint the tab connector and tab parts.

4. Expand the **Geometry** > **Parts** node.
5. Select all the **Battery Module** parts by holding the `Shift` key while selecting the first and last part.
6. Right-click one of the highlighted nodes and select **Repair Surface...**.



7. In the *Surface Repair Options* dialog, click **OK** to launch the surface repair tool for the selected parts.
8. In the surface repair *edit* panel, set the **Mode** to **Merge/Imprint**.
9. In the *Merge/Imprint Tool* box, set the **Imprint Mode** to **Multi-Part Imprint**.
10. Set the **Tolerance** to `1.0E-6 m`.
11. In the *Part Browser* box, click **Find Pairs**.



12. Once all the pairs are found, click **Imprint All** and then **Close**.
13. As with the parts and connector parts, right-click the **Batteries** > **Battery Modules** > **Battery Module** node and select **Generate Regions...**.



14. In the *Additional (non-Battery) Parts to Generate Regions* dialog, click (**Column Selection**) to deselect all parts.
15. Click **OK**.



STAR-CCM+ automatically creates interfaces between the tabs and the connectors, based on the imprinted parts.

16. Expand the **Interfaces** node and select each interface in turn.
    The selected interface is highlighted in the scene.

Generating the Volume Mesh

Define the battery mesh and set the appropriate properties.

Use the same mesh settings to generate the mesh for the three-cell battery.

To generate the volume mesh:

1. Click (**Generate Volume Mesh**).
2. Display *Mesh Scene 1* in the *Graphics* window.
3. Click **scene/plot**.
4. Right-click **Mesh Scene 1** > **Displayers** > **Mesh 1** > **Parts** and select **Edit**.
5. In the *Parts* dialog, expand the **Regions** node then click **Select Children** on the left of the **Regions** node.



6. Click **OK**.

The mesh scene is shown below.



Running the Simulation

Modify the simulation and run it again.

Before running the simulation, modify the scalar scene to display the battery volumetric heat on the middle cell of the stack.

You make this modification by assigning geometry parts to the scalar displayer.

1. Show *Scalar Scene 1* in the *Graphics* window.
2. Under **scene/plot**, right-click the **Scalar Scene 1** > **Displayers** > **Scalar 1** > **Parts** node and select **Edit**.
3. Click **Column Selection** to deselect all parts.
4. Expand the **Parts** > **Battery Module: Battery Module Cell 0,1: Stack** node and select **Stack Front**.



5. Click **OK**.

Make sure that the scene is displaying the volume mesh representation.

6. Select the **Scalar Scene 1 > Displayers > Scalar 1** node and make sure that the **Representation** is set to **Volume Mesh**.



7. Save the simulation under the name `batteryTutorial_threeCells.sim`.

You are now ready to run the simulation.

8. Click (**Run**) in the *Solution* toolbar.
9. When the simulation has finished running, save it.

Visualizing the Results

View the results of the simulation.

To view the results of the simulation:

1. Select *Scalar Scene 1*.

   The Scalar Scene is shown below.



The scalar scene demonstrates that battery heat generation does not display any changes when more cells are added in series.

As seen in Part 2 (Running a STAR-CCM+ Simulation with a Single Cell), the highest concentration of battery heat generation is seen near the tab roots area.

Summary

This section outlines the STAR-CCM+ features that are introduced by the Battery Simulation Module: Cell Thermal Analysis tutorial.

This tutorial covered the following STAR-CCM+ features:

- Setting up the battery physics model.
- Importing a battery cell from Battery Design Studio®.
- Loading program data from Battery Design Studio®.

- Generating a battery module by creating parts, connector parts, and regions.
- Generating a volume mesh using the battery mesher.
- Creating a circuit.
- Setting up reports, monitors, and plots.
- Post-processing voltage and battery heat generation per unit volume.

Cylindrical Cells: Cell Thermal Analysis

This tutorial demonstrates how to model a battery module in STAR-CCM+ using cylindrical cells. It simulates three cylindrical cells in series in a staggered configuration with an external casing.

The `tbm` file, `prg` file, a tab connector, and the external casing, are imported into STAR-CCM+.

Note: For battery modeling, STAR-CCM+ requires a `batterysim` license in addition to the other license features.

The three cylindrical cells are connected in series and are discharged for 200 seconds with a load of 22.11 A. The cells are cooled with a draft of 1 m/s. You model both the heat generation in the battery and the air flow in the external casing.



Contents

Prerequisites

Creating the Simulation
Selecting the Physics Models
Importing the Cylindrical Cell
Defining the Cell Configuration
Setting Up Tab Connectors
Generating the Battery Module
Imprinting the Tab Connectors
Defining the External Casing
Generating Regions
Generating the Volume Mesh
Selecting the Physics Models for the External Casing
Creating a Circuit
Setting Solver Parameters and Stopping Criteria
Reporting, Monitoring, and Plotting
Preparing a Scalar Scene
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Cylindrical Cells: Cell Thermal Analysis tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Creating the Simulation

Launch STAR-CCM+ and create a simulation.

To create the simulation:

Create a working directory and copy the files that are required for this tutorial.

1. Create a working directory called `CylindricalCell`.
2. Copy the following files from the `doc/startutorialdata/battery/data` subdirectory of your STAR-CCM+ installation directory into the `CylindricalCell` directory:
   - `TutorialCellCylindrical.tbm`
   - `Tutorial3CDischargeCylindrical.prg`
   - `post1Cylindrical.x_b`
   - `post2Cylindrical.x_b`
   - `strapCylindrical.x_b`
   - `externalCasingCylindrical.x_b`

Create the simulation.

3. Launch STAR-CCM+.
4. Start a simulation.
5. Save the simulation as `batteryTutorialCylindrical.sim` in the `CylindricalCell` directory.

Selecting the Physics Models

Physics models define the primary variables of the simulation and what mathematical formulation is used to generate the solution.

To select the physics models:

1. Create a physics continuum and rename the **Continua** > **Physics 1** node to **Battery Physics**.
2. For the physics continuum, **Continua** > **Battery Physics**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Implicit Unsteady**
Material	**Solid**
Optional Models	**Battery**
Energy	**Segregated Solid Energy**
Equation of State	**Constant Density**

3. Click **Close**.
4. To review the models, open the **Battery Physics** > **Models** node.



The battery pre-processing is carried out primarily under this node.

5. Save the simulation.

Importing the Cylindrical Cell

The battery file created in BDS is imported into STAR-CCM+ as a battery cell.

The battery file from Battery Design Studio®, `TutorialCellCylindrical.tbm`, is an idealized cylindrical cell. It uses the NTG3D model.

To import the cylindrical cell:

1. Expand the **Batteries** node.
2. Right-click **Batteries** > **Battery Cells** and select **Create from Tbm...**.



3. In the *Import Battery Data From File* dialog, navigate to your working directory, `CylindricalCell`, and select the file `TutorialCellCylindrical.tbm`.
4. Click **Open**.
5. In the *Import Battery Options* dialog, ensure that all objects are selected, then click **OK**.



During the import process, STAR-CCM+ builds an idealized model of the cell using the 3D-CAD module. The battery cell model appears in the *3D-CAD View 1* window.



Defining the Cell Configuration

Define the configuration of the cells within the battery module.

Configure the cylindrical battery cells using a staggered configuration with uniform spacing.

To define the cell configuration:

Create a battery module:

1. Right-click the **Batteries** > **Battery Modules** node and select **New**.
2. Select the **Batteries** > **Battery Modules** > **Battery Module** > **Module Configuration** node.
3. In the *Module Configuration - Properties* window to the right of **Battery Cells**, click ... (**Custom Editor**).
4. In the *Module Configuration - Battery Cells* dialog, select **Cylindrical Cell**.



5. Click **OK**.

Once you specify the battery cell to use for module configuration, STAR-CCM+ adds sub-nodes to the **Module Configuration** manager. These sub-nodes allow you to complete the setup.

6. Select the **Module Configuration > Battery Module Configuration** node.
7. In the *Battery Module Configuration - Properties* window, set *Option* to **Staggered**.



By default, the cell orientation is set to **Flip Every Other**, with the even cells flipped. Set the **Series Connection Flip Option**:

8. Select the **Module Configuration** > **Module Connection Values** > **Series Connection Flip Option** node.
9. In the *Series Connection Flip Option - Properties* window, make sure that *Option* is set to **Flip Starting With Second Cell**.

To set the number of cells in the module and the spacing between cells:

10. Select the **Module Configuration** > **Module Layout** > **Module Staggered Positioning** node.
11. In the *Module Staggered Positioning - Properties* window, set *Gap Between Cells* to `0.004 m`.



12. Select the **Module Configuration** > **Module Layout** > **Module Staggered Size** node and set *Number of Cells in Staggered Row* to `3`.

Setting Up Tab Connectors

Import and configure the tab connectors.

A tab connector has been prepared for this staggered case.

To import the tab connector into the simulation:

1. Select **File** > **Import** > **Import Surface Mesh** from the menu bar.
2. In the *Open* dialog, multi-select the `post1Cylindrical.x_b`, `post2Cylindrical.x_b`, and `strapCylindrical.x_b` files, and click **Open**.
3. In the *Import Surface Options* dialog, do the following:
   a) Make sure that the **Import Mode** is set to **Create New Part**.
   b) Deactivate **Open Geometry Scene After Import** under **Import Options**.
4. Click **OK**.

The imported tab connector includes: a strap, a positive post, and a negative post.

5. Right-click the **Batteries** > **Battery Modules** > **Battery Module** > **Tab Connections** node and select **Setup Tab Connections**.

   The **Setup Tab Connections Edit** tab appears.

The first post part is the negative post that connects the negative tab of the first battery cell to the positive tab of the second cell. The positive tab in the first battery cell always stays open. The tab connector is copied to connect the tabs of the second and third cell.

6. In the **Setup Tab Connections Edit** tab, click to the right of the *First Post Parts* property.
7. In the *Set Connection Geometry Parts* dialog, select **Post 1**.
8. Click **OK**.
9. Repeat the previous steps and assign the **Strap** part to **Strap Parts**, and the **Post 2** part to **Second Post Parts**.



10. Click **Apply** then **Close**.
11. Save the simulation.

Generating the Battery Module

Generate the battery module parts.

Generate the battery module parts that you have configured in the previous steps.

To generate the battery module:

1. Right-click the **Batteries** > **Battery Modules** > **Battery Module** node and select **Generate Battery Parts**.



The following messages are printed in the *Output* window:

```
Generating Battery Module Cell Parts 0,0

Generating Battery Module Cell Parts 0,1

Generating Battery Module Cell Parts 0,2
```

2. Right-click the **Battery Modules** > **Battery Module** node and select **Generate Connector Parts**.
3. In the **Vis** tool bar, click the **Create / Open Scenes** button and select **Geometry**.

STAR-CCM+ creates two tab connectors on the opposite sides of the cells, as shown in the example below.



Imprinting the Tab Connectors

To ensure that all the interfaces are properly formed when you generate the battery module regions, perform an imprint operation between the battery cells and the tab connectors.

To imprint the tab connectors and the battery cells:

1. Click (**Launch Surface Repair**).
2. In the *Surface Repair Options* dialog next to **Geometry Parts**, click **Select**.

   The *Select Objects* dialog appears.
3. Click (**Column Selection**) until all parts are deselected.
4. Select the following:
   - **Battery Module: Battery Module Cell 0,0: +Ve EndPlate**
   - **Battery Module: Battery Module Cell 0,0: -Ve EndPlate**
   - **Battery Module: Battery Module Cell 0,1: +Ve EndPlate**

- **Battery Module: Battery Module Cell 0,1: -Ve EndPlate**
- **Battery Module: Battery Module Cell 0,2: +Ve EndPlate**
- **Battery Module: Battery Module Cell 0,2: -Ve EndPlate**
- **Battery Module: Tab Connection 0,0: Post 1**
- **Battery Module: Tab Connection 0,0: Post 2**
- **Battery Module: Tab Connection 0,1: Post 1**
- **Battery Module: Tab Connection 0,1: Post 2**
- **Battery Module: Tab Connection 0,0: Strap**
- **Battery Module: Tab Connection 0,1: Strap**

5. Click the **Show All** drop down menu and select **Selected**. Ensure your selection appears as shown below.



6. Click **OK**.
7. Click **OK** in the *Surface Repair Options* dialog.
8. In the surface repair *Edit* panel, set the mode to **Merge/Imprint**.



9. Set the **Imprint Mode** to **Multi-Part Imprint**.
10. Set the **Tolerance** to `1.0E-6 m`.
11. Click **Priority Indexing**.
    The *Merge/Imprint Priority Indexing* dialog appears.
12. Click **Create New Set**.
13. Click (**Custom Editor**) for Priority 1.
14. In the *Select Parts* dialog, click **Column Selection** to select all parts.
15. Deselect the following:
    - **Post 1**
    - **Post 2**
    - **Strap**
16. Click **OK**.
17. Click **Create New Set**.
18. Click (**Custom Editor**) for Priority 2.
19. In the *Select Parts* dialog, select the following:
    - **Battery Module: Tab Connection 0,0: Strap**
    - **Battery Module: Tab Connection 0,1: Strap**
20. Click **OK**.
21. Click **Create New Set**.
22. Click (**Custom Editor**) for Priority 3.
23. In the *Select Parts* dialog, select the following:
    - **Battery Module: Tab Connection 0,0: Post 1**
    - **Battery Module: Tab Connection 0,0: Post 2**
    - **Battery Module: Tab Connection 0,1: Post 1**

- **Battery Module: Tab Connection 0,1: Post 2**

24. Click **OK**.



25. To exit the *Merge/Imprint Priority Indexing* dialog, click **OK**.
26. In the **Part Browser** box, click **Find Pairs**.
27. When all the pairs are found, click **Imprint All**.



28. When all the pairs are imprinted, **Close** the *Edit* panel.
29. Save the simulation.

Defining the External Casing

Define the external casing for the battery module.

A simple external casing part is provided with your STAR-CCM+ installation.

To define the external casing:

Import the geometry into the simulation.

1. Select **File > Import > Import Surface Mesh...** from the menu bar.
2. In the *Open* dialog, select the `externalCasingCylindrical.x_b` file, and click **Open**.
3. In the *Import Surface Options* dialog, do the following:
   a) Make sure that the **Import Mode** is set to **Create New Part**.
   b) Set **Mark Feature Edges** to **All CAD Edges**.
   c) Deactivate **Open Geometry Scene After Import** under **Import Options**.
4. Click **OK**.

Create an inlet and outlet for the external casing.

5. Right-click the **Geometry > Parts > Case > Surfaces > Faces** node and select **Split By Patch...**.

Note: You must have a scene open to use this feature.

The *Split Part Surface by Patch* panel appears.

6. In the **Part Surface Patches** box, select `3280`.
7. Set the **Part Surface Name** to `Outlet`.
8. Click **Create**.
9. Similarly, select `3278` and rename this part `Inlet`.
10. Click **Create**.
11. Click **Close**.

Select the external casing part.

12. Select the **Batteries > Battery Modules > Battery Module > External Casing Geometry > External Casing Specification** node.
13. In the *External Casing Specification - Properties* window, set *Battery External Casing Specification* to **Specify External Casing Using Part**.



14. Select the **Battery Module > Casing Values > External Casing Part** node.
15. In the *External Casing Part - Properties* window, in the **External Casing Part** property, click ... (**Custom Editor**).
16. In the *External Casing Part* dialog, select **Case**.



17. Click **OK**.

You can now generate the external casing part.

18. Right-click the **Battery Modules** > **Battery Module** node and select **Create External Casing Part**.



STAR-CCM+ creates a **Battery Module: External Casing** part. The other battery module parts have been subtracted from it.

19. Save the simulation.

Generating Regions

Create the regions of the battery module.

To create the regions of the battery module:

1. Right-click the **Battery Modules** > **Battery Module** node and select **Generate Regions...**.
2. In the *Additional (non-Battery) Parts to Generate Regions* dialog, click (**Column Selection**) twice to deselect all parts.
3. Click **OK**.

   STAR-CCM+ automatically creates regions and interfaces for the battery module.



Generating the Volume Mesh

Create the volume mesh continua for the battery module.

Create two separate mesh continua—one for the battery cells and the tab connectors, then one for the external casing.

To generate the volume mesh:

Set up the battery mesh continuum.

1. Right-click the **Batteries** > **Battery Modules** > **Battery Module** node.
2. Select **Setup Battery Meshing**.



A new **Mesh 1** node appears under the **Continua** node. This new battery mesh continuum already contains the **Surface Remesher** and **Thin Mesher** models. The Battery Mesh Continuum also contains a volumetric control for the strap and post parts.

3. Rename the new **Mesh 1** node to `Battery Mesh`.

To confirm the **Battery Mesh** continuum contains the correct meshing models:

4. Right-click the **Continua** > **Battery Mesh** > **Models** node and select **Select Meshing Models**.

The *Battery Mesh Model Selection* dialog appears as shown below.



5. Ensure the **Surface Remesher** and **Thin Mesher** models are selected.
6. Click **Close**.

To refine the volume mesh for the tab connectors:

7. Expand the **Battery Mesh > Volumetric Controls > Battery Cell Tab Volumetric Control 1** node.
8. Select the **Mesh Values > Custom Size > Relative Size** node.
9. In the *Properties* window, set *Percentage of Base* to `5.0`.



To create a second mesh continuum for the external casing:

10. Right-click the **Continua** node and select **New > Mesh Continuum**.
11. Rename the new **Mesh 1** node to `External Casing Mesh`.

To select the models for the external casing mesh:

12. Right-click the **Continua > External Casing Mesh** node and select **Select Meshing Models**.
13. In the *External Casing Mesh Model Selection* dialog, select the following:
    - **Surface Remesher** from the *Surface Mesh* group box.
    - **Polyhedral Mesher** from the *Volume Mesh* group box.



14. Click **Close**.
15. Select the **Continua > External Casing Mesh > Reference Values > Base Size** node and set *Value* to `0.005 m`.

To assign the external casing mesh continuum to the external casing region:

16. Select the **Regions > Battery Module: External Casing** node.
17. In the *Battery Module: External Casing - Properties* window, set the *Mesh Continuum* to **External Casing Mesh**.

To build the volume mesh:

18. Click (**Generate Volume Mesh**) in the tool bar.

Visualize the mesh.

19. In the **Vis** tool bar, click the **Create / Open Scenes** button and select **Mesh**.

20. Save the simulation.

Selecting the Physics Models for the External Casing

Physics models define the primary variables of the simulation and what mathematical formulation is used to generate the solution.

The external casing is modeled as an air continuum.

To select the physics models for the external casing:

1. Create a physics continuum and rename the **Continua** > **Physics 1** node to **External Casing**.
2. For the physics continuum, **Continua** > **External Casing**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Segregated Flow Gradients** (selected automatically)
Equation of State	**Constant Density**
Viscous Regime	**Turbulent Reynolds-Averaged Navier-Stokes** (selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
	Realizable K-Epsilon Two-Layer (Selected automatically)
	Two-Layer All y+ Wall Treatment (Selected automatically)
Optional Models	**Segregated Fluid Temperature**

3. Click **Close**.
4. To review the models, open the **External Casing** > **Models** node.



5. Select the **Regions** > **Battery Module: External Casing** node and set **Physics Continuum** to **External Casing**.
6. Select the **Regions** > **Battery Module: External Casing** > **Boundaries** node and set the following properties:

Node	Property	Setting
Battery Module: External Casing.Case.Inlet	*Type*	**Velocity Inlet**
Battery Module: External Casing.Case.Outlet	*Type*	**Pressure Outlet**

7. Save the simulation.

Creating a Circuit

Connect the electrical load to the battery module to create a circuit.

A program file containing the electrical load on the battery is provided.

To create a circuit:

1. Right-click the **Batteries** > **Program Files** node and select **Load...**.
2. In the *Open* dialog, select the `Tutorial3CDischargeCylindrical.prg` file.
3. Click **Open**.

The program file that you loaded in the previous step represents an electrical load on the battery. To simulate the response of the battery, the electrical load must be connected to the battery in a circuit. Connect the electrical load using the **Circuit Solver** tool.

4. Expand the **Circuit** > **Circuit Solver** node.

The **Circuit Solver** uses Elements and Connections. Each element is either a battery module or a program, and a connection joins elements together in the circuit.

First, you define the battery module and program as elements of the circuit.

5. Right-click the **Battery Modules** > **Battery Module** node and select **Create Circuit Element**.



6. Right-click the **Program Files** > **Tutorial3CDischargeCylindrical.prg** node and select **Create Circuit Element**.



Next, the positive and negative terminals of the elements are connected.

7. Expand the **Circuit** > **Circuit Solver** > **Circuit Elements** node.
8. Multi-select the **Battery Module** > **(+) Battery Module** node and the **Tutorial3CDischargeCylindrical.prg** > **(-) Tutorial3CDischarge.prg** node.
9. Right-click one of the nodes and select **Create Connection**.



10. Repeat for the remaining terminals: **(-) Battery Module** and **(+) Tutorial3CDischargeCylindrical.prg**.

    Each connection is shown under the **Circuit Connections** node.



To model the effect of ohmic heating, activate the ohmic heating models.

11. Right-click **Batteries** > **Battery Modules** > **Battery Module** and then select **Setup Connector Ohmic Heating**.

    The **Continua** > **Battery Module Connectors: Ohmic Heating** physics continuum node is added with the appropriate models activated.
12. Save the simulation.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

Run the simulation for 200 seconds, after which the state of charge of the battery is around 80%.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`1.0 s`

2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	*Enabled*	Activated
	Maximum Physical Time	`200 s`
Maximum Inner Iterations	*Enabled*	Activated
	Maximum Inner Iterations	`5`

Node	Property	Setting
Maximum Steps	*Enabled*	Deactivated

3. Save the simulation.

Reporting, Monitoring, and Plotting

Create reports, monitors, and plots, to visualize the results of the simulation.

Create maximum and minimum temperature reports, monitors, and plots to visualize the temperature rise of the cells.

To create reports, monitors, and plots:

Start with the maximum temperature report.

1. Right-click the **Reports** node and select **New Report** > **Maximum**.
   A new report, **Maximum 1**, is created.
2. Rename the **Maximum 1** node to `Maximum Temperature`.
3. Select the **Reports** > **Maximum Temperature** node and set *Scalar Field Function* to **Temperature**.
4. Set *Units* to **C**.
5. Click ... (**Custom Editor**) for the *Parts* property.
   The *Maximum Temperature - Parts* dialog appears.
6. Expand the **Regions** node.
7. Click **Battery Module: Cylindrical Cell: Jellyroll** twice.



8. Click **OK**.

   The *Maximum Temperature Properties* window appears as shown:



Set up the minimum temperature report.

9. Right-click the **Reports** node and select **New Report > Minimum**.

10. Rename the **Reports** > **Minimum 1** node to `Minimum Temperature`.
11. Select the **Reports** > **Minimum Temperature** node and set *Scalar Field Function* to **Temperature**.
12. Set *Units* to **C**.
13. Click (**Custom Editor**) for the *Parts* property.

    The *Minimum Temperature - Parts* dialog appears.
14. Expand the **Regions node**.
15. Click **Battery Module: Cylindrical Cell: Jellyroll** twice.
16. Click **OK**.

    The *Minimum Temperature Properties* window appears as shown:



Create a monitor and plot with the temperature reports:

17. Multi-select the **Reports** > **Maximum Temperature** and **Reports** > **Minimum Temperature** nodes.
18. Right-click one of the selected nodes and select **Create Monitor and Plot from Report**.
19. In the *Create Plot From Reports...* dialog, click **Single Plot**.

    The **Maximum Temperature Monitor**, **Minimum Temperature Monitor**, and **Reports Plot** nodes appear in the simulation tree.



20. Rename **Plots > Reports Plot** to `Temperature Monitor Plot`.
21. Right-click the **Plots > Temperature Monitor Plot** node and select **Open**.

    The *Temperature Monitor Plot* window opens.
22. Select the **Plots > Temperature Monitor Plot > Axes > Y Axis > Title** node.
23. In the *Title - Properties* window, set the *Title* to `Temperature (C)`.



Create a state of charge report, monitor, and plot to visualize the discharging of the cells.

24. Right-click the **Reports** node and select **New Report > Volume Average**.

    A new report, **Volume Average 1**, is created.
25. Rename the **Volume Average 1** node to `State of Charge`.
26. Select the **Reports > State of Charge** node and set *Scalar Field Function* to **Battery > Battery SOC**.
27. Click ... (**Custom Editor**) for the *Parts* property.

    The *State of Charge-Parts* dialog appears.
28. Expand the **Regions** node.
29. Select **Battery Module: Cylindrical Cell: Jellyroll**.



30. Click **OK**.

    The *State of Charge Properties* window appears as shown:



Create a monitor and plot with the state of charge report:

31. Right-click the **Reports** > **State of Charge** node and select **Create Monitor and Plot from Report**.
32. Right-click the **Plots** > **State of Charge Monitor Plot** node and select **Open**.
33. Save the simulation.

Preparing a Scalar Scene

Create a scalar scene to display the results of the simulation.

You create a scalar scene to display the temperature of the battery cans.

To prepare a scalar scene:

1. In the **Vis** tool bar, click the **Create / Open Scenes** button and select **Scalar**.
2. Rename the **Scenes** > **Scalar Scene 1** node to `Temperature Scene`.
3. Right-click the **Temperature Scene** > **Displayers** > **Scalar 1** > **Parts** node and select **Edit...**.
4. In the *Parts* dialog, expand the **Regions** node.
5. Select **Battery Module: Cylindrical Cell: Jellyroll**.



6. Click **OK**.
7. Select the **Temperature Scene > Displayers > Scalar 1 > Scalar Field** node.
8. In the *Scalar Field - Properties* window, set **Function** to **Temperature**.



9. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To run the simulation:

1. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

If left alone, the simulation continues until 200 seconds of physical time are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

```
Stopping criterion Maximum Physical Time satisfied.
```

2. When the simulation has finished running, save it.

Visualizing the Results

Visualize the effect of the draft in the external casing.

The scalar scene of temperature appears as follows.



Create a scalar scene to visualize the effect of the draft in the external casing.

To visualize the results:

1. Right-click the **Scenes** > **Temperature Scene** node and click **Copy**.
2. Right-click the **Scenes** node and click **Paste**.

   A new scene, **Copy of Temperature Scene**, appears in the simulation tree.
3. Rename **Copy of Temperature Scene** to `Temperature Plane Section Scene`.
4. Open the *Temperature Plane Section Scene*.
5. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane...**.

   The *Create Plane Section* panel appears.
6. Set **normal** to `[0 m, 1 m, 0 m]` as shown below.



7. Set **Display** to **No Displayer**.
8. Click **Create**, and then **Close**.
9. Select **Scenes** > **Temperature Plane Section Scene** > **Displayers** > **Scalar 1** > **Parts** and click ... (**Custom Editor**) for the *Parts* property.

   The *Parts* dialog appears.
10. Click the **Column selection** button twice to deselect all parts.
11. Expand the **Derived Parts** node and select **Plane Section**.



12. Click **OK**.

As expected, the temperature is lower near the inlet.



The battery state of charge has dropped to 80% as expected.



The temperature of the cells has risen by 5 C during the discharging process.



Summary

This tutorial demonstrated how to model a battery module in STAR-CCM+ using cylindrical cells.

The steps covered were:

- Setting up the battery physics model.
- Imprinting parts.
- Generating a battery module by creating parts, connector parts, and regions.
- Generating a volume mesh using the battery mesher.
- Creating a circuit.

Li-Ion Battery Cell Model: Cell Electrochemistry Analysis

In this tutorial, you model a Li-Ion battery cell at the microscale level. The simulation introduces the method for the prediction of the three-dimensional microstructure of porous battery electrodes.

In a production environment, such analysis results in speeding up the design process of a battery cell, and assessing the safety and life-span of a design.



Generally, the structure of the active materials is irregular and complex, and varies from one active material to another. The structure of the solid particles is often represented by a simpler, regular structure. In this tutorial, you are given randomly distributed spherical particles that represent a single structure, and reflect the desired porosity.

In this tutorial, you model the following:

- electric potential in active materials, electrolytes, and current collectors
- concentration of lithium in active materials and electrolytes
- special conditions on the SEI (Solid Electrolyte Interface)

Note: To activate the Li-Ion battery cell, concentration, and electric potential models that are necessary for modeling the transport of Lithium ions, run a double-precision version of STAR-CCM+.

The geometry consists of a sample across a pair of electrodes from a typical Li-ion cell. The included parts of the geometry are:

- a cathode collector—10 m thick aluminum foil
- an anode collector—10 m thick copper foil
- electrolytes—a 50:50 mixture of ethylene carbonate/ethyl methyl carbonate, and $LiFP_6$ (lithium hexafluorophosphate—a common electrolyte salt in lithium batteries)
- a separator—10 m thick, with a porosity of 40%, and MacMullin number of 2.5.

  The MacMullin number is the ratio of the resistance of the separator when it is filled with electrolyte to the resistance of the electrolyte alone. This number describes the tortuosity when defining a homogeneous porous medium.
- a cathode AM (Active Material)—35 m thick $LiMn_2O_4$ (lithium manganese oxide), with particle diameter equal to 18 m, and a target porosity of 40%
- an anode AM—35 m thick graphite, with particle diameter equal to 20 m, and a target porosity of 40%

Each electrode is formed from a solid active material and the surrounding fluid electrolyte. These regions within an electrode are modeled separately. Chemical reactions occur at the SEI (Solid-Electrolyte Interfaces).



Note: The tutorial uses a simplified idealized microstructure of a $LiMn_2O_4$ electrode. The electrode structure is only represented by a few aggregated spherical particles and the electrolyte. The voluntary omission of the binders and conductivity aids is to speed up the computation for this tutorial. This case does not represent a commercially used electrode. However, it is relevant enough to compute a consistent solution and demonstrate the capabilities of the model.

Contents

Prerequisites
Importing the Geometry and Naming the Simulation
Generating the Mesh
Setting Up the Physics Models
Assigning Physics Continua to Regions
Setting the Material Properties
Setting Initial Conditions
Setting Boundary and Interface Conditions
Setting Solver Parameters and Stopping Criteria
Preparing Scalar and Vector Scenes
Initializing the Solution and Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Li-Ion Battery Cell Model: Cell Electrochemistry Analysis tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Importing the Geometry and Naming the Simulation

Launch STAR-CCM+, import the supplied geometry, and save the simulation.

For this tutorial, you load a .sim file that contains the following objects:

- Imprinted geometry parts: in order to create a matching (or conformal) mesh across an interface, the STAR-CCM+ meshers require that geometry faces on both sides of the interface are identical.
- Regions and interfaces.
- Boundaries with the type set correctly: as you are modeling only a section of a single layer cell, symmetry boundary conditions are specified on all outer walls of the cell cross-section.
- Three material property tables: one table contains the electrical and concentration properties of the electrolytes. Two tables contain data for the relationship between the equilibrium potentials and stoichiometries of the electrodes.

To import the geometry and name the simulation:

1. Launch a double precision version of STAR-CCM+.
2. Select **File** > **Load Simulation** in the menu bar.
3. In the *Open* dialog, choose `doc/startutorialsdata/multiphaseFlow/data/3DEChemGeometry.sim` then click **OK**.

   STAR-CCM+ provides feedback on the import process, which takes a few seconds, in the *Output* window.
4. Save the simulation as `3DEChemTutorial.sim`.
5. Create a geometry scene.

The geometry with transparency turned on is shown below.



The active material spheres are represented by a single part for the anode, and a single part for the cathode. For each electrode, the ratio of the solid volume to the total volume is equivalent to the porosity.

Note: The imported geometry is based on a meter length scale instead of micrometer, which is the typical order of magnitude for these simulations. You scale the mesh in a later part of the tutorial.

Note: When using 3D-CAD to construct the geometry of a Li-Ion battery cell section, use the meter length scale. Due to the precision of the CAD tool (1.0E-8 m), you can experience difficulties while performing boolean operations and manipulating point coordinates if you use the micrometer scale.

Generating the Mesh

Use the polyhedral mesher to generate the volume mesh in all regions of the model.

As concentration gradients are highest across the SEI, you define prism layers to grow on either side of the interface. To ensure Li-ion/salt

conservation on the interfaces between the electrodes and the separator, it is also best to designate a conformal mesh at these interfaces.

To generate the mesh:

1. Create two mesh continua then rename them `AM Electrolyte Separator` and `Collector`.

   As the **AM Electrolyte Separator** mesh continuum is created first for this simulation, it is assigned to all regions by default.
2. Multi-Select the **Anode Collector** and **Cathode Collector** regions, and set *Mesh Continuum* to **Collector**.



Set the models for the **AM Electrolyte Separator** mesh continuum.

3. Right-click the **Continua > AM Electrolyte Separator** node and choose **Select Meshing models...**.
4. Select the following meshing models:
   - **Surface Remesher** from the *Surface Mesh* group box.
   - **Polyhedral Mesher** from the *Volume Mesh* group box.
   - **Prism Layer Mesher** from the *Optional Models* group box.
5. **Close** the dialog.

The complete list of selected meshing models is shown below.



6. Select the **AM Electrolyte Separator > Reference Values > Base Size** node and set *Value* to `0.25 m`.
7. Select the **AM Electrolyte Separator > Reference Values > Number of Prism Layers** node and set *Number of Prism Layers* to `10`.



Set the models for the **Collector** mesh continuum:

8. Right-click the **Continua** > **Collector** node and choose **Select Meshing models...**.
9. Select the following meshing models:
   - **Surface Remesher** from the *Surface Mesh* group box.
   - **Polyhedral Mesher** from the *Volume Mesh* group box.
10. **Close** the dialog.
11. Select **Collector** > **Reference Values** > **Base Size** and set *Value* to `0.1 m`.

Prism layers are required only at the SEI. Deactivate prism layers on all other boundaries, and activate them on the SEI:

12. Expand the **Regions** > **Anode AM** > **Boundaries** > **Anode AM/ Anode Collector** node.
13. Select the **Mesh Conditions** > **Customize Prism Mesh** node and set *Customize Prism Mesh* to **Disable**



14. Similarly, **Disable** the **Customize Prism Mesh** on these region boundaries:
    - **Anode Electrolyte > Boundaries > Anode Electrolyte/Anode Collector**
    - **Cathode AM > Boundaries > Cathode AM/Cathode Collector**
    - **Cathode Electrolyte > Boundaries > Cathode Electrolyte/ Cathode Collector**

Prism layers do not grow from interfaces by default. Turn on this option for the following interfaces:

15. Multi-select the following interfaces:
    - **Anode AM/Anode Electrolyte**
    - **Anode AM/Separator**
    - **Anode Electrolyte/Separator**
    - **Cathode AM/Cathode Electrolyte**
    - **Cathode AM/Separator**
    - **Cathode Electrolyte/Separator**
16. Right-click any selected node and select **Edit...**.



The *Multiple Objects* dialog appears.

17. Expand the **Mesh Conditions > Interface Prism Layer Option** node and activate *Grow Prisms From Interface*.



18. Click **Close**.
    All other properties for each model are left at their defaults.

You are ready to generate the volume mesh.

19. Click (**Generate Volume Mesh**) in the toolbar.
20. Select **Mesh > Scale Mesh...** from the menu bar.
21. In the *Scale Mesh* dialog, set the **Scale Factor** to `1E-5`.
22. Click **Apply** to scale down the mesh.



23. Click **Close**.
24. Create a Mesh Scene.



25. Save the simulation.

Setting Up the Physics Models

Physics models define the primary variables of the simulation and what mathematical formulation is used to generate the solution.

Physics continua are required for the following:

- anode and cathode active materials
- anode and cathode collectors
- electrolyte regions and separator, represented by one fluid continuum, with additional porosity defined for the separator.

STAR-CCM+ creates a physics continuum, **Physics 1**, automatically when you generate a mesh.

To select the physics models:

1. Rename the **Continua** > **Physics 1** node to **Anode AM Physics**.
2. For the physics continuum, **Continua** > **Anode AM Physics**, select the following models in order:

Group Box	Model
Time	**Implicit Unsteady**
Material	**Solid**
Optional Models	**Segregated Solid Energy**
Equation of State	**Constant Density**
Optional Models	**Cell Quality Remediation**
	Li-Ion Battery Cell
	Li-Ion Concentration
	Li-Ion Electric Potential

3. Click **Close**.
4. To review the models, open the **Anode AM Physics** > **Models** node.



To create physics continua for the other solid materials, copy the **Anode AM Physics** continuum:

5. Right-click the **Continua** > **Anode AM Physics** node and select **Copy**.
6. Right-click the **Continua** node and select **Paste**.

   STAR-CCM+ adds a **Copy of Anode AM Physics** node to the feature tree.

7. Repeat the paste process two more times.
8. Rename the physics continua nodes as follows:
   - **Copy of Anode AM Physics** to `Anode Collector Physics`
   - **Copy of Anode AM Physics 2** to `Cathode AM Physics`
   - **Copy of Anode AM Physics 3** to `Cathode Collector Physics`



Next, you create a fluid physics continuum for the electrolytes and separator:

9. Create a physics continuum.
10. Rename the **Continua** > **Physics 2** node to **Electrolyte Physics**.
11. For the physics continuum, **Continua** > **Electrolyte Physics**, select the following models in order:

Group Box	Model
Space	**Three-Dimensional**
Time	**Implicit Unsteady**
Material	**Liquid**
Flow	**Segregated Flow**
Viscous Regime	**Laminar**
Equation of State	**Constant Density**
Optional Models	**Segregated Fluid Temperature**
	Cell Quality Remediation
	Li-Ion Battery Cell
	Li-Ion Concentration
	Li-Ion Electric Potential

12. Click **Close**.

Assigning Physics Continua to Regions

Assign each physics continuum to its appropriate region.

1. Assign the physics continua to the appropriate regions, as follows:

Physics Continuum	Region
Anode AM Physics	**Anode AM**
Anode Collector Physics	**Anode Collector**
Cathode AM Physics	**Cathode AM**
Cathode Collector Physics	**Cathode Collector**
Electrolyte Physics	**Anode Electrolyte**
	Cathode Electrolyte
	Separator



The separator region acts as a porous region rather than a fluid one between the electrodes, with 40% porosity:

2. Select the **Regions > Separator** node and set *Type* to **Porous Region**.
3. Select the **Regions > Separator > Physics Values > Porosity > Constant** node and set *Value* to `0.4`.



The physics continua of the collectors do not use the Li-Ion concentration model:

4. For the **Anode Collector Physics** and **Cathode Collector Physics** continua, do the following:
   a) Right-click the node and choose **Select models...**.

b) Deactivate the **Li-Ion Concentration** model from the *Enabled Models* group box.

c) Click **Close**.

5. Save the simulation.

Setting the Material Properties

Define the material properties for each continuum.

To set the material properties:

Set the anode active material and current collector to graphite and copper, respectively.

1. Right-click the **Anode AM Physics** > **Models** > **Solid** > **Al** node and select **Replace with**.



2. In the *Replace Material* dialog, activate **C (Graphite)** and then click **OK**.
3. Repeat the process for the **Anode Collector Physics** continuum, but replace the material with **Cu (Copper)**.



Set the material properties for the cathode active material.

4. Expand the **Continua** > **Cathode AM Physics** > **Models** > **Solid** node.
5. Rename the **Al** node to **LiMn2O4**.
6. Select the **LiMn2O4** > **Material Properties** > **Electrical Conductivity** > **Constant** node and set *Value* to `0.01 S/m`.



The Lithium/Salt Diffusivity for LiMn2O4 is similar to the value for aluminum: therefore, the default value is fine. The other material properties are not important for the accuracy of this tutorial and are left unaltered.

The cathode collector remains as aluminum.

You treat the electrolyte as a mixture of LiPF6 salt in an equal mixture of ethylene carbonate and ethyl methyl carbonate.

Set the material properties for the electrical conductivity, Li+/salt diffusivity, and transference number using a material table. The simulation file contains this table. Each material property is interpolated using two independent variables: temperature, and salt concentration.

7. Rename **Continua > Electrolyte Physics > Models > Liquid > H2O** to **LiPF6**.
8. Select the **LiPF6 > Material Properties > Electrical Conductivity** node and set *Method* to **Tabular**.



9. Select the **Electrical Conductivity** > **Tabular** > **Table** node and set *Input table* to **electrolyteLiPF6**.
10. Select the **Table** > **Interpolated** node and set *Column title* to **Spec.Cond**.
11. Right-click the **Table** > **Variables** node and select **New**.



12. Select the **Variables > Variable 1** node and set **Scalar Function** to **Temperature**.

    The **Variable 1** node is automatically renamed to **Temperature**.
13. Set *Column title* to **Temp**.



14. Create one more variable, then set the *Scalar Function* to **Lithium/Salt Concentration**, and the *Column Title* to **c2**—the salt concentration in the liquid phase.

Set the tabular material property for the **Lithium/Salt Diffusivity** and **Lithium Cation Transference Number** material properties.

15. Repeat the previous steps to set the tabular material property for the **Lithium/Salt Diffusivity** and **Lithium Cation Transference Number** material properties.

    The input table for both properties is also **electrolyteLiPF6**, and the first independent variable is **Temperature**.

    Note the following differences:

    - For **Lithium/Salt Diffusivity**, set *Column title* of the **Interpolated** node to **Diff.Coef**.
    - For **Lithium Cation Transference Number**, set *Column title* of the **Interpolated** node to **t+**.

    The remaining material properties remain at their default values.



Setting Initial Conditions

Set the initial conditions for the physics continua.

The initial conditions of the active materials represent a 50% utilized electrode.

To set the initial conditions:

Set the initial conditions for the **Anode AM Physics** continuum.

1. Edit the **Initial Conditions** node and set the following property:

Node	Property	Setting
Lithium/Salt Concentration > Constant	*Value*	`13 kmol/m^3`



Set the initial conditions for the **Cathode AM Physics** continuum.

2. Edit the **Initial Conditions** node and set the following property:

Node	Property	Setting
Lithium/Salt Concentration > Constant	*Value*	`13 kmol/m^3`

Set the initial conditions for the **Electrolyte Physics** continuum.

3. Edit the **Initial Conditions** node and set the following property:

Node	Property	Setting
Lithium/Salt Concentration > Constant	*Value*	`1 kmol/m^3`

4. Save the simulation.

Setting Boundary and Interface Conditions

Define the boundary conditions at the free surfaces of the two current collectors.

To simulate an approximate discharge current of 10 A/m^2 across the cell segment, you define the current density at the anode collector, and a zero electric potential at the end of the cathode collector.



The discharge process is an electrochemical process: under an applied electric current, the lithium that is inserted in the anode active material is dissociated into a Li^+ ion and an electron. The Li^+ ion then migrates through the separator and is inserted into the cathode active material. Hence, the concentration of Li^+ ion decreases on the anode side and increases on the cathode side over time.

You specify the current density, in A/m^2, at the anode since both the current and the anode areas are known. To allow a null net flux, you set the electrical potential at the outlet to 0 V.

To set the boundary and interface conditions:

1. Expand the **Regions** > **Anode Collector** > **Boundaries** > **Free Surface** node.
2. Select the **Physics Conditions** > **Electric Potential Specification** node and make sure that *Method* is set to **Specific Electric Current**.



3. Select the **Physics Values > Specific Electric Current > Constant** node and set *Value* to `10 A/m^2`.
4. Expand the **Regions > Cathode Collector > Boundaries > Free Surface** node.
5. Select the **Physics Conditions > Electric Potential Specification** node and set *Method* to **Electric Potential**.

   Leave the default `0.0 V` for the corresponding value of electric potential.



At the SEI, the Butler-Volmer relation is used to compute the normal component of the electric current density at the SEI. The equilibrium potential of the active material term varies with the relative concentration in the electrode. Two tables are provided for you to use in the positive and negative electrodes.

6. Multi-select the **Interfaces > Anode AM/Anode Electrolyte** node and the **Anode AM/Separator** node.
7. Right-click any selected node and select **Edit...**.

   The *Multi-Edit* dialog appears.
8. Expand the **Physics Values > Butler-Volmer relation parameters** node and set *Electrode Equilibrium Potential* to **Tabular**.



The **Tabular Electrode Equilibrium Potential** node appears.

9. Expand the **Tabular Electrode Equilibrium Potential** node:
   - Set *Table* to **UeqNeg**.
   - Set *Relative Concentration in Electrode (y)* to **y**.
   - Set *Electrode Equilibrium Potential* to **- Ueq**.



10. Click **Close**.
11. Repeat the interface condition procedure for the **Cathode AM/ Cathode Electrolyte** and **Cathode AM/Separator** interfaces, but set *Table* to **UeqPos** and *Electrode Equilibrium Potential* to **Ueq**.

Setting Solver Parameters and Stopping Criteria

Set the appropriate solver parameters and stopping criteria for the simulation.

In this tutorial, charge and discharge currents have low values, and all variable changes are slow. Therefore, you can increase the time-step of the implicit solver to a large value.

The solver requirements for three-dimensional electrochemistry are higher than usual, because of the large differences in electrical and species conductivities in the different regions. These differences in conductivities affect the linear system. Therefore, you can change the expert properties of the AMG Linear Solver for Electric Potential.

To set the solver parameters and stopping criteria:

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	**Time-Step**	`10.0 s`
Electric Potential > AMG Linear Solver	*Enable direct-solver*	Activated
	Maximum direct-solver equations	`1024`



2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	*Enabled*	Activated
	Maximum Physical Time	`100 s`
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`

3. Save the simulation.

Preparing Scalar and Vector Scenes

Create scalar and vector scenes to display the simulation results.

You create two scalar scenes to display the Li-Ion/salt concentration across a plane section of the battery cell, and the specific electric current on the SEI.

To prepare the scalar and vector scenes:

1. Right-click the **Derived Parts** node and select **New Part** > **Section** > **Plane** .



2. In the *Edit* dialog, set the normal to `x = 0 m`, `y = 0 m`, and `z = 1 m`.



The plane scene looks as follows.



3. Click **Create**, and then **Close**.
4. Create a scalar scene.

5. Right-click the **Scalar Scene 1 > Displayers > Scalar 1 > Parts** node and select **Edit...**.
6. In the *Edit* dialog, activate **Derived Parts > plane section**.



7. Click **OK**.
8. Click (**Show All Meshes**).
9. Select the **Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Lithium/Salt Concentration**.



Prepare a vector scene to visualize the electronic pathways.

10. Create a vector scene.
11. Select the **Vector Scene 1 > Displayers > Vector 1 > Parts** node and click to the right part of the *Parts* property.
12. Activate **Derived Parts > plane section** and click **OK**.
13. Select the **Vector Scene 1 > Displayers > Vector 1 > Vector Field** node and set *Function* to **Electric Current Density**.

Initializing the Solution and Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

To initialize the solution and run the simulation:

1. Click (**Initialize Solution**) in the *Solution* toolbar or select **Solution** > **Initialize Solution** from the menu.

   The initial Lithium/Salt Concentration looks as follows:



2. Click (**Run**) in the *Solution* toolbar.

   The solution progress is displayed in the *Output* window. The *Residuals* display is created automatically in the *Graphics* window and shows the progress of the solvers.

   While the simulation is running, you can click the tabs at the top of the *Graphics* window to view scenes and plots.

   During the run, you can stop the process by clicking (**Stop**) on the toolbar. If you do stop the simulation, you can click (**Run**) to resume it.

   If left alone, the simulation continues until 100 seconds of physical time are complete. When the last iteration is complete, the following message is displayed in the *Output* window:

   `Stopping criterion Maximum Physical Time satisfied.`

3. When the simulation has finished running, save it.

Visualizing the Results

Visualize the Lithium/salt concentration in the active materials.

The scalar scene of the Li-Ion/salt concentration looks as follows.



To visualize the Lithium/salt concentration in the active materials:

1. Select the **Derived Parts > plane section** node and click ... (**Custom Editor**) for the *Parts* property.
2. In the *plane section* dialog, expand the **Regions** node and deactivate **Anode Electrolyte**, **Cathode Electrolyte**, and **Separator**.



3. Click **OK**.

The Lithium/salt concentration in the active materials looks as follows.



On the anode side (left-hand side in the above screenshot), the concentration distribution in the active material indicates a discharge

process. The lower level of concentration next to the SEI, compared with the core part, is consistent as the Li+ ions in this region naturally migrate in the electrolyte.

On the cathode side, higher levels of concentration are observed in the region next to the SEI, where the Li+ ions are inserted. The particle next to the current collector has the higher concentration level, which is a result of the high current density in this location.



By reviewing the Li+ ion concentration distribution on this structure, it is possible to observe whether large concentration gradients occur over time. These large concentration gradients are undesirable as they can lead to an increase in cell impedance, which results in large electric potential drop—hence adversely affecting the energy that is delivered by the cell.

Analyzing the current density distribution, especially on the anode side, can give good indications of potential occurrence of certain undesired side reactions, such as lithium deposition. These reactions are more likely to occur in regions of high current density.

This feature allows the cell manufacturer to understand the electrochemical process under a given load, and analyze the effects of various parameters—particle size, porosity, or total electrode thickness—on the cell performance.

Summary

This tutorial demonstrated how to model a Li-Ion battery cell at the microscale level.

The steps covered were:

- Setting up an electrochemistry model.
- Post-processing Lithium/Salt Concentration and SEI Specific Electric Current.

Casting

The tutorials in this set illustrate various STAR-CCM+ features for setting up casting:

Contents

Basic Gravity Sand Casting
STAR-Cast: Investment Casting

Basic Gravity Sand Casting

This tutorial demonstrates the setup, running, and post-processing of a basic gravity sand casting analysis using STAR-CCM+.

The mold is made of furan sand. The aluminum alloy AlSi7Mg0.3 is poured from a casting ladle into a sprue cup.

The casting scenario consists of a rectangular sand mold with a cavity composed of three cylindrical bodies and a gating system. The inflow condition assumes that the melt enters at constant pressure from the start of the simulation.

For this case, the sprue cup and casting regions are initially empty. The sprue cup outlet is set to atmospheric pressure. The mass flow into the pour cup corresponds to the flow rate that is required to fill the casting region in about 5–6 s.



This tutorial first simulates the filling of the casting region. You also simulate solidification after the mold is filled with the melt.

Depending on the specifications of your computer, the run time of this tutorial is of the order of several days. You are advised to run the tutorial in parallel to decrease the run time.

This tutorial uses the Criteria Functions model and therefore requires a double-precision version of STAR-CCM+.

Contents

Prerequisites
Loading the Geometry
Generating the Mesh
Setting Up the Physics Models
Selecting the Casting Materials
Setting Up the Sand Mold Physics Properties
Creating Eulerian Phases and Setting Phase Models
Setting Up the Casting Properties
Setting Region and Boundary Conditions
Setting Solver Parameters and Stopping Criteria
Creating a Solution History File
Visualizing the Solution
Running the Simulation
Visualizing the Results
Creating a Melt Filling Animation
Summary

Prerequisites

The instructions in this tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction
Using visualization tools, scenes, and plots	Introduction

Loading the Geometry

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied geometry.

1. Start up a double precision version of STAR-CCM+ in a manner that is appropriate to your working environment.

A file has been prepared that contains geometry parts for:

- The sand mold
- The sprue cup
- The casting region

To load the file:

2. In the menu bar, select **File** > **Load Simulation**.
3. In the *Load Simulation* dialog, click **Browse** and then navigate to the `doc/startutorialsdata/casting/data` subdirectory of your STAR-CCM+ installation directory.
4. Select `basicGravitySandCastingGeo.sim` and click **Open**.
5. Click **OK**.

Visualize the geometry and save the simulation:

6. Right-click **Scenes** > **Geometry Scene 1** and select **Open**.
7. Select **File** > **Save As** and save the simulation under the name `basicGravitySandCastingTutorial.sim`.

Generating the Mesh

The mesh models for this case are provided with the initial `.sim` file.

To review the mesh models for the sprue cup, casting, and mold regions:

1. Expand the **Geometry** > **Operations** > **Automated Mesh** node.



The mold, casting, and sprue cup regions use a polyhedral mesh.

To generate the mesh:

2. Click (**Generate Volume Mesh**) in the **Mesh** toolbar.

   The mesh generation process takes several seconds, depending on the specifications of your computer.

Upon completion of the mesh generation process, make the mesh visible.

3. Right-click the **Scenes** > **Geometry Scene 1** node and select **Apply Representation** > **Volume Mesh**.
4. In the *Geometry Scene 1 - Properties* window, set *Mesh Override* to **Show All Meshes**.
5. Click (**Make Scene Transparent**) in the **Vis** toolbar.
6. Position the parts to obtain a similar view to the one shown below.



7. Save the simulation.

Setting Up the Physics Models

The casting process is transient. The flow of the melt, degassing of the air that is initially present in the casting region, and the solidification of the melt are all of interest for this case.

Melt Continuum

For the casting and sprue cup regions, you set up an implicit unsteady, laminar, multi-phase model using the segregated VOF flow and multiphase temperature models. To access the casting material database and criteria functions, you additionally select the casting model.

To select the physics models for the melt:

1. Rename the **Continua** > **Physics 1** node to `Casting Physics`.
2. For the **Continua** > **Casting Physics** continuum, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase** **Multiphase Interaction** (Selected automatically)
Eulerian Multiphase Model	**Volume of Fluid (VOF)** **Gradients** (Selected automatically) **Multiphase Equation of State** (Selected automatically) **Segregated Flow** (Selected automatically)
Viscous Regime	**Laminar**
Optional Models	**Segregated Multiphase Temperature** **Gravity** **Casting** **Segregated Time Step** **Cell Quality Remediation**

3. Click **Close**.

The following screenshot shows the models that are used in the **Casting Physics** continuum.



Mold Continuum

The heat transfer is of interest in the sand mold region. To activate the models for the mold:

4. Right-click the **Continua** node and select **New > Physics Continuum**.

   A new physics continuum, **Physics 1**, appears in the model tree.
5. Rename the **Continua > Physics 1** node to `Mold Physics`.
6. For the **Continua > Mold Physics** continuum, select the following models in order:

Group Box	Model
Space	**Three Dimensional** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Solid**
Optional Models	**Segregated Solid Energy Gradients** (Selected automatically)
Equation of State	**User Defined EOS**

Group Box	Model
Optional Models	**Cell Quality Remediation** **Radiation**
Radiation	**Surface-to-Surface Radiation**
Radiation Spectrum	**Gray Thermal Radiation**

7. Click **Close**.

The following screenshot shows the models that are used in the **Mold Physics** continuum.



Make sure that the correct continuum is assigned to each region:

8. Multi-select the **Regions > Casting** and **Regions > SprueCup** nodes and make sure that *Physics Continuum* is set to **Casting Physics**.
9. Select the **Regions > SandMold** node and set *Physics Continuum* to **Mold Physics**.



10. Save the simulation .

Selecting the Casting Materials

You select the materials for the casting simulation from the casting material database. The aluminum alloy is AlSi7Mg0.3 and the mold is made of furan sand.

To open the casting material database dialog:

1. Right-click the **Tools > Casting > Material Database Editor** node and select **Casting Material Database Editor...**

   The dialog can take several seconds to load.

   In the *Disclaimer notice* dialog:
2. Read the disclaimer.
3. Activate **I fully and expressly agree to both sets of provisions specified above and the agreement concerning a contract penalty**.
4. Click **OK**.

   The *STAR-Cast mat* dialog opens with the *Material Selection* panel selected by default. The STAR-CCM+ GUI is not active when the casting material database dialog is opened.
5. Select the casting material:

Property	Setting
Material Class	**Casting metals**

Property	Setting
Material Group	**Aluminium alloys**
Material Subgroup	**Aluminium-silicon-alloys**
Material Designation	**EN AC-AlSi7Mg0.3**



To review the density of EN AC-AlSi7Mg0.3:

6. Click the *Physical Properties* tab.
7. Within the *Properties* group box, beneath the *Thermal* section, click **Density**.

   The *Datasets* panel opens and shows the graph of density against temperature for the selected material.



Similarly, you can review the other properties for the casting material.

To add EN AC-AlSi7Mg0.3 to the collection of materials to be exported to STAR-CCM+:

8. Click the *Export Data* tab.
9. Click **Add Material to Collection**.

   EN AC-AlSi7Mg0.3 is added to the *Collection of Materials*.



Select the material for the sand mold:

10. Click the *Material Selection* tab.

11. Set *Material Class* to **Mould Materials**.
12. Click the Icon List to the right of the Material Class menu.

    The *Available 'Mould materials'-materials* dialog appears.
13. Select **Furan Resin sand**.



14. Click **Select**.
15. Open the *Export Data* panel and click **Add Material to Collection**.



To export this material collection to STAR-CCM+:

16. Click **Export Collection to STAR-CCM+**.

    The casting material database dialog closes and STAR-CCM+ GUI becomes available.
17. Expand the **Tools** > **Material Databases** > **Selected_Casting_Materials** node.

    The materials that are required for this tutorial have been imported into STAR-CCM+.



18. Save the simulation .

Setting Up the Sand Mold Physics Properties

You must set material properties, physical model properties, and initial conditions for the sand mold.

To set up the properties of the sand mold:

1. Right-click the **Continua** > **Mold Physics** > **Models** > **Solid** > **Al** node and select **Replace with...**

   The *Replace Material* dialog appears.
2. Expand the **Material Databases** > **Selected_Casting_Materials** > **Mould materials** > **Lost moulds** > **Sand moulds** node and select **Furan Resin Sand (Furan Resin Sand)**.



3. Click **OK**.

Set the enthalpy and thermal conductivity of the furan sand using the table from the imported material collection:

4. Select the **Furan Resin Sand** > **Material Properties** > **Enthalpy** node and set *Method* to **Table, h(T)**.
5. Select the **Enthalpy** > **Table, h(T)** > **Table(T)** node and make sure that *Input table* is set to **Material Database Table**.
6. Make sure that *Table:Data* is set to **Enthalpy**.
7. Select the **Furan Resin Sand** > **Material Properties** > **Thermal Conductivity** node and set *Method* to **Table(T)**.
8. Select the **Thermal Conductivity** > **Table(T)** > **Table(T)** node and make sure that *Input table* is set to **Material Database Table**.
9. Make sure that *Table:Data* is set to **ThermalConductivity**.



Set the environment radiation temperature to 30 C:

10. Select the **Mold Physics** > **Models** > **Gray Thermal Radiation** > **Thermal Environments** > **Radiation Temperature** node and set *Radiation Temperature* to `30 C`.

Note: You can type the temperature value together with its unit.

Set the initial temperature of the sand mold to 30 C:

11. Select the **Mold Physics** > **Initial Conditions** > **Static Temperature** > **Constant** node and set *Value* to `30 C`.

12. Save the simulation .

Creating Eulerian Phases and Setting Phase Models

The two phases flowing through the casting and sprue cup regions are the melt and air.

These phases are defined by first creating them under the Eulerian phases model, and then selecting their corresponding material properties.

Melt Phase

The melt phase is created first:

1. Right-click the **Continua** > **Casting Physics** > **Models** > **Eulerian Multiphase** > **Eulerian Phases** node and select **New**.

   A new phase, **Phase 1**, appears in the simulation tree.
2. Rename **Phase 1** to `Melt`.
3. Right-click the **Eulerian Multiphase** > **Eulerian Phases** > **Melt** > **Models** node and choose **Select models...**

   In the *Melt Model Selection* dialog, select the following models in order:

Group Box	Model
Enabled Models	**Laminar** (Selected automatically) **Energy** (Selected automatically)
Material	**Liquid**
Equation of State	**User Defined EOS**
Optional Models	**Melting-Solidification** **Slurry Viscosity** **Mushy Zone Permeability** **Solidification Time** When the melt is in a partly solidified state, it is modeled as a slurry in cells where the solid volume fraction is low, and as a mushy zone for higher solid volume fractions. **Criteria Functions** The Criteria Functions model is used solely for post-processing

Group Box	Model
	but has to be activated before running the simulation.

4. Click **Close**.

   The following screenshot shows the models that are used in the **Melt** phase.



Set the material properties of the melt:

5. Right-click the **Melt** > **Models** > **Liquid** > **H2O** node and select **Replace with...**

   The *Replace Material* dialog appears.

6. Expand the **Material Databases** > **Selected_Casting_Materials** > **Casting metals** > **Aluminium alloys** > **Aluminium-silicon-alloys** node and activate **EN AC-AlSi7Mg0.3 (EN AC-AlSi7Mg0.3)**.



7. Click **OK**.

Use the imported material properties for EN AC-AlSi7Mg0.3.

To export the density table from the Selected_Casting_Materials database:

8. Right-click the **Tools** > **Material Databases** > **Selected_Casting_Materials** > **Casting metals** > **Aluminium alloys** > **Aluminium-silicon-alloys** > **EN AC-AlSi7Mg0.3 (EN AC-AlSi7Mg0.3)** > **Density** > **Table(T)** > **Material Database Table** node and select **Export...**.
9. In the *Save* dialog, enter `rho_T.csv` as **File Name** and click **Save**.
10. Right-click **Tools** > **Tables** and select **New Table** > **File Table**.
11. In the *Open* dialog, select `rho_T.csv` and click **Open**.

To enhance the stability of the simulation, an artificial compressibility of the melt phase is introduced. The density of EN AC-AlSi7Mg0.3 depends on pressure and temperature.

To make the melt phase compressible:

12. Select the **Melt** > **Models** > **User Defined EOS** node and activate *Compressible*.

You create a field function to set the density of EN AC-AlSi7Mg0.3:

13. Right-click **Tools** > **Field Functions** and select **New** > **Scalar**.
14. Rename **User Field Function 1** to **Variable Melt Density**.
15. Set *Function Name* to `VariableMeltDensity`.

16. Set *Dimensions* to **Mass/Volume** by clicking ... (**Custom Editor**).
17. In the *Variable Melt Density - Dimensions* dialog:
    a) Set *Mass* to `1`.
    b) Set *Volume* to `-1`.
    c) Click **OK**.
18. Set *Definition* to `max($MinimumMeltDensity, $InterpolatedMeltDensity + $MeltDensityPressureDerivative*$Pressure)`.

    This field function definition uses other user-defined field functions that you define next using the same procedure as described above:

Field Function Node	Dimensions	Function Name	Definition
Minimum Melt Density	Dimensionless	`MinimumMeltDensity`	1
Melt Density Pressure Derivative	**Mass/ Pressure-Volume**	`MeltDensityPressureDerivative`	`1e-5`
Interpolated Melt Density	Dimensionless	`InterpolatedMeltDensity`	`interpolateTable(@Table("rho_T"), "Temperature", LINEAR, "Density", $Temperature)`
Melt Density Temperature Derivative	Dimensionless	`MeltDensityTemperatureDerivative`	`interpolateTable(@Table("rho_T"), "Temperature", LINEAR, "Density", $Temperature + 1) - $InterpolatedMeltDensity`

Note: You can copy and paste the field function, Function Names, and Definitions into STAR-CCM+.

Use this variable density for the melt phase.

19. Select the **Continua > Casting Physics > Models > Eulerian Multiphase > Eulerian Phases > Melt > Models > Liquid > EN AC-AlSi7Mg0.3 > Material Properties > Density** node and set *Method* to **Field Function**.
20. Select the **Material Properties > Density > Field Function** node and set *Scalar Function* to **Variable Melt Density**.
21. Select the **Material Properties > Density pressure derivative** node and set *Method* to **Field Function**.
22. Select the **Material Properties > Density pressure derivative > Field Function** node and set *Scalar Function* to **Melt Density Pressure Derivative**.
23. Select the **Material Properties > Density temperature derivative** node and set *Method* to **Field Function**.
24. Select the **Material Properties > Density temperature derivative > Field Function** node and set *Scalar Function* to **Melt Density Temperature Derivative**.

To set the other material properties that depend on temperature:

25. Select the **Liquid > EN AC-AlSi7Mg0.3 > Material Properties > Dynamic Viscosity** node and set *Method* to **Table(T)**.
26. Select the **Dynamic Viscosity > Table(T) > Table(T)** node and make sure that *Table: Data* is set to **DynamicViscosity**.
27. Similarly, set the following material properties to the **Table(T)** method: **Fraction Solid Curve** and **Thermal Conductivity**.



Activate the use of the switching function for the Carman-Kozeny Mushy Zone Permeability model and the Metzner Slurry Viscosity model:

28. Select the **Melt > Models > Liquid > EN AC-AlSi7Mg0.3 > Material Properties > Mushy Zone Permeability > Carman-Kozeny** node and make sure that **Use Switching Function** is activated.
29. Select the **Melt > Models > Liquid > EN AC-AlSi7Mg0.3 > Material Properties > Slurry Viscosity > Metzner** node and make sure that **Use Switching Function** is activated.

The interesting physics does not always happen when the melt is completely solidified, but rather when a large amount of the melt is solidified and feeding is blocked by the dendrite arm network. You set

the critical temperature at which some quantities—temperature gradients, cooling rates, and local solidification times—are of interest. To set the critical temperature for EN AC-AlSi7Mg0.3:

30. Select the **Melt > Models > Criteria Functions** node and set *Input Units* to **C**.
31. Set *Critical Temperatures* to `[574.0]`

    This critical temperature is between the solidus temperature (555 C) and liquidus temperature (615 C) of the melt.

To stabilize the simulation, under-relax the solidification process:

32. Select the **Melt > Models > Melting-Solidification** node and set *Under-Relaxation Factor* to `0.7`.
33. Select the **Continua > Casting Physics > Models > Average Slurry Viscosity** node and set *Slurry Viscosity Under-Relaxation Factor* to `0.5`.
34. Select the **Continua > Casting Physics > Models > Average Mushy Zone Permeability** node and set *Mushy Zone Permeability Under-Relaxation Factor* to `0.5`.
35. Set the *Max. Porous Viscous Resistance* to `70000 kg/m^3-s`

The flow in the melt stops when the solid fraction is near 0.9. Set this flow stop value:

36. Select the **Melt > Models > Melting-Solidification** node and set *Flow Stop Relative Solid Fraction* to `0.95`.
37. Select the **Continua > Casting Physics > Models > Melting-Solidification Flow Stop** node and set *Flow Stop Solid Fraction* to `0.95`.

    The properties of the melt are determined the *Flow Stop Relative Solid Fraction*. The *Flow Stop Solid Fraction* is set to a higher value because the slurry viscosity and mushy zone permeability models already decelerate the flow for lower solid volume fractions. The flow stop functionality is only required when creeping flow is to be prevented, or when the solid needs to withstand very high pressure gradients. Hence, the *Flow Stop Solid Fraction* is generally set to a value close to 1.0.

Air Phase

Create the air phase:

38. Right-click the **Continua > Casting Physics > Models > Eulerian Multiphase > Eulerian Phases** node and select **New**.

    A new phase, **Phase 1**, appears in the simulation tree.

39. Rename **Phase 1** to `Air`.
40. Right-click the **Eulerian Multiphase** > **Eulerian Phases** > **Air** > **Models** node and choose **Select models...**

    In the *Air Model Selection* dialog, select the following models.

Group Box	Model
Enabled Models	**Laminar** (Selected automatically) **Energy** (Selected automatically)
Material	**Gas**
Equation of State	**Ideal Gas**
Optional Models	**Wall Porosity** Degassing occurs when the melt flows into the casting region and air is pushed into the sand mold.

41. Click **Close**.

    The following screenshot shows the models that are used in the **Air** phase.



Activate Density Limiting for the **Air** phase:

42. Select the **Eulerian Phases** > **Air** > **Models** > **Ideal Gas** node and activate *Density Limiting*.

43. Select the **Air** > **Models** > **Ideal Gas** > **Density Limits** node and set *Cut-off pressure* to `10000.0 Pa`.

Phase Interaction

Set up the phase interaction between the melt and the air:

44. Right-click the **Continua** > **Casting Physics** > **Models** > **Multiphase Interaction** > **Phase Interactions** node and select **New**.

    A new phase interaction, **Phase Interaction 1**, appears in the simulation tree.
45. Rename **Phase Interaction 1** to `Casting_Phase_Interaction`.
46. Right-click the **Casting_Phase_Interaction** > **Models** node and choose **Select models...**

    In the *Casting_Phase_Interaction Model Selection* dialog, select the following models.

Group Box	Model
Phase Interaction Topology	**VOF Phase Interaction**
Optional Models	**Surface Tension Force** **Multiphase Material** (activated automatically) **Interface Momentum Dissipation**

47. Click **Close**.
48. Select the **Casting_Phase_Interaction** > **Models** > **VOF Phase Interaction** node and set *Primary Phase* to **Melt**.
49. Set *Secondary Phase* to **Air**.
50. Select the **Casting_Phase_Interaction** > **Models** > **Multiphase Material** > **Multiphase Material** > **Material Properties** > **Surface Tension** > **Constant** node and set *Value* to `0.9 N/m`



51. Save the simulation .

Setting Up the Casting Properties

Here, you set up model properties such as the direction of gravity and the initial conditions of the simulation.

To improve the prediction of the melt-air interface between the VOF phases, set the VOF sharpening factor:

1. Select the **Continua** > **Casting Physics** > **Models** > **Volume of Fluid (VOF)** node and set *Sharpening Factor* to `0.2`.

Gravity acts in the -Z direction.

2. Select the **Continua** > **Casting Physics** > **Reference Values** > **Gravity** node and make sure that *Value* is set to `[0.0, 0.0, -9.81] m/s^2`.

The casting region is initially at a temperature of 30 C and filled with air. To set up the initial properties of the casting continua:

3. Select the **Continua** > **Casting Physics** > **Initial Conditions** > **Static Temperature** > **Constant** node and set *Value* to `30 C`.

4. Select the **Continua** > **Casting Physics** > **Initial Conditions** > **Volume Fraction** > **Constant** node and set *Value* to `[0.0, 1.0]`



Note: The first entry in the *Value* array denotes the melt fraction, and the second entry denotes the air fraction.

5. Save the simulation .

Setting Region and Boundary Conditions

Set the conditions in the sprue cup, casting, and sand mold regions.

The sprue cup is initially filled with air. The outlet surface is set to atmospheric pressure and a constant temperature that is lower than the initial melt temperature, yet significantly hotter than room temperature. The mass flow into the pour cup yields a filling time of around 5–6 s.

Outlet

Specify the outlet properties in the sprue cup:

1. Select the **Regions** > **SprueCup** > **Boundaries** > **Outlet** > **Physics Values** > **Static Temperature** > **Constant** node and set *Value* to `30 C`.
2. Select the **Volume Fraction** > **Constant** node and set *Value* to `[0.0, 1.0]`.



Mass Flow Inlet

Specify the mass flow inlet properties:

3. Select the **Regions** > **SprueCup** > **Boundaries** > **Inlet** > **Physics Values** > **Total Temperature** > **Constant** node and set *Value* to `690 C`.

During the first 0.05 s of the simulation, the mass flow rate is linearly ramped up to 0.35 kg/s and remains constant for 5.7 s. Then, the mass flow rate is linearly ramped down to 0 kg/s. This corresponds to an approximate filling time of 5.8 s. The melt is not poured into the domain after this interval. This time-dependent inlet melt mass flow is provided as a table.

Import the inlet mass flow table:

4. Right-click the **Tools** > **Tables** node and select **New Table** > **File Table**.
5. In the *Open* dialog, select **inletMassFlow.csv** and click **Open**.

Specify this table for the inlet mass flow rate:

6. Select the **Regions** > **SprueCup** > **Boundaries** > **Inlet** > **Physics Values** > **Mass Flow Rate** node and set *Method* to **Table (time)**.
7. Select the **Table (time)** node and set *Table* to **inletMassFlow**.
8. Set *Interpolation* to **Linear**.
9. Set *Table: Data* to **Mass Flow Rate**.



The volume fraction of melt at the inlet during the filling is 1.0. When the filling is complete, this fraction is set to 0.0. Use a table to specify this condition.

Import the inlet volume fraction table:

10. Right-click **Tools** > **Tables** and select **New Table** > **File Table**.
11. In the *Open* dialog, select **inletVolumeFraction.csv** and click **Open**.

Specify this table for the mass flow inlet volume fraction:

12. Select the **Regions** > **SprueCup** > **Boundaries** > **Inlet** > **Physics Values** > **Volume Fraction** node and set *Method* to **Composite**.
13. Multi-select the **Volume Fraction** > **Composite** > **Melt** and **Air** nodes and set *Method* to **Table (time)**.
14. Select the **Melt** > **Table (time)** node and in the *Properties* window:
    a) Set *Table* to **inletVolumeFraction**.
    b) Set *Interpolation* to **Linear**.
    c) Set *Table: Data* to **Melt**.
15. Select the **Air** > **Table (time)** node and in the *Properties* window:

a) Set *Table* to **inletVolumeFraction**.

b) Set *Interpolation* to **Linear**.

c) Set *Table: Data* to **Air**.



Interface Conditions

Set the interface conditions between the sprue cup and the sand mold. You specify the degassing conditions on both the interface and the wall. When the interface does not completely match the wall, the non-matching faces are impermeable to the gas.

Enter the degassing condition:

16. Select the **Regions > SprueCup > Boundaries > Walls > Phase Conditions > Air > Physics Conditions > Wall Porosity Specification** node and set *Method* to **Specified**.

17. Select the **Boundaries > Walls [SandMold/SprueCup] > Phase Conditions > Air > Physics Conditions > Wall Porosity Specification** node and set *Method* to **Specified**.
18. Beneath the **Boundaries** node, multi-select the **Walls > Phase Conditions > Air > Physics Values > Porous Viscous Resistance > Constant** and the **Walls [SandMold/SprueCup] > Phase Conditions > Air > Physics Values > Porous Viscous Resistance > Constant** nodes and set *Value* to `50000.0 m/s`.



Specify the melt wetting angle at the interface between the sprue cup and the casting region:

19. Select the **Walls [SandMold-SprueCup] > Phase Conditions > Casting_Phase_Interaction > Physics Values > Contact Angle > Constant** node and set *Value* to `115 deg`.

Similarly, set the contact angle of the sprue cup inlet wall to the constant melt angle value:

20. Select the **Walls > Phase Conditions > Casting_Phase_Interaction > Physics Values > Contact Angle > Constant** node and set *Value* to `115 deg`.

Set the interface conditions between the casting region and the sand mold. These conditions are identical to those of the Walls boundary.

21. Select the **Regions > Casting > Boundaries > Interface_SandMold > Phase Conditions > Air > Physics Conditions > Wall Porosity Specification** node and set *Method* to **Specified**.
22. Select the **Boundaries > Interface_SandMold > Phase Conditions > Air > Physics Values > Porous Viscous Resistance > Constant** node and set *Value* to `50000 m/s`.
23. Select the **Boundaries > Interface_SandMold [Casting/SandMold] > Phase Conditions > Air > Physics Conditions > Wall Porosity Specification** node and set *Method* to **Specified**.
24. Select the **Boundaries > Interface_SandMold [Casting/SandMold] > Phase Conditions > Air > Physics Values > Porous Viscous Resistance > Constant**node and set *Value* to `50000 m/s`.
25. Select the **Boundaries > Bottom_Interface_SandMold > Phase Conditions > Air > Physics Conditions > Wall Porosity Specification** node and set *Method* to **Specified**.
26. Select the **Boundaries > Bottom_Interface_SandMold > Phase Conditions > Air > Physics Values > Porous Viscous Resistance > Constant** node and set *Value* to `50000 m/s`.
27. Select the **Boundaries > Bottom_Interface_SandMold [Casting/SandMold 2] > Phase Conditions > Air > Physics Conditions > Wall Porosity Specification** node and set *Method* to **Specified**.
28. Select the **Boundaries > Bottom_Interface_SandMold [Casting/SandMold 2] > Phase Conditions > Air > Physics Values > Porous Viscous Resistance > Constant** node and set *Value* to `50000 m/s`.
29. Select the **Boundaries > Interface_SandMold [Casting/SandMold] > Phase Conditions > Casting_Phase_Interaction >**

Physics Values > **Contact Angle** > **Constant** node and set *Value* to `115 deg`.

30. Select the **Boundaries** > **Bottom_Interface_SandMold [Casting/SandMold 2]** > **Phase Conditions** > **Casting_Phase_Interaction** > **Physics Values** > **Contact Angle** > **Constant** node and set *Value* to `115 deg`.

Sand Mold Region

Set the radiation conditions such that the hot sand mold only radiates to the environment of the sand mold and not within the fluid domain.

31. Select **Regions** > **SandMold** > **Physics Conditions** > **Radiation Transfer Option** and deactivate *Internal Radiation Transfer*.

Reduce the number of radiation patches:

32. Select **Regions** > **SandMold** > **Physics Values** > **Patch/Face Proportion** and set *Patch/Face Proportion* to `20.0`.

Sand Mold Wall Boundaries

Set the thermal conditions:

33. Multi-select the **Regions** > **SandMold** > **Boundaries** > **Sides** > **Physics Conditions** > **Thermal Specification** and the **Regions** > **SandMold** > **Boundaries** > **Top** > **Physics Conditions** > **Thermal Specification** nodes and set *Condition* to **Environment**.
34. Multi-select the **Sides** > **Physics Values** > **Heat Transfer Coefficient** > **Constant** and the **Top** > **Physics Values** > **Heat Transfer Coefficient** > **Constant** nodes and set *Value* to `10 W/m^2-K`.
35. Save the simulation .

Setting Solver Parameters and Stopping Criteria

To enhance the convergence and stability of the simulation, you use the Segregated Time Step model. The Segregated Time Step model automatically adjusts the time step during run based on the CFL criteria that you specify.

You create two time step controls for the Segregated Time Step model:

1. Right-click the **Continua** > **Casting Physics** > **Models** > **Segregated Time Step** > **Time Step Controls** node and select **New**.
2. Rename the **Time Step Control 1** node to `domain`.
3. Create a second time step control and rename it `subset`.

You can specify the properties that are common to both time step controls simultaneously.

4. Multi-edit the **domain** and the **subset** nodes and set the following properties:

Property	Setting
Start Time Step	`5.0E-5 s`
Min. Time Step	`5.0E-6 s`
Flow-Based CFL Condition	**Activated**



The parameters for the condition profiles of the **domain** and the **subset** time controls differ.

5. Select the **domain** > **Flow-Based CFL Condition** node and set the following properties:

Property	Setting
Target Mean CFL Number	`0.26`
Max. Allowed CFL Number	`9.79`

You must first create the field function that is used by the subset condition profile. The field function selects only the cells in the vicinity of the melt-solid interface.

6. Right-click the **Tools > Field Functions** node and select **New > Scalar**.
7. Rename the new field function to `Melt Volume Fraction Condition`.
8. Select the **Melt Volume Fraction Condition** node and set the following properties:

Property	Setting
Function Name	`MeltVolumeFractionCondition`
Definition	`(${VolumeFractionMelt} >= 0.1 && ${VolumeFractionMelt} <= 0.9) ? 1 : 0`

9. Edit the **Time Step Controls > subset** node and set the following properties:

Node	Property	Setting
Condition Profile	*Method*	**Field Function**
Field Function	*Scalar Function*	**Melt Volume Fraction Condition**
Flow-Based CFL Condition	*Target Mean CFL Number*	`0.42`
	Max. Allowed CFL Number	`7.25`

Set an upper limit for the Segregated Time Step and the under-relaxation factor for pressure:

10. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	*Time-Step*	`0.1 s`
Segregated Flow > Pressure	*Under-Relaxation Factor*	`0.6`

Set the number of inner iterations per time step and the physical run time:

11. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`
Maximum Physical Time	*Maximum Physical Time*	`8.0 s`
Maximum Steps	*Enabled*	**Deactivated**

12. Save the simulation .

Creating a Solution History File

Create a solution history file to save the solved variables at different stages of the simulation.

You will use these saved states at a later stage for post-processing, the filling, and the solidification:

1. Right-click the **Solution Histories** node and select **New**.
2. In the *Save* dialog, set File Name to `Cast`, and click **Save**.

   The **Solution Histories** > **Cast** node appears in the simulation tree.
3. Select the **Cast** node and click (**Custom Editor**) for the *Scalar Functions* property.
4. In the *Scalar Functions* dialog, select the following variables: **Solid Volume Fraction**, **Temperature**, and **Volume Fraction**.



5. Click **OK**.

During the filling, the variables are saved at 0.1 s intervals:

6. Select the **Solution Histories** > **Cast** > **Update** node and set *Trigger* to **Delta Time**.
7. Select the **Update** > **Delta Time Frequency** node and set *Delta Time* to `0.1 s`.

   The filling starts at 0 s.
8. Set *Stop Time* to `5.8 s`.
9. Activate *Enable Stop*.



10. Save the simulation .

Visualizing the Solution

Create a scalar scene to display the melt volume fraction colored with temperature in the casting and sprue cup regions.

To create a melt volume fraction threshold:

1. Right-click **Derived Parts** and select **New Part** > **Threshold...**
2. Select the *Threshold* node and set the following properties:
   a) Set *Scalar Field* to **Volume Fraction** > **Volume Fraction of Melt**.
   b) Set *Mode* to **Above Max**.

c) Set *Range* to `[0.0, 0.5]`.

To create the melt fraction scene:

3. Create a scalar scene.
4. Rename **Scalar Scene 1** to `Melt Fraction Scene`.
5. Click **scene/plot**.
6. Select the **Melt Fraction Scene** > **Displayers** > **Outline 1** and set the following properties:
   a) Activate *Surface*.
   b) Set *Color Mode* to **Constant**.
   c) Set *Opacity* to `0.1`.
7. Select the **Melt Fraction Scene** > **Displayers** > **Outline 1** > **Parts** node and set *Parts* to **Regions** > **SandMold**.
8. Select the **Melt Fraction Scene** > **Displayers** > **Scalar 1** > **Parts** node and set *Parts* to **Derived Parts** > **threshold**.



9. Select the **Melt Fraction Scene** > **Displayers** > **Scalar 1** > **Scalar Field** node and set **Function** to **Temperature**.
10. Set *Units* to **C**.
11. Save the simulation .

Running the Simulation

You are now ready to run the simulation.

1. Click **Run** in the toolbar.

   A Residuals plot is created automatically, and shows the progress of the solver.

2. When the solver has finished running, save the simulation .

Visualizing the Results

Visualize the results after the simulation is completed.

The melt fraction scene after a simulation time of 8 s is shown below.



When the simulation stops after a physical time of 8 s, the melt has not completely solidified.

The tutorial stops here to keep the computer runtime to a minimum. If you wish, you can continue the run to achieve full solidification by increasing the physical run time of the simulation.

Creating a Melt Filling Animation

The solution history file contains the solution data from both the filling and solidification runs. Solution views are used to access this data for post-processing.

Add a solution time annotation to the scene to track the solution time as the scene animates:

1. Open the melt fraction scene to display the melt threshold.
2. Drag and drop the **Tools > Annotations > Solution Time** node onto the bottom left of the *Melt Fraction Scene*.

Create and set up a solution view for the animation:

3. Right-click the **Solution Histories > Cast** node and select **Create Recorded Solution View**.

   The **Solution Views > Cast** node appears in the simulation tree.
4. Select the **Solution Views > Cast > Animation** node and set *Animation Mode* to **Solution Time**.

   The **Solution Time Animation** sub-node appears.
5. Select the **Solution Time Animation** node and set *Cycle Time (seconds)* to `6.8`.
6. Set *End Solution Time* to `5.8`.
7. Set *Start Solution Time* to `0.0`.
8. To set the *Representation* property of all displayers and solution history annotations of the melt fraction scene to **Cast**, drag and drop the **Solution Views > Cast** node onto the *Melt Fraction Scene*.

To create the melt filling animation:

9. Click (**Write Movie**) in the **Animation** toolbar.
10. In the *Write animation* dialog, set Animation Length to `6.8`.
11. Set the *Size* to a resolution of your choice.
12. Set the *File Name* to `meltFilling.avi`.



13. Click **Save** to write the animation to disk.
14. Play back the animation using a player of your choice.

Summary

This tutorial demonstrated the setup, running, and post-processing of a basic gravity sand casting analysis.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up a basic sand casting simulation.
- Importing material properties from the casting material database.
- Interpolating tabular values with a field function.
- Creating a melt filling animation using a solution history file.

STAR-Cast: Investment Casting

This tutorial demonstrates the workflow for setting up an investment casting simulation.

Investment casting is an industrial process, which is also called lost-wax casting or lost-mold casting, because the cast mold is destroyed when the cast part is extracted. For precision casting, a wax model is dipped into a ceramic based slurry. The ceramic slurry dries and forms the mold for casting which is known as the shell. This process is repeated until the desired thickness is reached. The wax is then removed in an autoclave.

STAR-CCM+ provides a dedicated Offset Mesher model for the mesh that is required by this type of simulation. The Offset Mesher creates the geometry that represents the shell. The Offset Mesher mimics the manufacturing process and creates a mesh that is based on local curvature. By taking the shell into account for this type of simulation, you can simulate the full conjugate heat transfer of the melt interacting with the shell.

Note: The Offset Mesher is available using a macro. To obtain the macro, contact your local CD-adapco support representative.

The investment casting simulation is set up as a multiphase VOF flow simulation. The geometry consists of two regions, one represents the fluid cavity, and the other defines the top which contains the velocity inlet and pressure outlet. A third region is created to represent the shell.

Ideally, the multi-region volume mesh must have conformal interfaces between the regions. Conformal interfaces are created automatically when the parts are assigned to regions. All regions use the polyhedral volume mesher with the addition of the Offset Mesher for the shell region. To accurately compute the gradients at the wall and at the interfaces, create prism layers at these locations.

For the investment casting simulation here, the melt (a titanium alloy) enters the domain with a temperature of 1620 °C through a velocity inlet boundary. The inlet velocity is a function of time and is described using a table. When the melt impacts on the surface of the shell, it loses heat through the mold to the surroundings. The shell has an initial temperature of 900 °C and is bounded by convective wall boundaries. The hot melt cools off at the shell and solidifies.

The tutorial starts with the import of the investment casting geometry in Parasolid format. The geometry is shown below:



As the geometry is symmetric about a center plane, the simulation only requires half of the fluid cavity. The second half of cavity is accounted for by applying the symmetry plane boundary condition.

The melt is centrally poured into the mold at the top of the geometry. The air inside the mold escapes the geometry through the gap surrounding the melt. Simulating this condition would require an inflow boundary enclosed by an outflow boundary, on the same plane. Placing an inflow and outflow boundary together is numerically very unstable and must be avoided. Therefore, these two flow boundaries are separated by a wall boundary. The rounded shape is recommended as it avoids sharp corners and edges, which would reduce the robustness of the solution.

Note: You must use a double precision version of STAR-CCM+ for this tutorial.

Contents

Prerequisites
Importing the Geometry
Generating the Mesh
Setting up the Investment Casting Simulation
Defining the Casting Process
Selecting the Casting Materials
Setting Initial Conditions
Setting Boundary Conditions
Setting Interface Properties
Setting Physics Models
Specifying Melt-Air Surface Tension
Setting Solver Parameters
Setting up Automatic Inner Iterations
Setting up Criteria Functions
Setting Up Post Processing
Running the simulation
Visualizing the solution
Summary

Prerequisites

The instructions in this tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorials
Using visualization tools, scenes, and plots	Introduction

Importing the Geometry

The meshing process starts from a CAD geometry in Parasolid format. The imported geometry represents the fluid cavity bounded by the

inner walls of the shell, and the top, which contains the velocity inlet and pressure outlet.

To import the investment casting geometry of the sprue cup, the gating, the wheels, and the top:

1. Create a directory for the tutorial, called `investmentCasting`.
2. Start up a double precision version of STAR-CCM+ and create a new simulation by clicking **File > New Simulation...**
3. Select **File > Import > Import Surface Mesh...**. Navigate to the `doc/startutorialsdata/casting/data` subdirectory of your STAR-CCM+ installation directory. Select the file `castingAssembly.x_t` from the selection list and click **Open**.
4. In the *Import Surface Options* dialog, make sure that *Create Part Contact from Coincident Faces* is activated. Leave all other settings at their defaults and click **OK**.

   The imported mesh appears in *Geometry Scene 1*, the imported geometry appears under the **Geometry > Parts** node, and strong in-place contacts are created between the geometry parts.
5. Expand the **Parts > Cast > Surfaces** node and inspect the imported surfaces.
6. Save the simulation as `investmentCastingTutorial.sim`.

Generating the Mesh

This simulation requires a volume mesh for the fluid interior, a volume mesh which contains the inlet and outlet boundaries, and a volume mesh for the solid shell.

The imported surface for the investment cast part defines the outer surface of the fluid interior. As the quality of this surface is not adequate for subsequent operations, you remesh the surface to improve its quality. The mesh for the shell region is then grown from the remeshed surface of the fluid interior.

The following sections describe the sequence of operations for setting up mesh continua, regions, and boundaries, then generating the volume meshes.

Before generating the mesh, load the Offset Mesher model macro:

1. Select **File > Macro > Play Macro**.
2. Navigate to and select the macro, `OffsetMesher.java`. You can get this macro from your local CD-adapco support representative.

3. Click **OK**.

   A mesh continuum is created which has the following meshing models selected:

   - Polyhedral Mesher
   - Offset Mesher
   - Embedded Thin Mesher

Contents

Creating the Cast Mesh
Creating the Top Mesh
Creating the Mesh for the Shell Region
Generating the Volume Mesh

Creating the Cast Mesh

Follow this procedure to modify the mesh continuum which represents the cast part.

The existing mesh continuum contains the Offset Mesher model which grows the mesh for the shell region from the outer surface of the cast mesh. Before setting up the mesh for the shell region, the cast mesh must be defined and mesh settings applied.

Before creating the cast mesh continuum, assign parts to regions:

1. Multi-select the **Geometry** > **Parts** > **Cast** and **Top** nodes, right-click one of the nodes and select **Assign Parts to Regions...**
2. In the *Assign Parts to Regions* dialog:
   a) Select **Create a Region for Each Part.**
   b) Select **Create a Boundary for Each Part Surface**.
   c) Accept **Create One Feature Curve for All Part Curves**.
3. Click **Apply** and then **Close**.
   Interfaces are automatically created from the part contacts.

Set up a cast mesh continuum.

4. Rename **Continua** > **Mesh 1** to `Cast`.
5. Right-click the **Cast** > **Models** node and select **Select Meshing Models**.
6. In the *Cast Model Selection* dialog:
   a) Keep the meshing models that are already selected.

b) Activate **Surface Remesher**.

c) Activate **Prism Layer Mesher**.

d) Click **Close**.

7. Select the **Cast** > **Models** > **Surface Remesher** node and deactivate **Do Curvature Refinement**.
8. Select the **Models** > **Embedded Thin Mesher** node and activate **Customize Thickness Threshold**.
9. Expand the **Cast** > **Reference Values** node and set the following properties:

Node	Property	Value
Base Size	*Value*	`0.0025 m`
Offset Values	*Surface Size*	`0.005 m`
	Thickness	`0.007 m`
	Concavity Factor	`30`
	Convexity Factor	`1`
Thin Solid Thickness	*Size Type*	**Absolute**
Absolute Size	*Value*	`15.0 mm`

Define the boundaries on the Cast region.

10. Expand the **Regions** > **Cast** > **Boundaries** node and set the following properties:

Node	Property	Value
Symmetry	*Type*	**Symmetry Plane**
Cast-Top Interface		
Mesh Conditions		
Customize Prism Mesh	*Customize Prism Mesh*	**Disable**

Creating the Top Mesh

Follow this procedure to create a separate mesh continuum for the Top part, then define the Top region boundaries.

The Top part does not use the Offset Mesher model and requires a separate mesh continuum from the cast part. In this case, the prism layer mesher–which is normally deactivated on inlet and outlet boundaries–is

maintained. Maintaining a constant thickness prism layer along the coincident wall, inlet, and outlet boundaries ensures that good quality cells are created throughout the Top part.

To create a mesh continuum for the Top part:

1. Create a copy of the **Cast** mesh continuum and rename it **Top**.
2. Right-click the **Top** > **Models** node and select **Select Meshing Models**.
3. In the *Top Model Selection* dialog, deactivate *Offset Mesher*.

   The *Embedded Thin Mesher* model is automatically deactivated.
4. Click **Close**.

Define the boundaries on the Top Region.

5. Select the **Regions** > **Top** node and set *Mesh Continuum* to **Top**.
6. Expand the **Top** > **Boundaries** node and set the following properties:

Node	Property	Setting
Inlet	*Type*	**Velocity Inlet**
Mesh Conditions		
Custom Surface Size	*Custom Surface Size*	Activated
Customize Prism Mesh	*Customize Prism Mesh*	**Specify Custom Values**
Mesh Values		
Surface Size		
Relative Target Size	*Percentage of Base*	`40.0`
Pressure Outlet	*Type*	**Pressure Outlet**
Mesh Conditions		
Customize Prism Mesh	*Customize Prism Mesh*	**Specify Custom Values**
Symmetry	*Type*	**Symmetry Plane**
Top-Cast Interface		
Mesh Conditions		

Node	Property	Setting
▮▮▮▮**Customize Prism Mesh**	*Customize Prism Mesh*	**Disable**

Creating the Mesh for the Shell Region

Follow this procedure to define the shell mesh, region, and boundaries.

The shell mesh is grown from the interface between the Shell region and the Cast region, and uses mesh values that you specify in the Cast mesh continuum.

Create a region for the shell mesh:

1. Create a region and rename it `Shell`.
2. Select the **Regions** > **Shell** node and set *Type* to **Solid Region**.
3. Expand the **Shell** > **Mesh Conditions** node and set the following Properties:

Node	Property	Setting
Offset Options	*Mark as Offset Region*	Activated. In the dialog, click **Yes**.
	Input Boundaries	**Cast-Shell Interface**
Customize Thin Mesher Parameters	*Customize Thin Mesher Parameters*	**Specify Custom Values**

   An **Interfaces** > **Offset interface on Cast-Shell Interface** node appears.
4. Expand the **Regions** > **Shell** > **Mesh Values** node and set the following Properties:

Node	Property	Setting
Thin Mesher Layers	*Number of Layers*	`3`
Thin Solid Thickness	*Size Type*	**Absolute**
▮▮**Absolute Size**	*Value*	`12 mm`

Normally, prism layers are only required on the fluid side of an interface. In this case, due to the high temperature gradients at the interface, prism layers are created on both sides. Prism layers accurately predict the heat transfer at the interface between the liquid melt and the solid shell mesh. To create prism layers at the interface:

5. Expand the **Interfaces** > **Offset interface on Cast-Shell Interface** > **Mesh Conditions** node.
6. Select the **Interface Prism Layer Option** node and activate *Grow Prisms From Interface*.
7. Expand the **Regions** > **Shell** > **Boundaries** > **Offset Extrusion** > **Mesh Conditions** node.
8. Select the **Customize Prism Mesh** node and set *Customize Prism Mesh* to **Disable**.

Generating the Volume Mesh

You can now generate the volume mesh.

The **Cast** and **Top** mesh continua are set up and ready to mesh. The **Shell** region is grown from the **Offset interface on Cast-Shell Interface** on the outer surface of the **Cast** region.

Generate the volume mesh:

1. Select **Mesh** > **Generate Volume Mesh**.

To visualize the volume mesh:

2. Click *scene/plot*.
3. Select the **Displayers** > **Geometry 1** > **Parts** node.
4. In the *Parts - Properties* window, click **custom editor** ... next to *Parts*.
5. In the *Parts* dialog, select **Regions**. Then click **OK**.
6. Return to *simulation*.
7. Drag and drop the **Representations** > **Volume Mesh** node into the currently open *Geometry Scene 1*.
8. Save the simulation.

Setting up the Investment Casting Simulation

A dedicated Casting User Interface (CUI) guides you through the setup of the investment casting simulation.

Before you can use the CUI, you need to include it in the STAR-CCM+ toolbar.

For more information on how to access the CUI, see Accessing the Casting User Interface in the **User Guide** > **Modeling Physics** > **Modeling Casting** > **The Casting User Interface**.

Defining the Casting Process

In the *Casting Processes and Simulation Type* panel of the CUI, you define the casting process.

Before setting up the casting process, close all open scenes in the *Graphics* window.

To specify the investment casting process including filling and solidification:

1. In the **Casting** toolbar, click the STAR-Cast User Interface.
2. The CUI opens up with the *Condition Settings Import/Export* panel. Select the *Casting Processes and Simulation Type* panel.
   a) Set *Casting Process* to **Investment**.
   b) Set *Process Type* to **Gravity**.
   c) Set *Simulation Type* to **Filling and Solidification**.
   d) Set *Casting Temperature* to `1620 C`.
   e) Leave the settings for *Gravity, Reference Pressure, and Ambient Temperature* unchanged.



Selecting the Casting Materials

You must select the materials for the melt and the shell mold.

To open the casting material database:

1. Click the Material Data symbol.
2. In the *Material Data* panel, select **Launch Database...**
3. Acknowledge the disclaimer and click OK.

The *STAR-Cast mat* dialog opens with the *Material Selection* panel selected by default. The STAR-CCM+ GUI is not active when the casting material database dialog is opened.

To select the material for the casting metal:

4. Set the *Material Class* to **Casting metals**.
5. Set the *Material Group* to **Titanium alloys**.
6. Set the *Material Subgroup* to **Titanium-aluminium alloys**.
7. Set the *Material Designation* to **Ti-46Al-8Nb**.



Selected materials (and the associated material properties) from the Material database must be exported to STAR-CCM+. You can create a collection of several materials and export them to STAR-CCM+ at the same time.

To add the selected Titanium alloy to the Collection of Materials:

8. Click the *Export Data* tab.
9. Click **Add Material to Collection**.

   The Ti-46Al-8Nb alloy now appears in the list of Collection of Materials.

To select the material for the investment shell mold:

10. Click the *Material Selection* tab.
11. Set the *Material Class* to **Mould materials**.
12. Set the *Material Group* to **Lost moulds**.
13. Set the *Material Subgroup* to **Ceramic molds**.
14. Set the *Material Designation* to **Shell Mould for Gamma-TiAl-Casting**.



To add the investment shell mold material to the Collection of Materials:

15. Click the *Export Data* tab.
16. Click **Add Material to Collection**.

The Shell mold for Gamma-TiAl-Casting appears in the Collection of Materials list.

To export the Collection of Materials to Collection of Materials:

17. Select **Export Collection to STAR-CCM+** in the *Export Data* tab.

This automatically closes the *STAR-Cast mat* dialog. In the CUI, you can now see the two selected materials from the Materials Database.



Setting Initial Conditions

Here, you set the initial conditions for the fluid InvestmentCast region and the solid Shell region.

1. Click the Regions symbol.
2. Click the *Settings* tab. Here, you can define the initial conditions before you assign them to the appropriate regions.
3. Set the initial conditions for the fluid InvestmentCast region:
   a) Enter `Cast_Init` as the *Name* for this initial condition.
   b) Leave the *Type* as **Fluid Region**.
   c) Set *Temperature* to a constant value of `1100` ■`C`.
   d) Leave *Velocity* at the default value of `[0.0,0.0,0.0] m/s`.
   e) Set *Volume Fraction* to `[0.0,1.0]`. This means that the region is initially only filled with air.
   f) Click **Create.** `Cast_Init`appears in the list of Initial Conditions.



4. Set the initial conditions for the solid Shell region:
   a) Enter `Shell_Init` as *Name* for this initial condition.
   b) Set *Type* to **Solid Region**.
   c) Set *Temperature* to a constant value of `900■C`.
   d) Click **Create**.



To associate the initial conditions and the material properties with the appropriate regions:

5. Click the *Selections* tab. The three regions Cast, Shell, and Top are listed in the *Regions* column.
6. Click the empty *Materials* table cell next to the *Cast* table cell. Select **Ti-46Al-8Nb (Ti-46Al-8Nb)** as the material for this region.
7. Click the empty *Conditions* table cell for the Cast region. Select **Cast_Init** as initial condition.

   Cast_Init is offered as the only selection because it is the only fluid region condition defined.
8. Click the empty *Materials* table cell next to the *Top* table cell. Select **Ti-46Al-8Nb (Ti-46Al-8Nb)** as the material for this region.

9. Click the empty *Conditions* table cell for the Top region. Select **Cast_Init** as initial condition.
10. Click the empty *Materials* table cell next to the Shell region. Select **Shell mold for Gamma-TiAl-Casting**.
11. Click the empty *Conditions* table cell for the Shell region. Select **Shell_Init** as the initial condition.

    Shell_Init is offered as the only selection because it is the only solid region condition that is defined.
12. Save the simulation.

Setting Boundary Conditions

For this investment casting simulation, you define five different types of boundary conditions.

To set the boundary conditions:

1. Click the Boundaries symbol in the CUI.
2. Click the **Settings** tab.

   In this tab, you define the boundary conditions before associating them with the appropriate boundaries.
3. Set the Velocity Inlet Boundary:
   a) Enter `Melt_Inlet` as the *Name* for the boundary condition.
   b) Set *Type* to **Velocity Inlet Boundary**.
   c) At the melt inflow boundary, you specify time-dependent profiles for velocity, temperature and the volume fraction of melt and air:
   d) In the boundary properties table, set *Temperature* to **Table(Time)**.
   e) Set **Tabulate** selector to **Import .csv**.

      A **File** button appears to the right.
   f) Click **File**, navigate to the `doc/startutorialsdata/casting/data` subdirectory, and select the file `MeltInletLinear.csv`.
   g) Repeat these steps for *Velocity* and *Volume Fraction*. Select the file `MeltInletLinear.csv` for each variable.
   h) Click **Create**.



You can use a common text editor to open the file `MeltInletLinear.csv`. Inspect the format of the file. You can see that it is a multi-column file where all the inflow values are listed as columns depending on time as the first column. The headers in double quotes contain the keywords for the CUI to import the correct variable into the table editor.

4. Set the Pressure Outlet Boundary:

a) Enter `Gas_Outlet` as the *Name* for the boundary condition.

b) Set *Type* to **Pressure Outlet Boundary**.

c) Click *Temperature*. Set a constant value of `600 ■C`.

Note that you can change the temperature units from K to °C by typing a C after the temperature value.

d) Leave the default value of `0.0 Pa` for *Pressure*.

e) Click *Volume Fraction*. Set the constant values to [0.0,1.0]. This means that in case of inflow through the pressure boundary (due to recirculation), only air enters the fluid domain.

f) Click **Create**.



5. Set the Adiabatic Fluid Wall Boundary:

a) Enter `Melt_Wall` as the *Name* for the boundary condition.

b) Set *Type* to **Adiabatic Fluid Wall Boundary**.

c) Set *Contact Angle* to `90.0 deg`.

d) Leave all other properties at their default values.

e) Click **Create.**



6. Set the Solid Wall Boundary:

a) Enter `Shell_Wall` as the *Name* for the boundary condition.

b) Set *Type* to **Solid Wall Boundary**.

c) Set *Ambient Temperature* to `20` ■C.

d) Set *Heat Transfer Coefficient*to `10 W/(m↕K)`.

e) Set *Emissivity* to `0.5`

f) Click **Create.**



7. Set the Symmetry Plane Boundary:
   a) Enter `Symmetry` as the *Name* for the boundary condition.
   b) Set *Type* to **Symmetry Plane Boundary**.
   c) Click **Create.**

Proceed to associate the boundary conditions with the appropriate boundary surfaces.

8. Click the *Selections* tab.
9. For each row in the table, click the empty *Conditions* cell and set the conditions as follows:

Boundaries	**Conditions**
Cast: Symmetry	**Symmetry**
Top: Inlet	**Melt_Inlet**
Top: NoSlipWall	**Melt_Wall**

Boundaries	Conditions
Top: Symmetry	**Symmetry**
Top: PressureOutlet	**Gas_Outlet**
Shell: Offset Extrusion	**Shell_Wall**

Maintain the default *Conditions* for the remaining boundaries.

Setting Interface Properties

In this model, there are two interfaces: a contact interface between the fluid Cast region and the solid Shell region and an internal interface between the Top and the Cast region. At the contact interface, variable values are transferred between the two regions for a conjugate heat transfer analysis.

To specify the interface settings:

1. Click the Interfaces symbol, then the *Settings* tab.
2. Set the following properties for the interface between the melt and the shell:

Properties	Options	Values
Name		`Melt_Shell_Interface`
Type		**Fluid to Solid**
Contact Angle	**Constant**	`90 deg`
Heat Transfer Coefficient	**Table (Temp.)**	`HTC_T.csv`
Porous Inertia Coefficient	**Constant**	`0.0`
Porous Viscous Coefficient	**Constant**	`120000 m/s`

3. Click **Create**.
4. Set the following properties for the interface between the Cast and the Top region:

Properties	Options	Values
Name		`Cast_Top_Interface`
Type		**Flow Through**

5. Click **Create**.

Associate the interface condition with the appropriate interface surfaces of the model:

6. Click the *Selections* tab and set the following conditions:

Interfaces	Conditions
Cast/Top	**Cast_Top_Interface**
Shelling interface on Cast-Shell Interface	**Melt_Shell_Interface**

7. Save the simulation.

Setting Physics Models

The CUI sets the basic physics models for a multiphase, conjugate heat transfer simulation. There are optional physical models available that allow you to specify the melting-solidification properties of the melt.

For this tutorial you use:

- Melting-solidification flow stop model
- Mushy zone permeability model

To specify the properties of the optional melting-solidification models:

1. Click the Physical Models symbol in the CUI.
2. Click the *Settings* tab.
3. Set the Melting-Solidification Flow Stop parameters:
   a) Enter `FlowStop` as the *Name*.
   b) Select **Melting and Solidification** for the *Model*.
   c) Set the *Flow Stop Relative Solid Fraction* to `0.85`.
   d) Set the *Flow Stop Solid Fraction* to `0.85`.
   e) Click **Create**.



4. Set the Mushy Zone model parameters:
   a) Enter `MushyZone` as the *Name*.
   b) Select **Mushy Zone** for the *Model*.
   c) Set *Max. Viscous Resistance* to `50000 kg/(m→s)`.
   d) Set *Secondary Dendrite Arm Spacing* to `30.0 um`. (`um` represents ¼m.)
   e) Set *Critical Solid Fraction* to `0.5` and activate **Use Switching**.
   f) Click **Create**.



To assign the physical model conditions to the appropriate physical models:

5. Click the *Selections* tab.
6. For each row in the table, click the empty *Conditions* table cell and set the conditions as follows:
   a) Set *Melting and Solidification* to **FlowStop**.
   b) Set *Mushy Zone* to **MushyZone**.
7. Save the simulation.

Specifying Melt-Air Surface Tension

You use a field function to set up the surface tension between the melt and the air. Currently, setting the surface tension is only possible in the simulation tree.

To set the Surface Tension for the melt, switch to the STAR-CCM+ simulation tree:

1. In the simulation tree, right-click **Tools > Field Functions** and select **New > Scalar**.

2. Rename this field function to `Melt-Air Surface Tension Threshold`.
3. Select the **Melt-Air Surface Tension Threshold** node and set the following properties:

Property	Setting
Function Name	`MeltAir_SurfaceTensionThreshold`
Definition	`0.75`

4. Create another field function and rename it to `Melt-Air Surface Tension`.
5. Select the **Melt-Air Surface Tension** node and set the following properties:

Property	Setting
Function Name	`MeltAir_SurfaceTension`
Definition	`1.15*(0.5-atan(1000*(${RelativeSolidVolumeFractionTi-46Al-8Nb_Liquid} - ${MeltAir_SurfaceTensionThreshold}))/3.14159265)`

6. Edit the **Continua** > **Melt** > **Model** > **Multiphase Interaction** > **Phase Interactions** > **Casting_Phase_Interaction** > **Models** > **Multiphase Material** > **Material Properties** > **Surface Tension** node and set the following properties:

Property	Setting
Method	**Field Function**
▮▮*Scalar Function*	**Melt-Air Surface Tension**



Setting Solver Parameters

Solver parameters must be set for the investment casting simulation.

To access the Solvers panel, click the Solvers symbol. The *Solvers* panel contains three tabs:

- Process Phases
- Solver Settings
- Stopping Criteria

Set the parameters moving from right to left.

1. Set the Stopping Criteria for the first 0.8 s of the run:

Properties	Values
Name	`RunFor0.8s`
Maximum Physical Time	`0.8 s`
Maximum Inner Iteration	`50`

2. Click **Create**.
3. Set the Stopping Criteria for the final time period of the run:

Properties	Values
Name	`StopAfter3s`
Maximum Physical Time	`3 s`
Maximum Inner Iteration	`50`

4. Click **Create**.
5. Set the Solver Parameters:
   a) Click the *Solver Settings* tab.
   b) Enter `FirstSolverSettings` as the *Name* for the condition.
   c) Set the time step size and the under-relaxation factors:

Properties	Values
Time Steps	**Automatic**
Velocity URF	`0.6`
Melting/Solidification URF	`0.75`
Mushy Zone URF	`0.2`
Slurry Viscosity URF	`0.2`

   d) Select the **Time Steps** property and set the Automatic Time Step controls at the bottom of the panel:

Properties	Values
Start Time	`0 s`
End Time	`0.8 s`
Start Time Step	`5.0E-5 s`
Minimum Time Step	`1.0E-5 s`
Target Mean CFL Number	`1000`
Max. CFL Number	`1000`
Max. Time Step Change Factor	`1.5`

Properties	Values
Subset Volume Fraction Range	`[0.1, 1.0]`
Subset Target Mean CFL Number	`0.5`
Subset Max. CFL Number	`6.5`

e) Click **Create**.

f) Select **FillingSolverSettings** in the *Conditions* window and click **Copy**.

g) Select **FirstSolverSettings_1** and change the *Name* to `SecondSolverSettings`.

h) Set the Automatic Time Step controls at the bottom of the panel:

Properties	Values
Start Time	`0.8 s`
End Time	`3 s`
Start Time Step	`5.0E-5 s`
Minimum Time Step	`1.0E-5 s`
Target Mean CFL Number	`1000`
Max. CFL Number	`1000`
Max. Time Step Change Factor	`1.5`
Subset Volume Fraction Range	`[0.1, 1.0]`
Subset Target Mean CFL Number	`0.5`
Subset Max. CFL Number	`2.0`

i) Click **Modify**.

6. Set the Process Phases.

a) Click the *Process Phases* tab and enter `FirstPhase` as *Process Phase Name*.

b) Click **Insert**.

c) Similarly, create a second process phase and name it `SecondPhase`.

d) Set the following conditions:

Process Phase	Solver Settings	Stopping Criteria
FirstPhase	**FirstSolverSettings**	**RunFor0.8s**

Process Phase	Solver Settings	Stopping Criteria
SecondPhase	**SecondSolverSettings**	**StopAfter3s**

Switch to the simulation tree.

7. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `0.1 s`.
8. Save the simulation.

Setting up Automatic Inner Iterations

You optimize the computational time of your simulation by automatically adjusting the number of inner iterations per time step. The criterion to increase or decrease the number of inner iterations is a user-specified value for the relative mass error rate.

To set up the automatic inner iterations control, you must switch to the simulation tree.

1. Right-click **Reports** and select **New Report** > **Maximum**.
2. Rename this report `Melt_Relative_Mass_Error_Rate`.
3. Select the **Melt_Relative_Mass_Error_Rate** node and set the following properties:

Property	Setting
Scalar Field Function	**Liquid VOF Phase Relative Mass Error Rate**
Parts	**(Regions) [Cast, Top]**

4. Right-click the **Melt_Relative_Mass_Error_Rate** node and select **Create Monitor from Report**.
5. Select the **Monitors** > **Melt_Relative_Mass_Error_Rate Monitor** node and set *Trigger* to **Iteration**.
6. Right-click the **Stopping Criteria** node and select **Create New Criterion** > **From Monitor...** > **Melt_Relative_Mass_Error_Rate**.
7. Right-click the **Stopping Criteria** node and select **Create New Criterion** > **From Monitor...** > **Energy**.
8. Right-click the **Stopping Criteria** node and select **Create New Criterion** > **Minimum Inner Iterations**.
9. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Energy Criterion	*Criterion Option*	**Asymptotic**
	Logical Rule	**And**
Asymptotic Limit	*[Max - Min]*	`0.01`
	Number of Samples	`3`
Melt_Relative_Mass_Error_Rate Criterion	*Logical Rule*	**And**
Minimum Limit	*Minimum Value*	`1.5 /s`
Minimum Inner Iterations	*Minimum Inner Iterations*	`8`

Setting up Criteria Functions

By default, the CUI activates all criteria functions except **Liquid Residence Time**, which is optional.

For more information, see Criteria Functions Definitions in the **User Guide > Modeling Physics > Modeling Casting > Using Criteria Functions**.

To set the critical temperatures for which criteria functions are evaluated:

1. Click the Criteria Functions symbol.
2. In the **Temperature** box, enter `1490 C` and click **Add**. This adds the first temperature to the list of critical temperatures.

   Note that the *Criteria Temperatures* have to be specified in °C.
3. In the same way, add temperatures `1500 C`, `1510 C`, `1520 C`, and `1530 C`.



Setting Up Post Processing

You set up visualization of the results of the investment casting simulation from within the CUI.

To select the items for reporting and visualization:

1. Click the Post Processing symbol.
2. In the *Global* tab, activate the following post-processing items:
   - In the *Cavity Plausibility Checks* group box, select **Velocity** and **Temperature**.
   - In the *Filling/Solidification Progress group* box, select **Melt (Mass), Volume Ratio (Melt/Cavity)**, and **Volume Ratio (Solid/Liquid)**.
   - In the *Scalar Scenes* group box, select **Melt and Solidified**.



This completes the setup of the casting simulation using the CUI.

3. Save the simulation.

Running the simulation

After the filling process is complete, the flow stabilizes, and only melting/solidification takes place.

You are advised to run this tutorial on a multi-core machine with four cores or more. Depending on the specifications of the machine, this tutorial takes one to two days to complete.

To run the simulation:

1. Click the Solvers symbol.
2. In the *Process Phases* tab, select **FirstPhase** in the *Process Phase* table.
3. Click **Run**.

   The simulation starts running with the numerical settings of the first process phase. The run automatically continues with the second process phase.

Visualizing the solution

The Casting User Interface (CUI) provides two scalar scenes by default.

These scenes are called:

- Melt

  The Melt scene shows a plot of the liquid melt in the computational domain.

- Solidified

  The Solidified scene shows a plot of the melt that is already solidified.

You can use these scalar scenes to visualize the development of the solution while the simulation runs.

To visualize the liquid melt:

1. Right-click the **Scenes** > **Melt** node and choose **Open**.
2. Press the <f> key on your keyboard to obtain a front view of the plot.

   The following screenshot shows the position and temperature of the liquid melt at the end of the simulation:



To visualize the solidified material in the domain:

3. Right-click the **Scenes** > **Solidified** node and choose **Open**.
4. Press the <f> key on your keyboard to obtain a front view of the plot.

   The following screenshot shows the position and temperature of the solidified melt at the end of the simulation:.



After the simulation is finished, you can visualize the criteria functions.

To visualize the Nyama criterion:

5. Create a scalar scene.
6. Select the **Displayers > Scalar 1 > Parts** node.
7. Set *Parts* to **Derived Parts > VolumeFractionTi-46Al-8Nb_LiquidAboveMaxThreshold.**
8. Select the **Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Niyama > Ti-46Al-8Nb_Liquid > Niyama of Ti-46Al-8Nb_Liquid @ T_solidus**.
9. Press the <f> key on your keyboard to obtain a front view of the plot.

   The following screenshot shows the Niyama criterion after a physical time of 3 s:



To visualize the local solidification time at the solidus temperature:

10. Create a scalar scene.
11. Select the **Displayers > Scalar 1 > Parts** node.
12. Set *Parts* to **Derived Parts > VolumeFractionTi-46Al-8Nb_LiquidAboveMaxThreshold.**
13. Select the **Displayers > Scalar 1 > Scalar Field** node and set *Function* to **Local Solidification Time > Ti-46Al-8Nb_Liquid > Local Solidification Time of Ti-46Al-8Nb_Liquid @ T_solidus**.
14. Press the <f> key on your keyboard to obtain a front view of the plot.

    The following screenshot shows the local solidification time criterion after a physical time of 3 s:



Summary

This tutorial demonstrates the workflow for setting up an investment casting simulation.

This tutorial has introduced the following STAR-CCM+ features:

- Importing an investment casting geometry.
- Using the Offset Mesher to generate a shell mold for investment casting simulations.
- Manually preparing the geometry for volume meshing.
- Setting up an investment casting simulation in the Casting User Interface.
- Running an investment casting simulation.
- Visualizing the simulation results.

Automation

The tutorials in this set illustrate various STAR-CCM+ automation and macro features:

Contents

Simple Java Macros: Post-Processing Objects
Intermediate Java Macros: Run Multiple Simulations
Simulation Assistant: Internal Flow Assistant

Simple Java Macros: Post-Processing Objects

This tutorial is intended to demonstrate the steps required in recording, editing, and replaying a Java macro in STAR-CCM+.

The case considered is that of an incompressible flow through a venturi flowmeter, and the objective in creating a macro is to generate multiple line probes throughout the solution domain to observe recovery of the flow profile.

The venturi flowmeter is symmetric around its central axis and can be readily modeled using a two-dimensional axisymmetric mesh. In STAR-CCM+ this requires you to import an initial three-dimensional mesh and convert it to two-dimensions using the relevant operations.

The object model that is used by the STAR-CCM+ Java Client is available for use by you in Java macros. Documentation for all classes and methods within this object model are available from the **Help > Java API** menu.



Contents

Prerequisites
Importing the Volume Mesh
Converting the Mesh to Two-Dimensions
Selecting the Physics Models
Setting Boundary Conditions and Values
Setting Stopping Criteria
Visualizing the Simulation
Running the Simulation

Recording a Macro
Editing the Macro
Reviewing the Edited Macro
Running the Macro
Plotting Data Using Line Probes
Exporting XY Plot Data
Further Training
Summary

Prerequisites

The instructions in the Simple Java Macros tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+

Importing the Volume Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Volume Mesh**.
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/geometry/data`.
5. Select `venturi.ccm` then click **Close**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `venturi.sim`.

Converting the Mesh to Two-Dimensions

STAR-CCM+ can use the surface of a three-dimensional mesh to create a two-dimensional mesh if the surface to be used lies on the X-Y plane at Z = 0.

Convert the mesh to two-dimensions:

1. Select **Mesh** > **Convert To 2D** from the menu.



2. In the *Convert Regions To 2D* dialog, set the following properties:

Property	Setting
Delete 3D regions after conversion	**Activated**
Tolerance	`1.0E-15`

3. Click **OK**.

Some work is now required to clean up the regions and associated boundaries.

4. Expand the **Regions** node and delete the redundant region **fluid 2D**.
5. Rename the **fluid 2D 2** node to **Fluid**.
6. Expand the **Regions** > **Fluid** > **Boundaries** node to see the boundaries that remain after the conversion process.



Combine these boundaries to form a single boundary and then split it into four separate boundaries:

7. Multiselect the **Default_Boundary_Region** node and the **cyclic 2** node.
8. Right-click and select **Combine**.

   There is now a single **Default_Boundary_Region** node.
9. Right-click the **Default_Boundary_Region** node and select **Split By Angle**.



The default settings in the **Split Boundaries by Angle** dialog are sufficient.

10. Click **OK**.
    Four boundaries appear as shown below:



11. Rename the boundaries as follows:

Old Name	New Name
Default_Boundary_Region	`Axis`
Default_Boundary_Region 2	`Outlet`
Default_Boundary_Region 3	`Inlet`
Default_Boundary_Region 4	`Wall`

    At this stage, do not set the boundary types until you have selected the models for the analysis.

An additional physics continuum was created for the two-dimensional mesh during the mesh conversion process, which means that the original continuum for the three-dimensional mesh is not required.

12. Delete the **Continua** > **Physics 1** node.
13. Rename the **Physics 1 2D** node to `Physics 1`.
14. Save the simulation.

Selecting the Physics Models

Physics models define the physical variables and phenomena in the simulation. An incompressible isothermal turbulent flow is sufficient for this analysis. Use the segregated flow model with the k-omega turbulence model.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, in the *Enabled Models* group box, deactivate **Two-Dimensional**.
2. Select the following models in order:

Group Box	Model
Space	**Axisymmetric** **Gradients** (Selected automatically)
Time	**Steady**
Material	**Liquid**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Omega Turbulence**

3. Click **Close**.
4. To review the models, open the **Physics 1** > **Models** node.



5. Save the simulation.

Setting Boundary Conditions and Values

Specify a constant velocity in the X direction at the inlet, and a constant pressure at the outlet. Set the centerline boundary to be of type axis. This axis boundary must lie exactly on the Y = 0 line; project the axis to the appropriate plane.

1. Select the **Regions** > **Fluid** > **Boundaries** > **Axis** node and set **Type** to **Axis**.
2. Right-click the **Boundaries** > **Axis** node and select **Project to Axis** from the pop-up menu.



3. In the `Repair Axis Boundary` dialog, click **OK**.



The vertices move along the axis to Y = 0.

Specify conditions for the inlet and outlet boundaries:

4. Set the boundary types for the inlet and outlet:

Boundary	Type
Inlet	**Velocity Inlet**
Outlet	**Pressure Outlet**

5. Edit the **Boundaries** > **Inlet** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Turbulence Specification	*Method*	**Intensity + Length Scale**
Velocity Specification	*Method*	**Components**
Physics Values		
Turbulence Intensity > Constant	*Value*	`0.034`
Turbulence Length Scale > Constant	*Value*	`0.0107`
Velocity > Constant	*Value*	`[1.5, 0.0 0.0] m/s`

6. Edit the **Boundaries** > **Oultet** node and set the following properties:

Node	Property	Setting
Physics Conditions		

Node	Property	Setting
Turbulence Specification	*Method*	**Intensity + Length Scale**
Physics Values		
Turbulence Intensity > Constant	*Value*	`0.034`
Turbulence Length Scale > Constant	*Value*	`0.0107`

7. Save the simulation.

Setting Stopping Criteria

Set the stopping criteria for the simulation. By default, simulations are set to stop after 1000 iterations.

In this simulation, convergence is reached before then, so reduce the maximum number of iterations to 600.

1. Select the **Stopping Criteria > Maximum Steps** node and set *Maximum Steps* to `600`.
2. Save the simulation.

Visualizing the Simulation

Create a scalar scene to display the pressure field as the solution develops.

1. Create a scalar scene.
2. Navigate to the **Scalar Scene 1 > Displayers > Scalar 1 > Scalar Field** and set *Function* to **Pressure**.
3. Save the simulation.

Running the Simulation

Preparation of the simulation is now complete, and the simulation can be run.

1. Click (**Initialize Solution**) in the *Solution* toolbar.

2. Click (**Run**) in the *Solution* toolbar.
3. While the simulation is running, select the *Scalar Scene 1* tab at the top of the *Graphics* window to visualize the solution.
4. When the simulation has finished running, save it.

The scalar scene appears as follows:



Recording a Macro

Create a macro to generate line probes at equally spaced intervals along the axis of the venturi tube. This enables you to extract profiles of axial velocity that can be used to observe the recovery of the flow after the venturi section.

Follow the procedure:

1. Initiate recording of the macro.
2. Manually work through the steps to create the first line probe.
3. Manually add the line probe to the scalar scene.

4. Stop recording the macro.

The **macro** toolbar contains four actions as shown below.

Start recording the macro:

1. Click (**Start Recording...**) to initiate the recording process.
2. In the *Save* dialog, enter `createLineProbes` for the name of the java file to which the macro is saved.
3. Click **Save**.

   An extra page called *createLineProbes* appears in the *Output* window.

   From this point onwards, any recordable action you make in STAR-CCM+ generates lines of Java code in this page. Clicking the (**Pause Recording**) button pauses generation of the Java code. Clicking the same button resumes generation of the Java code.

   Clicking the (**Stop Recording**) button ends all generation of the Java code and write the macro to file. Do not click either of these buttons for now.

Work through the required actions:

4. Right-click the **Derived Parts** node and select **New Part** > **Probe** > **Line** from the pop-up menu.
5. In the *Create Line Probe* dialog, set the following properties:

Property	Setting
Input Parts	**Fluid** (default)
Point 1	`[0.04, 0.0, 0.0] m`
Point 2	`[0.04, 0.0762, 0.0] m`
Resolution	`20`
Display	**New Geometry Displayer** (default)

6. Click **Create** then **Close**.

7. Rename the **Derived Parts > line-probe** node to `X = 0.04`.
8. Select the **Scenes > Scalar Scene 1 > Displayers > Scalar 1 > Partsnode** and set *Parts* to **Fluid > X = 0.04**.

   This completes the steps you wish to include in the macro.

To stop the macro:

9. Click the (**Stop Recording**) button in the macro toolbar.

Editing the Macro

Examine the recorded macro to identify where you can insert a loop for the creation of multiple line probes.

Note: The layout of the macro presented in the following sections may be different from what appears in your editor.

- Initial lines of the macro
- Initial sections of the execute0 method
- Code that obtains the Scene object
- Code to be repeated multiple times
- Adding the loop
- Displaying the new parts

Contents

Initial Lines of the Macro
Initial Sections of the execute0 Method
Code that Obtains the Scene Object
Code to be Repeated Multiple Times
Adding the Loop
Displaying the New Parts

Initial Lines of the Macro

Open the file `createLineProbes.java` in an appropriate text editor. The first few lines of the macro specify the package to which the macro

class belongs and the imported libraries that STAR-CCM+ use to compile the macro.

```
// STAR-CCM+ macro: createLineProbes.java
package macro;

import java.util.*;

import star.common.*;
import star.base.neo.*;
import star.vis.*;
```

After this initial section, see the class specifier.

```
public class createLineProbes extends StarMacro {
```

A macro recorded by STAR-CCM+ is a fully specified Java class derived from the `StarMacro` base class. The parts of this class specifier are briefly explained as follows:

- `public` informs the Java compiler that there are no access restrictions on this class.
- `class` is the keyword used to define a class.
- `createLineProbes` is the name of the class. This name must match the name of the file containing the class.
- `extends` indicates that the class is derived from another class (and inherits all the public and protected methods of that class).
- `StarMacro` is the name of the base class from which this one is derived.
- `{` is the opening brace of the class. It is matched by a closing brace at the end of the class.

The `createLineProbes` class contains one method called `execute`. All macros must contain this method, and this is the method that STAR-CCM+ calls to run the macro:

```
public void execute() {
```

The parts of the method signature are as follows:

- `public` is the access specifier and indicates that there are no restrictions on calling this method.
- `void` indicates that this method does not return any value.
- `execute()` is the name and argument list of the method. There are no arguments for this method.

- `{` is the opening brace of the method, and must be matched by a closing brace at the end.

The body of the `execute` method simply calls the private method, `execute0`. Consideration of `execute0` now follows.

Initial Sections of the execute0 Method

Examine the initial sections of the `execute0` method.

The first two sections of the method are:

```
Simulation simulation_0 =
  getActiveSimulation();

Units units_0 =
  simulation_0
  .getUnitsManager()
  .getPreferredUnits(
    new IntVector(
      new int[] {0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0}
  )
);
```

where:

```
simulation_0
.getUnitsManager()
.getPreferredUnits(
  new IntVector(
    /* etc */
```

is equivalent to

```
simulation_0.getUnitsManager().getPreferredUnits(new
IntVector(/* etc */
```

The call to `getActiveSimulation()` returns the Java object that represents the simulation currently active in the STAR-CCM+ client window. This is the main object in the Java macro and provides access to many other object managers (for example `RegionManager` or `RepresentationManager`). The Java API documentation provides details on all methods available on the `Simulation` class.

The `simulation_0` object is used to obtain the `UnitsManager` object. The `UnitsManager` object is used to find the `Units` object that represents the default unit for length in the active simulation (through the call to `getPreferredUnits`). The default unit for length is required when specifying the two points used in creating line probes.

Code that Obtains the Scene Object

Examine the code that obtains the `scene` object.

There are several lines of code after the call on the `UnitsManager` object that obtains the `Scene` object (as `scene_0`) corresponding to *Scalar Scene 1* in the *Graphics* window. The **Fluid** region that is displayed in this scene is obtained using the `RegionManager` and is referred to in the code as `region_0` , which is a `Region` object.

```
Scene scene_0 =
  simulation_0
  .getSceneManager()
  .getScene("Scalar Scene 1");

Region region_0 =
  simulation_0
  .getRegionManager()
  .getRegion("Fluid");
```

Some additional code is present that defines the particular representation being used in the `PartDisplayer` :

```
PartDisplayer partDisplayer_1 =
  ((PartDisplayer)
    scene_0
    .getCreatorDisplayer()
  );

FvRepresentation fvRepresentation_0 =
  ((FvRepresentation)
    simulation_0
    .getRepresentationManager()
    .getObject("Volume Mesh")
);

partDisplayer_1
.setRepresentation(fvRepresentation_0);
```

Code to be Repeated Multiple Times

You now come to the code to be placed within a loop. Some Java comments have been added to help describe the operations taking place. Items to be replaced by variables are underlined, and sections to be removed are underlined.

```
// The LinePart object is the Derived Part you want to
// create. The PartManager on the simulation_0 object
// has the method that lets you do this.

LinePart linePart_0 =
```

```
  simulation_0
  .getPartManager()
  .createLinePart(
    new NeoObjectVector(
      new Object[] {region_0}
  ),
  new DoubleVector(
    new double[] {0.04, 0.0, 0.0}
  ),
  new DoubleVector(
    new double[] {0.04, 0.0762, 0.0}
  ),
  20
);

// The LinePart object must now be told
// what coordinate system it is using. You retrieve the
// LabCoordinateSystem object from the
// CoordinateSystemManager on the simulation_0 object.

LabCoordinateSystem labCoordinateSystem_0 =
  ((LabCoordinateSystem)
    simulation_0
    .getCoordinateSystemManager()
    .getLabCoordinateSystem();

// As you now have the labCoordinateSystem_0 object we
// can tell linePart_0 to use it.

linePart_0
.setCoordinateSystem(labCoordinateSystem_0);

// The next two sections set the units on each axis,
// and the position, of the first Point of linePart_0

Coordinate coordinate_0 =
  linePart_0.getPoint1Coordinate();

coordinate_0
.setCoordinate(
  units_0,
  units_0,
  units_0,
  new DoubleVector(
    new double[] {0.04, 0.0, 0.0}
  )
);

// The next two sections set the units on each axis,
// and the position, of the second Point of linePart_0.

Coordinate coordinate_1 =
  linePart_0
  .getPoint2Coordinate();
```

```
coordinate_1
.setCoordinate(
  units_0,
  units_0,
  units_0,
  new DoubleVector(
    new double[] {0.04, 0.0762, 0.0}
  )
);

// The following two sections are updating the visible
// scene and are not required for creating line parts.
// Note: The part number may differ slightly in your
// recording.

PartDisplayer partDisplayer_2 =
  ((PartDisplayer)
    scene_0
    .getHighlightDisplayer()
);

partDisplayer_2
.getParts()
.setObjects(linePart_0)

// Change the name of the linePart.

linePart_0.setPresentationName("X = 0.04");
```

This section is pruned to remove the redundant code and placed within a loop. Code shown as underlined is replaced by appropriate variables that are incremented within the loop.

The redundant code shown above is dealing with a `PartDisplayer` from `scene_0`, and has nothing to do with actually creating the line part. This can all safely be omitted.

Adding the Loop

Add a loop to the macro that renames the line probes that are created.

The additional code to be added to the macro is described in plain language as follows:

- Set the minimum X coordinate, XMIN.
- Set the maximum X coordinate, XMAX.
- Set the required X distance between line probes, XDIST.
- Initialize the current X location, XLOC = XMIN + XDIST.
- WHILE (XLOC is less than XMAX)

  - Create a line probe at XLOC
  - Rename the new line probe with XLOC in the title
  - XLOC = XLOC + XDIST
- END WHILE

To avoid having to set XMIN and XMAX directly in the macro, use the `getRegionExtents` method of the `Representation` class to set appropriate values. The method in the Java API documentation for the `Representation` class, shows that it requires the `Region` object for which extents are to be obtained.

In your editor window, replace the section of code considered previously with the modified code that now follows. Original code is shown in plain text and new or modified code is shown in bold. Enter the new code exactly as shown, as Java is a case-sensitive programming language.

```
// Start of execute0 function

Simulation simulation_0 =
  getActiveSimulation();

Units units_0 =
  simulation_0
  .getUnitsManager()
  .getPreferredUnits(
    new IntVector(
      new int[]
        {0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0}
    )
  );

Scene scene_0 =
  simulation_0
  .getSceneManager()
  .getScene("Scalar Scene 1");

PartDisplayer partDisplayer_1 =
  ((PartDisplayer) scene_0
  .getCreatorDisplayer());

FvRepresentation fvRepresentation_0 =
  ((FvRepresentation) simulation_0
  .getRepresentationManager()
  .getObject("Volume Mesh"));

partDisplayer_1
  .setRepresentation(fvRepresentation_0);

Region region_0 =
  simulation_0
```

```
  .getRegionManager()
  .getRegion("Fluid");

// New code begins here
// Define preliminary variables

double dXDIST = 0.1; // Distance between probes

// Use the FvRepresentation to find the extents of
region_0
// (which was retrieved from simulation_0 earlier).
// FvRepresentation is derived from the Representation
// class and provides it's own version of the
// getRegionExtents method specifically for volume
// regions.
// getRegionExtents returns an array of six doubles:
// [xmin,xmax,ymin,ymax,zmin,zmax]

double[] dExtents =
  fvRepresentation_0
  .getRegionExtents(region_0);

// Retrieve xmin and xmax from dExtents

double dXMIN = dExtents[0];
double dXMAX = dExtents[1];

// Set the initial X location

double dXLOC = dXMIN + dXDIST;

// Start the loop
while(dXLOC < dXMAX) {
// Comes the section from the recorded macro
// Indent it further by two spaces and change
// all references to 0.04 with the variable dXLOC

LinePart linePart_0 =
  simulation_0
  .getPartManager()
  .createLinePart(
    new NeoObjectVector(
      new Object[] {region_0}
  ),
  new DoubleVector(
    new double[] {dXLOC, 0.0, 0.0}
  ),
  new DoubleVector(
  new double[] {dXLOC, 0.0762, 0.0}
  ),
  20
);

LabCoordinateSystem labCoordinateSystem_0 =
  simulation_0
```

```
  .getCoordinateSystemManager()
  .getLabCoordinateSystem();

linePart_0
  .setCoordinateSystem(labCoordinateSystem_0);

Coordinate coordinate_0 =
  linePart_0.getPoint1Coordinate();

coordinate_0
.setCoordinate(
  units_0,
  units_0,
  units_0,
  new DoubleVector(
    new double[] {dXLOC, 0.0, 0.0}
  )
);

// The next two sections set the units on each axis,
// and the position, of the second Point of linePart_0.

Coordinate coordinate_1 =
  linePart_0
  .getPoint2Coordinate();

coordinate_1
.setCoordinate(
  units_0,
  units_0,
  units_0,
  new DoubleVector(
    new double[] {dXLOC, 0.0762, 0.0}
  )
);

// Redundant sections removed from here

// Change the name of the linePart.
// The new name is created using the current X
// location. The String.format method ensures
// a consistent appearance of the number. See
// the Java documentation for details.

String strPartName =
  String.format("X = %6.4f",dXLOC);

linePart_0.setPresentationName(strPartName);

// Update dXLOC and end the loop with a closing brace
dXLOC = dXLOC + dXDIST;
} // End of Loop
```

This completes the main section for creating line probes.

Displaying the New Parts

The remaining change is to display all of the line probes in the scalar scene. This is achieved by obtaining all of the derived parts from the `PartManager` and setting them on the appropriate displayer of object `scene_0` obtained previously. Delete the redundant line of code which is underlined below:

```
// Following line is redundant: comment it out as shown
// or else delete it.

// partDisplayer_2.getParts().setObjects();

// Following section should remain unchanged; it is
// retrieving the relevant displayer from scene_0

ScalarDisplayer scalarDisplayer_0 =
  ((ScalarDisplayer)
    scene_0
    .getDisplayerManager()
    .getDisplayer("Scalar 1")
  );

// Add more code to retrieve all Derived Parts
// For more information on the Java technology, see
// the Java documentation for Collections and Generics

PartManager partManager_0 =
  simulation_0.getPartManager();

Collection<Part> allDerivedParts =
  partManager_0.getObjects();

scalarDisplayer_0
.getParts()
.setObjects(allDerivedParts);

} // End of execute0()
} // End of createLineProbes
```

This completes modification of the recorded macro.

- Save the macro in your editor.

Reviewing the Edited Macro

The final macro is shown below:

```
// Start of execute0 function
```

```
Simulation simulation_0 =
  getActiveSimulation();

Units units_0 =
  simulation_0
  .getUnitsManager()
  .getPreferredUnits(
    new IntVector(
      new int[]
        {0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0}
      )
    );

Scene scene_0 =
  simulation_0
  .getSceneManager()
  .getScene("Scalar Scene 1");

PartDisplayer partDisplayer_1 =
  ((PartDisplayer) scene_0
  .getCreatorDisplayer());

FvRepresentation fvRepresentation_0 =
  ((FvRepresentation) simulation_0
  .getRepresentationManager()
  .getObject("Volume Mesh"));

partDisplayer_1
.setRepresentation(fvRepresentation_0);

Region region_0 =
  simulation_0
  .getRegionManager()
  .getRegion("Fluid");

double dXDIST = 0.1; // Distance between probes

double[] dExtents =
  fvRepresentation_0
  .getRegionExtents(region_0);


double dXMIN = dExtents[0];
double dXMAX = dExtents[1];

double dXLOC = dXMIN + dXDIST;

// Start the loop
while(dXLOC < dXMAX) {
  LinePart linePart_0 =
    simulation_0
    .getPartManager()
    .createLinePart(
      new NeoObjectVector(
        new Object[] {region_0}
```

```
      ),
      new DoubleVector(
        new double[] {dXLOC, 0.0, 0.0}
      ),
      new DoubleVector(
        new double[] {dXLOC, 0.0762, 0.0}
      ),
      20
    );

LabCoordinateSystem labCoordinateSystem_0 =
  simulation_0
  .getCoordinateSystemManager()
  .getLabCoordinateSystem();

linePart_0
  .setCoordinateSystem(labCoordinateSystem_0);

Coordinate coordinate_0 =
  linePart_0.getPoint1Coordinate();

coordinate_0
.setCoordinate(
  units_0,
  units_0,
  units_0,
  new DoubleVector(
    new double[] {dXLOC, 0.0, 0.0}
  )
);

Coordinate coordinate_1 =
  linePart_0
  .getPoint2Coordinate();

coordinate_1
.setCoordinate(
  units_0,
  units_0,
  units_0,
  new DoubleVector(
    new double[] {dXLOC, 0.0762, 0.0}
  )
);

String strPartName =
  String.format("X = %6.4f",dXLOC);

linePart_0.setPresentationName(strPartName);

// Update dXLOC and end the loop with a closing brace
dXLOC = dXLOC + dXDIST;
} // End of Loop

// partDisplayer_2.getParts().setObjects();
```

```
ScalarDisplayer scalarDisplayer_0 =
  ((ScalarDisplayer)
    scene_0
    .getDisplayerManager()
    .getDisplayer("Scalar 1")
  );

PartManager partManager_0 =
  simulation_0.getPartManager();

Collection<Part> allDerivedParts =
  partManager_0.getObjects();

scalarDisplayer_0
.getParts()
.setObjects(allDerivedParts);

  } // End of execute0()
} // End of createLineProbes
```

Running the Macro

Run the macro in the STAR-CCM+ client.

Before running the macro, delete the single-line probe created earlier:

1. Right-click the **Derived Parts > X = 0.04** node and select **Delete** from the pop-up menu.
2. Click **Yes** in the *Confirm Object Deletion* dialog that appears.

Run the macro:

3. In the *Macro* toolbar, click (**Play Macro...**).
4. In the *Open* dialog, select `createLineProbes.java`.
5. Click **Open**.

STAR-CCM+ first compiles the macro class, and if there are no errors reported, it calls the **execute** method. This results in the creation of 22 line probes, from **X = 0.1000** to **X = 2.2000**, which appear in *Scalar Scene 1*.



Plotting Data Using Line Probes

Use the line probes to create an X-Y plot of axial velocity profiles. Each data series presented on the plot corresponds to the data extracted along a single-line probe. The data contained in the plot are exported to a text file.

1. Right-click the **Plots** node and select **New Plot > X-Y** from the pop-up menu.
2. Rename the **Plots > XY Plot 1** node to `Velocity Profiles`.
3. Switch to the *scene/plot* pane by clicking the **scene/plot** button above the object tree.



The X-axis of the plot shows the Y position data from the line probe, and the Y-axis of the plot shows the axial velocity at those Y positions.

4. Edit the **Velocity Profiles** node and set the following properties:

Node	Property	Setting
Velocity Profiles	*Title*	`Velocity Profiles`
X Type > Vector Quantity	*Value*	`[0.0, 1.0, 0.0] m`
Y Type > Y Type 1 > Scalar Function	*Field Function*	**Velocity > Lab Reference Frame > Laboratory > i**
Legend	*Position*	**East**
	Layout	**Vertical**
	Font	**Lucida Sans 8 Plain**

5. Click **OK**.

The final operation is to set the **Parts** property for the plot.

6. Select the **Velocity Profiles** node at the top of the tree.
7. In the *Properties* window, edit the *Parts* property.
8. In the *Velocity Profiles - Parts* dialog, select the **Derived Parts** node.
9. Expand the node and make sure that all parts are selected.

10. Click **OK**.

The final plot is displayed in the *Velocity Profiles* window.



Exporting XY Plot Data

The data contained in the *Velocity Profiles* plot can be exported to a text file.

1. Right-click the background of the *Velocity Profiles* plot and select **Export** from the pop-up menu.



2. In the *Save* dialog, set *File Name* to `exported_velocities`.
3. Click **Save**.

All data contained in the *Velocity Profiles* plot are written to a Comma-Separated Variable (CSV) file. This type of file can be imported into spreadsheet software for further analysis of the data.

Further Training

CD-adapco provides an in-depth training course on the use of Java macros and the automation of tasks in STAR-CCM+. Contact your local support agent for details.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Importing a three-dimensional mesh intended for use in a two-dimensional simulation.
- Converting a three-dimensional mesh to two-dimensions.
- Cleaning up regions and boundaries after the conversion.

- Splitting a boundary by angle and renaming the subsequent nodes.
- Selecting the axisymmetric space model.
- Setting an axial boundary type and projecting it to the Y axis.
- Recording a macro while creating a single-line probe and adding it to a scalar scene.
- Editing a previously recorded macro to add a loop.
- Including code to add all derived parts to a scalar scene.
- Running a macro.
- Creating an XY plot based on line probes.
- Exporting the data from an XY plot to a CSV file.

Intermediate Java Macros: Run Multiple Simulations

This tutorial simulates a train in a cross-wind, to calculate the drag coefficient on the train for different angles of attack.

One way in which Java macros can be used effectively is to run a number of simulations in which a single variable or a set of variables is altered in each one.



The train is modeled under field test conditions. The resultant velocity of the wind as seen from the train, is calculated based on a 22 m/s head wind and a 6 m/s cross-wind at angle of attack. In addition, you set a constant velocity on the track that is equivalent to the velocity of the train. The following diagram shows how the resultant velocity is calculated:



A Java macro runs a set of simulations consecutively, each with a different set of boundary conditions. The macro performs the following:

Read the boundary condition values from an external text file;

Set the new boundary conditions in each simulation;

Run the simulation;

Export data, scenes, and plots for post-processing.

The tutorial introduces the use of multiple classes to organize your macro and includes a fully commented Java file for reference purposes.

Contents

Prerequisites
Importing the Surface Mesh and Naming the Simulation
Visualizing the Initial Surface

Setting Up the Mesh
Specifying Boundary Types
Generating the Mesh
Setting Up the Physics Continuum
Setting Global Initial Conditions
Specifying Inflow and Outflow Conditions
Creating a Drag Coefficient Report
Creating a Velocity Magnitude Scalar Scene
Creating a Scalar Scene with Streamlines
Designing the Macro
Recording the Initial Macro
Creating the Final Macro
Testing and Debugging
Running the Macro
Summary

Prerequisites

The instructions in the Intermediate Java Macros tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+
Recording and editing Java macros	Simple Java Macros: Post Processing

Importing the Surface Mesh and Naming the Simulation

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File** > **Import** > **Import Surface Mesh**.

4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/automation/data`.
5. Select `trainAndTrack.dbs` then click **Open**.
6. In the *Import Surface Options*, set *Import Mode* to **Create New Region**.
7. Click **OK**.
8. Save the simulation as `train.sim`.

Visualizing the Initial Surface

Visualize the initial surface of the geometry.

In this example, all boundaries are contained within a single region, **Region 1**. During the initial model preparation, the surface was divided into separate boundaries and these were given appropriate names.

1. Expand the **Regions** > **Region 1** > **Boundaries** node.
2. Select the **TrainEngine** node.

   The **TrainEngine** boundary is highlighted on the model.



3. Select each of the boundaries and observe which part is highlighted.

Remove two of the boundaries from the current scene to view the mesh on the internal features.

4. Edit the **Displayers** > **Geometry 1** > **Parts** node.

5. In the **Edit** dialog, expand the **Boundaries** nodes and deactivate the **Inflow** and **Top** boundaries.
6. Click **OK**.

It is possible to toggle the visibility of any displayer in a scene to disable a particular displayer without deleting it from a scene.

7. Right-click the **Outline 1** node and select **Toggle Visibility**.

   The outline is no longer displayed in the scene.

8. In the *Vis* toolbar, click (**Show All Meshes**) to display the mesh.

9. In the *Vis* toolbar, click (**Make Scene Transparent**) to restore the scene transparency.



This type of mesh is generated when neutral format data or native CAD data is imported into STAR-CCM+, and is known as a tessellated surface mesh.

Setting Up the Mesh

To calculate the drag coefficient of the train as it is subjected to cross-winds at different angles accurately there must be enough cells near to the train, which requires a relatively fine mesh.

You can minimize the total number of cells in the mesh by generating a fine mesh around the train, and having a coarse mesh in the remainder of the fluid domain. Using different mesh densities is a good practice for mesh generation, promoting faster solution convergence, and effective use of processing power. In external aerodynamics cases, it is common to have a high mesh cell count (>1 million cells). However, the mesh generated in this tutorial is much lower to decrease the run-time.

Contents

Inspecting the Initial Surface Mesh
Selecting Meshing Models
Setting the Global Mesh Reference Values
Setting up Wake Refinement
Specifying Boundary Mesh Settings

Inspecting the Initial Surface Mesh

Inspect the quality of the initial surface mesh before selecting meshing models.

1. Click **simulation** to return to the main object tree.
2. Expand the **Representations** node.
3. Right-click the **Import** node and select **Repair Surface**.



The *Surface Repair Options* dialog appears.



4. Click **OK** to accept the default settings.
5. In the inset *edit* panel, click **Manage Thresholds**
6. In the *Create/Modify Thresholds* dialog, make sure that all thresholds are activated.



7. Click **OK**.
   You can now assess the quality of the mesh.



There are no pierced faces, free edges, or non-manifold edges and vertices which means that the mesh is of sufficient quality to allow you to proceed. A tessellated surface mesh import, such as the one used in this example, often contains poor quality and close proximity faces. The remaining issues can be fixed using the Surface Remesher.

8. Click **Close**.

Selecting Meshing Models

Select the meshing models for the imported geometry.

Use the Surface Remesher to generate the surface mesh. Use the Trimmer and Prism Layer Mesher to generate the volume mesh.

1. For the mesh continuum, **Continua** > **Mesh 1**, select the following models in order:

Group Box	Models
Surface Mesh	**Surface Remesher**
Volume Mesh	**Trimmer**
Optional Models	**Prism Layer Mesher**

2. Click **Close**.

Setting the Global Mesh Reference Values

Set the global reference values.

In this example, these values set the size of the coarser region of the mesh.

1. Edit the **Mesh 1 > Reference Values** node and set the following properties:

Node	Property	Setting
Base Size	*Value*	`5 m`
Prism Layer Thickness > Relative Size	*Percentage of Base*	`1`
Template Growth Rate	*Default Growth Rate*	**Medium**
	Boundary Growth Rate	**Medium**

2. Click **Close**.

Setting up Wake Refinement

Use the wake refinement tool to refine the mesh in the region behind the train by specifying a boundary from which the wake extends. In this case, the boundary is **TrainEngineRear**.

1. Expand the **Regions > Region 1 > Mesh Values** node.
2. Right-click the **Trimmer Wake Refinement** node and select **New**.

   A new wake refinement set has been created called **Wake Refinement Set 1**.



Specify the properties of the mesh in the region of the wake.

3. Edit the **Wake Refinement Set 1** node and set the following properties:

Node	Property	Setting
Wake Refinement Set 1	*Boundaries/Features Curves*	**TrainEngineRear**
	Distance	`5.0 m`
	Direction	`[0, 1, 0]`
Relative Size	*Percentage of Base*	`6`

4. Click **Close**.

Specifying Boundary Mesh Settings

Apply custom mesh settings to the train engine, wheels, and rails boundaries to make sure a suitably fine mesh around the train and the rails is generated.

Mesh settings at the boundary level supersede settings at the global level. Custom mesh values for boundaries affect only the Surface Remesher.

1. Select the **Region 1** > **Boundaries** > **TrainEngine** > **Mesh Conditions** > **Custom Surface Size** node and activate *Custom surface size*.
2. Edit the **TrainEngine** > **Mesh Values** node and set the following properties:

Node	Property	Setting
Surface Size > **Relative Minimum Size**	*Percentage of Base*	`3`
Surface Size > **Relative Target Size**	*Percentage of Base*	`3`

3. Click **Close**.

Apply the same settings to the **Wheels** and **TrainEngineRear** boundaries using the drag-and-drop functionality.

4. Click-and-hold the **TrainEngine** node and drag the mouse pointer down onto the **Wheels** node. Make sure the mouse pointer changes to show the plus symbol before releasing the mouse button.



5. Repeat for the **TrainEngineRear** boundary. Drag-and-drop the **TrainEngine** node onto the **TrainEngineRear** node.

You can verify that the mesh settings have been copied by expanding the relevant nodes and inspecting the properties.

6. Expand the **Wheels** > **Mesh Values** > **Surface Size** node and make sure that the **Relative Minimum Size** and **Relative Target Size** are set to `3`.

Specify a custom surface size for the **Rails** boundary.

7. Select the **Rails** > **Mesh Conditions** > **Custom Surface Size** node and tick the box next to *Custom surface size*.
8. Expand the **Mesh Values** > **Surface Size** node and set the **Relative Minimum Size** and **Relative Target Size** to `7%`.

Specifying Boundary Types

Specify the boundary types for the simulation.

It is good practice to set boundary types before generating the mesh. The default boundary type is Wall, and when the Prism Layer Mesher is used, prism layers grow from wall boundaries. Prism layers are not generated for non-wall boundaries unless they are explicitly chosen using the **Mesh Conditions** node for each boundary.

1. Set the following boundary types:

Boundary	Type
Ground	**Symmetry Plane**
Top	**Symmetry Plane**
Inflow	**Velocity Inlet**
Outflow	**Pressure Outlet**

2. Save the simulation.

Generating the Mesh

The mesh parameters are fully defined and you can proceed to generate the mesh. It is good practice to generate the surface mesh and examine that before proceeding with the volume mesh.

1. Generate the surface mesh.

   The surface mesh near the train is shown.



2. Generate the volume mesh.



3. Save the simulation.

Setting Up the Physics Continuum

Physics models define the physical variables and phenomena in the simulation. In this tutorial, model the flow as Steady and Turbulent. Use the Segregated Flow solver with the K-Epsilon Turbulence model.

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Optional Physics Models	**Cell Quality Remediation**

2. Click **Close**.

3. To review the models, open the **Physics 1** > **Models** node.



4. Save the simulation.

Setting Global Initial Conditions

Set the initial flow field conditions. Retain the default value for the initial pressure condition, and specify custom values for turbulence.

1. Select the **Physics 1** > **Initial Conditions** > **Turbulence Specification** node and set *Method* to **Intensity + Length Scale**.
2. Select the **Turbulent Length Scale** > **Constant** node and set the *Value* to `2.5 m`.

Specifying Inflow and Outflow Conditions

Set the boundary conditions on the inflow, outflow, rails, and railway bed boundaries. The inflow velocity is different for each cross-wind flow angle and so is set in the Java macro.

1. Expand the **Regions** > **Region 1** > **Boundaries** node and set the following properties:
2. Edit the **Inflow** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Turbulence Specification	*Method*	**Intensity + Length Scale**
Velocity Specification	*Method*	**Components**
Physics Values		
Turbulent Length Scale > Constant	*Value*	`2.5 m`

Note: You set the velocity specification method at the inlet to record the steps to the macro for later use.

3. Edit the **Outlfow** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Physics Conditions	*Method*	**Intensity + Length Scale**
Backflow Direction Specification	*Method*	**Extrapolated**
Physics Values		
Turbulent Length Scale > Constant	*Value*	`2.5 m`

Note: You change the backflow direction specification method at the **Outflow** boundary to allow the fluid to flow parallel to the outlet boundary.

4. Edit the **Rails** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Tangential Velocity Specification	*Method*	**Vector**
Physics Values		

Node	Property	Setting
Velocity > Constant	*Value*	`[0, 22, 0] m/s`

5. Edit the **RailwayBed** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Tangential Velocity Specification	*Method*	**Vector**
Physics Values		
Velocity > Constant	*Value*	`[0, 22, 0] m/s`

6. Save the simulation.

Creating a Drag Coefficient Report

Create a drag coefficient report to obtain the value for the drag coefficient on the train at each cross-wind flow angle. Solution convergence can be monitored using this report also.

1. Right-click the **Reports** node and select **New Report > Force Coefficient**.
2. Rename the **Force Coefficient 1** node to `Drag Coefficient`.
3. Edit the **Drag Coefficient** node and set the following properties:

Property	Setting
Reference Velocity	`22 m/s`
Reference Area	`10 m^2`
Direction	`[0, 1, 0]`
Parts	**TrainEngine** **TrainEngineRear** **Wheels**

4. Right-click the **Drag Coefficient** node and select **Create Monitor and Plot from Report**.

A drag coefficient monitor is created along with a plot to show the drag coefficient calculated after each iteration.

Creating a Velocity Magnitude Scalar Scene

Create a scalar scene to visualize the velocity magnitude.

1. Create a section plane with the following properties:

Property	Setting
origin	`[0, 0, 0.8] m`
normal	`[0, 0, 1] m`
Display	**No Displayer**

2. Click **Create** then **Close**.
3. Create a scalar scene and rename it to `Velocity Magnitude`.
4. Click **scene/plot**.
5. Edit the **Velocity Magnitude** > **Displayers** node and set the following properties:

Node	Property	Setting
Outline 1	*Outline*	**Deactivated**
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Derived Parts > plane section**
Scalar Field	*Function*	**Velocity Magnitude**

6. Click **Close**.
7. Initialize the solution.

   By initializing manually you are able to view the initialized solution, which helps in the following steps.

Adjust the position of the components in the scene. It is useful to adjust the scene at this stage before saving a hardcopy using the Java macro.

8. Position the scene to match the screenshot:



9. Click **simulation**.
10. Save the simulation.

Creating a Scalar Scene with Streamlines

Create a second scalar scene to display streamlines over the train.

1. Create a scalar scene and rename it to `Streamlines`.
2. Create a streamline derived part with the following properties:

Property	Setting
Seed Mode	**Line Seed**
Point 1	`[-7, -73, 0.4] m`
Point 2	`[20, -52, 0.4] m`
Resolution	`50`
Direction	**Both**

3. Click **Create** then **Close**.
4. Click **scene/plot**.
5. Edit the **Streamlines** > **Displayers** node and set the following properties:

Node	Property	Setting
Outline 1	*Outline*	**Deactivated**
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Ground** **Rails** **RailwayBed** **TrainEngine** **TrainEngineRear** **Wheels**
Scalar Field	*Function*	**Pressure**
Color Bar	*Color Map*	**grayscale**
Section Stream 1 > Scalar Field	*Scalar Field*	**Velocity Magnitude**

6. Click **Close**.
7. Adjust the scene so that it appears as follows:



8. Close the *Streamlines* scene.
9. Click **simulation**.
10. Save the simulation.

The case is now set up correctly and you are ready to begin Designing the Java Macro.

Designing the Macro

Carry out the following steps when designing a macro:

- Define the aims of the macro
- Create a simple plan for the macro
- Plan the actions to be recorded in a macro
- Record a macro in STAR-CCM+
- Examine the recorded macro and pick out useful statements
- Finalize the structure of the macro
- Write the final macro

Contents

Defining the Aims of the Macro
Creating a Simple Plan

Defining the Aims of the Macro

The first step in designing a Java macro is to define in one or two sentences, what you want the macro to do.

For the tutorial, the aim of the macro is to run a set of simulations with different cross-wind angles of attack, and provide the value of the drag coefficient on the train.

Creating a Simple Plan

With the aims of the macro clearly defined, you can start planning.

First, write down the steps required. For example, set the velocity to 5 m/s and run the simulation. These steps form the basis of a plan. Next, expand the plan to include other processes such as extracting data from an input file, looping through to repeat a set of actions, and opening a panel. Write the plan in simple terms.

To simplify the process of scripting the macro, you can record a simple Java macro in STAR-CCM+ and extract sections of the code.

To meet the aims that were defined previously, the following steps are required (processes that can be obtained by recording a macro in STAR-CCM+ are marked with a *):

- Read cross-wind data from an external source
- Clear any previous solution*
- Set the inlet velocity value*
- Set the initial velocity value*
- Run the simulation for some iterations*
- Calculate the drag coefficient on the train*
- Save the drag coefficient value for each cross-wind angle of attack in an external data file
- Save the simulation*
- Repeat for each cross-wind angle of attack

You have marked a number of actions for which the Java code can be obtained by recording a macro.

Recording the Initial Macro

Before you record the initial macro, it is good practice to have a clear plan of which actions you perform when recording the macro. The resulting Java file is clearer and less cluttered, and identifying which statements refer to a particular action is much easier.

The following sections cover:

- How to plan the actions to record
- Recording a macro
- Understanding the recorded macro

Contents

Planning Actions to Record
Recording the Macro
Understanding the Initial Macro

Planning Actions to Record

At this stage you can consider any additional actions to perform in the final macro, such as exporting plots and scenes.

Include additional steps in the plan to help with post-processing the results of each simulation. Export hardcopies of the Residuals plot and *Velocity Magnitude* scalar scene, as well as a STAR-View+ file of the *Streamlines* scene. Using this technique, the results of each simulation can be compared without needing to open each one individually.

Record the following actions:

1. Set the initial velocity
2. Set the inflow velocity
3. Clear any previous solution
4. Run for x iterations
5. Get a value for the drag coefficient
6. Save a hardcopy of the *Velocity Magnitude* scene and Residuals plot
7. Save a STAR-View+ file of the *Streamlines* scene
8. Save the simulation

Recording the Macro

Record the planned actions.

Any command issued to STAR-CCM+ has an associated Java statement. The **macroRecording** tab of the **Output** window displays the statements. Do not record unnecessary commands to avoid the Java file being cluttered. Pause the recording to carry out actions without the associated statements being added to the macro.

1. In the *Macro* toolbar, click **(Start Recording...)**.
2. Save the file as `macroRecording.java`.

Set the initial velocity:

3. Expand the **Continua** > **Physics 1** > **Initial Conditions** > **Velocity** > **Constant** node.
4. Set *Value* to `[-10, 10, 0] m/s`.

Set the inflow velocity:

5. Select the **Regions** > **Region 1** > **Boundaries** > **Inflow** > **Physics Values** > **Velocity** > **Constant** node.
6. Set the *Value* to `[-5, 5, 0] m/s`.

Clear the previous solution and run the simulation:

7. From the menu, select **Solution** > **Clear Solution**.



8. Run the simulation.
9. Click (**Stop Iterating**) to stop the simulation running.

Run the **Drag Coefficient** report. When you create the final macro, obtain a single value for the drag coefficient rather than running the entire report. To obtain some of the statements required later on, however, run the report.

10. Right-click the **Reports** > **Drag Coefficient** node and select **Run Report**.



Save hardcopies of the scene:

11. Select the **Velocity Magnitude** tab in the **Graphics** window.
12. Right-click in the white area of the scene and select **Hardcopy**.



13. In the *Save To File* dialog, set the following properties:

Property	Setting
File Name	`velMag`
Use Screen Resolution	**deselected**
X Resolution	`1024`
Y Resolution	`768`

14. Click **Save**.
15. Repeat these steps to save a hardcopy of the *Residuals* plot; name the file `res`.
16. Right-click the **Scenes > Streamlines** node and select **Export 3D Visualization File**.



17. Name the file `Streamlines` and click **Save**.
18. Click the **Save As** button in the **File** menu and save the file under the name `trainMacro`.
19. Click **Save**.
20. Click the ■ (**Stop Recording**) button in the *Macro* toolbar to stop recording.

Understanding the Initial Macro

After recording the macro, it is good practice to open the file and noting which statements correspond to which actions.

To add notes to the Java file type `//"and your comment"` . The following steps examine the initial macro and describe which statements correspond to the actions you recorded in the previous section.

Open the `macroRecording.java` file in an appropriate text editor. Refer to the first section of the Java macro. The first set of statements in the macro conforms to the template of a standard STAR-CCM+ macro.

Code	Explanation

`// STAR-CCM+ macro: macroRecording.java`
`package macro;`
 	
`import java.util.*;`
 	
`import star.common.*;`
`import star.base.neo.*;`
`import star.vis.*;`
`import star.base.report.*;`
`import star.flow.*;`
 	
`public class macroRecording extends StarMacro {`
 	
`  public void execute() {`
`  execute0();`
`}`
 	
`  private void execute0()`
`{`
	• The class is defined as belonging to the `macro` package.
	• A number of libraries are imported which are required to compile the macro.
	• A class called `macroRecording` is declared, and is specified as a subclass of the `StarMacro` class.
	• The main method, `execute()`, is declared. When the macro is played in STAR-CCM+ it uses `execute()`, which calls `execute0()`.

Look at the body of statements that were added to the macro during the record operation. Examine the body of the `execute0()` method. Find each statement as it is described, and add a comment in the text editor to describe its function.

Code	Explanation
`Simulation simulation_0 = getActiveSimulation();{`	• This statement is essential for STAR-CCM+ macros that interact with and manipulate a simulation. A variable of type `Simulation` is declared, called `simulation_0`. A reference to the location of the active simulation object is obtained using the `getActiveSimulation()` method and stored in the `simulation_0` variable. Add

Code	Explanation
	the comment: "`// Retrieve the current simulation`".

`PhysicsContinuum physicsContinuum_0 = ((PhysicsContinuum) simulation_0 .getContinuumManager() .getContinuum("Physics 1") );`
 	
`VelocityProfile velocityProfile_0 = physicsContinuum_0 .getInitialConditions() .get(VelocityProfile.class );`
 	
`velocityProfile_0 .getMethod(ConstantVectorProfileMethod.class) .getQuantity() .setComponents(-10.0, 10.0, 0.0);`
	• These statements set the initial velocity. Setting values using a macro follows the same logical sequence as expanding nodes in the object tree. First you find the physics continuum **Physics 1** in the active simulation, `simulation_0`, and you store a reference to its location in the variable `physicsContinuum_0`. Then you look in `physicsContinuum_0` for the initial velocity condition, and store a reference to its location in `velocityProfile_0`. Finally you set the new velocity condition using the `setComponents()` method.

Each of the entities in STAR-CCM+ is an object, for example boundaries, continua and even the active simulation. To interact with such objects, such as setting a constant value, methods exist which allow you to locate the object in question. To locate an object, follow the hierarchy of the object tree. For example, the initial velocity object is contained in a physics continuum object, which is contained in a simulation object. To interact with such an object, create a variable of the relevant type to store a reference to the location of the object.

Code	Explanation

`Region region_0 = simulation_0 .getRegionManager() .getRegion("Region 1");`
 	
`Boundary boundary_0 = region_0 .getBoundaryManager() .getBoundary("Inflow");`
	• Here you set the inflow velocity. Find **Region 1** in the active simulation, then find the **Inflow** boundary within that region, the **Velocity** condition within that boundary, and finally set the new velocity. To find the region you use a `RegionManager` , and to find a

Code	Explanation

`VelocityProfile velocityProfile_1 = boundary_0 .getValues() .get(VelocityProfile.class);`
 	
`velocityProfile_1 .getMethod(ConstantVectorProfileMethod.class) .getQuantity() .setComponents(-5.0, 5.0, 0.0);`
	boundary in the region you use the `BoundaryManagerbelonging` to that region. A second variable of type `VelocityProfile` has been created and has been numbered incrementally to `1`.

`Solution solution_0 = simulation_0 .getSolution();`
 	
`solution_0.clearSolution();`
	• These statements retrieve any solution belonging to the current simulation and clear it by calling the `clearSolution()` method.

`ResidualPlot residualPlot_0 = ((ResidualPlot) simulation_0 .getPlotManager() .getObject("Residuals"));`
 	
`residualPlot_0.setTitleFont( new java.awt.Font( "SansSerif", 0, 12));`
	• This set of statements precedes the run simulation command, creating a *Residuals* plot in which to display the simulation progress.

`simulation_0.getSimulationIterator() .run();`	• This statement is used to run the simulation.

Code	Explanation

`ForceCoefficientReport forceCoefficientReport_0 =`
 	
`((ForceCoefficientReport)`
 	
`simulation_0.getReportManager()`
	• Here you find the drag coefficient report, and store a reference to its location in the variable `forceCoefficientReport_0`. The `printReport()` method is

Code	Explanation

`.getReport("Drag Coefficient")`
`);`
 	
`forceCoefficientReport_0.printReport();`
	then used to run the report within STAR-CCM+. You can use various alternative methods here, for example `getReportMonitorValue()` - this yields the value of the drag coefficient.

`Scene scene_0 =`
`simulation_0`
`.getSceneManager()`
`.getScene("Velocity Magnitude");`
 	
`CurrentView currentView_0 =`
`scene_0`
`.getCurrentView();`
`currentView_0.setInput(`
`new DoubleVector(`
`new double[]{-2.37, -46.74, -5.45}`
`), new DoubleVector(`
`new double[]{-2.37, -46.74, 152.76}`
`), new DoubleVector(`
`new double[]{0.0, 1.0, 0.0}`
`), -1.0, 0`
`);`
 	
`scene_0.printAndWait(`
`resolvePath("/myTutorials/velMag.png`
`), 1, 1024, 768`
`);`
	• This set of statements was generated when the hardcopy of the *Velocity Magnitude* scene was saved. The first statement finds the scene and assigns it to the variable `scene_0`. The lines that follow set the current view of the scene, however they are not necessary in this case and can be removed. The next statement uses the `printAndWait()` method to print the image to file. Notice that the resolution 1024x768 that was entered when the macro was recorded is included in this statement.

`residualPlot_0.encode(`
`resolvePath("/myTutorials/res.png"`
`), "png", 1024, 768`
`);`
	• This set of statements is used to save a copy of the *Residuals* plot.

`Scene scene_1 =`
`simulation_0`
`.getSceneManager()`
`.getScene("Streamlines")`
	• Here you save a STAR-View+ scene file of the `Streamlines` scene by using the

Code	Explanation

``;
 	
scene_1.open(true);
 	
scene_1.export3DSceneFileAndWait(
  resolvePath("/myTutorials/streamlines.sce"
  ), true
);```
`export3DSceneFileAndWait()` method.	

```simulation_0.saveState(resolvePath("/myTutorials/trainMacro.sim"));``` • Finally, the simulation is saved under the name `trainMacro.sim`.	

To quickly find what you need when writing the final macro, add comments above each statement or set of statements at this stage.

Save the `macroRecording.java` file.

Creating the Final Macro

You are ready to write the final macro.

Writing the final macro can be approached in a number of different ways, and can be as complex as the task requires. Study the Java macro written to automate the tutorial case in the following steps, taking care to examine how the macro has been constructed. Proceed to the next section for a list of keywords relevant to this chapter, or jump to the following section to see how to finalize the macro structure.

Contents

Keywords
Finalizing the Structure of the Macro
Class-by-Class Breakdown
Understanding the Main Method in the Macro

Keywords

This topic contains a list of key words that are used in this tutorial when describing elements in the Java code.

Keyword	Definition

Nested Class	A class that has been defined within another class. Nested classes take the following form in Java:
	`class OuterClass {`
	`  class NestedClass {`
	`  }`
	`}`

Member Variable	A variable belonging to a class.
Method	A subroutine belonging to a class.
Instance Variable	A variable contained in an object.
Instance Method	A subroutine contained in an object.

Constructor	A constructor is used to create new objects. By default it allocates memory and initializes instance variables. Every class has at least one constructor; if a constructor definition has not been written in the class the system provides a default constructor. The following example shows a constructor that is used to set the values of the instance variables to the values of some input parameters.
	`public class MyClass {`
	`  private int m_num1;`
	`  private int m_num2;`
	`  public MyClass(int param1, int param2) {`
	`  m_num1 = param1;`
	`  m_num2 = param2;`
	`  }`
	`}`

Instantiate	The process of creating an object from a class is called instantiation, since an object is an instance of a class. You can say that to create an object you must instantiate the relevant class.

Finalizing the Structure of the Macro

It is good practice to finalize the macro structure before writing the final macro. Deciding how best to organize the macro, for example what

methods to use, which actions to include in a loop, and the number of classes required. These decisions are based on necessity and your Java programming experience.

It is also useful at this stage to consider whether the macro needs to be modified in the future to include extra commands and processes, for example setting additional boundary conditions. If likely, it is recommended that the macro is divided into different subsections which carry out each task rather than having everything in a single continuous section. When you add extra processes to the macro, you can locate the section easily, and insert the command. The object-oriented nature of the Java language makes it suited to this kind of organization. In this tutorial, nested classes are used to distinguish between steps in the process.

The main processes are divided as follows:

- a class to read the input data
- a class from which to create objects which store input data
- a class to set values in STAR-CCM+ and run the simulation
- a class to carry out post-processing tasks
- an additional class to write any data to an external file, for example the drag coefficient

These classes are given the following names, respectively:

- `DataReader`
- `SimData`
- `SimRunner`
- `PostProcessor`
- `DataWriter`

One key advantage in organizing the macro in this way is that the main method in the macro is easy to understand as the statements follow a logical sequence.

Using an Input Data File

Consider is how the macro reads the wind data and passed to the simulation. You want to be able to alter the data as required, and run as many simulations as necessary. For this reason, the input data is saved in a simple text file, and read into the simulation using the macro.

- Navigate to the `doc/startutorialsdata/automation/data` subdirectory of your STAR-CCM+ installation directory. Copy the `trainInput.txt` file to your working directory.
- Open the `trainInput.txt` file using an appropriate text editor.

The file contains the following data:

- the angle of attack of the cross-wind in degrees, `AngDeg`
- the velocity of the train, `VelTrn`
- the velocity of the cross-wind, `VelWnd`
- the corresponding initial velocities, `InitVelTrn`, and `InitVelWnd`

Each line in the file represents a set of data. Use the macro to read each line and store the data in an object. Store a collection of these objects in an array. When each simulation is run, the data is read from an object and passed efficiently to the simulation, without the need to reopen the input file.

Class-by-Class Breakdown

A class-by-class breakdown is provided for you. The final macro is too complicated to describe how to write it from scratch, so it has been provided for you.

- Navigate to the `doc/startutorialsdata/automation/data` subdirectory of your STAR-CCM+ installation directory. Copy the `trainFlowAngles.java` file to your working directory.
- Open the `trainFlowAngles.java` file using an appropriate text editor.

A detailed description which includes Java coding follows in the following sections. The processes in each nested class are as follows:

SimData

The data passed to the `SimData` object includes the velocities of the train and wind. You must calculate the x and y components of the inflow velocity and initial velocity to set the velocities in the simulation. These calculations are made in the `SimData` class.

- Receive the data from `DataReader`.
- Calculate the velocity components.

DataReader

- Read a line of data from the input file

- Assign each value to the corresponding variable
- Create a `SimData` object for the set of variables
- Store the `SimData` object in a collection of `SimData` objects
- Repeat for each line in the input file.

DataWriter

- Create an output file and write the necessary table headings.
- Add a line to the file including the cross-wind angle and drag coefficient after each simulation has run.

SimRunner

- Retrieve the properties in STAR-CCM+related to the input data.
- Retrieve data from the `SimData` object and set the values of these properties.
- Clear any previous solution.
- Run the simulation.
- Obtain the value of the drag coefficient and store it in the `SimData` object.

PostProcessor

- Retrieve the scenes and plots in STAR-CCM+ to be exported as hardcopies.
- Save hardcopies of various scenes and plots.

The following sections describe each of these classes in detail, highlighting key features in the Java code.

Contents

Understanding the SimData Nested Class
Understanding the DataReader Nested Class
Understanding the DataWriter Nested Class
Understanding the SimRunner Nested Class
Understanding the PostProcessor Nested Class

Understanding the SimData Nested Class

The SimData class defines objects that hold the data associated with a single simulation. Objects of this type are used throughout the macro.

Description	Code

• Find the `SimData` class declaration in the macro. In the initial lines of the class, declare the member variables . It is good practice for any member variable in a class to be defined as `private` to prevent unauthorized modification of the variable. If necessary, the class provides a "getter" function that returns the value of the member variable. Define the "getter" function as `public` or `protected` (with `public` being the most general).	`public class SimData {`
	
	`private double m_angDeg = 0.0;`
	`private double m_velX = 0.0;`
	`private double m_initVelY = 0.0;`
	`private double m_dragC = 0.0;`

• The constructor definition follows the declaration of the member variables. By writing the constructor definition in this way you have specified that when a `SimData` object is created, a specific set of input parameters is required. Statements in the body of the constructor are executed when a `SimData` object is created, and in this case the constructor calculates velocity components and sets the value of the member variables.	`public SimData(`
	`double angDeg,`
	`double velTrn,`
	`double velWnd,`
	`double initVelTrn,`
	`double initVelWnd) {`
	
	`m_angDeg = angDeg;`
	
	`double angRad =`
	
	`Math.toRadians(angDeg);`
	
	`m_velX =`
	`-1 * velWnd *`
	`Math.sin(angRad);`
	
	`}`

• This class exists as a data store, and requires "getter" and "setter" methods so that the data can be accessed and modified. Include a "getter" for each value that to be read. Include a "setter" for each modifiable value to be modified.	`public double getAngle() {`
	`return m_angDeg;`
	`}`
	
	`public void`
	`setDrag(double dDrag) {`
	`m_dragC = dDrag;`
	`}`

Proceed to the next section to examine the `DataReader` class.

Understanding the DataReader Nested Class

The `DataReader` class reads data from the input file and creates a `SimData` object for each set of wind data.

Each new `SimData` object is added to an `ArrayList` collection.

Description	Code

• In the first lines of the class a member variable is declared. The `SimData` objects populate the variable.
• No constructor has been defined, so the system supplies a default constructor when a `DataReader` object is created.
	`public class DataReader {`
	
	`  private List<SimData> m_flows =`
	`  new ArrayList<SimData>();`

• Specify a method to carry out the tasks. The path to the input file is passed as an argument to the method. A try-catch construct is used to handle any exceptions when the file cannot be found. If an exception is thrown, the class opens a MessageDialog showing the error string.	`public void readInput(String fileToRead) {`
	
	`  try {`
	
	`  FileReader fr =`
	`  new FileReader(fileToRead);`
	`  BufferedReader br =`
	`  new BufferedReader(fr);`
	`  Scanner sc =`
	`  new Scanner(br);`
	
	`  } catch (Exception e) {`
	
	`JOptionPane.showMessageDialog(`
	`  null, e.toString()`
	`  );`
	`  }`
	`}`

The `JOptionPane` statement is included to aid in debugging; as described in Testing and Debugging.

Description	Code
• Look at the `while` loop which slots into the `try-catch` brackets. The `Scanner` class allows you to read each value in the file as it is encountered, as well as read the file line-by-line. By default the scanner recognizes white spaces, such as tabs, spaces, and carriage returns, as breaks between values. You can use this functionality to read each value in a line, and pass it to a `SimData` object. Once a `SimData` object has been created and stored in the collection, repeat the process for the next line in the file. A loop repeats the process "while the file has another line". Also use an additional check in the form of an `if` statement. A line is only read out if it contains a number, and not if it.	`while (sc.hasNextLine()) { sc.nextLine(); if (sc.hasNextDouble()) { double angDeg = sc.nextDouble(); double initVelWnd = sc.nextDouble(); SimData sd = new SimData( angDeg, , initVelWnd ); m_flows.add(sd); } }`
• Finally, a getter method is provided to allow access to the collection of `SimData` objects from outside the `DataReaderclass`.	`public List<SimData> getFlowDetails() { return m_flows; }`

Proceed to the next section to examine the `DataWriter` class.

Understanding the DataWriter Nested Class

The `DataWriter` class is used to create an output text file containing the drag coefficient for each cross-wind angle.

Description	Code
• The class contains a member variable of type `String`, which	`public DataWriter(String fileToWrite) {`

Description	Code

is used to specify the name of the output file.
• The `FileWriter` and `BufferedWriter` used to write data to an external file, belong to the Java language rather than the STAR-CCM+ API. Use the constructor to create an output file and write a line of table headings.
	`m_outputFile = fileToWrite;`
	
	`try {`
	`FileWriter fw = new FileWriter(m_outputFile);`
	`BufferedWriter bw = new BufferedWriter(fw);`
	`bw.write("Angle (deg), DragCoefficient");`
	`bw.newLine();`
	`bw.close();`
	`} catch (Exception e) {`
	`}`
	`}`

• The `writeDataLine()` method opens the output file, adds a line that includes the cross-wind angle and the drag coefficient, then closes the file. The `newLine()` method is included to move the cursor to the start of the next line in the output file. As the output data is contained in the `SimData` object, the relevant object is passed to the method in the argument list.
• When the `FileWriter` object is created in the method, you can see that one of the arguments is `true`. This results in new data being appended to the end of the file.
	`public void writeDataLine(SimData sD) {`
	
	`try {`
	`FileWriter fw = new FileWriter(m_outputFile, true);`
	`BufferedWriter bw = new BufferedWriter(fw);`
	
	`bw.write(sD.getAngle() + ", " + sD.getDrag());`
	`bw.newLine();`
	`bw.close();`
	`} catch (Exception e) {`
	`}`
	`}`

Proceed to the next section to examine the `SimRunner` class.

Understanding the SimRunner Nested Class

In the `SimRunner` class, you implement the statements that were obtained when the macro was recorded.

Description	Code
• You declare the properties that are set in the simulation as member variables of the `SimRunner` class. Include only variables that are set using the input data, such as inlet velocity, and variables that are extracted, such as drag coefficient. The type of each of the member variables can be found by referring to the `macroRecording.java` file, and to the section Understanding the Recorded Macro .	`private VelocityProfile m_inflowVel = null;`

• Locating an object in STAR-CCM+ follows the same path as expanding nodes in the object tree. Therefore, to perform an operation such as setting the inlet velocity, locate each of the upper objects first. Here this action is performed in the constructor, as the location of the objects do not change, and you want to find the objects for every `SimRunner` object that is created.
• The statements for locating each object can be copied from the recorded macro, taking care the change the variable names as necessary.
	`public SimRunner(Simulation theSim) {`
	
	`  m_sim = theSim;`
	
	`  PhysicsContinuum physics =`
	`  ((PhysicsContinuum) m_sim`
	`  .getContinuumManager()`
	`  .getContinuum("Physics 1"));`
	
	`  m_initVel =`
	`  ((VelocityProfile) physics`
	`  .getInitialConditions()`
	`  .get(VelocityProfile.class));`
	
	`}`

• In the next method, `runCase()`, implement the statements that set the necessary properties in the simulation. First, you extract the values from the `SimData`	`public void runCase(SimData sD, int iterations) {`
	
	`  double initX =`
	`  sD.getInitVelX();`

Description	Code
object; then you use this data to set the values of each property. Notice that you use the getter methods that were defined in the `SimData` nested class to extract the data.	`double velY = sD.getVelY();`

Description	Code

• The variables are inserted into the `x` and `y` positions in the `setComponents()` method for the inflow and initial velocities; the variables have been underlined.	`((ConstantVectorProfileMethod) m_initVel .getMethod()) .getQuantity() .setComponents(initX, initY, 0.0);`
	
	`((ConstantVectorProfileMethod) m_inflowVel .getMethod()) .getQuantity() .setComponents(velX, velY, 0.0);`

• Each of the statements in the method is carried out in sequence. Once the new values have been set, you want to make sure that any previous solution has been cleared.	`m_sim.clearSolution();`
• You can now run the simulation. The number of iterations is specified by passing an integer to the `run()` method, the value of which set when the `runCase()` method is called.	`m_sim.getSimulationIterator() .run(iterations);`

• Finally, you obtain the drag coefficient, and store it in the `SimData` object by using the setter method defined earlier.	`double cdValue = m_forceReport.getReportMonitorValue();`
	
	`sD.setDrag(cdValue); }`

Proceed to the next section to examine the `PostProcessor` class.

Understanding the PostProcessor Nested Class

The `PostProcessor` class is used to export hardcopies of the *Velocity Magnitude* scene and the *Residuals* plot. You can also export the *Streamlines* scene as a STAR-View+ file.

Description	Code

• In a similar fashion to the previous class, you set member variables for each scene or plot that you want to export.	`public class PostProcessor {`
	
	`  private Scene m_velMag = null;`

• Furthermore, you find the relevant objects in the simulation by placing the necessary statements in the constructor.	`public PostProcessor(Simulation theSim) {`
	
	`  m_sim = theSim;`
	
	`  m_velMag =`
	`  ((Scene) m_sim`
	`  .getSceneManager()`
	`  .getScene("Velocity Magnitude"));`
	
	`}`

• Following are methods which are used to perform the action of saving the relevant scene or plot. As before, the statements for carrying out these operations can be copied from the recorded macro. You pass the path of the output file to the method as an argument, `String sceneToSave`.	`public void saveVelMagScene(String sceneToSave) {`
	
	`  m_velMag.printAndWait(`
	`  sceneToSave, 1, 1024, 768`
	`  );`
	
	`}`

• When a hardcopy is saved, it is automatically opened in the *Graphics* window in the STAR-CCM+ GUI. In the Streamlines scene, make sure that the scene is closed before	`public void saveStreamlinesScene(String sceneToSave) {`
	
	`m_streamlines.export3DSceneFileAndWait(`

Description	Code

the next simulation is run, otherwise the simulation run time is dramatically increased. Use the `Scene.close()` method.	`sceneToSave, true`
	`);`
	
	`m_streamlines.close(true);`
	
	`}`

Understanding the Main Method in the Macro

The main method in the macro is the `execute()` method, and is executed when the macro is run within STAR-CCM+.

Here you instantiate objects for each of the nested classes and use the methods in each, to carry out the required tasks.

Description	Code

• A try-catch operation surrounds the entire block of statements in the `execute()` method. The try-catch is useful in case any other errors are thrown when you run the macro.	`public void execute() {`
	
	`try {`
	
	`} catch (Exception e) {`
	
	`jOptionPane.showMessageDialog(`
	`null, e.toString()`
	`);`
	`}`
	`}`

Within the try-catch operation, you can divide the statements into two main sections. The first section contains statements that are carried out once, for example retrieving the active simulation. The second section contains a loop in which you place statements that are repeated for each set of wind data.

Description	Code

• In the first section you retrieve the active simulation and instantiate the `DataReader`, `DataWriter`, `SimRunner`, and `PostProcessor`. In addition, use the `DataReader` object to	`Simulation theSim =`
	`getActiveSimulation();`
	
	`DataReader reader =`
	`new DataReader();`
	`reader.readInput(folder +`
	`"/trainInput.txt");`

Description	Code

call the `readInput()` method that reads the input file and create a collection of `SimData` objects. The `getFlowDetails()` method is called to obtain a reference to the `SimData` collection and is supplied to the loop.	`List<SimData> listCases = reader.getFlowDetails();`
	
	`PostProcessor postP = new PostProcessor(theSim);`

• Next, use a loop to run the simulation for each set of wind data. The syntax used is that of a for-each loop, and reads: for each `SimData` object in the collection, do the following. This functionality came about in Java 5, and offers a convenient way to iterate over arrays and other collections. You use the reference to the `SimData` collection, `listCases` , to define the for-each loop.	`for (SimData sD : listCases) {`
	
	`}`

Description	Code
• Within the loop, the first process is to use the `SimRunner` object to carry out the tasks such as setting various conditions, clearing the previous solution, or running the case. Recall that these processes were defined in the method `runCase()`. You call this method, and pass it the current `SimData` object `sD` and an integer to define the number of iterations to carry out when each simulation is run, in this case 5.	`runner.runCase(sD, 5);`

Description	Code

• Next, following the logical process of running a simulation, you extract the drag coefficient and save hardcopies of scenes and a plot. As before, call the methods that were defined in the relevant nested classes. The cross-wind angle was retrieved from the current `SimData` object to save each scene under a suitable name.	`writer.writeDataLine(sD);`
	
	`postP.saveVelMagScene(`
	`  folder + "/velMag" +`
	`sD.getAngle() + ".png"`
	`);`

• Finally, use another statement copied from the recorded macro to save the simulation after it has been run.	`theSim.saveState(`
	`  folder + "/train" +`
	`sD.getAngle() + ".sim"`
	`);`

Testing and Debugging

Test and debug your macro.

There are some useful tools that can help to debug the macro and find the source of the problems.

The first is used with the try-catch error handler, and has the following syntax:

```
try {

} catch (Exception e) {
  JOptionPane.showMessageDialog(
    null, e.toString()
  );
}
```

The purpose of this code is to create a window that displays the error message. The macro is paused until you accept the error by clicking the **OK** button. Afterwards, the macro attempts to continue running and if it cannot continue the process stops.



Another useful tool that can be used when debugging is the ability to print a line of text to the *Output* window in the STAR-CCM+GUI. Implement code that follows the following syntax:

```
Simulation mySim =
  getActiveSimulation();

mySim.println("Output Text");
```

The output line is passed to the `println()` method as an `Object`, and so can include various parameters. The following example has inputs of type `String`, `int`, and `double`:

```
String myStr = "initial conditions";
double myDbl = 19.84;
int myInt = 910;

Simulation mySim =
  getActiveSimulation;

mySim.println(
  "Setting " + myStr + ": " + myInt + " " + myDbl
);
```

This code results in the following display in the *Output* window.



You can use this tool to print a line of text to the *Output* window when each key step in the macro has been performed. In this way, if the macro comes to a halt prematurely then you know that the code before the last printed message is not responsible for the error.

Keep the macro run-time to a minimum during development, as it is likely that several attempts are required to get everything correct. The number of iterations can be increased to their final amount at the end.

Running the Macro

Modify the macro to use your working directory to find the input file:

1. With the `trainFlowAngles.java` file open in an appropriate text editor, locate the following line. Change the section of code that is shown in bold.

```
public static final String folder =
  "/myTutorials";
```

2. In place of `/myTutorials`, enter the absolute path to the location of the `trainInput.txt` file that was saved earlier in the tutorial.
3. When the new path has been set, save the file to your working directory. Retain the file name `trainFlowAngles.java`.
4. Close the file

You can now run the macro:

5. Return to the STAR-CCM+ GUI and make sure that the `trainMacro` simulation is active.
6. In the *Macro* toolbar, click **(Play Macro)** and select the `trainFlowAngles.java` file.

   If the macro does not run, make sure that:

   - The case has been set up correctly to receive the input data, for example the inlet velocity specification is set to components
   - The `trainInput.txt` file is located in the directory specified in the macro
   - Use the tools described in the Testing and Debugging section to identify the source of the problem.
7. Proceed to the next section if you wish to run the simulation for sufficient iterations to achieve convergence, or jump to the Summary to review the tutorial.

Contents

Running the Simulation to Convergence

Running the Simulation to Convergence

The following steps demonstrate how to alter the macro to run the set of simulations until convergence has been reached.

1. Open `trainFlowAngles.java` file using an appropriate text editor.

2. Locate the line in the macro where you call the `runCase()` method from the `SimRunner` object.

```
public void execute() {

  try {

    for (/**/) {

      runner.runCase(sD, 5);

    }
    } catch (/**/) {
    }
}
```

3. Change the number of iterations to `400`.

```
runner.runCase(sD, 400);
```

4. Save and close the file.
5. Make sure that the simulation `trainMacro.sim` is open.
6. Click the **(Play Macro)** button in the *Macro* toolbar and select the `trainFlowAngles.java` file.
7. If, at any time, the macro must be interrupted click the **(Abort)** button next to the status bar.



When the macro has finished running you get a set of drag coefficient data, as well as a collection of scenes and plots. The following are examples of the residual plot, and streamlines scene for the train with a cross-wind angle of 90 degrees, after 400 iterations.

Residual Plot:



Velocity Magnitude Scene:



Streamlines Scene:



Summary

This tutorial has introduced the following STAR-CCM+ features:

- Setting up wake refinement
- Defining the motion of the track
- Creating a drag coefficient report
- Creating a velocity magnitude scalar scene
- Creating a scalar scene with streamlines
- Recording a macro;
- Understanding the recorded macro
- Designing a macro that uses nested classes
- Reading data into STAR-CCM+from a text file

- Using a macro to run consecutive simulations with different boundary conditions
- Writing data to an output text file
- Using a macro to export hardcopies of scenes and plots
- Testing and debugging the macro
- Running the macro

Simulation Assistant: Internal Flow Assistant

The Simulation Assistant provides you with an API framework that allows you to create workflows.

When loaded into STAR-CCM+, workflows appear on the right-hand side of the user interface, as shown highlighted in red in the following screenshot. An assistant project consists of Java classes and XHTML text.



When developing an assistant, it is recommended you use an IDE (Integrated Development Environment). This tutorial uses the NetBeans 7.3.1 IDE. You can get a free copy of NetBeans SE from the official website.

This tutorial presents the steps to create a Simulation Assistant for a simple workflow. The specific workflow that is created is for modeling an internal pipe flow. The existing S-Bend tutorial is an example of such an internal flow.

When creating an assistant, it is useful to break down the workflow into logical tasks, and to group related actions together. For example, actions that are related to creating and setting up the physics of a simulation can be grouped together in a single task.

The assistant in this tutorial has six tasks:

- Importing the geometry
- Creating regions from the parts
- Setting up the physics
- Generating the mesh
- Visualizing the solution
- Running the simulation

In general, there are two methods that you can use to create the initial Java code for each task:

- Record a Java macro while working through an example simulation in the STAR-CCM+ GUI. Copy and paste the relevant code from the recorded macro to the Simulation Assistant project.
- Write your own code for the assistant project.

In this tutorial you start by creating a simple assistant that has basic functionality. When the basic assistant is complete, you add more advanced features to it. These are:

- Conditions—add an extra layer of control and guidance to the assistant, and act as safeguards when you move between tasks
- Information Tags—allow you to add more information about a task or a particular step without cluttering the assistant panel. By default, additional information is hidden and only appears when required.
- Object Tree Focus—allows you to manipulate the STAR-CCM+ object tree. When you click on a step, you can control whether the corresponding object in the STAR-CCM+ object tree is highlighted. If the object is hidden inside manager nodes, these nodes are automatically expanded. To further focus on the action, you can choose to collapse all other expanded nodes in the object tree.

Contents

Prerequisites
Creating the NetBeans Project
Creating the Assistant Class
Creating the First Task: Importing the Geometry
Building the Project
Loading the Assistant into STAR-CCM+
Creating the Second Task: Converting the Parts to Regions
Creating the Third Task: Continuum Physics
Creating the Remaining Tasks
Advanced Simulation Assistant Features
Expanding the Object Tree to Reveal Specific Objects
Adding an Information Tag to the Assistant
Summary

Prerequisites

The instructions in the 3D-CAD tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

You require an installation of NetBeans 7.3.1 or later to complete this tutorial. Before attempting this tutorial, make sure that you are familiar with the following techniques:

Technique	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes and plots	Introduction to STAR-CCM+
Recording and editing Java macros	Simple Java Macros: Post Processing

Creating the NetBeans Project

Create a NetBeans project.

The first step in creating a Simulation Assistant is to create a project in NetBeans. Each Simulation Assistant typically consists of an assistant class and numerous task classes. The assistant class provides the overall order for the tasks in your assistant, while the task classes contain the steps that perform actions in STAR-CCM+.

Note: If you require assistance in setting up NetBeans or your project, refer to *Creating a Simulation Assistant Project in NetBeans* in the STAR-CCM+ user guide.

Create and prepare a new project in NetBeans:

1. Launch NetBeans and create a new Java application project called `InternalFlowAssistant`.
2. While creating the project, do not create the main class.
3. Add two Java packages to the project's source package and name them:
   - `Assistant`
   - `XHTML`
4. Load the STAR-CCM+ libraries into your NetBeans project. You must import the `.jar` files that are located in the following directories:
   - `[STAR-CCM+_Installation]/star/lib/java/platform/core`
   - `[STAR-CCM+_Installation]/star/lib/java/platform/core/locale`
   - `[STAR-CCM+_Installation]/star/lib/java/platform/lib`
   - `[STAR-CCM+_Installation]/star/lib/java/platform/modules`
   - `[STAR-CCM+_Installation]/star/lib/java/platform/modules/ext`
   - `[STAR-CCM+_Installation]/star/lib/java/platform/modules/locale`

Note: If you are creating the project on a Windows platform, it is recommended that you copy and paste the STAR-CCM+ libraries to the root directory of your system, for example: `C:\STAR-CCM+ Libraries`. When building code, Windows has a limitation of how large the path to the libraries can be. In the default STAR-CCM+ installation directory, the STAR-CCM+ libraries would contain paths that are too long. When building the project, this would cause problems.

The project is displayed below.



Creating the Assistant Class

Create the assistant class.

The assistant class acts as the 'table of contents' for all the tasks in your assistant. It is good practise to create this class first, as it will allow you to test your tasks when you complete them. Initially, you create the 'template' of the class, and add the name of the task classes to this template as you complete them, in the order that you want them to appear.

Create the assistant class:

1. Right-click the Assistant package and select **New** > **Java Class**.
2. Add a new class to the **Assistant** package and name it `InternalFlowAssistant`.



3. Copy and paste the following code into the `InternalFlowAssitant` class, overwriting everything:

```
package Assistant;

import java.util.ArrayList;
import java.util.List;
import star.assistant.SimulationAssistant;
import star.assistant.Task;
import star.assistant.annotation.StarAssistant;

// Specifies the name of the Simulation Assistant in
the GUI.
@StarAssistant(display = "Internal Flow Assistant")
public class InternalFlowAssistant extends
SimulationAssistant {

    public InternalFlowAssistant() {
        // Creates a new array list for the list of
tasks.
        List<Task> tasks = new ArrayList<Task>();
        // Contains the list of tasks.
        tasks.add(new Task1Name());
        setOutline(tasks);
    }
}
```

where:

- `tasks.add(new Task1Name());` is the placeholder for the first task.
- `Task1Name` is the name of the task class. When you complete the first task, replace it with the name of the class.

When you copy and paste from a web browser into NetBeans, the text loses its formatting. NetBeans can reformat this for you:

4. Select **Source** > **Format**.

Creating the First Task: Importing the Geometry

The first task imports the geometry into STAR-CCM+. This task is designed to be interactive, meaning that when you click the link in the task, an import dialog appears.

To begin:

1. Add a new class to the Assistant package and name it `Task01ImportGeometry`.



2. Copy and paste the following code into the `Task01ImportGeometry` class, overwriting everything:

```
package Assistant;

import java.io.File;
import java.util.Collection;

import javax.swing.JFileChooser;

import star.assistant.Task;
import star.assistant.ui.FunctionTaskController;
import star.common.GeometryPart;
import star.common.Simulation;
import star.common.SimulationPartManager;
import star.meshing.PartImportManager;

public class Task01ImportGeometry extends Task {

    public Task01ImportGeometry() {
    }

    public class ImportGeometryTaskController
extends FunctionTaskController {

        public void importSurfaceMeshDialog() {

            //open a file chooser
            JFileChooser fileChooser = new
JFileChooser();
```

```
            if (fileChooser.showOpenDialog(null) !=
JFileChooser.APPROVE_OPTION) {
                return;
            }
            File cadFile =
fileChooser.getSelectedFile();

            //import the part
            Simulation simulation_0 =
getSimulation();
            PartImportManager partImportManager_0 =

simulation_0.get(PartImportManager.class);

            //use default import options

partImportManager_0.importCadPart(cadFile.getPath(),
"SharpEdges", 30.0, 2, true, 1.0E-5, true, false);

            //add the new part to the lookup
            Collection<GeometryPart> new_parts =
simulation_0.get(SimulationPartManager.class).getPart
s();
            if (!new_parts.isEmpty()) {

addToTaskLookup(new_parts.iterator().next());
            }

            //create a Geometry Scene

simulation_0.getSceneManager().createGeometryScene("G
eometry Scene", "Outline", "Geometry", 1);
            Scene scene_0 =

simulation_0.getSceneManager().getScene("Geometry
Scene 1");

            scene_0.initializeAndWait();
            PartDisplayer partDisplayer_1 =
                    ((PartDisplayer)
scene_0.getCreatorDisplayer());
            partDisplayer_1.initialize();
            PartDisplayer partDisplayer_0 =
                    ((PartDisplayer)
scene_0.getDisplayerManager().getDisplayer("Outline
1"));
            partDisplayer_0.initialize();
            PartDisplayer partDisplayer_2 =
                    ((PartDisplayer)
scene_0.getDisplayerManager().getDisplayer("Geometry
1"));
            partDisplayer_2.initialize();
            scene_0.open(true);
            CurrentView currentView_0 =
                    scene_0.getCurrentView();
```

```
            currentView_0.setInput(new
DoubleVector(new double[]{0.07000000029802322,
0.02000000048428774, 0.0}),

new DoubleVector(new double[]{-0.14607243684130425,
0.11741414039370678, 0.17860456694301519}),

new DoubleVector(new double[]{0.20083797287207372,
0.9415226723669483, -0.2705534440210218}),

0.07747426272518879, 0);
        }
    }
}
```

In this task, you use the standard Java function, `JFileChooser`, to launch an open dialog. The person using the assistant can then import the desired geometry file into STAR-CCM+.

Note: When you write your own code, NetBeans automatically suggests which classes to import. Lines that need your attention are flagged with a red exclamation mark. Press "`Alt`" + "`Enter`" on these lines to display the suggestions. For now, you can ignore any warnings.

Contents

Adding a Description to the Task Panel
Linking the XHTML Content to the Task Class
Reviewing the Code for Task 1

Adding a Description to the Task Panel

You can add descriptions to each task panel. Furthermore, you can link a particular word or phrase in a step to a staraction. The word effectively becomes a link, and when clicked, performs a pre-defined action. It is recommended that you store all text in an external `.xhtml` file. It is also good practise to store all `.xhtml` files in a separate package to the source package. For this tutorial, you already created the `XHTML` package in an earlier section.

Add an XHTML file to your project:

1. Right-click on the **XHTML** package and select **New > Other...**.



The *New File* dialog appears.

2. From the **Categories** list, select **Other**.
3. From the **File Types** list, select **XHTML**.



4. Click **Next** and enter `01_ImportGeometry` as the name.
5. Click **Finish**.

   The `.xhtml` file is added to the XHTML package.



6. Copy and paste the following code into the `01_ImportGeometry.xhtml` file:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.1//EN"
"http://www.w3.org/TR/xhtml11/DTD/xhtml11.dtd">
<html xmls="http://www.w3.org/1999/xhtml">
    <head>
        <meta http-equiv="CONTENT-TYPE"
```

```
content="text/html; charset=utf-8" />
    </head>
    <body>
        <p>
            <!--This is the description of the task
that appears in the task panel.-->
            This task imports a geometry part into
STAR-CCM+.
            <ul>
                <!--This is a list item that appears
in the task. The word "Import" is linked to a
staraction.-->
                <li><a
href="staraction:importSurfaceMeshDialog">Import</a>
the surface mesh.</li>
            </ul>
        </p>
    </body>
</html>
```

where:

- `staraction:importSurfaceMeshDialog` calls the `importSurfaceMeshDialog` function from the `Task01ImportGeometry` class.

7. Reformat the code in NetBeans.

Linking the XHTML Content to the Task Class

Link the XHTML content to the task class.

Now that the XHTML file is completed, link this file to the `Task01ImportGeometry` class:

1. Return to `Task01ImportGeometry` and locate the line beginning with `public class Task01ImportGeometry extends Task {`.
2. Insert the following code above that line:

```
@StarAssistantTask(display = "Import Geometry",
    contentPath = "XHTML/01_ImportGeometry.xhtml",
    controller =
Task01ImportGeometry.ImportGeometryTaskController.cla
ss)
```

Reviewing the Code for Task 1

Review the code for the first task.

After linking the XHMTL file to the task class, NetBeans adds a warning on the line beginning with `@StarAssistantTask`. This error appears due to a missing import for the assistant task.

To resolve the issue, import the appropriate class:

1. Click on the line with the warning and press <`Alt`> + <`Enter`>.
   A pop-up appears with suggestions to fixing the problem.
2. Select **Add import for star.assistant.annotation.StarAssistantTask**.



   The following line is added to the code and the error flag is removed:

```
import star.assistant.annotation.StarAssistantTask;
```

3. Select **Source** > **Fix Imports** to automatically add all the missing imports to the class.
4. Select **Source** > **Format** to reformat the code in NetBeans.

The final code for the first task is displayed below:

```
package Assistant;

import java.io.File;
import java.util.Collection;

import javax.swing.JFileChooser;

import star.assistant.Task;
import star.assistant.annotation.StarAssistantTask;
import star.assistant.ui.FunctionTaskController;
import star.base.neo.DoubleVector;
import star.common.GeometryPart;
import star.common.Simulation;
import star.common.SimulationPartManager;
import star.meshing.PartImportManager;
import star.vis.CurrentView;
import star.vis.PartDisplayer;
```

```
import star.vis.Scene;

@StarAssistantTask(display = "Import Geometry",
    contentPath = "XHTML/01_ImportGeometry.xhtml",
    controller =
Task01ImportGeometry.ImportGeometryTaskController.class)
public class Task01ImportGeometry extends Task {

    public Task01ImportGeometry() {
    }

    public class ImportGeometryTaskController extends
FunctionTaskController {

        public void importSurfaceMeshDialog() {

            //open a file chooser
            JFileChooser fileChooser = new
JFileChooser();
            if (fileChooser.showOpenDialog(null) !=
JFileChooser.APPROVE_OPTION) {
                return;
            }
            File cadFile =
fileChooser.getSelectedFile();

            //import the part
            Simulation simulation_0 = getSimulation();
            PartImportManager partImportManager_0 =

simulation_0.get(PartImportManager.class);

            //use default import options

partImportManager_0.importCadPart(cadFile.getPath(),
"SharpEdges", 30.0, 2, true, 1.0E-5, true, false);
            //add the new part to the lookup
            Collection<GeometryPart> new_parts =
simulation_0.get(SimulationPartManager.class).getParts()
;
            if (!new_parts.isEmpty()) {

addToTaskLookup(new_parts.iterator().next());
            }

            //create a Geometry Scene

simulation_0.getSceneManager().createGeometryScene("Geom
etry Scene", "Outline", "Geometry", 1);
            Scene scene_0 =

simulation_0.getSceneManager().getScene("Geometry Scene
1");
            scene_0.initializeAndWait();
            PartDisplayer partDisplayer_1 =
```

```
                ((PartDisplayer)
scene_0.getCreatorDisplayer());
            partDisplayer_1.initialize();
            PartDisplayer partDisplayer_0 =
                ((PartDisplayer)
scene_0.getDisplayerManager().getDisplayer("Outline
1"));
            partDisplayer_0.initialize();
            PartDisplayer partDisplayer_2 =
                ((PartDisplayer)
scene_0.getDisplayerManager().getDisplayer("Geometry
1"));
            partDisplayer_2.initialize();
            scene_0.open(true);
            CurrentView currentView_0 =
                scene_0.getCurrentView();
            currentView_0.setInput(new DoubleVector(new
double[]{0.07000000029802322, 0.02000000048428774,
0.0}),
                new DoubleVector(new double[]
{-0.14607243684130425, 0.11741414039370678,
0.17860456694301519}),
                new DoubleVector(new double[]
{0.20083797287207372, 0.9415226723669483,
-0.2705534440210218}),
                0.07747426272518879, 0);
        }
    }
}
```

Building the Project

Build the NetBeans project.

At this stage, you must update the assistant class with the name of the class for Task 1:

1. Return to the assistant class, `InternalFlowAssistant`, and locate the line:

   ```
   tasks.add(new Task1Name());
   ```

2. In that line, replace `Task1Name()` with the name of the Task 1 class:

   ```
   Task01ImportGeometry()
   ```

3. Save the project.

You can now build the Simulation Assistant project:

4. Select **Run > Build Project**.

When complete, `InternalFlowAssistant.jar` is placed in the `[Simulation_Assistant_Project]/dist` folder. This is the assistant file that you load into STAR-CCM+.

Loading the Assistant into STAR-CCM+

Load the assistant into STAR-CCM+:

1. Load STAR-CCM+ and create a new simulation.
2. Select **File** > **Load Simulation Assistant...**.



The *Open* dialog appears.

3. Navigate to the `.jar` file for your assistant and open it.

   The assistant loads into STAR-CCM+ and appears on the right-hand side of the user-interface. At this stage, the assistant only has a single task in it.



It is recommended that you rebuild the assistant project after adding each new task. Doing so gives you the opportunity to test that each task works as expected.

Testing the Simulation Assistant

Test that the assistant behaves as expected:

1. In the *Import Geometry* panel, click **Import**.
   The *Open* dialog appears.
2. Navigate to `[STAR-CCM+_Installation]/doc/startutorialsdata/simulationAssistant/data` and open `sBendGeometry.x_t`.
   The geometry is imported and placed in **Geometry** > **Parts**.
3. Close STAR-CCM+ without saving.

It is recommended that you rebuild the assistant project after adding each new task. Doing so gives you the opportunity to test that each task works as expected.

Creating the Second Task: Converting the Parts to Regions

The second task takes the imported part from the first task and creates a region and boundaries for it. Every new task has its own separate Java class.

For this tutorial, you are provided with the code for Task 2. The code makes sure that a geometry part exists in the simulation and, if found, creates a region for it. If no parts are found, nothing happens.

1. Add a new class to the Assistant package and name it `Task02CreateRegionFromPart`.



2. Copy the following code into the new class:

```
package Assistant;

import java.util.ArrayList;
import java.util.Collection;

import star.assistant.Task;
import star.assistant.annotation.StarAssistantTask;
import star.assistant.ui.FunctionTaskController;
import star.common.GeometryPart;
import star.meshing.CadPart;

@StarAssistantTask(display = "Create Region from
Part",
    contentPath = "XHTML/
02_CreateRegionFromPart.xhtml",
    controller =
Task02CreateRegionFromPart.RegionFromPartTaskControll
er.class)
public class Task02CreateRegionFromPart extends Task
{

    public Task02CreateRegionFromPart() {
    }

    public class RegionFromPartTaskController
extends FunctionTaskController {

        public void createRegion() {
            CadPart cadPart_1 =
lookupObject(CadPart.class);
            if (cadPart_1 != null) {
                Collection<GeometryPart> list = new
ArrayList<GeometryPart>();
                list.add(cadPart_1);

getSimulation().getRegionManager().newRegionsFromPart
s(list,
                    "OneRegionPerPart", null,
"OneBoundaryPerPartSurface", null,
```

```
                        "OneFeatureCurve", null, true);
                }
            }
        }
}
```

Contents

Adding the Text to Task 2
Building and Testing the Second Task

Adding the Text to Task 2

Add the text for the second task:

1. Add an XHTML file to the HMTL package and name it `02_CreateRegionFromPart`.



2. Copy and paste the following code into the new `02_CreateRegionFromPart.xhtml`:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.1//EN"
"http://www.w3.org/TR/xhtml11/DTD/xhtml11.dtd">
<html xmls="http://www.w3.org/1999/xhtml">
    <head>
        <meta http-equiv="CONTENT-TYPE"
content="text/html; charset=utf-8" />
    </head>
    <body>
        <p>
            This task creates the regions and
boundaries from the geometry part. Renamed surfaces
are used to create boundaries.
            <ul>
                <li><a
href="staraction:createRegion">Create</a> a region.</
```

```
li>
            </ul>
        </p>
    </body>
</html>
```

Building and Testing the Second Task

Build and test the second task.

Before you can build the project, you must add the second task to the assistant class:

1. Return to `InternalFlowAssistant` and add the following line below the line corresponding to the first task:

   ```
   tasks.add(new Task02CreateRegionFromPart());
   ```

   The final code is displayed below:

   ```
   public InternalFlowAssistant() {

       List<Task> tasks = new ArrayList<Task>(); //
   Creates a new array list for the list of tasks.
       // Contains the list of tasks.
       tasks.add(new Task01ImportGeometry());
       tasks.add(new Task02CreateRegionFromPart());
       setOutline(tasks);
       }
   ```

2. Save the project.
3. Rebuild the assistant project.
4. Load the rebuilt assistant into STAR-CCM+ and work through it.

   Task 1 imports the geometry into STAR-CCM+. Task 2 creates a region with the geometry from Task 1.

Creating the Third Task: Continuum Physics

Create the third task for the project.

The third task contains three steps:

- The first step creates a physics continuum and populates it with physics models that are appropriate for the simulation.
- The second step modifies the material properties of the air.
- The third step sets up initial and boundary conditions, including the boundary types.

For this task, you record a Java macro and perform the desired actions in the STAR-CCM+ user interface. You then copy and paste the code for each step into your class, under the relevant function. As the task has three steps, the class will have three functions.

First, create the class for Task 3:

1. Add a new class to the Assistant package and name it `Task03Physics`.



2. Copy and paste the following code into the class:

```
package Assistant;

import star.assistant.annotation.StarAssistantTask;
import star.assistant.Task;
import star.assistant.ui.FunctionTaskController;

@StarAssistantTask(display = "Create Physics",
    contentPath = "XHTML/03_Physics.xhtml",
    controller =
Task03Physics.PhysicsTaskController.class)
public class Task03Physics extends Task {

    public class PhysicsTaskController extends
FunctionTaskController {

        public void createPhysicsContinuum() {
            // code for Step 1: creating and
defining the physics continuum.
        }

        public void materialProperties() {
            // code for Step 2: modifying the
material properties of air.
        }
```

```
        public void
initialConditionsAndBoundarySettings() {
            // code for Step 3: defining the initial
conditions, boundary type, and boundary conditions.
        }
    }
}
```

This code contains three functions, one function for each step:

- `createPhysicsContinuum()`
- `materialProperties()`
- `initialConditionsAndBoundarySettings()`

Contents

Recording the Actions for Step 1
Recording the Actions for Step 2
Recording the Actions for Step 3
Reviewing the final Code for Task 3
Adding the Text to Task 3
Building and Testing the Third Task

Recording the Actions for Step 1

Record the actions for Step 1 in a macro:

1. Launch STAR-CCM+ and load the latest version of the assistant.
2. Work through the assistant until you reach the end.
3. Start recording a macro and name it `CreatingPhysicsContinuum.java`.
4. Create a physics continuum and rename it to `Physics`.
5. Edit the **Continuum** > **Physics** node and set the following properties:

Group Box	Model
Space	**Three Dimensional**
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**

Group Box	Model
Reynolds-Averaged Turbulence	**K-Omega Turbulence**

6. Click **Close**.
7. Stop recording the macro.

   When you record a macro, a macro window appears next to the **Output** window. The code for the macro is displayed there. The following code is displayed:

```
// begin macro: CreatePhysicsContinuum
Simulation simulation_0 =
    getActiveSimulation();
PhysicsContinuum physicsContinuum_0 =

simulation_0.getContinuumManager().createContinuum(Ph
ysicsContinuum.class);
physicsContinuum_0.setPresentationName("Physics");
physicsContinuum_0.enable(ThreeDimensionalModel.class
);
physicsContinuum_0.enable(SteadyModel.class);
physicsContinuum_0.enable(SingleComponentGasModel.cla
ss);
physicsContinuum_0.enable(SegregatedFlowModel.class);
physicsContinuum_0.enable(ConstantDensityModel.class)
;
physicsContinuum_0.enable(LaminarModel.class);
// end macro: CreatePhysicsContinuum
```

From here, you can copy and paste the code into the relevant function:

8. Return to NetBeans.
9. Copy and paste the code between `// begin macro: CreatePhysicsContinuum` and `// end macro` into the createPhysicsContinuum() function.
10. Reformat the macro code in NetBeans.
11. Select **Source > Fix Imports** to automatically add all the missing imports to the class.

Recording the Actions for Step 2

Record the actions for Step 2.

The second step involves modifying the material properties of air. Use the same technique as Step 1:

1. Return to STAR-CCM+ and record a macro named `ModifyMaterialProperties.java`.
2. Navigate to **Continua > Physics > Models > Gas > Air > Material Properties** and set the following properties:

Node	Property	Setting
Density	*Value*	`1.0 kg/m^3`
Dynamic	*Value*	`1.716E-5 Pa-s`

3. Stop recording the macro.

   The macro code for these actions is displayed below:

```
// begin macro: ModifyMaterialProperties
Simulation simulation_0 =
    getSimulation();
PhysicsContinuum physicsContinuum_0 =
    ((PhysicsContinuum)
simulation_0.getContinuumManager().getContinuum("Phys
ics"));
SingleComponentGasModel singleComponentGasModel_0 =

physicsContinuum_0.getModelManager().getModel(SingleC
omponentGasModel.class);
Gas gas_0 =
    ((Gas) singleComponentGasModel_0.getMaterial());
ConstantMaterialPropertyMethod
constantMaterialPropertyMethod_0 =
    ((ConstantMaterialPropertyMethod)
gas_0.getMaterialProperties()
.getMaterialProperty(ConstantDensityProperty.class).g
etMethod());
constantMaterialPropertyMethod_0.getQuantity().setVal
ue(1.0);
ConstantMaterialPropertyMethod
constantMaterialPropertyMethod_1 =
    ((ConstantMaterialPropertyMethod)
gas_0.getMaterialProperties()
.getMaterialProperty(DynamicViscosityProperty.class).
getMethod());
constantMaterialPropertyMethod_1.getQuantity().setVal
ue(1.716E-5);
// end macro: ModifyMaterialProperties
```

4. Copy and paste the code into the `materialProperties()` function.
5. Reformat the code in NetBeans.
6. Select **Source > Fix Imports**.

Recording the Actions for Step 3

Record the actions for Step 3.

The third step sets the initial conditions, defines the boundary types, and sets the velocity value on the Inlet boundary. The macro code for this step is provided for you:

1. Copy the following code into the `initialConditionsAndBoundarySettings()` function:

```
Simulation simulation_0 =
    getActiveSimulation();
PhysicsContinuum physicsContinuum_0 =
    ((PhysicsContinuum)
simulation_0.getContinuumManager().getContinuum("Phys
ics"));
VelocityProfile velocityProfile_0 =

physicsContinuum_0.getInitialConditions().get(Velocit
yProfile.class);
velocityProfile_0.getMethod(ConstantVectorProfileMeth
od.class).getQuantity().setComponents(0.429, 0.0,
0.0);
Region region_0 =

simulation_0.getRegionManager().getRegion("Fluid");
Boundary boundary_0 =

region_0.getBoundaryManager().getBoundary("Inlet");
boundary_0.setBoundaryType(InletBoundary.class);
VelocityMagnitudeProfile velocityMagnitudeProfile_0 =

boundary_0.getValues().get(VelocityMagnitudeProfile.c
lass);
velocityMagnitudeProfile_0.getMethod(ConstantScalarPr
ofileMethod.class).getQuantity().setValue(0.429);
Boundary boundary_1 =

region_0.getBoundaryManager().getBoundary("Outlet");
boundary_1.setBoundaryType(PressureBoundary.class);
```

2. Reformat the code in NetBeans.
3. Select **Source > Fix Imports**.
4. In the *Fix All Imports* dialog, click **OK**.

Reviewing the final Code for Task 3

Review the code for the third task.

The final code for Task 3 is displayed below:

```
package Assistant;

import star.assistant.annotation.StarAssistantTask;
import star.assistant.Task;
import star.assistant.ui.FunctionTaskController;
import star.common.Boundary;
import star.common.ConstantScalarProfileMethod;
import star.common.ConstantVectorProfileMethod;
import star.common.InletBoundary;
```

```
import star.common.PhysicsContinuum;
import star.common.PressureBoundary;
import star.common.Region;
import star.common.Simulation;
import star.common.SteadyModel;
import star.flow.ConstantDensityModel;
import star.flow.ConstantDensityProperty;
import star.flow.DynamicViscosityProperty;
import star.flow.LaminarModel;
import star.flow.VelocityMagnitudeProfile;
import star.flow.VelocityProfile;
import star.material.ConstantMaterialPropertyMethod;
import star.material.Gas;
import star.material.SingleComponentGasModel;
import star.metrics.ThreeDimensionalModel;
import star.segregatedflow.SegregatedFlowModel;

@StarAssistantTask(display = "Create Physics",
    contentPath = "XHTML/03_Physics.xhtml",
    controller =
Task03Physics.PhysicsTaskController.class)
public class Task03Physics extends Task {

    public class PhysicsTaskController extends
FunctionTaskController {

        public void createPhysicsContinuum() {

            // code for Step 1: creating and defining
the physics continuum.
            Simulation simulation_0 =
                getActiveSimulation();
            PhysicsContinuum physicsContinuum_0 =

simulation_0.getContinuumManager().createContinuum(Physi
csContinuum.class);

physicsContinuum_0.setPresentationName("Physics");

physicsContinuum_0.enable(ThreeDimensionalModel.class);

physicsContinuum_0.enable(SteadyModel.class);

physicsContinuum_0.enable(SingleComponentGasModel.class)
;

physicsContinuum_0.enable(SegregatedFlowModel.class);

physicsContinuum_0.enable(ConstantDensityModel.class);

physicsContinuum_0.enable(LaminarModel.class);
        }

        public void materialProperties() {
```

```
            // code for Step 2: modifying the material
properties of air.
            Simulation simulation_0 = getSimulation();
            PhysicsContinuum physicsContinuum_0 =
                ((PhysicsContinuum)
simulation_0.getContinuumManager().getContinuum("Physics
"));
            SingleComponentGasModel
singleComponentGasModel_0 =

physicsContinuum_0.getModelManager().getModel(SingleComp
onentGasModel.class);
            Gas gas_0 =
                Gas)
singleComponentGasModel_0.getMaterial());
            ConstantMaterialPropertyMethod
constantMaterialPropertyMethod_0 =
                ConstantMaterialPropertyMethod)
gas_0.getMaterialProperties()

getMaterialProperty(ConstantDensityProperty.class).getMe
thod());

constantMaterialPropertyMethod_0.getQuantity().setValue(
1.0);
            ConstantMaterialPropertyMethod
constantMaterialPropertyMethod_1 =
                ConstantMaterialPropertyMethod)
gas_0.getMaterialProperties()

getMaterialProperty(DynamicViscosityProperty.class).getM
ethod());

constantMaterialPropertyMethod_1.getQuantity().setValue(
1.716E-5);
        }

        public void
initialConditionsAndBoundarySettings() {

            // code for Step 3: defining the initial
conditions, boundary type, and boundary conditions.
            Simulation simulation_0 =
                getActiveSimulation();
            PhysicsContinuum physicsContinuum_0 =
                PhysicsContinuum)
simulation_0.getContinuumManager().getContinuum("Physics
"));
            VelocityProfile velocityProfile_0 =

physicsContinuum_0.getInitialConditions().get(VelocityPr
ofile.class);

velocityProfile_0.getMethod(ConstantVectorProfileMethod.
class).getQuantity().setComponents(0.429, 0.0, 0.0);
```

```
            Region region_0 =

simulation_0.getRegionManager().getRegion("Fluid");
            Boundary boundary_0 =

region_0.getBoundaryManager().getBoundary("Inlet");

boundary_0.setBoundaryType(InletBoundary.class);
            VelocityMagnitudeProfile
velocityMagnitudeProfile_0 =

boundary_0.getValues().get(VelocityMagnitudeProfile.clas
s);

velocityMagnitudeProfile_0.getMethod(ConstantScalarProfi
leMethod.class).getQuantity().setValue(0.429);
            Boundary boundary_1 =

region_0.getBoundaryManager().getBoundary("Outlet");

boundary_1.setBoundaryType(PressureBoundary.class);
        }
    }
}
```

Adding the Text to Task 3

Add the descriptions to the third task.

Now that Task 3 is complete, you can add the text and link the star actions to the relevant functions:

1. Add an XHTML file to your XHTML package and name it `03_Physics`.



2. Copy the following code into `03_Physics.xhtml`:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.1//EN"
"http://www.w3.org/TR/xhtml11/DTD/xhtml11.dtd">
<html xmls="http://www.w3.org/1999/xhtml">
    <head>
        <meta http-equiv="CONTENT-TYPE"
content="text/html; charset=utf-8" />
    </head>
    <body>
        <p>
            <!--Description in task panel.-->
            This task creates a physics continuum
and populates it
            with the appropriate physics models.
            <ol>
                <!--First list item.-->
                <li><a
href="staraction:createPhysicsContinuum">Create</a>
a physics continuum.</li>
                <!--Second list item.-->
                <li><a
href="staraction:materialProperties">Define</a> the
material properties.</li>
                <!--Third list item.-->
                <li><a
href="staraction:initialConditionsAndBoundarySettings
">Set</a>
                    the initial conditions and
boundary settings.</li>
            </ol>
        </p>
    </body>
</html>
```

Note: Each step is defined in a new list item. This task has three steps, so there are three list items, each of which is linked to a star action.

Building and Testing the Third Task

Build and test the third task.

Before building the project, add the third task to the assistant class:

1. Return to `InternalFlowAssistant` and add the following line below the line corresponding to the second task:

```
tasks.add(new Task03Physics());
```

2. Save the project.
3. Rebuild the project.

4. Test the assistant in STAR-CCM+.

Creating the Remaining Tasks

Create the remaining three tasks for the project.

The remaining tasks and associated XHTML files are already created for you to import into the project. An overview of what the tasks do is provided below.

Task 4

The fourth task deals with the meshing. This task is split into four steps:

- Step 1 creates a mesh continuum and populates it with the appropriate mesh models.
- Step 2 sets up the mesh parameters, such as base size.
- Step 3 generates the surface mesh.
- Step 4 generates the volume mesh.

Task 5

The fifth task creates a scalar scene that contains a plane-section along the center of the imported geometry.

Task 6

The sixth task is split into two steps:

- Step 1 sets the stopping criteria.
- Step 2 runs the simulation.

Contents

Adding the Remaining Tasks to Your Project
Test the Complete Assistant

Adding the Remaining Tasks to Your Project

Add the remaining tasks to your project:

1. In the file system, navigate to `[STAR-CCM+_Installation]/doc/startutorialsdata/simulationAssistant/data`.

2. Copy the following Java files into the `[Simulation_Assistant_Project]/src/Assistant` folder:
   - `Task04Meshing.java`
   - `Task05CreateScalarScene.java`
   - `Task06RunTheSimulation.java`
3. Copy and paste the following XHTML files into the `[Simulation_Assistant_Project]/src/XHTML` folder:
   - `04_Meshing.xhtml`
   - `05_ScalarScene.xhtml`
   - `06_RunSimulation.xhtml`
4. Save the project.

The tasks are automatically added to the NetBeans project.



Test the Complete Assistant

Perform a final test on the complete assistant:

1. Return to `InternalFlowAssistant` and add the missing tasks to the array list:

```
tasks.add(new Task04Meshing());
tasks.add(new Task05CreateScalarScene());
tasks.add(new Task06RunTheSimulation());
```

2. Save the project.
3. Rebuild the project.
4. Launch STAR-CCM+ and load the rebuilt assistant.
5. Work through the assistant and make sure that it behaves as expected:
   - Task 1 loads a geometry part into STAR-CCM+.
   - Task 2 creates a region from the imported geometry part.
   - Task 3 creates a physics continuum and sets up the physics models, initial conditions, and boundary conditions.
   - Task 4 creates a mesh continuum and sets up the mesh parameters. This task also covers the steps to generate the surface and volume meshes.
   - Task 5 creates a scalar scene.
   - Task 6 sets the stopping criteria and runs the STAR-CCM+ solver.

Advanced Simulation Assistant Features

There are multiple advanced Simulation Assistant features that you can use to improve the code. The following features are described in this section:

- Adding conditions to the task.
- Expanding the object tree to reveal specific objects.
- Adding information buttons within a task panel.

Contents

Adding Conditions to the Assistant
Adding a Post-Condition to Task 1
Adding a Pre-Condition to Task 2
Testing the Conditions

Adding Conditions to the Assistant

Add a post-condition to Task 1 and a pre-condition to Task 2.

In a Simulation Assistant, conditions impose requirements that must be met at specific points in the workflow. There are two types of conditions:

- Pre-conditions—define what is required before a task can begin.
- Post-conditions—check that the requirements of a task are met.

In this tutorial, you add a post-condition to Task 1, and a pre-condition to Task 2. The condition checks whether a geometry file part exists in the simulation. If no part exists, Task 1 is not complete and Task 2 is disabled.

It is good practice to add the code for all conditions in a separate class:

1. Add a new class to the Assistant package and name it `InternalFlowConditions`.



2. Copy the following code into the `InternalFlowConditions` class. Comments are included in the code:

```
package Assistant;

import java.util.Collections;
```

```
import star.assistant.CSOCondition;
import star.assistant.CSOLookupConditionTrigger;
import star.common.GeometryPart;
import star.common.filters.Predicate;

/**
 * This class contains all conditions used in the
Internal Flow Assistant.
 */
public class InternalFlowConditions {

    public static synchronized
CSOCondition<GeometryPart> createPartCondition() {

        // Creates a new condition
        CSOCondition<GeometryPart> partCondition =
new CSOCondition<GeometryPart>();
        // Sets the text description of the condition
        partCondition.setDesc("A geometry part must
be present.");
        // Creates a new condition trigger that goes
off when a cadpart is added to the lookup
        CSOLookupConditionTrigger<GeometryPart>
partConditionTrigger = new
CSOLookupConditionTrigger<GeometryPart>(GeometryPart.
class);
        // Sets the list of triggers to the one
created above.

partCondition.setTriggers(Collections.singleton(partC
onditionTrigger));
        // Creates a new predicate (true/false
evaluation) with an evaluate method and evaluates
whether an object satisfies the condition.
        partCondition.setPredicate(new
Predicate<GeometryPart>() {
            @Override
            public boolean evaluate(GeometryPart
part) {
                // You could check for specific
attributes of the part here.
                return true;
            }
        });
        return partCondition;
    }
}
```

Adding a Post-Condition to Task 1

The following code adds the post-condition to Task 1. The post-condition checks whether a geometry part exists in the **Geometry > Parts** manager node. If a part exists, the condition is satisfied.

1. Return to the `Task01ImportGeometry` class.
2. Locate the `public Task01ImportGeometry()` constructor.
3. Add the following code to this constructor:

```
// Makes sure that a geometry part is created.
setPostconditions(Collections.singleton(InternalFlowC
onditions.createPartCondition()));
```

   where:

   - `setPostconditions()` sets the post-condition for the task.
   - `Collections.singleton(InternalFlowConditions.createP artCondition())` is the condition that checks for a geometry part.
4. Add the missing import.

Adding a Pre-Condition to Task 2

Add a pre-condition to Task 2.

You cannot create a region from a part if there are no parts in the simulation, so it makes sense to add a pre-condition to Task 2, that disables that task if no geometry parts exist.

The following code adds a pre-condition to Task 2 to check whether a part exists in the `Geometry > Parts` manager node. If no part exists, the task is disabled (greyed out). If a part exists, the condition is satisfied and the task is enabled. Pre-conditions only affect the task that they are placed in.

1. Return to the `Task02CreateRegionFromPart` class.
2. Locate the `public Task02CreateRegionFromPart()` constructor.
3. Add the following code to this constructor:

```
// Makes sure that a geometry exists before enabling
this task.
setPreconditions(Collections.singleton(InternalFlowCo
nditions.createPartCondition()));
```

   where:

   - `setPreconditions()` sets the pre-condition for the task.
   - `Collections.singleton(InternalFlowConditions.createP artCondition())` is the condition that checks for a geometry part.
4. Add the missing import.
5. Save and build your project.

Testing the Conditions

Test the conditions to make sure that they behave as expected:

1. Load the assistant into STAR-CCM+.
2. Task 1 and Task 2 have condition tabs added to their panels.



3. Click on the **Post-Condition** tab in Task 1 and the **Pre-Condition** tab in Task 2 to expand them and display more information about the condition requirements.



4. Import the geometry part into STAR-CCM+.

   The conditions for both tasks are satisfied and Task 2 becomes available. When a Post-Condition is satisfied, a green checkmark is displayed on the task panel title bar. The object that satisfies the condition is also displayed in the condition tab.



Expanding the Object Tree to Reveal Specific Objects

Add an exclusive and an inclusive focus to the assistant.

When working through the workflow of an assistant, it is helpful to see what part of the tree an action is associated with. When you click on a link in the assistant, you can cause the relevant node in the STAR-CCM+ object tree to expand.

There are two types of object tree focus:

- Exclusive focus: collapses all nodes in the simulation tree and only expands the node that you select.
- Inclusive focus: expands the node for the item that you select. All other nodes remain unchanged.

In this tutorial, you add an object tree focus to Task 3 > Step 1 and Task 3 > Step 2.

Adding an Exclusive Focus to Step 1

Add an object tree focus so that the **Continua** > **Physics** > **Models** manager node expands in the simulation tree when you click the link in the Simulation Assistant:

1. Return to the `Task03Physics` class.
2. Within the `createPhysicsContinuum()` method, insert the following code after the last statement:

```
// Object tree focus that expands the physics model
manager node. All other open nodes are collapsed.
selectAndExpandNodeExclusive(physicsContinuum_0.getMo
delManager());
```

The final code for the first step in Task 3 is displayed below:

```
public void createPhysicsContinuum() {

    // code for Step 1: creating and defining the
physics continuum.
    Simulation simulation_0 =
        getActiveSimulation();
    PhysicsContinuum physicsContinuum_0 =

simulation_0.getContinuumManager().createContinuum(Ph
ysicsContinuum.class);

physicsContinuum_0.setPresentationName("Physics");

physicsContinuum_0.enable(ThreeDimensionalModel.class
);
    physicsContinuum_0.enable(SteadyModel.class);

physicsContinuum_0.enable(SingleComponentGasModel.cla
ss);

physicsContinuum_0.enable(SegregatedFlowModel.class);

physicsContinuum_0.enable(ConstantDensityModel.class)
;
    physicsContinuum_0.enable(LaminarModel.class);
    //Object tree focus that expands the physics
model manager node. All other open nodes are
collapsed.

selectAndExpandNodeExclusive(physicsContinuum_0.getMo
delManager());
}
```

Adding an Inclusive Focus to Step 2

Add an object tree focus that so that the **Continua** > **Physics** > **Materials** > **Dynamic Viscosity**node expands in the simulation tree when you click the link in the Simulation Assistant:

1. Within the `materialProperties()` method, insert the following code after the last statement:

```
//Object tree focus that expands the Dynamic
Viscosity node without affecting other open nodes.
selectAndExpandNode(gas_0.getMaterialProperties()
    .getMaterialProperty(DynamicViscosityProperty.cla
ss).getMethod());
```

The final code for the first step in Task 3 is displayed below:

```
public void materialProperties() {

    // code for Step 2: modifying the material
properties of air.
    Simulation simulation_0 = getSimulation();
    PhysicsContinuum physicsContinuum_0 =
        PhysicsContinuum)
simulation_0.getContinuumManager().getContinuum("Phys
ics"));
    SingleComponentGasModel
singleComponentGasModel_0 =

physicsContinuum_0.getModelManager().getModel(SingleC
omponentGasModel.class);
    Gas gas_0 =
        Gas)
singleComponentGasModel_0.getMaterial());
    ConstantMaterialPropertyMethod
constantMaterialPropertyMethod_0 =
        ConstantMaterialPropertyMethod)
gas_0.getMaterialProperties()
        .getMaterialProperty(ConstantDensityProperty.
class).getMethod());

constantMaterialPropertyMethod_0.getQuantity().setVal
ue(1.0);
    ConstantMaterialPropertyMethod
constantMaterialPropertyMethod_1 =
        ConstantMaterialPropertyMethod)
gas_0.getMaterialProperties()
        .getMaterialProperty(DynamicViscosityProperty
.class).getMethod());

constantMaterialPropertyMethod_1.getQuantity().setVal
ue(1.716E-5);
    Object tree focus that expands the Dynamic
Viscosity node without affecting other open nodes.
    selectAndExpandNode(gas_0.getMaterialProperties()
```

```
        .getMaterialProperty(DynamicViscosityProperty
.class).getMethod());
}
```

Testing the Object Tree Focus

Rebuild the project and test the modified assistant in STAR-CCM+. The first link in Task 3 expands the **Continua** > **Physics** > **Models** manager node.



The link in the second step expands the **Dynamic Viscosity** node without collapsing other open nodes.



Adding an Information Tag to the Assistant

Add an information tag to the assistant.

You can add information tags to a panel description or step. Information tags can contain more information about a specific step or task, but are not necessary for the task to be complete. Information tags are minimized by default, and can be expanded when required. The extra information is placed inside an information element in the XHTML file:

```
<information>Additional information goes here.</
information>
```

For this tutorial, you add the names of the physics models which are created in the first step of Task 3.

To add an information tag:

1. Return to `03_Physics.xhtml`.
2. Add the information tag to the first list item:

```
<ol>
    <!--First list item.-->
    <li><a
href="staraction:createPhysicsContinuum">Create</a>
a physics continuum.
        <!--Adds an information tag to first list
item.-->
        <information>The following physics models
are suitable for this case:
            <ul>
                <li>Three Dimensional</li>
                <li>Gradients</li>
                <li>Steady</li>
                <li>Gas</li>
                <li>Segregated Flow</li>
                <li>Constant Density</li>
                <li>Laminar</li>
            </ul>
        </information></li>
```

3. Save your project and rebuild it.

Contents

Testing the Information Tag

Testing the Information Tag

Test the information tag:

1. Launch STAR-CCM+ and test the new feature.

   A icon appears next to the first step in Task 3.



2. Click ⓘ to display the additional information.



3. Click ⓘ again to collapse the additional information.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Creating a Simulation Assistant project in NetBeans.
- Creating Simulation Assistant tasks, links and descriptions.
- Formatting the Simulation Assistant code and building the project.
- Loading and testing the Simulation Assistant code in STAR-CCM+.
- Adding conditions to the Simulation Assistant.
- Adding an object tree focus to the Simulation Assistant.
- Adding collapsible information to the Simulation Assistant.

Coupling with CAE Codes

The tutorials in this set illustrate various STAR-CCM+ features for coupling with CAE codes:

Contents

Abaqus File-Based Coupling: Exhaust Manifold
Abaqus Co-Simulation: Thermal Coupling
Abaqus Co-Simulation: Mechanical Coupling
STAR-CCM+ to STAR-CCM+ Coupling: Heat Transfer in a Chimney
GT-POWER Co-Simulation: 1D Coupling
OLGA Co-Simulation: 1D Coupling
AMESim Co-Simulation: 1D Coupling
Co-Simulation API: Spindle Valve

Abaqus File-Based Coupling: Exhaust Manifold

The purpose of this tutorial is to demonstrate the typical workflow in running a thermal fluid-structure interaction case using the file-based coupling method with Abaqus.

In the file-based coupling method, the thermal conduction in the solid structure is computed using Abaqus and the two separate simulations in Abaqus and STAR-CCM+ are coupled by exchanging data for the fluid-structure interface using files on the hard drive.

STAR-CCM+ passes local heat transfer coefficients and reference temperatures to Abaqus and Abaqus supplies the wall temperature to STAR-CCM+. In Abaqus terminology, the heat transfer coefficient is referred to as film coefficient and the reference temperature is referred to as sink temperature. The goal is to exchange data between the two codes for enough times so that you reach a converged solution. This can be accomplished using the GUI by manually performing the import-map-export operations, and launching Abaqus runs after each exchange of data. It is also possible to carry out the STAR-CCM+ steps automatically by using Java macros, which is demonstrated in this tutorial.

The case is an exhaust manifold subject to 10 m/s inlet gases at 500K. The objective is to find the steady-state solution for the flow and temperature in the exhaust gas, and the temperature in the solid manifold. Abaqus leads the process by providing the initial steady-state solution for the wall-temperature on the interior surface of the manifold, based on estimations for the thermal conditions in the manifold. From the initial Abaqus solution, STAR-CCM+ extracts the wall temperatures, map them onto the fluid model, and use them to set a boundary condition for the energy equation in the fluid region. From the STAR-CCM+ solution, the local heat transfer coefficients and reference temperatures are mapped and exported to Abaqus, and in the next iteration Abaqus uses this data to obtain a set of wall temperature data. These data exchanges are repeated until the wall temperatures, heat transfer coefficients, and reference temperatures have converged.



This type of coupling algorithm is referred to as a *loose coupling algorithm*, that is, where files on the hard disk are used to exchange data. STAR-CCM+ supports this method of coupling by allowing users to generate output files in a number of different formats that can be read by third-party software. In this tutorial, the heat transfer coefficients and reference temperatures are exported to a text file that can be read into Abaqus. This text file uses the keyword `*FILM` to specify the type of field data within the file. More information on the `*FILM` keyword can be found in the Abaqus Keywords Reference Manual.

An input file for the Abaqus analysis has been provided for this tutorial, which contains the manifold geometry and the settings required for the Abaqus job. A surface mesh for the fluid in the manifold has also been prepared for the tutorial; this is provided in `.dbs` format.

Note: You could also use the conjugate heat transfer capabilities within STAR-CCM+ to solve the solid manifold temperatures. The setup for such a case is demonstrated in the Solid Stress: Exhaust Manifold tutorial.

Contents

Prerequisites

Running the Initial Abaqus Job
Preparing the STAR-CCM+ Model
Importing the Abaqus Model
Importing Temperature Data
Mapping Imported Temperature Data
Preparing Scalar Scenes
Running the Simulation Using Imported Data
Reaching a Converged Solution
Summary

Prerequisites

The instructions in the Abaqus File-Based Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Running the Initial Abaqus Job

Run the Abaqus job to obtain the initial solution for the solid manifold.

First, copy the Abaqus input file required for this tutorial from the STAR-CCM+ installation directory into a separate directory. When the Abaqus job is run, a number of files are generated by Abaqus into this directory:

1. Create a directory for the tutorial called `fileBased`.
2. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
3. Copy the file `exhaust.inp` to the `fileBased` directory.

The Abaqus input file contains model and history data for the exhaust manifold. Run the Abaqus job first to obtain the initial solution for the solid manifold, based on the initial conditions and thermal loads specified in the input file. This allows you to use the initial Abaqus solution as a boundary condition in the fluid domain later on:

4. In a shell/command prompt, run the initial Abaqus job by typing:

```
% abaqus interactive job=exhaust
```

   where `abaqus` is the command to run Abaqus.

   Monitor the output in the *shell/command prompt* window. When the job is finished, the line `Abaqus JOB exhaust COMPLETED` is displayed.

During the Abaqus analysis, a number of files are created, including: `exhaust.com`, `exhaust.dat`, `exhaust.msg`, `exhaust.odb`, `exhaust.sim`, and `exhaust.sta`. The output database file, `exhaust.odb`, contains the results of the analysis. This file can also be used in Abaqus/CAE to visualize the results of the job.

Preparing the STAR-CCM+ Model

Prepare the STAR-CCM+ simulation for the fluid inside the manifold.

Preparing the STAR-CCM+ model involves importing the surface mesh and creating the volume mesh, and also adding an extrusion to the outlet to improve convergence. Define the physics in the fluid domain by selecting the appropriate physics models, and specify boundary conditions. When the model preparation is complete, import the Abaqus model into STAR-CCM+. You can then begin coupling the initial Abaqus solution with the fluid model.

1. Importing the Fluid Surface Mesh.
2. Setting Up the Mesh.
3. Generating the Volume Mesh.
4. Selecting Physics Models.
5. Defining Boundary Conditions.
6. Setting Up a Monitor.
7. Setting Stopping Criteria.

Contents

Importing the Fluid Surface Mesh
Setting Up the Mesh
Generating the Volume Mesh
Selecting Physics Models
Defining Boundary Conditions

Setting Up a Monitor
Setting Stopping Criteria

Importing the Fluid Surface Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied surface mesh.

1. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
2. Copy the file containing the fluid mesh, `manifoldFluid.dbs`, into the `/fileBased` directory created for the tutorial.
3. Launch STAR-CCM+.
4. Start a new simulation.
5. Select **File > Import > Import Surface Mesh**.
6. Navigate to `/fileBased`, select `manifoldFluid.dbs`, then click **Open**.
7. For *Import Mode*, select **Create New Region**, then click **OK**.

   A geometry scene is created in the *Graphics* window.



A new region, **Region 1**, is added within the **Regions** node.

8. Rename the **Region 1** node to `Fluid`.
9. Save the simulation as `manifold-gas.sim`.

Setting Up the Mesh

Set up the mesh for the imported geometry. This includes selecting meshing models, setting the global mesh reference values, and setting the mesh extrusion parameters.

Due to the geometry of the manifold, the gas flowing through the part is highly turbulent at the outlet. As this could limit the likelihood of the solution converging, use the mesh extruder feature in STAR-CCM+ to add a length of pipe at the manifold outlet. This aids good flow characteristics and helps the solution to converge. Although it is unlikely that such a length of pipe represents an exhaust pipe in reality, it is appropriate for a case such as this where you are concerned with thermal solution at the fluid-solid interface.

1. Selecting Meshing Models
2. Selecting the Global Mesh Reference Values
3. Setting the Mesh Extrusion Parameters

Contents

Selecting Meshing Models
Setting the Global Mesh Reference Values
Setting the Mesh Extrusion Parameters

Selecting Meshing Models

Select meshing models for the imported geometry, and set prism layer parameters.

Use the mesh extruder to add a length of pipe at the manifold outlet. Use the polyhedral mesher and prism layer mesher to create a volume mesh. Also use the surface remesher tool to specify a different mesh size.

1. For the mesh continuum, **Continua** > **Mesh 1**, select the following meshing models:

Group Box	Meshing Models
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**

Group Box	Meshing Models
	Extruder

Set the Prism Layer Mesher to use the thickness of the cell next to the wall as the defining parameter. STAR-CCM+ distributes the remaining cell layers based on a geometric progression to the overall thickness of the prism layer.

2. Select the **Mesh 1 > Models > Prism Layer Mesher** node and set *Stretching Mode* to **Wall Thickness**.

Setting the Global Mesh Reference Values

Define the global reference values for the volume mesh.

1. Edit the **Continua > Mesh 1 > Reference Values** node and set the following properties:

Node	Property	Value
Base Size	*Value*	`0.007 m`
Number of Prism Layers	*Number of Prism Layers*	`3`
Prism Layer Thickness	*Size type*	**Absolute**
▮▮ **Absolute Size**	*Value*	`0.0035 m`
Surface Size > Relative Minimum Size	*Percentage of Base*	`50`

Setting the Mesh Extrusion Parameters

Set up the mesh extruder to add a straight section of pipe at the manifold outlet.

1. Edit the **Regions > Fluid > Boundaries > Outlet** node and set the following properties:

Node	Property	Setting
Mesh Conditions > Extrusion Type	*Extrusion Type*	**Constant rate normal**
Mesh Values > Normal Extrusion Parameters	*Number of Layers*	`18`
	Stretching	`2.0`

Node	Property	Setting
	Magnitude	`0.1 m`
	Specify a New Region	**None** (default)

Note: **Specify a New Region** is set to **None**. This means that the extrusion is added to the Fluid region.

2. Save the simulation.

Generating the Volume Mesh

The mesh setup is complete, and you can now run the polyhedral volume mesher.

The boundary types of the inlet and outlet boundaries were automatically set to Velocity Inlet and Flow-Split Outlet when the geometry was imported, so on this occasion you do not need to specify the boundary types before generating the mesh. Had the boundary types not been set automatically, you would need to set the inlet and outlet boundary types to make sure that prism layers are not generated these faces.

1. In the *Mesh Generation* toolbar, click (**Generate Volume Mesh**).
2. When the mesh generation is complete, save the simulation.
3. Visualize the volume mesh with a method of your choice, such as creating a mesh scene.

The mesh appears as shown below.



Selecting Physics Models

Select physics models for the Fluid region.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following physics models in order:

Group box	Model
Time	**Steady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Ideal Gas**
Energy	**Segregated Fluid Enthalpy**
Viscous Regime	**Turbulent**
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**

2. To review the models, open the **Physics 1** > **Models** node.



Defining Boundary Conditions

The gas entering the manifold is traveling at 10 m/s with a temperature of 500 K. Specify these boundary conditions under the inlet boundary node, and also select a different turbulence specification.

To set the boundary conditions:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Inlets** node.
2. Set the following properties:

Node	Property	Setting
Physics Conditions		
Turbulence Specification	*Method*	**Intensity + Length Scale**
Physics Values		
Static Temperature > Constant	*Value*	`500 K`

Node	Property	Setting
Velocity Magnitude > Constant	*Value*	`10 m/s`

3. Save the simulation.

Setting Up a Monitor

Create a report, monitor, and plot to monitor the temperature at a point inside the fluid domain, close to the solid-fluid interface.

By default, STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the fluid solution. You also need a way to assess convergence of the outer coupled solution between Abaqus and STAR-CCM+. Monitor the temperature at a point inside the fluid domain, close to the solid-fluid interface. When the outer coupled solution is converged, this temperature comes to a steady value.

1. Within the **Derived Parts** node, create a new point probe with the following settings:

Property	Setting
Input Parts	**Fluid**
Point	`[0.251, 0.107, -0.0405] m`
Display	**No Displayer**

A new node, **point**, appears within the **Derived Parts** node.

Use the point probe to monitor the temperature in the fluid:

2. Right-click the **Reports** node and select **New Report > Maximum**.
3. Rename the **Reports > Maximum 1** node to `Temperature Point`.
4. Edit the **Reports > Temperature Point** node and set the following properties:

Property	Setting
Scalar Field Function	**Temperature**
Parts	**point**

To monitor the convergence of the solution while the simulation is running:

5. Right-click the **Reports > Temperature Point** node and select **Create Monitor and Plot from Report**.

   Two nodes, **Temperature Point Monitor** and **Temperature Point Monitor Plot**, appear underneath the **Monitors** and **Plots** nodes, respectively.



6. Open the **Temperature Point Monitor Plot**.

Setting Stopping Criteria

Set the stopping criteria so that the simulation stops running when the variance of temperature across the outlet boundary falls below 0.001 K^2.

You want STAR-CCM+ to stop iterating when the variance of temperature across some part of the simulation falls to a sufficiently low value. Monitor the variance of temperature across the outlet boundary.

Assume that the solution is sufficiently converged when this variance falls below 0.001 K^2.

1. Right-click the **Monitors** node and select **New Monitor > Field Variance**.
2. Rename the **Field Variance 1** node to `T Variance Monitor`.
3. Edit the **T Variance Monitor** node and set the following properties:

Property	Setting
Publish Mean	**Activated**
Field Function	**Temperature**
Trigger	**Iteration**
Enable sliding sample window	**Activated**
Parts	**Regions > Fluid > Outlet**

When *Field Function* is set to **Temperature**, a new field function, called **Variance of Temperature**, is created within the **Tools > Field Functions** node.

When *Enable sliding sample window* is activated, a new node, called **Sliding Window**, appears within the **T Variance Monitor** node.

4. Select the **T Variance Monitor > Sliding Window** node and set *Sliding sample window size* to `10`.



To use the **Variance of Temperature** field function as a stopping criterion, you must:

- Create a maximum value report to return the function value on the outlet boundary.
- Create a monitor to track the value of this report. Also, create a plot so that the value can be visualized.
- Create a stopping criterion from the monitor.

5. Right-click the **Reports** node and select **New Report > Maximum**.
6. Rename the **Maximum 1** node to `Outlet T Variance`.
7. Edit the **Reports > Outlet T Variance** node and set the following properties:

Property	Setting
Scalar Field Function	**Variance of Temperature**
Parts	**Regions > Fluid > Outlet**

8. Right-click the **Outlet T Variance** node and select **Create Monitor and Plot from Report**.

   Two new nodes are created: **Monitors > Outlet T Variance Monitor**, and **Plots > Outlet T Variance Monitor Plot**.
9. Open the **Plots > Outlet T Variance Monitor Plot**.

To create the stopping criterion based on the variance report monitor:

10. Right-click the **Monitors > Outlet T Variance Monitor** node and select **Create Stopping Criterion from Monitor**.
11. Edit the **Stopping Criteria > Outlet T Variance Monitor Criterion** node and set the following properties:

Node	Property	Setting
Outlet T Variance Monitor Criterion	*Logical Rule*	**And**
▮▮ **Minimum Limit**	*Minimum Value*	`0.001 K^2`

   Setting the logical rule to **And** means that the solution only stops when this criterion is satisfied with any other criterion whose logical rule is set to **And**.

To prevent the solution from stopping at the first iteration (when the variance of temperature across the outlet is zero), the Maximum Steps criterion is set to 10 iterations and its logical rule is also set to And:

12. Edit the **Stopping Criteria > Maximum Steps** node and set the following properties:

Property	Setting
Maximum Steps	`10`
Logical Rule	**And**

13. Save the simulation.

Importing the Abaqus Model

Import the solid model into the STAR-CCM+ simulation using the Abaqus `.inp` file.

When preparing a model in Abaqus for file-based coupling, include a surface definition (the *SURFACE keyword in the *.inp file). This command identifies the face of the 3D elements that lie at the fluid-structure interface. In this case, the interface surface has been called INNER, and two other surfaces have been defined to identify the faces for the flanges and the manifold exterior.

1. Select **File > Import > Import CAE Model...**.
2. In the *Open* dialog, select file `exhaust.inp` and click **Open**.
3. In the *Import Options* dialog, deactivate *Open geometry scene after import*.



When the *Import Mode* is set to **Create New Region(s) and Imported Model(s)**, a new node called **Imported Models**, is added to the

simulation tree. This node allows you to carry out file-based coupling operations.

4. Click **OK**.

   The imported Abaqus model is stored in the simulation under the **Regions** node. The model surfaces and elements (that is, shells and volumes), are also listed under the **Imported Models** node.

The solid model can be visualized in a geometry scene, as any other region in STAR-CCM+. To visualize the solid model:

5. Create a geometry scene.
6. Select the **Geometry Scene 2** > **Displayers** > **Geometry 1** > **Parts** node, and set *Parts* to **Regions** > **manifold-1**.
7. Select the **Geometry Scene 2** > **Displayers** > **Outline 1** > **Parts** node, and set *Parts* to **Regions** > **manifold-1**.



To highlight the imported surfaces in the scene, select their nodes under the **Imported Models** > **Imported Surfaces** node. Alternatively, select the corresponding boundaries within the **Regions** node.

Importing Temperature Data

CAE solution data is imported in a separate step to the model. This tutorial requires that you import the temperature data from Abaqus.

To import the Abaqus temperature data on the **INNER** surface of the Abaqus model:

1. Right-click the **Imported Models** > **Abaqus: exhaust** node and select **Import CAE Solution Data...**



2. In the *Open* dialog, navigate to the `/fileBased` directory, select the file `exhaust.odb` and click **Open**.
3. In the *Import Data Options* dialog that appears, do the following:
   a) In the *Abaqus Command* box enter the command to launch Abaqus, for example, `abaqus`. On Windows, enter `abaqus.bat`. If your environment is not set to include the command automatically from a command prompt, enter the full path.
   b) Accept the remaining default settings by clicking **OK**.

Note: During this operation Abaqus is launched in the background to read the data, therefore you must have access to Abaqus.



The message `Importing: Finished` is displayed in the status bar when the operation is complete.



Contents

Visualizing the Imported Temperature

Visualizing the Imported Temperature

Visualize the imported temperature data in a scalar scene.

During the process of importing the temperature data, STAR-CCM+ created a field function called **Imported Temperature**.

Use the **Imported Temperature** field function to visualize the temperature data:

1. Create a new scalar scene.
2. Rename the **Scalar Scene 1** node to `Imported Temperature`.
3. Edit the **Displayers** > **Outline 1** > **Parts** node.
4. In the *Parts* dialog, deselect the **Regions** node, then click **OK**.

5. Expand the **Displayers** > **Scalar 1** node and set the following properties:

Node	Property	Setting
Parts	*Parts*	**Imported Models > Abaqus: exhaust > Imported Surfaces > INNER**
Scalar Field	*Function*	**Imported Temperature**



Mapping Imported Temperature Data

Once the data has been imported into STAR-CCM+, map it onto the corresponding boundary in the Fluid region.

To map the imported temperature data:

1. Right-click the **Imported Models** > **Abaqus: exhaust** node and select **Map data** > **Surface Data...**



The *Map Imported Model Data* dialog appears.

In this case, the **INNER** surface temperature data is mapped to the **Fluid-Wall** boundary.

2. In the *Map Imported Model Data* dialog, set the following properties:

Group box	Property	Setting
Map Direction	*Imported Model to STAR-CCM+*	**Activated**
Mapping Inputs	*Imported Surfaces*	**Imported Surfaces > INNER**
	Boundaries	**Fluid-Wall**
	Data to map	**Imported Temperature** (default)

The completed dialog is shown below.



3. Click **Apply**, then **Close**.

The *Output* window displays `Done` when the mapping process is complete. A node, **Surface Data Mapper 1**, is added under the **Tools >**

Data Mappers node. This node contains all the settings that you entered in the mapping dialog in the previous steps.

It would also have been possible to create an empty data mapper separately, specify the properties, then right-click the data mapper and select **Map Data**. The right-click menu options are a quicker way of setting up a data mapper.

Contents

Visualizing the Mapped Imported Temperature

Visualizing the Mapped Imported Temperature

Visualize the mapped imported temperature in a scalar scene.

In the same way as for the imported temperature data, STAR-CCM+ created a **Mapped Imported Temperature** field function during the mapping operation. Create a scalar scene and use this field function to visualize the mapped temperature data on the fluid wall boundary.

The new scalar scene shares a number of settings with the *Imported Temperature* scene, so create a duplicate of this scene to use as a starting point:

1. Copy the **Scenes > Imported Temperature** node and paste it onto the **Scenes** node.
2. Rename the **Copy of Imported Temperature** node to `Mapped Imported Temperature`.
3. Edit the **Displayers > Scalar 1 > Parts** node.
4. In the *Edit* dialog, deselect the **Imported Models** node, then select the **Regions > Fluid > Boundaries > Fluid-Wall** node, and click **OK**.
5. Select the **Scalar 1 > Scalar Field** node, and set *Function* to **Mapped Imported Temperature**.

A scene similar to the one shown below is displayed.



Comparing the imported temperature contour with the mapped imported temperature contour, it is clear that the imported temperature data from Abaqus has been mapped successfully onto the fluid wall boundary.

Preparing Scalar Scenes

Prepare two scalar scenes to visualize the contours of fluid temperature and velocity on a plane section as the simulation runs.

To create a plane section through the fluid region:

1. Activate *Mesh Scene 1* by clicking the corresponding tab at the top of the *Graphics* window.
2. Within the Derived Parts node, create a new plane section.
3. In the *Create Section* dialog, set the following properties:

Property	Setting
Input Parts	**Fluid**
origin	`[0, 0, 0] m`
normal	`[0, 0, 1] m`
Display	**No Displayer**

4. Click **Create** and then **Close**.
   A node, called **plane section**, appears within the **Derived Parts** node.

To highlight this part in the active scene:

5. Expand the **Derived Parts** node and select **plane section**.



Use the plane section to visualize the contours of fluid temperature and velocity:

6. Creating a Fluid Velocity Scalar Scene
7. Creating a Fluid Temperature Scalar Scene

Contents

Creating a Fluid Velocity Scalar Scene
Creating a Fluid Temperature Scalar Scene

Creating a Fluid Velocity Scalar Scene

Use the first scalar scene to display the velocity magnitude contour on the plane section.

1. Create a new scalar scene.
2. Rename the **Scalar Scene 1** node to `Fluid Velocity`.
3. Select the **Fluid Velocity > Displayers > Outline 1 > Parts** node and set *Parts* to **Derived Parts > plane section**.
4. Expand the **Fluid Velocity > Displayers > Scalar 1** node and set the following properties:

Node	Property	Setting
Parts	*Parts*	**Derived Parts > plane section**
Scalar Field	*Function*	**Velocity > Magnitude**

5. Click (**Save-Restore-Select views**) and select **Views > -Z > Up +Y**
6. Zoom in and position the part as shown below.



Creating a Fluid Temperature Scalar Scene

Use the second scalar scene to display the fluid temperature on the plane section.

1. Copy the **Scenes** > **Fluid Velocity** node and paste it onto the **Scenes** node.
2. Rename the **Copy of Fluid Velocity** node to `Fluid Temperature`.
3. Expand the **Scenes** > **Fluid Temperature** > **Displayers** > **Scalar 1** node.
4. Select the **Scalar Field** node and set *Function* to **Temperature**.

   The scene appears as follows.



5. Save the simulation.

Running the Simulation Using Imported Data

The STAR-CCM+ simulation runs using the imported temperature data as a boundary condition.

The temperature data that has been imported from the Abaqus model and mapped onto the fluid boundary can be used to define a Temperature boundary condition. This fixes the temperature at the Fluid-Wall boundary, that is, the interface between the fluid and solid domains, to be the temperature that is calculated in the Abaqus analysis. It is expected that this initial temperature solution from Abaqus is a good starting point for the coupled simulation.

Contents

Specifying a Temperature Boundary Condition
Running the Simulation
Visualizing Initial Results

Specifying a Temperature Boundary Condition

Set the temperature at the fluid-wall to the surface temperature imported from the Abaqus simulation.

To set the temperature at the fluid wall:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Fluid-Wall** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Temperature**
Physics Values		
Static Temperature	*Method*	**Field Function**
Field Function	*Scalar Function*	**Mapped Imported Temperature**

2. Save the simulation.

Running the Simulation

The simulation is now fully set up and ready to be run using the imported data.

1. Run the simulation

   The iteration progress can be monitored in the *Output* window. Monitor the convergence by viewing the **Temperature Point Plot** and the **Residuals** using the visualization tabs in the **Graphics** window.

   When the convergence criterion is satisfied, the message `Stopping criteria Maximum Steps and Outlet T Variance Monitor Criterion satisfied` is displayed in the *Output* window.

2. Save the simulation.

Visualizing Initial Results

Visualize the fluid velocity and fluid temperature scalar scenes that you prepared earlier.

1. Open the scenes to visualize the results from the initial run:





The **Temperature Point Monitor Plot** shows that the solution temperature solution for the initial run has converged.



Reaching a Converged Solution

Perform multiple data exchange cycles between Abaqus and STAR-CCM+ to reach a converged solution.

Often, more iterations are required to obtain a fully converged solution. The number of iterations necessary between STAR-CCM+ and Abaqus depend on how close the initial conditions are to the final temperature solution. Perform the import and mapping operations each time data is passed from Abaqus to STAR-CCM+. When data is passed from STAR-CCM+ to Abaqus, map the data to the solid model, and export it.

To facilitate this process, record two macros for the following tasks:

- Mapping boundary data onto Abaqus surface and exporting it to Abaqus.
- Importing data from Abaqus and mapping it onto a boundary in STAR-CCM+.

These macros can be played for each required coupling iteration rather than going through the steps one at a time.

Contents

Mapping Boundary Data onto the Abaqus Surface
Exporting Mapped Data
Modifying the Abaqus Input File and Running the Abaqus Job
Importing and Mapping Temperature Data
Running the Simulation
Running Abaqus and STAR-CCM+ Sequentially

Mapping Boundary Data onto the Abaqus Surface

Record a macro to map STAR-CCM+ boundary data to the Abaqus surface.

When the fluid solution has converged, the thermal data at the solid-fluid interface, that is, the Fluid-Wall boundary, can be mapped from STAR-CCM+ to the INNER surface on the imported model. This data can then be exported in a format that can be used to define a thermal load in the Abaqus model.

Before carrying out the mapping and export processes, start recording the first macro. Once the recording has started, each action you perform in STAR-CCM+ is included in the macro. Use this macro later on to repeat the map and export processes.

To start recording the first macro:

1. Click (**Start Recording**) to initiate the recording process.
2. In the *Save* dialog, enter `loopMapExport` for the name of the Java macro.
3. Click **Save**.

From this point onwards, any recordable action you make in

STAR-CCM+ generates lines of Java code in this page. Clicking (**Pause Recording**) pauses generation of the Java code; clicking the same

button resumes generation of the Java code. Clicking (**Stop Recording**) ends all generation of the Java code and write the macro to file. Do not click either of these buttons for now.

Proceed by mapping the STAR-CCM+ solution data onto the Abaqus model surface:

4. Right-click the **Imported Models** > **Abaqus: exhaust** node and select **Map data** > **Surface Data...**



The *Map Imported Model Data* dialog appears, which allows you to set up the mapping operation.

5. In the *Map Imported Model Data* dialog, set the following properties:

Group box	Property	Setting
Map Direction	*STAR-CCM+ to Imported Model*	**Activated**
Mapping Inputs	*Imported Surfaces*	**Imported Surfaces > INNER**
	Boundaries	**Fluid-Wall**
	Data to map	**Local Heat Transfer Coefficient** (default)
Reference Temperature Evaluation Option	*Use Local Heat Transfer Reference Temperature*	**Selected** (default)

The completed dialog is shown below.



6. Click **Apply**, then **Close**.
   The *Output* window displays the time to map the data, and `Done` when the mapping process is complete.

Exporting Mapped Data

Once the data has been mapped onto the imported Abaqus surface, it can be exported to an .inp file, which then can be read into the Abaqus job. Stop the macro recording when the export process is complete.

1. Right-click the **Imported Models** > **Abaqus: exhaust** node and select **Export Mapped Data...**



In the *Export File* dialog:

2. Click the *Imported Surfaces* tab and select the **INNER** node.
3. Click the *FieldFunctions* tab and enter the following settings:

Property	Setting
Export Data as	**Heat Transfer Coefficient**
Data to Export	**Mapped Local Heat Transfer Coefficient**
Export Units	**W/m^2-K**
Tref to Export	**Mapped Local Heat Transfer Reference Temperature**
Export Mode	**Surface Data**
Append to Existing File	**Deactivated**
File Name	**mappedHTC**
Files of Type	**Abaqus input files (*.inp)**



4. Click **Save**.

   The export process is now complete, and the Java macro contains all the necessary steps for the map-export process.

5. Click (**Stop Recording**) to end the recording and write the macro to the file.

Modifying the Abaqus Input File and Running the Abaqus Job

Use the solution data from the fluid solution to define a thermal load in the Abaqus model.

Add a command in the Abaqus input file:

1. Open the `exhaust.inp` file in a script editor of your choice.
2. Scroll down to the section that begins with the following:

```
**
** STEP: Step-1
**
```

This section of the file, up to `*End Step`, contains the history data for the Abaqus job, including solver parameters, boundary conditions, loads, and outputs.

3. To cause Abaqus to read in the heat transfer coefficients from the `mappedHTC.inp` file, add the following lines that are shown in bold text:

```
*Output, field, variable=PRESELECT
*Output, history, frequency=0
**
*Include, input=mappedHTC.inp
**
*End Step
```

4. Save the file under the name `exhaust-mod.inp` and exit the script editor.
5. In a shell/command prompt, run the Abaqus job using the modified input file by typing the following command:

```
> abaqus interactive job=exhaust-mod
```

where abaqus is the command you use to launch Abaqus.

The set of files created after running this job are named `exhaust-mod.*`

Importing and Mapping Temperature Data

The results from the Abaqus analysis, exhaust-mod, contain an updated set of temperature data. Import the set into STAR-CCM+ and map it onto the Fluid-Wall boundary using the same technique as described previously.

In a similar way as for the map-export process, record a macro for the process of importing the latest data from Abaqus, and mapping it onto the boundary. Include some additional steps in this macro, so that the simulation begins running once the data has been mapped.

To start recording the second macro:

1. Click (**Start Recording**) to initiate the recording process.
2. In the *Save* dialog, enter `loopImportMap` for the name of the Java macro.
3. Click **Save**.

Proceed by importing the temperature data from the latest Abaqus analysis:

4. Right-click the **Imported Models** > **Abaqus: exhaust** node and select **Import CAE Solution Data**.
5. In the *Open* dialog, select the `exhaust-mod.odb` file and click **Open**.
6. In the *Import Data Options* dialog:
   a) In the *Abaqus Command* box enter the command to launch Abaqus, for example, `abaqus`. On Windows, enter `abaqus.bat`. If your environment is not set to include the command automatically from a command prompt, enter the full path.
   b) Accept the remaining default settings by clicking **OK**.

Note: During this operation Abaqus is launched in the background to read the data, therefore you must have access to Abaqus.



The message `Importing: Finished` is displayed in the status bar when the operation is complete. The new set of imported temperature data can be visualized in the *Imported Temperature* scene you created earlier.

7. Right-click the **Imported Models** > **Abaqus: exhaust** node and select **Map data** > **Surface Data...**

   The *Map Imported Model Data* dialog appears. This dialog lets you carry out mapping operations.

In this case, the **INNER** surface temperature data is mapped to the **Fluid-Wall** boundary.

8. In the *Map Imported Model Data* dialog, set the following properties:

Group box	Property	Setting
Map Direction	*Imported Model to STAR-CCM+*	**Activated**
Mapping Inputs	*Imported Surfaces*	**Imported Surfaces > INNER**
	Boundaries	**Fluid-Wall**
	Data to map	**Imported Temperature** (default)

The completed dialog is shown below.



9. Click **Apply**, then **Close**.

The *Output* window displays `Done` when the mapping process is complete. The results of the mapping operation can be visualized in the *Mapped Imported Surface Temperature* scene.

Running the Simulation

Following the import and mapping operations in the previous section, you can proceed by running the simulation.

To make sure that the **Temperature Point Monitor** stopping criterion is reset, first perform a single iteration:

1. From the menu, select **Solution** > **Step**.



2. Run the simulation.
3. When the simulation has finished running, save the simulation.
4. Click (**Stop Recording**) to end recording and to write the macro to the file.

In the *Temperature Point Plot*, you can observe a "jump" in the temperature solution when the next set of temperature data from Abaqus was imported. This jump represents the exchange of solution data between the two codes. As the Abaqus solution converges with the STAR-CCM+ solution, these jumps become smaller, until you can say that the solution is fully converged.

Running Abaqus and STAR-CCM+ Sequentially

Following the completion of the second run in STAR-CCM+, you can now pass thermal data back to Abaqus and run the Abaqus job. You can then return to STAR-CCM+, import the latest set of temperature data, map the data, and run the simulation.

Overwrite the old data in the `mappedHTC.inp` file with the latest thermal data from STAR-CCM+. Recall that the first macro you recorded,

`loopMapExport.java`, performs these actions automatically, so you can proceed by playing this macro.

1. Click (**Play Macro**).
2. In the *Open* dialog, select `loopMapExport.java` and click **Open**.

   When the macro has finished running, a message beginning `Export Mapped Data:MappedLocal` is displayed in the *Output* window. This signifies that the new set of thermal data is ready to be used in the Abaqus analysis.

Proceed by running the Abaqus job:

3. In a shell/command prompt, run the Abaqus job by typing:

   ```
   > abaqus interactive job=exhaust-mod
   ```

   where abaqus is the command you use to launch Abaqus. When prompted, confirm that it is OK to overwrite the old job files.

When the Abaqus job is complete, you can return to STAR-CCM+, import the latest set of temperature data, map the data, and run the simulation. These steps can be carried out automatically by using the macro `loopImportMap.java`:

4. In STAR-CCM+, click (**Play Macro**).
5. In the *Open* dialog, select `loopImportMap.java` and click **Open**.

   Recall that this macro has finished running when the stopping criterion has been reached and the simulation is saved.

Each time the file-based coupling occurs, you move closer to reaching the final fluid-solid thermal solution. Continue to run Abaqus and STAR-CCM+ sequentially by repeating the following steps:

6. In STAR-CCM+, click (**Play Macro**) and run the `loopMapExport.java` macro.
7. In a shell/command prompt, run the Abaqus job by typing:

   ```
   > abaqus interactive job=exhaust-mod
   ```

8. In STAR-CCM+, click (**Play Macro**) and run `loopImportMap.java` macro.

   The model converges after approximately four coupling iterations, and the *Temperature Point Monitor Plot* appears as shown below.



The thermal solution in the solid model can be post-processed in Abaqus/CAE using the `exhaust-mod.odb` file.

Summary

This tutorial has introduced the following STAR-CCM+ features:

- Creating an asymptotic stopping criterion from a monitor
- Importing an Abaqus model
- Importing and visualizing temperature data for the Abaqus model
- Mapping the imported temperature data
- Using the imported data to define a boundary condition
- Running a simulation using imported temperature data
- Mapping and exporting STAR-CCM+ solution data
- Using exported data in an Abaqus job
- Recording macros for import-map-export processes
- Reaching a converged solution between Abaqus and STAR-CCM+

Abaqus Co-Simulation: Thermal Coupling

This tutorial demonstrates how to run a thermal co-simulation case using STAR-CCM+ and Abaqus. Co-simulation involves a strong coupling between the two codes, in which data exchange occurs at frequent intervals, that are known as coupling steps. In this case, data is exchanged at the end of each time-step, and a relatively large time-step is used to approximate a steady-state solution.

The deep level of communication between the solvers allows you to obtain a full solution across the solid-fluid interface. The solution includes the effects of the interaction between the solid and fluid domains. STAR-CCM+ passes heat flux data to Abaqus, and Abaqus passes temperatures back to STAR-CCM+. In each case, the thermal data defines a boundary condition in the model, to be used in the next set of iterations.

The coupling is carried out using the SIMULIA Co-Simulation Engine, which exchanges data between the two models automatically during the co-simulation. This technique differs from the file-based method in that the coupling between the two codes is much stronger; the exchange of data is carried out automatically, and STAR-CCM+ and Abaqus run simultaneously.

To prepare a co-simulation case, two models are required. The Abaqus model is for the solid domain, and is used to create an Abaqus input file. In STAR-CCM+ the fluid model is set up, with the appropriate boundary conditions defined. For each model, specify a set of co-simulation parameters which define the boundary/surface at the fluid-structure interface, the field functions that are imported and exported at this interface, the frequency of coupling steps, and so on. With the co-simulation parameters defined, Abaqus can be launched from the STAR-CCM+ GUI, and the co-simulation run.

Before running the co-simulation, it is good practice to run standalone cases for the fluid and solid models to make sure that they converge independently. In this example, the Abaqus model could be run standalone using a surface film condition, `*SFILM`, on the INNER surface with estimated constant values for the sink temperature and film coefficient. Similarly, in STAR-CCM+, the simulation could be run using an estimated temperature on the Fluid-Wall boundary.

The case is an exhaust manifold subject to 10 m/s inlet gases at 500 K. A solution for the flow and temperature in the exhaust gas, and the temperature in the solid manifold is required.



This tutorial is divided into three sections. In the first section, the Abaqus model and input file are created using Abaqus/CAE. In the second section, the fluid model is prepared in STAR-CCM+. The third section involves editing the Abaqus input file for co-simulation using a script editor, defining the co-simulation session in STAR-CCM+, and running the co-simulation. The third section also demonstrates how to start and close Abaqus via the STAR-CCM+ user interface.

Note: The figures in this tutorial were obtained using Abaqus 6.10-2; images may differ slightly in subsequent releases.

Contents

Prerequisites
Preparing the Abaqus Model
Preparing the STAR-CCM+ Model
Setting Up and Running the Co-Simulation
Summary

Prerequisites

The instructions in the Abaqus Thermal Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Preparing the Abaqus Model

In this section, the model for the exhaust manifold is prepared in Abaqus/CAE, culminating in the creation of the Abaqus input file (`*.inp`).

It is recommended that you save the files for each application in separate subdirectories within an empty directory, to help you to organize the files that are created during the co-simulation.

1. In your working directory, create a directory `/thermalCoSimTutorial`.
2. In that directory, create two additional directories: `/abaqus` and `/starccm+`.

   The files for the Abaqus part of the co-simulation are saved in `/abaqus` and the files for STAR-CCM+ are saved in `/starccm+`.

   If you are unfamiliar with Abaqus/CAE, it is recommended that you work through the examples in the Getting Started with Abaqus: Interactive Edition manual before attempting this part of the tutorial, in particular "Abaqus Basics", section 2.

   If you do not wish to work through the steps involving Abaqus/CAE you can jump directly to the STAR-CCM+ section of the tutorial in which the fluid model is prepared. A copy of the input file that is created using Abaqus/CAE has been provided in the STAR-CCM+ installation, and can be used later on when setting up the co-simulation. To start the tutorial from the STAR-CCM+ section, proceed to Preparing the STAR-CCM+ Model.

   To work through the steps to set up the model in Abaqus/CAE, proceed to the next section.

Contents

Starting Abaqus/CAE and Importing Geometry
Creating a Material
Defining and Assigning Section Properties
Creating a Partition

Creating an Assembly
Defining Surfaces and Sets
Defining the Analysis Steps
Defining Initial Conditions
Defining Boundary Conditions
Meshing the Model
Creating the Abaqus Input File

Starting Abaqus/CAE and Importing Geometry

Import the manifold geometry into Abaqus/CAE as a part.

An exhaust manifold geometry has been prepared for this tutorial and saved in the STEP (Standard for the Exchange of Product model data) file format. Begin by copying this file from the STAR-CCM+ installation directory to the directory created for the Abaqus files.

1. Navigate to the `[INSTALL_DIR]/doc/startutorialsdata/coupling/data` directory and copy the file `manifold.stp` into the `/thermalCoSimTutorial/abaqus/` directory.

To import the manifold geometry into Abaqus/CAE as a part:

2. Launch Abaqus/CAE and select the option to Create Model Database **With Standard/Explicit Model** in the *Start Session* dialog.
3. Select **File** > **Import** > **Part**.
4. In the *Import Part* dialog, change **File Filter** to **STEP** using the drop-down list.
5. Navigate to the `/thermalCoSimTutorial/abaqus/` directory and select the file `manifold.stp`.
6. Click **OK**.
7. In the *Create Part from STEP File* dialog, accept the default options by clicking **OK**.

If a message warning that the part contains imprecise geometry is displayed, you can ignore it. In this case, imprecise geometry does not affect the analysis, since the part is meshed using tetrahedral elements.

8. If a warning message regarding imprecise geometry is displayed, click **Dismiss**.

The geometry import process takes a few minutes, and when it is complete the manifold part is displayed in the viewport.

Creating a Material

Create a material definition for the manifold and specify the material properties that are required for a thermal analysis.

In general, given the type of analysis, you specify the following material properties:

- For steady-state heat transfer analyses, you specify:
  - Thermal Conductivity
- For transient heat transfer analyses, you also specify:
  - Density
  - Specific Heat
- For thermal-structural analyses, you also specify:
  - Young's Modulus
  - Poisson Ratio
  - Coefficient of Thermal Expansion

In your own cases, you can also specify other properties depending on the material and analysis requirements.

As this is a thermal analysis, you only specify the material thermal conductivity, density, and specific heat.

To set the material properties for this case, create a material definition and specify numerical values for density, thermal conductivity, and specific heat, without entering units. In this tutorial, the units used by Abaqus are automatically set to SI units.

1. In the *Model Tree*, double-click the **Materials** container to create a material definition.
2. In the material editor that appears:
   a) Name the material `GRAY CAST IRON`.
   b) Select **General** > **Density** and enter a value of `7817` (kg/m^3).
   c) Select **Thermal** > **Conductivity** and enter a value of `55` (W/mK).
   d) Select **Thermal** > **Specific Heat** and enter a value of `446` (J/kg K).

Note: In Abaqus, all numerical values are cast in the same system of units. When coupling STAR-CCM+ to Abaqus, you can control what units Abaqus uses from your STAR-CCM+ simulation, through the **Co-Simulations** > **Abaqus Co-Simulation 1** > **Abaqus Units** node. In this

tutorial, you use SI units, that correspond to the default setting for the **Abaqus Units** node.

3. Click **OK** to close the dialog.

Defining and Assigning Section Properties

Create a section from the material and assign it to the manifold part.

To create a section definition:

1. In the Model Tree, double-click the **Sections** container.
2. In the *Create Section* dialog, change the name of the section to `ManifoldSection`.
3. Make sure that the category and type are set to **Solid** and **Homogeneous**, and click **Continue...**



4. In the *Edit Section* dialog, make sure that the **Material** is set to **GRAY CAST IRON** and click **OK**.



5. In the Model Tree, expand the **manifold** item underneath the **Parts** container and double-click **Section Assignments** in the list of part attributes.
6. To assign the section to the part, select the part in the viewport and click **Done** in the prompt area.
7. In the *Edit Section Assignment* dialog that appears, make sure that **Section** is set to **ManifoldSection** and click **OK**.

Creating a Partition

As the manifold geometry is complex, a partition is created to aid in the meshing process.

Note: Partition before defining surfaces and sets on the manifold to make sure that the correct faces are referenced when these features are defined.

1. Switch to the Mesh module by double-clicking **Mesh (Empty)** underneath the **Parts > manifold** container.
2. Select **Tools > Partition**.
3. In the *Create Partition* dialog, change *Type* to **Cell**, and select **Use n-sided patch** in the *Method* box.
4. At the bottom of the viewport, choose the option to **Select Edges**.
5. Hover the mouse over the edge that is shown in the following screenshot:



6. Left-click to select the edge; the whole loop is selected.



7. Click **Create Partition**, at the bottom of the viewport, to create the partition.

   When the outlet section has been partitioned, it is displayed in yellow in the viewport.
8. Close the *Create Partition* dialog and press `<Esc>` to exit the partitioning tool.

Creating an Assembly

An assembly contains all the geometry that forms the final finite element model. In this case, the assembly contains only the manifold geometry, which is added to the assembly by creating an instance of the part.

To create an instance:

1. In the Model Tree, expand the **Assembly** container and double-click **Instances**.
2. In the *Create Instance* dialog, make sure that **manifold** is selected, keep the default options and click **OK**.

Defining Surfaces and Sets

Specify the FSI interface as the inner wetted surface of the manifold model, and define two additional surfaces to specify thermal conditions on the manifold exterior and the flanges.

Surfaces and sets identify particular regions of the model with a unique name. They can be used when defining loads and boundary conditions, and must be used in co-simulation cases to identify the interface between the fluid and solid domains. In this case, specify the FSI interface as the inner wetted surface of the manifold model, and define two additional surfaces to specify thermal conditions on the manifold exterior and the flanges. Create two sets to apply temperature boundary conditions.

To create a surface:

1. Double-click **Surfaces** underneath the **Assembly** container in the Model Tree.
2. In the *Create Surface* dialog name the surface `INNER` and click **Continue**.



To select the inner faces of the model in the viewport:

3. At the bottom of the window, change the **Select the regions for the surface** option to **by angle**.
4. Enter an angle of `89` and press `Enter`.
5. Hover the mouse pointer over one of the inner faces, and after a moment the inner surfaces are highlighted. Click to confirm this selection.



6. Once the inner faces have been highlighted, click **Done** at the bottom of the viewport.

   The surface is shown under the **Surfaces** container.



7. Using the technique described above, create another surface called `FLANGES`.
8. In the viewport, hold the `<Shift>` key and click the four flat faces on the bottom of the inlet flanges.



9. Click **Done** at the bottom of the viewport.
10. Finally, create a third surface called `EXTERIOR`.
11. In the viewport, use `<Shift>`-click to select all the exterior faces on the manifold part, including any face on the model that was not part of the FLANGES or INNER surfaces.



12. Click **Done** to create the surface.

Next, define two geometry sets:

13. Double-click **Sets** underneath the **Assembly** container.
14. In the *Create Set* dialog, name the set `INLET` and click **Continue**.



15. Select the continuous set of edges around each inlet as shown below.



16. Click **Done** when you have selected these edges.

The second set includes the entire model and is used to define the initial temperature across the manifold:

17. Using the same technique as for the previous set, create another geometry set called `ALL`.
18. In the viewport, select the entire model by clicking in the space above and to the left of the model, and dragging the mouse down and to the right.

    A box is drawn around the model, and everything within the box is selected when you release the mouse button.



19. Click **Done** to confirm the set.

Defining the Analysis Steps

Define the heat transfer analysis step.

The analysis steps for the model include:

- An initial step, in which the initial temperature of the manifold is defined.
- A general heat transfer step, in which convection conditions are defined on the manifold exterior and flanges, and the temperature at the inlet is fixed at 500K. This step also includes the import of heat flux data from STAR-CCM+ via the SIMULIA Co-Simulation Engine.

By default, the Abaqus model already contains an initial step, which can be viewed by expanding the **Steps** container in the model tree.

To create the heat transfer step:

1. Double-click the **Steps** container in the model tree.
2. In the *Create Step* dialog, keep the default step name, `Step-1`, and select **Heat transfer** from the list.
3. Click **Continue...**



4. In the *Edit Step* dialog, select the *Basic* tab and set the following properties:

Property	**Setting**
Description	`Co-Simulation of Hot Exhaust Gases in the Manifold`
Response	**Transient** (default)
Time Period	`10000`



5. Select the *Incrementation* tab and set the following properties:

Property	Setting
Type	**Fixed**
Maximum number of increments	`2000`
Increment size	`250`



6. Click **OK** to confirm these settings and create the step.

Defining Initial Conditions

Define the manifold initial temperature.

To specify the initial temperature through the manifold using a predefined field:

1. Double-click the **Predefined Fields** container in the Model Tree.
2. In the *Create Predefined Field* dialog, set the following properties:

Property	Setting
Name	`InitialTemp`
Step	**Initial**
Category	**Other**
Type	**Temperature**



This predefined field must now be applied to a region of the model, and in this case apply it to the whole model via the ALL geometry set:

3. Click **Sets** to the bottom right of the viewport.
4. In the *Region Selection* dialog, select **ALL** and click **Continue**.
5. In the *Edit Predefined Field* dialog, set **Magnitude** to `500` and click **OK**.



Defining Boundary Conditions

In the heat transfer analysis step, the convection conditions and fixed temperature at the inlet become active. Specify the convection properties as interactions, since the manifold transfers heat to the surroundings. Specify the temperature condition at the inlet as a boundary condition.

1. Double-click the **Interactions** container in the Model Tree.
2. In the *Create Interaction* dialog, set the following properties:

Property	Setting
Name	`SURFFILM`
Step	**Step-1**
Type	**Surface film condition**

3. Click **Continue...**
4. In the *Region Selection* dialog, select **EXTERIOR** and click **Continue**.
5. In the *Edit Interaction* dialog, set the following properties:

Property	Setting
Film coefficient	`6.667`
Sink temperature	`300`

6. Click **OK**.
7. Repeat the steps above, this time using the following settings:

Property	Setting
Name	`SURFFILM-2`
Step	**Step-1** (default)
Type	**Surface film condition** (default)
Region	**FLANGES**
Film coefficient	`333.3`
Sink temperature	`500`

Specify the inlet temperature boundary condition next:

8. Double-click the **BCs** container in the Model Tree.
9. In the *Create Boundary Condition* dialog, set the following properties:

Property	Setting
Name	`Temp-BC`
Step	**Step-1** (default)
Type	**Temperature** (default)

10. Click **Continue...**
11. In the *Region Selection* dialog, select **INLET** and click **Continue...**
12. In the *Edit Boundary Condition* dialog, set *Magnitude* to `500` and click **OK**.

Meshing the Model

The next step involves defining and generating a mesh for the manifold part and is carried out within the Mesh module in Abaqus.

1. To switch to the Mesh module, expand the **Parts > manifold** container and double-click **Mesh (Empty)**.

The color of the manifold indicates which method Abaqus uses to mesh the part. Currently, most of the manifold is colored orange indicating that it cannot be meshed using the default element shape. Select a more appropriate element shape and use a free method to mesh the part:

2. Click **Assign Mesh Controls** in the mesh module toolbar or select **Mesh > Controls**.
3. In the viewport, hold the `<Shift>` key and select the main part of the manifold and the partitioned outlet section. When both parts are selected, click **Done**.
4. In the *Mesh Controls* dialog change the *Element Shape* to **Tet**. Keep the default options and click **OK**.



5. Press `Esc` to exit the mesh controls tool.

Choose the element type for the manifold part. In this case, 4-node linear tetrahedral elements are used:

6. Click **Assign Element Type** or select **Mesh > Element Type**.
7. In the same way as described previously, select both parts of the manifold in the viewport and click **Done**.
8. In the *Element Type* dialog, set the following properties:

Property	Setting
Element Library	**Standard** (default)
Geometric Order	**Linear**
Family	**Heat Transfer**

The *Tet* tab is selected by default. The element selection `DC3D4` is displayed at the bottom of the dialog.

9. Click **OK**.



To determine the size of the mesh, specify a target global element size:

10. Click **Seed Part** or select **Seed > Part**.
11. In the *Global Seeds* dialog, set **Approximate global size** to `0.005`, then click **OK**.

    Seeds appear on each of the model edges.
12. Click **Done** at the bottom of the viewport to confirm the assignment of these seeds to the part.



13. Generate the mesh by clicking **Mesh Part** or by selecting **Mesh** > **Part**. Click **Yes** at the bottom of the viewport.



14. Save the model database by selecting **File** > **Save As**.
15. In the dialog that appears, navigate to the location to save the model, enter the name `manifold`, and click **OK**.

Creating the Abaqus Input File

The model preparation is now complete, and you can proceed to writing the input file.

Create an analysis job and the input file:

1. Double-click **Jobs** underneath the **Analysis** container.
2. In the *Create Job* dialog, set *Name* to `manifold`, then click **Continue...**
3. In the *Edit Job* dialog, set **Description** to `Exhaust manifold heat transfer`, then click **OK**.

   The job is added under the **Jobs** container.
4. Right-click the **manifold** job and select **Write Input**.



For the purposes of this tutorial, output history is not required. To generate the input file without requesting history output:

5. In the *warning* dialog, click **Yes**.

   The file `manifold.inp` is created in the working directory.
6. Click **Save** to save the model database.

Preparing the STAR-CCM+ Model

Set up the model in STAR-CCM+ for the fluid inside the manifold.

A surface mesh for the fluid part of the model has been prepared for this tutorial.

1. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data` and copy the file containing the fluid mesh, `manifoldFluid.dbs`, into your `/thermalCoSimTutorial/starccm+/` directory.
2. Launch STAR-CCM+.
3. Start a new simulation.
4. Select **File** > **Import** > **Import Surface Mesh**.
5. In the *Open* dialog, navigate to `/thermalCoSimTutorial/starccm+/`.
6. Select `manifoldFluid.dbs`, then click **Open**.
7. In the *Import Surface Options* dialog, set *Import Mode* to **Create New Region**.
8. Click **OK**.
9. Save the simulation as `manifold-gas-cosim.sim`.

Contents

Visualizing the Imported Geometry
Setting Up the Mesh Continuum
Setting the Global Mesh Reference Values
Setting the Mesh Extrusion Parameters
Generating the Volume Mesh
Selecting Physics Models
Defining Boundary Conditions
Setting Up a Monitor
Setting Solver Variables
Setting Stopping Criteria
Setting Up a Plane Section
Creating a Fluid Velocity Scalar Scene
Creating a Fluid Temperature Scalar Scene

Visualizing the Imported Geometry

Examine the imported geometry to make sure that each boundary has been imported correctly.

The model currently has the default region name, change this to a more suitable name:

1. Rename the **Regions** > **Region 1** node to `Fluid`.

Check the location of each boundary by selecting the boundary node in the object tree and observing which part is highlighted in the geometry scene:

2. In the *Graphics* window, activate *Geometry Scene 1* and press the `<T>` key.
3. Expand the **Fluid** > **Boundaries** node and select the **Fluid-Wall** node.



4. Select the **Inlets** and **Outlet** nodes in turn and observe their location.

Setting Up the Mesh Continuum

Select the meshing models for the mesh continuum, including the mesh extruder model.

Due to the shape of the manifold, the gas flowing through the part is highly turbulent at the outlet. As the turbulence could limit the likelihood of the solution converging, use the mesh extruder feature in STAR-CCM+ to add a length of pipe at the manifold outlet. The pipe

aids good flow characteristics and helps the solution to converge. Although it is unlikely that such a length of pipe represents an exhaust pipe in reality, it is appropriate where you are concerned with temperature concentrations in the solid region.

Use the polyhedral mesher and prism layer mesher to create a volume mesh. Use the surface remesher tool to specify a different mesh size.

1. For the mesh continuum, **Continua** > **Mesh 1**, select the following meshing models:

Group Box	Meshing Models
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**
	Extruder

2. Click **Close**.

Setting the Global Mesh Reference Values

Define the global reference values for the volume mesh.

1. Edit the **Continua** > **Mesh 1** > **Reference Values** node, and set the following properties:

Node	Property	Value
Base Size	*Value*	`0.0035 m`
Number of Prism Layers	*Number of Prism Layers*	`4`
Prism Layer Thickness	*Size type*	**Absolute**
▮▮ **Absolute Size**	*Value*	`0.0035 m`
Surface Size > **Relative Minimum Size**	*Percentage of Base*	`100`

Setting the Mesh Extrusion Parameters

Set up the mesh extruder to add a straight section of pipe at the manifold outlet.

The extrusion extends the fluid domain past the edge of the solid manifold. For co-simulation cases, the interfacing surfaces, must be the same size and be located in the same place. Therefore, make sure that the surfaces the mesh extruder creates are not added on to the Fluid-Wall boundary, which is the default behavior. STAR-CCM+ maintains continuity between the main fluid domain and the extrusion by automatically creating an interface between the two regions. Since the outlet boundary must be at the end of the extruded section, it becomes part of the new region.

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Outlet** node and set the following properties:

Node	Property	Setting
Mesh Conditions > Extrusion Type	*Extrusion Type*	**Constant rate normal**
Mesh Values > Normal Extrusion Parameters	*Magnitude*	`0.25 m`
	Number of Layers	`40`
	Stretching	`5.0`
	Specify a New Region	**Yes, with Extended Type Boundaries**

2. Save the simulation.

Generating the Volume Mesh

The mesh setup is complete, and you can now run the polyhedral volume mesher.

1. Click (**Generate Volume Mesh**).
   The output from the meshers is shown in the *Output* window and the mesh generation status is shown in the status bar.
2. When the mesh generation is complete, save the simulation.

To visualize the volume mesh:

3. Create a mesh scene.
4. In the *Graphics* window, activate *Mesh Scene 1* and press the `<T>` key.



The outlet boundary has been moved into the extrusion region, **Extrusion from Outlet**. Additional boundaries have also been created automatically to define the two sides of the interface between the main fluid region and the extrusion region.

Selecting Physics Models

A physics continuum was created automatically at the end of the mesh generation process. Select physics models for this continuum and make sure that it is assigned to the Fluid region.

1. For the physics continuum, **Continua** > **Physics 1**, select the following physics models in order:

Group box	Model
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Ideal Gas**
Energy	**Segregated Fluid Enthalpy**
Viscous Regime	**Turbulent**

Group box	Model
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Optional Models	**Cell Quality Remediation**
Optional Models	**Co-Simulation**
Co-Simulation	**Abaqus Co-Simulation**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



Defining Boundary Conditions

Specify the fluid temperature, fluid velocity, and the turbulence settings at the manifold inlets.

The gas entering the manifold is traveling at 10 m/s with a temperature of 500 K. Specify these boundary conditions under the inlet boundary node, as well as selecting a different turbulence specification.

To set the boundary conditions:

1. Edit the **Regions** > **Fluid** > **Boundaries** > **Inlets** node.
2. Set the following properties:

Node	Property	Setting
Physics Conditions		
Turbulence Specification	*Method*	**Intensity + Length Scale**
Physics Values		
Static Temperature > **Constant**	*Value*	`500 K`
Velocity Magnitude > **Constant**	*Value*	`10 m/s`

3. Save the simulation.

Setting Up a Monitor

By default STAR-CCM+ produces a residuals plot which can be used to monitor the convergence of the solution. It is useful to create other monitors to help in assessing convergence of the solution. Monitor the temperature at a point inside the fluid domain, close to the solid-fluid interface.

1. Right-click the **Derived Parts** node and create a new point probe with the following settings:

Property	Setting
Input Parts	**Fluid**
	Extrusion from Outlet
Point	`[0.251, 0.107, -0.0405] m`
Display	**No Displayer**

2. Click **Create** then **Close**.

Use this point to monitor the temperature in the fluid. Before you can create the monitor, first create a report. In this case use a Maximum report:

3. Right-click the **Reports** node and select **New Report > Maximum**.
4. Rename the **Reports > Maximum 1** node to `Temperature Point`.

Reports can be created for all derived parts and boundaries. In this case, select the point:

5. Edit the **Reports > Temperature Point** node and set the following properties:

Property	Setting
Scalar Field Function	**Temperature**
Parts	**Derived Parts > Point**

   The report settings are now complete and you can proceed by creating a monitor and plot from the report.
6. Right-click the **Temperature Point** node and select **Create Monitor and Plot from Report**.

   The monitor and plot are named **Temperature Point Monitor** and **Temperature Point Monitor Plot**, respectively.
7. Open the **Temperature Point Monitor Plot**.

Setting Solver Variables

Set the solver parameters to specify the simulation time-step and to display detailed output about data exchanges with Abaqus.

Response times to changes in the temperature in the solid and fluid domains are relatively slow, so specify a large time-step. In a co-simulation case, data is exchanged between STAR-CCM+ and Abaqus at each time-step, hence a time-step is also referred to as a coupling step.

1. Edit the **Solvers** node and set the following properties:

Node	Property	Setting
Implicit Unsteady	*Time-Step*	`250 s`
Abaqus Co-Simulation	*Verbosity*	**High**

2. Save the simulation.

Setting Stopping Criteria

Set the maximum physical time and create an additional stopping criterion from the temperature point monitor.

In this tutorial, the co-simulation converges within 40 time steps, so use the maximum physical time to impose this limit:

1. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `10000 s`.

Note: Note: the maximum physical time does not need to match the Time Period setting in Abaqus. During the co-simulation, you can terminate the run using either application, depending on your specified stopping criteria.

In the early stages of the simulation, you can expect the number of inner iterations STAR-CCM+ requires to be much higher, since STAR-CCM+ is effectively solving the initial steady-state solution in the fluid. However, as the simulation progresses, the number of inner iterations per time step reduces, as the solution converges much more quickly. To accelerate the progress of the co-simulation, either decrease the number of inner iterations manually (by observing the temperature monitor and adjusting the **Maximum Inner Iterations** stopping criterion) or use an additional stopping criterion. The latter technique is described.

Create a stopping criterion from a temperature monitor such that when the temperature solution has converged for each time step, STAR-CCM+ stops iterating. Increase the number of iterations for the Maximum Inner Iterations stopping criterion so that enough iterations are carried out in the initial steps:

2. Copy the **Monitors** > **Temperature Point Monitor** node and paste it onto the **Monitors** node.
3. Rename **Temperature Point Monitor 3** to `Temperature Point Monitor_Iteration`.
4. Select the **Temperature Point Monitor_Iteration** node and set *Trigger* to **Iteration**.
5. Right-click the **Temperature Point Monitor_Iteration** node and select **Create Stopping Criterion from Monitor**.

   A new node named **Temperature Point Monitor_Iteration Criterion** is created under the **Stopping Criteria** node.

The temperature solution can be said to have converged when the temperature variations per iteration have dropped below 0.05 K. To specify an asymptotic stopping criterion:

6. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`60`
Temperature Point Monitor_Iteration Criterion	*Criterion Option*	**Asymptotic**
Asymptotic Limit	|*Max-Min*|	`0.05 K`

7. Save the simulation.

Setting Up a Plane Section

Create a plane section through the fluid region.

During the process of defining the plane section, a preview is shown in the *Graphics* window if a scene (rather than a plot) is currently on display. Begin by selecting a scene for this preview.

1. Activate *Mesh Scene 1* by clicking the corresponding tab at the top of the *Graphics* window.
2. Right-click the **Derived Parts** node and create a new plane section with the following settings:

Property	Setting
Input Parts	**Fluid**
	Extrusion from Outlet
origin	`[0, 0, 0] m`
normal	`[0, 0, 1] m`
Display	**No Displayer**

3. Click **Create** then **Close**.
   The plane section derived part has been created.

To highlight this part in the active scene:

4. Expand the **Derived Parts** node and select **plane section**.



Creating a Fluid Velocity Scalar Scene

Use a scalar scene to display the velocity magnitude on the plane section.

1. Create a scalar scene.
2. Rename the **Scenes** > **Scalar Scene 1** node to `Fluid Velocity`.
3. Edit the **Scenes** > **Fluid Velocity** > **Displayers** > **Outline 1** > **Parts** node.
4. In the *selection* dialog, deselect **Regions**, then select **Derived Parts** > **plane section**.
5. Click **OK**.
6. Edit the **Scenes** > **Fluid Velocity** > **Displayers** node.
7. Set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Derived Parts** > **plane section**

Node	Property	Setting
Scalar Field	*Function*	**Velocity > Magnitude**

8. Click (**Save-Restore-Select views**) and select **Views > -Z > Up +Y**. Use the mouse to zoom in on the section and position it in the center of the scene.

To display the solution time in the scalar scene, add an annotation:

9. Expand the **Tools > Annotations** node. Select the **Solution Time** node and drag it into the scalar scene. Release the mouse button to add the annotation to the scene.



10. Click the annotation and drag it into a position as shown below.



Creating a Fluid Temperature Scalar Scene

Create a second scalar scene to display the fluid temperature on the plane section.

Since many of the settings are the same as for the Fluid Velocity scene, use a duplicate of this scene as the starting point.

1. Copy the **Scenes** > **Fluid Velocity** node and paste it onto the **Scenes** node.
2. Rename the **Copy of Fluid Velocity** node to `Fluid Temperature`.
3. Select the **Scenes** > **Fluid Temperature** > **Displayers** > **Scalar 1** > **Scalar Field** node and set *Function* to **Temperature**.

   The scene appears as shown below.



4. Save the simulation.

Setting Up and Running the Co-Simulation

Preparation of the solid and fluid models is complete, and you can now set up and run the co-simulation.

The final step in setting up the solid model in Abaqus/CAE was to create the input file, `manifold.inp`. This text-based file can be edited using a script editor.

If you skipped the section involving Abaqus/CAE, a copy of the `manifold.inp` file has been provided with the STAR-CCM+ installation. Before proceeding, navigate to the `doc/`

`startutorialsdata/coupling/data` subdirectory of the STAR-CCM+ installation directory and copy the file `manifold.inp` into your `/thermalCoSimTutorial/abaqus/` directory.

Note: For keyword lines containing many parameters, you can split the line at a comma. If the last character in a line is a comma, the next line is interpreted as a continuation of the line.

To prepare the Abaqus model for co-simulation, a set of keywords and commands must be added to the input file to:

- Identify the Abaqus co-simulation step
- Identify the interfacing surface on the Abaqus model and the fields that are imported and exported to/from this surface

Contents

Setting Up the Co-Simulation Session
Running the Simulation
Visualizing the Results

Identifying the Abaqus Co-Simulation Step

The co-simulation is carried out during the heat transfer analysis step that you defined earlier. Add the co-simulation commands within the Step definition of the input file.

1. Open the `manifold.inp` file in a script editor of your choice.
2. Scroll down to the section that begins with the following:

```
**
** STEP: Step-1
**
```

   This section of the file, up to `*End Step`, contains the history data for the Abaqus job, including solver parameters, boundary conditions, loads, and outputs.

3. To identify this step as the Abaqus co-simulation step, add the following lines that are shown in bold text:

```
*Output, field, variable=PRESELECT
*Output, history, frequency=0
**
*CO-SIMULATION, NAME=FSI_Thermal,
PROGRAM=MULTIPHYSICS
**
*End Step
```

This line specifies that co-simulation is carried out in this step, and defines the co-simulation name, `FSI_Thermal`.

Identifying the Interfacing Surface and Import/Export Fields

1. To identify the interfacing surface and the fields that are imported and exported during the co-simulation, add the following lines:

```
*Output, field, variable=PRESELECT
*Output, history, frequency=0
**
*CO-SIMULATION, NAME=FSI_Thermal,
PROGRAM=MULTIPHYSICS
*CO-SIMULATION REGION, TYPE=SURFACE, EXPORT
ASSEMBLY_INNER, NT
*CO-SIMULATION REGION, TYPE=SURFACE, IMPORT
ASSEMBLY_INNER, CFL
**
*End Step
```

   The first line added here indicates that you are identifying the region in the model, in this case a surface, from which data are exported. In the next line, the surface is identified by entering the name of the assembly in which it is defined, `Assembly`, and then the name of the surface, `Inner`. The syntax of the surface identifier is `<ASSEMBLY NAME>_<SURFACE NAME>`. The next line defines the export field as the nodal temperature, `NT`. The next two lines use the same format to specify that heat flux, `CFL`, are imported into Abaqus.

2. Save the file as `manifold-mod.inp`.

Setting Up the Co-Simulation Session

Set up the co-simulation in STAR-CCM+.

In STAR-CCM+, the properties of the co-simulation are defined within a co-simulation session. This includes details of the location of the Abaqus input file, the boundary defining the fluid-structure interface, the type of data that is exchanged at each coupling step, and a number of coupling controls.

A co-simulation session, **Abaqus Co-Simulation 1**, is automatically added under the **Co-Simulations** node when the Abaqus Co-Simulation physics model is activated. Co-simulation settings are specified using this node.

1. Edit the **Co-Simulations > Abaqus Co-Simulation 1 > Values** node, and set the following properties:

Node	Property	Setting
Abaqus Execution	*Current Job Name*	`co-simulation_manifold`
	Input File	`../abaqus/manifold-mod.inp`
	Executable Name	name you use to launch Abaqus, for example `abaqus`
Coupling Algorithm	*Constant Coupling Time Step*	`250 s`

To specify the location of the input file, you type the path directly into the box. For example, if your working directory is `/thermalCoSimTutorial/starccm+/` and your input file is contained in the folder `/thermalCoSimTutorial/abaqus/`, the path to the input file would be `../abaqus/manifold-mod.inp`.



Define which field functions are imported and exported during the co-simulation. As this tutorial is purely a heat transfer case, you only deal with scalar quantities. STAR-CCM+ exports boundary heat flux data to Abaqus, and imports temperature data. Also set the coupled boundaries. In this case, the boundary at which the coupling takes place is the fluid wall boundary.

2. Edit the **Abaqus Co-Simulation 1** > **Zones** > **Zone 1** node, and set the following properties:

Node	Property	Setting
Values		
Fields Exported To Abaqus	*Exported Field Functions*	**Boundary Heat Flux**
Fields Imported From Abaqus	*Imported Field Functions*	**Temperature**

To specify the fluid wall as the coupled boundary:

3. Expand the **Regions > Fluid > Bundaries > Fluid-Wall** node.
4. Edit the **Physics Conditions > External Code Coupling Specification** node and set *External Code Coupling Specification* to **Zone 1**.

   The *Boundaries* property for the **Co-Simulations > Abaqus Co-Simulation 1 > Zones > Zone 1 > Conditions > Coupled Boundaries** node is automatically set to **Fluid-Wall**.

   The thermal specification for the Fluid-Wall boundary is automatically set to a temperature field function condition.
5. Save the simulation.

Running the Simulation

You are now ready to run the co-simulation.

During this process, STAR-CCM+ automatically launches Abaqus in the background and starts communicating with Abaqus via the SIMULIA Co-Simulation Engine.

1. Run the simulation.

   The co-simulation libraries are loaded and the co-simulation settings registered. The corresponding messages are displayed in the *Output* window.

   The iteration progress can be monitored in the *Output* window. Monitor the convergence by viewing the Temperature Point Plot and the Residuals in the *Graphics* window. You can also view the *Fluid Velocity* and *Fluid Temperature* scenes to visualize the solution while the co-simulation is running.

   At the end of each coupling step, data is exchanged between the two codes, and details are displayed in the *Output* window.

2. When the co-simulation is complete, save the simulation.

Visualizing the Results

Visualize the fluid velocity and temperature.

The results for the fluid solution can be observed by examining the solution that is displayed on the plane section in the *Fluid Velocity* and *Fluid Temperature* scalar scenes. To display each scene in the Graphics window, click the relevant tab at the top of the window.

You can see from the Temperature Point Plot that the fluid-solid solution has converged. Each time data was exchanged via the SIMULIA Co-Simulation Engine, that is, at each coupling step, the temperature at the point "jumped" in response to the imported temperature data. You can observe that as the temperature difference at the fluid-solid interface tended towards zero, the temperature at the monitor point close to the fluid wall converged.



The results for the solid solution can be post-processed by opening the file `/thermalCoSimTutorial/abaqus/co-simulation_manifold.odb` in Abaqus/CAE.

Summary

This tutorial demonstrated how to run a thermal co-simulation case using STAR-CCM+ and Abaqus.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up a solid model in Abaqus, including:
  - Importing geometry from a STEP file.
  - Defining model properties.
  - Defining surfaces.
  - Specifying initial and boundary conditions.
  - Generating a mesh.
  - Creating the input file.
- Setting up a fluid model in STAR-CCM+, including:
  - Importing a surface mesh.
  - Selecting physics models (including the co-simulation model).
  - Setting up an additional monitor and defining a stopping criterion from this monitor.
  - Specifying the time-step size.
  - Creating scalar scenes.
- Running a co-simulation, including:
  - How to edit the Abaqus input file for co-simulation.
  - Setting up a co-simulation session in STAR-CCM+.
  - Running the co-simulation.

Abaqus Co-Simulation: Mechanical Coupling

This tutorial demonstrates how to carry out a mechanical co-simulation case using STAR-CCM+ and Abaqus. Co-simulation involves a strong coupling between the two codes. Data is exchanged at frequent intervals that are called coupling steps. This level of communication between the solvers allows you to obtain a full solution across the fluid-solid interface.

In a mechanical coupling, STAR-CCM+ passes traction loads to Abaqus (pressure + wall shear stress), and Abaqus passes displacements to STAR-CCM+. In Abaqus, the traction loads are applied to the surface of the solid structure. In STAR-CCM+, the displacements are used as an input to the mesh morpher.

During the co-simulation, data is automatically exchanged between STAR-CCM+ and Abaqus using the SIMULIA Co-Simulation Engine. This technique differs from the file-based method in that the coupling between the codes is much stronger, and STAR-CCM+ and Abaqus run simultaneously.

To set up a co-simulation case, separate models must be prepared for the fluid and the solid domains. The solid model is set up in Abaqus, and defined within an Abaqus input file. The fluid model is set up in a STAR-CCM+ simulation where the appropriate boundary conditions and physics are defined. For each model, a set of co-simulation parameters must be specified. The parameters identify the boundaries/surfaces that make up the fluid-structure interface, the field data that is imported/exported at the interface, the frequency of coupling steps, and so on. STAR-CCM+ is used to run the co-simulation when the setup is complete.

Before running the co-simulation, it is good practice to run standalone cases for the fluid and solid models to make sure that they converge independently. In this example, the Abaqus model could be run standalone using an applied pressure load, *DSLOAD, on the plate surface with an estimated constant value for the pressure. The STAR-CCM+ simulation could be run without mesh morphing, using a fixed plate.

This tutorial simulates an elastic flap that is pinned at the base, and then subjected to a 10 m/s flow of air. The mass and stiffness of the plate are such that the first bending mode has a frequency near 4 Hz and the first twisting mode has a frequency near 20 Hz. With a fixed plate, the vortex

shedding at a 10 m/s inlet velocity is near 20 Hz, so you expect resonance between the vortex shedding and the first twisting mode.



In this tutorial, the fluid domain is initialized by running a quasi-steady simulation. The initial flow field is then used as the initial condition for the co-simulation. Once the initial flow field is established, the solver settings are changed to the values that are required for the co-simulation. When the co-simulation is run, the plate is released and bends in the wind. In the tutorial, the co-simulation is run for 0.125s, which is about half the fundamental period of the first bending mode and about 2.5 vortex shedding periods. Running the case for longer (around 4 s) shows that the first bending mode is damped, and the first twisting mode dominates plate motion.

Contents

Prerequisites
Obtaining Input Files
Reviewing the Abaqus Model
Preparing the STAR-CCM+ Model
Setting Up the Co-Simulation
Obtaining an Initial Flow Field
Running the Co-Simulation
Summary

Prerequisites

The instructions in the Abaqus Mechanical Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Also, make sure that you have access to an Abaqus installation and that your system has been set up for co-simulation.

Obtaining Input Files

Copy the files that are required to complete the tutorial into your working directory.

When running a co-simulation case, it is good practice to create separate directories for the STAR-CCM+ and Abaqus files:

1. Create a directory called `/mechanicalCoSimTutorial`.
2. Within this directory, create two more directories: `/abaqus` and `/starccm+`.

Two files have been prepared for this tutorial: an Abaqus input file, and a fluid surface mesh. Copy these files into the `/abaqus` and `/starccm+` directories:

3. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
4. Copy the file `plate.inp` into your `/mechanicalCoSimTutorial/abaqus/` directory.
5. Copy the file `plateFluid.dbs` into your `/mechanicalCoSimTutorial/starccm+/` directory.

Reviewing the Abaqus Model

Review the features of the Abaqus model.

The Abaqus model for this tutorial is a flat plate that is pinned at the base. The mesh and plate dimensions are shown in the diagram below.



The plate is made from a material with a Young's modulus of 38.4 MPa, a Poisson's ratio of 0.3, and a density of 4096 kg/m^3. Rayleigh damping is also introduced with a beta value of 1e-6. For the purposes of this tutorial, the Young's Modulus of the plate material has been specified as an elastic material to make sure that the plate is compliant in response to the air flow.

The plate is meshed using C3D8R elements, and the size of the grid is 16x20x5 (width x height x thickness).

Loading is defined using a single dynamic load step. The maximum time-step is specified as 5.0e-4 s and the total time is 0.125 s.

Preparing the STAR-CCM+ Model

Prepare the fluid model simulation in STAR-CCM+.

The STAR-CCM+ model defines a fluid domain surrounding the solid plate. In the simulation, air enters the fluid domain at 10 m/s. The outlet boundary is modeled as a Pressure Outlet, and the top and sides of the

domain are specified as Symmetry boundaries. Flow is defined as incompressible, and turbulence is modeled using the K-Omega model.



The dimensions of the fluid domain are shown in the above diagram.

- Importing the Fluid Surface Mesh
- Setting Up the Mesh Continuum
- Generating the Volume Mesh
- Selecting Physics Models
- Defining Boundary Conditions
- Setting Up the Morpher
- Setting Up Monitors
- Creating a Plane Section
- Preparing a Scalar Scene
- Setting Up Automatic Scene Export

Contents

Importing the Fluid Surface Mesh
Setting Up the Mesh Continuum
Generating the Volume Mesh
Selecting Physics Models
Defining Boundary Conditions
Setting Up the Morpher
Setting Up Monitors
Creating a Plane Section

Preparing a Scalar Scene
Setting Up Automatic Scene Export

Importing the Fluid Surface Mesh

Launch STAR-CCM+ for co-simulation, and import the supplied fluid surface mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Surface Mesh...**
4. In the *Open* dialog, navigate to `/mechanicalCoSimTutorial/starccm+/`.
5. Select `plateFluid.dbs`, then click **Open**.
6. In the *Import Surface Options* dialog:
   a) Set **Import Mode** to **Create new Region**.
   b) Click **OK**.

   The import process takes a few seconds, and STAR-CCM+ provides feedback on the process in the **Output** window. The imported geometry is displayed in a geometry scene in the **Graphics** window.
7. Click (**Make Scene Transparent**).
8. In **Geometry Scene 1**, use the mouse to rotate the view and examine the geometry.
9. Save the simulation as `plate-cosim.sim`.

Setting Up the Mesh Continuum

A mesh continuum was created automatically when the geometry was imported. Use this continuum to set up the mesh.

1. For the mesh continuum, **Continua > Mesh 1**, select the following meshing models:

Group box	Meshing Models
Surface Mesh	**Surface Remesher**
Volume Mesh	**Polyhedral Mesher**
Optional Models	**Prism Layer Mesher**

2. Click **Close**.
3. Select the **Mesh 1 > Reference Values > Base Size** node and set *Value* to `0.01 m`.

To refine the mesh around the plate:

4. Edit the **Regions** > **Region 1** > **Boundaries** > **Plate_FSI** node and set the following properties:

Node	Property	Setting
Mesh Conditions		
Custom Surface Size	*Custom surface size*	**Activated**
Mesh Values		
Surface Size > **Relative Target Size**	*Percentage of Base*	50

Generating the Volume Mesh

Generate the volume mesh for this simulation.

1. Click (**Generate Volume Mesh**).
   Output from the meshing process is shown in the *Output* window.
2. When the mesh generation process is complete, save the simulation.

To visualize the volume mesh using a mesh scene:

3. Create a new mesh scene.
   The volume mesh is displayed in **Mesh Scene 1**.
4. Remove some of the boundaries from the scene to see the plate inside the fluid domain:
   a) Edit the **Scenes** > **Mesh Scene 1** > **Displayers** > **Mesh 1** > **Parts** node.
   b) In the *Edit* dialog, deselect the **Top** and **Front** boundaries, then click **OK**.
5. Use the mouse to rotate the view and examine the mesh.
   The mesh appears as shown below.



Selecting Physics Models

Following mesh generation, STAR-CCM+ automatically created a physics continuum. Use this continuum to define the physics models for the fluid domain.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically)
	Gradients (Selected automatically)
Time	**Implicit Unsteady**
Material	**Gas**
Flow	**Segregated Flow**
Equation of State	**Constant Density**
Viscous Regime	**Turbulent**

Group Box	Model
Reynolds-Averaged Turbulence	**K-Omega Turbulence**
Optional Models	**Co-Simulation**
Co-Simulation	**Abaqus Co-Simulation**
Optional Models	**Cell Quality Remediation**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



Initialize the air velocity in the domain at 10 m/s perpendicular to the plate:

4. Select the **Continua** > **Physics 1** > **Initial Conditions** > **Velocity** > **Constant** node and set *Value* to `[10, 0, 0] m/s`.

Defining Boundary Conditions

Define the boundary type and specify the fluid velocity at the inlet.

1. Edit the **Regions** > **Region 1** > **Boundaries** node and set the boundary *Type*:

Boundary	Type
Outlet	**Pressure Outlet**
Back	**Symmetry Plane**
Front	**Symmetry Plane**
Top	**Symmetry Plane**

2. Select the **Inlet** > **Physics Values** > **Velocity Magnitude** > **Constant** node and set *Value* to `10 m/s`.

Setting Up the Morpher

Use the Morpher to allow the fluid mesh to deform in response to the imported nodal displacements from Abaqus.

The first step in setting up the morpher is to create a morphing motion and enable this motion in the fluid region:

1. Right-click the **Tools** > **Motions** node and select **New** > **Morphing**.

   A new **Morphing** node is added under the **Motions** node. By default, the morpher is set to automatically remove vertices on boundaries while the simulation runs. The *Automatic Thin-out Cl* factor determines the number of control vertices that are used in each iteration. As the level of deformation increases, the morpher automatically uses more control vertices.

Decrease the *Automatic Thin-out Cl* factor to increase the sensitivity of the morpher to the mesh deformation:

2. Select the **Motions** > **Morphing** node and set *Automatic Thin-out Cl* to `0.5`.

To apply this motion to the fluid region:

3. Select the **Regions** > **Region 1** > **Physics Values** > **Motion Specification** node and set *Motion* to **Morphing**.

   The morpher specification for the FSI boundary will be automatically set to **Abaqus Co-Simulation**, once that the co-simulation session is set up. This setting allows the FSI boundary to move according to the imported displacements.

Change the morpher method for the inlet and outlet boundaries from **Fixed** to **Fixed Plane**. This morpher method is suitable for the selected boundaries and is more computationally efficient:

4. Expand the **Boundaries** node.
5. Multi-select the **Inlet** and **Outlet** nodes.
6. Right-click one of the selected boundaries and select **Edit**.
7. Select the **Physics Conditions** > **Morpher** node and set *Method* to **Fixed Plane**.

To prevent inaccuracies from accumulating during the simulation (due to interpolation), change the mesh morpher settings:

8. Select the **Solvers** > **Mesh Morpher** node and activate *Morph From Zero*.



Setting Up Monitors

Create reports to monitor the displacement of the upper corners of the plate in the x-direction, the traction load, and the moment about the z-axis, during the co-simulation.

To monitor the displacements of the upper-left and upper-right corners of the plate, two point-probes must be created that follow the motion of the corners as the plate moves in the flow:

1. Within the **Derived Parts** node, create a new point probe with the following settings:

Property	Setting
Input Parts	**Region 1** > **Plate FSI**
Point	`[0.0, 0.04, 0.1]`
Display	**No Displayer**

A node, **point**, is added beneath the **Derived Parts** node.

2. Create another point probe with the following settings:

Property	Setting
Input Parts	**Region 1 > Plate FSI**
Point	`(0.0, -0.04, 0.1)`
Display	**No Displayer**

A node, **point 2**, is added beneath the **Derived Parts** node.

3. Rename the **Derived Parts > point** node to `Upper +Y`.
4. Rename the **Derived Parts > point 2** node to `Upper -Y`.

To make the points move with the plate when the morpher moves the mesh:

5. Select the **Derived Parts > Upper +Y** node and activate *Follow Motion*.



6. Repeat the previous step for the **Upper -Y** node.

Having created point probes on the upper corners of the plate, you can now create reports that report the displacement of these corners:

7. Right-click the **Reports** node and select **New Report > Maximum**.
8. Rename the **Reports > Maximum 1** node to `X Displacement Upper +Y`.
9. Set the following properties:

Property	Setting
Scalar Field Function	**Nodal Displacement > Laboratory > i**
Parts	**Derived Parts > Upper +Y**

10. Create another **Maximum** report for the other corner.
11. Rename the report that you just created, **Reports > Maximum 1**, to `X Displacement Upper -Y`.
12. Set the following properties:

Property	Setting
Scalar Field Function	**Nodal Displacement > Laboratory > i**
Parts	**Derived Parts > Upper -Y**

Use the displacement reports to create monitors and a plot:

13. Multi-select the **X Displacement Upper +Y** and **X Displacement Upper -Y** reports.
14. Right-click one of the reports and select **Create Monitor and Plot from Report**.

    The *Create Plot From Reports* dialog appears.
15. Select **Single Plot**.

    The report monitors are displayed under the **Monitors** node, and the plot is shown under the **Plots** node.
16. Rename the **Plots > Reports Plot** node to `Displacement Plot`.

Create two more reports to monitor the traction load and moment on the plate:

17. Create a **Force** report.
18. Select the **Reports > Force 1** node and set the following properties:

Property	Setting
Parts	**Regions > Plate_FSI**
Force Option	**Pressure + Shear** (default)
Direction	`[1.0, 0.0, 0.0]` (default)

19. Right-click the **Reports** node and select **New Report > Moment**.
20. Select the **Reports > Moment 1** node and set the following properties:

Property	Setting
Parts	**Regions > Plate_FSI**
Axis	`[0.0, 0.0, 1.0]` (default)
Force Option	**Pressure + Shear** (default)

Create two separate plots for the force and moment reports:

21. Multi-select the **Force 1** and **Moment 1** reports.
22. Right-click and select **Create Monitor and Plot from Reports**.
23. In the dialog that appears, select **Multiple Plots (one per report)**.



Finally, verify the trigger for monitoring the data is set to time-step:

24. Expand the **Monitors** node and multi-select the four report monitors.
25. Verify that *Trigger* is set to **Time Step**.



26. Open all plots.
27. Save the simulation.

Creating a Plane Section

Create a plane section through the plate, to visualize the solution and mesh deformation.

1. Click the **Mesh Scene 1** tab in the *Graphics* window.
2. Within the **Derived Parts** node, create a new plane section with the following settings:

Property	Setting
origin	`[0.05, 0, 0.1] m`
normal	`[0, 1, 0] m`
Display	**Existing Displayer**

Modify the parts that are displayed in the scene to show the mesh on the plane section and the solid plate:

3. Edit the **Displayers** > **Mesh 1** > **Parts** node.

4. In the *Edit* dialog, deselect all the boundaries in **Region 1**, apart from **Plate_FSI**, then click **OK**.
5. Use the mouse to rotate the view so that it appears as shown below.



Preparing a Scalar Scene

Prepare a scalar scene to display the pressure field around the plate. Include the plane section in a separate scalar displayer so that the opacity setting can be adjusted independently.

1. Create a scalar scene.
2. Right-click the **Scalar Scene 1 > Displayers** node and select **New Displayer > Scalar**.



3. Edit the **Scalar Scene 1 > Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Regions > Region 1 > Bottom**
		Regions > Region 1 > Plate_FSI
Scalar Field	*Function*	**Pressure**
Scalar 2	*Opacity*	`0.5`
	Contour Style	**Smooth Filled**
Parts	*Parts*	**Derived Parts > plane section**
Scalar Field	*Function*	**Pressure**
Color Bar	*Visible*	**Deactivated**

4. Position the scene as shown below.



Add an annotation to the scene to show the current solution time:

5. Expand the **Tools** > **Annotations** node.
6. Click the **Solution Time** node and drag it into the scene. Release the mouse button anywhere in the scene.
7. Move the annotation to the top-right of the scene using click-and-drag.



Add a solution time annotation to the mesh scene:

8. Display the mesh scene in the **Graphics** window by selecting the **Mesh Scene 1** tab.
9. Using the same technique as described above, add a Solution Time annotation to the scene.
10. Save the simulation.

Note: In transient cases, it is often useful to export copies of scalar scenes while the simulation is running. These image files can then be used to visualize the solution in an animation. The next page in this tutorial runs through the steps for setting up automatic scene export. If you do not wish to export scenes automatically during the co-simulation, you can jump directly to Setting Up the Co-Simulation.

Setting Up Automatic Scene Export

In the following steps, set up automatic export of the mesh and pressure scenes.

The automatic export of scenes leads to the generation of a number of `.png` files on your hard drive while the simulation runs. These image files can then be used to create an animation.

1. Create an `/images` directory in your `/mechanicalCoSimTutorial/starccm+/` directory as a location to save the image files.

Many of the settings for the scenes are the same. You can apply the settings to both scenes simultaneously using the Multi-edit dialog.

2. Multi-select the **Scenes** > **Scalar Scene 1** and **Scenes** > **Mesh Scene 1** nodes.
3. Right-click one of the selected nodes and select **Edit**.
4. Set the following properties:

Node	Property	Setting
Mesh Scene 1, **Scalar Scene 1**		
	Save To File	**Activated**
Attributes > **Update**	*Output Directory*	`/mechanicalCoSimTutorial/starccm+/images`
	Trigger	**Time Step**
Time-Step Frequency	*Time Step Frequency*	`5`
Hardcopy	*Output Width*	`800`
	Output Height	`600`

A red * is shown next to the icon for each scene to indicate that copies are saved automatically when the co-simulation is run.

Finally, specify unique base filenames for STAR-CCM+ to use when saving the image files:

5. Select the **Scenes** > **Mesh Scene 1** > **Attributes** > **Update** node and set *Base Filename* to `mesh`.
6. Using the same technique, set the base filename for **Scalar Scene 1** to `pressure`.

Setting Up the Co-Simulation

Modify the Abaqus input file to set it up for co-simulation.

Before running the co-simulation, it is necessary to insert a set of keywords and commands to:

- Identify the Abaqus co-simulation step

- Identify the interfacing surface on the Abaqus model and the fields that are imported and exported to/from this surface

In the User Guide, see *STAR-CCM+ Properties and Equivalent Abaqus Keywords* for a full summary of properties and matching keywords.

The Abaqus input file for this tutorial, `plate.inp`, already includes the necessary commands for co-simulation. In the following steps, locate these commands within the text file, and review the settings.

Note: For keyword lines containing many parameters, you can split the line at a comma. If the last character in a line is a comma, the next line is interpreted as a continuation of the line.

Contents

Setting Up the STAR-CCM+ Co-Simulation Session

Identifying the Abaqus Co-Simulation Step

The co-simulation is carried out during the load step. Co-simulation commands have been added within the Step definition.

1. Using a script editor of your choice, open the file `plate.inp`, located in your `/mechanicalCoSimTutorial/abaqus` directory.
2. Scroll down to the section that begins with the following text:

```
**
** STEP: Step-1
**
```

This section of the file, up to `*End Step`, contains the history data for the Abaqus job, including solver parameters, boundary conditions, loads, and outputs.

3. Within the Step definition, locate the following line of text (shown in bold):

```
**
*Output, history, variable=PRESELECT
**
*CO-SIMULATION, NAME=Plate, PROGRAM=MULTIPHYSICS
```

This line specifies that the co-simulation takes place in this step, and defines an arbitrary name for the co-simulation, Plate.

Identifying the Interfacing Surface and Import/Export Fields

1. The next set of commands identifies the interfacing surface and the fields that are imported and exported during the co-simulation:

```
**
*Output, history, variable=PRESELECT
*CO-SIMULATION, NAME=Plate, PROGRAM=MULTIPHYSICS
*CO-SIMULATION REGION, TYPE=SURFACE, EXPORT
ASSEMBLY_FSI_INTERFACE, U
*CO-SIMULATION REGION, TYPE=SURFACE, IMPORT
ASSEMBLY_FSI_INTERFACE, CF
```

   The first line identifies a region in the model, in this case a surface, from which data is exported. In the next line, the surface is identified by entering the name of the assembly in which it is defined, `Assembly`, and then the name of the surface, `FSI_Interface`. The syntax of the surface identifier is `<ASSEMBLY NAME>_<SURFACE NAME>`. The export field is defined as the nodal displacement, `U`. The next two lines follow the same pattern to specify that traction vectors, `CF`, are imported into Abaqus.

Setting Up the STAR-CCM+ Co-Simulation Session

Specify co-simulation settings in STAR-CCM+.

In STAR-CCM+ the co-simulation includes details of the Abaqus input file location, the boundary defining the fluid-structure interface, the data that is exchanged at each coupling step, and a number of coupling controls.

A co-simulation session, **Abaqus Co-Simulation 1**, is automatically added within the **Co-Simulations** node when the Abaqus Co-Simulation physics model is activated. Co-simulation settings are specified using this node.

1. Select the **Co-Simulations > Abaqus Co-Simulation 1 > Values > Abaqus Execution** node and set the following properties:

Property	Setting
Current Job Name	`co-simulation_plate`
Executable Name	command you use to launch Abaqus, for example: `abaqus`
Input File	`../abaqus/plate.inp`

   To specify the location of the input file, you type the path directly into the box. For example, if your working directory is `/`

`mechanicalCoSimTutorial/starccm+/` and your input file is contained in the folder `/mechanicalCoSimTutorial/abaqus/`, the path to the input file would be `../abaqus/plate.inp`.

The complete list of properties of the **Abaqus Execution** node is shown below.



Select the coupled boundaries. Also, specify which field functions are imported and exported during the co-simulation. In this case, traction data (pressure and wall shear stress) is exported to Abaqus and nodal displacements are imported into STAR-CCM+.

To select the coupled boundaries:

2. Expand the **Regions** > **Region1** > **Boundaries** > **Plate_FSI** node.
3. Select the **Physics Conditions** > **External Code Coupling Specification** node and set *External Code Coupling Specification* to **Zone 1**.

   The **Plate_FSI** boundary is automatically added as coupled boundary under the **Co-Simulations** > **Abaqus Co-Simulation 1** > **Zones** > **Zone 1** > **Conditions** > **Coupled Boundaries** node.

To specify the exported and imported fields:

4. Edit the **Co-Simulations** > **Abaqus Co-Simulation 1** > **Zones** > **Zone 1** > **Values** node and set the following properties:

Node	Property	Setting
Fields Exported To Abaqus	*Exported Field Functions*	**Wall Shear Stress**
		Pressure

Node	Property	Setting
Fields Imported From Abaqus	*Imported Field Functions*	**Nodal Displacement**

When *Imported Field Functions* is set to **Nodal Displacement**, the morpher specification for the FSI boundary is automatically set to **Abaqus Co-Simulation**, so that the boundary can move according to the imported displacements.

5. Select the **Co-Simulations** > **Abaqus Co-Simulation 1** > **Values** > **Coupling Algorithm** node and set *Constant Coupling Time Step* to `5.0E-4`.



6. Save the simulation.

Obtaining an Initial Flow Field

Run the simulation in a quasi-steady state to obtain an initial flow field. The solution from this initialization step is used as the initial condition for the co-simulation.

To solve for a steady-state flow about a rigid plate, freeze the co-simulation and mesh morpher solvers. The number of inner iterations per time-step are increased to allow a flow field to develop within a single time-step. Then, run the simulation for one time-step. As you are only searching for a steady-state solution, you do not need to be time-accurate. Therefore, use a large time-step of 1.5 s and a 1st-order temporal discretization setting.

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time Step* to `1.5 s`.
2. Multi-select the **Solvers** > **Abaqus Co-Simulation** and **Solvers** > **Mesh Morpher** nodes.
3. Activate *Solver Frozen*.



4. Select the **Stopping Criteria** > **Maximum Inner Iterations** node and set *Maximum Inner Iterations* to `400`.

The quasi-steady simulation only runs for one step, so disable the maximum physical time and maximum steps stopping criteria. Before the co-simulation is run, re-enable the maximum physical time:

5. Multi-select the **Stopping Criteria** > **Maximum Physical Time** and **Stopping Criteria** > **Maximum Steps** nodes.
6. Deactivate *Enabled*.
7. From the menu, select **Solution** > **Step**.

When the initial simulation is run, the pressure field is not displayed in **Scalar Scene 1**. To display the solution:

8. Refresh **Scalar Scene 1**.
9. Save the simulation.
10. Also, save a copy of the simulation at this stage called `plate-cosim-initial.sim`.

The solution data is retained in the simulation, and provides the initial conditions for the co-simulation. Reset the time to zero by clearing the solution history. Clearing the solution fields would cause you to lose the initial solution that has been computed:

11. From the menu, select **Solution** > **Clear Solution**.
12. In the *Clear Solution* dialog, deactivate *Fields* and *Reset Mesh*, then click **OK**.



Running the Co-Simulation

Specify the solver settings, stopping criteria, and run the simulation.

The co-simulation settings are registered with the Co-Simulation Engine, and feedback is shown in the **Output** window.

To prepare for the co-simulation, specify solver settings and stopping criteria suitable for the unsteady simulation:

1. Select the **Solvers** > **Implicit Unsteady** node and set the following properties:

Property	Setting
Time Step	`5e-4 s`
Temporal Discretization	**2nd-order**

2. Multi-select the **Solvers** > **Abaqus Co-Simulation** and **Solvers** > **Mesh Morpher** nodes.
3. Deactivate *Solver Frozen*.
4. Select the **Solvers** > **Abaqus Co-Simulation** node and set *Verbosity* to **High**.
5. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`

Node	Property	Setting
Maximum Physical Time	*Maximum Physical Time*	`0.125 s`
	Enabled	**Activated**

You are now ready to run the co-simulation. During this process, Abaqus runs in the background, and STAR-CCM+ communicates with Abaqus using the SIMULIA Co-Simulation Engine:

6. Run the simulation.

   The co-simulation settings are registered with the Co-Simulation Engine, and feedback is shown in the *Output* window.

   While the co-simulation is running, you can monitor progress using the plots and scenes created earlier. At each coupling step solution data is exchanged between Abaqus and STAR-CCM+, and is reported in the **Output** window.
7. When the co-simulation is complete, save the simulation.

Contents

Visualizing the Results

Visualizing the Results

Examine the mesh deformation and the Displacement Plot.

In *Mesh Scene 1*, you can see that the oncoming air flow has caused the plate to bend backwards. The morpher has deformed the mesh throughout the fluid region.



In the *Displacement Plot*, you can see that after 0.125s the plate has reached its maximum displacement in the x-direction, around 0.042m.



If this case is run for 4 s, the displacement plot shows that the first bending mode, where both corners of the plate would be in phase, is damped out after eight first bending mode periods or 4 Hz. The remaining mode, the vortex shedding excites, is the first twisting mode at a frequency of about 20 Hz. In this mode, the corners of the plate are out of phase.



Summary

This tutorial demonstrated how to carry out a mechanical co-simulation case using STAR-CCM+ and Abaqus.

This tutorial has introduced the following STAR-CCM+ features:

- Setting up the STAR-CCM+ model:
  - Setting up a physics continuum for co-simulation.
  - Defining boundary conditions.
  - Setting up the morpher.
  - Setting up monitors and plots.
  - Setting up automatic scene export.
- Setting up the co-simulation:
  - Editing the Abaqus input file for co-simulation.
  - Creating the co-simulation session in STAR-CCM+.
- Obtaining an initial flow field.

- Running the co-simulation.

STAR-CCM+ to STAR-CCM+ Coupling: Heat Transfer in a Chimney

This tutorial demonstrates the process for modeling conjugate heat transfer using STAR-CCM+ to STAR-CCM+ thermal co-simulation.

Using STAR-CCM+ to STAR-CCM+ co-simulation to solve multi-physics problems has the advantages that:

- Regions can be developed independently.
- Each case can be modeled with a relevant time-scale.

The diagram shows the case to be modeled. Hot exhaust gas enters a metal chimney with a temperature of 600 K and an average velocity of 0.2 m/s. The chimney has internal diameter of 0.4 m, a wall thickness of 10 mm and is made of T316 austenitic stainless steel. The external environment is assumed to be a perfect heat sink at a temperature of 283 K.



The tutorial is divided into three sections:

1. In the first section, you create the solid model of the chimney.
2. In the second section, you prepare the fluid model.
3. In the third section, you couple both simulations to model the complete fluid-solid system.

Contents

Prerequisites
Preparing the Solid Model Simulation
Preparing the Fluid Model Simulation
Coupling the Simulations
Summary

Prerequisites

The instructions in the STAR-CCM+ to STAR-CCM+ Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Preparing the Solid Model Simulation

Prepare the simulation for the solid chimney and run it to convergence.

Follow these steps to prepare the solid simulation:

1. Importing the Solid Volume Mesh.
2. Selecting Solid Physics Models.
3. Defining the Solid Material.
4. Defining the Solid Boundary and Initial Conditions.
5. Preparing a Temperature Scene for the Solid.
6. Setting Stopping Criteria for the Solid Simulation.
7. Running the Solid Simulation.
8. Visualizing the Decoupled Solid Results.

Contents

Importing the Solid Volume Mesh
Selecting Solid Physics Models
Defining the Solid Material
Defining the Solid Boundary and Initial Conditions
Preparing a Temperature Scene for the Solid
Setting Stopping Criteria for the Solid Simulation
Running the Solid Simulation
Visualizing the Decoupled Solid Results

Importing the Solid Volume Mesh

To set up the STAR-CCM+ simulation for the solid, launch a simulation and import the supplied volume mesh representing the chimney.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data/`.
5. Select `chimney.ccm`, then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `Chimney.sim`.

Selecting Solid Physics Models

Select physics models to solve for the solid region.

To simulate heat conduction through a solid, choose an energy transport model and an equation of state.

1. Rename the **Continua > Physics 1** node to `Solid`.
2. For the physics continuum, **Continua > Solid**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically)
Time	**Implicit Unsteady**
Material	**Solid**
Optional Models	**Segregated Solid Energy**
	Gradients (Selected automatically)
Equation of State	**Constant Density**

3. Click **Close**.
4. To review the models, open the **Solid > Models** node.



5. Save the simulation.

Defining the Solid Material

Set the solid material to stainless steel.

To set the solid material:

1. Right-click the **Continua** > **Solid** > **Models** > **Solid** > **Al** node and select **Replace with...**
2. In the *Replace Material* dialog, expand **Material Databases** > **Standard** > **Solids** and select *UNSS31600 (Austenitic Stainless Steel)*.



3. Click **OK**.

   The **Al** material node is replaced with **UNSS31600**.

4. Expand the **Material Properties** node to view the density, specific heat, and thermal conductivity properties.



The default values are appropriate for this tutorial.

Defining the Solid Boundary and Initial Conditions

Set boundary conditions for the solid region and the initial temperature.

For the independent simulation of the chimney, set the outer boundary to lose heat by convection, and the inner wall boundary to maintain a constant temperature. The **OuterSurf** boundary is exposed to air at 283 K with a heat transfer coefficient of 0.12 W/m^2-K.

1. Expand the **Regions** > **Chimney** > **Boundaries** node.

To apply the convection boundary:

2. Edit the **OuterSurf** node, and set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Convection**
Physics Values		
Ambient Temperature > Constant	*Value*	`283.0 K`
Heat Transfer Coefficient > Constant	*Value*	`0.12 W/m^2-K`

To apply the constant temperature boundary:

3. Edit the **InnerSurf** node, and set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Temperature**
Physics Values		
Static Temperature > Constant	*Value*	`380.0 K`

Leave the **Default** boundary as adiabatic.

4. Select the **Continua** > **Solid** > **Initial Conditions** > **Static Temperature** > **Constant** node, and set *Value* to `283.0 K`.

Preparing a Temperature Scene for the Solid

Create a scene of the temperature for the entire chimney and create a file to record the transient solution history of the chimney.

To create a temperature scene:

1. Create a scalar scene.
2. Rename the **Scalar Scene 1** node to `Wall Temperature`.
3. Edit the **Wall Temperature** > **Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1		
Parts	*Parts*	**Default**
		InnerSurf
		OuterSurf
Scalar Field	*Function*	**Temperature**
	Clip	**Off**
	Min	`280`
	Max	`400`

Add an annotation to the scene to display each time step:

4. Select the **Tools** > **Annotations** > **Time Step** node and drag it onto the *Wall Temperature* scene.
5. Click the annotation text in the *Wall Temperature* scene and move it to the upper right corner of the solution scene.
6. Position the chimney as shown below:



Set update trigger to time step:

7. Select the **Scenes** > **Wall Temperature** > **Attributes** > **Update** node and set *Trigger* to **Time Step**.



Record the transient solution history of the chimney:

8. Right-click the **Solution Histories** node and select **New...**
9. In the *Save* dialog, save the solution history to your working directory as `ChimneyHistory.simh`.
10. Edit the **Solution Histories** node and set the following properties:

Node	Property	Setting
ChimneyHistory	*Scalar Functions*	**Temperature**
Update	*Trigger*	**Time Step**

11. Save the simulation.

Setting Stopping Criteria for the Solid Simulation

Set solver parameters and stopping criteria for the simulation.

Run the case for 5 minutes with a time step of 20 seconds:

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `20 s`.
2. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`
Maximum Physical Time	*Maximum Physical Time*	`300 s`

Running the Solid Simulation

Run the simulation for the solid chimney.

1. Run the simulation.
2. When the solution finishes, save the simulation.

Visualizing the Decoupled Solid Results

Visualize the temperature distribution on the chimney outer wall.

The temperature distribution on the outer wall is uniform along the outer wall of the pipe.



These results can be compared against the coupled results later.

1. Save and close the simulation.

Preparing the Fluid Model Simulation

Prepare the simulation for the fluid inside the chimney and run it to convergence.

Follow these steps to prepare the fluid simulation:

1. Importing the Fluid Volume Mesh.
2. Selecting Fluid Physics Models.
3. Specifying Fluid Material Properties.
4. Defining the Fluid Boundary and Initial Conditions.
5. Preparing a Temperature Scene for the Fluid.
6. Setting Stopping Criteria and Running the Fluid Simulation.

7. Visualizing the Decoupled Fluid Results.

Contents

Importing the Fluid Volume Mesh
Selecting Fluid Physics Models
Specifying Fluid Material Properties
Defining the Fluid Boundary and Initial Conditions
Preparing a Temperature Scene for the Fluid
Setting Stopping Criteria and Running the Fluid Simulation
Visualizing the Decoupled Fluid Results

Importing the Fluid Volume Mesh

To set up the STAR-CCM+ simulation for the fluid, launch a simulation and import the supplied volume mesh representing the fluid inside the chimney.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
5. Select `flue.ccm` and click **Open**.
   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
6. Save the simulation as `Flue.sim`.

Selecting Fluid Physics Models

Select physics models to describe the fluid inside the chimney, and specify material properties.

To select the physics models:

1. Rename the **Continua > Physics 1** node to `Fluid`.
2. For the physics continuum, **Continua > Fluid**, select the following models in order:

Group Box	Model
Enabled Models	**Three Dimensional** (Selected automatically)
Time	**Steady**

Group Box	Model
Material	**Gas**
Flow	**Segregated Flow**
	Gradients (Selected automatically)
Equation of State	**Ideal Gas**
Optional Models	**Segregated Fluid Temperature**
Viscous Regime	**Laminar**

3. Click **Close**.
4. To review the models, open the **Fluid** > **Models** node.



5. Save the simulation.

Specifying Fluid Material Properties

Specify the properties of the exhaust gas as having the same properties as air at 600 K.

To set the gas properties:

1. Expand the **Continua** > **Fluid** > **Models** > **Gas** > **Air** node.

2. Edit the **Material Properties** node and set the following properties:

Node	Property	Setting
Dynamic Viscosity > Constant	*Value*	`3.058E-5 Pa-s`
Molecular Weight > Constant	*Value*	`28.57 kg/kmol`
Specific Heat > Constant	*Value*	`1051.0 J/kg-K`
Thermal Conductivity > Constant	*Value*	`0.0469 W/m-K`

3. Save the simulation.

Defining the Fluid Boundary and Initial Conditions

Specify conditions for the boundaries of the fluid region. Also, set the fluid initial temperature and velocity.

The gas entering the chimney is traveling at 0.2 m/s with a temperature of 600 K. The Reynolds number for this scenario is 1518, thus justifying the laminar flow model. Specify these boundary conditions under the inlet boundary node:

1. Expand the **Regions > Flue > Boundaries** node.
2. Edit the **Inlet** node and set the following properties:

Node	Property	Setting
Physics Values		
Static Temperature > Constant	*Value*	`600 K`
Velocity Magnitude > Constant	*Value*	`0.2 m/s`

Set the outlet boundary to be at a static temperature of 283 K:

3. Select the **Outlet > Physics Values > Static Temperature > Constant** node and set *Value* to `283 K`.

For the decoupled case, assume that the inner wall is at a constant 380 K:

4. Edit the **HotGas** node and set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Temperature**
Physics Values		
Static Temperature > Constant	*Value*	`380 K`

Set initial conditions of 283 K and a constant velocity of 0.2 m/s in the Y direction:

5. Edit the **Continua** > **Fluid** > **Initial Conditions** node and set the following properties:

Node	Property	Setting
Static Temperature > Constant	*Value*	`283 K`
Velocity > Constant	*Value*	`[0.0, 0.2, 0.0] m/s`

6. Save the simulation.

Preparing a Temperature Scene for the Fluid

Prepare a scalar scene to display the temperature within the fluid region.

To display the temperature within the fluid, first create a plane section through the fluid region:

1. Open *Geometry Scene 1*.
2. Within the **Derived Parts** node, create a plane section with the following properties:

Property	Setting
Input Parts	**Flue**
normal	`[0, 0, 1] m`
Display	**No Displayer**

3. Click **Create**, then **Close** the *Create Section* dialog.



Create a scalar scene to display the temperature on the plane section:

4. Create a scalar scene.
5. Rename the **Scalar Scene 1** node to `Fluid Temperature`.
6. Edit the **Fluid Temperature** > **Displayers** node and set the following properties:

Node	Property	Setting
Scalar 1	*Contour Style*	`Smooth Filled`
Parts	*Parts*	**Derived Parts** > **plane section**
Scalar Field	*Function*	**Temperature**

7. Save the simulation.

Setting Stopping Criteria and Running the Fluid Simulation

Set the stopping criteria, then run the fluid simulation.

Reduce the maximum number of iterations stopping criteria to 100:

1. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `100`.

2. Run the simulation.

   The solution converges within 100 iterations.
3. Save the simulation.

Visualizing the Decoupled Fluid Results

The temperature distribution in the gas is not uniform, even for this low speed flow. This means that the assumption of constant wall temperature used in `Chimney.sim` is not valid and an improved solution of the wall temperature distribution can be obtained by coupling the simulations.



These results can be compared against the coupled results later.

1. Save, then close the simulation.

Coupling the Simulations

Specify the co-simulation settings and run the co-simulation.

To set up and run the coupled simulations interactively:

1. Load each partner case in a single STAR-CCM+ session.
2. Define the co-simulation properties for the lagging partner.
3. Define the co-simulation session properties for the leading partner.
4. Run the partners from the lead partner simulation.

Contents

Loading the Simulations
Modifying the Physics Models for the Lagging Simulation
Preparing the Simulation that Lags
Modifying the Physics Models for the Leading Simulation
Preparing the Simulation that Leads
Running the Co-Simulation
Visualizing the Results

Loading the Simulations

Load both the simulations, that you created for the solid and fluid models, in the same instance of STAR-CCM+.

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to the directory where you saved the simulation for the fluid model.
5. Select `Flue.sim`, then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.
7. Save the simulation as `FluidC.sim`.
8. Repeat this process for the Chimney.sim file and save it as `ChimneyC.sim`.

   Both cases appear in the *Explorer Pane*:



Note: It is not necessary to have both simulations in the same instance of STAR-CCM+. It is equally valid to have the leading and lagging simulations in separate instances of STAR-CCM+ that can communicate with each other.

Modifying the Physics Models for the Lagging Simulation

Modify the solid simulation physics models to allow coupling.

In a coupled analysis, there are two partner simulations: a leading partner and a lagging partner.

Activate coupling in the case that lags:

1. In the *Explorer Pane*, make sure that the ChimneyC simulation is highlighted.
2. For the physics continuum, **Continua** > **Solid**, select the following models in order:

Group box	Model
Optional Models	**Co-Simulation**
Co-Simulation	**STAR-CCM+ Explicit Co-Simulation**

3. Click **Close**.

   A new physics continuum, **Physics 1**, is added within the **Continua** node to represent the partner simulation.

   The models appear as shown below:



Preparing the Simulation that Lags

Prepare the solid simulation for co-simulation.

When you activate a co-simulation model, STAR-CCM+ adds a **Co-Simulations** node to the simulation tree. This node allows you to specify global co-simulation settings, such as the *Concurrency Mode* and the *Connect Method*, as well as define and set up the co-simulation coupling zones.

A coupling zone defines a set of boundaries where fields are exchanged with the partner simulation. STAR-CCM+ supports multiple coupling zones, that allow to specify different conditions on different sets of coupled boundaries.

As this is a standard thermal co-simulation analysis, you set up a single coupling zone defining the solid-fluid interface, and you specify solid thermal import and export conditions.

Start by assigning a suitable name to the co-simulation session that STAR-CCM+ has created for you:

1. Rename the **Co-Simulations** > **STAR-CCM+ Co-Simulation 1** node to `Fluid Flue`.

Specify the inner surface of the chimney as the coupled boundary:

2. Expand the **Regions** > **Chimney** > **Boundaries** > **InnerSurf** > **Physics Conditions** node.
3. Select the **External Code Coupling Specification** node and set *External Code Coupling Specification* to **Zone 1**.



The *Method* for the **Physics Values** > **Ambient Temperature** node is automatically set to **Reference Temperature applied by Co-Simulation**.

STAR-CCM+ updates the conditions for **Zone 1** by adding **InnerSurf** to the **Zone 1** coupled boundaries.



In this tutorial, the solid simulation imports convective fluxes from the fluid simulation and exports temperature. You specify the fields that the solid simulation sends to the fluid simulation through the coupling zone **Conditions** node.

The default export settings for this simulation are appropriate:

4. Expand the **Co-Simulations** > **Fluid Flue** > **Zones** > **Zone 1** > **Conditions** node.

Node	Property	Setting
Solid Thermal Export Specification	*Method*	**Temperature Export Field** (default setting)
Temperature Export Field	*Temperature*	**Temperature** (default setting)

Set the solid simulation to be the lagging simulation:

5. Select the **Co-Simulations** > **Fluid Flue** > **Conditions** > **Concurrency mode** and set *Method* to *Lag*.

Select a coupling negotiation method:

6. Select the **Solvers** > **Co-Simulation Explicit Unsteady** > **Coupling Step Size Negotiation Method** > **Constant Negotiation Method node** and set *Constant Time Step* to `20 s`.



Increase the run time for this solution:

7. Select the **Stopping Criteria** > **Maximum Physical Time** node and set *Maximum Physical Time* to `1800 s`.

Clear the results from the previous simulation:

8. Select **Solution** > **Clear Solution**.
9. In the *Clear Solution* dialog, click **OK**.
10. Save the simulation.

Modifying the Physics Models for the Leading Simulation

Modify the fluid simulation physics models to allow coupling.

Activate coupling in the case that leads:

1. In the *Explorer Pane*, make sure that the FluidC simulation is highlighted.
2. For the physics continuum, **Continua** > **Fluid**, select the following models in order:

Group box	Model
Optional Models	**Co-Simulation**
Co-Simulation	**STAR-CCM+ Steady Co-Simulation**

3. Click **Close**.

A new physics continuum, **Physics 1**, is added within the **Continua** node to represent the partner simulation.

The models appear as shown below:



Preparing the Simulation that Leads

Prepare the fluid simulation for co-simulation.

Select the coupled boundaries and use the **Co-Simulations > STAR-CCM+ Co-Simulation 1** node to specify global co-simulation settings and to set up the coupling zone.

1. Rename the **STAR-CCM+ Co-Simulation 1** node to `Solid Chimney`.

To specify the **HotGas** boundary as the coupled boundary:

2. Expand the **Regions > Flue > Boundaries > HotGas > Physics Conditions** node.
3. Select the **External Code Coupling Specification** node and set *External Code Coupling Specification* to **Zone 1**.

The *Method* for the **Physics Values > Static Temperature** node is automatically set to **Temperature applied by Co-Simulation**.

STAR-CCM+ updates the conditions for **Zone 1** by adding **HotGas** to the **Zone 1** coupled boundaries.



Specify the fields that the fluid simulation exports to the solid simulation:

4. Expand the **Co-Simulations > Solid Chimney > Zones > Zone 1 > Conditions** node.
5. Select the **Fluid Thermal Export Specification** node and set *Method* to **Convective Flux Export Fields**.

Set the fluid simulation to be the leading simulation:

6. Select the **Co-Simulations > Solid Chimney > Conditions > Concurrency mode** and set *Method* to *Lead*.

To connect to the launched partner co-simulation, specify the connect method between the two cases:

7. Select the **Co-Simulations > Solid Chimney > Conditions > Connect Method > Connect to Partner Simulation** node.

8. Set *Connect to STAR-CCM+ Server Process* to the STAR-CCM+ server representing the solid simulation:

   a) Click ... (**Custom Editor**) to open the *Connect to Partner Simulation* dialog.

   b) In the *Connect to Partner Simulation* dialog, select the **ChimneyC** server and click **OK**.



To set the number of iterations before exchanging fields:

9. Select **Solvers > Co-Simulation Steady > Iterations per Exchange** and set *Number of Iterations per Exchange* to `100`.



The final step is to increase the number of iterations for this solution.

10. Select the **Stopping Criteria** > **Maximum Steps** node and set *Maximum Steps* to `9000`.

    This is the maximum physical time for Chimney.sim (1800), divided by the number of time steps per second (20), times the number of iterations per update (100).
11. Clear the results from the previous simulation.
12. Save the simulation.

Running the Co-Simulation

The simulations are fully set up and the co-simulation is ready to be run.

Running a co-simulation interactively requires an instruction for the lead partner to start solving. Before the solution begins, the lead partner connects to its lagging partner.

1. Make sure the leading simulation, FluidC, is highlighted in the *Explorer Pane*.
2. Click (**Run**).

   During the run, the leading simulation connects to the lagging simulation. The simulations take turns solving their part of the problem before updating the boundary conditions at the coupled boundary. When its partner simulation is solving, the progress bar on the waiting partner displays the message *Waiting for partner code to finish*.
3. Open the *Residuals* plots for each case to check convergence behavior.
4. When the co-simulation is complete, save each simulation.

Visualizing the Results

Visualize the temperature scenes in the solid and the fluid simulation.

The final results for the FluidC case appear as shown below:



The final results for the ChimneyC case appear as shown below:



The coupling has improved the quality of the solution of both simulations. However, of the two, the temperature distribution in the metal chimney has changed the most. The chimney shows hot spots near the inlet and on the wall furthest away from the apex of the bend.

Display the ChimneyC solution at an interim time-step:

1. Right-click the **Solution Histories** > **ChimneyHistory** node and select **Create Recorded Solution View**.

   The **ChimneyHistory** node appears.
2. Click the *ChimneyC-Wall Temperature* tab in the *Graphics* window.
3. Drag and drop the **Solution Views** > **ChimneyHistory** node onto a blank area in the scene window.
4. Select the **Solution Views** > **ChimneyHistory** node and set *Time Step* to `45`.

   The scene updates, as shown below:



The difference between the temperature distribution at Time Step 45 and Time Step 90 shows that the metal is heating gradually as a result of the hot exhaust gases.

5. Save the simulation.

Summary

This tutorial demonstrated the process for modeling conjugate heat transfer using STAR-CCM+ to STAR-CCM+ thermal co-simulation.

This tutorial has demonstrated the following features:

- Setting up a STAR-CCM+ to STAR-CCM+ co-simulation.
- Creating the independent solid simulation and checking for convergence.
- Creating the independent fluid simulation and checking for convergence.
- Coupling two STAR-CCM+ simulations.

- Viewing intermediate results.

GT-POWER Co-Simulation: 1D Coupling

This tutorial demonstrates the typical workflow when running a coupled simulation with the engine analysis code, GT-POWER from Gamma Technologies.

The GT-POWER model simulates a four-cylinder, four-stroke engine running at 3000 rpm on Indolene fuel. It also includes the effect of Exhaust Gas Recirculation (EGR). GT-POWER uses a one-dimensional method for simulating flow and combustion within an engine. By coupling GT-POWER with STAR-CCM+, you can examine detailed three-dimensional flow in a component, and use the one-dimensional code to calculate flow in the rest of the powertrain. In this tutorial, STAR-CCM+ calculates the three-dimensional flow within an intake manifold that forms part of an overall engine analysis that GT-POWER solves.



The GT-POWER coupling feature exchanges species mass fractions and physics values between STAR-CCM+ and GT-POWER once per time-step. If flow is entering the STAR-CCM+ region, the coupling model applies the physics values and mass fractions from GT-POWER uniformly at the STAR-CCM+ boundary. If flow is leaving the STAR-CCM+ region, the coupling model volume-averages the physics values and mass fractions in a zone and sends them to GT-POWER.

Coupling exchanges the following data at each time step:

- Species Mass Fractions
- Density (Mass Flow inlet only)
- Normal Velocity (Mass Flow inlet only)
- Pressure (Pressure outlet only)
- Mass Flow
- Temperature

In this tutorial, you obtain initial conditions for the STAR-CCM+ simulation from an uncoupled, transient GT-POWER solution.

During the simulation, the GT-POWER solver runs 20 pre-cycles before transferring data to STAR-CCM+, which allows the GT-POWER model to converge to a periodic steady-state. Once coupling begins, the solvers exchange data every time-step for three more engine cycles.

Contents

Prerequisites
Preparing the File Structure
Setting Up the GT-POWER Model
Setting Up the STAR-CCM+ Simulation
Coupling STAR-CCM+ to GT-POWER
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the GT-POWER Co-Simulation - 1D Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

This tutorial also assumes that you are familiar with the GT-ISE interface and workflow. If you have no experience with GT-ISE, complete the standard GT-POWER tutorial. You can access this tutorial in GT-ISE from **Help** > **Tutorials** > **Modeling_Applications** > **Engine_Performance** > **Engine_Performace-tutorials.pdf**.

This tutorial describes the procedure using GT-POWER v7.3.

Preparing the File Structure

To couple STAR-CCM+ to a GT-POWER simulation, you specify the location of the GT-POWER case file. Therefore, arrange a simple directory structure so the case is easy to find.

The following files are provided:

- GT-POWER model file, `CFDCoupling_Start.gtm`, representing an engine in operation.
- A STAR-CCM+ mesh, `manifoldGT-POWER.ccm`, representing the three-dimensional intake manifold.

To prepare the file structure:

1. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
2. Copy the `manifoldGT-POWER.ccm` file to your current working directory.
3. In the current working directory, create a subdirectory called `caseFilesGT-POWER`.
4. Copy the `CFDCoupling_Start.gtm` file to the new directory.

This structure makes it simpler when defining the GT-POWER case file to use in the coupled simulation. The diagram below illustrates the required file structure.



Setting Up the GT-POWER Model

Set up the case in GT-POWER.

In this tutorial, the procedure to set up the GT-POWER model is as follows:

1. Loading the GT-POWER Case File:

   You begin setting up the GT-POWER model by opening a supplied model file and saving a working copy. You modify this file so that it is suitable for coupling.
2. Checking Model Validity and Convergence Criteria:

   Analyze GT-POWER results before coupling. It is poor practice to begin a coupled simulation before an analysis of standalone GT-POWER results.
3. Obtaining Initial Pressure, Temperature, and Mass Fractions:

Obtain STAR-CCM+ initial conditions using uncoupled GT-POWER results to speed up the coupled model convergence.

4. Adding CFD Objects to the Project Tree:

   Make CFD objects available that couple the model with the three-dimensional simulation in STAR-CCM+.

5. Adding CFDInterfaces to the Model Map:

   Create and add CFDInterfaces, which switch the flow solution from the one-dimensional components to STAR-CCM+ when coupling begins.

6. Adding CFDConnections to the Model Map:

   Create and add CFDConnections, which specify the point at which coupling occurs.

7. Adding a CFDComponent to the Model Map:

   Create and add a CFDComponent, which represents the three-dimensional STAR-CCM+ simulation in the GT-POWER model.

8. Defining Species Mapping:

   Define a species map that reduces the computational costs in STAR-CCM+ by solving fewer gas components.

9. Modifying the Run Setup:

   Make sure the GT-POWER simulation duration contains the coupled cycles. This practice allows the GT-POWER solver to collect results for the coupled cycles for post-processing within GT-POST.

10. Creating the .dat file:

    Create a `.dat` file that provides an input file for the GT-POWER solver that defines the case setup.

Contents

Loading the GT-POWER Case File
Checking Model Validity and Convergence Criteria
Obtaining Initial Conditions
Adding CFD Objects to the Project Tree
Adding CFDInterfaces to the Model Map
Adding CFD Connections to the Model Map
Adding a CFD Component to the Model Map
Defining Species Mapping
Modifying the Run Setup
Creating the .dat file

Loading the GT-POWER Case File

Load the GT-POWER case file in GT-ISE.

You begin setting up the GT-POWER model by opening a supplied model file and saving a working copy. You then modify this file so that it is suitable for coupling.

To begin:

1. Launch GT-ISE.
2. Load the `CFDCoupling_Start.gtm` model file.
3. Save a working copy of the model file called `CFDCoupling_Setup.gtm`.

Checking Model Validity and Convergence Criteria

Run an independent GT-POWER analysis. It is poor practice to begin a coupled simulation before an analysis of standalone GT-POWER results.

Analyze GT-POWER results before coupling. This practice makes sure that the model is valid and determines how many cycles are required to achieve convergence. It also reduces the risk of error and inaccuracies when running the coupled simulation.

First, make sure that the solver terminates when the periodic, steady-state criteria are met:

1. From the menu bar, select **Run** > **Run Setup**.
2. In the *Run Setup* dialog, select the *TimeControl* tab.
3. Make sure that *Automatic Shut-Off When Steady-State* is set to **on**.



4. Click **OK**.

Make sure that GT-POWER writes Time RLT output data once per engine cycle:

5. From the menu bar, select **Run > Output Setup**.
6. In the *Output Setup* dialog, select the *Data_Storage* tab.
7. Make sure that *Time RLT Storage Multiple (Periods)* is set to `1`.



8. Click **OK**.

To check that the simulation is valid and obtain initial conditions for the STAR-CCM+ simulation, run the GT-POWER solver:

9. In the toolbar, click **Start Simulation**.

10. In the *Simulation Wizard* dialog, click **Finish**.
11. When the solver is complete, close the solver window.

Using GT-POST, you can now analyze the results to assess convergence.

12. In the toolbar, click **Open GT-POST** and select the `CFDCoupling_Setup.gx` post data file.

13. Select the *Time RLT* tab.
14. Right-click the **EngineCrankTrain > engine > Performance > BMEP** plot, and select **View**.

    As shown below, the simulation reaches a steady-state solution after 18 cycles, when the solver terminates. Therefore, in a later stage you run 20 pre-cycles before allowing GT-POWER to begin sending data to STAR-CCM+.



Obtaining Initial Conditions

Obtain STAR-CCM+ initial conditions, using uncoupled GT-POWER results to speed up the coupled model convergence.

Note: The initial conditions in STAR-CCM+ do not need to match the GT-POWER model exactly, as the influence of the initial conditions reduces during the transient solution. However, a close approximation speeds up convergence during the first few time-steps.

- Obtaining the Initial Pressure Condition
- Obtaining the Initial Temperature Condition

- Obtaining the Initial Mass Fractions

Contents

Obtaining the Initial Pressure
Obtaining the Initial Temperature
Obtaining the Initial Mass Fractions

Obtaining the Initial Pressure

Use uncoupled GT-POWER results to obtain an initial pressure for the simulation in STAR-CCM+.

To obtain the initial pressure:

1. Select the *Plots* tab.
2. From the menu bar, select **Macro > RLT Plots**.
3. In the Run *RLT Macro* dialog, select **Time RLT** and click **Next**.
4. In the *Data* tab, select the **Plot1**, *X* cell.
5. In the *RLT Selection* group box, select the **Rlt tree > Special** node and select the **Cycle Number** data.
6. In the *Data* tab, select the **Plot1**, *Y1* cell.
7. In the *RLT Selection* group box, select the **FlowSplitTRight > InMan01 > Pressure** node, and select the **Pressure at End of Cycle** data.
8. Repeat the above steps and select the following data for the remaining components.

Cell	Node	Data
Y2	**FlowSplitTRight > InMan02**	**Pressure at End of Cycle**
Y3	**FlowSplitTRight > InMan03**	**Pressure at End of Cycle**
Y4	**FlowSplitTRight > InMan04**	**Pressure at End of Cycle**
Y5	**FlowSplitGeneral > egr-fs-cfd**	**Pressure at End of Cycle**
Y6	**PipeRound > InPipe-cfd**	**Pressure at End of Cycle (Outlet)**

Data | Constraints

Plot...	Cases	X	Y1	Y2	Y3	Y4	Y5	Y6
Plot1	1	period	pinst:InMan01	pinst:InMan02	pinst:InMan03	pinst:InMan04	pinst:egr-fs-cfd	pinstr:InPipe-cfd
Plot2								
Plot3								
Plot4								
Plot5								

9. Click **Next**.
10. In the *Dataset Label and Display Options* box, activate **Part Name**.
11. Deactivate **Use Short Case Legend** and **Constraint RLT, Min/Max**.
12. Click **Finish** to display the plot.

    When complete, GT-POST displays the following plot, which shows that the initial pressure is approximately 2.13 bar.



Obtaining the Initial Temperature

Use a similar procedure as detailed in the previous section to determine an approximate initial temperature.

To obtain the initial temperature:

1. Set the x- and y-axis data for the RLT plot as summarized below:

Cell	Node	Data
X	**Rlt tree > Special**	**Cycle Number**
Y1	**FlowSplitTRight > InMan01 > Thermal**	**Temperature at End of Cycle**
Y2	**FlowSplitTRight > InMan02 > Thermal**	**Temperature at End of Cycle**
Y3	**FlowSplitTRight > InMan03 > Thermal**	**Temperature at End of Cycle**
Y4	**FlowSplitTRight > InMan04 > Thermal**	**Temperature at End of Cycle**
Y5	**FlowSplitGeneral > egr-fs-cfd > Thermal**	**Temperature at End of Cycle**
Y6	**PipeRound > InPipe-cfd > Thermal**	**Temperature at End of Cycle (Outlet)**

When complete, GT-POST displays the following plot, which shows that the initial temperature is approximately 315 K.



Obtaining the Initial Mass Fractions

Use a similar procedure as detailed in the previous section to determine initial mass fractions.

To obtain the initial mass fractions:

1. Select the *TimeRLT* tab.
2. Right-click the **Components** > **SpeciesSampler** > **Cylinder_Runner_1** > **Main** > **n2-vap Mass Fraction** node and select **View**.



When complete, GT-POST displays the following plot, which shows that the initial n2-vap mass fraction is approximately 0.60.



3. Repeat the above procedure for each gas component, to find the remaining initial mass fractions.

   Trace components (<1%) are not included in the initial composition. All trace components are included in the n2-vap mass fraction, as it is the largest of the gas components and the total mass fraction must equal 1.

   The initial mass fractions are summarized in the table below:

Gas Component	Mass Fraction
n2-vap (Fresh Nitrogen)	0.601
o2-vap (Fresh Oxygen)	0.183
prod_n2 (Burnt Nitrogen)	0.154
prod_co2 (Carbon Dioxide Products)	0.043
prod_h2o (Water Products)	0.016
prod_o2 (Burnt Oxygen)	0.001

4. Close GT-POST.

Adding CFD Objects to the Project Tree

Make CFD objects available that couple the model with the three-dimensional simulation in STAR-CCM+.

Note: In your own cases, make sure that there are no RLT dependencies or sensors that are connected to the components that represent the three-dimensional region. When coupling commences, these components are deactivated and the flow is solved in STAR-CCM+, so the RLT dependencies and sensors become redundant.

To add CFD objects to the GT-POWER project tree:

1. In the toolbar, click the **Tile with Libraries** button.

2. Make sure that the *Flow* tab is selected and open the **General Flow** > **Components** node.
3. Drag and drop the **CFDComponent** node from the *Template Library* to the *Project Tree*.



4. Similarly, drag and drop the **General Flow** > **Components** > **CFDInterface** node, and the **General Flow** > **Connections** > **CFDFlowConn** node.



Adding CFDInterfaces to the Model Map

Create and add CFDInterfaces, which switch the flow solution from the one-dimensional components to STAR-CCM+ when coupling begins.

During the uncoupled cycles, GT-POWER solves the flow in the one-dimensional components. When coupling commences, the GT-POWER solver supplies boundary conditions to STAR-CCM+ and no longer solves the flow in the one-dimensional components.

Note: There is no overlap at the interface between the one-dimensional and three-dimensional models. Therefore, in your own cases, split pipes in the GT-POWER model to accommodate the three-dimensional region.

To add CFDInterfaces to the model map:

1. In the six 1D/3D sections, delete all default orifices, as highlighted in the following image.



2. Replace the default orifices with a **CFDInterface > def** component.



3. Join the components on either side of the CFDInterface with the **Start A New Link** tool, ensuring that the arrows point in the direction of flow.



4. Save the GT-POWER model.

Adding CFD Connections to the Model Map

Add CFDFlowConn connections to the model to set data exchange between STAR-CCM+ and GT-POWER.

CFDFlowConn connections determine the point where data is exchanged between STAR-CCM+ and GT-POWER. The settings require

you to define the boundary type, diameters, averaging zone length and turbulence parameters.

Note: GT-POWER permits the use of pipes with circular or rectangular cross-sections. For cross-sections that do not meet these requirements, calculate the effective diameter. You can find this effective diameter by using a Sum report of the Area field function at the coupled boundary in STAR-CCM+.

To add CFDFlowConn connections to the model:

1. In the project tree, right-click the **Connections** > **CFDFlowConn** node and select **Add Object**.
2. In the *Create Object* dialog, set the following attributes:
   a) Set *Object* to `CFDConnection` to name the CFD connection.
   b) Make sure that *Boundary Type* is set to **inlet**, to transfer mass flow data (as opposed to pressure data at a pressure boundary).
   c) Make sure that *Orifice Diameter* is set to **def**, which matches the value to the diameter of the adjacent pipe in the GT-POWER model.
   d) Make sure that *CFD Domain Diameter* is set to **def**, as there is no difference between the diameter in the GT-POWER and STAR-CCM+ model.
   e) Set *CFD Domain DX* to `[dxi]`, to define the STAR-CCM+ volume-averaging zone length with the same value as the discretization length in GT-POWER. Generally, use a value greater than `25 mm` to prevent oscillations at the coupled boundary.
   f) Set *Turbulence Intensity* to `0.05`
   g) Set *Turbulence Length Scale* to `5 mm`



3. Click **OK**.
4. Drag and drop the **CFDFlowConn** > **CFDConnection** connection on to the project map and place it near the CFDInterface between **InPipe04** and **InPipe-cfd**.
5. Rename the CFDConnection as `Intake`.
6. Join the components from the **CFDInterface** to the **CFDConnection**.



7. Create another CFDConnection, with the settings shown below:



8. Place and join the **CFDConnEGR** connection between the **egr_pipe-1** and **egr_pipe-2**, and rename it as EGR.



9. Create a final CFDConnection, with the settings shown below:



10. Place and join four **CFDConnCyl** connections between the remaining 1D/3D sections (for example, between InRunner-cfd-1 and InRunner01), and rename them as `Cyl1`, `Cyl2`, `Cyl3`, and `Cyl4`.



11. Save the GT-POWER model.

Adding a CFD Component to the Model Map

Create and add a CFDComponent, which represents the three-dimensional STAR-CCM+ simulation in the GT-POWER model.

Note: In versions of GT-POWER before 7.1, it was necessary to specify a *CFD Domain Volume*. For this case, the CFDInterfaces allow you to model the manifold using one-dimensional components during the uncoupled cycles.

To add a CFDComponent to the model:

1. Right-click the **Components** > **CFDComponent** node and select **Add Object**.
2. In the *Create Object* dialog, set *Object* to `Manifold` to name the CFD component.
3. In the *Simulation Phases* tab, set *Engine Cycles Before CFD Started* to `20`.



4. Click **OK**.
5. Drag and drop the **CFDComponent** > **Manifold** connection on to the project map and link all six CFDConnections to it. After resizing the **Manifold** component, the map appears as shown below.



Defining Species Mapping

Define species mapping, so that GT-POWER can map STAR-CCM+ gas components to its own species.

GT-POWER maps the gas components from STAR-CCM+ to its own species. The advantage to using species mapping is that you do not have to represent every GT-POWER species in the STAR-CCM+ simulation,

which reduces the computational cost. Using species mapping also provides greater flexibility for you to choose which individual species are fully tracked in STAR-CCM+.

In this GT-POWER model, there are four species that define the fresh air-fuel mixture (intake charge) entering the cylinder. In addition to these species, there are 13 products of combustion, which are typical for GT-POWER v7.3 simulations. Therefore, a total of 17 species are present in the GT-POWER simulation. The intake charge species are always individually represented in STAR-CCM+. However, it is not necessary to track all burned species in STAR-CCM+, as some are only present in trace amounts. The lump_N2 scalar represents these negligible species in STAR-CCM+.

You map the species as illustrated in the diagram below:



In this case, one species is always mapped to one cluster. This behavior is not a limitation of the mapping facility. For example, you could map all harmful gases to a cluster called Pollutants.

To implement the above species mapping:

1. Right-click the **Components** > **CFDComponent** > **Manifold** > **Manifold-1** node and select **Edit Properties**.
2. In the *Edit Part: Manifold-1* dialog, select the *Coupled Cycles* tab.
3. Click the ellipsis (...) next to the *Species Mapping Object* attribute.
4. In the *Value Selector* dialog, select **CFDSpeciesMap** from the *Create New Object* group box and click **OK**.
5. In the *Create Object* dialog, set *Object* to **CFDSpeciesMap**.
6. Set *User Defined Species Map* to `def`.
7. Click the ellipsis (...) next to the *prod_co2's CFD Species* attribute.
8. In the *Value Selector* dialog, select **CFDSpecies** > **prod_co2** from the *GT-SUITE Library* tab and click **OK**.
9. Use a similar method to set the remaining attributes, as shown below:



10. Click **OK**.

11. Save the GT-POWER model.

Modifying the Run Setup

Modify the run setup to allow the GT-POWER simulation to run 20 pre-cycles, and then 3 coupled cycles.

Make sure that the GT-POWER simulation duration contains the coupled cycles. This practice allows the GT-POWER solver to collect results for the coupled cycles for post-processing within GT-POST.

Previously, you ensured the solver would stop when the solution was steady-state. However, you begin coupling to STAR-CCM+ after the GT-POWER solution reaches a periodic steady-state. Allow the solver to continue after the steady-state criteria are met, and set the simulation to run 20 pre-cycles, plus 3 coupled cycles, which requires 23 cycles in total.

To modify the run setup:

1. From the menu bar, select **Run > Run Setup**.
2. In the *Run Setup* dialog, select the *Time Control* tab.
3. To run the solver for the required number of engine cycles, set *Maximum Simulation Duration* to `23`.

   This value gives enough cycles for the GT-POWER simulation to converge (20), plus the required number of coupled cycles (3).
4. Set *Automatic Shut-Off When Steady-State* to **Off**, so that the solver runs for the maximum number of cycles, without terminating prematurely.



5. Click **OK**.
6. Save the GT-POWER model.

Creating the .dat file

Create a .dat file that provides an input file for the GT-POWER solver that defines the case setup.

To create a .dat file:

1. From the menu bar, click **Run** > **Run-Simulation Utilities** > **Create .dat file**.
2. In the *.dat File Warning* dialog, click **Yes** to overwrite the existing results.

To make sure that the model is valid, run the simulation preprocessing feature. This action generates a file containing materials properties of all species in the GT-POWER model, which can be used in STAR-CCM+ to define identical materials.

3. From the toolbar, click **Start Simulation Preprocessing**.

4. In the *.dat File Warning* dialog, click **Yes** to overwrite the existing results.

   If the model is valid, the solver preprocessor exits with no errors or warnings.
5. Make sure that the `CFDCoupling_Setup.chemkin_GT.dbs` file is present in the current working directory, as this file is required at a later stage of the tutorial.
6. Close GT-ISE.

Setting Up the STAR-CCM+ Simulation

Prepare the simulation in STAR-CCM+.

In this tutorial, the procedure to set up the STAR-CCM+ simulation is as follows:

1. Importing the Volume Mesh:

   You begin setting up the STAR-CCM+ simulation by launching a session and importing a supplied volume mesh.
2. Selecting Physics Models:

   Physics models determine the behavior of the flow in the simulation. In this tutorial, you activate a physics model that enables coupling between STAR-CCM+ and GT-POWER.
3. Importing the GT-POWER Material Properties:

   Create a material database and import the material properties from the GT-POWER simulation.
4. Setting Up the Species Representation:

   Define gas components and passive scalars under the relevant physics models. To reduce computational costs, you use six gas components and two passive scalars to represent the species that GT-POWER transfers to STAR-CCM+.
5. Setting the Reference Pressure:

   STAR-CCM+ uses a reference pressure to reduce round-off error, with all other pressures in the simulation defined relative to its value. However, in GT-POWER, all pressures are absolute.
6. Setting Initial Conditions:

Set the initial conditions to match the conditions in the GT-POWER model at the start of coupling. This practice speeds up convergence at the start of the STAR-CCM+ simulation.

7. Defining Coupled Boundaries:

   Set all flow boundaries as coupled with GT-POWER. This setting defines the boundaries available for the transfer of data between STAR-CCM+ and GT-POWER.

8. Setting the Time Step:

   Set the time step to a value of 5.55556E-5 seconds, which in this case, is a one degree crank angle rotation.

9. Setting Stopping Criteria:

   Set the stopping criteria to terminate the simulation after three engine cycles.

10. Plotting the Mass Flow for Each Cylinder:

    Create a mass flow report, monitor, and plot to analyze the fluid flow in and out of each cylinder.

11. Plotting the Mass Fraction of Fresh Air at Cylinder 1:

    Create a surface average report, monitor, and plot for the mass fraction of fresh air at Cylinder 1.

12. Setting up a Vector Scene:

    Create a vector scene that displays the velocity field at the mass flow boundaries.

Contents

Importing the Volume Mesh
Visualizing the Volume Mesh
Selecting Physics Models
Importing the GT-POWER Material Properties
Setting Up the Species Representation
Setting the Reference Pressure
Setting Initial Conditions
Defining Coupled Boundaries
Setting the Time Step
Setting Stopping Criteria
Plotting the Mass Flow for Each Cylinder
Plotting the Mass Fraction of Fresh Air at Cylinder 1
Setting up a Vector Scene

Importing the Volume Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Import the `manifoldGT-POWER.ccm` volume mesh from your working directory.

Note: When creating this volume mesh, the Extruder was used to generate orthogonal cells in each averaging zone. This mesh helps to capture the one-dimensional nature of the flow in these zones.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.
4. Save the simulation as `manifold-gt.sim`.

Visualizing the Volume Mesh

Display the imported volume mesh and store a convenient view.

1. Create a mesh scene.
   A new scene, *Mesh Scene 1*, is displayed in the *Graphics* window.
2. Click (**Save-Restore-Select views**) and select **Views > -X+Y-Z > Up +Y**.
3. Position the geometry as shown below.



4. Click (**Save-Restore-Select views**) and select **Store Current View**, which saves this view for later use.
5. Save the simulation.

Selecting Physics Models

Select the physics models that solve for the fluid in STAR-CCM+. Also, activate a physics model that enables coupling between STAR-CCM+ and GT-POWER.

Physics models determine the behavior of the flow in the simulation. In this tutorial, activate the physics models for the fluid region and the GT-POWER Coupling model.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group box	Model
Enabled Models	**Three-Dimensional** (selected automatically)
Time	**Implicit Unsteady**

Group box	Model
Material	**Multi-Component Gas**
Reaction Regime	**Non-reacting**
Flow	**Segregated Flow**
Enabled Models	**Segregated Species** (selected automatically)
	Gradients (selected automatically)
Equation of State	**Ideal Gas**
Energy	**Segregated Fluid Temperature**
Viscous Regime	**Turbulent**
Enabled Models	**Reynolds-Averaged Navier-Stokes** (selected automatically)
Reynolds-Averaged Turbulence	**K-Epsilon Turbulence**
Enabled Models	**Realizable K-Epsilon Two-Layer** (selected automatically)
Enabled Models	**Two-Layer All y+ Wall Treatment** (selected automatically)
Optional Models	**Co-Simulation**
Co-Simulation	**GT-POWER Coupling**
Optional Models	**Passive Scalar**

2. Click **Close**.
3. To review the models, open the **Physics 1 > Models** node.



4. Save the simulation.

Importing the GT-POWER Material Properties

Create a material database and import the material properties from the GT-POWER simulation.

When preprocessing the GT-POWER solver, you created a file, `CFDCoupling_Setup.chemkin_GT.dbs`, that contains a database of material properties from the GT-POWER simulation. You can import these properties into STAR-CCM+, so that you use identical materials in both codes.

To import the GT-POWER material properties:

1. Right-click the **Tools > Material Databases** node and select **New Material Database**.
2. Rename the **matl** node as `GT-POWER Materials`.

3. Right-click the **GT-POWER Materials** node, and select **Import Properties**.
4. In the *Open* dialog, navigate to the `caseFilesGT-POWER` directory and select `CFDCoupling_Setup.chemkin_GT.dbs`.
5. In the *Import Properties* dialog, select **Import properties for all materials** and click **OK**.
6. Open the **GT-POWER Materials** > **Gases** node.

   The GT-POWER material properties appear as shown below:



Setting Up the Species Representation

Define gas components and passive scalars under the relevant physics models.

To reduce computational costs, you use six mixture components and two passive scalars to represent the species that GT-POWER transfers to STAR-CCM+.

The passive scalars represent species that are only found in small amounts (that is, less than 1% of the volume average). Therefore, it is safe to assume that they have little influence on the bulk fluid and can be modeled as passive.

The gas components and passive scalars are summarized below:

Gas Components	Passive Scalars
n2-vap (Fresh Nitrogen)	indolene-vap (Gaseous Fuel)
o2-vap (Fresh Oxygen)	indolene-combust (Liquid Fuel)

Gas Components	Passive Scalars
lump_N2 (Burnt Nitrogen and lumped, trace species)	
prod_o2 (Burnt Oxygen)	
prod_co2	
prod_h2o	

1. Defining Gas Components
2. Defining Passive Scalars

Contents

Defining Gas Components
Defining Passive Scalars

Defining Gas Components

Define and reorder the gas components, and change the specific heat so that it uses the thermodynamic polynomial data from the GT-POWER materials database file.

In the following steps, you define and reorder the gas components so that the material with the largest expected volume fraction, in this case n2-vap, is last in the list. This ordering improves numerical accuracy, as described in Ordering of Mixture Components.

To define the gas components:

1. Expand the **Continua** > **Physics 1** > **Models** node.
2. Right-click **Multi-Component Gas** > **Gas Components** and select **Select Mixture Components**.
3. In the *Select Mixture Components* dialog, expand the **Material Databases** > **GT-POWER Materials** > **Gases** node.
4. Activate the following components:
   - lump_N2
   - n2-vap
   - o2-vap
   - prod_co2
   - prod_h2o
   - prod_o2
5. Click **Apply** then **Close**.

To reorder the gas components:

6. Right-click the **Gas Components** node and select **Reorder Mixture Components**.
7. Click and drag the components so that they are in the following order:



Change the specific heat so that it uses the thermodynamic polynomial data from the GT-POWER materials database file, as specific heat is defined this way in GT-POWER.

8. For each gas component, select the **Material Properties > Specific Heat** node, and set *Method* to **Thermodynamic Polynomial Data**.



9. Save the simulation.

Defining Passive Scalars

Define passive scalars to represent species that are only found in small amounts (that is, less than 1% of the volume average).

To define passive scalars:

1. Right-click the **Models** > **Passive Scalar** > **Passive Scalars** node and select **New**.
2. Rename the **Passive Scalar** as `indolene-vap`.
3. Create another passive scalar called `indolene-combust`.

   The nodes corresponding to these passive scalars are as shown below.



4. Save the simulation.

Setting the Reference Pressure

Set the reference pressure to the value obtained from the independent GT-POWER run.

STAR-CCM+ uses a reference pressure to reduce round-off error, with all other pressures in the simulation defined relative to its value. However, in GT-POWER, all pressures are absolute.

To minimize the numerical round-off, set the reference pressure to the GT-POWER pressure at the start of coupling. The reference pressure that you use in this tutorial comes from the independent GT-POWER run at the start of this tutorial.

To set the reference pressure:

1. Select the **Reference Values** > **Reference Pressure** node and set *Value* to `2.13 bar`.



Setting Initial Conditions

Set the initial conditions to match the conditions in the GT-POWER model at the start of coupling. This speeds up convergence at the start of the STAR-CCM+ simulation.

The initial conditions that you use in the following steps come from the independent GT-POWER run at the start of this tutorial.

As stated in the previous section, pressure values in STAR-CCM+ are relative to the reference pressure. Therefore, you set the initial pressure to 0.0 bar, as you have already set the reference pressure accordingly.

To set the initial pressure, temperature, and turbulence settings:

1. Edit the **Continua** > **Physics 1** > **Initial Conditions** node, and set the following properties:

Node	Property	Setting
Pressure > **Constant**	*Value*	`0 bar`
Static Temperature > **Constant**	*Value*	`315 K`
Turbulence Specification	*Method*	**Intensity + Length Scale**

Node	Property	Setting
Turbulence Intensity > Constant	*Value*	`0.1`
Turbulence Length Scale > Constant	*Value*	`1 mm`

To set the initial mass fractions:

2. Select the **Initial Conditions > Species Mass Fraction > Constant** node, and set *Value* as shown below:

Gas Component	Mass Fraction
n2-vap	`0.61`
o2-vap	`0.18`
lump_N2	`0.15`
prod_co2	`0.03`
prod_h2o	`0.02`
prod_o2	`0.01`

3. Save the simulation.

Defining Coupled Boundaries

Define the boundaries available for the transfer of data between STAR-CCM+ and GT-POWER.

Set all flow boundaries as coupled with GT-POWER, to make these boundaries available for the transfer of data between STAR-CCM+ and GT-POWER.

As GT-POWER supplies all the boundary conditions at a coupled boundary, no other boundary specifications are required.

To define the coupled boundaries:

1. Open the **Regions > Manifold > Boundaries** node.
2. Multi-select the **Cyl1**, **Cyl2**, **Cyl3**, **Cyl4**, **EGR**, and **Intake** nodes, and set *Type* to **Mass Flow Inlet**.



3. Right-click any selected boundary and select **Edit**.
4. Set the following properties:

Node	Property	Setting
Physics Conditions		
1D Coupling Specification	*1D Coupling Specification*	**Coupled**
Turbulence Specification	*Method*	**Intensity + Length Scale**

Node	Property	Setting
Physics Values		
Turbulence Intensity	*Method*	**Constant**
Constant	*Value*	`0.05`

5. Click **Close**.
6. Edit the **Intake** node and set the following properties:

Node	Property	Setting
Physics Values > Turbulent Length Scale	*Method*	**Constant**
Constant	*Value*	`0.005 m`

7. Edit the **EGR** node and set the following properties:

Node	Property	Setting
Physics Values > Turbulent Length Scale	*Method*	**Constant**
Constant	*Value*	`0.0032 m`

8. Multi-select the **Cyl1**, **Cyl2**, **Cyl3**, and **Cyl4** nodes, right-click one of them and select **Edit**.
9. Set the following properties:

Node	Property	Setting
Physics Values > Turbulent Length Scale	*Method*	**Constant**
Constant	*Value*	`0.0035 m`

10. Save the simulation.

Setting the Time Step

Set the time step to a value of 5.55556E-5 seconds, which in this case is a one degree crank angle rotation.

The time step that GT-POWER uses is not constant, but the STAR-CCM+ time step can be greater than the GT-POWER time step during the simulation. This difference is possible because GT-POWER

can run several time steps while STAR-CCM+ is running one time step. However, for numerical reasons, it is recommended that the STAR-CCM+ time step is as close as possible to the GT-POWER time step. As a general guideline, set the STAR-CCM+ time-step size to no more than three times the GT-POWER time-step size.

It is recommended you round up the STAR-CCM+ time-step, so the elapsed duration of one cycle in STAR-CCM+ slightly exceeds the actual duration of the engine cycle. This rounding ensures the GT-POWER solver completes the engine cycle before STAR-CCM+ terminates it.

To set the time step:

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `5.55556E-5` s.

Setting Stopping Criteria

Set the stopping criteria to terminate the simulation after three engine cycles.

As the engine is running at 3000 rpm, each time step is a one degree crank angle rotation. Therefore, for three engine cycles, 2160 time steps are required. However, you run the simulation for a few extra time steps, as rounding errors in the time step duration can cause 720 time steps to be slightly less than a full cycle. When an incomplete cycle occurs, the GT-POWER solver is stopped prematurely and output files are not generated correctly. Therefore, you use a total of 2162 time steps.

Also, the STAR-CCM+ solver runs iterations between time steps to solve the effect of transient boundary conditions. In this case, 5 inner iterations are adequate to guarantee convergence between time-steps.

To set the stopping criteria:

1. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Steps	*Maximum Steps*	`2162`
Maximum Inner Iterations	*Maximum Inner Iterations*	`5`

2. Save the simulation.

Plotting the Mass Flow for Each Cylinder

Create a mass flow report, monitor, and plot to analyze the fluid flow in and out of each cylinder.

To plot the mass flow for each cylinder:

1. Right-click the **Reports** node and select **New Report > Mass Flow**.
2. Rename the **Mass Flow 1** node to `Mass Flow Cyl1`.
3. Select the **Reports > Mass Flow Cyl1** node and set *Parts* to **Cyl1**.
4. Similarly, create three more reports to monitor mass flow in cylinders two, three and four, with the following names and boundaries:

Report	Boundary
Mass Flow Cyl2	**Cyl2**
Mass Flow Cyl3	**Cyl3**
Mass Flow Cyl4	**Cyl4**



5. Multi-select the **Mass Flow Cyl1**, **Mass Flow Cyl2**, **Mass Flow Cyl3**, and **Mass Flow Cyl4** nodes.
6. Right-click one of the selected nodes and select **Create Monitor and Plot from Report**.
7. In the *Create Plot From Reports* dialog, click **Single Plot**.
8. Save the simulation.

Plotting the Mass Fraction of Fresh Air at Cylinder 1

Create a surface average report, monitor, and plot for the mass fraction of fresh air at Cylinder 1.

First, define a field function to calculate the fresh air mass fraction:

1. Right-click the **Tools > Field Functions** node and select **New > Scalar**.
2. Rename the **User Field Function 1** node to `Fresh Air`.
3. Edit the **Fresh Air** node and set the following properties:

Property	Setting
Function Name	`FreshAir`
Definition	`${MassFractionn2-vap}+${MassFractiono2-vap}`

The braces in the definition are required as the field function names use hyphens, which the code interprets as a minus operation.

To plot the mass fraction of fresh air at cylinder 1:

4. Right-click the **Reports** node and select **New Report > Surface Average**.
5. Rename the **Surface Average 1** node to `Fresh Air`.
6. Edit the **Fresh Air** node and set the following properties:

Property	Setting
Parts	**Cyl1**
Scalar Field Function	**Fresh Air**

7. Right-click the **Fresh Air** report and select **Create Monitor and Plot from Report**.
8. Save the simulation.

Setting up a Vector Scene

Create a vector scene that displays the velocity field at the mass flow boundaries.

To set up a vector scene for the velocity field:

1. Create a vector scene.
2. Edit the **Displayers** node and set the following properties:

Node	Property	Setting
Outline 1	*Surface*	**Activated**
	Opacity	`0.3`
Vector 1 > Parts	*Parts*	**Cyl1**
		Cyl2
		Cyl3
		Cyl4
		EGR

Node	Property	Setting
		Intake

3. Click (**Save-Restore-Select views**) and select **Restore View > View 1**.



4. Save the simulation.

Coupling STAR-CCM+ to GT-POWER

Couple the STAR-CCM+ and GT-POWER simulations.

To couple the simulations, the following steps are required:

1. Starting the GT-POWER solver:

   The first step in coupling STAR-CCM+ to GT-POWER is to define the directory and case name of the GT-POWER model. You can then start GT-POWER, which launches the GT-POWER solver in server mode.

2. Assigning gas components and passive scalars to GT-POWER species:

   Assign a STAR-CCM+ mixture component or passive scalar to each GT-POWER species.

3. Assigning boundaries to GT-POWER zones:

Assign STAR-CCM+ boundaries to the corresponding zones in GT-POWER.

4. Visualizing the volume-averaging zones:

   The **GT-POWER Coupling** feature provides the average zone flag field function, as a means of visualizing the volume where physics values and mass fractions are averaged and sent to GT-POWER. Inspect the zone and determine whether the flow is relatively one dimensional.

Contents

Starting the GT-POWER Solver
Assigning Gas Components and Passive Scalars to GT-POWER Species
Assigning Boundaries to GT-POWER Zones
Visualizing the Volume-Averaging Zones

Starting the GT-POWER Solver

Define the GT-POWER co-simulation session properties and start the GT-POWER solver.

The first step in coupling STAR-CCM+ to GT-POWER is to define the directory and case name of the GT-POWER model. You can then start GT-POWER, which launches the GT-POWER solver in server mode.

Note: If you intend to couple to a different version of GT-POWER, set the *GT-SUITE Version* property.

To start the GT-POWER solver:

1. Select the **Co-Simulations** > **GT-POWER 1** node.
2. Set *Directory Name* to `caseFilesGT-POWER`.

   Directory Name must be set to either the relative or absolute path to the directory that contains the GT-POWER case files. In this tutorial, *Directory Name* is specified relative to the current working directory (that is, the directory from where STAR-CCM+ was launched), according to the file structure prepared in Preparing the File Structure. If you encounter issues when coupling, the most likely cause is an incorrect *Directory Name*.
3. Set *Case Name* to `CFDCoupling_Setup`.



4. Right-click the **GT-POWER 1** node and select **Start GT-POWER**.



The *Output* window displays `GT-POWER Initialized` when the process is complete.

At this stage, the GT-POWER solver is started in the background, waiting for STAR-CCM+ to begin running its solver. Clearing the

solution stops the GT-POWER solver and breaks the connection to STAR-CCM+. If you clear the solution, restart the GT-POWER solver as outlined above.

The **GT-POWER Species** and **GT-POWER Zones** nodes contain the species and zones from the GT-POWER case.



5. Save the simulation.

Assigning Gas Components and Passive Scalars to GT-POWER Species

Assign a STAR-CCM+ mixture component or passive scalar to each GT-POWER species.

The GT-POWER model determines the species that are passed to STAR-CCM+. Similarly, the GT-POWER setup handles the mapping of species from STAR-CCM+ back to GT-POWER.

Note: Assign all species from GT-POWER to a unique gas component or passive scalar.

To assign gas components and passive scalars to GT-POWER species:

1. Select the **GT-POWER 1 > GT-POWER Species > indolene-combust** node and set *Mixture Component/Passive Scalar* to **indolene-combust**.



2. Similarly, assign the remaining species to their counterpart gas components and passive scalars with the same name.
3. Save the simulation.

Assigning Boundaries to GT-POWER Zones

Assign STAR-CCM+ boundaries to the corresponding zones in GT-POWER.

This setting defines the location of data exchange between the one-dimensional and three-dimensional codes. The **GT-POWER Zones** node presents a list of STAR-CCM+ boundaries on which you have set **1D Coupling Specification** to coupled.

Note: When selecting GT-POWER zones, the `(inlet)` next in the zone name refers to the type of boundary that is used in the GT-POWER simulation. Fluid can flow both in and out of any coupled boundary.

To assign boundaries to GT-POWER zones:

1. Select the **Co-Simulations** > **GT-POWER 1** > **GT-POWER Zones** > **Cyl1 (inlet)** node and set *Boundary* to **Cyl1**.
2. Similarly, assign the remaining zones to their counterpart boundaries with the same name.
3. Save the simulation.

Visualizing the Volume-Averaging Zones

Visualize the zones where the flow is approximately one dimensional.

The **GT-POWER Coupling** feature provides the average zone flag field function, as a means of visualizing the volume where physics values and mass fractions are averaged and sent to GT-POWER. Inspect the zone and determine whether the flow is relatively one dimensional.

To visualize the volume-averaging zone:

1. Create a scalar scene.
2. Edit the **Scalar Scene 1** > **Displayers** > **Scalar 1** node and set the following properties:

Node	Property	Setting
Outline 1	*Surface*	**Activated**
	Mesh	**Activated**
Scalar 1		
Parts	*Parts*	**Cyl1**
		Cyl2
		Cyl3

Node	Property	Setting
		Cyl4
		EGR
		Intake
		Wall
Scalar Field	*Function*	**AvgZoneFlag**
	Min	1
	Max	6
Color Bar	*Levels*	6

3. Click (**Save-Restore-Select views**) and restore **View 1**.
4. Initialize the Solution.

   A scene similar to the one shown below appears.



   The scalar values, which relate to which boundary was created first in STAR-CCM+, are of no significance and simply highlight the individual zones.
5. Save the simulation.

Running the Simulation

Run the simulation and monitor the iterative progress output.

The simulation is now ready to be run.

1. Run the simulation.

   The iterative progress can be monitored in the *Output* window. The mass flow and mass fraction monitor plots show the state of the simulation up to the current time-step. The velocity vector scene shows the flow in and out of the boundaries at the current time-step.

   There is a slight delay before any iterations are executed, as the GT-POWER solver runs for 20 cycles before sending data to STAR-CCM+.

Note: In your own cases, activate the Autosave facility in STAR-CCM+ so you can easily perform a restart if the solver is unintentionally stopped or more cycles are required for convergence.

Visualizing the Results

Visualize the flow velocity field and the mass flow throughout the manifold.

The engine cycle in GT-POWER starts when the Cylinder 1 Intake Valve closes, which the model defines as 110 degrees before Cylinder 1 Top Dead Center Firing (TDCF). Therefore, the results in STAR-CCM+ display the engine cycle from 110 degrees before Cylinder 1 TDCF, over a further 2160 degrees of crank rotation (three cycles).

A velocity vector scene similar to the one shown below is displayed, which shows velocity vectors during the last time-step. It shows that GT-POWER has opened the valve on Cylinder 3, with high velocity air flow.



To view the mass flow plot over the last full engine cycle:

1. Right-click the **Plots** > **Reports Plot** and select **Open**.
2. Click **scene/plot**.
3. Select the **Axes** > **X Axis** node and set *Minimum* to `0.08` and *Maximum* to `0.12`.



4. Similarly, view the fresh air mass fraction monitor over the last full engine cycle.



Summary

This tutorial has introduced the following STAR-CCM+ features:

- Preparing the GT-POWER model for coupling to STAR-CCM+.
- Obtaining initial conditions from an independent run of GT-POWER.
- Setting up the STAR-CCM+ simulation for coupling:
  - Selecting physics models for coupling to GT-POWER
  - Importing GT-POWER material properties
  - Setting up the species representation
  - Defining coupled boundaries
  - Setting solver and stopping criteria
- Coupling STAR-CCM+ to GT-POWER
- Visualizing the solution

OLGA Co-Simulation: 1D Coupling

The purpose of this tutorial is to show the typical workflow in running a coupled simulation with the Windows-based Oil & Gas analysis code OLGA. As OLGA uses a one-dimensional method for predicting flow within oil lines, it requires much shorter computational run times when compared to a multi-dimensional analysis of the same system. In this tutorial, STAR-CCM+ calculates the three-dimensional flow through a 90 degree bend, which forms part of an overall pipeline analysis simulated by OLGA.

The OLGA coupling feature exchanges species volume fractions and physics values between STAR-CCM+ and OLGA. Data is exchanged either once per time step, or once per inner iteration, depending on the solver properties. If flow is entering the STAR-CCM+ region, the physics values and mass fractions are applied stratified at the boundary. If flow is leaving the STAR-CCM+ region, the physics values and volume fractions are averaged across a zone and sent to OLGA.

The following data is exchanged:

- Mass Flux
- Pressure (Pressure outlet only)
- Temperature (Pressure outlet only)
- Eulerian Component Velocities (Velocity inlet only)
- Eulerian Component Densities (Velocity inlet only)

This tutorial concentrates on the STAR-CCM+ side of the simulation. Initial conditions were obtained from an independent run of OLGA (not discussed here). A velocity inlet is used to couple to an OLGA boundary, and a pressure outlet is used to couple to an OLGA source.

Contents

Prerequisites
Preliminaries
Importing the Volume Mesh
Selecting Physics Models
Defining Eulerian Phases
Setting Reference Values and Initial Conditions
Setting Boundary Conditions
Setting Solver Parameters
Setting Stopping Criteria

Assigning Eulerian Phases to OLGA Species
Assigning Boundaries to OLGA Zones
Coupling STAR-CCM+ to OLGA
Reporting, Monitoring and Plotting
Setting up a Visualization Scene
Running the Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the OLGA Co-Simulation - 1D Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

This tutorial describes the procedure using OLGA v5.3.2.

Preliminaries

Make a note of the following system-specific details, which are required later:

- The host name of the machine for the OLGA run. For example: `mercury`.
- The directory where the OLGA executable is installed. For example: `C:\Program Files\SPT Group\OLGA 5.3.2\OlgaExecutables_5_3_2\olga-5.3.2.exe`.
- The OLGA **service name** and **port number**, which can be found using the following steps:
  1. Navigate to `C:\WINDOWS\system32\drivers\etc\`.
  2. Open the `service` file with a text editor.

The service names and port numbers are available depending on the number of licenses. Typically, they follow the following format:

Service Name	Port Number
olga	12340
olga_1	12341
olga_2	12342

Make sure that the chosen service is not currently in use.

To successfully couple to OLGA, the following files are required:

- A STAR-CCM+ mesh representing the three-dimensional region (`90degBend.ccm`).
- OLGA server batch file (`olgaServer.bat`), which runs the OLGA solver in batch mode and waits for input from STAR-CCM+.
- OLGA input file (`OLGAInput.geninp`), representing an oil pipe line, and providing the transient boundary conditions in STAR-CCM+. Locate this file on the same machine as `olgaServer.bat`.
- OLGA properties table (`3phase.tab`), which contains information on the physical properties of the OLGA species. Locate this file in the same directory as `olgaServer.bat` and the STAR-CCM+ simulation.

Importing the Volume Mesh

To set up the STAR-CCM+ simulation, launch a simulation and import the supplied volume mesh.

1. Launch STAR-CCM+.
2. Start a simulation.
3. Select **File > Import > Import Volume Mesh...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
5. Select `90degBendOLGA.ccm` then click **Open**.

   STAR-CCM+ provides feedback on the import process in the *Output* window and a mesh region is created under the **Regions** node.

   You can visualize the imported volume mesh with a method of your choice, such as creating a mesh scene. The mesh appears as shown below.



6. Save the simulation as `olgaCoupling.sim`.

Selecting Physics Models

Select physics models to define the behavior of the fluid, and to allow coupling between STAR-CCM+ and OLGA.

To select the physics models:

1. For the physics continuum, **Continua** > **Physics 1**, select the following models in order:

Group Box	Model
Enabled Models	**Three-Dimensional** (Selected automatically)
	Gradients (Selected automatically)
Time	**Implicit Unsteady**
Material	**Eulerian Multiphase**
Enabled Models	**Multiphase Interaction** (Selected automatically)
Multiphase Flow Model	**Volume of Fluid (VOF)**
Enabled Models	**Segregated Flow** (Selected automatically)
	Multiphase Equation of State (Selected automatically)
Viscous Regime	**Turbulent**

Group Box	Model
Enabled Models	**Reynolds-Averaged Navier-Stokes** (Selected automatically)
Reynolds-Averaged Turbulence	**K-Omega Turbulence**
Enabled Models	**SST (Menter) K-Omega** (Selected automatically)
	All y+ Wall Treatment (Selected automatically)
Optional Models	**Co-Simulation**
Co-Simulation	**OLGA Coupling**
Enabled Models	**Gravity** (Selected automatically)
	Segregated Multiphase Temperature (Selected automatically)

2. Click **Close**.

   The color of the **Physics 1** node has turned from gray to blue to indicate that fluid models have been activated and that the continuum is valid.

3. To review the models, expand the **Physics 1** > **Models** node.



4. Save the simulation.

Defining Eulerian Phases

Define eulerian phases in STAR-CCM+ to represent the OLGA species.

OLGA uses three species to represent the oil, water and gas mixture within the one-dimensional simulation. Three Eulerian phases represent these species in the STAR-CCM+ three-dimensional simulation.

To define the **Gas** phase:

1. Expand the **Continua** > **Physics 1** > **Models** > **Eulerian Multiphase** node.
2. Right-click the **Eulerian Phases** node and select **New**.

   A new node, **Phase 1**, is added within the **Eulerian Phases** node.
3. For the eulerian phase, **Phase 1**, select the following models:

Group box	Models
Enabled Models	**Turbulent** (selected automatically)
	Energy (selected automatically)
Materials	**Gas**
Enabled Models	**OLGA EOS** (selected automatically)

4. Click **Close**.

To define the gas phase material:

5. Expand the **Phase 1 > Models** node.
6. Right-click the default gas material, **Air**, and select **Replace with...**



7. In the *Replace Material* dialog, expand the **Material Databases > Standard > Gases** node.
8. Select **CH4 (Methane)**, as it has similar material properties to the OLGA gas mixture species.
9. Click **OK**.

The OLGA solver defines several of the material properties. Modify the remaining properties, to provide a more accurate representation of the OLGA species:

10. Expand the **CH4 > Material Properties** node.



OLGA defines the following properties:

- Density
- Density pressure derivative
- Density temperature derivative
- Dynamic Viscosity
- Thermal Conductivity

11. Edit the **Material Properties** node, and set the remaining properties as follows:

Node	Property	Setting
Molecular Weight > Constant	*Value*	`17.1793 kg/kmol`
Specific Heat > Constant	*Value*	`2934.729 J/kg-K`
Speed of Sound > Constant	*Value*	`347.0 m/s`

The values were sourced from OLGA property tables.

12. To denote the species it represents in the coupled simulation, rename the phase as `OLGA Gas`.

    The object tree appears as shown below when complete.



13. Create another phase representing the OLGA Oil phase, using a similar process as shown above. For this phase, use the following settings:

Step	Setting
Phase Name	OLGA Oil
Model selection	Liquid
Material selection	C12H26 (Dodecane)

Step	Setting
Molecular Weight	122.182 kg/kmol
Specific Heat	2027.008 J/kg-K
Speed of Sound	1500.0 m/s

14. Create another phase representing the OLGA Water phase, using the following settings:

Step	Setting
Phase Name	OLGA Water
Model selection	Liquid
Material selection	H20 (Water)
Molecular Weight	18.01534 kg/kmol
Specific Heat	4821.635 J/kg-K
Speed of Sound	1500.0 m/s

15. Save the simulation.

Setting Reference Values and Initial Conditions

Set up the reference values and initial conditions of the continuum. This process defines the physical properties and multiphase composition in the pipe at the start of the simulation.

The reference pressure is used in STAR-CCM+ to reduce the effect of numerical round-off, with all other pressures in the simulation defined relative to the reference pressure. However, in OLGA, all pressures are absolute. To reduce the risk of error, set the reference pressure value to the average pressure in OLGA. The reference pressure of 80 bar used in this example has been obtained from an independent run of OLGA. Also, set gravity to act in the -y direction.

To set the reference values:

1. Edit the **Continua** > **Physics 1** > **Reference Values** node.
2. Set the following properties:

Node	Property	Setting
Reference Pressure	*Value*	`80.0 bar`
Gravity	*Value*	`[0.0, -9.81, 0.0] m/s^2`

As initial conditions, specify the volume fraction of the OLGA phases, the initial temperature, and turbulence settings.

Initialize the continua as 100% gas by setting the OLGA Gas volume fraction to 1, to provide an interesting visualization of the flow through the three-dimensional section. This approach also gives better convergence, as the gas is replaced in the region more easily than if oil or water were used.

To set the initial volume fraction:

3. Expand the **Continua** > **Physics 1** > **Initial Conditions** node.
4. Select the **Volume Fraction** > **Constant** node, and set *Value* to `[1.0, 0.0, 0.0]`.



Temperature is an absolute value in both STAR-CCM+ and OLGA, with the value set as 20.0 C. As turbulence settings, use the values defined at the coupled boundaries, which in this example are a turbulence intensity of 0.05, and a turbulent viscosity ratio of 50.

To set the initial temperature and turbulence settings:

5. Edit the **Initial Conditions** node, and set the following properties:

Node	Property	Setting
Initial Conditions		
Static Temperature > Constant	*Value*	`20.0 C`
Turbulence Intensity > Constant	*Value*	`0.05`

Node	Property	Setting
Turbulent Viscosity Ratio > Constant	*Value*	`50`

6. Save the simulation.

Setting Boundary Conditions

Set the three-dimensional model wall boundaries to have the same properties as the OLGA model. Also, couple the inlet and outlet boundaries with OLGA to allow the OLGA solver to define the transient boundary conditions.

To set the wall boundary conditions:

1. Edit the **Regions** > **90degBend** > **Boundaries** > **Wall** node, and set the following properties:

Node	Property	Setting
Physics Conditions		
Thermal Specification	*Condition*	**Convection**
Physics Values		
Ambient Temperature > Constant	*Value*	`40.0 C`
Heat Transfer Coefficient > Constant	*Value*	`200 W/m^2-K`

Couple the inlet and outlet boundaries with OLGA to allow the OLGA solver to define the transient boundary conditions. In all cases, couple a STAR-CCM+ velocity inlet boundary to an OLGA boundary, and a STAR-CCM+ pressure boundary to an OLGA source.

Following these criteria, set the boundary types as follows:

2. Edit the **Regions** > **90degBend** > **Boundaries** node, and set the following properties:

Node	Property	Setting
OLGA OUTLET	*Type*	**Velocity Inlet**

Node	Property	Setting
OLGA SOUR-2	*Type*	**Pressure Outlet**

Set the inlet and outlet boundary conditions to allow direct coupling to OLGA, and also specify turbulence values:

3. Multi-select the **OLGA OUTLET** and **OLGA SOUR-2** nodes.
4. Right-click one of the selected nodes, select **Edit**, and set the following properties:

Node	Property	Setting
Physics Conditions > 1D Coupling Specification	*1D Coupling Specification*	**Coupled**
Physics Values		
Turbulence Intensity > Constant	*Value*	`0.05`
Turbulent Viscosity Ratio > Constant	*Value*	`50`

5. Click **Close**.
6. Save the simulation.

Setting Solver Parameters

Set the Implicit Unsteady solver time-step to be smaller than the average time-step in OLGA.

Generally, use a smaller time-step in STAR-CCM+ than in OLGA. Since the time step in an OLGA simulation is not constant, use the average OLGA time step as a reference value. In this example, the average time step from an independent run of OLGA is 0.002983 seconds, and the time step in STAR-CCM+ is set to 0.0025 seconds.

1. Select the **Solvers > Implicit Unsteady** node and set *Time-Step* to `0.0025 s`.

Setting Stopping Criteria

As stopping criteria, set the maximum physical time and the maximum inner iteration.

To keep the processing time low, while still providing a detailed evaluation of multiphase flow with OLGA coupling, set the simulation to run for 4 s.

1. Edit the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Physical Time	*Maximum Physical Time*	`4.0 s`
Maximum Steps	*Enabled*	**Deactivated**
Maximum Inner Iterations	*Maximum Inner Iterations*	`10`

2. Save the simulation.

Assigning Eulerian Phases to OLGA Species

Assign the OLGA species to the eulerian phases that you defined within the Eulerian Multiphase model node.

1. Select the **Co-Simulations > OLGA5 > OLGA Species > Gas** node and set *Eulerian Phase* to **OLGA Gas**.



2. Similarly, assign the remaining OLGA species to the corresponding eulerian phases:

OLGA Species	Eulerian Phase
Oil	**OLGA Oil**
Water	**OLGA Water**

Assigning Boundaries to OLGA Zones

Add OLGA zones and assign them to STAR-CCM+ boundaries.

During the solution, data is exchanged between the coupled zones in OLGA and boundaries in STAR-CCM+.

First, add an OLGA zone to couple the STAR-CCM+ velocity inlet boundary to the OLGA boundary. To add an OLGA zone:

1. Right-click the **OLGA5 > OLGA Zones** node, and select **Add New OLGA Zone**.



A new node, **NewOlgaZone**, is added within the **OLGA Zones** node.

To define the first OLGA zone:

2. Rename the **NewOlgaZone** node to `OUTLET`.
3. Enter the following settings, which must match exactly and are case-sensitive:

Property	Setting
OLGA Pos Name	`POS-OUTLET`
Boundary	**OLGA OUTLET**



Create another zone for the OLGA source, and assign the OLGA source to the STAR-CCM+ pressure boundary.

4. Within the **OLGA Zones** node, add a new OLGA zone and rename it to `SOUR-2`.
5. Set the following properties:

Property	Setting
OLGA Pos Name	`POS-INLET1`
Boundary	**OLGA SOUR-2**



6. Save the simulation.

Coupling STAR-CCM+ to OLGA

Couple the STAR-CCM+ simulation to the OLGA server.

The procedure for coupling to OLGA is outlined below:

1. Start the OLGA server.
2. Define the host name and port number.
3. Define the input and property files.
4. Connect to the OLGA server.

Note: The examples in the following sections are for illustration only. Use the specific details for your system, as described earlier.

Contents

Start the OLGA Server

Define the Host Name, Port Number and Files
Connect to the OLGA Server

Start the OLGA Server

Launch the OLGA server in batch mode.

The `olgaServer.bat` file, which was copied to the working directory earlier, contains the following information:

```
C:\"Program Files"\"SPT Group"\"OLGA
5.3.2"\OlgaExecutables_5_3_2\olga-5.3.2.exe -server
[olga]
olgaServer.bat
```

Modify the text in bold for your OLGA installation directory and version. Change the text within the square brackets to match the OLGA service name.

To launch the OLGA server:

1. Make sure that `OLGAInput.geninp` is on the same machine as `olgaServer.bat`.
2. Make sure that `3phase.tab` is in both the same directory as `olgaServer.bat` and the STAR-CCM+ simulation file.
3. Edit the `olgaServer.bat` file as required.
4. Run `olgaServer.bat` in the usual Windows manner (single or double-click depending on system set-up).

The following command prompt window opens:



The OLGA solver is now running in batch mode, waiting for input from STAR-CCM+. When the solution is obtained, information from the OLGA solver is printed in this window.

Note: This batch file runs the OLGA server continuously, which uses an OLGA license. Therefore, to release the license, terminate this batch file.

Define the Host Name, Port Number and Files

To allow STAR-CCM+ to communicate with the OLGA solver, set the host name and port number. Also, set the input and property files, which are used as input when STAR-CCM+ runs the OLGA solver.

To set the host name and port number:

1. Select the **Co-Simulations** > **OLGA5** node, and set *Host Name* to the name or IP address of the machine where the OLGA server is running. If OLGA is running locally, use the name of the local machine.
2. Set *Port Number* according to the service used in the `olgaServer.bat` file.

To set the input and property files:

3. Select the **Co-Simulations** > **OLGA5** node, and set the *Input File* to `OLGAInput.geninp`, including its absolute path on the machine where `olgaServer.bat` is running.
4. Set *Property File* to `3phase.tab`. For this property, an absolute path is not required, as the property file must be in both the same directory as `olgaServer.bat` and the STAR-CCM+ simulation.



Connect to the OLGA Server

Start the OLGA server from the STAR-CCM+ interface.

The STAR-CCM+ simulation is now connected to the OLGA server.

1. Right-click the **OLGA5** node and select **Start OLGA**.



When coupled, the *Output* window displays: `OLGA Initialized`.

Reporting, Monitoring and Plotting

Create reports, monitors, and a plot to monitor the volume fractions of the three phases throughout the transient simulation.

Create a volume average report of the entire region to summarize the composition of the three-phase mixture within the pipe:

1. Right-click the **Reports** node and select **New Report > Volume Average**.
2. Rename the **Volume Average 1** report to `Volume Fraction Oil`.
3. Select the **Volume Fraction Oil** node, and set the following properties:

Property	Setting
Parts	**90degBend**
Scalar Field Function	**Volume Fraction > OLGA Oil**

4. Create a volume averaged, volume fraction report for the **OLGA Gas** and **OLGA Water** phases, using a similar method as above.

   When complete, the tree appears as follows:



These reports are used to create monitors and a plot, which displays the volume fractions of the three-phases throughout the transient simulation:

5. Multi-select the three reports. Right-click one of the selected nodes and select **Create Monitor and Plot from Report**.



6. In the *Create Plots from Reports* dialog, select **Single Plot**.
   Three new monitors are created, and assigned to a single plot.



7. Save the simulation.

Setting up a Visualization Scene

Create a volume fraction scene to display a representation of the three-phase mixture within the fluid region, during the transient simulation. Export images of the scene at each time step, which can be used to create an animation.

1. Create a geometry scene.

Create two isosurfaces to represent the Eulerian phases in the visualization:

2. Right-click the **Derived Parts** node and select **New Part > Isosurface**.
3. In the *Create Isosurface* window, set the following properties:

Property	Setting
Scalar	**Volume Fraction > OLGA Oil**

Property	Setting
isovalue	`0.5`
Display	**New Geometry Displayer**

The *isovalue* is set to 0.5 to create a surface where the volume fraction of oil is 50%.

4. Click **Create**.
5. In the same dialog, create an isosurface for the **OLGA Water** species, with the following settings:

Property	Setting
Scalar	**Volume Fraction > OLGA Water**
isovalue	`0.5`
Display	**New Geometry Displayer**

6. When the creation of both isosurfaces is complete, click **Close** in the *Create Isosurface* window.

Expand the **Geometry Scene 2 > Displayers** node.

7. Rename the **Isosurface Geometry 1** node to `oilIso`, and the **Isosurface Geometry 2** node to `waterIso`.

Modify **Geometry Scene 2** to provide a suitable visual representation of the flow through the pipe:

8. Edit the **Displayers** node, and set the following properties:

Node	Property	Setting
Geometry 1	*Opacity*	`0.3`
Parts	*Parts*	**Regions > 90degBend > Wall**
oilIso	*Color Mode*	**Constant**
waterIso	*Color Mode*	**Constant**

A new node, **Color**, is added within both the **oilIso** and **waterIso** node.

9. Select the **oilIso > Color** node, and click the ellipsis next to the *Color* properties.
10. In the *Color* dialog, select the color that is labeled **51, 51, 51,**, and click **OK**.



11. Repeat the above steps for the **waterIso** node, setting *Color* to **0, 0, 255** (blue).

The isosurface only represents the phase within the fluid region. Therefore, add scalar displayers to the region boundaries to show the phases on the surfaces of the pipe:

12. Within the **Geometry Scene 2** > **Displayers** node, create a scalar displayer.
13. Rename the **Scalar 1** displayer node to `oilScalar`.
14. Expand the **oilScalar** node and set the following properties:

Node	Property	Setting
oilScalar	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Regions**
Scalar Field	*Function*	**Volume Fraction > OLGA Oil**
	Min	`0.5`
	Max	`2.0`
Color Bar	*Levels*	`2`
	Visible	**Deactivated**

Node	Property	Setting
	Color Map	**grayscale**

Setting a relatively high value for the *Max* property ensures the 0.5 to 1.0 volume fraction values for the oil phase are in lower half of the color bar spectrum. The high volume fraction allows for accurate representation of the oil phase on the region boundaries.

15. Similarly, create a scalar displayer named `waterScalar`, with the following settings:

Node	Property	Setting
waterScalar	*Contour Style*	**Smooth Filled**
Parts	*Parts*	**Regions**
Scalar Field	*Function*	**Volume Fraction > OLGA Water**
	Min	`0.5`
	Max	`2.0`
Color Bar	*Levels*	`2`
	Visible	**Deactivated**
	Color Map	**blue-red**

Set the attributes of the scene to export an image of the visualization at every time step:

16. Edit the **Attributes** node and set the following properties:

Node	Property	Setting
Update	*Trigger Type*	**Time-Step**
	Save To File	**Activated**
Time-Step Frequency	*Time-step Frequency*	`10`

Add an annotation to the scene to display the solution time:

17. Expand the **Tools > Annotations** node.
18. Click and drag the **Solution Time** node from the object tree to the *Visualization* window.



19. Save the simulation.

Running the Simulation

Run the simulation and monitor the iterative progress.

The simulation is now ready to be run:

1. Click (**Run**).

The iterative progress can be monitored in the *Output* window. The volume fraction plot displays the composition of the multiphase mixture in the pipe for the duration of the solution. The visualization scene displays the oil and water entering from the velocity inlet at the start of the solution, and flowing towards the pressure outlet. The phases continue to flow through the region for the entire solution.

Visualizing the Results

In the geometry scene that you prepared, visualize the three phases mixture flowing through the pipe. Also, visualize the volume fractions of the oil, gas, and water phases.

When the solver is complete, *Geometry Scene 2* in the *Visualization* window appears as shown below.



The scene shows the oil, water and gas phases flowing through the region. A series of images have also been created in the current working directory, named Scene1_imagexxxxx.png, where xxxxx refers to the time-step number. These images show the flow through the region during the transient solution, and can be used to create an animation.

The volume fractions of the oil, water and gas phases have also been plotted for the duration of the solution, and can be visualized by opening the **Plots** > **Reports Plot** plot.



Summary

This tutorial demonstrated the typical workflow in running a coupled simulation with the Windows-based Oil & Gas analysis code OLGA.

This tutorial has introduced the following STAR-CCM+ features:

- Selecting fluid physics models for coupling to OLGA 5.
- Defining Eulerian phases.
- Assigning Eulerian phases to OLGA Species.
- Assigning Boundaries to OLGA Zones.
- Coupling STAR-CCM+ to OLGA 5.
- Monitor phase volume fractions throughout a transient simulation.
- Setting up visualization scenes.
- Setting solver and stopping criteria for coupling to OLGA.

AMESim Co-Simulation: 1D Coupling

This tutorial demonstrates a typical workflow when running a coupled simulation with AMESim.

With STAR-CCM+ to AMESim co-simulation, you can obtain more accurate data on the amount of force acting on the plate to feed back into your AMESim model. The data is transferred between the two codes every 5.0E-4 seconds. The total simulation time is 0.05 seconds.

STAR-CCM+ solves the fluid flow through the check valve during a transient simulation. The AMESim model simulates a check valve, where a linear spring controls the plate displacement. The following list summarizes the characteristics of the check valve:

- Plate Diameter: 50 mm
- Plate Mass: 0.2 kg
- Plate Higher Displacement Limit: 10.21 mm. This property limits the amount that the plate can move in the specified direction, and simulates the distance to spring compression.
- Spring Rate: 2000 N/m

The following image shows the three-dimensional geometry and highlights the plate valve components:



Index	Description
1	**Inlet**
2	**Overset Region**
3	**Plate**
4	**Plate Support**
5	**Outlet**

For this tutorial, you are provided with a simulation file that contains predefined data. The following items are predefined for you in the simulation file:

- The mesh continuum with mesh settings. To keep the file size to a minimum, the volume mesh is not provided.
- The physics continuum with physics models not related to AMESim.
- The initial conditions (in the physics continuum and at the inlet boundary).
- The regions and boundaries (including the boundary types).
- A mesh scene.

The AMESim Model

With AMESim coupling, you can exchange any field function data between STAR-CCM+ and AMESim. The AMESim model for this tutorial is shown in the following diagram.



In this tutorial, STAR-CCM+ sends the mass flow data to AMESim. AMESim controls the pressure drop across the check valve. In AMESim, the inlet pressure is set to climb from 1 bar to 10 bar for the first 0.01 seconds, and then is set to maintain 10 bar for 0.04 seconds. AMESim maintains the outlet pressure at 1 bar for the entire simulation. AMESim sends the calculated pressure data to STAR-CCM+, where it is applied to the corresponding boundaries (the inlet and outlet). The following diagram illustrates these pressures:



The force that the plate experiences due to the flow is sent from STAR-CCM+ to AMESim. AMESim calculates the displacement of the plate using the force input signal (along with the specified plate mass and spring force), and sends the displacement output signal back to STAR-CCM+. STAR-CCM+ applies the displacement data to the overset region using a user-defined vertex motion. The new position of the plate affects the flow, and the new data is sent back to AMESim. This cycle repeats for each time step.

The following diagram illustrates how data is exchanged between STAR-CCM+ and AMESim in this tutorial:



Index	Description	Data Transfer
1	Outlet mass flow	STAR-CCM+ Field Function to AMESim Pin 1
2	Force on the plate	STAR-CCM+ Field Function to AMESim Pin 2
3	Outlet pressure	AMESim Pin 3 to STAR-CCM+ Outlet Boundary
4	Inlet pressure	AMESim Pin 4 to STAR-CCM+ Inlet Boundary
5	Plate displacement	AMESim Pin 5 to STAR-CCM+ Morpher

Contents

Prerequisites
Loading the Initial Simulation
Generating the Volume Mesh
Selecting Physics Models
Creating Reports for the AMESim Field Function Zones
Starting the AMESim Solver
Defining AMESim Field Function Zones
Coupling to AMESim
Defining Boundary Conditions
Setting Up the User Defined Vertex Motion

Setting the Time Step and Stopping Criteria
Plotting the Plate Displacement
Running the Simulation
Visualizing the Solution on a Vector Scene
Summary

Prerequisites

The instructions in the AMESim Coupling tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

This tutorial describes the procedure using AMESim v12.

Loading the Initial Simulation

For this tutorial, you are provided with a simulation file that contains pre-defined objects.

Load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
5. Select `checkValve-AMESim.sim`, then click **Open**.
6. In the *Load Simulation* dialog, click **OK**.

Generating the Volume Mesh

Generate the volume mesh using the predefined mesh settings.

To generate the mesh:

1. Click (**Generate Volume Mesh**).
   The surface mesh and volume mesh are generated with the predefined settings.

To visualize the mesh:

2. Open *Mesh Scene 1*.
3. Right-click in a blank area of the scene and select **Apply Representation > Volume Mesh**.
   The volume mesh is shown in the following screenshot.



4. Save the simulation.

Selecting Physics Models

In this tutorial, you activate the physics models that enable coupling between STAR-CCM+ and AMESim. All other physics models that are not related to AMESim are preselected in the simulation.

To select the AMESim physics models:

1. For the physics continuum, **Continua > Physics 1**, select the following models in order:

Group Box	Model
Optional Models	**Co-Simulation**
Co-Simulation	**AMESim Coupling**

2. Click **Close**.
3. To review the models, open the **Physics 1** > **Models** node.



The initial conditions for this simulation are predefined.

4. Save the simulation.

Creating Reports for the AMESim Field Function Zones

Create reports for the data exchanged between STAR-CCM+ and AMESim.

Each user-defined report in STAR-CCM+ has a field function that is associated with it. STAR-CCM+ communicates with AMESim through

these field functions. Create the reports for the data that you want to send to AMESim.

To define a report for the mass flow rate through the outlet:

1. Right-click the **Reports** node and select **New Report > Mass Flow**.
2. Select the **Mass Flow 1** node and set *Parts* to **Regions > Pipe > Outlet**.

To define a report for the force acting on the surface of the plate:

3. Right-click the **Reports** node and select **New Report > Force**.

Property	Setting
Force Option	**Pressure + Shear**
Direction	`[1.0, 0.0, 0.0]`
Parts	**Regions > Plate > Plate**

Starting the AMESim Solver

Before STAR-CCM+ can communicate with AMESim, you start the AMESim solver in server mode.

When an AMESim model contains a `dynamic_cosim_socket`, the AMESim solver waits until the STAR-CCM+ solver runs before it begins solving the model.

To start the AMESim solver:

1. Launch AMESim and load the `AMESim_checkValve.ame` model file. The file is located in the `[INSTALL_DIR]/doc/startutorialsdata/coupling/data` directory.
2. Click to switch to simulation mode.
3. Click to start the AMESim solver.

Defining AMESim Field Function Zones

Define the AMESim field function zones, that determine the data that STAR-CCM+ sends to AMESim from any available field functions.

In this case, you set up the STAR-CCM+ simulation to send the outlet mass flow rate and force on the plate to Pin 1 and Pin 2 respectively.

To define an AMESim field function zone:

1. Right-click the **Co-Simulations** > **AMESim** > **AMESim FF Zones** node and select **Add New AMESim Zone**.



   A node, **NewAmesimFFZone**, is added beneath **AMESIM FF Zones**.
2. Rename the new node to `Mass Flow to Pin 1`.
3. Select the **Mass Flow to Pin 1** node and set the following properties:

Property	Setting
Scalar Field Function	**Reports** > **Mass Flow 1**
AMESim I/O number	`1`

4. Create another AMESim field function zone and rename it to `Force to Pin 2`.
5. Select the **Force to Pin 2** node and set the following properties:

Property	Setting
Scalar Field Function	**Reports** > **Force 1**
AMESim I/O number	`2`

Coupling to AMESim

Specify the coupling setting and start AMESim.

Coupling determines which AMESim pins receive data from STAR-CCM+ and which pins send data to STAR-CCM+. Before

coupling, you specify the host and port on which the AMESim solver is running.

At this stage, you also set the Couple From Step property to run two time steps in STAR-CCM+ before sending data to AMESim. This setting allows STAR-CCM+ to calculate an initial value for the force on the plate.

To couple to AMESim:

1. Select the **Co-Simulations** > **AMESim** node and set *AMESim Host* to the name or IP address of the machine where AMESim is running.
2. Set the remaining properties as follows:

Property	Setting
AMESim Port	60000
AMESim Pins	5
Couple From Step	2

AMESim Port specifies the port through which STAR-CCM+ and AMESim communicate. *AMESim Pins* is set to 5 as there are five pins that send and receive data between STAR-CCM+ and AMESim.



3. Right-click the **Co-Simulations** > **AMESim** node and select **Start AMESim**.

The communication between AMESim and STAR-CCM+ is established, and the number of input and output pins are determined. This information is displayed in the *Output* window.

4. Save the simulation.

Defining Boundary Conditions

Specify boundary conditions for the inlet and outlet boundaries.

When using field functions to send and receive data between STAR-CCM+ and AMESim, you can apply the field function as a boundary condition to control the pressure drop in the system. In this tutorial, AMESim Pin 3 determines the outlet pressure and AMESim Pin 4 determines the inlet pressure. The boundary types are already set in the simulation—here you define the boundary conditions only.

To set the boundary conditions:

1. Expand the **Regions** > **Pipe** > **Boundaries** node.
2. Select the **Inlet** node and set the following properties:

Node	Property	Setting
Physics Values		
Total Pressure	*Method*	**Field Function**
Field Function	*Scalar Function*	**AMESimTransfer_4**

3. Select the **Outlet** node and set the following properties:

Node	Property	Setting
Physics Values		
Pressure	*Method*	**Field Function**
Field Function	*Scalar Function*	**AMESimTransfer_3**

4. Save the simulation.

Setting Up the User Defined Vertex Motion

In this tutorial, STAR-CCM+ receives the plate displacement from AMESim Pin 5. Use this displacement with a user defined vertex motion to move the overset region (plate) during the simulation.

To set up the user defined vertex motion:

1. Right-click the **Tools** > **Motions** node and select **New** > **User Defined Vertex**.
2. Select the **Regions** > **Plate** > **Physics Values** > **Motion Specification** node and set *Motion* to **User Defined Vertex**.
3. Edit the **Regions** > **Plate** node and set the following properties:

Node	Property	Setting
Physics Conditions > **User Defined Vertex Option**	*Method*	**Displacement**
Physics Values > **LinearDisplacement**	*Total Displacement*	**Activated**
	Method	**Composite**
Composite		
X Component	*Method*	**Field Function**
Field Function	*Scalar Function*	**AMEsimTransfer_5**

Note: The *Total Displacement* option specifies that the displacement data sent by AMESim corresponds to the total amount the plate has moved. If this option is deactivated, STAR-CCM+ assumes that the displacement value corresponds to the amount the plate has moved in the current time step.

Setting the Time Step and Stopping Criteria

Match the time step in STAR-CCM+ with the corresponding value in AMESim.

In this tutorial, the time step is 5E-4 seconds. This setting uses all the data that AMESim sends to STAR-CCM+ to analyze the three-dimensional flow characteristics.

To set the time step:

1. Select the **Solvers** > **Implicit Unsteady** node and set *Time-Step* to `5.0E-4 s`.

To set the stopping criteria:

2. Expand the **Stopping Criteria** node and set the following properties:

Node	Property	Setting
Maximum Inner Iterations	*Maximum Inner Iterations*	`5`
Maximum Physical Time	*Maximum Physical Time*	`0.05 s`
Maximum Steps	*Enabled*	**Deactivated**

The simulation is set up to run for a total of 0.05 seconds, which is enough time for the plate to settle after the increase in inlet pressure.

Note: When running a STAR-CCM+ to AMESim co-simulation, the total time in AMESim must be set to a higher value than the total time in STAR-CCM+. STAR-CCM+ keeps AMESim data for two time steps and interpolates between them.

3. Save the simulation.

Plotting the Plate Displacement

Plot the plate displacement and analyze the dynamic response of the plate valve to the increased inlet pressure.

To plot the plate displacement:

1. Right-click the **Reports** node and select **New Report > Expression**.
2. Rename the **Expression 1** node to `Plate Displacement`.
3. Select the **Reports > Plate Displacement** node and set *Definition* to `$AMESimTransfer_5`.
4. Right-click the **Plate Displacement** node and select **Create Monitor and Plot from Report**.

Configure the plot to display the appropriate values:

5. Open the **Plate Displacement Monitor Plot**.
6. Edit the **Plot > Plate Displacement Monitor Plot > Axes** node and set the following properties:

Node	Property	Setting
X Axis > Labels	*Maximum*	`0.05`
Y Axis > Labels	*Maximum*	`0.015`

Running the Simulation

Visualize the Plate Displacement Monitor Plot.

Solution progress can be monitored in the *Output* window. The Plate Displacement Monitor Plot shows the solution status during the simulation.

1. Run the simulation.
2. When the solver is complete, save the simulation.

   At the end of the simulation, the Plate Displacement Monitor Plot appears as follows.



Visualizing the Solution on a Vector Scene

Create a vector scene that shows the flow on a center plane within the plate valve. Use line integral convolution to visualize the flow velocity.

To visualize the solution on a vector scene:

1. Create a new empty scene.
2. Within the **Derived Parts** node, create a new plane section with the following properties:

Property	Setting
origin	`[0, 0, 0] m`
normal	`[0, 0, 1] m`
Display	**New Vector Displayer**

3. Click **scene/plot**.
4. Select the **Displayers > Section Vector 1** node and set *Display Mode* to **Line Integral Convolution**.
5. Zoom in to the plate area, as shown in the following screenshot.



6. Save the simulation.

Summary

This tutorial demonstrated a typical workflow when running a coupled simulation with AMESim.

This tutorial has introduced the following STAR-CCM+ features:

- Selecting physics models for AMESim coupling.
- Defining AMESim field function zones.

- Coupling to AMESim.
- Setting boundary conditions for AMESim coupling.
- Setting up the User Defined Vertex motion.
- Setting the time step and stopping criteria.
- Reporting, monitoring and plotting.
- Visualizing a solution on a vector scene.

Co-Simulation API: Spindle Valve

This tutorial demonstrates how to use the STAR-CCM+ Co-Simulation API (Application Programming Interface), in order to couple a partner program to a STAR-CCM+ simulation. The partner simulation models the behavior of a spindle ball valve in response to a multiphase flow.

To exchange data with a running STAR-CCM+ simulation, the partner code makes calls on the Co-Simulation API library that creates its own companion STAR-CCM+ simulation. Internally, the partner code exchanges data with its companion STAR-CCM+ API simulation that in turn exchanges data with the fluid simulation. The coupling between the companion STAR-CCM+ API simulation and the STAR-CCM+ fluid simulation behaves in the same way as a STAR-CCM+ to STAR-CCM+ co-simulation.

In this tutorial, you model the flow of water and gas through the opening between the ball and seat of a spindle valve. A numerical code is provided that models the position of the ball in response to the fluid forces. The fluid simulation uses the position data from the ball valve code to move the ball relative to the valve surface.



During the co-simulation, the ball valve code receives traction (pressure plus shear wall stress) from STAR-CCM+ and returns the ball displacement. On receiving the ball valve displacements, STAR-CCM+ applies the mesh morpher to move the surface of the ball valve in response.

Contents

Prerequisites
The STAR-CCM+ Fluid Simulation
The Ball Valve Partner Code Simulation
Running the Co-Simulation
Visualizing the Results
Summary

Prerequisites

The instructions in the Co-Simulation API: Spindle Ball Valve tutorial assume that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+
STAR-CCM+ to STAR-CCM+ Co-Simulation	STAR-CCM+ to STAR-CCM+ Coupling: Heat Transfer in a Chimney

The STAR-CCM+ Fluid Simulation

Prepare the STAR-CCM+ simulation that models the fluid flow in the spindle ball valve.

A pressure relief valve is designed to control the pressure in a system by relieving the pressurised fluid when the pressure exceeds preset limits. When the system pressure exceeds the design pressure limit, the valve is forced open and part of the fluid flows out of the system through an auxiliary outlet, lowering the system pressure.

In this simulation, STAR-CCM+ solves for the traction loads (pressure + shear wall stress) exerted by the fluid on the solid ball valve.

The turbulent multiphase flow of water and gas enters the valve inlet at a constant velocity of 0.2 m/s in the valve axial direction z.

The physics models and conditions for the fluid simulation are provided in the initial simulation file.



Contents

Loading the Initial Simulation
Generating the Volume Mesh
Activating the Morpher
Specifying Co-Simulation Settings
Creating the Input Mesh File for the Partner Code
Monitoring the Solution

Loading the Initial Simulation

For this tutorial, you are provided with a simulation file that contains pre-defined objects.

Note: You are recommended to run this tutorial in parallel, using at least 4 cores.

To load the initial simulation:

1. Launch STAR-CCM+.
2. Select **File > Load Simulation...**
3. In the *Load Simulation* dialog, click **Browse...**
4. In the *Open* dialog, navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
5. Select `cosimApi_spindleValve.sim` then click **Open**.
6. In the *Load Simulation* dialog:
   a) Within *Process Options*, select **Parallel on Local Host** and set the number of *Compute Processes* to at least 4.
   b) Click **OK**.
7. Save the simulation as `spindleValve_fluid.sim`.

Generating the Volume Mesh

You are ready to generate the volume mesh for the fluid. All the mesh settings, including the meshing models and reference values, are provided in the initial simulation.

To generate the mesh:

1. Click (**Generate Volume Mesh**).
2. When the mesh generation is complete, save the simulation.
3. Visualize the volume mesh with a method of your choice, such as creating a mesh scene.

The mesh appears as shown below.

Activating the Morpher

Activate the mesh morpher to deform the fluid mesh according to the spindle ball displacements.

STAR-CCM+ exports traction loads (pressure + wall shear stress) to the partner code, and the partner code solves for the solid surface displacements. In STAR-CCM+, the displacements are used as inputs to the mesh morpher.

To activate the mesh morpher:

1. Expand the **Tools** node.
2. Right-click the **Motions** node and select **New** > **Morphing**.

Assign the morphing motion to the fluid region:

3. Expand the **Regions** > **Spindle Valve Fluid** > **Physics Values** node.
4. Select the **Motion Specification** node and set *Motion* to **Morphing**.

Select the appropriate morpher method for the relevant boundaries:

5. Expand the **Regions** > **Spindle Valve Fluid** > **Boundaries** > **slidingwall** > **Physics Conditions** node.
6. Select the **Morpher** node and set *Method* to **Slide on Guide Surface**.

7. Select the **slidingwall** > **Physics Values** > **Morpher: Guide Surface** node and set *Guide Surface* to **Initial Surface**.

   Setting a guide surface makes sure that the ball only moves in the direction of the spindle axis.

   The morpher method for the spindle ball boundary is automatically set to **STAR-CCM+ Co-Simulation** when you specify this boundary as the coupled boundary.

Specifying Co-Simulation Settings

Prepare the fluid simulation for co-simulation. Setting up a co-simulation requires you to activate a co-simulation model and to specify the coupled boundaries and the fields exchanged with the partner code.

The fluid simulation couples with the STAR-CCM+ API Server associated with the partner code. As the STAR-CCM+ API Server has all the features of a STAR-CCM+ simulation, you specify the co-simulation settings as you would do for a standard STAR-CCM+ to STAR-CCM+ co-simulation.

To enable co-simulation, activate a STAR-CCM+ Co-Simulation model:

1. For the physics continuum, **Continua** > **Fluid Physics**, select the following models in order:

Group box	Model
Optional Models	**Co-Simulation**
Co-Simulation	**STAR-CCM+ Explicit Co-Simulation**

2. Click **Close**.

   A new physics continuum, **Physics 1**, is added within the **Continua** node to represent the partner simulation.

When you activate a co-simulation model, STAR-CCM+ adds a **Co-Simulations** node to the simulation tree. This node allows you to specify global co-simulation settings, such as the *Concurrency Mode*, as well as to define and set up the co-simulation coupling zones.

In this simulation you set up a single coupling zone, as you have a single solid-fluid interface across which the two simulations exchange data.

To set the STAR-CCM+ server as the leading simulation:

3. Expand the **Co-Simulations** > **STAR-CCM+ Co-Simulation 1** > **Conditions** node.

4. Select the **Concurrency Mode** node and set *Method* to **Lead**.

To specify the external surface of the spindle ball as the coupled boundary:

5. Expand the **Regions** > **Spindle Valve Fluid** > **Boundaries** > **spindle_ball** > **Physics Conditions** node.
6. Select the **External Code Coupling Specification** node and set *External Code Coupling Specification* to **Zone 1**.

   STAR-CCM+ automatically sets the Morpher method for this boundary to **STAR-CCM+ Co-Simulation**. Once you associate a boundary to a coupling zone, STAR-CCM+ updates the **Coupled Boundaries** condition for that zone.



To specify the fields exported to the partner code:

7. Expand the **Co-Simulations** > **STAR-CCM+ Co-Simulation 1** > **Zones** > **Zone 1** > **Conditions** node.
8. Select the **Fluid Mechanical Export Specification** node and set *Method* to **Traction Export Fields**.

   The **Fluid Mechanical Import Specification** is already set to **Displacements**.

Specify a coupling negotiation method:

9. Expand the **Solvers** > **Co-Simulation Explicit Unsteady** > **Coupling Step Size Negotiation Method**.
10. Select the **Constant Negotiation Method** node and set *Constant Time Step* to `5.0E-6 s`.

Creating the Input Mesh File for the Partner Code

Extract the spindle ball boundary surface to use as the input surface mesh in the partner code.

In this tutorial, the partner numerical code that solves for the ball position reads input mesh files with .stl coded format. You create this surface mesh by exporting the spindle ball boundary from the fluid simulation.

1. In your working directory, create a new folder named `CoSimAPI_Mesh`.
2. In your simulation, right-click the **Representations** > **Volume Mesh** node and select **Extract Boundary Surface...**
3. In the *Extract Boundary Surface* dialog, click **OK**.

   A new node, **Extracted Surfaces**, is added to the simulation tree.
4. Expand the **Representations** > **Extracted Surfaces** > **Regions** > **Spindle Valve Fluid** > **Boundaries** node.
5. Right-click the **Spindle-ball** node and select **Export Boundaries...**
6. In the *Save* dialog:
   a) Navigate to `CoSimAPI_Mesh`.
   b) Set *Files of Type* to **Stereolithography File (*.stl)**.
   c) Click **Save**.
7. In the *STL Export Options* dialog, set *File Type* to **Coded** and click **OK**.
8. Save the simulation.

Monitoring the Solution

Prepare scenes and plots to monitor the solution progress while the co-simulation is running.

In the fluid simulation, you are provided with scenes and plots of some of the quantities that are relevant to this analysis. You can use these tools to monitor the solution progress during the co-simulation with the partner code.



Now that the fluid simulation is set up for co-simulation, complete the set up of the plots and scenes that are provided in the simulation:

1. Expand the **Reports** node.
2. Select the **Spindle Ball Displacement** node and set *Scalar Field Function* to **Displacement Imported By Co-Simulation - Zone Index 0 > Laboratory > k**.
3. Multi-select all the reports and set *Representation* to **Volume Mesh**.
4. Expand the **Scenes > Exported Pressure > Displayers > Scalar 1** node.
5. Select **Scalar Field** and set *Function* to **Pressure Exported By Co-Simulation - Zone Index 0**.
6. Multi-select the **Absolute Pressure** and **Exported Pressure** scenes. Right-click one of the selected scenes and select **Edit**.
7. In the *Multiple Objects* dialog, expand the **Displayers > Scalar 1** node and set *Representation* to **Volume Mesh**.
8. Save the simulation.

The Ball Valve Partner Code Simulation

In this tutorial, you are provided with a sample partner code that solves the equation of motion for the ball valve. Compile the code and launch the ball valve solid simulation.

A spring connected to the ball valve allows the valve motion in response to the fluid forces. The ball valve is constrained to move along its axial direction, that in this tutorial corresponds to the z direction. The partner

code calculates the position of the valve treating the model as a 1-DoF spring-mass-damping system, with equation of motion:

$$m\ddot{z}(t) + c\dot{z}(t) + kz(t) = F_z(t) \tag{11}$$

where m is the ball valve mass, c is the damping coefficient, k is the spring stiffness, and F_z is the z component of the traction field imported from the STAR-CCM+ fluid simulation, integrated over surface. In this tutorial damping effects are neglected, so $c = 0$.

The equation of motion is solved using the Newmark-beta numerical integration method.

The code that is provided in this tutorial consists of several .cpp and .h files. The main file, `SpindleValve.cpp`, outlines the recommended workflow when using the Co-Simulation API, from loading the Co-Simulation API library to setting up the co-simulation loop. You can modify the ball valve simulation time-step, coupling time-step, and maximum physical time values in this file. You can modify the spring-mass-damping system parameters in the file, `SpringMassSolver.cpp`.

Compiling the code generates an executable file, that you use to launch the ball valve solid simulation.

Contents

Compiling the Partner Code
Launching the Ball Valve Simulation

Compiling the Partner Code

Compile the sample partner code using the compiler helper script that is included in the STAR-CCM+ installation.

Start by copying the sample partner code to your working directory:

1. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/coupling/data`.
2. Copy the entire `CoSimApi_Code` folder to your working directory.

   The files that are contained in this folder make up the partner code that solves for the ball valve motion.

Based on the platform that you are using, follow the instructions for running the `buildapi` helper script for Linux or the `buildapi.bat` helper script for Windows.

Linux

To use the `buildapi` script to compile the sample partner code:

1. Open a terminal.
2. Navigate to the `CoSimApi_Code` folder that contains the sample code files:

   ```
   % cd [PATH]/CoSimApi_Code
   ```

   where `[PATH]` is the path to your working directory.
3. Enter the following command:

   ```
   % [INSTALLDIR]/star/bin/buildapi spindleValve_solid
   ```

   where `[INSTALLDIR]` is the path to your STAR-CCM+ installation directory.

   This command compiles all the `.cpp` files that are located in the `CoSimApi_Code` folder and creates an executable output file, `spindleValve_solid`, that simulates the ball valve motion. Unless you specify a different file path in the command, the `spindleValve_solid` executable is saved in the `CoSimApi_Code` folder.
4. Close the terminal.

Windows

To use the `buildapi.bat` script to compile the sample partner code:

1. Open a *Command Prompt*.
2. Navigate to the `CoSimApi_Code` folder that contains the sample code files:

   ```
   > cd [PATH]\CoSimApi_Code
   ```

   where `[PATH]` is the path to your working directory.
3. Call your Intel C++ Compiler by entering the following command:

   ```
   > call "[INSTALLDIR]\bin\compilervars.bat" intel64
   ```

   where `[INSTALLDIR]` is the path to the compiler installation directory.
4. To run the `buildapi.bat` helper script, enter the following command:

   ```
   > [INSTALLDIR]\star\bin\buildapi.bat
   spindleValve_solid
   ```

where `[INSTALLDIR]` is the path to your STAR-CCM+ installation directory.

This command compiles all the `.cpp` files that are located in the `CoSimApi_Code` folder and creates an executable output file, `spindleValve_solid.exe`, that simulates the ball valve motion. Unless you specify a different file path in the command, the `spindleValve_solid.exe` executable is saved in the `CoSimApi_Code` folder.

5. Close the *Command Prompt*.

Launching the Ball Valve Simulation

Launch the ball valve simulation using the `starapi` helper script for Linux or the `starapi.bat` helper script for Windows. These scripts are included in the STAR-CCM+ installation.

The `starapi` helper script launches the partner program, and allows the partner code to load the Co-Simulation API library and spawn the companion STAR-CCM+ API Server.

You can run the script from a command console or from the STAR-CCM+ fluid simulation. In this tutorial you run the script from a command console. This method requires two STAR-CCM+ licenses, one for the STAR-CCM+ fluid simulation and one for the STAR-CCM+ API Server. However, if you do not wish to monitor the output from the ball valve simulation, you can use the second method for sharing a single STAR-CCM+ license. For more information, see the section **Interacting with CAD/CAE Products > Co-Simulation API > Running a Co-Simulation with Partner Code**.

Based on the platform that you are using, follow the instructions for running the `starapi` script for Linux or the `starapi.bat` script for Windows.

Linux

To launch the partner program:

1. Open a terminal.
2. Navigate to your working directory:

```
% cd [PATH]/[WORKDIR_NAME]
```

where `[PATH]` is the path to your working directory and `[WORKDIR_NAME]` is the name of your working directory.

3. Enter the following command, specifying the relative path to the `spindleValve_solid` executable and to the input mesh file that you created in the previous sections of this tutorial:

   ```
   % [INSTALLDIR]/star/bin/starapi [FILEPATH]/
   spindleValve_solid [FILEPATH]/spindle_ball.stl
   ```

   where `[INSTALLDIR]` is the path to your STAR-CCM+ installation directory.

   For example, if you saved all the files as instructed in this tutorial, the command complete with the path to the files is:

   ```
   % [INSTALLDIR]/star/bin/starapi CoSimAPI_Code/
   spindleValve_solid CoSimAPI_Mesh/spindle_ball.stl
   ```

Information on the launching process is displayed in the terminal.

The message `spindleValve_solid: Waiting for incoming fields...` tells you that the STAR-CCM+ API Server has been launched and the partner code, `spindleValve_solid`, is ready to import fields from the STAR-CCM+ `spindleValve_fluid` simulation.

Note: Leave the terminal open. Closing the terminal terminates the partner program process.

Windows

To launch the partner program:

1. Open a *Command Prompt*.
2. Navigate to your working directory:

   ```
   > cd [PATH]\[WORKDIR_NAME]
   ```

   where `[PATH]` is the path to your working directory and `[WORKDIR_NAME]` is the name of your working directory.
3. Enter the following command, specifying the relative path to the `spindleValve_solid.exe` executable and to the input mesh file that you created in the previous sections of this tutorial:

   ```
   > [INSTALLDIR]\star\bin\starapi.bat [FILEPATH]
   \spindleValve_solid.exe [FILEPATH]\spindle_ball.stl
   ```

   where `[INSTALLDIR]` is the path to your STAR-CCM+ installation directory.

For example, if you saved all the files as instructed in this tutorial, the command complete with the path to the files is:

```
> [INSTALLDIR]\star\bin\starapi.bat CoSimAPI_Code
\spindleValve_solid.exe CoSimAPI_Mesh
\spindle_ball.stl
```

Information on the launching process is displayed in the *Command Prompt*.

The message `spindleValve_solid: Waiting for incoming fields...` tells you that the STAR-CCM+ API Server has been launched and the partner code, `spindleValve_solid`, is ready to import fields from the STAR-CCM+ `spindleValve_fluid` simulation.

Note: Leave the *Command Prompt* open. Closing the *Command Prompt* terminates the partner program process.

Running the Co-Simulation

Couple the STAR-CCM+ fluid simulation to the ball valve simulation and run the co-simulation.

1. Expand the **Co-Simulations > STAR-CCM+ Co-Simulation 1 > Conditions > Connect Method** node.
2. Select the **Connect to Partner Simulation** node and set *Connect to STAR-CCM+ Server Process* to the running STAR-CCM+ API Server process:
   a) Click ... (**Custom Editor**).
   b) In the *Connect to Partner Simulation* dialog, select the STAR-CCM+ server named `spindleValve_solid`, then click **OK**.



The fluid simulation is now connected to the partner code through the STAR-CCM+ API Server.

3. In the fluid simulation, click (**Run**).
   Information on the co-simulation process is displayed in both the fluid simulation *Output* window and the terminal where you launched the partner code.

To monitor the solution progress during the co-simulation:

4. Right-click the **Plots** node and select **Open All Plots**.

Visualizing the Results

Visualize the scenes and plots that are relevant to this mechanical co-simulation analysis.

When the co-simulation is complete, the plots for the spindle ball displacement, the traction loads, and the mass flow at the outlet as functions of time appear as shown below:


Spindle Ball Displacement Monitor Plot
Spindle Ball Displacement Monitor
Spindle Ball Displacement Monitor (m)
0.00018
0.00016
0.00014
0.00012
0.0001
8e-05
6e-05
4e-05
2e-05
0
0.0005 0.001 0.0015 0.002 0.0025 0.003 0.0035 0.004 0.0045
Physical Time (s)


Traction Load Monitor Plot
Traction Load Monitor
Traction Load Monitor (N)
2000
1800
1600
1400
1200
1000
800
600
400
200
0
0.0005 0.001 0.0015 0.002 0.0025 0.003 0.0035 0.004 0.0045
Physical Time (s)



Visualize the scalar scenes displaying the pressure in the fluid domain and the pressure exported to the ball valve surface.

1. Open the **Absolute Pressure** and **Exported Pressure** scenes.

   The **Absolute Pressure** scene appears as shown below.



Visualize the exported pressure data on the ball valve surface, that is imported under the **Regions** node when you run the co-simulation.

2. Edit the **Exported Pressure** > **Displayers** > **Scalar 1** > **Parts** node.
3. In the *Parts* dialog, remove the selected boundary, select the **Spindle Ball Surface** > **Boundaries** > **Spindle Ball Surface** node then click **OK**.

   The **Exported Pressure** scene appears as shown below.



4. Save the simulation.

Summary

This tutorial demonstrated how to use the STAR-CCM+ Co-Simulation API to couple partner code with a running STAR-CCM+ simulation.

This tutorial has demonstrated the following features:

- Activating the morpher for a mechanical co-simulation analysis.
- Specifying co-simulation settings for coupling with partner code.
- Creating the input mesh for the partner code.
- Compiling the partner code using the `buildapi` helper script.
- Launching the STAR-CCM+ API Server using the `starapi` helper script.
- Coupling a STAR-CCM+ simulation with the partner code.
- Running the co-simulation.

CAD Client Contents

Using CAD Clients 1

Working With STAR-Works 2

Introducing STAR-Works . . . 2
Transferring Geometry to STAR-CCM+ . . . 34
Understanding the STAR-Works Process . . . 41
Defining Materials and Physics . . . 46
Setting Initial Conditions . . . 75
Setting Boundary Conditions . . . 78
Meshing the Model . . . 90
Obtaining a Solution . . . 118
Exporting Files and Data . . . 131
Resetting STAR-Works Settings . . . 135
Best Practice Guidelines . . . 136

Working With STAR-CAT5 178

Introducing STAR-CAT5 . . . 178
Transferring Geometry to STAR-CCM+ . . . 207
Understanding the STAR-CAT5 Process . . . 216
Defining Materials and Physics . . . 221
Setting Initial Conditions . . . 256
Setting Boundary Conditions . . . 259
Meshing the Model . . . 272
Obtaining a Solution . . . 306
Exporting Files and Data . . . 328
Resetting STAR-CAT5 Settings . . . 334
Best Practice Guidelines . . . 334

Working With STAR-NX 393

Introducing STAR-NX . . . 393
Transferring Geometry to STAR-CCM+ . . . 430

Understanding the STAR-NX Process 439
Defining Materials and Physics 443
Setting Initial Conditions. 471
Setting Boundary Conditions 474
Meshing the Model 486
Obtaining a Solution. 515
Exporting Files and Data 529
Resetting STAR-NX Settings 534
Best Practice Guidelines 535

Working With STAR-Pro/E 572

Introducing STAR-Pro/E 572
Transferring Geometry to STAR-CCM+ 610
Understanding the STAR-Pro/E Process 617
Defining Materials and Physics 622
Setting Initial Conditions. 650
Setting Boundary Conditions 653
Meshing the Model 666
Obtaining a Solution. 701
Exporting Files and Data 715
Resetting STAR-Pro/E Settings 720
Best Practice Guidelines 721

Working With STAR-Creo 746

Introducing STAR-Creo 746
Transferring Geometry to STAR-CCM+ 779
Understanding the STAR-Creo Process 786
Defining Materials and Physics 790
Setting Initial Conditions. 818
Setting Boundary Conditions 820
Meshing the Model 833
Obtaining a Solution. 867
Exporting Files and Data 879
Resetting STAR-Creo Settings 883

Working With STAR-Inventor 885

Introducing STAR-Inventor 885
Transferring Geometry to STAR-CCM+ 918
Understanding the STAR-Inventor Process 926
Defining Materials and Physics 930

Setting Initial Conditions . . . 959
Setting Boundary Conditions . . . 962
Meshing the Model . . . 975
Obtaining a Solution. . . . 1011
Exporting Files and Data . . . 1023
Resetting STAR-Inventor Settings . . . 1028
Best Practice Guidelines . . . 1029

Working With STAR-SpaceClaim 1072
Introducing STAR-SpaceClaim . . . 1072
Transferring Geometry to STAR-CCM+ . . . 1101
Understanding the STAR-SpaceClaim Process . . . 1108
Defining Materials and Physics . . . 1113
Setting Initial Conditions. . . . 1141
Setting Boundary Conditions . . . 1144
Meshing the Model . . . 1155
Obtaining a Solution. . . . 1184
Exporting Files and Data . . . 1195
Resetting STAR-SpaceClaim Settings . . . 1199

CAD Client Tutorials 1201

STAR-Works CAD Client 1202
Cold Flow Through a Motorcycle Exhaust Manifold . . . 1203
Porous Media Flow . . . 1231
Aircraft Wing Deicing System . . . 1258
Rotating Flow inside a Graphics Card . . . 1306
External Flow Over a Car . . . 1345

STAR-CAT5 CAD Client 1384
Cold Flow Through a Motorcycle Exhaust Manifold . . . 1385
Porous Media Flow . . . 1415
Aircraft Wing Deicing System . . . 1446
Rotating Flow inside a Graphics Card . . . 1497
External Flow Over a Car . . . 1538

STAR-NX CAD Client 1582
Cold Flow Through a Motorcycle Exhaust Manifold . . . 1583
Porous Media Flow . . . 1610
Aircraft Wing Deicing System . . . 1640
Rotating Flow inside a Graphics Card . . . 1688

External Flow Over a Car 1729

STAR-Pro/E CAD Client 1771

Cold Flow Through a Motorcycle Exhaust Manifold 1772
Porous Media Flow 1799
Aircraft Wing Deicing System 1828
Rotating Flow inside a Graphics Card 1875
External Flow Over a Car 1915

STAR-Creo CAD Client 1955

Cold Flow Through a Motorcycle Exhaust Manifold 1956
Porous Media Flow 1982
Aircraft Wing Deicing System 2010
Rotating Flow inside a Graphics Card 2056
External Flow Over a Car 2095

STAR-Inventor CAD Client 2135

Cold Flow Through a Motorcycle Exhaust Manifold 2136
Porous Media Flow 2164
Aircraft Wing Deicing System 2190
Rotating Flow inside a Graphics Card 2235
External Flow Over a Car 2273

STAR-SpaceClaim CAD Client 2314

Cold Flow Through a Motorcycle Exhaust Manifold 2315
Porous Media Flow 2341
Aircraft Wing Deicing System 2369
Rotating Flow inside a Graphics Card 2415
External Flow Over a Car 2453

Using CAD Clients

CD-adapco provide a number of CAD Clients integrated with major CAD software. The currently available clients are:

- STAR-Works
- STAR-CAT5
- STAR-NX
- STAR-Pro/E
- STAR-Creo
- STAR-Inventor
- STAR-SpaceClaim

Working With STAR-Works

This part of the documentation details how to create, set up and run a CFD simulation in STAR-Works. It assumes that you have previous experience in using SolidWorks, and only describes the additional features available as part of STAR-Works. The following sections are included:

- Introducing STAR-Works
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-Works Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-Works Settings
- Best Practices Guidelines

You use STAR-Works to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-Works

STAR-Works provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the SolidWorks environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-Works
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-Works enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-Works provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-Works interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

  - Ideal Gas
- Porous Media
  - Isotropic
  - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-Works provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the SolidWorks environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
  - Tetrahedral core mesh generation
  - Polyhedral core mesh generation
  - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-Works allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-Works supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-Works interface, and basic concepts. Topics include:

- Launching STAR-Works
- The STAR-Works Info App Window
- Modifying the config.xml File
- Understanding the STAR-Works Interface
- The STAR-Works Object Tree
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-Works Settings While the STAR-CCM+ Client is Running

Launching STAR-Works Manually

By default STAR-Works should start automatically when you launch SolidWorks.

If this does not happen, you can manually start STAR-Works by following the steps outlined below:

1. Start SolidWorks as usual.

2. Click Tools > Add Ins...



3. Activate the STAR-Works checkboxes, and click OK.

The STAR-Works feature manager tree appears on the left of the SolidWorks window when you load or create a geometry file.

The STAR-Works Info App Window

The *STAR-Works Info App* window automatically appears when you start SolidWorks. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-Works. More tabs appear as you create or load geometry files in STAR-Works, allowing you to perform more actions.



You cannot close the *STAR-Works Info App* window while SolidWorks or STAR-Works are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-Works license checked out successfully. The version of STAR-Works is also displayed in this tab. This tab appears by default when you launch SolidWorks without a geometry file.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-Works.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-Works searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-Works user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-Works is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-Works Info App. This can be set to one of the following:
  - `Normal` — This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose` — This will display a moderate amount of information in the Output panel.
  - `Log` — This will display all available information in the Output panel.

Understanding the STAR-Works Interface

The STAR-Works user interface is embedded within the SolidWorks environment, with the addition of the STAR-Works feature tree, ribbon, and toolbar.



The majority of the STAR-Works operations are executed within the feature tree and ribbon. Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-Works interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings
- Region settings
- Add new boundary
- Boundary conditions

- Mesh Pipeline settings
- Add new mesh controls

Ribbon

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
- Solution
  - Build and Run
  - Run Solver
  - Run For More Iterations
  - Clear Solution
- Model
  - Check Model For Errors
  - Transfer CAE Model
- Export
  - Surface Mesh (.dbs)
  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Geometry
  - Transfer Geometry
- STAR-Works Help
- About STAR-Works

Toolbar

The operations provided in the toolbar include:

- Definition of Units System
- Solution Controls
- Clear All
- Preferences
- Generate Surface Mesh
- Generate Volume Mesh
- Check Model For Errors
- Build and Run
- Run Solver
- Run For More Iterations
- Show/hide Mesh
- Transfer CAE Model
- STAR-Works Help
- About STAR-Works
- Transfer Geometry

Split Panel Display in SolidWorks

The panel to the left of the graphics area should not be split when working with the STAR-Works object tree.

The STAR-Works Object Tree

STAR-Works uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-Works. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-Works applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-Works Interface section.

Clicking on the Region node within the object tree will highlight all the relevant faces on the geometry model for that particular region. If different parts are assigned to multiple regions, then clicking on a region will only highlight the part and faces associated with that specific region. For example, clicking on the Fluid node below will highlight all the faces associated with the Fluid region.





Similarly, clicking on the Solid node highlights all the faces associated with the Solid region.





Clicking on a boundary node in the STAR-Works object tree highlights all the faces on the CAD geometry model that are associated with it.Note that clicking on the *Default* boundary node will not highlight any faces on the CAD geometry model. To stop highlighting the faces, you can either click on a blank space in the geometry window, or click on the root node in the STAR-Works object tree.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select Preferences from the ribbon.

2. Alternatively, click on the Preferences button on the *STAR-Works* toolbar.

The *Preferences* dialog appears.

3. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

4. To run STAR-Works in parallel, select Parallel from the *Run Mode* box. The *Parallel Options* box appears.



5. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.
6. Click ✓ to close the dialog.

   If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-Works and STAR-CCM+ need to be closed for the changes to take effect.



7. Click OK to proceed.
8. Close the active STAR-Works and STAR-CCM+ sessions.
9. Restart STAR-Works.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-Works.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-Works.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-Works.

To change the units system:

1. Bring up the *Units System* dialog:
   - Right-click the STAR-Works node within the object tree and select Units System.



   - Alternatively, from the ribbon, select Units System.



- You can also adjust the units system by clicking on the Units System button on the *STAR-Works* toolbar.

The *Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click ✓ to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers
- Angstroms
- Nanometers
- Mils
- Microinches

- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-Works environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the ribbon.



- Alternatively, from the toolbar, click the Transfer CAE Model button.

The STAR-CCM+ version affiliated with STAR-Works launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Using SolidWorks Mouse Controls in STAR-CCM+

STAR-CCM+ offers the ability to set the mouse interactor style to match that of SolidWorks; this could make geometry manipulation in STAR-CCM+ more natural if you are used to working with SolidWorks.

To change the mouse interactor style:

1. Launch STAR-CCM+.

2. Click *Tools* and select Options from the menu.



The *Options* dialog appears.

3. Expand the STAR-CCM+ node and select Visualization.
4. In the *Properties* window, select Interactor Style and set this to Solidworks using the drop-down menu.



5. Click Close.

STAR-CCM+ will now use this mouse interactor style each time you launch it.

Editing STAR-Works Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-Works while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-Works, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Common Functions

The following section details common operations used within the STAR-Works environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-Works
- Removing surfaces and parts from objects in STAR-Works

Renaming an Object

To rename an object:

1. Right-click the object node within the object tree and select Rename.



2. Enter the name.



3. Click on a blank space in the object tree.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, no pop-up menu appears.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select Edit...



   The settings dialog for the object appears.

2. Make the changes that you want.
3. Click ✓ to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-Works

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-Works object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

- On the CAD geometry, click the surfaces that you want to assign to the object. You can add multiple surfaces by clicking more surfaces on the CAD geometry.

The selected surfaces are highlighted on the CAD geometry.



The selected surfaces are also displayed in the *Face Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to assign to the object. You can add multiple parts by clicking more parts on the CAD geometry.

The selected part is highlighted on the CAD geometry.



The selected parts are displayed in the *Part Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-Works

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-Works object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

1. In the *Face Selection* box, select the surface you want to delete.

The edge of the selected surface becomes highlighted.



2. Right-click on the selection and select Delete.



The surface is removed from the *Face Selection* box. You can select multiple surfaces for deletion by holding down the <`Ctrl`> key, and remove them in one operation. The features of the native CAD package also work with STAR-Works, so if you want to remove all surfaces from the *Face Selection* box, you can right-click anywhere in the *Face Selection* box and select Clear Selections.

Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

1. In the *Part Selection* box, select the part you want to remove.
2. Right-click on the selection and select Delete.



The part is removed from the *Part Selection* box. You can select multiple parts for deletion by holding down the <`Ctrl`> key, and remove them in one operation. The features of the native CAD package also work with STAR-Works, so if you want to remove all parts from the *Part Selection* box, you can right-click anywhere in the *Part Selection* box and select Clear Selections.

Alternatively, on the CAD geometry, click the part that you wish to remove. This part is removed from the *Part Selection* box and is no longer highlighted on the CAD geometry.

Transferring Geometry to STAR-CCM+

With STAR-Works, you can transfer your CAD geometry directly to STAR-CCM+ from within the SolidWorks environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. An additional benefit of using this method is that the CAD face names and assembly structure are maintained in the imported geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

To transfer the geometry to STAR-CCM+:

1. Click Transfer Geometry on the STAR-CCM+ ribbon.



   STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-Works. This action ensures that the simulation file remains associated with the SolidWorks file. You can also safely rename and move the file using this method.

- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the SolidWorks file.

Resuming a Simulation

To resume a simulation after closing STAR-Works and STAR-CCM+:

1. Launch STAR-Works.
2. Open the CAD geometry in STAR-Works.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click Design Parameters on the STAR-Works ribbon.



The *Design Parameters* dialog appears.



2. Select the design parameters that you want to transfer from the list.You can select multiple items at once.

Note: STAR-Works only recognizes native CAD features as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Transfer to STAR-CCM+ option.
4. Click ✔ .
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-Works, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-Works, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

The geometry updates according to the design parameter changes.



6. Return to SolidWorks. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-Works and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

If you close the geometry part in SolidWorks, but leave SolidWorks running, the geometry part is re-opened in the existing SolidWorks session when

you click Update Model in STAR-CCM+. After the changes are made to the geometry, the geometry closes and the interactive session is left as it was before.

Updating the Geometry in STAR-CCM+ with SolidWorks Closed

You can close SolidWorks and continue working with the simulation in STAR-CCM+. Doing so, however, breaks the link between the two software packages. Any further updates that you make to the geometry in STAR-CCM+ are not transferred back to the geometry in SolidWorks. Instead, these changes are stored in memory, which allows you to include the design parameters in your optimization scripts.

This workflow also allows you to release the SolidWorks and STAR-Works licenses when performing tasks in STAR-CCM+. In this case, the SolidWorks and STAR-Works licenses are temporarily checked out only when the updates to the geometry are stored in memory.

Transferring the geometry from STAR-Works will overwrite the geometry in STAR-CCM+.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from SolidWorks to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-Works root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click ✓.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.
Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.
Transfer Sheet Geometry	Specifies whether to transfer sheet bodies as part of the transfer process. This option is turned off by default.

Understanding the STAR-Works Process

In STAR-Works the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume
3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity

5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings
7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-Works and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-Works

STAR-Works only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-Works.

Suppressed or deactivated parts and features are not supported by STAR-Works and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-Works.

STAR-Works does not support Large Assembly Mode. If you wish to use large assemblies within STAR-Works, disable Large Assembly Mode in SolidWorks before opening the assembly.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

  For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-Works includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-Works includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-Works is completely integrated within the SolidWorks environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-Works model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

  Including material type and properties.
- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.
- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.

- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-Works automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select New.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



To create an additional region:

- Right-click on the Regions manager node in the object tree and select New.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-Works). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-Works model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. Then right-click on the region node under the Region manager node in the object tree and select Edit... In the *Region Settings* dialog, select the alternative continuum using the *Continuum* drop-down menu and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-Works each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By

default the material type is set to Air, and this can be changed to any one of a number of predefined solids, liquids or gases that are available in STAR-Works. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-Works object tree and select Edit.



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material* box, you can select a material type.
- In the *Material Properties* box, you can define each of the material properties.

Continuum Settings
Continuum Name
Continuum
Material
Gas
Liquid
Solid
Acetylene(V)
Air(V)
Aluminum Oxide(V)
Aluminum(V)
Ammonia(V)
Argon(V)
Butane(V)
CH2(S)Group(V)
Flow Characteristics
Equation of State
Optional Physics Models
Material Properties
Density
1.18415
kg/m^3
Specific Heat
1003.62
J/kg-K
Molecular Viscosity
1.85508e-005
Pa-s
Conductivity
0.0260305
W/m-K
Reference Properties
Gravity
Initialisation

Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-Works offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the list. When you have selected a material, the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-Works are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.
Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-Works, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)
- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select Edit...



The physics models available in the *Continuum Settings* dialog are determined by the material type.



Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.

Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.



The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-Works are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Locate the *Gravity* box and tick the checkbox in the title bar. This will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. To define reference properties for a continuum, right-click the Continuum node in the object tree and select Edit. In the *Continuum Settings* dialog, locate the *Reference Properties* box and expand , if necessary.



The reference properties are described below:

- Reference Pressure

  The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density ρ, Reference Velocity, Reference Area.

  These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-Works, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select Edit...



The physics models can then be enabled for the region by ticking the corresponding checkbox and entering values as necessary.

Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the *Porous* checkbox.

There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select the radio button labeled *Isotropic*.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, α and β respectively.

Uni-Directional

To specify uni-directional porous resistance, select the radio button labeled *Uni-Directional*.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, α and β respectively. The cross-flow resistance is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, $\mathbf{i}$ $\mathbf{j}$ $\mathbf{k}$.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the *Moving Reference Frame* checkbox.

In the *Moving Reference Frame* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the ω box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, i j k. The units used are the same as the file defaults.

Alternatively, click on the Choose Point... button and select a vertex on the geometry to define the origin of the axis of rotation. The direction vector of the axis is defined using the i j k fields in the *Moving Reference Frame* panel.

Note that clicking on the Choose Point... button, will activate the vertex selection mode (the *Part Selection* box will be deselected and changes from blue to white); this will enable you to select vertices on the geometry model. To cancel the vertex selection mode and return to the part selection mode, click in the *Part Selection* box.

The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.

There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select the radio button labeled *Energy Input per Unit Volume*.



Enter the value for the energy input value per unit volume in the box.

Total Energy Input

To define the total energy input, select the radio button labeled *Total Energy Input*.



Enter the value for the total energy input in the box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.

2. Create a streamlines derived part.
3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-Works.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-Works will use default values to initialize the solution. In STAR-Works, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-Works object tree and select Edit.



This opens the *Continuum Settings* dialog, in which the *Initialization* box can be used to specify the following values:



- U V W Initial velocity
- Initial pressure
- Initial temperature
- Int len Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial velocity, open the *Continuum Settings* dialog, expand the *Initialization* box, and enter the values in the U, V, and W boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, locate the *Initialization* box, and enter a value in the box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the Turbulence Intensity Int and Turbulent Length Scale len boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-Works boundary is made up of one or more CAD surfaces within the simulation which, along with other boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select New from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-Works). When you assign a surface to another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-Works, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Velocity Inlet.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity, with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors U V W, tick the checkbox labeled Specify velocity components.



Cylindrical vectors V_r V_θ V_z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i j k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Mass Flow Inlet.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.



The mass flow rate , must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors i, j, k, tick the checkbox labeled Specify Direction.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Stagnation Inlet.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure , must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors i, j, k, tick the checkbox labeled Specify Components.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-Works, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Flow-Split Outlet.

The only parameter required is the flow-split ratio 60↗ 40↘, expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.

Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Pressure Outlet.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.

The pressure , must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Wall.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U V W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity ω. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i j k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux , to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature , and heat transfer coefficient .



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and click the radio button labeled Symmetry Plane.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-Works to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-Works and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature*.



Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Turbulence*.



Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity Int

- Turbulence Length Scale

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Passive Scalar*.



This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-Works provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-Works* ribbon), it is possible to execute only single stages from the STAR-Works ribbon. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the SolidWorks parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from SolidWorks.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the SolidWorks *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-Works creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-Works appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-Works appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all SolidWorks parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose Edit.



The *Mesh Pipeline Settings* dialog appears. This dialog has several collapsible sections for the general settings.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click ✓ to save them and close the dialog.
- Click ✗ to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.



Selecting the trimmer option will activate the Trimmer tab..

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-Works Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.



The base size can be set as either a relative or absolute value. When it is set as a relative value (by ticking the Relative to Geometry option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* box of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of prism layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.



This panel allows you to change the following parameters:

- Relative Minimum Surface Size

- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-Works will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer panel.

The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the SolidWorks parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when they are created. A part can only belong to one mesh pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select New.



During the creation of a new mesh pipeline, it is necessary to assign to it the SolidWorks parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

  This mesh control is applied to one or more faces of the SolidWorks model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose New Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select New Volumetric Mesh Control from the pop-up menu.



The *Volumetric Mesh Control* dialog appears.



The *Volumetric Mesh Control* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape
- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the SolidWorks geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.

X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the SolidWorks window. To use an existing point, click the Choose Point 1... or Choose Point 2... button and use the SolidWorks point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Choose Point... button and using the SolidWorks point selection procedure to choose an existing point in the SolidWorks window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update Selected button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select New Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the Custom Surface Mesh Size checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-Works* ribbon.



- Alternatively, click the Generate Surface Mesh button located in the *STAR-Works* toolbar.

Output from the surface mesher will be displayed in the *STAR-Works Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-Works* ribbon.



- Alternatively, click the Generate Volume Mesh button located in the *STAR-Works* toolbar.

Information on the progress of the volume mesher will be displayed in the *STAR-Works Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Clear Mesh from the ribbon.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-Works* ribbon.



When the mesh is displayed, the Mesh View button appears pressed , otherwise the standard icon is shown .

Alternatively, click the Show/Hide Mesh button on the *STAR-Works* toolbar.

Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both

sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-Works will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-Works Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Running in Batch Mode
- Clearing the Solution
- Visualizing the Solution

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the STAR-Works node within the object tree and select Solution Controls.



- Alternatively, from the ribbon, select Solution Controls.



- You can also access the *Solution Control* dialog by clicking on the Solution Controls button on the *STAR-Works* toolbar.

The *Solution Controls* dialog appears.



2. Here, you can:
   - Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.
   - Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.
3. Click ✔ to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:



1. In the Number of Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:
   - First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.
   - Second Order—This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:
   - Momentum
   - Pressure
   - Turbulence (with turbulence modelling defined in Continuum Settings)
   - Temperature (with temperature modelling defined in Continuum Settings)
     - Fluid Temperature Under-Relaxation Factor
     - Solid Temperature Under-Relaxation Factor
5. If you want to use additional STAR-CCM+ physics models with STAR-Works, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros
6. Click to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-Works by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-Works:

1. Open the *Solution Controls* dialog.
2. Activate the Use Java File option.

   A text entry field appears.



3. Click the Select Input File button. This action opens a browser dialog.
4. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
5. Click ✔ to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-Works.

Using STAR-CCM+ Continuum Physics with STAR-Works

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-Works and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.

2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.
5. Return to STAR-Works, open the *Solution Controls* dialog and activate the Use Java File checkbox.
6. Select the Java macro that you created earlier.
7. Run the case from STAR-Works.

Checking the Model for Errors

STAR-Works monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.
- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.



2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-Works object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-Works root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check in the *STAR-Works* ribbon.



- Alternatively click the Check Model for Errors button in the *STAR-Works* toolbar.

If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-Works Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-Works, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-Works* ribbon.



- Alternatively click the Build and Run button in the *STAR-Works* toolbar.

This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-Works Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Run Solver in the *STAR-Works* ribbon.



- Alternatively, click the Run button in the *STAR-Works* toolbar.

STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Run For More Iterations in the *STAR-Works* ribbon.

- Alternatively you can click the Run for More Iterations button in the *STAR-Works* toolbar.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Running in Batch Mode

To run STAR-Works in batch mode, you must first add the paths to the `STAR-Works.dll` and `STARWorksBatch.exe` files to your Path user environment variable. These paths are added to the list of paths that may already exist on your machine. For more information on how to set the `Path` variable please see the Setting the Path User Environment Variable section.

A typical path is shown below.

```
C:\Program Files\CD-adapco\STAR-CAD<star-ccm+_version>\
STAR-Works\win64\bin
```

To run the solver in batch mode:

1. Choose or create a directory from which to launch STAR-Works in batch mode. This is referred to as the Batch Runtime Folder.
2. Set up the desired case in STAR-Works, assigning all relevant continuum, region, boundary, mesh pipeline, and solution control settings.
3. Place the geometry file in the Batch Runtime Folder.
4. Launch a Command Prompt.
5. In the Command Prompt, change to the Batch Runtime Folder that contains the CAD geometry with the saved STAR-Works settings.

   An example code is shown below:

   ```
   C:\>cd "C:\Batch Runtime Folder"_
   ```

6. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

   ```
   StarWorksBatch.exe -doc <docname> -ver <vernumb> -mode
   copy
   ```

   where:

   - `-doc` is the input option to specify the input geometry file. It has to be followed by the `<docname>` argument which is the full path to the input SolidWorks part/assembly file (including the extension).

- `-ver` is the SolidWorks version number. Refer to the release notes to see which versions of SolidWorks are supported.
- `-mode copy` is an optional parameter that creates a copy of the SolidWorks file. A file named `exhaustManifold.sldprt` is copied to `exhaustManifold_0.sldprt`. STAR-Works uses the copy instead of the original file.

An example code for running STAR-Works in batch mode using SolidWorks 2012 is shown below.

```
StarWorksBatch.exe -doc "C:\Batch Runtime Folder\exhaus
tManifold.SLDPRT" -ver 2012
```

7. Press <`Enter`> to run the simulation in batch mode.

STAR-Works will work in the background, meshing the geometry and running the solvers. A simulation file will be saved in the Batch Runtime Folder once the solver has finished running. If you used the `-mode copy` argument, the simulation file will share the same name as the copy: `exhaustManifold_0.sim`. If there is an existing simulation file in the Batch Runtime Folder, a backup file is created with a `.sim~` extension. A log file (`starcad.log`) is also created and saved in the Batch Runtime Folder—this log file contains a record of all the messages that appeared in the command prompt during batch processing.

Setting the Path User Environment Variable

You need to set the `Path` user environment variable in order for STAR-Works to run in batch mode. This variable will have the full path to the folder containing the `STAR-Works.dll` and STARWorksBatch`.exe` files. To create the `Path` user environment variable:

1. Open the Environment Variables dialog:
   - For Windows 7:
     i. Click Start and type `env` into the search field near the bottom of the *Start* menu.
     ii. Select `Edit environment variables for your account` from the list. The *Environment Variables* dialog appears.
     iii. If the Path user environment variable already exists in the *User variables* box click on Edit..., otherwise click on New... to create it. The *User Variable* dialog appears.
2. Set the Variable name to `Path`.
3. Set the Variable value to:

```
C:\Program Files\CD-adapco\STAR-CAD<star-ccm+_version>\
STAR-Works\win64\bin
```

4. Click OK to close the *User Variable* dialog.

5. Click OK to close the *Environment Variables* dialog.
6. Click OK to close the *System Properties* dialog.

Launching a Command Prompt

To launch a command prompt in Windows 7:

1. Click Start.
2. Type `cmd` in the search field at the bottom of the *Start* menu and select `cmd.exe` from the search results.

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-Works case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the ribbon, select Clear Solution.



- Alternatively, click on the Clear Solution button on the *STAR-Works* toolbar.

The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and

vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-Works.

- Exporting a Surface Mesh
- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-Works as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-Works ribbon, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-Works as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-Works ribbon, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-Works ribbon, select Export > Engineering Report (.html,.csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)

  - Density (kg/m^3)
- Force data, at each boundary:
  - Net Force [Fx, Fy, Fz]
  - Pressure Force [Fx, Fy, Fz]
  - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
  - Coefficient of Net Force [Cfx, Cfy, Cfz]
  - Coefficient of Pressure Force [Cfx, Cfy, Cfz]
  - Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
  - Physics models applied to the continuum
  - Materials properties
  - Turbulence properties
  - Segregated flow properties
  - Reference values
  - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
  - Polyhedral mesher properties
  - Prism layer mesher properties (if activated)
  - Surface remesher properties
  - Reference values
  - Volumetric controls
- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values
- Solver parameters

- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-Works to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-Works, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-Works ribbon, select Export > Simulation (.sim).



Resetting STAR-Works Settings

STAR-Works saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-Works settings:

- From the ribbon, select Clear All.



Alternatively, click on the Clear All button on the *STAR-Works* toolbar.

This action returns the geometry to its original state.

Best Practice Guidelines

This section aims to provide guidance and assistance on making a geometry part suitable for CFD analysis using STAR-Works. The following items are covered:

- Creating the Fluid Domain:
  - Internal Flow
  - External Flow
  - Multi-Region Flow
- Imprinting:
  - Imprinting Parts in an Assembly
  - Imprinting Surfaces on a Part for Boundary Region Applications

Prior to using these methods, ensure that the components in assemblies are all solid bodies and are in contact with each other. Any gaps between components will need to be closed.

A basic knowledge and understanding of both CAD and CFD are assumed, although most basic functions are covered.

Creating the Fluid Domain

In the majority of CAD models, the fluid domain for analysis must be extracted from the original part or assembly. The following sections demonstrate how to prepare the fluid domain for:

- an internal flow
- an external flow
- a multi-region flow

Internal Flow

In general, the steps to follow in creating the fluid domain for an internal flow are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create additional parts that close off any openings in the original geometry, we will call these parts the “caps”.
3. Create an overall block part that encloses all the other parts, the “enclosure”. This will form the internal fluid volume.
4. Subtract all parts from the enclosure--the remaining geometry is the fluid domain.

An example case is shown below. In this case the assembly contains two parts, the valve (shown in red) and the valve housing (shown in blue).



1. If the geometry to be analyzed is a part then this will need to be inserted into an assembly. A quick way of doing this is by clicking on *File* > *Make Assembly from Part*.



If the geometry is already an assembly then this step can be skipped.

We will insert a new part into the assembly. This will be used to cap any openings in the geometry.

2. Insert a new part, *Insert* > *Component* > *New Part...*



You may need to confirm this action by clicking on the screen--the mouse pointer will have a green tick next to it to indicate this is required.

3. Edit the new part within the assembly by right-clicking on it in the feature tree and selecting Edit Part.



This will make the selected part the active part. By default, all inactive parts in an assembly will become transparent. The transparency of the these parts can be controlled by using the Assembly Transparency button in the *Features* ribbon.



The edges and faces of inactive parts can be selected and used with the active part.

4. Create a sketch on one of the surfaces where there is an opening, such as the one shown below.



We will now convert the outer edge of the opening to a sketch entity in the current sketch.

5. Select the Convert Entities function from the toolbar

and choose the outermost edge of the opening's surface.



6. Exit the sketch.

7. Extrude the sketch to close off the opening. The depth of the extrusion is not important.



8. Repeat these steps for all other openings.



We will now insert a new part into the assembly. A solid block that encloses all the other parts will be created. We will then subtract the other parts from this solid block--the result will be the internal fluid volume.

9. Return to assembly mode and insert another part into the assembly, as described in step 1.

10. Edit this new part within the assembly and create a new sketch that will enclose the entire geometry when extruded.



11. Extrude the sketch so that it encloses all the parts, as shown above.

12. Use the Cavity feature, *Insert* > *Features* > *Cavity...*, to extract the internal volume.



The *Cavity* panel will appear.

13. Expand the feature tree and select all the inactive parts from here as shown below. The parts to select should be the original assembly components and the caps.



The *Bodies to Keep* dialog will now appear. Here, you will need to tell SolidWorks which bodies you want to keep after the operation. The aim is to keep the cavity which represents the internal fluid volume of the geometry.

14. Click on the Selected Bodies radio button in the *Bodies* box. This will show a list of available bodies. In this example there are two bodies available: one containing all the solid bodies (the valve housing, the

valve, the caps, and the enclosure), whilst the other contains the internal fluid volume or “cavity”. A preview is shown in the geometry window, in this example, the body we want is Body 2.



The preview is shown below, the blue area corresponds to the volume that will be kept after the operation, this is the internal fluid volume.



15. Return to the assembly and suppress any parts that are not needed in the analysis, in this case the original assembly components and the caps. To suppress a part, right-click on the part in the feature tree and select Suppress.

The final extracted volume is shown below. This internal volume will automatically update itself if the original geometry is modified, or if the assembly constraints (such as the valve angle) is changed.



The geometry is now ready for the CAD client analysis. Suppressed parts in an assembly will not be included in the analysis. In the above example, it is recommended that you extend the inlet and outlet regions to appropriate lengths so that a fully developed flow can be achieved before reaching the valve and to help reduce recirculation at the outlet.

External Flow

The generation of a domain for external aerodynamics is generally easier than for internal flow cases but the sizing of the domain is more critical. As a guideline (assuming a specific situation is not being modelled i.e. flow in a wind tunnel) the domain should be sized as follows:

Inlet: 2 x part length from the front of the part

Outlet: 4 x part length from the rear of the part

Sides: 3 x part width from the side of the part

Ceiling: 4 x part height from the top of the part

In general, the steps to follow in creating the fluid domain for an external flow are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an overall block part that encloses all the other parts. The shape and size should correspond to the desired fluid domain.
3. Subtract all parts from the overall block part - the remaining geometry is the fluid domain.

An example case is shown below. In this example a geometry of a car will be used.



1. If the geometry to be analyzed is a part then this will need to be inserted into an assembly. A quick way of doing this is by clicking on *File* > *Make Assembly from Part*.



Click on a blank space in the geometry window to place the car in the assembly. If the geometry is already an assembly then this step can be skipped.

2. If the case is symmetrical, cut it in half as shown below. The easiest way to do this is by performing an extruded cut along the symmetry plane.



We will now add a new part to the assembly. This will become the fluid domain.

3. Insert a new part, *Insert* > *Component* > *New Part...*



Click on a blank area in the geometry window to place the part into the assembly--the mouse the pointer will have a green tick next to it to indicate this is required.

4. Edit the new part within the assembly by right-clicking on it and selecting Edit Part.



This will make the selected part the active part. By default, all non-active parts in an assembly will become transparent. The transparency of the non-active parts can be controlled by using the Assembly Transparency button in the *Features* ribbon.



The edges and faces of inactive parts can be selected and used with the active part.

The domain extents will be created below.

5. Create a sketch, with suitable dimensions for the analysis, on the part's symmetry plane.



6. On completion, extrude the sketch to an appropriate thickness for the case.



We will now subtract the part from the fluid domain.

7. Use the Cavity feature, *Insert* > *Features* > *Cavity...*, to subtract the car geometry from the domain.



The *Cavity* panel will appear.

8. Expand the feature tree and select all the inactive parts from here as shown below.



9. Expand the feature tree and select the part to subtract from the fluid domain, as shown below. This part will be added to the *Design Components* group box. The default settings for this operation are suitable.

10. Return to the assembly and suppress the original geometry part. This can be done by right-clicking on the part in the feature tree and selecting Suppress.

The final fluid domain is shown below.



The geometry is now ready for the cad client analysis. Suppressed parts are not included in the analysis. If you make any changes to the car geometry, these will automatically be passed on to the fluid region, giving you an up-to-date fluid domain.

Multi-Region Flow

This section shows how to create the fluid domain for multiple regions such as rotating machinery and porous media flow.

For the numerical assumptions made using the Moving Reference Frame (MRF) technique to hold true, a minimum of two cells need to be present between the rotating-static interface and any wall boundaries (for example the fan blades), and the rotating region must be circular in section. Where there is a large clearance between the rotating object and any static objects, it is advisable to avoid placing the interface too close to either hence facilitating a small cell size on the interface. Where there is a small clearance between the rotating and static parts, there is little choice but to place the interface approximately midway between the two.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an additional part that will enclose the rotating region (for example a fan) and subtract it from the part. In the case of a porous region, a separate part will be required for this as well.

3. Create an additional part for the main static domain (the inlet and outlet regions) and subtract any obstructions/features that may be present there (for example the fan motor support arms on a fan casing).



1. If the geometry to be analyzed is a part file then this will need to be inserted into an assembly. A quick way of doing this is to click on File > Make Assembly from Part.



If the geometry is already an assembly, this step can be skipped.

We will insert a new part into the assembly--this part will form the rotating fluid region. The fan geometry will be subtracted from this part using the Cavity feature in SolidWorks. In a porous media case, this region would form the porous region of the fluid domain.

2. Insert a new part into the assembly, *Insert* > *Component* > *New Part...*



Click on a blank area in the geometry window to place the new part. SolidWorks will indicate that this is required by having a green tick symbol next to the mouse pointer.

3. Edit the new part within the assembly by right-clicking on it in the feature tree and selecting Edit Part.



This will make the selected part the active part. By default, all inactive parts in an assembly will become transparent. The transparency of inactive parts can be controlled by clicking on the Assembly Transparency button in the *Features* ribbon.



The edges and faces of non-active parts can be selected and used with the active part.

4. If the assembly is transparent, change it to opaque as previously described.
5. Hide the rotating component (the fan) by right-clicking on it in the workspace and selecting Hide Components.

To make a part visible again you can right-click on it in the feature tree and select Show Component. All parts, even inactive parts, can be shown or hidden using this method.

6. Create a sketch on the contact face of the fan casing (where the fan and the fan casing are coincident).



7. Draw a circle that shares the same center point as the fan casing. The circle's edges should be coincident with the fan casing's inner wall. To do this, click on the circle sketch and on the innermost edge on the fan casing, as shown below, and set the constraint to Coradial.



8. Hide the fan casing and show the fan again.

9. Extrude the sketch until it completely encloses the fan.



We will now subtract the fan from the rotating region.

10. Use the Cavity feature, *Insert* > *Features* > *Cavity...*, to subtract the fan geometry from the rotating region.



The *Cavity* panel will appear.

11. Expand the feature tree and select the fan part from here as shown below. The selection will be added to the *Design Components* group box.



12. Hide the fan part so that only the rotating region remains visible. The fan part will no longer be required from this point.
13. Right-click on the rotating region part and select Change Transparency. The fan has been subtracted from the rotating region.



We will add a new part to the assembly--this will represent the inlet region. The diameter of this part will match the diameter of the rotating region.

14. Insert a new part into the assembly and edit it, as previously described.
15. Set the Assembly Transparency to Opaque.
16. Create a sketch on the rotating region's front face.
17. Select the Convert Entities function from the toolbar and click on the rotating region's outermost edge.

This tool projects the selected edges of an existing part onto the current sketch. These projections will automatically update to match any changes made to the parent part.

18. Extrude the sketch to a suitable length.



19. Return to the assembly.

We will add a new part to the assembly--this will form the outlet region. In this example, the support arms of the fan casing will need to be subtracted from the outlet region.

20. Insert another part into the assembly and edit it.
21. Create a sketch on the back face of the rotating region and convert the outermost edge, as described in the previous steps.

22. Extrude the sketch to a suitable length.



23. Show the fan casing.
24. Use the Cavity feature to subtract the fan casing from the outlet region, as described previously.
25. Return to the assembly.

The parts which are not required in the simulation (the fan and fan casing), should be suppressed.

26. Right-clicking on the parts that are not required in the feature tree and select Suppress.

A screenshot of the final fluid domain is shown below. The parts were made transparent to show the cavities of the fan and fan casing.



A cross-section that cuts through all the parts is shown below. The fan is shown in red, the fan casing is shown in white, the inlet and outlet regions are shown in grey, and the rotating region is shown highlighted in blue.



The geometry is now ready for cad client analysis. Suppressed parts will not be included in the analysis. The three regions that form the fluid domain all have a conformal match, so it will be possible to create the interfaces and obtain a topologically valid mesh in STAR-CCM+ without any further modifications to the geometry.

Imprinting Parts

Most CAD model assemblies are not suitable for meshing as the parts in the assembly are disjointed, where the coincident faces are overlapping. Attempting to generate a mesh with such an assembly will result in a meshing error. To rectify this problem, the parts in an assembly must be imprinted together.

Imprinting can also be used to split surfaces for boundary region applications. An example of such an application would be to model the area that a fan would occupy on the inside of computer case, such as in the third STAR-Works tutorial, where a smaller section of the surface is needed to represent the fan, rather than the entire inner surface of the computer case.

Imprinting Parts in an Assembly

In cases containing multiple adjacent parts, such as in Conjugate Heat Transfer cases, it is essential to imprint the profiles of the coincident faces of the smaller parts onto the face of the larger part. This will ensure that the volume mesh created by STAR-CCM+ contains interfaces between the adjacent parts.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work within the assembly, making the part that requires imprinting the work part.
2. Create a sketch on the face you want to imprint, and project the contact edges of the other components in the assembly onto this sketch.
3. Split the face you want to imprint using the sketch created in the previous step.

The simple circuit board shown below requires imprinting before the mesh can be generated.



In this example, the capacitors and heat sinks are not imprinted to the circuit board. This would result in a failure when attempting to generate a mesh.

To imprint the capacitors and heat sinks onto the circuit board follow the steps outlined below:

1. Load the assembly file in SolidWorks.
2. Edit the part that needs to be imprinted (in this case the circuit board) by right-clicking on it and selecting Edit.



This will make the circuit board the active part. By default, all non-active parts in an assembly will become transparent. The transparency of the can be controlled by clicking on the Assembly Transparency button in the *Features* ribbon.



The lines and faces of non-active parts can be selected and used with the active part.

3. Create a new sketch on the surface that requires imprinting.
4. Use the Intersection Curve tool from the toolbar and select all the parts that need to be imprinted onto the circuit board.



A quick way of selecting these parts is through the feature tree. Once selected, these parts will become highlighted in blue.



The result of this operation is shown below.



The Intersection Curve feature will automatically find any intersections from the selected parts and project the profiles onto the current sketch.

5. Use the Split Line... function, *Insert* > *Curve* > *Split Line...*, to imprint the sketch containing the intersection curves onto the face.



The Split Line dialog will appear.

6. In the *Type of Split* group box, ensure that Projection is selected.
7. Add the sketch that contains the profiles for the split in the Sketch to Project field.
8. Add the face you want to imprint in the Faces to Split field.

The completed *Split Line* dialog is shown below.



9. Return to the assembly and click Rebuild. It is good practice to rebuild assemblies that have been modified.

Completing the operation will imprint the base board with the profiles of the capacitors and heat sinks. Hiding the aforementioned parts shows the imprinted edges.



A high quality mesh with fully defined interfaces can now be successfully generated for this assembly. Suppressed parts are not included in the STAR-CCM+ analysis, so for CHT cases, such as the above example, ensure that all relevant parts are active.

Imprinting Surfaces on Parts for Boundary Region Applications

Faces in a part body can be split into smaller faces for boundary application.

The general procedure to follow is:

1. Create a sketch on the surface you want to split. The shape of the sketch should correspond to the shape of the new face after it is split.
2. Split the face that requires imprinting with the sketch containing the profiles.

A simple example of an aeration tank is shown below. The image shows the tank, the two air injectors need to be imprinted onto the base of the tank.



1. Load the part file into SolidWorks.
2. Create a sketch on the face you want to imprint.

3. Draw the profile of the boundaries you want to create. In this example the air injectors contain a circular profile.



4. Use the Split Line... function, *Insert* > *Curve* > *Split Line...*



The Split Line dialog will appear.

5. Add the sketch containing the profile(s) to be imprinted in the Sketch to Project field.
6. Add the face to be imprinted in the Faces to Split field.



Completing this operation will imprint the selected face with the profiles in the sketch. The new faces can be selected independently of each other and can be used when defining boundary conditions. The two new faces are highlighted below.



The separate faces can now be defined as boundary regions.

Working With STAR-CAT5

This part of the documentation details how to create, set up and run a CFD simulation in STAR-CAT5. It assumes that you have previous experience in using CATIA V5, and only describes the additional features available as part of STAR-CAT5. The following sections are included:

- Introducing STAR-CAT5
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-CAT5 Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-CAT5 Settings
- Best Practices Guidelines

You use STAR-CAT5 to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-CAT5

STAR-CAT5 provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the CATIA V5 environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-CAT5
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-CAT5 enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-CAT5 provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-CAT5 interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

  - Ideal Gas
- Porous Media
  - Isotropic
  - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-CAT5 provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the CATIA V5 environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
  - Tetrahedral core mesh generation
  - Polyhedral core mesh generation
  - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-CAT5 allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-CAT5 supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-CAT5 interface, and basic concepts. Topics include:

- Launching STAR-CAT5
- The STAR-CAT5 Info App Window
- Modifying the config.xml File
- Understanding the STAR-CAT5 Interface
- The STAR-CAT5 Object Tree
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-CAT5 Settings While the STAR-CCM+ Client is Running
- Manually Creating a STAR-CAT5 Shortcut

Launching STAR-CAT5

The steps for launching STAR-CAT5 are as follows:

1. Start CATIA V5 using the STAR-CAT5 desktop icon created by the STAR-CCM+ installer. For version 8.0x.0xx, this will be `STAR-CAT5 Rxx 8 0x 0xx`. In the unlikely event that a desktop shortcut was not created during the installation process, you can create a new shortcut manually.
2. Open the part or product you want to analyze using the `File > Open` menu.

3. Select Start > Analysis & Simulation > STAR-CAT5.



The STAR-CAT5 tree is located on the top-left of the window, and the *STAR-CAT5 Info App* window launches.

The STAR-CAT5 Info App Window

The *STAR-CAT5 Info App* window automatically appears when you start STAR-CAT5. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-CAT5. More tabs appear as you create or load geometry files in STAR-CAT5, allowing you to perform more actions.



You cannot close the *STAR-CAT5 Info App* window while CATIA V5 or STAR-CAT5 are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-CAT5 license checked out successfully. The version of STAR-CAT5 is also displayed in this tab. This tab appears by default when you launch STAR-CAT5 without a geometry file.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-CAT5.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-CAT5 searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-CAT5 user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-CAT5 is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-CAT5 Info App. This can be set to one of the following:
  - `Normal` — This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose` — This will display a moderate amount of information in the Output panel.
  - `Log` — This will display all available information in the Output panel.

Understanding the STAR-CAT5 Interface

The STAR-CAT5 user interface is embedded within the CATIA V5 environment, with the addition of the STAR-CAT5 feature tree, menu, and toolbar.



The majority of the STAR-CAT5 operations are executed within the feature tree and toolbar. Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-CAT5 interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings
- Region settings
- Add new boundary
- Boundary conditions

- Mesh Pipeline settings
- Add new mesh controls

Toolbar

The operations provided in the toolbar include:

- Generate Surface Mesh
- Generate Volume Mesh
- Check Model For Errors
- Build and Run
- Run Solver
- Run For More Iterations
- Show/hide Mesh
- Transfer CAE Model
- Transfer Geometry

Menu

The operations provided in the menu include:

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Transfer CAE Model
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
- Solution
  - Build and Run
  - Run Solver
  - Run For More Iterations
  - Clear Solution

- Check Model For Errors
- Export
  - Surface Mesh (.dbs)
  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Transfer Geometry
- STAR-CAT5 Help
- About STAR-CAT5

The STAR-CAT5 Object Tree

STAR-CAT5 uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-CAT5. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-CAT5 applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-CAT5 Interface section.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select *STAR-CAT5* > *Preferences* from the menu.



The *Preferences* dialog appears.

2. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

3. To run STAR-CAT5 in parallel, select Parallel from the *Run Mode* box. The *Parallel Options* box appears.



4. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.
5. Click OK to close the dialog.

   If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-CAT5 and STAR-CCM+ need to be closed for the changes to take effect.



6. Click OK to proceed.
7. Close the active STAR-CAT5 and STAR-CCM+ sessions.
8. Restart STAR-CAT5.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-CAT5.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-CAT5.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-CAT5.

To change the units system:

1. Bring up the *Choose Units System* dialog appears:
   - Right-click the STAR-CAT5 > CAE Model Case > Solution node within the object tree and select Units System.



- Alternatively, from the menu, select *STAR-CAT5* > *Settings* > *Units System*.



The *Choose Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click OK to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers

- Angstroms
- Nanometers
- Mils
- Microinches
- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-CAT5 environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the toolbar.

- Alternatively, from the menu, select *STAR-CAT5* > *Transfer CAE Model*.



The STAR-CCM+ version affiliated with STAR-CAT5 launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Editing STAR-CAT5 Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-CAT5 while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-CAT5, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Manually Creating a STAR-CAT5 Shortcut

If a STAR-CAT5 shortcut is not created by the installer, you can manually generate one yourself by following the steps outlined below.

1. Open a Command Prompt.

   To launch Command Prompt in Windows XP:

   - Click on Start > Run.
   - Type `cmd` in the *Run* dialog and press the <return> key.

   To launch Command Prompt in Windows Vista and Windows 7:

   - Click on Start.
   - Type `cmd` in the search field at the bottom of the *Start* menu and select `cmd.exe` from the search results.

2. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

```
<setcatenv> -e "<shortcutname>" -p
"<catiainstalldir>;<starcadinstalldir>" -a global -icon no
-menu yes -cs CATIA
```

where:

- `<setcatenv>` is the full path to `setcatenv.exe`. The default location for a 64-bit installation is:

  ```
  C:\Program Files\Dassault Systemes\Bxx\win_b64\code\
  bin\setcatenv.exe
  ```

- `<shortcutname>` is the name of the new shortcut that will be generated.
- `<catiainstalldir>` is the full path to the CATIA V5 Bxx directory. The default path is:

  ```
  C:\Program Files\Dassault Systemes\Bxx
  ```

- `<starcadinstalldir>` is the full path to the STAR-CAT5 Bxx directory. The default path for a 64-bit installation is:

  ```
  C:\Program Files\CD-adapco\STAR-CAD7.0x.0xx\STAR-CAT
  5\win64\bin\Bxx
  ```

  where `x` is the relevant version of your installation.

An example code is shown below. In this case a new shortcut for a 64-bit installation of STAR-CAT5 7.06.006 is shown. In this case the shortcut name is `STAR-CAT5_R21_TEST`.

```
"C:\Program Files\Dassault Systemes\B21\win_b64\code\bi
n\setcatenv.exe" -e "STAR-CAT5_R21_TEST" -p
"C:\Program Files\Dassault Systemes\B21;C:\Program File
s\CD-adapco\STAR-CAD7.06.006\STAR-CAT5\win64\bin\B21"
-a global -icon no -menu yes -cs CATIA
```

A new shortcut will be generated and placed in the *Start* menu of CATIA V5, the default location is *Start* > *CATIA*.

Common Functions

The following section details common operations used within the STAR-CAT5 environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-CAT5
- Removing surfaces and parts from objects in STAR-CAT5

Renaming an Object

To rename an object:

1. Right-click the object node in the object tree, and select <entity> object > Definition...



The settings dialog for the object appears.

2. Enter the name in the Name box.
3. Click OK to close the dialog.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, the Delete option is disabled in the pop-up menu.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select <entity> object > Definition...



The settings dialog for the object appears.

2. Make the changes that you want.
3. Click OK to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-CAT5

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-CAT5 object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

- On the CAD geometry, click the surfaces that you want to assign to the object. You can add multiple surfaces by clicking more surfaces on the CAD geometry.

The selected surfaces are highlighted on the CAD geometry. By default, CATIA only highlights the edges of selected surfaces. This behavior can be changed so that the surface is highlighted as well, as is shown in the following screenshot.



The number of selected surfaces are displayed in the *Supports* box of the active dialog.

Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to assign to the object. You can add multiple parts by clicking more parts on the CAD geometry.

The selected part is highlighted on the CAD geometry. By default, CATIA only highlights the edges of selected parts. This behavior can be changed so that the surfaces are highlighted as well, as is shown in the following screenshot.



The number of selected parts are displayed in the *Supports* box of the active dialog.

Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-CAT5

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-CAT5 object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

- On the CAD geometry, click the face that you want to remove.

Removed surfaces are no longer highlighted on the CAD geometry. Additionally, you can use the *Supports* box to monitor how many surfaces are currently assigned to a particular object.

Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to remove.

Removed parts are no longer highlighted on the CAD geometry. Additionally, you can use the *Supports* box to monitor how many parts are currently assigned to a particular object.

Transferring Geometry to STAR-CCM+

With STAR-CAT5, you can transfer your CAD geometry directly to STAR-CCM+ from within the CATIA V5 environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

To transfer the geometry to STAR-CCM+:

1. Select *STAR-CCM+* > *Transfer Geometry.*



STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-CAT5. This action ensures that the simulation file remains associated with the CATIA V5 file. You can also safely rename and move the file using this method.
- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the CATIA V5 file.

Resuming a Simulation

To resume a simulation after closing STAR-CAT5 and STAR-CCM+:

1. Launch STAR-CAT5.

2. Open the CAD geometry in STAR-CAT5.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click STAR-CAT5 > Choose Design Parameters.



The *Design Parameters* dialog appears.



2. Select the design parameters that you want to transfer from the list.You can select multiple items at once.

Note: STAR-CAT5 only recognizes native CAD features as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Transfer to STAR-CCM+ option.
4. Click OK.
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-CAT5, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-CAT5, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

The geometry updates according to the design parameter changes.



6. Return to CATIA V5. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-CAT5 and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

Updating the Geometry in STAR-CCM+ with CATIA V5 Closed

You can close CATIA V5 and continue working with the simulation in STAR-CCM+. Doing so, however, breaks the link between the two software packages. Any further updates that you make to the geometry in STAR-CCM+ are not transferred back to the geometry in CATIA V5. Instead, these changes are stored in memory, which allows you to include the design parameters in your optimization scripts.

This workflow also allows you to release the CATIA V5 and STAR-CAT5 licenses when performing tasks in STAR-CCM+. In this case, the CATIA V5 and STAR-CAT5 licenses are temporarily checked out only when the updates to the geometry are stored in memory.

Transferring the geometry from STAR-CAT5 will overwrite the geometry in STAR-CCM+.

Transferring CATIA Publications to STAR-CCM+

You can transfer published entity names from CATIA V5 to STAR-CCM+ using the Transfer Geometry feature of STAR-CAT5. When you transfer geometry that contains CATIA publication data, the published names are transferred instead of the default entity name.

You can transfer the published names for the following entities:

- Surfaces—by default, surface names are extracted from the last feature that created them, for example `Pad.1`. Transferring the published name allows you to control the names of the surfaces.
- Design Parameters—by default, design parameters are transferred with a generic name that describes what they are, for example `Length.19`. If the design parameter is published, then the name of the published design parameter is transferred to STAR-CCM+. In STAR-CCM+, design parameter names are concatenated with the part or assembly name that they come from, for example `Part\Design Parameter`.
- Features—by default, the feature name is transferred to STAR-CCM+. If the feature contains published data, then this name is transferred instead.
- Bodies—by default, body names are not used in STAR-CCM+. When a part contains more than one body, the bodies are transferred to STAR-CCM+ as parts inside a composite part. The published name of these parts will be transferred to STAR-CCM+.

If the same entity, feature, or body is published more than once in the assembly tree, then the top-most publication is used for that entity, feature, or body respectively.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from CATIA V5 to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-CAT5 root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click OK.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.
Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.
Transfer Hidden Geometry	Specifies whether to transfer hidden parts. By default, hidden parts in CATIA V5 are not transferred.
Transfer Sheet Geometry	Specifies whether to transfer sheet bodies as part of the transfer process. This option is turned off by default.

Understanding the STAR-CAT5 Process

In STAR-CAT5 the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume

3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity
5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings
7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-CAT5 and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-CAT5

STAR-CAT5 only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-CAT5.

Suppressed or deactivated parts and features are not supported by STAR-CAT5 and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-CAT5.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

  For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-CAT5 includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-CAT5 includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-CAT5 is completely integrated within the CATIA V5 environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-CAT5 model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

  Including material type and properties.
- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.
- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.

- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-CAT5 automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select Add Continuum.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



- Right-click on the Regions manager node in the object tree and select Add Region.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-CAT5). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-CAT5 model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. Then right-click on the region node under the Region manager node in the object tree and select <region name> object > Definition... In the *Region Settings* dialog select the alternative continuum using the *Select Continuum* drop-down menu and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-CAT5 each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By default the material type is set to Air, and this can be changed to any one of a number of predefined solids, liquids or gases that are available in STAR-CAT5. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-CAT5 object tree and select <continuum name> object > Definition....



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material*, you can select a material type.
- In the *Material Properties*, you can define each of the material properties.



Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-CAT5 offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the drop-down menu. When you have selected a material, the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-CAT5 are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.
Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-CAT5, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)

- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select <continuum name> object > Definition...



The physics models available in the *Continuum Settings* dialog are determined by the material type.



Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.



Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.



The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-CAT5 are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Gravity* checkbox. This will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the *Magnitude* box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. To define reference properties for a continuum, right-click the Continuum node in the object tree and select [continuum_name] object > Definition... In the *Continuum Settings* dialog, click on the Reference tab.



The reference properties are described below:

- Reference Pressure

  The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density, Reference Velocity, Reference Area.

  These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-CAT5, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select <region name> object > Definition...



The physics models can then be enabled for the region by ticking the corresponding checkbox and entering values as necessary.



Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the *Porous* checkbox.

This activates the Porous tab. There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select Isotropic from the *Type* menu.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, Alpha and Beta respectively.

Uni-Directional

To specify uni-directional porous resistance, select Uni-Directional from the *Type* menu.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, Alpha and Beta respectively. The cross-flow resistance is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, `[i, j, k]`.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the *Moving Reference Frame (MRF)* checkbox.



In the *MRF* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the Omega box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, `[i, j, k]`. The units used are the same as the file defaults.

The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.

This action activates the Energy Source tab. There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select Per Unit Volume from the *Energy Input* menu.



Enter the value for the energy input value per unit volume in the Energy Input box.

Total Energy Input

To define the total energy input, select Total from the *Energy Input* menu.



Enter the value for the total energy input in the Energy Input box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.

2. Create a streamlines derived part.
3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-CAT5.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-CAT5 will use default values to initialize the solution. In STAR-CAT5, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-CAT5 object tree and select <continuum name> object > Definition....



This opens the *Continuum Settings* dialog, in which the *Initialization* tab can be used to specify the following values:



- Initial velocity
- Initial pressure
- Initial temperature
- Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial velocity, open the *Continuum Settings* dialog, click on the Initialization tab, and enter the values in the U, V and W Component boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, click on the Initialization tab and enter a value in the Pressure box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, click on the Initialization tab, and enter a value in the Temperature box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, click on the Initialization tab and enter a value in the Turb. Intensity and Turb. Length Scale boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-CAT5 boundary is made up of one or more CAD surfaces within the simulation which, along with other boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select Add Boundary from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-CAT5). When you assign a surface to

another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-CAT5, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Velocity Inlet from the drop-down menu.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity, with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors U, V, W, tick the checkbox labeled Specify velocity components.



Cylindrical vectors R, Theta, Z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i, j, k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Mass Flow Inlet from the drop-down menu.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.



The mass flow rate must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors i, j, k, tick the checkbox labeled Specify Components.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Stagnation Inlet from the drop-down menu.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure, must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors i, j, k, tick the checkbox labeled Specify Components.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-CAT5, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Flow-Split Outlet from the drop-down menu.

The only parameter required is the flow-split ratio, expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.

Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Pressure Outlet from the drop-down menu.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.



The pressure must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Wall from the drop-down menu.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U, V, W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i, j, k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature and heat transfer coefficient.



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Boundary Type* box and select Symmetry Plane from the drop-down menu.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-CAT5 to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-CAT5 and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature*.

Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Turbulence*.



Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity
- Turbulence Length Scale

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Passive Scalar*.

This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-CAT5 provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-CAT5* menu), it is possible to execute only single stages from the STAR-CAT5 > Mesh menu. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the CATIA V5 parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from CATIA V5.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the CATIA V5 *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-CAT5 creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per

boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-CAT5 appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-CAT5 appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all CATIA V5 parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose <mesh pipeline name> object > Definition...


STAR-CAT5
Nodes and Elements
CAE Model Case
Solution
Continuum
Region
Boundaries
Default
Wall 1
Mesh Pipeline
Sensors
Center graph
Reframe On
Hide/Show
Properties Alt+Enter
Open Sub-Tree
Cut Ctrl+X
Copy Ctrl+C
Paste Ctrl+V
Paste Special...
Delete Del
Mesh Pipeline object
Definition...

The *Mesh Pipeline Settings* dialog appears.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click OK to save them and close the dialog.
- Click Cancel to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.

Selecting the trimmer option will activate the Trimmer Mesh Settings tab.

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-CAT5 Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.

The base size can be set as either a relative or absolute value. When it is set as a relative value (by clearing the Use Absolute Value option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* box of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.



This panel allows you to change the following parameters:

- Relative Minimum Surface Size
- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-CAT5 will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer Mesh Settings panel.



The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the CATIA V5 parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when they are created. A part can only belong to one mesh pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select Add Mesh Pipeline.



During the creation of a new mesh pipeline, it is necessary to assign to it the CATIA V5 parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

This mesh control is applied to one or more faces of the CATIA V5 model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose Add Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select Add Volumetric Mesh Control from the pop-up menu.



The *Volumetric Control Settings* dialog appears.



The *Volumetric Control Settings* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape
- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the CATIA V5 geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.

X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the CATIA V5 window. To use an existing point, click the Choose Point 1... or Choose Point 2... button and use the CATIA V5 point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Choose Point... button and using the CATIA V5 point selection procedure to choose an existing point in the CATIA V5 window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select Add Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the Custom Mesh Size checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-CAT5* toolbar.

- Alternatively, select Mesh > Generate Surface Mesh from the *STAR-CAT5* menu.



Output from the surface mesher will be displayed in the *STAR-CAT5 Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-CAT5* toolbar.

- Alternatively, select Mesh > Generate Volume Mesh from the *STAR-CAT5* menu.



Information on the progress of the volume mesher will be displayed in the *STAR-CAT5 Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Mesh > Clear Mesh from the *STAR-CAT5* menu.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-CAT5* toolbar.

Alternatively, select Mesh > Show/Hide Mesh from the *STAR-CAT5* menu.



Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-CAT5 will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-CAT5 Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Running in Batch Mode
- Clearing the Solution
- Visualizing the Solution
- Setting Up Sensors

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the Solution node within the object tree and select Solution Controls.



- Alternatively, from the menu, select *STAR-CAT5* > *Settings* > *Solution Controls*.



The *Solution Controls* dialog appears.



2. Here, you can:

- Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.
- Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.

3. Click OK to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:

1. In the *Solution Controls* dialog, click the More... button.

   The advanced options appear.



1. In the Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:

- First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.
- Second Order—This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:
   - Momentum
   - Pressure
   - Turbulence (with turbulence modelling defined in Continuum Settings)
   - Temperature (with temperature modelling defined in Continuum Settings)
     - Fluid Temperature Under-Relaxation Factor
     - Solid Temperature Under-Relaxation Factor
5. If you want to use additional STAR-CCM+ physics models with STAR-CAT5, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros
6. Click OK to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-CAT5 by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-CAT5:

1. Open the *Solution Controls* dialog.
2. Expand the advanced solver options.
3. Activate the Use Java File option.

   A text entry field appears.



4. Click the Select Input File button. This action opens a browser dialog.
5. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
6. Click OK to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-CAT5.

Using STAR-CCM+ Continuum Physics with STAR-CAT5

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-CAT5 and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.
2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.
5. Return to STAR-CAT5, open the *Solution Controls* dialog and activate the Use Java File checkbox.

6. Select the Java macro that you created earlier.
7. Run the case from STAR-CAT5.

Checking the Model for Errors

STAR-CAT5 monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.
- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.



2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-CAT5 object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-CAT5 root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check Model for Errors in the *STAR-CAT5* menu.



- Alternatively click the Check Model for Errors button in the *STAR-CAT5* toolbar.

If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-CAT5 Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-CAT5, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-CAT5* menu.



- Alternatively click the Build and Run button in the *STAR-CAT5* toolbar.

This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-CAT5 Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Solution > Run Solver in the *STAR-CAT5* menu.



- Alternatively, click the Run button in the *STAR-CAT5* toolbar.

STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in

the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Solution > Run for More Iterations in the *STAR-CAT5* menu.



- Alternatively you can click the Run for More Iterations button in the *STAR-CAT5* toolbar.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Running in Batch Mode

To run the solver in batch mode:

1. Choose or create a directory from which to launch STAR-CAT5 in batch mode. This is referred to as the Batch Runtime Folder.

2. Set up the desired case in STAR-CAT5, assigning all relevant continuum, region, boundary, mesh pipeline, and solution control settings.
3. Place the geometry file in the Batch Runtime Folder.
4. Create a text file and copy the following sample Visual Basic macro into it:

```
Language="VBSCRIPT"

' Main function - the role of this function is to load
in the document and run the analysis

Sub CATMain()

' Open the analysis model

Set analysisDocument1 =
CATIA.Documents.Open("C:\Batch Runtime Folder\Analysis1
.CATAnalysis")

' Run STAR-CAT5

RunSTARCAT5(analysisDocument1)

analysisDocument1.Save()

analysisDocument1.Close()

CATIA.Quit()

End Sub


' This subroutine runs STAR-CAT5 given an analysis
document

Sub RunSTARCAT5(analysisDocument1)

' do not edit this - the following code determines the
structure of the analysis and provides

' access to the meshes and the run objects

Set analysisManager1 = analysisDocument1.Analysis

Set analysisModels1 = analysisManager1.AnalysisModels

Set analysisModel1 = analysisModels1.item(1)

Set meshManager = analysisModel1.MeshManager

Set meshes = meshManager.AnalysisMeshParts

Set analysisCases1 = analysisModel1.AnalysisCases

Set analysisCase = analysisCases1.item(1)

' end of no need to change

' Create the surface mesh

meshes.Item(1).update

' Create the Volume Mesh
```

```
meshes.Item(2).update
' Run the calculation'
analysisCase.compute
End Sub
```

where:

- `analysisDocument1.save()` saves the analysis file, overwriting the current file. If you do not want to save the analysis file, delete this line from the script.
- `Set analysisDocument1 = CATIA.Documents.Open` should point to the saved `CATAnalysis` file in the Batch Runtime Folder.

5. If you wish to edit or modify the part that the analysis is to be run on, you can add the following additional code:

```
` Open the linked document for modification
Set analysisManager1 = analysisDocument1.Analysis
Set linkedDocuments = analysisManager1.LinkedDocuments
Set linkedDocument1 = linkedDocuments.Item(1)
```

where:

- `linkedDocument` will be either a `CATPart` or a `CATProduct` document. From here you can access the PartBody or subproducts to modify any settings.

6. Save the text file as a `.CATvbs` macro file and exit.
7. Place the macro file in the Batch Runtime Folder.
8. Launch a Command Prompt.
9. In the Command Prompt, change to the Batch Runtime Folder that contains the CAD geometry, the CATAnalysis file, and the Visual Basic macro.

   An example code is shown below:

```
C:\>cd "C:\Batch Runtime Folder"_
```

10. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

```
"<catia_install_dir>\B<catia_version>\win_b64\code\bin\
CNEXT.exe" -env STAR-CAT5_R<catia_version>_<star-ccm+_v
ersion> -macro "mymacro.CATvbs" -batch
```

where:

- `<catia_install_dir>` is the installation directory of your CATIA V5 release. The default path is:

  `C:\Program Files\Dassault Systemes`

Note: The full path to the `CNEXT.exe` executable is required.

- `<catia_version>` is the version of CATIA V5 that you want to run. Refer to the release notes to see which versions of CATIA V5 are supported.
- `<star-ccm+_version>` is the version of STAR-CCM+ that you want to use.
- `-env` is the CATIA V5 environment that STAR-CAT5 is run from.
- `"mymacro.CATvbs"` is the name of the Visual Basic macro. If this file is not in the Runtime Batch Folder, the full path is required.
- `-batch` tells CATIA V5 to run in batch mode.

An example code for running STAR-CAT5 9.06.004 in batch mode using CATIA V5 R21 is shown below.

```
"C:\Program Files\Dassault Systemes\B21\win_b64\code\bi
n\CNEXT.exe" -env STAR-CAT5_R21_9_06_004 -macro "mymacr
o.CATvbs" -batch
```

11. Press <`Enter`> to run the simulation in batch mode.

STAR-CAT5 will work in the background, meshing the geometry and running the solvers. A simulation file will be saved in the Batch Runtime Folder once the solver has finished running. If there is an existing simulation file in the Batch Runtime Folder, a backup file is created with a `.sim~` extension. A log file (`starcad.log`) is also created and saved in the Batch Runtime Folder—this log file contains a record of all the messages that appeared in the command prompt during batch processing.

Launching a Command Prompt

To launch a command prompt in Windows 7:

1. Click Start.
2. Type `cmd` in the search field at the bottom of the *Start* menu and select `cmd.exe` from the search results.

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-CAT5 case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the menu, select *STAR-CAT5* > *Solution* > *Clear Solution*.



The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Setting Up Sensors

You can define sensors on boundaries in STAR-CAT5. A sensor is an object that can be used to monitor solution data that is created by STAR-CCM+. The sensor creates a CATIA parameter which can then be used within optimization loops.

Defining Sensors

You can create a sensor either before or after running a case. If you create it before running the case, the sensor will be evaluated when the case is run. If you create it after running the case, you will need to update the sensor.

To create a sensor:

1. Right-click Sensors.1 in the object tree and select Create CFD Sensor.



The *Create CFD Sensor* dialog appears.



2. Select the type of report and the quantity to monitor.
3. Click Add.

A sensor object is created and added to the object tree.

Note: You can create multiple objects by clicking Add multiple times.

4. Click OK to close the dialog.

   The sensor object contains an ! icon—this icon indicates that the sensor definition is incomplete.

5. Double click the sensor to edit it.

   The *Sensor Properties* dialog appears.



6. Rename the sensor to something more appropriate in the Sensor Name property.
7. Select the boundaries that you want to associate with the sensor from the STAR-CAT5 object tree.
8. Click OK.

The sensor icon in the object tree will show if it requires an update. The sensor will be evaluated when the case is run or when an update is requested.

To update the sensor:

1. Right-click on the sensor in the object tree and select Update Sensor.



2. Click OK to dismiss the warning dialog.

The data for the sensor is calculated and a new report sub-node is added to the sensor.

Note: The units that are used in the sensor are CATIA working units, not STAR-CAT5 units.

To view the sensor data after the solver has finished running:

1. Double-click on the new sensor report object in the feature tree.



The *Edit Parameter* dialog appears, displaying the report value.



2. Rename the report.
3. Click OK to close the report.

Working with Parameters

You can use CATIA's formula editor to add numerical data together. You can find this feature in CATIA's Knowledge toolbar.

To create a formula:

1. Click the Formula button.

   The *Formulas* dialog appears. In the following example, pressure loss is calculated using the `Area Averaged Pressure` at the Inlet and Outlet.



2. Next to New Parameter of type, select Pressure from the drop-down list.
3. Click on New Parameter of type to add a new pressure parameter to the list.
4. Give this new parameter a suitable name, such as `Pressure Loss`.
5. Click Add Formula.

The *Formula Editor* dialog appears.



6. Enter the desired formula, in this case the pressure at the Inlet - the pressure at the outlet.



Note: You can add sensor report objects to the formula by double-clicking on them in the Members of All list.

7. Click OK to close the *Formula Editor* dialog.
8. Click OK to close the *Formulas* dialog.

Displaying Parameters in the Analysis Feature Tree

You can display parameters that you created in the analysis feature tree:

1. Select *Tools* > *Options*.

The Options dialog appears.

2. Navigate to the *Analysis & Simulation* > *General* tab.
3. Activate the Show Parameters checkbox.
4. Click OK to close the *Options* dialog.

   A new node, Parameters, is added to the feature tree.
5. Double-click the parameter to view it.



6. Click OK.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-CAT5.

- Exporting a Surface Mesh
- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-CAT5 as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-CAT5 menu, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-CAT5 as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-CAT5 menu, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-CAT5 menu, select Export > Engineering Report (.html, .csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
  - Pressure (Pa)

    - Total Pressure (Pa)
    - Pressure Coefficient
    - Total Pressure Coefficient
    - Turbulent Kinetic Energy (J/kg)
    - Turbulent Dissipation Rate (m^2/s^3)
    - Turbulent Viscosity (Pa-s)
    - Density (kg/m^3)
- Force data, at each boundary:
    - Net Force [Fx, Fy, Fz]
    - Pressure Force [Fx, Fy, Fz]
    - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
    - Coefficient of Net Force [Cfx, Cfy, Cfz]
    - Coefficient of Pressure Force [Cfx, Cfy, Cfz]
    - Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
    - Physics models applied to the continuum
    - Materials properties
    - Turbulence properties
    - Segregated flow properties
    - Reference values
    - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
    - Polyhedral mesher properties
    - Prism layer mesher properties (if activated)
    - Surface remesher properties
    - Reference values

  - Volumetric controls
- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values
- Solver parameters
- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-CAT5 to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-CAT5, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-CAT5 menu, select Export > Simulation (.sim).



Resetting STAR-CAT5 Settings

STAR-CAT5 saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-CAT5 settings:

- From the menu, select Settings > Clear All.



This action returns the geometry to its original state.

Best Practice Guidelines

This section aims to provide guidance and assistance on making a geometry part suitable for CFD analysis using STAR-CAT5. The following items are covered:

- Creating the Fluid Domain:
  - Internal Flow
  - External Flow

  - Multi-Region Flow
- Imprinting:
  - Imprinting Parts in an Assembly
  - Imprinting Surfaces on a Part for Boundary Region Applications

Prior to using these methods, ensure that the components in assemblies are all solid bodies and are in contact with each other. Any gaps between components will need to be closed.

A basic knowledge and understanding of both CAD and CFD are assumed, although most basic functions are covered.

Creating the Fluid Domain

In the majority of CAD models, the fluid domain for analysis must be extracted from the original part or assembly. The following sections demonstrate how to prepare the fluid domain for:

- an internal flow
- an external flow
- a multi-region flow

Internal Flow

In general, the steps to follow in creating the fluid domain for an internal flow are:

1. Work in a product. If the original geometry contains only a single part, create a new product that contains the part.
2. Create additional parts that close off any openings in the original geometry, we will call these parts the “caps”.
3. Create an overall block part that encloses all the other parts, the “enclosure”. This will form the internal fluid volume.
4. Subtract all parts from the enclosure. This will result in a solid block that contains a cavity.
5. Delete the outer solid body--the remaining geometry is the fluid domain.

An example case is shown below. In this case the assembly contains two parts, the valve (shown in red) and the valve housing (shown in blue).



1. Load the part or product into CATIA V5. If the geometry is a part, add it to a product.

2. Insert a new part into the assembly by selecting *Insert* > *New Part*.



3. Click on the root node in the assembly's specification tree to insert the part into that assembly.

A dialog will appear asking whether a new origin for this part should be defined or not. Choose either option.

4. Give the new part an appropriate name, in this example the part is called `Caps`.

We will now close all the openings on the geometry--in this example two caps are needed as there are two openings.

5. Make the Caps part the active part by double-clicking on its Part Body node in the Specification Tree.



The part will become active and the *Part Design* workbench appears. When parts are active these will be highlighted with a blue background in the specification tree.

6. Create a new sketch on a surface that contains an opening, as shown below.



7. In the *Sketcher* workbench, select the outermost edge of the face



and use the Project 3D Elements tool.

This will project the edge of the existing body onto the current sketch. The yellow color of the sketch means that it is associative, so if the diameter of the part changes, for example, the cap will automatically update to match the change.

8. Exit the Sketcher workbench and create a Pad from the sketch. The length of the pad is not important, however ensure that the pad extends away from the geometry. If the pad is going inwards, use the Reverse Direction button in the *Pad Definition* dialog. In this example a Length of `5 mm` is sufficient.



The result of this operation is shown below.



9. Repeat the same steps for the other opening(s).



Note that all the caps should be in the same part.

10. Double-click the root node in the specification tree to return to the *Assembly Design* workbench.

We will now insert a new part into the product. This part will become a composite of all the current parts, in this example the valve, the valve body, and the caps.

11. Insert a new part into the product as described previously. Give this new part a suitable name. In this example, the part was named `Composite Part`.

12. Double-click on the part file to activate it, as shown below.



13. Expand the nodes for all the parts to reveal the PartBody nodes and select them as shown below. Hold down the <`ctrl`> key to select multiple items.



Right-click on one of the highlighted PartBody nodes and select Copy.



14. Paste these into the active Composite Part by right-clicking on it's PartBody node as shown below, and selecting Paste Special...



The *Paste Special* dialog appears.

15. In the *Paste Special* dialog, select As Result With Link and click OK.



This will paste all the selected parts into the active part without breaking the link to the original parts, meaning that if the original parts change, the pasted parts will also change--this is useful when analyzing different configurations in the product such as valve angle in this example.

The original parts are no longer required and will be hidden.

16. Return to the Assembly Design workbench by double-clicking on the root node in the specification tree.

17. Select the valve, valve body, and caps parts in the specification tree. Right-click on one of the nodes and select Hide/Show.



The parts' nodes will become greyed out indicating they are hidden.

We will now insert a new part into the product. This will ultimately become the internal fluid volume. To begin with, we will create an enclosure that completely encloses the existing parts. The caps and valve assembly will then be subtracted from the enclosure which will result in two solid bodies, the outer enclosure and the internal fluid volume. The outer enclosure will then be deleted, leaving only the desired internal fluid volume.

18. Insert a new part into the assembly as described previously and rename it to `Fluid Region`.

19. Activate the Fluid Region part by double-clicking on it.



20. Create a sketch on the top view plane.



21. Draw a rough sketch that encloses the entire assembly as shown below.



Exit the Sketcher workbench.

22. Create a pad from the sketch, ensuring that it encloses the geometry completely.



The pad can be extruded in both directions by ticking the Mirrored extent option in the *Pad Definition* dialog.



The view can be changed to Shading with Edges and Hidden Edges to show the hidden part inside the new pad.

23. Select the copied Body parts (in this case Body.2, Body.3 and Body.4) from the Composite Part specification tree.



and select *Insert* > *Boolean Operations* > *Remove...*



This will remove the Composite Part from the Fluid Region using a boolean operation. The result will be a cavity. In the *Remove* dialog, ensure that the three parts from the Composite Part are listed in the Remove field and that the From field contains the Fluid Region's PartBody as shown below.



Click OK to close the dialog.

The Fluid Region part now contains a cavity separating the internal fluid volume from the bulk solid.

24. Hide the Composite Part.

The outer bulk solid will be removed using the Remove Lump operation, resulting in only the internal fluid volume.

25. Select *Insert* > *Boolean Operations* > *Remove Lump*.



26. Click on any surface on the outer solid to bring up the *Remove Lump Definition* dialog.

27. Click in the Faces to Remove field and select any surface belonging to the outer body.



28. Click in the Faces to Keep field and select an internal face on the geometry. The face has to be a wetted face, such as the inside surface of the valve housing. To select hidden surfaces use `<Ctrl> + <F11>`.
29. Click OK to close the dialog.

30. Deactivate all parts that you do not want to take part in the analysis, such as the original parts and the composite part, by right-clicking on each node and selecting <Instance name> object > Activate/Deactivate Component.



The internal fluid volume will appear as shown below.



The geometry is now ready for the CAD client analysis. Deactivated parts in an assembly will not be included in the analysis. In the above example, it is recommended that you extend the inlet and outlet regions to appropriate lengths so that a fully developed flow can be achieved before reaching the valve and to help reduce recirculation at the outlet.

External Flow

The generation of a domain for external aerodynamics is generally easier than for internal flow cases but the sizing of the domain is more critical. As a guideline (assuming a specific situation is not being modelled i.e. flow in a wind tunnel) the domain should be sized as follows:

Inlet: 2 x part length from the front of the part

Outlet: 4 x part length from the rear of the part

Sides: 3 x part width from the side of the part

Ceiling: 4 x part height from the top of the part

In general, the steps to follow in creating the fluid domain for an external flow are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an overall block part that encloses all the other parts. The shape and size should correspond to the desired fluid domain.
3. Subtract all parts from the overall block part - the remaining geometry is the fluid domain.

An example case is shown below. In this example a geometry of a car will be used.



1. Load the product file in CATIA. If the geometry is a part, insert it into a product.
2. If the part is symmetrical, this should be cut in half, as shown below, to reduce computational costs during the analysis.



We will now insert a new part into the product. This will become the fluid domain.

3. Insert a new part into the product, *Insert* > *New Part*.



Click on the Assembly root node in the specification tree to insert the part into the product. A dialog will appear asking whether a new origin should be defined for the part or not--you may choose either option.

4. Create a new sketch on any plane. In this example the sketch was created on the symmetry plane of the car.

5. Draw a profile with suitable dimensions--this sketch represents the profile of the fluid domain.



6. Exit the Sketcher workbench and create a Pad from the sketch.

   The length of the pad will correspond to the width of the fluid domain, as shown below.



We will now remove the car from the fluid domain using a boolean operation.

7. Select *Insert* > *Boolean Operations* > *Remove...*



In the *Remove* dialog that appears, select the car's part body for the Remove field and the fluid domain for the From field..



8. Return to the assembly after the operation is complete and hide the car part, right-click the car node in the specification tree and select Show/Hide.

The area around the car in the final domain is shown below.



The geometry is now ready for the cad client analysis. Deactivated parts are not included in the analysis. If you make any changes to the car geometry, these will automatically be passed on to the fluid region, giving you an up-to-date fluid domain.

Multi-Region Flow

This section shows how to create the fluid domain for multiple regions such as rotating machinery and porous media flow.

For the numerical assumptions made using the Moving Reference Frame (MRF) technique to hold true, a minimum of two cells need to be present between the rotating-static interface and any wall boundaries (for example the fan blades), and the rotating region must be circular in section. Where there is a large clearance between the rotating object and any static objects, it is advisable to avoid placing the interface too close to either hence facilitating a small cell size on the interface. Where there is a small clearance between the rotating and static parts, there is little choice but to place the interface approximately midway between the two.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an additional part that will enclose the rotating region (for example a fan) and subtract it from the part. In the case of a porous region, a separate part will be required for this as well.
3. Create an additional part for the main static domain (the inlet and outlet regions) and subtract any obstructions/features that may be present there (for example the fan motor support arms on a fan casing).



1. Load the assembly in CATIA

We will now insert a new part into the product. This will form the rotating region of the fluid domain.

2. Insert a new part into the assembly, *Insert* > *New Part*.



3. Click on the Assembly node in the feature tree to insert the part into the assembly. A dialog will pop up asking whether a new origin should be defined or not. Choose either option.

4. Rename the new part by right-clicking on it in the Specification Tree and selecting Properties.



In the *Properties* dialog enter an appropriate name for the new part (in this case it will be called `Rotating Region`) and click OK to close the dialog.

5. Double-click the Rotating Region part body. The part will become active and the Part Design workbench appears. When parts are active the part node will be highlighted in blue.



6. Hide the fan casing, right-click on the feature tree or on the geometry and select Hide/Show.



7. Create a sketch on the contact face of the fan casing (where the fan and the fan casing are coincident). If the part is a native CATIA geometry then this will be possible by simply clicking on the desired face, as shown below.



If the part is an imported geometry, a new plane will need to be created on the desired face before a sketch can be created.

8. Create a circle with the same center point as the fan casing - the diameter of the circle is not important at this stage as it will be made coincident with the fan casing in the next step.

9. Select the circle (shown in orange) and the inner-most edge on the fan casing (shown in red).



10. Click on the Constraints Defined in Dialog Box button.

The *Constraint Definition* dialog will appear.

11. In the *Constraint Definition* dialog, tick the Coincidence checkbox.



Notice that the sketch automatically snaps to the selected edge on the fan casing. The green color of the sketch entity indicates that it is fully defined.



12. Exit the sketch.
13. Hide the fan casing part and show the fan part (using the Show/Hide option).
14. Create a pad from the sketch so that it encloses the fan, as shown below. The fan can be visualized by changing the view to Shading with Edges and Hidden Edges



We will now remove the fan from this new pad using a boolean operation.

15. Change the view to Shaded with Edges.

16. Select *Insert* > *Boolean Operations* > *Remove...*



17. In the *Remove* dialog, select the fan part (from the tree) for the Remove field and the pad for the From field as shown below.



Click Preview to show a preview of the operation.



18. Click OK to perform the operation and close the dialog.
19. Hide the fan part as it is no longer required.

We will now insert a new part into the product. This will become the inlet region.

20. Return to the assembly by double-clicking on the root node in the specification tree.
21. Insert another part into the assembly and make it the active part as previously described.
22. Enter the sketch workbench and select the front face of the Rotating Region.

23. Select the outermost edge of the face



and use the Project 3D Elements tool.

This will project the selected edge from the existing body onto the current sketch. The yellow color of the sketch means that it is associative, so if the diameter of the part changes (for example), the sketch will automatically update to match the change.

24. Exit the Sketcher workbench and create a Pad of appropriate length, from the sketch.



We will now insert a new part into the product. This will become the outlet region. The fan casing contains the fan motor's support arms--these will need to the removed from outlet region.

25. Return to the assembly.
26. Insert another part in the assembly and make it the active part.
27. Create a sketch on the back face of the Rotating Region part,.
28. Selecting the outermost edge of the rotating region and project it to the current sketch, as described previously.

29. Create a pad, of appropriate length, that extends downstream. If the direction of the pad is pointing the wrong way, use the Reverse Direction button in the *Pad Definition* dialog.



30. Show the fan casing part (right-click on the node in the specification tree and select Show/Hide).

31. Remove the fan casing from the Outlet Region, as described previously steps. A preview of this operation is shown below.



32. Return to the assembly and hide the fan casing.

The final fluid domain is shown below. The view was changing to Shading with edges and hidden edges.



A cross-section of the fluid domain is shown below. Note that all components were made visible.



The fan is shown in red, the fan casing is shown in white, the Inlet and Outlet fluid regions are shown in grey, and the Rotating Region is shown highlighted in orange.

33. Deactivate all parts that you do not want to take part in the analysis, such as the fan and fan casing, by right-clicking on each node and selecting <Instance name> object > Activate/Deactivate Component.



The geometry is now ready for cad client analysis. Deactivated parts will not be included in the analysis. The three regions that form the fluid domain all have a conformal match, so it will be possible to create the interfaces and obtain a topologically valid mesh in STAR-CCM+ without any further modifications to the geometry.

Imprinting Parts

Most CAD model assemblies are not suitable for meshing as the parts in the assembly are disjointed, where the coincident faces are overlapping. Attempting to generate a mesh with such an assembly will result in a meshing error. To rectify this problem, the parts in an assembly must be imprinted together.

Imprinting can also be used to split surfaces for boundary region applications. An example of such an application would be to model the area that a fan would occupy on the inside of computer case, such as in the third

STAR-CAT5 tutorial, where a smaller section of the surface is needed to represent the fan, rather than the entire inner surface of the computer case.

Imprinting Parts in an Assembly

In cases containing multiple adjacent parts, such as in Conjugate Heat Transfer cases, it is essential to imprint the profiles of the coincident faces of the smaller parts onto the face of the larger part. This will ensure that the volume mesh created by STAR-CCM+ contains interfaces between the adjacent parts.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work within the product, making the part that requires imprinting the work part.
2. Extract the contact surfaces from the other components in the product.
3. Sew the extracted surfaces onto the face you want to imprint.

The simple circuit board shown below requires imprinting before the mesh can be generated.



In this example, the capacitors and heat sinks are not imprinted to the circuit board. This would result in a failure when attempting to generate a mesh.

To imprint the capacitors and heat sinks onto the circuit board follow the steps outlined below:

1. Load the product file in CATIA V5.

   In this example, the green circuit board needs to be imprinted.

2. Right-click Base Board and select Define in Work Object.



3. Switch to the *Generative Shape Design* workbench, by clicking on Start > Shape > Generative Shape Design.



4. `Hide` the base board by right-clicking on it in the specification tree and select Hide/Show.
5. Select all the contact faces on the other parts in the product, holding down the <`Ctrl`> key. In this example the bottom surfaces of the capacitors and heat sinks are selected.



We will now extract the selected faces.

6. Click on the Extract button.



Highlighting multiple faces before using the Extract function will automatically perform the same operation for each selection. If the Extract function was used first this would only allow one face to be selected at a time.

7. Switch back to the Part Design workbench, *Mechanical Design* > *Part Design*.



8. `Show` the circuit board and `hide` the capacitors and heat sinks. The extracted faces will appear in yellow and overlap the existing top face of the circuit board.

We will now sew the extracted surfaces onto the face that we want to imprint.

9. Use the Sew function to sew one of the extracted surfaces to the base board.

Ensure that the Simplify Geometry option in the *Sew Surface Definition* dialog is cleared.

Also note that the sew direction arrow should be pointing towards the surface of the body to be sewn.



10. Repeat the same steps for the other extracted surfaces.

The base board face is imprinted with the contact edges of the capacitors and heat sinks, as shown below.



A high quality mesh with fully defined interfaces can now be successfully generated for this assembly. Deactivated parts are not included in the STAR-CCM+ analysis, so for CHT cases, such as the above example, ensure that all relevant parts are active.

Imprinting Surfaces on Parts for Boundary Region Applications

Faces in a part body can be split into smaller faces for boundary application.

The general procedure to follow is:

1. Create a sketch on the surface you want to split. The shape of the sketch should correspond to the shape of the new face after it is split.
2. Fill the sketch so that is becomes a surface.
3. Sew the surface onto the face that requires imprinting.

A simple example of an aeration tank is shown below. The image shows the tank, the two air injectors need to be imprinted onto the base of the tank.



1. Load the geometry part into CATIA V5.

2. Create a sketch on the face that requires imprinting. In this case the air injectors are sketched onto the tank's bottom face.



3. Switch to the Generative Shape Design workbench, by clicking on Start > *Shape* > *Generative Shape Design*.



4. Use the Fill tool to create surfaces from the sketch profiles.

The *Fill Surface Definition* dialog appears.

5. Select the sketch that contains the profiles you want to imprint.
6. Click OK.
7. In the *Multi-Result Management* dialog, select the keep all the sub-elements radio button and click *OK*. The sketches are now filled in with a planar surface as shown below.



8. Switch back to the Part Design workbench, *Mechanical Design* > *Part Design*.



9. Use the Sew function to sew the surfaces to the bottom face of the tank.

Ensure that the Simplify Geometry option in the *Sew Surface Definition* dialog is cleared.



Also note that the sew direction arrow should be pointing towards the surface of the body to be sewn.

10. Hide the fill feature and the sketch on the feature tree by right-clicking and selecting Show/Hide.

The base of the tank is now imprinted with the air injector profiles. Selecting the bottom face shows that the injectors are now separate faces that can be selected independently from each other.



The separate faces can now be defined as boundary regions.

Working With STAR-NX

This part of the documentation details how to create, set up and run a CFD simulation in STAR-NX. It assumes that you have previous experience in using NX, and only describes the additional features available as part of STAR-NX. The following sections are included:

- Introducing STAR-NX
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-NX Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-NX Settings
- Best Practices Guidelines

You use STAR-NX to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-NX

STAR-NX provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the NX environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-NX
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-NX enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-NX provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-NX interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

    - Ideal Gas
- Porous Media
    - Isotropic
    - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-NX provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the NX environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
    - Tetrahedral core mesh generation
    - Polyhedral core mesh generation
    - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-NX allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-NX supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-NX interface, and basic concepts. Topics include:

- Launching STAR-NX
- The STAR-NX Info App Window
- Modifying the config.xml File
- Understanding the STAR-NX Interface
- The STAR-NX Object Tree
- Editing Parts within NX Assemblies
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-NX Settings While the STAR-CCM+ Client is Running
- Switching Between NX Versions
- Disabling STAR-NX

Launching STAR-NX Manually

By default, STAR-NX should start automatically once you launch NX. The STAR-NX object tree icon appears on the left hand side of the screen when you open a part/assembly in NX.

If this icon does not appear, follow the steps outlined below to manually create it:

1. Start NX as usual.
2. Select Preferences > Palettes from the menu bar.
3. In the *Palettes* dialog, click Open Palette File.



4. Browse to and open the `starnx.pax` palette file, usually located in `C:\Program Files\CD-adapco\STAR-CAD8.0x.0xx\STAR-NX\win64\bin\star-nx7\startup`.

5. Click OK to load the palette.



The palette file will appear in the *Palettes* dialog.



6. Click Close.

The STAR-NX icon should now appear on the left of the screen, when you load or create a geometry file. Click on this icon to display the STAR-NX object tree.

Note: You may need to click on the icon twice for the tree to appear if you just loaded the palette.

The STAR-NX Info App Window

The *STAR-NX Info App* window automatically appears when you start NX. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-NX. More tabs appear as you create or load geometry files in STAR-NX, allowing you to perform more actions.



You cannot close the *STAR-NX Info App* window while NX or STAR-NX are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-NX license checked out successfully. The version of STAR-NX is also displayed in this tab. This tab appears by default when you launch NX without a geometry file.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-NX.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-NX searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-NX user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-NX is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-NX Info App. This can be set to one of the following:
  - `Normal` — This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose` — This will display a moderate amount of information in the Output panel.
  - `Log` — This will display all available information in the Output panel.

Understanding the STAR-NX Interface

The STAR-NX user interface is embedded within the NX environment, with the addition of the STAR-NX feature tree, menu, and toolbar.



The majority of the STAR-NX operations are executed within the feature tree and toolbar. Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-NX interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings
- Region settings
- Add new boundary
- Boundary conditions
- Mesh Pipeline settings
- Add new mesh controls

Toolbar

The operations provided in the toolbar include:

- Generate Surface Mesh
- Generate Volume Mesh
- Check Model For Errors
- Build and Run
- Run Solver
- Run For More Iterations
- Show/hide Mesh
- Transfer CAE Model
- Transfer Geometry

Menu

The operations provided in the menu include:

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Transfer CAE Model
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
- Solution
  - Build and Run
  - Run Solver
  - Run For More Iterations
  - Clear Solution
- Check Model For Errors
- Export
  - Surface Mesh (.dbs)

  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Transfer Geometry
- STAR-NX Help
- About STAR-NX

The STAR-NX Object Tree

STAR-NX uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-NX. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-NX applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-NX Interface section.

Clicking on the Region node within the object tree will highlight all the relevant faces on the geometry model for that particular region. If different parts are assigned to multiple regions, then clicking on a region will only

highlight the part and faces associated with that specific region. For example, clicking on the Fluid node below will highlight all the faces associated with the Fluid region.





Similarly, clicking on the Solid node highlights all the faces associated with the Solid region.





Clicking on a boundary node in the STAR-NX object tree highlights all the faces on the CAD geometry model that are associated with it.Note that clicking on the *Default* boundary node will not highlight any faces on the CAD geometry model. To stop highlighting the faces, you can either click on a blank space in the geometry window, or click on the root node in the STAR-NX object tree.

Editing Parts Within NX Assemblies

When you edit a part within an NX assembly (right-clicking and selecting Make Work Part), all STAR-NX functions are deactivated and appear greyed out. You can continue editing the part as you normally would.

To reactivate STAR-NX:

- Right-click the assembly root node and select Make Work Part.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select *STAR-NX* > *Preferences* from the menu.



   The *Preferences* dialog appears.

2. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

3. To run STAR-NX in parallel, select Parallel from the *Run Mode* box. The Compute Processes field appears.



4. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.
5. Click OK to close the dialog.

   If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-NX and STAR-CCM+ need to be closed for the changes to take effect.



6. Click OK to proceed.
7. Close the active STAR-NX and STAR-CCM+ sessions.
8. Restart STAR-NX.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-NX.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-NX.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-NX.

To change the units system:

1. Bring up the *Units System* dialog:
   - Right-click the STAR-NX node within the object tree and select Set Units System.



   - Alternatively, from the menu, select *STAR-NX* > *Settings* > *Units System*.



The *Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click OK to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers

- Angstroms
- Nanometers
- Mils
- Microinches
- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-NX environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the toolbar.

- Alternatively, from the menu, select *STAR-NX* > *Transfer CAE Model*.



The STAR-CCM+ version affiliated with STAR-NX launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Using NX Mouse Controls in STAR-CCM+

STAR-CCM+ offers the ability to set the mouse interactor style to match that of NX; this could make geometry manipulation in STAR-CCM+ more natural if you are used to working with NX.

To change the mouse interactor style:

1. Launch STAR-CCM+.
2. Click *Tools* and select Options from the menu.



The *Options* dialog appears.

3. Expand the STAR-CCM+ node and select Visualization.
4. In the *Properties* window, select Interactor Style and set this to pro-E / NX using the drop-down menu.



5. Click Close.

STAR-CCM+ will now use this mouse interactor style each time you launch it.

Editing STAR-NX Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-NX while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-NX, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Switching Between NX Versions

If you have multiple versions of NX installed on your machine, STAR-NX is installed to work with the latest version. Attempting to run a different version of NX will result in an error message with the ugraf.exe executable.



Note: The error message may differ depending on what version of NX you are trying to run.

This error indicates that STAR-NX is not set to work with the version of NX that you are trying to run. Clicking OK will dismiss the message and allow you to work in that version of NX without STAR-NX.

You can choose which version of NX STAR-NX works with from the `custom_dirs.dat` file. Typically, this file is located in:

```
C:\Program Files\CD-adapco\STAR-CAD8.0x.0xx\STAR-NX\
```

You can find the `custom_dirs.dat` file by checking the `UGII_CUSTOM_DIRECTORY_FILE` environment variable.

To specify which version of NX to use with STAR-NX:

1. Open `custom_dirs.dat` with a text editor of your choice.

   Example `custom_dirs.dat` content is shown below. Depending on what versions of NX are installed on your machine, the contents of this file may differ:

   ```
   // If uncommented, the following line supports STAR-NX with
   NX 8.

   C:\Program Files\CD-adapco\STAR-CAD8.04.004\STAR-NX\win64
   \bin\STAR-NX8

   // If uncommented, the following line supports STAR-NX with
   NX 7.

   // C:\Program Files\CD-adapco\STAR-CAD8.04.004\STAR-NX\wi
   n64\bin\STAR-NX7
   ```

2. Uncomment the line that corresponds to the path of the version of NX that you want to use. In this example, NX7.5.

3. Comment out the lines that corresponds to the path of any other versions of NX.

   The final `custom_dirs.dat` content appears as follows:

   ```
   // If uncommented, the following line supports STAR-NX with
   NX 8.

   // C:\Program Files\CD-adapco\STAR-CAD8.04.004\STAR-NX\wi
   n64\bin\STAR-NX8

   // If uncommented, the following line supports STAR-NX with
   NX 7.

   C:\Program Files\CD-adapco\STAR-CAD8.04.004\STAR-NX\win64
   \bin\STAR-NX7
   ```

4. Save the file and exit the text editor.

Disabling STAR-NX

If for any reason you would like to temporarily disable STAR-NX without having to uninstall it, do the following:

- Change the variable name of the `UGII_CUSTOM_DIRECTORY_FILE` environmental variable. It is recommended that the variable value is not changed.

  or

- Navigate to the STAR-NX installation folder, typically:

  `C:\CD-adapco\STAR-CAD8.0x.00x\STAR-NX`

  and rename the `custom_dirs.dat` file.

Where `x` is the relevant version number of your installation.

NX will now launch without STAR-NX. To enable STAR-NX again undo the changes listed above.

Common Functions

The following section details common operations used within the STAR-NX environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-NX
- Removing surfaces and parts from objects in STAR-NX

Renaming an Object

To rename an object:

1. Right-click the object node within the object tree and select Rename.



2. Enter the name.



3. Click on a blank space in the object tree.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, the Delete option is disabled in the pop-up menu.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select Edit...



The settings dialog for the object appears.

2. Make the changes that you want.
3. Click OK to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-NX

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-NX object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

- On the CAD geometry, click the surfaces that you want to assign to the object. You can add multiple surfaces by clicking more surfaces on the CAD geometry.

The selected surfaces are highlighted on the CAD geometry.



The selected surfaces are also displayed in the *Face Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to assign to the object. You can add multiple parts by clicking more parts on the CAD geometry.

The selected part is highlighted on the CAD geometry.

The selected parts are displayed in the *Part Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-NX

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-NX object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

1. In the *Face Selection* box, select the surface you want to delete.

The selected surface on the CAD geometry is highlighted in blue.



2. Click the Remove Selected Face button.



The surface is removed from the *Face Selection* box. You can select multiple surfaces for deletion by holding down the <`Ctrl`> key, and remove them in one operation.

Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

1. In the *Part Selection* box, select the part you want to remove.

The selected part on the CAD geometry is highlighted in dark blue.



2. Click the Remove Selected Part button.



The part is removed from the *Part Selection* box. You can select multiple parts for deletion by holding down the <`Ctrl`> key, and remove them in one operation.

Transferring Geometry to STAR-CCM+

With STAR-NX, you can transfer your CAD geometry directly to STAR-CCM+ from within the NX environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

To transfer the geometry to STAR-CCM+:

1. Select *STAR-CCM+* > *Transfer Geometry*.



STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-NX. This action ensures that the simulation file remains associated with the NX file. You can also safely rename and move the file using this method.
- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the NX file.

Resuming a Simulation

To resume a simulation after closing STAR-NX and STAR-CCM+:

1. Launch STAR-NX.
2. Open the CAD geometry in STAR-NX.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click STAR-NX > Parameters.



The *Expression Selection* dialog appears.



2. Select the design parameters that you want to transfer from the list.

Note: STAR-NX only recognizes native CAD features as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Select option.
4. Click OK.
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-NX, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-NX, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

   The geometry updates according to the design parameter changes.



6. Return to NX. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-NX and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

Updating the Geometry in STAR-CCM+ with NX Closed

You can close NX and continue working with the simulation in STAR-CCM+. Doing so, however, breaks the link between the two software packages. Any further updates that you make to the geometry in STAR-CCM+ are not transferred back to the geometry in NX. Instead, these changes are stored in memory, which allows you to include the design parameters in your optimization scripts.

This workflow also allows you to release the NX and STAR-NX licenses when performing tasks in STAR-CCM+. In this case, the NX and STAR-NX licenses are temporarily checked out only when the updates to the geometry are stored in memory.

Transferring the geometry from STAR-NX will overwrite the geometry in STAR-CCM+.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from NX to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-NX root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click OK.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.
Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.

Transfer Sheet Geometry	Specifies whether to transfer sheet bodies as part of the transfer process. This option is turned off by default.

Understanding the STAR-NX Process

In STAR-NX the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume
3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity
5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings
7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-NX and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-NX

STAR-NX only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-NX.

Suppressed or deactivated parts and features are not supported by STAR-NX and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-NX.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

  For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-NX includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-NX includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-NX is completely integrated within the NX environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-NX model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

  Including material type and properties.
- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.
- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.

- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-NX automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select Add Continuum.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



- Right-click on the Regions manager node in the object tree and select Add Region.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-NX). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-NX model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. Then right-click on the region node under the Region manager node in the object tree and select Edit... In the *Region Settings* dialog, select the alternative continuum from the *Continuum Selection* box and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-NX each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model

physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By default the material type is set to Air, and this can be changed to any one of a number of predefined solids, liquids or gases that are available in STAR-NX. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-NX object tree and select Edit.



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material* box, you can select a material type.
- In the *Material Properties* box, you can define each of the material properties.



Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-NX offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the drop-down menu. When you have selected a material,

the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-NX are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.
Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-NX, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)
- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select Edit...



The physics models available in the *Continuum Settings* dialog are determined by the material type.



Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.



Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.



The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-NX are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Locate the *Gravity* box and tick the *Enable* checkbox. This will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. To define reference properties for a continuum, right-click the Continuum node in the object tree and select Edit. In the *Continuum Settings* dialog, locate the *Reference Properties* box and expand , if necessary.



The reference properties are described below:

- Reference Pressure

  The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density , Reference Velocity , Reference Area .

These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-NX, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select Edit...



The physics models can then be enabled for the region by ticking the corresponding checkbox and entering values as necessary.



Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the *Porous* checkbox.

This activates the Porous Settings box. There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select the radio button labeled *Isotropic*.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, α and β respectively.

Uni-Directional

To specify uni-directional porous resistance, select the radio button labeled *Uni-Directional*.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, α and β respectively. The cross-flow resistance R is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, $\mathbf{i}$ $\mathbf{j}$ $\mathbf{k}$.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the *Moving Reference Frame* checkbox.



In the *Moving Reference Frame* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the ω box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, i j k. The units used are the same as the file defaults.

Alternatively, click on to bring up the *Vector* dialog, where you can set the Axis Origin and Vector of the axis of rotation using native NX functionality. The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.



This action activates the Energy Source Settings box. There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select the radio button labeled *Energy Input per Unit Volume*.



Enter the value for the energy input value per unit volume in the box.

Total Energy Input

To define the total energy input, select the radio button labeled *Total Energy Input*.



Enter the value for the total energy input in the box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.
2. Create a streamlines derived part.

3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-NX.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-NX will use default values to initialize the solution. In STAR-NX, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-NX object tree and select Edit.



This opens the *Continuum Settings* dialog, in which the *Initialization* box can be used to specify the following values:



- U V W Initial velocity
- Initial pressure
- Initial temperature
- Int len Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial velocity, open the *Continuum Settings* dialog, expand the *Initialization* box, and enter the values in the U, V, and W boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, locate the *Initialization* box, and enter a value in the box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the Turbulence Intensity Int and Turbulent Length Scale len boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-NX boundary is made up of one or more CAD surfaces within the simulation which, along with other boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select Add Boundary from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-NX). When you assign a surface to another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-NX, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Velocity Inlet from the drop-down menu.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity , with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors , select Cartesian Components from the drop-down menu.



Cylindrical vectors V_r V_θ V_z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i j k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Mass Flow Inlet from the drop-down menu.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.



The mass flow rate , must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors , tick the checkbox labeled Specify Direction.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Stagnation Inlet from the drop-down menu.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure , must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors , tick the checkbox labeled Specify Direction.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-NX, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Flow-Split Outlet from the drop-down menu.

The only parameter required is the flow-split ratio , expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.



Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Pressure Outlet from the drop-down menu.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.



The pressure , must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Wall from the drop-down menu.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U V W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity ω. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i j k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux, to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature , and heat transfer coefficient HTC.



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Symmetry Plane from the drop-down menu.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-NX to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-NX and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature*.



Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Turbulence*.

Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity Int
- Turbulence Length Scale len

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Passive Scalar*.

This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-NX provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-NX* menu), it is possible to execute only single stages from the STAR-NX > Mesh menu. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the NX parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from NX.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the NX *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-NX creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-NX appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-NX appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all NX parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose Edit.



The *Mesh Pipeline Settings* dialog appears. This dialog has several collapsible sections for the general settings.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click OK to save them and close the dialog.
- Click Cancel to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.



Selecting the trimmer option will activate the Trimmer tab..

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-NX Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.



The base size can be set as either a relative or absolute value. When it is set as a relative value (by selecting the Relative to Geometry option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* box of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of prism layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.

This panel allows you to change the following parameters:

- Relative Minimum Surface Size

- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-NX will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer panel.



The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the NX parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when they are created. A part can only belong to one mesh pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select Add Mesh Pipeline.



During the creation of a new mesh pipeline, it is necessary to assign to it the NX parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

  This mesh control is applied to one or more faces of the NX model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose Add Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select Add Volume Mesh Control from the pop-up menu.



The *Volumetric Control Settings* dialog appears.



The *Volumetric Control Settings* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape

- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the NX geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.



X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the NX window. To use an existing point, click the Set Point button and use the NX point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Set Point button and using the NX point selection procedure to choose an existing point in the NX window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete selected shape button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select Add Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-NX* toolbar.

- Alternatively, select Mesh > Generate Surface Mesh from the *STAR-NX* menu.



Output from the surface mesher will be displayed in the *STAR-NX Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-NX* toolbar.

- Alternatively, select Mesh > Generate Volume Mesh from the *STAR-NX* menu.



Information on the progress of the volume mesher will be displayed in the *STAR-NX Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Mesh > Clear Mesh from the *STAR-NX* menu.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-NX* toolbar.

Alternatively, select Mesh > Show/Hide Mesh from the *STAR-NX* menu.



Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-NX will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-NX Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Running in Batch Mode
- Clearing the Solution
- Visualizing the Solution

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the STAR-NX node within the object tree and select Set Solution Controls.



- Alternatively, from the menu, select *STAR-NX* > *Settings* > *Solution Controls*.



The *Solution Controls* dialog appears.



2. Here, you can:
   - Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.
   - Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.
3. Click OK to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:

1. In the *Solution Controls* dialog, expand the *Advanced* box.

   The advanced options appear.



1. In the Number of Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:
   - First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.
   - Second Order—This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:
   - Momentum
   - Pressure
   - Turbulence (with turbulence modelling defined in Continuum Settings)
   - Temperature (with temperature modelling defined in Continuum Settings)
     - Fluid Temperature Under-Relaxation Factor
     - Solid Temperature Under-Relaxation Factor
5. If you want to use additional STAR-CCM+ physics models with STAR-NX, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros
6. Click OK to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-NX by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-NX:

1. Open the *Solution Controls* dialog.
2. Expand the advanced solver options.
3. Activate the Use Java File option.

   A text entry field appears.



4. Click the Select Input File button. This action opens a browser dialog.
5. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
6. Click OK to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-NX.

Using STAR-CCM+ Continuum Physics with STAR-NX

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-NX and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.
2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.
5. Return to STAR-NX, open the *Solution Controls* dialog and activate the Use Java File checkbox.
6. Select the Java macro that you created earlier.
7. Run the case from STAR-NX.

Checking the Model for Errors

STAR-NX monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.
- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.



2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-NX object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-NX root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check Model for Errors in the *STAR-NX* menu.



- Alternatively click the Check Model for Errors button in the *STAR-NX* toolbar.

If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-NX Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-NX, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-NX* menu.



- Alternatively click the Build and Run button in the *STAR-NX* toolbar.

This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-NX Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Solution > Run Solver in the *STAR-NX* menu.



- Alternatively, click the Run button in the *STAR-NX* toolbar.

STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Solution > Run for More Iterations in the *STAR-NX* menu.



- Alternatively you can click the Run for More Iterations button in the *STAR-NX* toolbar.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Running in Batch Mode

To run STAR-NX in batch mode, you must first add the paths to the `starnx.dll` and `StarNXBatch7.exe` files to your Path user environment variable. These paths are added to the list of paths that may already exist on your machine. For more information on how to set the `Path` variable please see the Setting the Path User Environment Variable section.

A typical path is shown below.

```
C:\Program Files\CD-adapco\STAR-CAD<star-ccm+_version>\
STAR-NX\win64\bin\star-nx<nx_version>\startup
```

Each version of NX has a different path.

To run the solver in batch mode:

1. Choose or create a directory from which to launch STAR-NX in batch mode. This is referred to as the Batch Runtime Folder.
2. Set up the desired case in STAR-NX, assigning all relevant continuum, region, boundary, mesh pipeline, and solution control settings.
3. Place the geometry file in the Batch Runtime Folder.
4. Launch the NX Command Prompt.
5. In the Command Prompt, change to the Batch Runtime Folder that contains the CAD geometry with the saved STAR-NX settings.

   An example code is shown below:

   ```
   C:\>cd "C:\Batch Runtime Folder"_
   ```

6. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

   ```
   StarNXBatch<nx_version>.exe GeometryName.prt
   ```

   where:

   - `StarNXBatch<nx_version>.exe` is the batch executable.
   - `<nx_version>` is the version of NX that you want to run. Refer to the release notes to see which versions of NX are supported.
   - `GeometryName.prt` is the name of the input geometry file (including extension).

   An example code for running STAR-NX in batch mode using NX 7.5 is shown below.

   ```
   StarNXBatch7.exe exhaustManifold.prt
   ```

7. Press <`Enter`> to run the simulation in batch mode.

STAR-NX will work in the background, meshing the geometry and running the solvers. A simulation file will be saved in the Batch Runtime Folder once the solver has finished running. If there is an existing simulation file in the Batch Runtime Folder, a backup file is created with a `.sim~` extension. A log file (`starcad.log`) is also created and saved in the Batch Runtime Folder—this log file contains a record of all the messages that appeared in the command prompt during batch processing.

Setting the Path User Environment Variable

You need to set the `Path` user environment variable in order for STAR-NX to run in batch mode. This variable will have the full path to the folder containing the `StarNX<nx_vesion>.dll` and `StarNXBatch<nx_version>.exe` files. To create the `Path` user environment variable:

1. Open the Environment Variables dialog:

- For Windows 7:
  i. Click Start and type `env` into the search field near the bottom of the *Start* menu.
  ii. Select `Edit environment variables for your account` from the list. The *Environment Variables* dialog appears.
  iii. If the Path user environment variable already exists in the *User variables* box click on Edit..., otherwise click on New... to create it. The *User Variable* dialog appears.

2. Set the Variable name to `Path`.
3. Set the Variable value to:

   ```
   C:\Program Files\CD-adapco\STAR-CAD<star-ccm+_version>\
   STAR-NX\win64\bin\star-nx<nx_version>\startup
   ```

   Where `<nx_version>` corresponds to the version of NX that you want to use, for example `7.5`.
4. Click OK to close the *User Variable* dialog.
5. Click OK to close the *Environment Variables* dialog.
6. Click OK to close the *System Properties* dialog.

Launching the NX Command Prompt

The NX command prompt is located in the `NX Tools` folder in the *Start* menu:

- The default location for NX 7.5 is:

  ```
  Start > All Programs > UGS NX 7.5 > NX Tools >
  Command Prompt
  ```

- The default location for NX 8.0 is:

  ```
  Start > All Programs > Siemens NX 8.0 > NX Tools >
  Command Prompt
  ```

- The default location for NX 9.0 is:

  ```
  Start > All Programs > Siemens NX 9.0 > NX Tools >
  Command Prompt
  ```

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-NX case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the menu, select *STAR-NX* > *Solution* > *Clear Solution*.



The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-NX.

- Exporting a Surface Mesh

- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-NX as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-NX menu, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-NX as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-NX menu, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-NX menu, select Export > Engineering Report (.html, .csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)

    - Pressure Coefficient
    - Total Pressure Coefficient
    - Turbulent Kinetic Energy (J/kg)
    - Turbulent Dissipation Rate (m^2/s^3)
    - Turbulent Viscosity (Pa-s)
    - Density (kg/m^3)
- Force data, at each boundary:
    - Net Force [Fx, Fy, Fz]
    - Pressure Force [Fx, Fy, Fz]
    - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
    - Coefficient of Net Force [Cfx, Cfy, Cfz]
    - Coefficient of Pressure Force [Cfx, Cfy, Cfz]
    - Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
    - Physics models applied to the continuum
    - Materials properties
    - Turbulence properties
    - Segregated flow properties
    - Reference values
    - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
    - Polyhedral mesher properties
    - Prism layer mesher properties (if activated)
    - Surface remesher properties
    - Reference values
    - Volumetric controls

- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values
- Solver parameters
- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-NX to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-NX, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-NX menu, select Export > Simulation (.sim).



Resetting STAR-NX Settings

STAR-NX saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-NX settings:

- From the menu, select Settings > Clear All.



This action returns the geometry to its original state.

Best Practice Guidelines

This section aims to provide guidance and assistance on making a geometry part suitable for CFD analysis using STAR-NX. The following items are covered:

- Creating the Fluid Domain:
  - Internal Flow
  - External Flow
  - Multi-Region Flow
- Imprinting:
  - Imprinting Parts in an Assembly
  - Imprinting Surfaces on a Part for Boundary Region Applications

Prior to using these methods, ensure that the components in assemblies are all solid bodies and are in contact with each other. Any gaps between components will need to be closed.

A basic knowledge and understanding of both CAD and CFD are assumed, although most basic functions are covered.

Creating the Fluid Domain

In the majority of CAD models, the fluid domain for analysis must be extracted from the original part or assembly. The following sections demonstrate how to prepare the fluid domain for:

- an internal flow
- an external flow
- a multi-region flow

Internal Flow

In general, the steps to follow in creating the fluid domain for an internal flow are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an additional part that will close off any openings in the original geometry, these are called “caps”.
3. Unite all parts with the caps to obtain one solid part.
4. Create an overall block part that encloses all the other parts, this is the “enclosure”.
5. Subtract all parts from the enclosure and delete the outer solid body--the remaining geometry is the fluid domain.

An example case is shown below. In this case the assembly contains two parts, the valve (shown in red) and the valve housing (shown in blue).



1. Load the assembly into NX. If the geometry is a part, insert it into an assembly.

We will now insert a new component into the assembly. This part will close any openings in the assembly using caps.

2. Insert a new component into the assembly.



Give this component an appropriate name and then accept the default placement settings. In this case the new component was given the name `cap`.

3. Make the new component the work part.



4. Select the Extrusion tool and select the outer edge of one of the openings on the geometry, as shown below.



Ensure that the Type Filter is set to No Selection Filter and that the Selection Scope is set to Entire Assembly, as shown below, otherwise you will only be able to select features from the active part.

5. Perform the same actions for all other openings.



6. Use the Unite operation to join all the parts and form one solid.



In the *Unite* dialog, select one of the caps you created in the previous step as the Target. Select the other parts (the other cap, valve and valve housing) as the Tool.

7. Hide the original valve and valve housing parts by going to the *Assembly Navigator* and clicking on the red tick besides each part. These should become greyed out as shown below.



The part will appear as shown below.



8. Return the to the assembly and make the assembly the work part.

A new component will be inserted into the assembly. A sketch that encloses the entire assembly will then be created and extruded. The valve, valve housing, and caps will then be subtracted from this, leaving two bodies. The outer body will be deleted, leaving only the internal fluid volume.

9. Insert another component into the assembly and make it the work part.
10. Create a sketch on any plane and sketch a bounding box that will, when extruded, enclose the entire assembly, as shown below.



11. Subtract the cap part from the internal fluid part.



Select the fluid box as the Target and the caps as the Tool. To be able to select parts through solids, click on the Highlight Hidden Edges button. Moving the mouse over the geometry will then highlight internal bodies in brown (the default color), as shown below.



12. Use the Extract Bodies tool to extract the internal fluid region.



With the *Extract Body* dialog active, click on any face on the outer body to select it. Note that there is another free-floating solid body enclosed within this body, which will not get selected -- this is the internal fluid volume.

13. Hide the caps geometry.

14. Go to the Part Navigator and suppress the Extracted Body feature. This is done by clicking on the tick next to the feature--when suppressed the tick besides the node will disappear.



15. Suppress the components that you do not want to take part in the analysis, such as the valve, valve housing and caps, by right-clicking on these and selecting Suppression...

16. In the *Suppression* dialog that appears set the component to Always Suppressed and click on OK.



The extracted internal fluid volume is shown below.



The geometry is now ready for the CAD client analysis. Suppressed parts in an assembly will not be included in the analysis. In the above example, it is recommended that you extend the inlet and outlet regions to appropriate lengths so that a fully developed flow can be achieved before reaching the valve and to help reduce recirculation at the outlet.

External Flow

The generation of a domain for external aerodynamics is generally easier than for internal flow cases but the sizing of the domain is more critical. As a guideline (assuming a specific situation is not being modelled i.e. flow in a wind tunnel) the domain should be sized as follows:

Inlet: 2 x part length from the front of the part

Outlet: 4 x part length from the rear of the part

Sides: 3 x part width from the side of the part

Ceiling: 4 x part height from the top of the part

In general, the steps to follow in creating the fluid domain for an external flow are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an overall block part that encloses all the other parts. The shape and size should correspond to the desired fluid domain.
3. Subtract all parts from the overall block part - the remaining geometry is the fluid domain.

An example case is shown below. In this example a geometry of a car will be used.



1. Load the assembly in NX. If the geometry is a part file, insert this into an assembly.

2. If the geometry is symmetrical, cut this in half to save on computation costs. The easiest way of doing this is to perform an extrusion to remove the unwanted half.



3. Insert a new component into the assembly.



Give this component an appropriate name and then accept the default placement settings. In this example, the component was called `domain`.

4. Make the new component the active part, right click on it in the feature tree and select Make Work Part.



5. Create a sketch on the symmetry plane of the geometry. This sketch will represent the fluid domain profile.



6. Extrude the sketch to an appropriate length to complete the domain creation.



7. Subtract the car geometry from the domain geometry. This is done using the Subtract feature.



In the *Subtract* dialog, select the domain as the Target and the car as the Tool.

Ensure that the Type Filter is set to No Selection Filter and that the Selection Scope is set to Entire Assembly, as shown below, otherwise you will only be able to select the active part.

8. Hide the car part, by clicking on the tick next to the part. Once done, the tick will be greyed out.

The final result is shown below.



The geometry is now ready for the cad client analysis. Suppressed parts are not included in the analysis.

Multi-Region Flow

This section shows how to create the fluid domain for multiple regions such as rotating machinery and porous media flow.

For the numerical assumptions made using the Moving Reference Frame (MRF) technique to hold true, a minimum of two cells need to be present between the rotating-static interface and any wall boundaries (for example the fan blades), and the rotating region must be circular in section. Where there is a large clearance between the rotating object and any static objects, it is advisable to avoid placing the interface too close to either hence facilitating a small cell size on the interface. Where there is a small clearance between the rotating and static parts, there is little choice but to place the interface approximately midway between the two.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work in an assembly. If the original geometry contains only a single part, create a new assembly that contains the part.
2. Create an additional part that will enclose the rotating region (for example a fan) and subtract it from the part. In the case of a porous region, a separate part will be required for this as well.

3. Create an additional part for the main static domain (the inlet and outlet regions) and subtract any obstructions/features that may be present there (for example the fan motor support arms on a fan casing).



1. Load the assembly into NX.

2. Insert a new empty component into the assembly.



Give this component an appropriate name and then accept the default placement settings. In this example the new component was given the name `Rotating Region`.

3. Make the new component the work part by right-clicking on its node in the assembly navigator and selecting Make Work Part.



4. Hide the Fan part by going to the assembly navigator and clicking on the red tick besides the part. The tick will turn grey and the part will not be visible in the geometry window.

5. Create a sketch on the contact face of the fan casing (where the fan and the fan casing are coincident).



6. If the continuous auto dimensioning feature of NX is turned on, disable it by clicking on the Continuous Auto Dimensioning button in the toolbar.



7. Draw a circle that shares the same center point as the sketch plane face. The size of the circle is not important at this stage as it will be made coincident with the innermost edge of the fan casing in a later step.

8. Ensure that the Type Filter is set to No Selection Filter and that the Selection Scope is set to Entire Assembly, as shown below, otherwise you will not be able to select edges from other parts.

9. Click on the Project Curve button in the toolbar



If this button is missing from the toolbar, you may need to click on the (Open in Sketch Task Environment) button to open the active sketch in the sketch task environment.

The *Project Curve* dialog appears.

10. Select the innermost edge of the fan casing.



and exit the dialog.

A message will appear warning you that the projection will not be associative. Click on Yes to accept and close the dialog.



11. Right-click on the projected edge and select Convert to Reference.
12. Click on the Constraints button in the toolbar.



The *Constraints* dialog will appear.

13. Select the circle and the innermost edge of the fan casing, as shown below. This will add new options in the *Constraints* dialog.
14. Click on the Equal Radius button to make the circle and the projected edge have the same radius.



The circle will snap to the projected edge, as shown below.



15. Hide the fan casing and show the fan.

16. Extrude the sketch so that it fully encloses the fan.



Ensure that no boolean operations are set to take place in the *Extrude* dialog - these might result in an error during the extrusion.

17. Use the boolean subtract operation to remove the fan from the Rotating Region.



You can use the See-Thru view with the Highlight Hidden Edges button to see and select the fan inside the Rotating Region.

Hiding the fan and using the See-Thru view will show the result of this operation.



18. Return to the assembly and make the assembly the work part.
19. Hide the fan part; this will no longer be required.
20. Insert another component into the assembly as previously described and make that the work part.
21. Use the Extrude tool and select the existing outer edge on the front face of the Rotating Region. To make it easier to see which is the front face, use the see-thru tool.

The preview of the extrude is shown below.



The extrusion depth should be suitable for the case you are investigating.

22. Return to the assembly and make this the work part.

The outlet region will be created using the same technique used for the inlet region.

23. Insert another part into the assembly. Make this the working part.

24. Create an extrusion using the outer edge of the back face of the Rotating Region.



25. Hide the Inlet Region and the Rotating Region, leaving only the Outlet Region visible.
26. Show the fan casing part; this will be subtracted from the Outlet Region.

27. Subtract the fan casing from the outlet region. In the *Subtract* dialog select the outlet region as the Target and the fan casing as the Tool. The result of this operation is shown below (using the See-Thru view).



28. Show the Inlet Region and the Outlet Region.
29. Suppress the components that you do not want to take part in the analysis, such as the fan and casing, by right-clicking on these and selecting Suppression...

30. In the *Suppression* dialog that appears set the component to Always Suppressed and click on OK.



The final fluid domain is shown below.



The geometry is now ready for cad client analysis. Suppressed parts will not be included in the analysis. The three regions that form the fluid domain all have a conformal match, so it will be possible to create the interfaces and obtain a topologically valid mesh in STAR-CCM+ without any further modifications to the geometry.

Imprinting Parts

Most CAD model assemblies are not suitable for meshing as the parts in the assembly are disjointed, where the coincident faces are overlapping. Attempting to generate a mesh with such an assembly will result in a meshing error. To rectify this problem, the parts in an assembly must be imprinted together.

Imprinting can also be used to split surfaces for boundary region applications. An example of such an application would be to model the area that a fan would occupy on the inside of computer case, such as in the third STAR-NX tutorial, where a smaller section of the surface is needed to represent the fan, rather than the entire inner surface of the computer case.

Imprinting Parts in an Assembly

In cases containing multiple adjacent parts, such as in Conjugate Heat Transfer cases, it is essential to imprint the profiles of the coincident faces of the smaller parts onto the face of the larger part. This will ensure that the volume mesh created by STAR-CCM+ contains interfaces between the adjacent parts.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work within the assembly, making the part that requires imprinting the work part.
2. Create a sketch on the face you want to imprint, and project the contact edges of the other components in the assembly onto this sketch.
3. Split the face you want to imprint using the sketch created in the previous step.

The simple circuit board shown below requires imprinting before the mesh can be generated.



In this example, the capacitors and heat sinks are not imprinted to the circuit board. This would result in a failure when attempting to generate a mesh.

To imprint the capacitors and heat sinks onto the circuit board follow the steps outlined below:

1. Load the assembly file in NX.

2. Make the part that will be imprinted the Work Part, right click on the part's node in the feature tree. In this example this is the base board.



3. Use the See-Thru button from the toolbar to make the other parts translucent.

4. Create a new sketch on the surface that requires imprinting.
5. Click on the Project Curve button on the toolbar



If this button is missing from the toolbar, you may need to click on the (Open in Sketch Task Environment) button to open the active sketch in the sketch task environment.

The *Project Curve* dialog will appear.

6. Click on the contact face of the other parts to select the contact edges, as shown below.



7. Clicking OK in the *Project Curve* dialog will bring up message informing you that the projected edges will not be associative. Click on Yes to accept and continue.



8. Click on the Divide Face button in the toolbar.



The *Divide Face* dialog will appear.

9. In the *Faces to Divide* box, select the face you want to imprint.
10. In the *Dividing Objects* box, select the sketch that contains the contact edges of the other parts.
11. Ensure that the projection direction (the arrow in the geometry window) is pointing towards the face to be imprinted. This can be reversed by clicking on the Reverse Direction button in the *Projection Direction* box.
12. Return to the assembly and make the assembly the work part.

The contact edges of the capacitors and heat sinks have now been imprinted onto the base board. This can be verified by the parts in the assembly.



A high quality mesh with fully defined interfaces can now be successfully generated for this assembly. Suppressed parts are not included in the STAR-CCM+ analysis, so for CHT cases, such as the above example, ensure that all relevant parts are active.

Imprinting Surfaces on Parts for Boundary Region Applications

Faces in a part body can be split into smaller faces for boundary application.

The general procedure to follow is:

1. Create a sketch on the surface you want to split. The shape of the sketch should correspond to the shape of the new face after it is split.
2. Split the face that requires imprinting with the sketch containing the profiles.

A simple example of an aeration tank is shown below. The image shows the tank, the two air injectors need to be imprinted onto the base of the tank.



1. Load the part into NX.
2. Create a sketch on the face that requires imprinting. In this case the bottom face of the tank needs imprinting.

3. Draw the profile of the boundary you want imprinted. In this case, the air injectors are circular in profile.



You can use `<Shift>` + `<F8>` to change the view so that it is normal to the sketch plane.

4. Use the Divide Face function to split the face.



Select the face you want to divide for the Faces to Divide field, and the sketch profile(s) you created in the previous step as the Dividing Objects.

The face is now imprinted; each face can be selected independently of each other, and can be used when defining boundaries.



The separate faces can now be defined as boundary regions.

Working With STAR-Pro/E

This part of the documentation details how to create, set up and run a CFD simulation in STAR-Pro/E. It assumes that you have previous experience in using Pro/E, and only describes the additional features available as part of STAR-Pro/E. The following sections are included:

- Introducing STAR-Pro/E
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-Pro/E Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-Pro/E Settings
- Best Practices Guidelines

You use STAR-Pro/E to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-Pro/E

STAR-Pro/E provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the Pro/E environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-Pro/E
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-Pro/E enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-Pro/E provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-Pro/E interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

  - Ideal Gas
- Porous Media
  - Isotropic
  - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-Pro/E provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the Pro/E environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
  - Tetrahedral core mesh generation
  - Polyhedral core mesh generation
  - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-Pro/E allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-Pro/E supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-Pro/E interface, and basic concepts. Topics include:

- Launching STAR-Pro/E
- The STAR-Pro/E Info App Window
- Modifying the config.xml File
- Understanding the STAR-Pro/E Interface
- The STAR-Pro/E Object Tree
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-Pro/E Settings While the STAR-CCM+ Client is Running
- Disabling STAR-Pro/E
- Manually Creating the STAR-Pro/E Toolbar

Launching STAR-Pro/E Manually

By default STAR-Pro/E automatically starts when you launch Pro/E. The STAR-Pro/E Info App window appears, along with the STAR-Pro/E object tree and menu.

If these items are missing, you can launch STAR-Pro/E manually by following the steps outlined below:

1. Start Pro/E as usual.
2. Select Tools > Auxiliary Applications.
3. In the *Auxiliary Applications* dialog, click Register...

4. Browse to the `protk.dat` file in your STAR-CAD series installation directory, typically:

   `C:\Program Files\CD-adapco\STARCAD7.0x.0xx\STAR-ProE\protk.dat`

5. Select this file and click Open. An entry stating `star-proe.dll Not Running` will appears in the *Auxiliary Applications* dialog.
6. Select this entry and click Start.



The *STAR-Pro/E Info App* window will appear, and the checks for licensing will begin. If everything is in order, the `Licensing successful` message appears in the Status page of the *STAR-Pro/E Info App* window, and the `star-proe.dllNot Running` message will vanish from the *Auxiliary Applications* dialog.

7. Click Close in the *Auxiliary Applications* dialog.

You can now access the STAR-Pro/E features from the *STAR-Pro/E* object tree and menu.

The STAR-Pro/E Info App Window

The *STAR-Pro/E Info App* window automatically appears when you start Pro/E. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-Pro/E. More tabs appear as you create or load geometry files in STAR-Pro/E, allowing you to perform more actions.



You cannot close the *STAR-Pro/E Info App* window while Pro/E or STAR-Pro/E are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-Pro/E license checked out successfully. The version of STAR-Pro/E is also displayed in this tab. This tab appears by default when you launch Pro/E without a geometry file.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-Pro/E.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-Pro/E searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-Pro/E user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-Pro/E is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-Pro/E Info App. This can be set to one of the following:
  - `Normal` — This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose` — This will display a moderate amount of information in the Output panel.
  - `Log` — This will display all available information in the Output panel.

Understanding the STAR-Pro/E Interface

The STAR-Pro/E user interface is embedded within the Pro/E environment, with the addition of the STAR-Pro/E feature tree, menu, and toolbar.



The majority of the STAR-Pro/E operations are executed within the feature tree and toolbar. The toolbar has to be manually created. Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-Pro/E interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings
- Region settings

- Add new boundary
- Boundary conditions
- Mesh Pipeline settings
- Add new mesh controls

Toolbar

Please note that you need to manually create the toolbar in Pro/E. The operations provided in the toolbar include:

- Generate Surface Mesh
- Generate Volume Mesh
- Check Model For Errors
- Build and Run
- Run Solver
- Run For More Iterations
- Show/hide Mesh
- Transfer CAE Model
- Transfer Geometry

Menu

The operations provided in the menu include:

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Transfer CAE Model
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
- Solution
  - Build and Run

  - Run Solver
  - Run For More Iterations
  - Clear Solution
- Check Model For Errors
- Export
  - Surface Mesh (.dbs)
  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Transfer Geometry
- STAR-Pro/E Help
- About STAR-Pro/E

The STAR-Pro/E Object Tree

STAR-Pro/E uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-Pro/E. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-Pro/E applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-Pro/E Interface section.

Clicking on the Region node within the object tree will highlight all the relevant faces on the geometry model for that particular region. If different parts are assigned to multiple regions, then clicking on a region will only highlight the part and faces associated with that specific region. For example, clicking on the Fluid node below will highlight all the faces associated with the Fluid region.





Similarly, clicking on the Solid node highlights all the faces associated with the Solid region.





Clicking on a boundary node in the STAR-Pro/E object tree highlights all the faces on the CAD geometry model that are associated with it.Note that clicking on the *Default* boundary node will not highlight any faces on the CAD geometry model. To stop highlighting the faces, you can either click on a blank space in the geometry window, or click on the root node in the STAR-Pro/E object tree.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select *STAR-Pro/E* > *Preferences* from the menu.



The *Preferences* dialog appears.

2. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

3. To run STAR-Pro/E in parallel, select Parallel from the *Run Mode* property. The Compute Processes field appears.



4. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.

5. Click OK to close the dialog.

   If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-Pro/E and STAR-CCM+ need to be closed for the changes to take effect.



6. Click OK to proceed.
7. Close the active STAR-Pro/E and STAR-CCM+ sessions.
8. Restart STAR-Pro/E.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-Pro/E.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-Pro/E.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-Pro/E.

To change the units system:

1. Bring up the *Units System* dialog:
   - Right-click the STAR-Pro/E node within the object tree and select Set Units System.



   - Alternatively, from the menu, select *STAR-Pro/E* > *Settings* > *Units System*.



   The *Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click OK to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers
- Angstroms
- Nanometers
- Mils
- Microinches
- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-Pro/E environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the toolbar.

- Alternatively, from the menu, select *STAR-Pro/E* > *Transfer CAE Model*.



The STAR-CCM+ version affiliated with STAR-Pro/E launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Using Pro/E Mouse Controls in STAR-CCM+

STAR-CCM+ offers the ability to set the mouse interactor style to match that of Pro/E; this could make geometry manipulation in STAR-CCM+ more natural if you are used to working with Pro/E.

To change the mouse interactor style:

1. Launch STAR-CCM+.
2. Click *Tools* and select Options from the menu.



The *Options* dialog appears.

3. Expand the STAR-CCM+ node and select Visualization.
4. In the *Properties* window, select Interactor Style and set this to pro-E / NX using the drop-down menu.



5. Click Close.

STAR-CCM+ will now use this mouse interactor style each time you launch it.

Editing STAR-Pro/E Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-Pro/E while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-Pro/E, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Disabling STAR-Pro/E

If for any reason you would like to temporarily disable STAR-Pro/E without having to uninstall it, do the following:

1. Go to the text folder in the Pro/E installation directory, typically:

   `C:\Program Files\proeWildfire 5.0\text`

2. Locate and edit the `config.pro` file using a text editor.
3. Comment out the following line with an "!":

   ```
   protkdat C:\Program Files\CD-adapco\STAR-CAD8.0x.0xx\ST
   AR-ProE\protk.dat
   ```

   so that it looks like this:

   ```
   !protkdat C:\Program Files\CD-adapco\STAR-CAD8.0x.0xx\S
   TAR-ProE\protk.dat
   ```

Where `x` is the relevant version number of your installation.

Pro/E will now launch without STAR-Pro/E. To enable STAR-Pro/E again undo the changes listed above.

Manually Creating the STAR-Pro/E Toolbar

If the STAR-Pro/E toolbar does not appear in Pro/E, you can manually create it by following the steps outlined below:

1. Launch STAR-Pro/E.
2. Select *Tools* > *Customize Screen...* from the menu bar.



The *Customize* dialog appears.

3. Click the Toolbars tab.

4. Scroll down to the bottom of the list and activate any blank toolbar by activating the corresponding checkbox.



A blank toolbar is added to the top of the Pro/E workspace.

5. Click the Commands tab.

6. In the *Categories* pane, scroll down to the bottom of the list and select *Foreign applications*. The STAR-Pro/E icons are located in the *Commands* pane.



7. Drag the first icon into the blank toolbar.



8. Drag the remaining STAR-Pro/E icons into the toolbar.

9. Click OK to close the *Customize* dialog.

Common Functions

The following section details common operations used within the STAR-Pro/E environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-Pro/E
- Removing surfaces and parts from objects in STAR-Pro/E

Renaming an Object

To rename an object:

1. Right-click the object node within the object tree and select Rename.



2. Enter the name.



3. Click on a blank space in the object tree.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, the Delete option is disabled in the pop-up menu.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select Edit...



The settings dialog for the object appears.

2. Make the changes that you want.
3. Click OK to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-Pro/E

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-Pro/E object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

1. Click the Select Face button .

   This button activates the Pro/E surface selection mechanism and allows you to select the faces on the CAD geometry.

2. On the CAD geometry, click the surfaces that you want to assign to the object. You can select multiple faces by holding down the `<Ctrl>` key whilst clicking the surfaces.

The selected surfaces are highlighted on the CAD geometry.



The selected surfaces are added to the *Included Surfaces* box in the *Surface Sets* dialog.



3. Click OK to close the *Surface Sets* dialog.

The surfaces are then added to the *Face Selection* box of the active dialog.

Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

1. Click the Select part button .

This button activates the Pro/E part selection mechanism and allows you to select parts on the CAD geometry.

2. On the CAD geometry, click the part that you want to assign to the object. You can select multiple parts by holding down the `<Ctrl>` key whilst clicking the parts.

The selected part is highlighted on the CAD geometry.



The selected parts are added to memory in the Pro/E part selection mechanism.

3. Click OK to close the *Select* dialog.

The parts are added to the *Part Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-Pro/E

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-Pro/E object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

1. In the *Face Selection* box, select the surface you want to delete.

The selected surface on the CAD geometry is cleared.



2. Click the Remove a face button.



The surface is removed from the *Face Selection* box. You can select multiple surfaces for deletion by holding down the `<Ctrl>` key, and remove them in one operation.

Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

1. In the *Part Selection* box, select the part you want to remove.

The selected part on the CAD geometry is highlighted in white.



2. Click the Remove a part button.



The part is removed from the *Part Selection* box. You can select multiple parts for deletion by holding down the <`Ctrl`> key, and remove them in one operation.

Transferring Geometry to STAR-CCM+

With STAR-Pro/E, you can transfer your CAD geometry directly to STAR-CCM+ from within the Pro/E environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

To transfer the geometry to STAR-CCM+:

1. Select *STAR-CCM+* > *Transfer Geometry*.



STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-Pro/E. This action ensures that the simulation file remains associated with the Pro/E file. You can also safely rename and move the file using this method.
- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the Pro/E file.

Resuming a Simulation

To resume a simulation after closing STAR-Pro/E and STAR-CCM+:

1. Launch STAR-Pro/E.
2. Open the CAD geometry in STAR-Pro/E.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click STAR-Pro/E > Select Parameters.



The *List of Expressions* dialog appears.



2. Select the design parameters that you want to transfer from the list.

Note: STAR-Pro/E only recognizes native CAD features as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Select/Deselect option.
4. Click OK.
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-Pro/E, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-Pro/E, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

   The geometry updates according to the design parameter changes.



6. Return to Pro/E. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-Pro/E and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

Updating the Geometry in STAR-CCM+ with Pro/E Closed

You can close Pro/E and continue working with the simulation in STAR-CCM+. Doing so, however, breaks the link between the two software packages. Any further updates that you make to the geometry in STAR-CCM+ are not transferred back to the geometry in Pro/E. Instead, these changes are stored in memory, which allows you to include the design parameters in your optimization scripts.

This workflow also allows you to release the Pro/E and STAR-Pro/E licenses when performing tasks in STAR-CCM+. In this case, the Pro/E and STAR-Pro/E licenses are temporarily checked out only when the updates to the geometry are stored in memory.

Transferring the geometry from STAR-Pro/E will overwrite the geometry in STAR-CCM+.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from Pro/E to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-Pro/E root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click OK.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.
Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.
Transfer Sheet Geometry	Specifies whether to transfer sheet bodies as part of the transfer process. This option is turned off by default.

Understanding the STAR-Pro/E Process

In STAR-Pro/E the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume

3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity
5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings
7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-Pro/E and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-Pro/E

STAR-Pro/E only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-Pro/E.

Suppressed or deactivated parts and features are not supported by STAR-Pro/E and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-Pro/E.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

  For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-Pro/E includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-Pro/E includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-Pro/E is completely integrated within the Pro/E environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-Pro/E model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

  Including material type and properties.
- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.
- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.

- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-Pro/E automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select Add Continuum.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



To create an additional region:

- Right-click on the Regions manager node in the object tree and select Add Region.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-Pro/E). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-Pro/E model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. Then right-click on the region node under the Region manager node in the object tree and select Edit... In the *Region Settings* dialog select the alternative continuum using the *Continuum* drop-down menu from the *Continuum Selection* box and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-Pro/E each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By default the material type is set to Air, and this can be changed to any one of a number of predefined solids, liquids or gases that are available in STAR-Pro/E. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-Pro/E object tree and select Edit.



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material* box, you can select a material type.
- In the *Material Properties* box, you can define each of the material properties.



Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-Pro/E offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the drop-down menu. When you have selected a material,

the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-Pro/E are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.

Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-Pro/E, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)
- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select Edit...



The physics models available in the *Continuum Settings* dialog are determined by the material type.



Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.



Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.



The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-Pro/E are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Locate the *Gravity* box and tick the *Enable* checkbox. This will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. To define reference properties for a continuum, right-click the Continuum node in the object tree and select Edit. In the *Continuum Settings* dialog, locate the *Reference Properties* box and expand , if necessary.



The reference properties are described below:

- Reference Pressure

  The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density , Reference Velocity , Reference Area .

  These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-Pro/E, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select Edit...



The physics models can then be enabled for the region by ticking the corresponding checkbox and entering values as necessary.



Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the *Porous* checkbox.

There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select the radio button labeled *Isotropic*.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, α and β respectively.

Uni-Directional

To specify uni-directional porous resistance, select the radio button labeled *Uni-Directional*.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, α and β respectively. The cross-flow resistance $R_{\perp}$ is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, $\mathbf{i}$ $\mathbf{j}$ $\mathbf{k}$.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the *Moving Reference Frame* checkbox.

In the *Moving Reference Frame* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the ω box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, i j k. The units used are the same as the file defaults.

Alternatively, click on and select two vertices on the geometry to define the axis of rotation. Holding down `<Ctrl>`, select two vertices on the geometry; the first selection will act as the origin of the axis, and the second selection will define the direction vector of the axis. The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.

There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select the radio button labeled *Energy Input per Unit Volume*.



Enter the value for the energy input value per unit volume in the box.

Total Energy Input

To define the total energy input, select the radio button labeled *Total Energy Input*.



Enter the value for the total energy input in the box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.
2. Create a streamlines derived part.
3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-Pro/E.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-Pro/E will use default values to initialize the solution. In STAR-Pro/E, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-Pro/E object tree and select Edit.



This opens the *Continuum Settings* dialog, in which the *Initialization* box can be used to specify the following values:



- U V W Initial velocity
- Initial pressure
- Initial temperature
- Int len Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial velocity, open the *Continuum Settings* dialog, expand the *Initialization* box, and enter the values in the U, V, and W boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, locate the *Initialization* box, and enter a value in the box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the Turbulence Intensity Int and Turbulent Length Scale len boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-Pro/E boundary is made up of one or more CAD surfaces within the simulation which, along with other boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select Add Boundary from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-Pro/E). When you assign a surface to another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-Pro/E, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Velocity Inlet from the drop-down menu.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity, with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors U V W, select Cartesian Components from the drop-down menu.



Cylindrical vectors V_r V_θ V_z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i j k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Mass Flow Inlet from the drop-down menu.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.



The mass flow rate , must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors , tick the checkbox labeled Specify Direction.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Stagnation Inlet from the drop-down menu.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure , must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors , tick the checkbox labeled Specify Direction.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-Pro/E, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Flow-Split Outlet from the drop-down menu.

The only parameter required is the flow-split ratio, expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.



Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Pressure Outlet from the drop-down menu.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.



The pressure , must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Wall from the drop-down menu.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U V W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity ω. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i j k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux, to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature , and heat transfer coefficient HTC.



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Symmetry Plane from the drop-down menu.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-Pro/E to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-Pro/E and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature, Turbulence*.



Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Temperature, Turbulence*.



Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity Int
- Turbulence Length Scale len

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Temperature, Turbulence, Passive Scalar*.



This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-Pro/E provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-Pro/E* menu), it is possible to execute only single stages from the STAR-Pro/E > Mesh menu. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the Pro/E parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from Pro/E.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the Pro/E *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-Pro/E creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-Pro/E appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-Pro/E appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all Pro/E parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose Edit.



The *Mesh Pipeline Settings* dialog appears. This dialog has several collapsible sections for the general settings.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click OK to save them and close the dialog.
- Click Cancel to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.



Selecting the trimmer option will activate the Trimmer tab..

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-Pro/E Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.



The base size can be set as either a relative or absolute value. When it is set as a relative value (by selecting the Relative to Geometry option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* box of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of prism layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.



This panel allows you to change the following parameters:

- Relative Minimum Surface Size
- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-Pro/E will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer panel.



The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the Pro/E parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when

they are created. A part can only belong to one mesh pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select Add Mesh Pipeline.



During the creation of a new mesh pipeline, it is necessary to assign to it the Pro/E parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

  This mesh control is applied to one or more faces of the Pro/E model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose Add Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select Add Volume Mesh Control from the pop-up menu.



The *Volumetric Control Settings* dialog appears.



The *Volumetric Control Settings* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape
- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the Pro/E geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.

X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the Pro/E window. To use an existing point, click the Set Point button and use the Pro/E point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Set Point button and using the Pro/E point selection procedure to choose an existing point in the Pro/E window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete selected shape button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select Add Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-Pro/E* toolbar.

- Alternatively, select Mesh > Generate Surface Mesh from the *STAR-Pro/E* menu.



Output from the surface mesher will be displayed in the *STAR-Pro/E Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-Pro/E* toolbar.

- Alternatively, select Mesh > Generate Volume Mesh from the *STAR-Pro/E* menu.



Information on the progress of the volume mesher will be displayed in the *STAR-Pro/E Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Mesh > Clear Mesh from the *STAR-Pro/E* menu.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-Pro/E* toolbar.

Alternatively, select Mesh > Show/Hide Mesh from the *STAR-Pro/E* menu.



Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-Pro/E will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-Pro/E Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Running in Batch Mode
- Clearing the Solution
- Visualizing the Solution

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the STAR-Pro/E node within the object tree and select Set Solution Controls.

- Alternatively, from the menu, select *STAR-Pro/E* > *Settings* > *Solution Controls*.



The *Solution Controls* dialog appears.



2. Here, you can:
   - Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.
   - Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.
3. Click OK to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:

1. In the *Solution Controls* dialog, expand ▶ the *Advanced* box.

The advanced options appear.



1. In the Number of Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:
   - First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.

- Second Order — This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:
   - Momentum
   - Pressure
   - Turbulence (with turbulence modelling defined in Continuum Settings)
   - Temperature (with temperature modelling defined in Continuum Settings)
     - Fluid Temperature Under-Relaxation Factor
     - Solid Temperature Under-Relaxation Factor
5. If you want to use additional STAR-CCM+ physics models with STAR-Pro/E, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros
6. Click OK to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-Pro/E by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-Pro/E:

1. Open the *Solution Controls* dialog.
2. Expand the advanced solver options.
3. Activate the Use Java File option.

   A text entry field appears.



4. Click the Select Input File button. This action opens a browser dialog.
5. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
6. Click OK to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-Pro/E.

Using STAR-CCM+ Continuum Physics with STAR-Pro/E

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-Pro/E and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.
2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.
5. Return to STAR-Pro/E, open the *Solution Controls* dialog and activate the Use Java File checkbox.
6. Select the Java macro that you created earlier.
7. Run the case from STAR-Pro/E.

Checking the Model for Errors

STAR-Pro/E monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.

- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.



2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-Pro/E object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-Pro/E root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check Model for Errors in the *STAR-Pro/E* menu.



- Alternatively click the Check Model for Errors button in the *STAR-Pro/E* toolbar.

If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-Pro/E Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-Pro/E, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-Pro/E* menu.



- Alternatively click the Build and Run button in the *STAR-Pro/E* toolbar.

This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-Pro/E Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Solution > Run Solver in the *STAR-Pro/E* menu.



- Alternatively, click the Run button in the *STAR-Pro/E* toolbar.



STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Solution > Run for More Iterations in the *STAR-Pro/E* menu.



- Alternatively you can click the Run for More Iterations button in the *STAR-Pro/E* toolbar.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Running in Batch Mode

To run the solver in batch mode:

1. Choose or create a directory from which to launch STAR-Pro/E in batch mode. This is referred to as the Batch Runtime Folder.
2. Set up the desired case in STAR-Pro/E, assigning all relevant continuum, region, boundary, mesh pipeline, and solution control settings.
3. Place the geometry file in the Batch Runtime Folder.
4. Launch Pro/E again, open the previously saved CAD geometry file and set the Number of Iterations in the *Solution Controls* dialog to 1 iteration.
5. Build and Run the case and wait until the solver has finished. When done, close STAR-CCM+ and STAR-Pro/E without saving.

6. This process is done in order to obtain the new trail file that contains the instructions to set the number of iterations and build and run the case. This trail file will be edited below. The default trail file should be called `trail.txt.1`, providing no other trail files already exist in the same location.
7. Copy the new trail file from your current working directory to the Batch Runtime Folder and rename it to `runbatch.txt`.
8. Open `runbatch.txt` in a text editor and locate the following lines:

```
~ Update `cda_solutionControls` `iter` \
1
```

9. Replace `1` in the second line with the number of iterations the case is to run for in batch mode, for example, `100`.

```
~ Update `cda_solutionControls` `iter` \
100
```

10. Save and close the text file.
11. Launch a Command Prompt.
12. In the Command Prompt, change to the Batch Runtime Folder that contains the CAD geometry with the saved STAR-Pro/E settings.

    An example code is shown below:

```
C:\>cd "C:\Batch Runtime Folder"_
```

13. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

```
"C:\Program Files\proeWildfire 5.0\bin\proe.exe" +Batch
runbatch.txt -g:no_graphics
```

    where:

    - `proeWildfire 5.0` is the installed version of Pro/E. Refer to the release notes to see which versions of Pro/E are supported.
    - `runbatch.txt` is the name of the training file.
    - `-g:no_graphics` disables the graphics window of Pro/E.
14. Press <`Enter`> to run the simulation in batch mode.

STAR-Pro/E will work in the background, meshing the geometry and running the solvers. A simulation file will be saved in the Batch Runtime Folder once the solver has finished running. If there is an existing simulation file in the Batch Runtime Folder, a backup file is created with a `.sim~` extension. A log file (`starcad.log`) is also created and saved in the Batch Runtime Folder—this log file contains a record of all the messages that appeared in the command prompt during batch processing.

Launching a Command Prompt

To launch a command prompt in Windows 7:

1. Click Start.
2. Type `cmd` in the search field at the bottom of the *Start* menu and select `cmd.exe` from the search results.

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-Pro/E case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the menu, select *STAR-Pro/E* > *Solution* > *Clear Solution*.



The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and

vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-Pro/E.

- Exporting a Surface Mesh
- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-Pro/E as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-Pro/E menu, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-Pro/E as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-Pro/E menu, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-Pro/E menu, select Export > Engineering Report (.html, .csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:

    - Pressure (Pa)
    - Total Pressure (Pa)
    - Pressure Coefficient
    - Total Pressure Coefficient
    - Turbulent Kinetic Energy (J/kg)
    - Turbulent Dissipation Rate (m^2/s^3)
    - Turbulent Viscosity (Pa-s)
    - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
    - Pressure (Pa)
    - Total Pressure (Pa)
    - Pressure Coefficient
    - Total Pressure Coefficient
    - Turbulent Kinetic Energy (J/kg)
    - Turbulent Dissipation Rate (m^2/s^3)
    - Turbulent Viscosity (Pa-s)
    - Density (kg/m^3)
- Force data, at each boundary:
    - Net Force [Fx, Fy, Fz]
    - Pressure Force [Fx, Fy, Fz]
    - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
    - Coefficient of Net Force [Cfx, Cfy, Cfz]
    - Coefficient of Pressure Force [Cfx, Cfy, Cfz]
    - Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
    - Physics models applied to the continuum

    - Materials properties
    - Turbulence properties
    - Segregated flow properties
    - Reference values
    - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
    - Polyhedral mesher properties
    - Prism layer mesher properties (if activated)
    - Surface remesher properties
    - Reference values
    - Volumetric controls
- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values
- Solver parameters
- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-Pro/E to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-Pro/E, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-Pro/E menu, select Export > Simulation (.sim).



Resetting STAR-Pro/E Settings

STAR-Pro/E saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-Pro/E settings:

- From the menu, select Settings > Clear All.



This action returns the geometry to its original state.

Best Practice Guidelines

This section aims to provide guidance and assistance on making a geometry part suitable for CFD analysis using STAR-Pro/E. The following items are covered:

- Imprinting:
  - Imprinting Parts in an Assembly
  - Imprinting Surfaces on a Part for Boundary Region Applications

Prior to using these methods, ensure that the components in assemblies are all solid bodies and are in contact with each other. Any gaps between components will need to be closed.

A basic knowledge and understanding of both CAD and CFD are assumed, although most basic functions are covered.

Imprinting Parts

Most CAD model assemblies are not suitable for meshing as the parts in the assembly are disjointed, where the coincident faces are overlapping. Attempting to generate a mesh with such an assembly will result in a meshing error. To rectify this problem, the parts in an assembly must be imprinted together.

Imprinting can also be used to split surfaces for boundary region applications. An example of such an application would be to model the area that a fan would occupy on the inside of computer case, such as in the third STAR-Pro/E tutorial, where a smaller section of the surface is needed to represent the fan, rather than the entire inner surface of the computer case.

Imprinting Parts in an Assembly

In cases containing multiple adjacent parts, such as in Conjugate Heat Transfer cases, it is essential to imprint the profiles of the coincident faces of the smaller parts onto the face of the larger part. This will ensure that the volume mesh created by STAR-CCM+ contains interfaces between the adjacent parts.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work within the assembly, making the part that requires imprinting the active part.
2. Project the contact edges of the other components in the assembly.
3. Create a copy of the surface that requires imprinting, this will become a quilt.
4. Cut out the contact area of the other parts from the new quilt
5. Sew the quilt onto the surface that requires imprinting.

The simple circuit board shown below requires imprinting before the mesh can be generated.



In this example, the capacitors and heat sinks are not imprinted to the circuit board. This would result in a failure when attempting to generate a mesh.

To imprint the capacitors and heat sinks onto the circuit board follow the steps outlined below:

1. Load the assembly into Pro/E.

2. Activate the base board by right-clicking on it in the Model Tree and selecting Activate.



The inactive parts will become transparent.

We will make a copy of the contact edges of the other parts. In this example these are the capacitors and heat sinks.

3. Select Edit > Project from the menu.



4. Move the mouse pointer over an edge so that it becomes highlighted.



5. Right-click on the edge to cycle through the selections. When all the contact edges are highlighted, as shown below, click to select them.



6. Once selected, the edges will turn red.



7. Repeat the same steps for all the other contact edges. Hold down the <`Ctrl`> key to select multiple items.



8. In the panel at the top of the screen, click in the Surfaces field

and click on the top face of the base board.

9. In the same panel, click in the Direction field next to the Along Direction option



and click on an edge that is normal to the face you selected in the previous step. Ensure that the direction of the arrow is pointing towards the face to be imprinted.



Click on the arrow to switch direction if needed.

10. Apply these settings.

We will now make a copy of the surface that we want to imprint.

11. Select Insert > Advanced > Surface Free Form... from the menu.



12. In the *SURFACE: Free Form* dialog, select the face you want to copy, this is the face you want to imprint. Pro/E will display a grid with a number of points in the vertical and then horizontal directions--accept all default settings. A new quilt will then be created and added to the feature tree; this will be named Surface id xx.



Tip: to show features in the Model Tree, click on the Settings button, and select Tree Filters from the drop-down menu.

The *Model Tree Items* dialog appears.

13. Tick the checkbox next to Features.



We will now trim the quilt to remove the contact area of each part.

14. Select the newly created quilt in the feature tree, and click on Edit > Trim.



15. Select one of the profiles you created in the projection.



A grid will appear in the geometry window; the shaded area will be the region that is kept after the operation. Ensure that the arrow is pointing outwards so that the larger area is kept. If necessary, flip the direction of the trim by clicking on the arrow in the geometry window. Only one profile can be selected at a time.

16. Repeat the same steps for all remaining profiles in the projection feature--always start by clicking on the quilt feature in the Model Tree. Note that the quilt is automatically updated after each trim operation. The final result will be a quilt with all the contact areas removed.

We will now merge the quilt onto the existing face we want to imprint--this will split the face.

17. Click on the quilt feature in the Model Tree again, and select Edit > Solidify.



18. In the panel at the top of the screen, ensure that the (Replaces portion of surface with quilt) button is selected.

19. Ensure that the direction of the operation is pointed towards the surface you want to split. The direction can be controlled by clicking on the arrow in the geometry window.



The face is now imprinted with the contact edges of the other parts.

20. In the feature tree, hide the sketch profile and the quilt by right-click on these and select Hide.
21. Return to the assembly level by right-clicking on the assembly root node and select Activate.

Hiding the other parts and selecting the base board face shows that this is now imprinted.



A high quality mesh with fully defined interfaces can now be successfully generated for this assembly. Suppressed parts are not included in the STAR-CCM+ analysis, so for CHT cases, such as the above example, ensure that all relevant parts are active.

Imprinting Surfaces on Parts for Boundary Region Applications

Faces in a part body can be split into smaller faces for boundary application.

The general procedure to follow is:

1. Create a sketch on the surface you want to imprint. Each closed profile will correspond to a new surface.
2. Create a copy of the surface that requires imprinting--this will result in a quilt.
3. Trim the quilt using the profiles you created.
4. Sew the quilt onto the surface that requires imprinting.

A simple example of an aeration tank is shown below. The image shows the tank, the two air injectors need to be imprinted onto the base of the tank.



1. Load the part into Pro/E.
2. Create a sketch on the face that requires imprinting, in this example the bottom face.
3. Draw the profile of the boundary you want to imprint. In this example the air injectors are circular in profile.



4. Create any additional profiles in new sketches.

We will now make a copy of the surface that requires imprinting.

5. Select Insert > Advanced > Surface Free Form... from the menu.



6. In the *SURFACE: Free Form* dialog, select the face you want to copy, this is the face you want to imprint on. Pro/E will display a grid with a number of points in the vertical and then horizontal directions -- accept all default settings.



A new quilt is created, shown above in dark grey, and a new feature is added to the feature tree; this will be named Surface id xx.



We will now trim the quilt using the profiles in the sketches we created earlier.

7. Select the newly created quilt in the feature tree, and click on Edit > Trim.



8. Select one of the profiles you created earlier.



A grid will appear in the geometry window; the shaded area will be the region that is kept after the operation. Ensure that the arrow is pointing outwards so that the larger area is kept. If necessary, flip the direction of the trim by clicking on the arrow in the geometry window.

9. Click on the quilt surface again and perform the same operation for all other profiles on the same face. The quilt is automatically updated after each operation so that the end result will be a final quilt with all trimmed areas removed.



We will now sew the quilt onto the face we want to imprint.

10. Click on the quilt feature in the Model Tree, and select Edit > Solidify.



11. In the panel at the top of the screen, ensure that the (Replaces portion of surface with quilt) button is selected.

12. Ensure that the direction of the operation is pointed towards the surface you want to split. The direction can be controlled by clicking on the arrow in the geometry window.



13. In the feature tree, hide the sketch profile and the quilt by right-clicking on these and selecting Hide.



The face is now imprinted; selecting one of the new imprinted faces will select all the faces whose profile was defined in the sketch. If multiple sketches were created (one for each face profile), then selecting a face will only highlight that one face.



The separate faces can now be defined as boundary regions.

Working With STAR-Creo

This part of the documentation details how to create, set up and run a CFD simulation in STAR-Creo. It assumes that you have previous experience in using Creo, and only describes the additional features available as part of STAR-Creo. The following sections are included:

- Introducing STAR-Creo
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-Creo Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-Creo Settings

You use STAR-Creo to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-Creo

STAR-Creo provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the Creo environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-Creo
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-Creo enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-Creo provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-Creo interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

  - Ideal Gas
- Porous Media
  - Isotropic
  - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-Creo provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the Creo environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
  - Tetrahedral core mesh generation
  - Polyhedral core mesh generation
  - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-Creo allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-Creo supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-Creo interface, and basic concepts. Topics include:

- Launching STAR-Creo
- The STAR-Creo Info App Window
- Modifying the config.xml File
- Understanding the STAR-Creo Interface
- The STAR-Creo Object Tree
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-Creo Settings While the STAR-CCM+ Client is Running

Launching STAR-Creo Manually

By default, STAR-Creo automatically starts once Creo launches. The STAR-Creo Info App window appears, along with the STAR-Creo object tree and ribbon.

If this does not happen, you can manually launch STAR-Creo by following the steps outlined below:

1. Start Creo as usual.
2. Create a blank part or load an existing part.

3. In the ribbon, select Auxiliary Applications from the Tools tab. The *Auxiliary Applications* dialog appears.
4. In the *Auxiliary Applications* dialog, click Register...



5. Browse to the `StarCrerotk.dat` file in your STAR-CAD series installation, typically:

   `C:\Program Files\CD-adapco\STARCAD7.0x.0xx\STAR-Creo\protk.dat`

6. Select this file and click Open. An entry stating `StarCreo.dllNot Running` appears in the *Auxiliary Applications* dialog.
7. Select this entry and click Start.



The *STAR-Creo Info App* window appears, and the checks for licensing begins. If everything is in order, the `Licensing successful` message appears in the Status page of the *STAR-Creo Info App* window, and the `"StarCreo.dll Not Running"` message will vanish from the *Auxiliary Applications* dialog.

8. Click Close in the *Auxiliary Applications* dialog.
9. Close the open geometry file (File > Close) and reopen it again.

You can now access STAR-Creo features from the *STAR-Creo* object tree and ribbon.

The STAR-Creo Info App Window

The *STAR-Creo Info App* window automatically appears when you start Creo. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-Creo. More tabs appear as you create or load geometry files in STAR-Creo, allowing you to perform more actions.



You cannot close the *STAR-Creo Info App* window while Creo or STAR-Creo are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-Creo license checked out successfully. The version of STAR-Creo is also displayed in this tab.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-Creo.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-Creo searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-Creo user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-Creo is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-Creo Info App. This can be set to one of the following:
  - `Normal`—This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose`—This will display a moderate amount of information in the Output panel.
  - `Log`—This will display all available information in the Output panel.

Understanding the STAR-Creo Interface

The STAR-Creo user interface is embedded within the Creo environment, with the addition of the STAR-Creo feature tree and ribbon.



The majority of the STAR-Creo operations are executed within the feature tree and ribbon. Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-Creo interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings

- Region settings
- Add new boundary
- Boundary conditions
- Mesh Pipeline settings
- Add new mesh controls

Ribbon

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
- Solution
  - Build and Run
  - Run Solver
  - Run For More Iterations
  - Clear Solution
- Model
  - Check Model For Errors
  - Transfer CAE Model
- Export
  - Surface Mesh (.dbs)
  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Geometry
  - Transfer Geometry
- STAR-Creo Help

- About STAR-Creo

The STAR-Creo Object Tree

STAR-Creo uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-Creo. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-Creo applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-Creo Interface section.

Clicking on the Region node within the object tree will highlight all the relevant faces on the geometry model for that particular region. If different parts are assigned to multiple regions, then clicking on a region will only highlight the part and faces associated with that specific region. For example, clicking on the Fluid node below will highlight all the faces associated with the Fluid region.





Similarly, clicking on the Solid node highlights all the faces associated with the Solid region.





Clicking on a boundary node in the STAR-Creo object tree highlights all the faces on the CAD geometry model that are associated with it. Note that clicking on the *Default* boundary node will not highlight any faces on the CAD geometry model. To stop highlighting the faces, you can either click on a blank space in the geometry window, or click on the root node in the

STAR-Creo object tree.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select Preferences from the ribbon.



   The *Preferences* dialog appears.

2. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

3. To run STAR-Creo in parallel, select Parallel from the *Run Mode*. The Compute Processes field appears.



4. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.

5. Click OK to close the dialog.

   If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-Creo and STAR-CCM+ need to be closed for the changes to take effect.



6. Click OK to proceed.
7. Close the active STAR-Creo and STAR-CCM+ sessions.
8. Restart STAR-Creo.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-Creo.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-Creo.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-Creo.

To change the units system:

1. Bring up the *Units System* dialog:
   - Right-click the STAR-Creo node within the object tree and select Set Units System.



   - Alternatively, from the ribbon, select Units System.



   The *Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click OK to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers
- Angstroms
- Nanometers
- Mils
- Microinches
- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-Creo environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the ribbon.



The STAR-CCM+ version affiliated with STAR-Creo launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Using Creo Mouse Controls in STAR-CCM+

STAR-CCM+ offers the ability to set the mouse interactor style to match that of Creo; this could make geometry manipulation in STAR-CCM+ more natural if you are used to working with Creo.

To change the mouse interactor style:

1. Launch STAR-CCM+.
2. Click *Tools* and select Options from the menu.



The *Options* dialog appears.

3. Expand the STAR-CCM+ node and select Visualization.
4. In the *Properties* window, select Interactor Style and set this to pro-E / NX using the drop-down menu.



5. Click Close.

STAR-CCM+ will now use this mouse interactor style each time you launch it.

Editing STAR-Creo Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-Creo while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-Creo, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Disabling STAR-Creo

If for any reason you would like to temporarily disable STAR-Creo without having to uninstall it, do the following:

1. Go to the text folder in the Creo installation directory, typically:

   `C:\Program Files\PTC\Creo 1.0\Common Files\M0x0\text`

2. Locate and edit the `config.pro` file using a text editor.
3. Comment out the following line with an "!":

   ```
   protkdat C:\Program Files\CD-adapco\STAR-CAD8.0x.0xx\ST
   AR-Creo\protk.dat
   ```

   so that it looks like this:

   ```
   !protkdat C:\Program Files\CD-adapco\STAR-CAD8.0x.0xx\S
   TAR-Creo\protk.dat
   ```

Where `x` is the relevant version number of your installation.

Creo will now launch without STAR-Creo. To enable STAR-Creo again undo the changes listed above.

Common Functions

The following section details common operations used within the STAR-Creo environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-Creo
- Removing surfaces and parts from objects in STAR-Creo

Renaming an Object

To rename an object:

1. Right-click the object node within the object tree and select Rename.



2. Enter the name.



3. Click on a blank space in the object tree.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, the Delete option is disabled in the pop-up menu.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select Edit...



The settings dialog for the object appears.

2. Make the changes that you want.
3. Click OK to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-Creo

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-Creo object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

1. Click the Select Face button .

   This button activates the Creo surface selection mechanism and allows you to select the faces on the CAD geometry.

2. On the CAD geometry, click the surfaces that you want to assign to the object. You can select multiple faces by holding down the `<Ctrl>` key whilst clicking the surfaces.

The selected surfaces are highlighted on the CAD geometry.



The selected surfaces are added to the *Included Surfaces* box in the *Surface Sets* dialog.



3. Click OK to close the *Surface Sets* dialog.

The surfaces are then added to the *Face Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

1. Click the Select part button .

This button activates the Creo part selection mechanism and allows you to select parts on the CAD geometry.

2. On the CAD geometry, click the part that you want to assign to the object. You can select multiple parts by holding down the `<Ctrl>` key whilst clicking the parts.

The selected part is highlighted on the CAD geometry.



The selected parts are added to memory in the Creo part selection mechanism.

Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-Creo

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-Creo object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

1. In the *Face Selection* box, select the surface you want to delete.

The selected surface on the CAD geometry is cleared.



2. Click the Remove a face button.



Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

1. In the *Part Selection* box, select the part you want to remove.

The selected part on the CAD geometry is highlighted in black.



2. Click the Remove a part button.



The part is removed from the *Part Selection* box. You can select multiple parts for deletion by holding down the <`Ctrl`> key, and remove them in one operation.

Transferring Geometry to STAR-CCM+

With STAR-Creo, you can transfer your CAD geometry directly to STAR-CCM+ from within the Creo environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

To transfer the geometry to STAR-CCM+:

1. Click Transfer Geometry on the STAR-CCM+ ribbon.

STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-Creo. This action ensures that the simulation file remains associated with the Creo file. You can also safely rename and move the file using this method.
- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the Creo file.

Resuming a Simulation

To resume a simulation after closing STAR-Creo and STAR-CCM+:

1. Launch STAR-Creo.
2. Open the CAD geometry in STAR-Creo.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click Select Parameters on the STAR-Creo ribbon.



The *List of Expressions* dialog appears.



2. Select the design parameters that you want to transfer from the list.

Note: STAR-Creo only recognizes native CAD features as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Select/Deselect option.
4. Click OK.
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-Creo, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



   Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-Creo, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

The geometry updates according to the design parameter changes.



6. Return to Creo. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-Creo and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

Updating the Geometry in STAR-CCM+ with Creo Closed

You can close Creo and continue working with the simulation in STAR-CCM+. Doing so, however, breaks the link between the two software packages. Any further updates that you make to the geometry in STAR-CCM+ are not transferred back to the geometry in Creo. Instead, these changes are stored in memory, which allows you to include the design parameters in your optimization scripts.

This workflow also allows you to release the Creo and STAR-Creo licenses when performing tasks in STAR-CCM+. In this case, the Creo and STAR-Creo licenses are temporarily checked out only when the updates to the geometry are stored in memory.

Transferring the geometry from STAR-Creo will overwrite the geometry in STAR-CCM+.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from Creo to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-Creo root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click OK.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.

Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.
Transfer Sheet Geometry	Specifies whether to transfer sheet bodies as part of the transfer process. This option is turned off by default.

Understanding the STAR-Creo Process

In STAR-Creo the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume
3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity
5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings
7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-Creo and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-Creo

STAR-Creo only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-Creo.

Suppressed or deactivated parts and features are not supported by STAR-Creo and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-Creo.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-Creo includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-Creo includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-Creo is completely integrated within the Creo environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-Creo model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

Including material type and properties.

- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.

- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new

regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.
- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-Creo automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select New.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



- Right-click on the Regions manager node in the object tree and select New.

  Right-click on the Regions manager node in the object tree and select Add Region.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-Creo). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-Creo model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. Then right-click on the region node under the Region manager node in the object tree and select Edit... In the *Region Settings* dialog, select the alternative continuum from the *Continuum Selection* box and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-Creo each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By default the material type is set to Air, and this can be changed to any one of a number of predefined solids, liquids or gases that are available in STAR-Creo. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-Creo object tree and select Edit.



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material* box, you can select a material type.
- In the *Material Properties* box, you can define each of the material properties.



Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-Creo offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the drop-down menu. When you have selected a material,

the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-Creo are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.

Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-Creo, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)
- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select Edit...

The physics models available in the *Continuum Settings* dialog are determined by the material type.



Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.



Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.

The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-Creo are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Locate the *Gravity* box and tick the *Enable* checkbox. This will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. To define reference properties for a continuum, right-click the Continuum node in the object tree and select Edit. In the *Continuum Settings* dialog, locate the *Reference Properties* box and expand , if necessary.



The reference properties are described below:

- Reference Pressure

  The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density , Reference Velocity , Reference Area .

  These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-Creo, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select Edit...



The physics models can then be enabled for the region by ticking the corresponding checkbox and entering values as necessary.



Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the *Porous* checkbox.



There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select the radio button labeled *Isotropic*.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, α and β respectively.

Uni-Directional

To specify uni-directional porous resistance, select the radio button labeled *Uni-Directional*.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, α and β respectively. The cross-flow resistance R is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, $\mathbf{i}$ $\mathbf{j}$ $\mathbf{k}$.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the *Moving Reference Frame* checkbox.



In the *Moving Reference Frame Settings* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, . The units used are the same as the file defaults.

Alternatively, click on and select two vertices on the geometry to define the axis of rotation. Holding down `<Ctrl>`, select two vertices on the geometry; the first selection will act as the origin of the axis, and the second selection will define the direction vector of the axis. The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.

This action activates the Energy Source tab. There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select the radio button labeled *Energy Input per Unit Volume*.



Enter the value for the energy input value per unit volume in the box.

Total Energy Input

To define the total energy input, select the radio button labeled *Total Energy Input*.



Enter the value for the total energy input in the box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.
2. Create a streamlines derived part.
3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-Creo.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-Creo will use default values to initialize the solution. In STAR-Creo, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-Creo object tree and select Edit.



This opens the *Continuum Settings* dialog, in which the *Initialization* tab can be used to specify the following values:



- U V W Initial velocity
- Initial pressure
- Initial temperature
- Int len Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial velocity, open the *Continuum Settings* dialog, expand the *Initialization* box, and enter the values in the U, V, and W boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, locate the *Initialization* box, and enter a value in the box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the Turbulence Intensity Int and Turbulent Length Scale len boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-Creo boundary is made up of one or more CAD surfaces within the simulation which, along with other

boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select Add Boundary from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-Creo). When you assign a surface to another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-Creo, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Velocity Inlet from the drop-down menu.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity, with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors U V W, select Cartesian Components from the drop-down menu.



Cylindrical vectors V_r V_θ V_z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i j k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Mass Flow Inlet from the drop-down menu.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.

The mass flow rate , must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors i, j, k, tick the checkbox labeled Specify Direction.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Stagnation Inlet from the drop-down menu.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure , must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors i, j, k, tick the checkbox labeled Specify Direction.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-Creo, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Flow-Split Outlet from the drop-down menu.

The only parameter required is the flow-split ratio, expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.

Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Pressure Outlet from the drop-down menu.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.



The pressure , must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Wall from the drop-down menu.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U V W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity ω. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i j k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux , to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature , and heat transfer coefficient .



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Symmetry Plane from the drop-down menu.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-Creo to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-Creo and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature, Turbulence*.



Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Temperature, Turbulence*.



Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity
- Turbulence Length Scale

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Passive Scalar*.



This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-Creo provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-Creo* ribbon), it is possible to execute only single stages from the STAR-Creo ribbon. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the Creo parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from Creo.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the Creo *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-Creo creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-Creo appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-Creo appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all Creo parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose Edit.



The *Mesh Pipeline Settings* dialog appears. This dialog has several collapsible sections for the general settings.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click OK to save them and close the dialog.
- Click Cancel to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.



Selecting the trimmer option will activate the Trimmer tab..

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-Creo Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.



The base size can be set as either a relative or absolute value. When it is set as a relative value (by selecting the Relative to Geometry option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* box of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.



This panel allows you to change the following parameters:

- Relative Minimum Surface Size
- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-Creo will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer panel.



The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the Creo parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when

they are created. A part can only belong to one mesh pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select Add Mesh Pipeline.



During the creation of a new mesh pipeline, it is necessary to assign to it the Creo parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

  This mesh control is applied to one or more faces of the Creo model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose Add Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select Add Volumetric Mesh Control from the pop-up menu.



The *Volumetric Control Settings* dialog appears.



The *Volumetric Control Settings* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape
- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the Creo geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.



X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the Creo window. To use an existing point, click the Set Point button and use the Creo point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Set Point button and using the Creo point selection procedure to choose an existing point in the Creo window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select Add Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-Creo* ribbon.



Output from the surface mesher will be displayed in the *STAR-Creo Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-Creo* ribbon.



Information on the progress of the volume mesher will be displayed in the *STAR-Creo Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Clear Mesh from the ribbon.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-Creo* ribbon.



Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-Creo will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-Creo Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Running in Batch Mode
- Clearing the Solution
- Visualizing the Solution

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the STAR-Creo node within the object tree and select Set Solution Controls.



- Alternatively, from the ribbon, select Solution Controls.



The *Solution Controls* dialog appears.



2. Here, you can:
   - Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.
   - Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.
3. Click OK to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:

1. In the *Solution Controls* dialog, expand ▶ the *Advanced* box.

   The advanced options appear.



1. In the Number of Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:
   - First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.
   - Second Order—This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:
   - Momentum
   - Pressure
   - Turbulence (with turbulence modelling defined in Continuum Settings)
   - Temperature (with temperature modelling defined in Continuum Settings)
     - Fluid Temperature Under-Relaxation Factor
     - Solid Temperature Under-Relaxation Factor
5. If you want to use additional STAR-CCM+ physics models with STAR-Creo, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros
6. Click OK to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-Creo by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-Creo:

1. Open the *Solution Controls* dialog.
2. Expand the advanced solver options.
3. Activate the Use Java File option.

   A text entry field appears.



4. Click the Select Input File button. This action opens a browser dialog.
5. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
6. Click OK to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-Creo.

Using STAR-CCM+ Continuum Physics with STAR-Creo

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-Creo and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.
2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.
5. Return to STAR-Creo, open the *Solution Controls* dialog and activate the Use Java File checkbox.
6. Select the Java macro that you created earlier.
7. Run the case from STAR-Creo.

Checking the Model for Errors

STAR-Creo monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.
- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.

2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-Creo object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-Creo root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check Model for Errors in the *STAR-Creo* ribbon.



If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-Creo Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-Creo, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-Creo* ribbon.



This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-Creo Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Run Solver in the *STAR-Creo* ribbon.



STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Run For More Iterations in the *STAR-Creo* ribbon.



Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Running in Batch Mode

To run the solver in batch mode:

1. Choose or create a directory from which to launch STAR-Creo in batch mode. This is referred to as the Batch Runtime Folder.
2. Set up the desired case in STAR-Creo, assigning all relevant continuum, region, boundary, mesh pipeline, and solution control settings.

3. Place the geometry file in the Batch Runtime Folder.
4. Launch Creo again, open the previously saved CAD geometry file and set the Number of Iterations in the *Solution Controls* dialog to 1 iteration.
5. Build and Run the case and wait until the solver has finished. When done, close STAR-CCM+ and STAR-Creo without saving.
6. This process is done in order to obtain the new trail file that contains the instructions to set the number of iterations and build and run the case. This trail file will be edited below. The default trail file should be called `trail.txt.1`, providing no other trail files already exist in the same location.
7. Copy the new trail file from your current working directory to the Batch Runtime Folder and rename it to `runbatch.txt`.
8. Open `runbatch.txt` in a text editor and locate the following lines:

```
~ Update `cda_solutionControls` `iter` \
1
```

9. Replace `1` in the second line with the number of iterations the case is to run for in batch mode, for example, `100`.

```
~ Update `cda_solutionControls` `iter` \
100
```

10. Save and close the text file.
11. Launch a Command Prompt.
12. In the Command Prompt, change to the Batch Runtime Folder that contains the CAD geometry with the saved STAR-Creo settings.

    An example code is shown below:

```
C:\>cd "C:\Batch Runtime Folder"_
```

13. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

```
"C:\Program Files\PTC\Creo
1.0\Parametric\bin\parametric.exe" +Batch runbatch.txt
-g:no_graphics
```

    where:

    - `Creo 1.0` is the installed version of Creo. Refer to the release notes to see which versions of Creo are supported.
    - `runbatch.txt` is the name of the training file.
    - `-g:no_graphics` disables the graphics window of Creo.

14. Press <`Enter`> to run the simulation in batch mode.

STAR-Creo will work in the background, meshing the geometry and running the solvers. A simulation file will be saved in the Batch Runtime Folder once the solver has finished running. If there is an existing simulation file in the Batch Runtime Folder, a backup file is created with a `.sim~` extension. A log file (`starcad.log`) is also created and saved in the Batch Runtime Folder—this log file contains a record of all the messages that appeared in the command prompt during batch processing.

Launching a Command Prompt

To launch a command prompt in Windows 7:

1. Click Start.
2. Type `cmd` in the search field at the bottom of the *Start* menu and select `cmd.exe` from the search results.

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-Creo case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the ribbon, select Clear Solution.



The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and

vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-Creo.

- Exporting a Surface Mesh
- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-Creo as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-Creo ribbon, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-Creo as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-Creo ribbon, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-Creo ribbon, select Export > Engineering Report (.html,.csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)

    - Turbulent Viscosity (Pa-s)
    - Density (kg/m^3)
- Force data, at each boundary:
    - Net Force [Fx, Fy, Fz]
    - Pressure Force [Fx, Fy, Fz]
    - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
    - Coefficient of Net Force [Cfx, Cfy, Cfz]
    - Coefficient of Pressure Force [Cfx, Cfy, Cfz]
    - Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
    - Physics models applied to the continuum
    - Materials properties
    - Turbulence properties
    - Segregated flow properties
    - Reference values
    - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
    - Polyhedral mesher properties
    - Prism layer mesher properties (if activated)
    - Surface remesher properties
    - Reference values
    - Volumetric controls
- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values

- Solver parameters
- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-Creo to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-Creo, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-Creo ribbon, select Export > Simulation (.sim).



Resetting STAR-Creo Settings

STAR-Creo saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-Creo settings:

- From the ribbon, select Clear All.



This action returns the geometry to its original state.

Working With STAR-Inventor

This part of the documentation details how to create, set up and run a CFD simulation in STAR-Inventor. It assumes that you have previous experience in using Inventor, and only describes the additional features available as part of STAR-Inventor. The following sections are included:

- Introducing STAR-Inventor
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-Inventor Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-Inventor Settings
- Best Practices Guidelines

You use STAR-Inventor to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-Inventor

STAR-Inventor provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the Inventor environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-Inventor
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-Inventor enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-Inventor provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-Inventor interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

  - Ideal Gas
- Porous Media
  - Isotropic
  - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-Inventor provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the Inventor environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
  - Tetrahedral core mesh generation
  - Polyhedral core mesh generation
  - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-Inventor allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-Inventor supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-Inventor interface, and basic concepts. Topics include:

- Launching STAR-Inventor
- The STAR-Inventor Info App Window
- Modifying the config.xml File
- Understanding the STAR-Inventor Interface
- The STAR-Inventor Object Tree
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Saving Files in Inventor
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-Inventor Settings While the STAR-CCM+ Client is Running

Launching STAR-Inventor Manually

By default STAR-Inventor should start automatically once Inventor is launched.

If this does not happen then you can launch STAR-Inventor manually by following the steps outlined below:

1. Start Inventor as usual.

2. Click on the Tools tab to reveal the *Tools* ribbon.
3. Click the Add-Ins button on the *Tools* ribbon. This will launch the *Add-In Manager 20xx* dialog.
4. Select `STAR-Inventor: CFD Analysis` from the list. The *Load Behavior* group box indicates whether STAR-Inventor will load when Inventor is launched.
5. Activate the Loaded/Unloaded checkbox to enable STAR-Inventor for the current session only.

6. Activate the Load On Startup checkbox to automatically load STAR-Inventor every time you start Inventor.



7. Click OK to apply the settings.

The STAR-Inventor feature manager tree will appear to the left of the Inventor window when you load or create a geometry file.

The STAR-Inventor Info App Window

The *STAR-Inventor Info App* window automatically appears when you start Inventor. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-Inventor. More tabs appear as you create or load geometry files in STAR-Inventor, allowing you to perform more actions.



You cannot close the *STAR-Inventor Info App* window while Inventor or STAR-Inventor are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-Inventor license checked out successfully. The version of STAR-Inventor is also displayed in this tab. This tab appears by default when you launch Inventor without a geometry file.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-Inventor.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-Inventor searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-Inventor user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-Inventor is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-Inventor Info App. This can be set to one of the following:
  - `Normal` — This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose` — This will display a moderate amount of information in the Output panel.
  - `Log` — This will display all available information in the Output panel.

Understanding the STAR-Inventor Interface

The STAR-Inventor user interface is embedded within the Inventor environment, with the addition of the STAR-Inventor feature tree and ribbon.



The majority of the STAR-Inventor operations are executed within the feature tree and ribbon. Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-Inventor interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings
- Region settings
- Add new boundary
- Boundary conditions

- Mesh Pipeline settings
- Add new mesh controls

Ribbon

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
  - Depth
- Solution
  - Build and Run
  - Run Solver
  - Run For More Iterations
  - Clear Solution
- Model
  - Check Model For Errors
  - Transfer CAE Model
- Export
  - Surface Mesh (.dbs)
  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Geometry
  - Transfer Geometry
- STAR-Inventor Help
- About STAR-Inventor

The STAR-Inventor Object Tree

STAR-Inventor uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-Inventor. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-Inventor applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-Inventor Interface section.

Clicking on the Region node within the object tree will highlight all the relevant faces on the geometry model for that particular region. If different parts are assigned to multiple regions, then clicking on a region will only highlight the part and faces associated with that specific region. For example, clicking on the Fluid node below will highlight all the faces associated with the Fluid region.





Similarly, clicking on the Solid node highlights all the faces associated with the Solid region.


STAR-Inventor
STAR-Inventor
Continua
Fluid
Solid
Regions
Fluid
Boundaries
Default
Solid
Boundaries
Default
Mesh Pipelines
Mesh Pipeline
Mesh Controls



Clicking on a boundary node in the STAR-Inventor object tree highlights all the faces on the CAD geometry model that are associated with it.Note that clicking on the *Default* boundary node will not highlight any faces on the CAD geometry model. To stop highlighting the faces, you can either click on a blank space in the geometry window, or click on the root node in the STAR-Inventor object tree.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select Preferences from the ribbon.



   The *Preferences* dialog appears.

2. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

3. To run STAR-Inventor in parallel, select Parallel from the *Run Mode* box. The *Parallel Options* box appears.



4. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.
5. Click OK to close the dialog.

   If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-Inventor and STAR-CCM+ need to be closed for the changes to take effect.



6. Click OK to proceed.
7. Close the active STAR-Inventor and STAR-CCM+ sessions.
8. Restart STAR-Inventor.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-Inventor.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-Inventor.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-Inventor.

To change the units system:

1. Bring up the *Units System* dialog:
   - Right-click the STAR-Inventor node within the object tree and select Units System.



   - Alternatively, from the ribbon, select Units System.



   The *Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click OK to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers
- Angstroms
- Nanometers
- Mils
- Microinches
- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Saving Files in Inventor

Using the Save As option to save a geometry file in Inventor results in any mesh displayed on the geometry model, and information in the output window, to be cleared. This occurs due to Inventor closing the old document file after saving the new document. All saved STAR-Inventor settings, such as boundary settings, are not affected by this operation and are retained. The mesh data is not lost. To display the mesh on the geometry model again, click on the Generate Volume Mesh button in the ribbon.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-Inventor environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the ribbon.



The STAR-CCM+ version affiliated with STAR-Inventor launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Editing STAR-Inventor Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-Inventor while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-Inventor, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Common Functions

The following section details common operations used within the STAR-Inventor environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-Inventor
- Removing surfaces and parts from objects in STAR-Inventor

Renaming an Object

To rename an object:

1. Right-click the object node within the object tree and select Rename.



2. Enter the name.



3. Click on a blank space in the object tree.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, no pop-up menu appears.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select Edit...



The settings dialog for the object appears.

2. Make the changes that you want.
3. Click OK to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-Inventor

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-Inventor object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

- On the CAD geometry, click the surfaces that you want to assign to the object. You can add multiple surfaces by clicking more surfaces on the CAD geometry.

The selected surfaces are highlighted on the CAD geometry.



The selected surfaces are also displayed in the *Face Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to assign to the object. You can add multiple parts by clicking more parts on the CAD geometry.

The selected part is highlighted on the CAD geometry.



The selected parts are displayed in the *Part Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-Inventor

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-Inventor object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

1. In the *Face Selection* box, select the surface you want to delete.

The selected surface on the CAD geometry is highlighted in light blue.



2. Click the Remove button.



The surface is removed from the *Face Selection* box. You can select multiple surfaces for deletion by holding down the `<Ctrl>` key, and remove them in one operation.

Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

1. In the *Part Selection* box, select the part you want to remove.

The selected part on the CAD geometry is highlighted in light blue.



2. Click the Remove button.



The part is removed from the *Part Selection* box. You can select multiple parts for deletion by holding down the `<Ctrl>` key, and remove them in one operation.

Transferring Geometry to STAR-CCM+

With STAR-Inventor, you can transfer your CAD geometry directly to STAR-CCM+ from within the Inventor environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

To transfer the geometry to STAR-CCM+:

1. Click Transfer Geometry on the STAR-CCM+ ribbon.



   STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-Inventor. This action ensures that the simulation file remains associated with the Inventor file. You can also safely rename and move the file using this method.
- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the Inventor file.

Resuming a Simulation

To resume a simulation after closing STAR-Inventor and STAR-CCM+:

1. Launch STAR-Inventor.

2. Open the CAD geometry in STAR-Inventor.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click Choose Design Parameters on the STAR-Inventor ribbon.



The *Design Parameters* dialog appears.



2. Select the design parameters that you want to transfer from the list. You can select multiple items at once.

Note: STAR-Inventor only recognizes native CAD features as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Transfer to STAR-CCM+ option.
4. Click OK.
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-Inventor, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-Inventor, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

   The geometry updates according to the design parameter changes.



6. Return to Inventor. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-Inventor and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

Updating the Geometry in STAR-CCM+ with Inventor Closed

You can close Inventor and continue working with the simulation in STAR-CCM+. Doing so, however, breaks the link between the two software packages. Any further updates that you make to the geometry in STAR-CCM+ are not transferred back to the geometry in Inventor. Instead, these changes are stored in memory, which allows you to include the design parameters in your optimization scripts.

This workflow also allows you to release the Inventor and STAR-Inventor licenses when performing tasks in STAR-CCM+. In this case, the Inventor and STAR-Inventor licenses are temporarily checked out only when the updates to the geometry are stored in memory.

Transferring the geometry from STAR-Inventor will overwrite the geometry in STAR-CCM+.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from Inventor to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-Inventor root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click OK.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.
Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.

Transfer Sheet Geometry	Specifies whether to transfer sheet bodies as part of the transfer process. This option is turned off by default.

Understanding the STAR-Inventor Process

In STAR-Inventor the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume
3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity
5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings
7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-Inventor and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-Inventor

STAR-Inventor only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-Inventor.

Suppressed or deactivated parts and features are not supported by STAR-Inventor and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-Inventor.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

  For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-Inventor includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-Inventor includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-Inventor is completely integrated within the Inventor environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-Inventor model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

  Including material type and properties.
- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.
- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.

- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-Inventor automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select New.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



To create an additional region:

- Right-click on the Regions manager node in the object tree and select New.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-Inventor). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-Inventor model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If

multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. Then right-click on the region node under the Region manager node in the object tree and select Edit... In the *Region Settings* dialog, select the alternative continuum using the *Continuum* drop-down menu and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-Inventor each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By default the material type is set to Air, and this can be changed to any one of

a number of predefined solids, liquids or gases that are available in STAR-Inventor. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-Inventor object tree and select Edit.



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material* box in the Physics tab, you can select a material type.
- In the *Material Properties* box in the Material tab, you can define each of the material properties.





Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-Inventor offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the list. When you have selected a material, the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-Inventor are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.
Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-Inventor, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)
- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select Edit...



The physics models available in the *Continuum Settings* dialog are determined by the material type.





Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.



Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.



The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-Inventor are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Click on the Reference tab, locate and tick the *Gravity* checkbox; this will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. To define reference properties for a continuum, right-click the Continuum node in the object tree and select Edit. In the *Continuum Settings* dialog, click on the Reference tab.



The reference properties are described below:

- Reference Pressure

The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density , Reference Velocity , Reference Area .

  These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-Inventor, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select Edit...



The physics models can then be enabled for the region by ticking the corresponding checkbox on the tabs and entering values as necessary.



Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the checkbox on the *Porous* tab.



There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select the radio button labeled *Isotropic*.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, α and β respectively.

Uni-Directional

To specify uni-directional porous resistance, select the radio button labeled *Uni-Directional*.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, α and β respectively. The cross-flow resistance $R_{\perp}$ is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, $\mathbf{i}$ $\mathbf{j}$ $\mathbf{k}$.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the checkbox on the *Moving Reference Frame* tab.



In the *Moving Reference Frame* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the ω box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, **i** **j** **k**. The units used are the same as the file defaults.

The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.



This action activates the Energy Source Settings box. There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select the radio button labeled *Energy Input per Unit Volume*.



Enter the value for the energy input value per unit volume in the box.

Total Energy Input

To define the total energy input, select the radio button labeled *Total Energy Input*.



Enter the value for the total energy input in the box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.
2. Create a streamlines derived part.
3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-Inventor.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-Inventor will use default values to initialize the solution. In STAR-Inventor, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-Inventor object tree and select Edit.



This opens the *Continuum Settings* dialog, in which the *Initialization* tab can be used to specify the following values:



- U V W Initial velocity
- Initial pressure
- Initial temperature
- Int len Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial pressure, open the *Continuum Settings* dialog, click on the Initialization tab and enter a value in the U, V, and W boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, click on the Initialization tab and enter a value in the Pressure box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, click on the Initialization tab, and enter a value in the Temperature box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, click on the Initialization tab and enter a value in the Turbulence Intensity and Turbulent Length Scale boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-Inventor boundary is made up of one or more CAD surfaces within the simulation which, along with other boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select New from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-Inventor). When you assign a surface to another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-Inventor, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Velocity Inlet from the drop-down menu.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity, with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors U V W, tick the checkbox labeled Specify velocity components.



Cylindrical vectors V_r V_θ V_z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i j k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Mass Flow Inlet from the drop-down menu.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.



The mass flow rate , must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors i, j, k, tick the checkbox labeled Specify Components.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Stagnation Inlet from the drop-down menu.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure, must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors i, j, k, tick the checkbox labeled Specify Components.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-Inventor, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Flow-Split Outlet from the drop-down menu.

The only parameter required is the flow-split ratio 60↗ 40↘, expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.



Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Pressure Outlet from the drop-down menu.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.



The pressure , must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Wall from the drop-down menu.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U V W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity ω. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i j k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux , to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature , and heat transfer coefficient .



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Symmetry Plane from the drop-down menu.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-Inventor to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-Inventor and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature*.



Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Turbulence*.



Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity Int
- Turbulence Length Scale len

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Passive Scalar*.



This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-Inventor provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-Inventor* ribbon), it is possible to execute only single stages from the STAR-Inventor ribbon. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the Inventor parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from Inventor.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the Inventor *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-Inventor creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-Inventor appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-Inventor appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all Inventor parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose Edit.



The *Mesh Pipeline Settings* dialog appears. This dialog has several for the general settings. This dialog has several tabs for the general settings.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click OK to save them and close the dialog.

- Click Cancel to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.

Selecting the trimmer option will activate the Trimmer tab..

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-Inventor Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.



The base size can be set as either a relative or absolute value. When it is set as a relative value (by ticking the Relative to bodies option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be

calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* tab of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.



This panel allows you to change the following parameters:

- Relative Minimum Surface Size
- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-Inventor will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer panel.



The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the Inventor parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when they are created. A part can only belong to one mesh

pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select New.



During the creation of a new mesh pipeline, it is necessary to assign to it the Inventor parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

  This mesh control is applied to one or more faces of the Inventor model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose New Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select New Volumetric Mesh Control from the pop-up menu.



The *Volumetric Control* dialog appears.



The *Volumetric Control* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape
- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the Inventor geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.

X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the Inventor window. To use an existing point, click the Choose Point 1... or Choose Point 2... button and use the Inventor point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Choose Point... button and using the Inventor point selection procedure to choose an existing point in the Inventor window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select New Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-Inventor* ribbon.



If the Surface Mesh button is not displayed, click on the expand arrow and click on the Generate Surface Mesh button.



Output from the surface mesher will be displayed in the *STAR-Inventor Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-Inventor* ribbon.



If the Volume Mesh button is not displayed, click on the expand arrow and click on the Generate Volume Mesh button.



Information on the progress of the volume mesher will be displayed in the *STAR-Inventor Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Clear Mesh from the ribbon.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-Inventor* ribbon.
When the mesh is displayed, the Mesh View button appears pressed Show/Hide Mesh, otherwise the standard icon is shown Show/Hide Mesh .

Controlling the Mesh Depth

The Depth option controls the mesh visualization (mesh depth). If the mesh appears to be hidden behind the geometry surface and is hard to see (particularly when zooming in closer, as shown in the screenshot below), changing the mesh Depth option may help improve mesh visualization.



Increasing the Depth value will improve the mesh definition.



To change the Depth settings, click on the Depth = Default button in the STAR-Inventor ribbon and select the desired depth from the drop-down list.



The higher the number the more pronounced the mesh will be, however increasing this value too much will cause the mesh from behind the geometry to start appearing (particularly near the edges). For thin geometries, a high depth number should be avoided as the entire mesh will be visible, including the mesh from behind the geometry.

The default setting (Depth = Default) should work well with most geometries. The original Inventor value is stored in the Depth = Original setting--this tends to hide the mesh when zooming in and is not recommended for most cases.

Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-Inventor will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-Inventor Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted
- Faces May Require Splitting

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Faces May Require Splitting

In some cases specific geometric shapes may not be able to mesh in STAR-Inventor, and an error message such as the one shown below will be displayed in the output window.

```
Creating Surface Mesh
Estimated Cell count = 795 Global scale fact = 1
Surfaces repaired successfully
Preparing surfaces
Meshing surfaces
Could not orient the edges of inner trimming loop 1 of surface 5.
     Winding angle should be 360 degrees, calculated 2.66759e-006 degrees
1 surfaces did not close properly

The faulty faces are highlighted.
Surface Mesh Error: -
49 bad edges detected. Surface Meshing Error
The surfaces 1 3 4 containing bad edges are highlighted.
```

The face that is causing the problem will be highlighted in the geometry window.



This issue can be overcome by splitting the highlighted face in Inventor - the simplest way of achieving this is by creating a plane that dissects the aforementioned face and use it as the splitting tool. The result is shown below.



A mesh can now be generated.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Running in Batch Mode
- Clearing the Solution
- Visualizing the Solution

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the node within the object tree and select Solution Controls.



- Alternatively, from the ribbon, select Solution Controls.



The *Solution Controls* dialog appears.



2. Here, you can:
   - Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.
   - Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.
3. Click OK to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:

1. In the *Solution Controls* dialog, click the Advanced tab.

The advanced options appear.



1. In the Number of Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:
   - First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.
   - Second Order—This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:

- Momentum
- Pressure
- Turbulence (with turbulence modelling defined in Continuum Settings)
- Temperature (with temperature modelling defined in Continuum Settings)
  - Fluid Temperature Under-Relaxation Factor
  - Solid Temperature Under-Relaxation Factor

5. If you want to use additional STAR-CCM+ physics models with STAR-Inventor, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros

6. Click OK to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-Inventor by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-Inventor:

1. Open the *Solution Controls* dialog.
2. Expand the advanced solver options.
3. Activate the Use Java File option.

   A text entry field appears.



4. Click the Select Input File button. This action opens a browser dialog.
5. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
6. Click OK to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-Inventor.

Using STAR-CCM+ Continuum Physics with STAR-Inventor

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-Inventor and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.
2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.

5. Return to STAR-Inventor, open the *Solution Controls* dialog and activate the Use Java File checkbox.
6. Select the Java macro that you created earlier.
7. Run the case from STAR-Inventor.

Checking the Model for Errors

STAR-Inventor monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.
- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.



2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-Inventor object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-Inventor root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check in the *STAR-Inventor* ribbon.



If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-Inventor Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-Inventor, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-Inventor* ribbon.



This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-Inventor Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Run Solver in the *STAR-Inventor* ribbon.



STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Run For More Iterations in the *STAR-Inventor* ribbon.



Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Running in Batch Mode

To run STAR-Inventor in batch mode, you must first add the paths to the `STAR-Inventor.dll` and `STARInventorBatch.exe` files to your Path user environment variable. These paths are added to the list of paths that may already exist on your machine. For more information on how to set the `Path` variable please see the Setting the Path User Environment Variable section.

A typical path is shown below.

```
C:\Program Files\CD-adapco\STAR-CAD<star-ccm+_version>\
STAR-Inventor\win64\bin
```

To run the solver in batch mode:

1. Choose or create a directory from which to launch STAR-Inventor in batch mode. This is referred to as the Batch Runtime Folder.
2. Set up the desired case in STAR-Inventor, assigning all relevant continuum, region, boundary, mesh pipeline, and solution control settings.
3. Place the geometry file in the Batch Runtime Folder.
4. Launch a Command Prompt.
5. In the Command Prompt, change to the Batch Runtime Folder that contains the CAD geometry.

   An example code is shown below:

   ```
   C:\>cd "C:\Batch Runtime Folder"_
   ```

6. Enter the following line in the Command Prompt, taking care to preserve the correct spacing:

   ```
   STARInventorBatch.exe -doc <docname> -ver <inventorexec
   utablefolder>
   ```

   where:

   - `-doc` is the input option to specify the input geometry file. It has to be followed by the `docname` argument which is the full path to the input Inventor part/assembly file (including the extension).
   - `-ver` is an optional input used to specify the version of Inventor you want to use when you have more than one version of Inventor on your machine. Follow this option with the `<inventorexecutable>` argument which is the full path to the executable file of the Inventor version you wish to use. If this option is not used, the latest version of Inventor will be used. Refer to the release notes to see which versions of Inventor are supported.

   An example code for running STAR-Inventor in batch mode using Inventor 2012 is shown below:

   ```
   STARInventorBatch.exe -doc "C:\Batch Runtime Folder\exh
   austManifold.ipt" -ver "C:\Program
   Files\Autodesk\Inventor 2012\Bin"
   ```

   If only one version of Inventor is installed on your machine, you can omit the `-ver` argument when running in batch mode:

   ```
   STARInventorBatch.exe -doc "C:\Batch Runtime Folder\exh
   austManifold.ipt"
   ```

7. Press <`Enter`> to run the simulation in batch mode.

STAR-Inventor will work in the background, meshing the geometry and running the solvers. A simulation file will be saved in the Batch Runtime Folder once the solver has finished running. If there is an existing

simulation file in the Batch Runtime Folder, a backup file is created with a `.sim~` extension. A log file (`starcad.log`) is also created and saved in the Batch Runtime Folder—this log file contains a record of all the messages that appeared in the command prompt during batch processing.

Setting the Path User Environment Variable

You need to set the `Path` user environment variable in order for STAR-Inventor to run in batch mode. This variable will have the full path to the folder containing the `STAR-Inventor.dll` and STARInventorBatch`.exe` files. To create the `Path` user environment variable:

1. Open the Environment Variables dialog:
   - For Windows 7:
     i. Click Start and type `env` into the search field near the bottom of the *Start* menu.
     ii. Select `Edit environment variables for your account` from the list. The *Environment Variables* dialog appears.
     iii. If the Path user environment variable already exists in the *User variables* box click on Edit..., otherwise click on New... to create it. The *User Variable* dialog appears.
2. Set the Variable name to `Path`.
3. Set the Variable value to:

   ```
   C:\Program Files\CD-adapco\STAR-CAD<star-ccm+_version>\
   STAR-Inventor\win64\bin
   ```
4. Click OK to close the *User Variable* dialog.
5. Click OK to close the *Environment Variables* dialog.
6. Click OK to close the *System Properties* dialog.

Launching a Command Prompt

To launch a command prompt in Windows 7:

1. Click Start.
2. Type `cmd` in the search field at the bottom of the *Start* menu and select `cmd.exe` from the search results.

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-Inventor case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the ribbon, select Clear Solution.



The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-Inventor.

- Exporting a Surface Mesh

- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-Inventor as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-Inventor ribbon, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-Inventor as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-Inventor ribbon, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-Inventor ribbon, select Export > Engineering Report (.html,.csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Force data, at each boundary:
  - Net Force [Fx, Fy, Fz]
  - Pressure Force [Fx, Fy, Fz]
  - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
  - Coefficient of Net Force [Cfx, Cfy, Cfz]

- Coefficient of Pressure Force [Cfx, Cfy, Cfz]
- Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
  - Physics models applied to the continuum
  - Materials properties
  - Turbulence properties
  - Segregated flow properties
  - Reference values
  - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
  - Polyhedral mesher properties
  - Prism layer mesher properties (if activated)
  - Surface remesher properties
  - Reference values
  - Volumetric controls
- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values
- Solver parameters
- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-Inventor to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-Inventor, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-Inventor ribbon, select Export > Simulation (.sim).



Resetting STAR-Inventor Settings

STAR-Inventor saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-Inventor settings:

- From the ribbon, select Clear All.



This action returns the geometry to its original state.

Best Practice Guidelines

This section aims to provide guidance and assistance on making a geometry part suitable for CFD analysis using STAR-Inventor. The following items are covered:

- Creating the Fluid Domain:
  - Internal Flow
  - External Flow
  - Multi-Region Flow
- Imprinting:
  - Imprinting Parts in an Assembly
  - Imprinting Surfaces on a Part for Boundary Region Applications

Prior to using these methods, ensure that the components in assemblies are all solid bodies and are in contact with each other. Any gaps between components will need to be closed.

A basic knowledge and understanding of both CAD and CFD are assumed, although most basic functions are covered.

Creating the Fluid Domain

In the majority of CAD models, the fluid domain for analysis must be extracted from the original part or assembly. The following sections demonstrate how to prepare the fluid domain for:

- an internal flow
- an external flow
- a multi-region flow

Internal Flow

In general, the steps to follow in creating the fluid domain for an internal flow are:

1. Work in a part. Import the original geometry into the new part (by making the new part derived), ensuring that all parts are maintained as separate solid bodies.
2. Close any openings by creating “caps”.
3. Create an overall block that encloses all the other geometry parts, the “enclosure”.

4. Subtract all parts from the enclosure and delete the outer solid body--the remaining geometry is the fluid domain.
5. Create a new assembly from the multi body derived part.

An example case is shown below. In this case the assembly contains two parts, the valve (shown in red) and the valve housing (shown in blue).



1. Create a new part file in Inventor. Name this new part `Internal Fluid Volume`. By default Inventor will open a sketch--close this sketch.

We will make this part file derived and import the assembly we want to extract the internal fluid volume from. During the import operation we can choose to import each part of the assembly as separate bodies. This is useful if you want to visualize the original components or perform a conjugate heat transfer analysis.

2. Click on the Derive button in the *Create* panel of the ribbon.

3. Locate and select the assembly or part file where you want to create the internal volume.

The *Derived Assembly* dialog appears.

4. In the *Derived Assembly* dialog, click on the (Maintain each body as a solid body) button. This will keep the solids in the assembly as separate bodies. The default option is to merge all solids into one solid body.



Notice that two bodies were imported into the part file under the Solid Bodies node in the feature tree.

5. Rename the imported bodies to something more representative--in this example the two bodies were renamed to `Valve Housing` and `Valve` respectively. Hovering the mouse pointer over the body nodes will highlights the corresponding body in the geometry window, this makes it easier to identify the bodies when renaming them.

We will now cap any openings.

6. Create a sketch on one of the faces containing an opening, as shown below.



By default, Inventor automatically converts any edges that lie on the sketch plane into sketch entities. These will appear in orange.

7. In this case we only require the outer edge, so click on the inner edge and delete it by pressing <`Delete`>. If the automatic projection option was turned off, you will need to manually project the outer edge of the face onto the sketch.



8. Create an extrusion from the sketch to cap the opening. The length of the extrusion is not important. In the Extrude dialog, click on the (New solid) button to make the cap a new body rather than merge it with the existing body it is in contact with.



A new body will be added to the Solid Bodies node in the feature tree; rename this to `Cap1`.

9. Using the same technique, cap any other openings on the geometry, and rename the new bodies in the Solid Bodies node in the feature tree.



We will now create a solid block that completely encloses the other parts. This block will be added as a separate solid body in the part file.

10. Create a sketch that completely encloses the assembly, as shown below.



11. Extrude the sketch so that it completely encloses the geometry. In the *Extrusion* dialog, click on the New Solid button, otherwise the geometry contained within the extrusion would be merged with the extrusion.



Note that a new solid is added to the Solid Bodies node in the feature tree, rename this to `Internal Fluid Volume`.

We will now subtract the original assembly and caps from the solid block. The result will be a solid block that contains a cavity separating the outer solid from the inner solid (the internal fluid volume).

12. Click on Combine in the *Modify* panel of the ribbon.



The *Combine* dialog will appear.

13. Click on the Cut button, and select the enclosure as the Base.
14. Now select the other bodies (the original assembly parts and the caps) as the Toolbody. A dashed outline will appear when the mouse pointer hovers over hidden parts, as shown below. Click to select these parts.



15. In the *Combine* dialog, tick the Keep Toolbody option. By default the toolbody parts are consumed by this operation, however we want to keep them.

We will now delete the outer solid. The remaining solid is the internal fluid volume.

16. Click on the Delete Face button in the *Surface* panel of the ribbon.



The *Delete Face* dialog will appear.

17. Click on the (Select lump or void) button, and select the solid body in the geometry window--this will delete the outer solid body and leave the internal fluid volume.

The internal fluid volume of the geometry is shown below.



The last step is to export the bodies from this multi body part into an assembly, where they can be selected for use in STAR-Inventor.

18. Save the current part file using a suitable name.
19. Make the original assembly components visible by right-clicking on these in the Solid Bodies node in the feature tree and selecting Visibility.



20. Click on the Manage tab of the ribbon.
21. Click on the Make Components button in the *Layout* panel of the ribbon.



The *Make Components : Selection* dialog appears.

22. Select the internal fluid volume and the original components from the geometry window. These will be added to the dialog.
23. Click Next.

The *Make Components : Bodies* dialog appears.

24. Give the new assembly (and parts) a suitable name and click OK.

The new assembly will be opened automatically.

25. Save the new assembly.

The derived part is associative, meaning that if the original part/assembly is modified, the derived part can be updated to reflect this. The up-to-date derived part can then be exported to an assembly, as described above, where it can be used with STAR-Inventor.

The geometry is now ready for the CAD client analysis. Suppressed parts in an assembly will not be included in the analysis. In the above example, it is recommended that you extend the inlet and outlet regions to appropriate lengths so that a fully developed flow can be achieved before reaching the valve and to help reduce recirculation at the outlet.

External Flow

The generation of a domain for external aerodynamics is generally easier than for internal flow cases but the sizing of the domain is more critical. As a guideline (assuming a specific situation is not being modelled i.e. flow in a wind tunnel) the domain should be sized as follows:

Inlet: 2 x part length from the front of the part

Outlet: 4 x part length from the rear of the part

Sides: 3 x part width from the side of the part

Ceiling: 4 x part height from the top of the part

In general, the steps to follow in creating the fluid domain for an external flow are:

1. Work in a part file. If the original geometry contains an assembly then create a new part and make it derived, loading the assembly into this part.
2. Create an overall block that encloses all the other parts. Ensure that this block is a separate solid body in the part.
3. Subtract all parts from the overall block part - the remaining geometry is the fluid domain.

An example case is shown below. In this example a geometry of a car will be used.



1. Create a new part file and make this a derived part.

Load the assembly into the part file, using the default import options. The assembly parts will be merged into one solid body.

2. If the geometry is symmetrical, as is the case with the car in this example, cut it in half to save on computational costs.



We will now create the fluid domain for the case.

3. Create a sketch on a suitable plane, in this example the symmetry plane of the car will do.
4. Draw the bounding box of the fluid domain using appropriate dimensions. In this case the sketch will represent the profile of the fluid domain, as shown below.



5. Extrude the sketch to the appropriate thickness for the domain.



The *Extrude* dialog appears.

6. click on the (New Body) button to create a new solid body from this extrusion--the default option would merge all solids together. Note that a new solid body is added to the Solid Bodies node in the feature tree.
7. Click on Combine in the *Modify* panel of the ribbon.



The *Combine* dialog appears.

8. Click on the (Cut) button, and select the domain as the Base, and the car as the Toolbody. A dashed outline will appear when the mouse pointer hovers over hidden parts, as shown below; click on this to select the hidden body.



The fluid domain is now complete. A close-up of the car in the fluid domain is shown below.



The geometry is now ready for the cad client analysis. Suppressed parts are not included in the analysis.

Multi-Region Flow

This section shows how to create the fluid domain for multiple regions such as rotating machinery and porous media flow.

For the numerical assumptions made using the Moving Reference Frame (MRF) technique to hold true, a minimum of two cells need to be present between the rotating-static interface and any wall boundaries (for example the fan blades), and the rotating region must be circular in section. Where there is a large clearance between the rotating object and any static objects, it is advisable to avoid placing the interface too close to either hence facilitating a small cell size on the interface. Where there is a small clearance between the rotating and static parts, there is little choice but to place the interface approximately midway between the two.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Create a new derived part, importing the original geometry into this part. Ensure that all parts are maintained as separate solid bodies.

2. Create a new body that encloses the rotating component, and subtract the rotating component from the solid. This is now the rotating region.
3. Create additional bodies that will form the inlet and outlet regions, and subtract any obstructions/features that may be present. An example would be the support arms of the fan motor in a fan casing.
4. Create a new assembly from the multi body derived part.



1. Create a new part file in Inventor. Name this new part `Fluid Domain`.

Make this part derived and import the assembly for which you want to create the fluid domain for.

2. Click on the Derive button in the *Create* panel of the ribbon.

3. Locate and select the assembly or part file you want to create the fluid domain for.
4. In the *Derived Assembly* dialog, click on the (Maintain each body as a solid body) button. This will keep the solids in the assembly as separate bodies. The default option is to merge all solids into one solid body.



Notice that two bodies were imported into the part file under the Solid Bodies node in the feature tree.

5. Rename the imported bodies to something more representative--in this example the two bodies were renamed to `Casing` and `Fan` respectively. Hovering the mouse pointer over the body nodes highlights the corresponding body in the geometry window, this makes it easier to identify the bodies when renaming them.

We will now create a solid body that will enclose the rotating component, this body will become the rotating region.

6. Hide the fan by right-clicking on its corresponding body sub-node in the Solid Bodies node and clear the checkbox next to Visibility.



To make a body visible again, right-click on it in the feature tree and tick the Visibility option.

7. Create a sketch on the circular face of the fan casing.



By default Inventor automatically adds any edges that lie on the sketch plane to the current sketch--these appear in orange. Delete this edge.

8. Draw a circle that shares the same center point as the fan casing. The edge of the circle should coincide with the innermost edge of the fan casing. To achieve this, select both the circle sketch and the innermost edge on the fan casing and click on the Equal button in the *Constrain* panel of the ribbon.



The result of this operation is shown below--the circle will snap to the selected edge.



9. Hide the fan casing and show the fan again.

10. Extrude the sketch until it completely encloses the fan.



The *Extrude* dialog appears.

11. Click on the (New Body) button to create a new solid body from this extrusion--the default option would merge all solids together.
12. Upon completing the operation, a new solid body will be added to the Solid Bodies node in the feature tree. Rename this body to `Rotating Region`.

We will now subtract the fan from the rotating region.

13. Click on Combine in the Modify panel of the ribbon.



The *Combine* dialog will appear.

14. Click on the (Cut) button, and select the rotating region as the Base, and the fan as the Toolbody. A dashed outline appears when the mouse pointer hovers over hidden parts, as shown below; click on this to select the hidden body.



15. In the *Combine* dialog, tick the Keep Toolbody checkbox to keep the toolbody. If this is not selected, the toolbody will be consumed by this operation. Keeping the bodies is useful if you want to import these into STAR-CCM+ for visualization purposes.

We will now create a new body that will extend upstream of the rotating region. This will be the inlet region.

16. Create a sketch on the front face of the rotating region.



By default Inventor automatically adds any edges that lie on the sketch plane to the current sketch; these appear in orange. Keep this sketch and exit the sketch. If this feature of Inventor was turned off, you can select the outermost edge on the body and click on the Project Geometry button on the *Draw* panel of the ribbon.



17. Extrude the sketch to a suitable length for the case.



In the *Extrude* dialog, click on the (New Body) button to create a new solid body from this extrusion--the default option would merge all solids together.

18. Note that a new solid body was added to the Solid Bodies node in the feature tree,> Rename this to `Inlet Region`.

We will now create a new body that will extend downstream of the rotating region. This will be the outlet region.

19. Create a sketch on the back face of the rotating region.



By default Inventor automatically adds any edges that lie on the sketch plane to the current sketch; these appear in orange. We only need the outer edge, so the inner edge can be deleted. If this option was turned off, you can manually project the outermost edge as previously described.

20. Extrude the sketch to a suitable length for the case.



The *Extrude* dialog appears.

21. Click on the (New solid) button to create a new solid body from this extrusion--the default option would merge all solids together.
22. Note that a new solid body was added to the Solid Bodies node in the feature tree, rename this to `Outlet Region`.

In this example, the fan casing contains support arms that hold the base of the motor. We need to subtracted these from the outlet region.

23. Hide the rotating region and the inlet region, leaving only the outlet region visible.
24. Make the fan casing body visible.
25. Click on Combine in the *Modify* panel of the ribbon.



The *Combine* dialog appears.

26. Click on the (Cut) button, select the outlet region as the Base, and the fan casing as the Toolbody.

27. In the *Combine* dialog, tick the Keep Toolbody checkbox to keep the toolbody. The toolbody will be hidden automatically after the operation completes.
28. Show the rotating region and the inlet region.

The multi-region fluid domain is complete. A cross-section of the fluid domain is shown below (note that the fan and fan casing were set to visible for this screenshot), where the fan is shown in red, the fan casing is shown in white, and the multi-region fluid domain is shown in grey.



The last step is to export the bodies from this multi body part into an assembly, where they can be selected for use in STAR-Inventor.

29. Save the current part file using a suitable name.
30. Make the original assembly components visible by right-clicking on these in the Solid Bodies node in the feature tree and selecting Visibility.
31. Click on the Manage tab of the ribbon.
32. Click on the Make Components button in the *Layout* panel of the ribbon.



The *Make Components : Selection* dialog appears.

33. Select the internal fluid volume and the original components from the geometry window. These will be added to the dialog.
34. Click Next.

The *Make Components : Bodies* dialog appears.

35. Give the new assembly (and parts) a suitable name and click OK.

The new assembly will be opened automatically.

36. Save the new assembly.

The derived part is associative, meaning that if the original part/assembly is modified, the derived part can be updated to reflect this. The up-to-date derived part can then be exported to an assembly, as described above, where it can be used with STAR-Inventor.

The geometry is now ready for cad client analysis. Suppressed parts will not be included in the analysis. The three regions that form the fluid domain all have a conformal match, so it will be possible to create the interfaces and obtain a topologically valid mesh in STAR-CCM+ without any further modifications to the geometry.

Imprinting Parts

Most CAD model assemblies are not suitable for meshing as the parts in the assembly are disjointed, where the coincident faces are overlapping. Attempting to generate a mesh with such an assembly will result in a meshing error. To rectify this problem, the parts in an assembly must be imprinted together.

Imprinting can also be used to split surfaces for boundary region applications. An example of such an application would be to model the area that a fan would occupy on the inside of computer case, such as in the third STAR-Inventor tutorial, where a smaller section of the surface is needed to represent the fan, rather than the entire inner surface of the computer case.

Imprinting Parts in an Assembly

In cases containing multiple adjacent parts, such as in Conjugate Heat Transfer cases, it is essential to imprint the profiles of the coincident faces of the smaller parts onto the face of the larger part. This will ensure that the volume mesh created by STAR-CCM+ contains interfaces between the adjacent parts.

In general, the steps to follow in creating the fluid domain for multiple regions are:

1. Work within the assembly, making the part that requires imprinting the work part.
2. Create a sketch on the face you want to imprint, and project the contact edges of the other components in the assembly onto this sketch.
3. Split the face you want to imprint using the sketch created in the previous step.

The simple circuit board shown below requires imprinting before the mesh can be generated.



In this example, the capacitors and heat sinks are not imprinted to the circuit board. This would result in a failure when attempting to generate a mesh.

To imprint the capacitors and heat sinks onto the circuit board follow the steps outlined below:

1. Load the assembly into Inventor.

2. Double-click on the node of the part that you want to imprint. This will make it the active part; all other parts in the assembly will become transparent.



3. Create a sketch on the face that requires imprinting.



We will now project the intersecting edges of the other components onto the active sketch.

4. Click on the Project Cut Edges button in the *Draw* panel of the ribbon.



5. Hover the mouse pointer over one of the transparent parts, this will become highlighted in white--click to select it.



The contact edges of the part are added to the sketch; these appear colored in blue.



6. Exit the sketch when done.

We will imprint the face with the sketch containing the contact edges of the other components.

7. Click on the Split button in the *Modify* panel of the ribbon.



The *Split* dialog appears.

8. Select the sketch containing the contact edges to act as the Split Tool.



Note that the sketch containing the profile will be consumed by this operation.

9. Select the face you want to imprint using the Faces to be Split tool.
10. Repeat the above steps (obtaining the profiles and splitting the face) for all the other components in the assembly.
11. Return to the assembly by double-clicking on the root node in the feature tree.

To verify that the face was correctly imprinted, we can hide the other parts by right-clicking on the imprinted part and selecting Isolate.



A high quality mesh with fully defined interfaces can now be successfully generated for this assembly. Suppressed parts are not included in the STAR-CCM+ analysis, so for CHT cases, such as the above example, ensure that all relevant parts are active.

Imprinting Surfaces on Parts for Boundary Region Applications

Faces in a part body can be split into smaller faces for boundary application.

The general procedure to follow is:

1. Create a sketch on the surface you want to split. The shape of the sketch should correspond to the shape of the new face after it is split.
2. Split the face that requires imprinting with the sketch containing the profiles.

A simple example of an aeration tank is shown below. The image shows the tank, the two air injectors need to be imprinted onto the base of the tank.



1. Load the part into Inventor.
2. Create a sketch on the face that requires imprinting. By default Inventor will automatically project edges that lie on the sketch plane--these will appear in orange. Delete these sketch entities.

3. Draw the profile of the faces you want to imprint.



In this example, the air injector profiles are circular in shape.

4. Click on the Split button in the *Modify* panel of the ribbon.



The *Split* dialog appears.

5. Select one of the profiles you created in the previous step to act as the Split Tool.



Note that the sketch containing the profile will be consumed by this operation.

6. Select the face you want to imprint using the Faces to be Split tool.
7. Repeat the above steps (creating the profiles and completing the split) for all the new boundaries you want to add.

The face is now imprinted; each face can be selected independently of each other, and can be used when defining boundaries.



The separate faces can now be defined as boundary regions.

Working With STAR-SpaceClaim

This part of the documentation details how to create, set up and run a CFD simulation in STAR-SpaceClaim. It assumes that you have previous experience in using SpaceClaim, and only describes the additional features available as part of STAR-SpaceClaim. The following sections are included:

- Introducing STAR-SpaceClaim
- Transferring Geometry to STAR-CCM+
- Understanding the STAR-SpaceClaim Process
- Defining Materials and Physics
- Setting Initial Conditions
- Setting Boundary Conditions
- Meshing the Model
- Obtaining a Solution
- Exporting Files and Data
- Resetting STAR-SpaceClaim Settings

You use STAR-SpaceClaim to transfer your CAD geometry directly into STAR-CCM+ without transferring the CAE settings. This method is an alternate way of importing geometry to using the STAR-CCM+ CAD translators.

Introducing STAR-SpaceClaim

STAR-SpaceClaim provides you with straightforward access to CD-adapco's STAR-CCM+ solver from within the SpaceClaim environment.

Standard physics models, including turbulence, ideal gas, heat transfer, and gravity are supported for single body cases. Multiple body simulations involving conjugate heat transfer, rotating reference frames, and porous media are also supported. Advanced physics models can be added to the simulation by either launching STAR-CCM+ and selecting them directly, or by implementing additional Java macro commands.

This introductory section contains the following:

- a list of capabilities of STAR-SpaceClaim
- a getting started guide
- a common functions section
- a geometry transfer section

What are the Capabilities?

STAR-SpaceClaim enables you to:

- Define flow characteristics for the simulation
- Specify mesh parameters and controls
- Execute the solver within STAR-CCM+, with specified solution controls
- Post-process results in STAR-CCM+
- Use parallel processing

Problem Definition

STAR-SpaceClaim provides an interface for the specification of continua, regions and boundaries in order to define the CFD simulation. This information is then used by STAR-CCM+ to calculate a solution and present results.

The particular solver options that you can currently select through the STAR-SpaceClaim interface include:

- Turbulence
  - k-ε Model
  - Laminar
- Temperature
  - Thermal Modelling
  - Conjugate Heat Transfer
- Equation of State
  - Constant Density

    - Ideal Gas
- Porous Media
    - Isotropic
    - Uni-Directional
- Moving Reference Frames
- Energy Source
- Gravity
- Passive Scalar

Meshing

STAR-SpaceClaim provides you with a transparent interface to CD-adapco's surface and volume meshing tools. This allows you to generate high-quality meshes using a solid model that was built within the SpaceClaim environment. The specific meshing functionality includes:

- Surface meshing using the surface remesher for optimum mesh quality
- Prismatic subsurface meshing in the volume mesh for capturing near-wall boundary layers
- Volume meshing of parts with either:
    - Tetrahedral core mesh generation
    - Polyhedral core mesh generation
    - Trimmed hexahedral core mesh generation
- Automatic handling of region-to-region interfaces

Running the Solver

STAR-SpaceClaim allows the generation of solutions with varying accuracy and speed by providing the following solver controls:

- Number of iterations
- Stopping criterion at which to assume convergence
- Discretization scheme (1st or 2nd order)
- Solver under-relaxation factors
- Additional commands using Java macros

Post-processing

As post-processing is carried out in STAR-CCM+, a wide range of tools are available for analysis. These include:

- Scalar scenes
- Vector scenes
- Streamlines
- Isosurfaces
- Plane sections
- Reports
- Plots

Parallel Processing

STAR-SpaceClaim supports parallel processing for mesh generation and solver calculations. You can set up parallel processing in the Preferences dialog. The different STAR-CCM+ licensing options can also be selected here.

Getting Started

The following section provides information on familiarization with the STAR-SpaceClaim interface, and basic concepts. Topics include:

- Launching STAR-SpaceClaim
- The STAR-SpaceClaim Info App Window
- Modifying the config.xml File
- Understanding the STAR-SpaceClaim Interface
- The STAR-SpaceClaim Object Tree
- Setting Up Parallel Processing and Licensing
- Changing the Unit System
- Transferring the CAE Model to STAR-CCM+
- Editing STAR-SpaceClaim Settings While the STAR-CCM+ Client is Running

Launching STAR-SpaceClaim Manually

By default STAR-SpaceClaim should start automatically when you launch SpaceClaim.

If this does not happen, you can manually start STAR-SpaceClaim by following the steps outlined below:

1. Start SpaceClaim as usual.

2. Click File > SpaceClaim Options.

   The *SpaceClaim Options* dialog appears.

3. Click the Add-Ins tab.
4. Activate the STAR-SpaceClaim option.



5. Click OK.

When you load or create a geometry file, the STAR-SpaceClaim ribbon appears across the top of the screen, and the STAR- Tree tab appears on the left hand side.

The STAR-SpaceClaim Info App Window

The *STAR-SpaceClaim Info App* window automatically appears when you start SpaceClaim. Initially, this window only contains the Status tab which displays the version and licensing information for STAR-SpaceClaim. More tabs appear as you create or load geometry files in STAR-SpaceClaim, allowing you to perform more actions.



You cannot close the *STAR-SpaceClaim Info App* window while SpaceClaim or STAR-SpaceClaim are running—attempting to close the window by clicking the X button returns a message in the Status tab informing you of this. You can maximize and minimize the window by clicking the - and _ buttons respectively.

Use the Clear button to clear all the information that is displayed in a particular tab.

Use the Abort button to stop active processes, such as surface/volume meshing, or the solver. This button only appears when you can stop a process—it is grayed out at all other times.

A clock is displayed in the bottom-left of the window. This clock displays the current time on your machine. A progress bar is located next to the clock and reports the progress of various processes, such as mesh generation and the solver.

More information on the various tabs is provided in the following sections:

- The Status tab
- The Output tab
- The Model Checker tab
- The Residuals tab

The Status Tab

The Status tab displays the licensing information for your current session, informing you whether the STAR-SpaceClaim license checked out successfully. The version of STAR-SpaceClaim is also displayed in this tab.



The Output Tab

The Output tab appears once you create or open a geometry file in STAR-SpaceClaim.



During the meshing process, this tab displays status information along with a brief report on the generated mesh. If errors were encountered during the meshing process, they are reported here with suggestions as to what can be causing them.

When the solver is running, the Output tab behaves like the output window in STAR-CCM+, displaying the data for each residual per iteration. When the solver finishes running, the same information is also displayed in the STAR-CCM+ *Output* window.

You can control the amount of information that is displayed in the Output tab by editing the verbosity in the `config.xml` file.

To abort the mesher or solver, click the Abort button.

The Model Checker Tab

The Model Checker tab displays information from the live model checker, where the case setup is actively monitored for any problems. The following example reports that the case has no inlet boundary defined.



The Residuals Tab

The Residuals tab appears when you launch the solver. Use the residuals to monitor the convergence of the solution being calculated.



Typically, a drop in three to five orders of magnitude in the residuals indicates that the solution has converged to an acceptable level. The residuals that are displayed in this tab also appear in STAR-CCM+ under the Plots node.

You can zoom in to the residuals by clicking and holding the left mouse button in the graph area and drawing a box around the area you want to view. To zoom out, right-click in the graph area. You can hide individual residuals from the Residuals tab by clicking their label at the top of the plot window.

Modifying the config.xml File

STAR-CCM+ requires a configuration file to enable it to find the appropriate CD-adapco components to run correctly. This configuration file is created during installation. An explanation of the file and settings within it is given here.

The configuration file is called `config.xml` and the STAR-SpaceClaim searches for it in the location given by the `STARCADSERIES` environment variable.

Typical `config.xml` content is shown below. In this example STAR-CCM+ version 9.04.004 is shown.

```
<CAD>

  <STAR_HOME>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004</STAR_HOME>


<SOLVER>C:\Program Files\CD-adapco\STAR-CCM+9.04.004\star
\bin\starccm+.exe</SOLVER>

  <LOGFILE>starcad.log</LOGFILE>

  <WWHELP>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\win64\Intel11.1\wwhelp.
exe</WWHELP>

  <MANDIR>C:\Program
Files\CD-adapco\STAR-CCM+9.04.004\doc\en\online</MANDIR>

 <VERBOSITY>Normal</VERBOSITY> <!-- Can be Verbose, Log or
Normal. Default Normal -->

</CAD>
```

The following lines can be edited:

- `<SOLVER>` This provides the location of the STAR-CCM+ solver executable. STAR-CCM+ licensing options are set in the Preferences dialog in the STAR-SpaceClaim user interface.
- `<LOGFILE>` is the logfile that will automatically be created and saved when STAR-SpaceClaim is run. A path of your choice can be used.
- `<VERBOSITY>` is the verbosity of the output window in the STAR-SpaceClaim Info App. This can be set to one of the following:
  - `Normal` — This is the default setting and will display the least amount of information in the Output panel.
  - `Verbose` — This will display a moderate amount of information in the Output panel.
  - `Log` — This will display all available information in the Output panel.

Understanding the STAR-SpaceClaim Interface

The STAR-SpaceClaim user interface is embedded within the SpaceClaim environment, with the addition of the STAR-SpaceClaim feature tree and ribbon.



Certain actions can also be carried out using multiple methods. A summary of the features available in each part of the STAR-SpaceClaim interface are detailed below.

Object Tree

The operations provided in the object tree include:

- Solution Controls
- Definition of Units System
- Continuum settings
- Region settings
- Add new boundary
- Boundary conditions
- Mesh Pipeline settings

- Add new mesh controls

Ribbon

- Settings
  - Definition of Units System
  - Solution Controls
  - Clear All
- Preferences
- Mesh
  - Generate Surface Mesh
  - Generate Volume Mesh
  - Clear Generated Meshes
  - Show/hide Mesh
- Solution
  - Build and Run
  - Run Solver
  - Run For More Iterations
  - Clear Solution
- Model
  - Check Model For Errors
  - Transfer CAE Model
- Export
  - Surface Mesh (.dbs)
  - Volume Mesh (.ccm)
  - Engineering Report (.html, .csv)
  - Simulation (.sim)
- Geometry
  - Transfer Geometry
- STAR-SpaceClaim Help
- About STAR-SpaceClaim

The STAR-SpaceClaim Object Tree

STAR-SpaceClaim uses an object tree as the primary method for setting parameters. The tree consists of nodes, each one of which represents a particular component of the simulation. This approach to the user interface complements the method used by the STAR-CCM+ package, allowing for an easy transition between the two.



The object tree appears slightly different in assembly mode as it must accommodate multiple parts, but the principles of operation remain the same.

Nodes are arranged in an order that follows the recommended workflow for setting up a simulation in STAR-SpaceClaim. At each stage, you typically right-click on a particular node in the tree and choose the appropriate action from the pop-up menu.

Certain functionality within STAR-SpaceClaim applies to the simulation as a whole, rather than to individual objects. These operations are provided in either the menu or toolbar.

A complete list of all functions and their respective location is included in the Understanding the STAR-SpaceClaim Interface section.

Clicking on the Region node within the object tree will highlight the relevant bodies on the geometry model for that particular region. If different parts are assigned to multiple regions, then clicking on a region will only highlight the bodies that are associated with that specific region. For example, clicking on the Static Region node below will highlight the body that is associated with the static region.


STAR Tree
STAR-SpaceClaim
Continua
Continuum
Regions
Static Region
Boundaries
Default
Rotating Region
Boundaries
Default
Mesh Pipelines
Mesh Pipeline
Mesh Controls
Structure
Layers
Selection
Groups
Views
STAR Tree



Similarly, clicking on the Rotating Region node highlights the body that is associated with the rotating region.


STAR Tree
STAR-SpaceClaim
Continua
Continuum
Regions
Static Region
Boundaries
Default
Rotating Region
Boundaries
Default
Mesh Pipelines
Mesh Pipeline
Mesh Controls
Structure
Layers
Selection
Groups
Views
STAR Tree



Clicking on a boundary node in the STAR-SpaceClaim object tree highlights all the faces on the CAD geometry model that are associated with it. Note that clicking on the *Default* boundary node will not highlight any faces on the CAD geometry model.To stop highlighting the faces, click on a blank space in the geometry window.

Setting Up Parallel Processing and Licensing

The licensing options and the options for setting up parallel processing are located in the *Preferences* dialog. To access the *Preferences* dialog:

1. Select Preferences from the ribbon.



The *Preferences* dialog appears.

2. Select the license type that you want:
   - Power Session—The Power Session uses the Power Session license feature, which is designed for parallel simulations. You can use an unlimited number of cores in this session.
   - Lite Session—If you have purchased STAR-CCM+ Lite, you should activate this option when launching a simulation. A limited set of features are provided when you activate this option.

Note: If you select no licensing option, the default CAD client license is used.

3. To run STAR-SpaceClaim in parallel, select Parallel from the *Run Mode*.



4. Set the number of cores that you want to use in the Compute Processes field. You can use the up/down buttons or you can type in the number of cores that you want to use.
5. Click OK to close the dialog.

If the STAR-CCM+ server is running when you modify the preferences, a message box appears informing you that both STAR-SpaceClaim and STAR-CCM+ need to be closed for the changes to take effect.



6. Click OK to proceed.
7. Close the active STAR-SpaceClaim and STAR-CCM+ sessions.
8. Restart STAR-SpaceClaim.

Note: The STAR-CCM+ server only starts when it is required, such as during surface meshing or volume meshing. It does not start when you launch STAR-SpaceClaim.

The settings in the *Preferences* panel are stored and remembered the next time you start STAR-SpaceClaim.

If an option in the *Preferences* panel conflicts with an option in the config.xml file, the option set in the *Preferences* panel will take precedence. An example of such a situation is if the Compute Processes property is set to `4` in the *Preferences* panel, whilst it is set to `2` in the `config.xml` file; the simulation will run using four processors.

Note: MPICH2 or IBM Platform MPI needs to be installed on your local machine for you to run in parallel. These programs are included with the STAR-CCM+ package installer.

Changing the Units System

You can use either the SI (International Standard) or USCS (United States Customary System) units system in STAR-SpaceClaim.

To change the units system:

1. Bring up the *Units System* dialog:

- Right-click the STAR-SpaceClaim node within the object tree and select Set Units System.



- Alternatively, from the ribbon, select Units System.



The *Units System* dialog appears.

2. Select the desired system from the drop-down menu.



3. Click OK to close the dialog.

Automatic Unit Conversion in Dialog Panels

Changing the default units system converts existing values to the new units system. Ensure that you select the desired units system *before* you enter the value, otherwise you will need to re-enter it.

Note: Specific values that you have already modified the units on are not affected by a change of units system.

Transferring Units to STAR-CCM+

Units that are used in the CAD geometry file are also used in the STAR-CCM+ simulation. Less common units, such as those listed below, are not supported in STAR-CCM+:

- Micrometers
- Angstroms
- Nanometers
- Mils
- Microinches
- Footinches (this unit may be converted to inches in STAR-CCM+)
- Miles

Note: This list may not be exhaustive and is subject to change.

These units are converted to meters when the CAD geometry is imported into STAR-CCM+.

Transferring the CAE Model to STAR-CCM+

You can transfer your geometry and the simulation setup to STAR-CCM+ at any time from within the STAR-SpaceClaim environment. If you want to transfer only the geometry to STAR-CCM+, refer to Transferring Geometry to STAR-CCM+.

To transfer the geometry with the simulation set up:

- Click the Transfer CAE Model button in the ribbon.



The STAR-CCM+ version affiliated with STAR-SpaceClaim launches.

Note: A surface mesh is required in order to visualize the geometry in STAR-CCM+. If you have not generated a mesh before clicking Transfer CAE Model, the geometry is tessellated with a medium density.

Editing STAR-SpaceClaim Settings While the STAR-CCM+ Client is Running

If you make a change in STAR-SpaceClaim while the STAR-CCM+ client is running, such as adding or deleting a boundary, you may need to update the volume mesh. Scenes in the STAR-CCM+ session will not display the model correctly unless you keep the volume mesh up to date. A list of changes that require you to re-mesh the geometry is provided below:

- For boundaries and regions: changing the type, adding or deleting
- For mesh pipelines and mesh controls: editing, adding or deleting

When adding or removing regions and boundaries in STAR-SpaceClaim, the scenes in STAR-CCM+ may be missing some parts, as shown below, even after re-meshing the geometry.



To update the scene, you need to manually add the modified regions and boundaries to the scene in the STAR-CCM+ client. You can do this by dragging and dropping the region or part from the STAR-CCM+ object tree,

or by adding the region or part using the Part Selector. An alternative method would be to create a new scene, as this will automatically pick up the new regions and boundaries added to the simulation.



Common Functions

The following section details common operations used within the STAR-SpaceClaim environment. These are listed below:

- Renaming an object
- Deleting an object
- Editing an object
- Assigning surfaces and parts to objects in STAR-SpaceClaim
- Removing surfaces and parts from objects in STAR-SpaceClaim

Renaming an Object

To rename an object:

1. Right-click the object node within the object tree and select Rename.



2. Enter the name.



3. Click on a blank space in the object tree.

Note: You can also assign entity names at the point where you create them, for example, when adding boundaries, continua, regions, mesh pipelines and mesh controls. When adding a new object, locate the Name text box in the settings dialog and enter the required name.

Deleting an Object

To delete an object:

- Right-click the object node within the object tree and select Delete.



To delete multiple entities:

- Select the desired nodes within the object tree (by holding down the `<Ctrl>` key), and right-click on any one of the selected nodes, then select Delete from the pop-up menu.

Note: Some combinations of selections cannot be deleted together using this method. In these cases, the Delete option is disabled in the pop-up menu.

Editing an Object

To edit an object:

1. Right-click on the desired node in the object tree and select Edit...



The settings dialog for the object appears.

2. Make the changes that you want.
3. Click OK to close the dialog.

Assigning Surfaces and Parts to Objects in STAR-SpaceClaim

Regions, boundaries, mesh pipelines, and mesh controls are all represented as objects on the STAR-SpaceClaim object tree. You can assign parts and surfaces to these objects when their property dialog is open.

Assigning Surfaces to Objects

To assign a surface to an object (when its property dialog is open):

- On the CAD geometry, click the surfaces that you want to assign to the object. You can select multiple faces by holding down the `<Ctrl>` key whilst clicking the surfaces.

The selected surfaces are highlighted on the CAD geometry.



The selected surfaces are also displayed in the *Face Selection* box of the active dialog.

Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Assigning Parts to Objects

To assign a part to an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to assign to the object. You can select multiple parts by holding down the `<Ctrl>` key whilst clicking the parts.

The selected part is highlighted on the CAD geometry.



The selected parts are displayed in the *Part Selection* box of the active dialog.



Parts that are assigned to other regions are automatically hidden on the CAD geometry.

Removing Surfaces and Parts from Objects in STAR-SpaceClaim

Regions, boundaries, mesh pipelines and mesh controls are all represented as objects on the STAR-SpaceClaim object tree. You can remove parts and surfaces from these objects when their property dialog is open. A surface that is removed from a boundary is returned to the default boundary. Similarly, removing a part from a region returns it to the default region.

Removing Surfaces from Objects

To remove a surface from an object (when its property dialog is open):

- On the CAD geometry, click the face that you want to remove.

Removed surfaces are no longer highlighted on the CAD geometry. The surface is removed from the *Face Selection* box.

Removing Parts from Objects

To remove a part from an object (when its property dialog is open):

- On the CAD geometry, click the part that you want to remove.

Removed parts are no longer highlighted on the CAD geometry. The part is removed from the *Part Selection* box.

Transferring Geometry to STAR-CCM+

With STAR-SpaceClaim, you can transfer your CAD geometry directly to STAR-CCM+ from within the SpaceClaim environment. The Transfer Geometry feature imports your geometry into STAR-CCM+. This function provides an alternative method to the CAD translator add-ons for importing geometry. This method of transferring geometry can provide a faster workflow when performing a design analysis.

Note: The Transfer Geometry feature of STAR-SpaceClaim requires you to install and license the SpaceClaim Data Exchange II module.

To transfer the geometry to STAR-CCM+:

1. Click Transfer Geometry on the STAR-CCM+ ribbon.



STAR-CCM+ launches and imports the CAD geometry. You are now at the same stage in the workflow as if you imported the surface mesh directly into STAR-CCM+ using the CAD translators (*File* > *Import* > *Import Surface Mesh...*).

2. Set up your case in STAR-CCM+.

Caution: Do not click Transfer CAE Model if you have previously used the Transfer Geometry method and set up your case in STAR-CCM+. This action results in the loss of your settings in STAR-CCM+.

Saving Simulation Files

To save your simulation files:

- Save your CAD geometry in STAR-SpaceClaim. This action ensures that the simulation file remains associated with the SpaceClaim file. You can also safely rename and move the file using this method.
- Alternatively, save your simulation file in STAR-CCM+. Do not rename the simulation file when using this method, as it will lose the associativity with the SpaceClaim file.

Resuming a Simulation

To resume a simulation after closing STAR-SpaceClaim and STAR-CCM+:

1. Launch STAR-SpaceClaim.
2. Open the CAD geometry in STAR-SpaceClaim.
3. Click Transfer Geometry.

Transferring Geometry with Design Parameters

You can transfer CAD features as design parameters to STAR-CCM+ along with the geometry. Design parameters control the feature dimensions, for example extrusion depth, and are useful for modifying the geometry in STAR-CCM+ after running the analysis.

To transfer design parameters to STAR-CCM+:

1. Click Select Parameters on the STAR-SpaceClaim ribbon.



The *Parameter List* dialog appears.



2. Select the design parameters that you want to transfer from the list. You can only select a single parameter at a time. The value of the parameter is displayed in the Value option.

Note: STAR-SpaceClaim only recognizes Driving Dimensions as design parameters. Only recognized design parameters are displayed in the list.

3. Activate the Select/Deselect option.
4. Click OK.
5. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

The geometry, along with the specified design parameters, are transferred to STAR-CCM+.

Modifying Geometry from Within STAR-CCM+

Once you transfer a design parameter to STAR-CCM+, you can modify a geometry from within STAR-CCM+ with respect to that parameter. This feature is useful for quickly modifying a geometry after performing an analysis in STAR-CCM+ without needing to return to the original CAD software. In the following example, the geometry of a diffuser is modified.



To modify a geometry in STAR-CCM+:

1. In STAR-SpaceClaim, select the design parameters that you want to transfer to STAR-CCM+.
2. Transfer the geometry to STAR-CCM+. See Transferring Geometry to STAR-CCM+.

   Transferred geometries are located under the Geometry > Parts node.



Transferred design parameters are located under the Geometry > CAD Client Models > [part_name] > Design Parameters node.



3. In STAR-SpaceClaim, save the geometry file. The simulation file is also saved.
4. In STAR-CCM+, select the design parameter that you want to modify and set Parameter Value to the desire value.



5. Update the part in STAR-CCM+:
   - To update the part and save the updated parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model and Save the CAD File.
   - To update the part without updating the parameters in the CAD file, right-click Geometry > CAD Client Models > [part_name] and select Update Model.

   The geometry updates according to the design parameter changes.



6. Return to SpaceClaim. The changes to the geometry are transferred back to the native CAD geometry.

This workflow maintains an active link between STAR-SpaceClaim and STAR-CCM+. When working with Geometry Transfer, you can modify the geometry in either package and the geometry in the relevant package is updated accordingly.

Caution: Do not modify the geometry in STAR-CCM+ if you have closed the geometry part but have left SpaceClaim running. Currently, doing so can cause issues.

Controlling How the Geometry is Transferred to STAR-CCM+

You can control how the geometry is transferred from SpaceClaim to STAR-CCM+ through the *Geometry Transfer Settings* dialog.

To access the *Geometry Transfer Settings* dialog:

1. Right-click the STAR-SpaceClaim root node and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.



2. Set the transfer parameters. See Geometry Transfer Settings Dialog.
3. Click OK.

Geometry Transfer Settings Dialog

Use the following properties to control how the geometry is transferred to STAR-CCM+.

Coincident Faces Detection Tolerance	Specifies the tolerance, in meters, for geometry in adjacent faces to be considered coincident. The default value is `1.0E-3`.
Create Part Contacts from Coincident Faces	Controls whether the geometry transfer feature searches for pairs of coincident faces on the CAD faces. For each pair found, part contact data is created and stored in the Contacts property (in STAR-CCM+). See the Creating Part Contacts from Coincident Entities section in the STAR-CCM+ user guide.
Transfer Multi-Body Parts as Composites	Specifies whether to transfer a single part containing multiple solid bodies as a composite part. Part contacts are created for each pair of coincident faces. Leaving this option deactivated merges all solid bodies in the part.

Understanding the STAR-SpaceClaim Process

In STAR-SpaceClaim the typical workflow for running a simulation is as follows:

1. Ensure that the CAD model is suitable for analysis
2. Extract the fluid volume
3. Define the material properties for each part of the model, for example, air, water, aluminum
4. Define model physics, for example, turbulence, temperature, gravity
5. Specify boundary conditions, for example, inlets, outlets, symmetry planes
6. Specify mesh settings

7. Define solution controls
8. Run the case
9. Post-process the results in STAR-CCM+

Once you have obtained an initial solution you may wish to re-run the case with some different settings, or with some changes to the geometry. To do this you can:

10. Return to STAR-SpaceClaim and change settings or modify the geometry
11. Re-run the case

Using Suitable CAD Models in STAR-SpaceClaim

STAR-SpaceClaim only supports the use of solid bodies. If a geometry is made up entirely of surfaces then it will not be possible to define any boundary conditions, and the meshing options will be greyed out.

It is good practice to create copies of the original geometry file(s) before importing them into STAR-SpaceClaim.

Suppressed or deactivated parts and features are not supported by STAR-SpaceClaim and will have no effect on the analysis. Hidden parts, however, are used in the analysis. Ensure that all desired parts and features are active before using them with STAR-SpaceClaim.

Extracting the Fluid Volume

A CAD model usually represents a component made from solid material. However, in the case of flow simulation, a model representing the fluid volume is needed, requiring a different CAD model to be created. This is illustrated below:

Initial CAD model



Extracted fluid volume



In some instances, a simulation will require multiple regions, such as:

- Moving Reference Frames

  For rotating flow, such as turbines.
- Porous Media

  For representing regions with fine geometric structures
- Energy Source

  For simulating thermal specifications, such as conjugate heat transfer

These types of simulation require each region to be represented by individual contiguous parts within an assembly.

Defining Material Properties

Continua determines the characteristics of the fluid and solid entities used within the simulation. The behavior of a single continuum is primarily defined by the chosen material type and physics models.

Specification of the material type used in the simulation provides properties such as density, viscosity and conductivity, which are all used by the solver to calculate a solution. STAR-SpaceClaim includes a database of common materials and allows you to modify material properties where necessary.

Defining Model Physics

Physics models determine the nature of the fluid, and can be described as either:

- Turbulent or Laminar
- Constant Density or Ideal Gas

Additional models can also be activated to meet the needs of various simulation requirements:

- Passive scalar
- Temperature
- Gravity

A region is defined as a volume domain in a simulation, surrounded by boundaries. Each region in the simulation must be represented by a part in the native CAD package. When multiple regions are required, multiple parts must be used within an assembly.

All regions must be assigned to a continuum, as they represent a volume of material in the simulation. Region settings can be used to define porous media, moving reference frames, and an energy source.

Specifying Boundary Conditions

To define boundary conditions you first need to select the surface or surfaces on the model to assign to a separate boundary.

A boundary is defined as a surface, which along with other boundaries, surround and characterize a region. Boundary conditions determine how the boundary will affect the material within a region, and consequently, the simulation as a whole.

Inlet boundaries determine how much fluid enters the region, using a velocity, mass flow rate or stagnation inlet. Outlet boundaries define the process for fluid leaving the region, which are defined with either a pressure or flow-split ratio. Wall boundaries contain the fluid, and can also be used to define a thermal specification that will influence the transfer of energy to or from the region. Symmetry boundaries can be used to avoid modeling the entire region where appropriate, and so reduce calculation time.

As boundary types can affect the way in which the volume mesh is constructed, it is important to set them before generating the mesh.

Specifying Mesh Settings

The STAR-CCM+ solver uses the finite volume method to obtain an approximate solution to the underlying equations of fluid flow. This requires a mesh to be created, which represents the fluid volume as a collection of much smaller volumes known as cells.

A surface mesh is first created based on the boundaries defining the regions in the simulation. Thereafter a volume mesh is created to discretize the entire model.

The accuracy of the results is directly affected by mesh quality. If it is necessary to improve the quality of the mesh, some of the ways in which this can be done are to:

- Increase the mesh density, that is, the number of cells
- Add further sub-surface prism layers to capture near wall flow
- Refine the mesh in areas where flow characteristics are expected to change rapidly

As volume mesh quality is largely governed by the surface mesh quality, it is common practice to create and inspect the surface mesh before creating a volume mesh.

Defining Solution Controls

Solution controls determine the level of accuracy the solver will work towards. As a high accuracy solution will take longer, it may be necessary to compromise on this in order to obtain a result in a shorter time.

STAR-SpaceClaim includes five predefined solver settings ranging from fast (with low accuracy) to slow (with high accuracy). Advanced options are also available to provide further control.

Running the Case

When the simulation has been fully defined, the volume mesh is sent to STAR-CCM+ along with all the necessary data for running the simulation. The time to complete a simulation will vary depending on the mesh density and solution controls.

As the solver is running, residuals showing the rate at which individual equations are converging are shown in a plot. A scalar scene displaying absolute pressure is displayed on completion of the simulation.

Post-Processing the Results in STAR-CCM+

Post-processing of simulations is carried out using the STAR-CCM+ interface. This allows the results to be taken from the solver and displayed in a variety of scenes and plots.

Making Modifications to the Settings or Geometry

As STAR-SpaceClaim is completely integrated within the SpaceClaim environment, it is very simple to perform parametric studies that examine the response of the fluid to changes in geometry. Once the first analysis is complete, the boundary conditions, run settings and all other properties of the CFD model are stored within the CAD file. Any changes in the flow field due to design changes can be recalculated by generating a new mesh and running the solver.

Defining Materials and Physics

Defining the physical properties of a STAR-SpaceClaim model involves specifying the materials, flow characteristics, operating conditions and physics models. These definitions and settings are accomplished through the use of continua and regions, which are described in the following section, Working with Continua and Regions.

Further sections describe how to define materials, continuum physics and region physics:

- Defining materials

  Including material type and properties.

- Defining continuum physics

  Including laminar or turbulent flow, density model, temperature, passive scalars and gravity.

- Defining region physics

  Including porosity, moving frame of reference, or energy source.

This part of the User Guide also includes details for setting reference properties for the model.

Working with Continua and Regions

Continua and regions are used together to capture the physical and material properties of a model. A continuum represents a material in combination with a set of physics models that will be used to simulate the response of the material to the conditions imposed upon it. A continuum also provides the facility to define reference and initial values.

A region represents one or more volumes within the overall model, not necessarily contiguous, which will share the same continuum and set of boundary conditions. Regions are defined using CAD parts, and will be discretized by the volume mesher.



As well as holding a set of boundary conditions, regions can also be used to model specific physical behavior, including porous resistance, moving reference frames and energy sources.

Parts can be imported into STAR-CCM+ for visualization purposes without affecting the analysis.

Continua are assigned to single or multiple regions, which can contain single or multiple parts. In part mode, only one continuum and one region are available, as there is only one part. In assembly mode, new continua can be added to represent different materials present in the simulation, and new regions can be added to represent the parts composed of these materials. Continua and regions can be deleted if they are not required, although the simulation must always contain at least one of each.

There is a well defined relationship between continua, regions, boundary conditions and CAD parts.

- Every region must be associated with a continuum.
- Several regions can be associated with the same continuum.
- Every CAD part must belong to a region.
- Several CAD parts can be associated with the same region.
- Boundary conditions belong to regions, and cannot be shared between regions.
- Every CAD surface must have a boundary condition associated with it. A default wall boundary is applied to CAD surfaces that have not been specifically associated with a user-defined boundary condition.

The following series of illustrations depict the relationship between CAD parts, continua, and regions, for a circular steel duct containing a rotating fan, and surrounded by a layer of insulation.

- Initially we have a collection of CAD parts in an assembly. These must be imprinted where appropriate to ensure that interfaces between parts can be created correctly.



Assembly of CAD Parts (imprinted where necessary)

- Three continua are required to represent the three different materials present in the simulation: the fluid (for example, air), the steel pipe, and the insulation.



View of Continua Required for Simulation

- Two regions are required for the fluid portion of the simulation, as a rotating part is defined by applying a rotating frame of reference to a separate region surrounding the fan blades. Further regions are required for each of the solid parts (the steel pipe and insulating layer), which means a total of four separate regions overall.



View of Regions Required for Simulation

- A clearer picture can be given of the fan blade itself.



For an analysis containing a single part, for example a fluid-only case involving fluid flow through a converging-diverging pipe, the model will only require one continuum and one region in order to fully define the physical and material properties of the model. In this case, STAR-SpaceClaim automatically assigns the default continuum to the default region, and no further assignments are necessary.

The following sections describe how to create and use additional continua and regions (assemblies only):

- Creating additional continua
- Creating additional regions
- Assigning a continuum to a region

Creating Additional Continua

To define different materials and associated physics for use in an assembly, additional continua are required. In each continuum you can specify a unique set of properties, such as a different material type or temperature. In assembly mode, continua are stored under a Continua manager node, as shown below.



To create an additional continuum:

- Right-click on the Continua manager node in the object tree and select Add Continuum.



This will open the *Continuum Settings* dialog, in which you can define the properties of the new continuum and also rename it.

The continuum will be created when you click OK, or you can select Cancel to exit the dialog without creating the continuum. Once a continuum has been created, you can assign it to a region or regions, as necessary.

Creating Additional Regions

Additional regions are required when parts must belong to different continua, or when specific behavior such as porous resistance is required for a particular part within the same continuum. Parts from an assembly are assigned to regions, the properties of which are controlled by the assigned continuum and the region physics.

In assembly mode, regions are stored under the Regions manager node in the object tree, as shown below.



- Right-click on the Regions manager node in the object tree and select Add Region.



This will open the *Region Settings* dialog in which you can assign the region parts and a continuum, define the region physics, and also rename the region.



The parts that are assigned to a region define the geometric extent of the region. Within this geometry, properties and behavior of the material is governed by the physics models and material properties defined in the region and continuum settings. Each region must be assigned at least one part, otherwise it will be ignored during the CFD analysis.

By default, all the parts in an assembly are assigned to the default region (this is the initial region created by STAR-SpaceClaim). When you assign a part to another region, it is automatically removed from the default region. If you remove a part from a region, it is automatically reassigned back to the default region.

Click OK to create the new region. New regions that you create are added to the object tree under the Regions manager node.

Assigning a Continuum to a Region

Assigning a continuum to a region causes the region to inherit the physical and material properties defined in the continuum. By default, every STAR-SpaceClaim model contains one continuum, and this is automatically assigned to the default region and each additional region that is created. If multiple continua are defined in the model, you can choose a different continuum to assign to each region. This enables you to control the physical properties of each region, including varying material types.

To assign a continuum to a region, first create and define the additional continua as necessary. In the *Region Settings* dialog, select the alternative continuum from the *Continuum Selection* box and click OK to confirm the change. Continuum assignment can also be performed in the *Region Settings* dialog when you are creating a new region. In the following example, two continua were defined in the model: Fluid and Solid.



Defining Materials

In STAR-SpaceClaim each continuum in the model is assigned a material. This defines the material properties used by the CFD solver when the case is run. The choice of material will affect the options available for the model physics. For example if you select a material from the Gases or Liquids lists, you will be able to specify whether the flow is turbulent or laminar. By default the material type is set to Air, and this can be changed to any one of a number of predefined solids, liquids or gases that are available in STAR-SpaceClaim. If necessary, it is possible to define an entirely different material by entering different material properties.

Only one material can be selected per continuum, so for a model containing a number of materials you will need to create additional continua. This is only possible for assemblies.

To define the material properties of a continuum, edit the continuum by right-clicking on the Continuum node in the STAR-SpaceClaim object tree and select Edit.



This opens the *Continuum Settings* dialog, in which you can define the material as follows:

- In the *Material* box, you can select a material type.
- In the *Material Properties* box, you can define each of the material properties.



Selecting Material Type

Every continuum is assigned a material, and the default material type is Air. STAR-SpaceClaim offers an extensive library of material types for you to choose from, categorized under gases, liquids and solids. To use a material other than those listed, select a material that is the closest match and alter the material properties.

To select a material type, open the *Continuum Settings* dialog. In the *Materials* box, select either the Gas, Liquid or Solid radio button, then select the material from the drop-down menu. When you have selected a material, the options available in the *Material Properties* box will automatically display the properties for that material. These can be edited as necessary, or you can proceed to define the continuum physics.



Defining Material Properties

When you select a material type in the *Material* box, the property values in the *Material Properties* box will automatically be updated to those of the chosen material. If you wish to enter different material properties, you can do so in the Material Properties box.

To define the material properties for the current material, locate the *Material Properties* box in the *Continuum Settings* dialog and modify the values as necessary.



The material properties that can be specified in STAR-SpaceClaim are listed below.

Density	Available for all materials. For gases, this property is available when the *Constant Density* equation of state is used.
Molecular Weight	Only applies to gases. This property is available when the *Ideal Gas* equation of state is used.
Specific Heat	Available for all materials when the temperature physics model is enabled.
Molecular Viscosity	Available for fluids only.
Conductivity	Available for all materials when the temperature physics model is enabled.

Defining Continuum Physics

Continuum physics define the characteristics and behavior of a particular substance, for example turbulent air, and are set by selecting a set of physics models within the continuum. In STAR-SpaceClaim, the physics models that can be defined are as follows:

- Flow characteristics (fluids only)
- Equation of state (gases only)
- Passive scalars
- Temperature
- Gravity
- Reference properties

Additional physics can be specified for each region, such as porosity.

Physics models are defined via the *Continuum Settings* dialog. To access this, right-click the Continuum node in the object tree and select Edit...



The physics models available in the *Continuum Settings* dialog are determined by the material type.



Flow Characteristics

For a gas or liquid, you can specify whether the flow is laminar or turbulent. When turbulent flow is selected, the Realizable Two-Layer K-Epsilon turbulence model is enabled in STAR-CCM+. This model is suitable for a wide variety of CFD cases, and typically provides a good initial solution for turbulence.

To select laminar or turbulent flow, first ensure that a gas or liquid material type has been selected, then open the *Continuum Settings* dialog. Locate the *Flow Characteristics* box and select either the Laminar or Turbulent radio button.



Selecting the Turbulent physics model will enable you to define turbulent conditions at boundaries. It will also be possible to specify initial values for the turbulent flow.

Equation of State

The equation of state model is used to define the density of the continuum substance. The density can either be specified as a constant value, or computed with respect to temperature and pressure, using the ideal gas law. An equation of state model can only be selected if you are working with a gas.

To select an equation of state model, first ensure that a gas material type has been selected, then open the *Continuum Settings* dialog. Locate the *Equation of State* box and select either the Constant Density or Ideal Gas radio button.



The Ideal Gas model requires that you also enable the temperature physics model.

Optional Physics Models

The optional physics models available in STAR-SpaceClaim are models that are not essential in order to obtain a solution for the model, but can be enabled if necessary. They include passive scalar and temperature.

Passive Scalar

The passive scalar model can be used to define a passive scalar on an inlet or pressure outlet boundary. A passive scalar injected into a simulation behaves in much the same way as a dye injected into water. The movement of the passive scalar can be used to identify how fluid entering a volume from a specific point will be transported downstream.

The passive scalar is defined by a concentration factor at each boundary. Its concentration throughout the flow field can be subsequently visualized at the end of the simulation.

To select the passive scalar physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Passive Scalar* checkbox.



Once enabled, the concentration of the passive scalar at each boundary can be specified.

Temperature

The temperature physics model enables the energy equation in STAR-CCM+ and allows you to define thermal physics conditions in the model. This includes thermal material properties, initial temperatures, reference temperatures and boundary conditions. The temperature model must be enabled if you have selected the ideal gas model for the equation of state.

To select the temperature physics model, open the *Continuum Settings* dialog. Locate the *Optional Physics Models* box and tick the *Temperature* checkbox.



Gravity

The gravity physics model enables the inclusion of the body force due to gravity in the momentum equations in STAR-CCM+.

To select the gravity physics model, open the *Continuum Settings* dialog. Locate the *Gravity* box and tick the *Enable* checkbox. This will activate the *Gravity* box.



Gravitational acceleration is defined by specifying a magnitude in the box, and a direction vector with components `[i, j, k]`.

Setting Reference Properties

Reference properties are specified in the *continuum* and are available depending on the physics models that you select. In the *Continuum Settings* dialog, locate the *Reference Properties* box and expand , if necessary.



The reference properties are described below:

- Reference Pressure

  The pressure values defined in the simulation are all relative to the reference pressure. See the section, What is the Reference Pressure?, for more information.

- Reference Density, Reference Velocity, Reference Area.

  These three quantities are not used as part of the CFD calculation, but are necessary for the calculation of force coefficients in post-processing.

To set any of the reference values, enter the new value in the corresponding box, then click OK to confirm the change(s) and exit the *Continuum Settings* dialog.

Defining Region Physics

In STAR-SpaceClaim, the physics models that can be defined in a region are as follows:

- Porosity
- A moving frame of reference
- Energy source

These settings apply to a whole region, and are defined in addition to continuum physics. A region can also be imported into STAR-CCM+ for visualization purposes without it affecting the analysis.

Physics models are defined via the *Region Settings* dialog. To access this, right-click the Region node in the object tree and select Edit...



The physics models can then be enabled for the region by ticking the corresponding checkbox and entering values as necessary.



Defining a Porous Region

Porous resistance affects fluid flow through a region in a particular direction. To define a porous region, open the *Region Settings* dialog and tick the *Porous* checkbox.

There are two types of porous resistances available, *Isotropic* and *Uni-Directional*.

Isotropic

To specify isotropic porous resistance, select the radio button labeled *Isotropic*.



The properties for isotropic porous resistance are the coefficients of inertial and viscous resistance, α and β respectively.

Uni-Directional

To specify uni-directional porous resistance, select the radio button labeled *Uni-Directional*.



The properties for uni-directional porous resistance are the coefficients of inertial and viscous resistance, α and β respectively. The cross-flow resistance is also required in order to determine the coefficients of inertial and viscous resistance in a direction perpendicular to the specified vector, **i** **j** **k**.

Using Moving Reference Frames

Moving reference frames can be used to simulate region motion. To define a moving reference frame, open the *Region Settings* dialog and tick the *Moving Reference Frame* checkbox.



In the *Moving Reference Frame Settings* box, enter values to define the properties of the moving reference frame.



Enter the angular velocity in the box. Define the axis of rotation by specifying the position of the origin, `[X, Y, Z]`, and a direction vector, i j k. The units used are the same as the file defaults.

Alternatively, click on and select two vertices on the geometry to define the axis of rotation. Holding down `<Ctrl>`, select two vertices on the geometry; the first selection will act as the origin of the axis, and the second selection will define the direction vector of the axis.The Use Assembly Coordinate System checkbox controls whether the properties of the axis of rotation are defined using the assembly's coordinate system or the selected part's coordinate system. If multiple parts are selected for the new region, the Use Assembly Coordinate System checkbox will be enabled and greyed out.

Defining an Energy Source

The energy source model allows you to define a direct source or sink of energy to or from a region. It is only available if the temperature physics model has been enabled in the continuum.

To specify that a region is an energy source, open the *Region Settings* dialog and tick the *Energy Source* checkbox.

This action activates the Energy Source Settings box. There are two types of energy sources available: energy input per unit volume, and total energy input.

Energy Input per Unit Volume

To define the energy input per unit volume, select the radio button labeled *Energy Input per Unit Volume*.



Enter the value for the energy input value per unit volume in the box.

Total Energy Input

To define the total energy input, select the radio button labeled *Total Energy Input*.



Enter the value for the total energy input in the box.

Importing Geometry Into STAR-CCM+ for Visualization

CAD parts that do not form part of the analysis can still be imported into STAR-CCM+ for visualization purposes. These parts should be added to a separate region that is excluded from the analysis. This feature is useful for showing the original geometry in a post-processing scene in STAR-CCM+.

Excluding a Region from the Analysis

To exclude CAD parts from the analysis, create a new region and tick the *Exclude From Analysis* checkbox in the *Advanced* panel in the *Region Settings* dialog. Add the CAD parts to be excluded to this region.



A surface mesh is generated for this region using the settings specified in the Mesh Pipeline. As is the case for all assemblies, the parts that are used for visualization purposes need to be imprinted onto any surface that they cut, otherwise the surface mesh will not be generated.

An example of this feature is shown in the next section, where the original valve and valve body geometry are imported into STAR-CCM+ for visualization.

Setting Up a STAR-CCM+ Scene to Show the Imported Geometry

When the solver has finished running, create a new scene in STAR-CCM+ and set it up to display what you desire. The original geometry can then be added to the scene.

In the example below, a *Geometry Scene* is created to display streamlines through a valve and valve housing.

1. First create a new geometry scene.
2. Create a streamlines derived part.
3. Click on the scene/plot button and double-click on the Geometry Scene > Displayers > Geometry > Parts node.



4. In the *Edit* dialog:
   a. Deactivate Regions
   b. Expand the Parts folder and activate the part that corresponds to the valve body. This is the region that you chose to exclude from the analysis in STAR-SpaceClaim.
5. Navigate to the Geometry node.

6. In the *Properties* window, set Representation to Geometry.



The imported geometry only contains a surface mesh and does not appear by default (the Representation is set to Volume). This is why you must change the Representation to Geometry. Finally, tick the Outline checkbox in the *Property* window to change the color of the outlines from red to black.

The finished scene is shown below.



Displaying Imported Geometry in Scalar and Vector Scenes in STAR-CCM+

To display an imported geometry in a Scalar or Vector scene, a new geometry displayer needs to be added to the existing scene. The geometry can then be displayed following steps 3 to 5 described above.

Setting Initial Conditions

Specifying initial conditions that are close to the expected solution can reduce the number of iterations required to reach a solution. The initial conditions may also be set similar to the inlet conditions of the fluid domain. Initial conditions are specified at a continuum level; if no initial conditions are specified, STAR-SpaceClaim will use default values to initialize the solution. In STAR-SpaceClaim, you can enter initial values for the velocity, pressure, temperature and turbulence within a fluid, and temperature in a solid.

To define the initial conditions of a continuum:

- Edit the continuum by right-clicking on the Continuum node in the STAR-SpaceClaim object tree and select Edit.



This opens the *Continuum Settings* dialog, in which the *Initialization* box can be used to specify the following values:



- U V W Initial velocity
- Initial pressure
- Initial temperature
- Int len Initial turbulence

Specifying the Initial Velocity

The initial velocity can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial velocity defines the starting velocity throughout any region to which the continuum is assigned. It is specified as a velocity vector, with components `[U, V, W]`. The default initial velocity is zero.

To define an initial velocity, open the *Continuum Settings* dialog, expand the *Initialization* box, and enter the values in the U, V, and W boxes.

Specifying the Initial Pressure

The initial pressure can be specified for a continuum in which a fluid material type has been selected. The physical value of the initial pressure defines the starting pressure in the regions to which the continuum is assigned. The initial pressure is specified with respect to the reference pressure. The default initial pressure is `0 Pa`.

To define an initial pressure, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the box.

Specifying the Initial Temperature

The initial temperature can be specified in any continuum in which the temperature physics model has been enabled. This physical value of the initial temperature defines the starting temperature throughout any region to which the continuum is assigned. The default initial temperature is `300 K`.

To define an initial temperature, open the *Continuum Settings* dialog, locate the *Initialization* box, and enter a value in the box.

Specifying the Initial Turbulence

The initial turbulence can be specified for a continuum in which a fluid material type has been selected, and for which the turbulent flow physics model has been enabled. When specifying the initial turbulence, values can be entered for the turbulence intensity and the turbulent length scale. These values define the turbulent conditions throughout any region to which the continuum is assigned. The default value for the initial turbulence intensity is `0.05`, or `5%`, and for the initial turbulent length scale is `0.01 m`.

To define an initial turbulence, open the *Continuum Settings* dialog, expand the *Initialization* box and enter a value in the Turbulence Intensity Int and Turbulent Length Scale len boxes.

Setting Boundary Conditions

Boundary conditions tell the STAR-CCM+ solvers how the substance within the region interacts with its environment. A velocity inlet boundary, for example, tells the solvers that an amount of substance is being added to the fluid domain at a certain rate. A STAR-SpaceClaim boundary is made up of one or more CAD surfaces within the simulation which, along with other boundaries, surround and define a region.



The boundary is discretized by the surface mesher, which is then used by the volume mesher to mesh the entire region.

It is important to set the boundary conditions before creating the volume mesh, as some parts of the volume mesh, such as prism layers, depend on the type of boundary.

By default, each region contains only one boundary. To define more boundaries you need to add boundaries to the region. You can modify the properties of existing boundaries by editing its boundary node.

Some boundaries may require the specification of turbulence quantities, temperature and passive scalar concentration, depending on whether these modeling options have been activated in Continuum Settings.

Creating a new Boundary

To create a new boundary, right-click the Boundaries manager within the object tree and select Add Boundary from the pop up menu.



The *Boundary Settings* dialog appears.



Choose the boundary type and set the associated settings.



Proceed by assigning surfaces to the boundary. By default, all the surfaces on the CAD geometry are assigned to the default boundary (this is the initial boundary that is created by STAR-SpaceClaim). When you assign a surface to another boundary, it is automatically removed from the default boundary. If you remove a surface from a boundary, it is automatically reassigned back to the default boundary.

Click OK to create the new boundary. New boundaries that you create are added to the object tree under the Boundaries manager node.

Specifying an Inlet Boundary

There are three types of inlet boundary available in STAR-SpaceClaim, which differ depending on the properties used to define the inlet flow:

- Velocity
- Mass flow rate
- Stagnation inlet

Using a Velocity Inlet

To set a velocity inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Velocity Inlet from the drop-down menu.

The values required to define a velocity inlet are detailed below.

Velocity

Defined in the box labeled *Velocity Inlet*.



The default input is normal velocity, with the drop-down list used to change units.

To enter the velocity in terms of Cartesian vectors U V W, select Cartesian Components from the drop-down menu.



Cylindrical vectors V_r V_θ V_z can also be used, selected from the drop-down menu. In addition to vector quantities, cylindrical components require a user defined axis origin, X Y Z, and a direction vector, i j k.



Using a Mass Flow Inlet

To set a mass flow inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Mass Flow Inlet from the drop-down menu.

The values required to define a mass flow inlet are detailed below.

Mass Flow

Defined in the box labeled *Mass Flow Inlet*.



The mass flow rate , must be specified, with the drop-down list used to change units.

To enter the mass flow rate in terms of direction vectors , , , tick the checkbox labeled Specify Direction.



Using a Stagnation Inlet

To set a stagnation inlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Stagnation Inlet from the drop-down menu.

The values required to define a stagnation inlet are detailed below:

Stagnation Pressure

Defined in the box labeled *Stagnation Inlet*.



The stagnation pressure , must be specified, with the drop-down list used to change units.

To enter the pressure in terms of direction vectors i, j, k, tick the checkbox labeled Specify Direction.



Specifying an Outlet Boundary

There are two types of outlet boundary available in STAR-SpaceClaim, which differ depending on the desired flow properties at the outlet:

- Flow-split
- Pressure

Using a Flow-Split Outlet

To set a flow-split outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Flow-Split Outlet from the drop-down menu.

The only parameter required is the flow-split ratio 60/40, expressed as a value between 0 and 1. The sum of flow-split ratio values for all flow-split outlet boundaries in the region must be equal to 1.



Using a Pressure Outlet

To set a pressure outlet, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Pressure Outlet from the drop-down menu.

The values required to define a pressure outlet are detailed below:

Pressure

Defined in the box labeled *Pressure Outlet*.



The pressure , must be specified, with the drop-down list used to change units.

Specifying a Wall Boundary

To set a wall boundary, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Wall from the drop-down menu.

A wall boundary can be defined as stationary, or having a particular motion, such as translating or rotating. This will affect the behavior of fluid adjacent to the wall boundary. You can also indicate slip conditions at the wall.

Wall boundaries can also be specified as having thermal properties.

Defining Wall Motion

When you define wall motion, this will not cause the boundary to move in space, rather it will impart this motion on the fluid adjacent to the wall boundary. By default a wall is defined as stationary. Choose the wall motion from the *Wall Type* drop-down menu.



All the wall motion options are listed below.

Stationary

No other boundary conditions are required for a stationary wall.

Translating

The velocity of the translating wall is then entered in vector form, U V W.



Units are specified in the drop-down list.

Rotating

The rotating wall requires a rotational velocity ω. Cylindrical components also require a user defined axis origin, X Y Z, and a direction vector, i j k.



Units are specified in the drop-down list.

Slip

No other boundary conditions are required for a slip wall.

Defining Wall Thermal Properties

By default, wall boundaries are defined as adiabatic.

Thermal Specifications can only be defined when the temperature model is activated.

Choose the wall thermal properties from the *Thermal Specification* drop-down menu.



The available wall thermal specification options are listed below.

Adiabatic

No other boundary conditions are required for an adiabatic wall.

Temperature

This requires the wall temperature to be specified.



Units are specified in the drop-down list.

Heat Flux

This requires the wall heat flux , to be specified.



Units are specified in the drop-down list.

Convection

This requires an ambient temperature , and heat transfer coefficient .



Units are specified in the drop-down list.

Specifying a Symmetry Boundary

To set a symmetry plane, first add a new boundary, opening the *Boundary Settings* dialog. Then locate the *Type* box and select Symmetry Plane from the drop-down menu.

No boundary conditions are required for a symmetry plane.

Specifying a Custom Boundary

The custom boundary type can be used in STAR-SpaceClaim to act as a placeholder for a boundary that will be customized in STAR-CCM+. Additional boundary types are available in STAR-CCM+, so when the model `.sim` file is exported from STAR-SpaceClaim and opened in STAR-CCM+, you can select a different boundary type.

The only properties that can be specified for a Custom boundary are the boundary name, and selected faces.

Defining a Temperature Boundary Condition

The temperature of a boundary is defined in the box labeled *Temperature, Turbulence*.



Temperature settings are only required when the temperature physics model is activated.

Units are specified in the drop-down list.

Defining Turbulent Boundary Conditions

Defined in the box labeled *Temperature, Turbulence*.



Turbulence settings are only required when the turbulence physics model is activated.

The required input for turbulence is:

- Turbulence Intensity Int
- Turbulence Length Scale len

Units are specified in the drop-down list.

Defining a Passive Scalar

Defined in the box labeled *Temperature, Turbulence, Passive Scalar*.



This option allows the user to define the concentration of passive scalar used at the boundary.

Passive scalar settings are only required when the passive scalar physics model is activated.

Meshing the Model

STAR-SpaceClaim provides options for generating polyhedral or trimmer volume meshes, with or without near-wall prism layers. By default, prism layers are generated on all wall surfaces. The provision of a prism layer depends on the type of boundary that is applied to a face. As a consequence, it is recommended that you only generate the mesh after setting up the boundaries and regions.

The meshing process first creates a triangulated surface (the surface mesh) and thereafter converts it to the desired volume mesh (polyhedral or trimmer). As the quality of the volume mesh depends to a large extent on the quality of the surface mesh, it is prudent to check the surface mesh before proceeding with the volume mesh.

Although the entire meshing process can be launched with a single command (the Build And Run command from the *STAR-SpaceClaim* ribbon), it is possible to execute only single stages from the STAR-SpaceClaim ribbon. Settings that are made in the Mesh Pipeline node govern the mesh generation process.



A mesh pipeline contains all the settings that are applied to the surface and volume meshes created on the SpaceClaim parts to which it is applied. In single part mode, only one mesh pipeline is available. In assembly mode, multiple mesh pipelines can be created if necessary.

When mesh generation commences, each mesh pipeline will:

1. Extract geometry from SpaceClaim.
2. Transfer the model to STAR-CCM+.
3. Execute the surface remesher.
4. Execute the polyhedral or trimmer volume mesher with the prism layer mesher (if selected).
5. Create a mesh representation in the SpaceClaim *Graphics* window.

Controls on the top-level mesh pipeline node allow you to set the base sizes for the surface mesh, and control the prism layer. Further refinement to specific faces, parts, or volume shapes can be made by adding additional controls to the mesh pipeline.

Creation of Additional Surfaces in STAR-CCM+

STAR-CCM+ creates geometry part surfaces during the geometry transfer stage. Initially, STAR-SpaceClaim creates one part surface for each face in the CAD model. These part surfaces are then combined so that one surface per boundary remains. This outcome, however, is not possible for boundaries that contain CAD surfaces from more than one part. The final part surfaces appear in the Geometry > Parts node in STAR-CCM+.

The following items affect the number of part surfaces that are created in STAR-CCM+:

- The number of boundaries present in the model, including the default boundary. Each boundary in STAR-SpaceClaim appears as a separate part surface in STAR-CCM+. If, however, a boundary contains CAD surfaces from multiple parts, the part surfaces in STAR-CCM+ cannot be combined.
- The number of mesh pipeline controls used (if any). Each CAD surface that is used to define a face mesh control in STAR-SpaceClaim appears as a separate part surface in STAR-CCM+.
- The number of interfaces that are created during the meshing process (for assemblies only). All the CAD surfaces that define one interface between two parts in an assembly are grouped together under one part surface in STAR-CCM+. One part surface exists for each interface in STAR-CCM+.

Boundary Splitting due to Interface Generation

Interfaces are automatically created between two parts in an assembly during the mesh generation. A new boundary is created for each interface on the surfaces that share the interface.



Boundary Splitting due to Face Mesh Controls and Part Mesh Controls

If you use a face mesh control or a part mesh control in your simulation, some boundaries may be split, depending on what faces are included in the mesh control. If a boundary and mesh control share the same faces, no

splitting occurs. If a mesh control contains faces from multiple boundaries, splitting occurs and a new boundary is created. The name of the new boundary is derived from the original boundary.



Working with Mesh Pipelines

When working with a single part, you will only have one mesh pipeline defined in the model tree. This will specify the mesh settings for the single part available. In assembly mode, multiple mesh pipelines can be created to give specific controls for different parts, and parts must be specifically associated with a mesh pipeline.

Setting General Mesh Pipeline Parameters

The mesh pipeline node contains the general mesh settings that will apply to all SpaceClaim parts associated within the pipeline in the absence of any further controls.

To access the basic mesh settings:

- Right-click on the Mesh Pipeline node and choose Edit.



The *Mesh Pipeline Settings* dialog appears. This dialog has several collapsible sections for the general settings.



These settings are described in the relevant sections below:

- Mesh Type
- Base Size
- Surface Mesh Size
- Prism Layer
- Advanced Settings

- When you have made the necessary settings click OK to save them and close the dialog.
- Click Cancel to exit without making any changes.

Specifying the Mesh Type

The mesh type specifies what type of volume mesh will be generated. There are two options available in the *Mesh Type* panel: Polyhedral and Trimmer.



Selecting the trimmer option will activate the Trimmer tab..

The trimmer mesh should not be applied to more than one part in an assembly as this will generate a non-conformal mesh at the interfaces and a warning message will appear in the Checker tab in the STAR-SpaceClaim Info App window when attempting to generate a volume mesh.

Specifying the Base Size

The base size is a reference length used in other meshing parameters. It may be set to equal the length of the largest triangle you expect to have on your surface mesh, or to some other characteristic dimension.



The base size can be set as either a relative or absolute value. When it is set as a relative value (by selecting the Relative to Geometry option), the reference size will be the diagonal of the bounding box enclosing the entire geometry. In the case of multiple parts in an assembly, the reference size will be calculated based on the diagonal length of the bounding box enclosing the entire assembly. This includes all parts from all mesh pipelines. The unit specification menu will become active when setting an absolute value.

Specifying the Surface Mesh Size

The surface mesh size is the base edge length of triangles on the surface mesh. Note that the ultimate length of some (or all) triangles on the surface mesh will depend on other factors such as surface curvature and edge proximity.



This parameter can be set relative to the base size or as an absolute length. The unit specification menu will become active for an absolute value.

Controlling the Prism Layer

A prism layer mesh is used to improve the simulation of flow characteristics near a non-slip wall, particularly when a turbulence model is active.



Three parameters are provided to control the creation of prism layers:

- Layer thickness
- Number of layers
- Thickness ratio

These are all provided in the *Prism Layer* box of the *Mesh Pipeline Settings* dialog. Creation of a prism layer mesh can be enabled or disabled by ticking or clearing the checkbox.



Layer Thickness

The layer thickness defines the overall depth of the prism layer, which includes all of the individual layers, as illustrated below.



The thickness can either be set relative to the base size, or as an absolute value with user specified length units.

Number of Layers

The number of layers defines how many individual cell layers will be included in the overall prism layer, as illustrated below.



Thickness Ratio

The thickness ratio is used to size the first and last layers within the prism layer.



Setting Advanced Surface Mesh Parameters

The advanced panel provides further customization of meshing parameters.



This panel allows you to change the following parameters:

- Relative Minimum Surface Size

- Growth Ratio
- Number of Points on a Circle

Relative Minimum Surface Size

The relative minimum surface size defines the minimum triangle size permitted on a feature curve. The minimum triangle size is set as a percentage of the surface mesh size. The default value is set to 25%. You must set the relative minimum surface size to a value greater than zero. If you set the relative minimum surface size to 100%, STAR-CCM+ will not modify the triangle size when it detects curvature.

The figure below illustrates this behavior. Note that the base size was set to 20% in this example to emphasize the refinement on the curve.



Having a large value for the number of points on a circle property increases the mesh resolution on curved surfaces.

The transition between the fine mesh and coarse mesh areas is controlled by the growth ratio.

Growth Ratio

The growth ratio controls the rate at which triangle sizes change when moving from one triangle size to another. In the example below, the end faces of the block were set to have a smaller triangle size than the four side

faces of the block. When the growth ratio is set to a low value, the transition between triangle size is very gradual. When a higher value is used, the transition is far more rapid.



The growth factor must be set to a value greater than 1.0. There is no upper limit for the growth factor.

Number of Points per Circle

The Number of Points per Circle parameter affects the number of triangle edges that appear on a curved surface, as shown below.



16 Points/Circle



32 Points/Circle

Regardless of the number of points per circle defined, STAR-SpaceClaim will not permit the triangle edge length to shrink below 25% of the base size. The curvature length scale may also affect the number of edges on a curved surface.

Controlling the Trimmer Mesh

The trimmer mesh is controlled using the Trimmer panel.



The parameters available are: Default Growth Rate and Boundary Growth Rate.

Default Growth Rate

The trimmer Default Growth Rate option controls the stepping of the refinement from one cell size to the next. Four settings are available, as shown below:



By default, the Default Growth Rate is set to Fast, meaning that the fastest transition from the smallest cell size to the largest will be used. This usually equates to having only one layer of each cell size for each transition, which will result in the least number of cells being created for the supplied input values. The Medium, Slow and Very Slow options further slow the transition from one cell size to the next in the mesh interior and therefore increase the overall cell count. In general, the medium option results in a minimum of two equal sized cell layers per transition, slow in four layers and the very slow option in eight layers.

Boundary Growth Rate

The trimmer Boundary Growth Rate option determines the number of cell layers to grow next to surfaces in the template. Five settings are available, as shown below:



By default, the Boundary Growth Rate is set to None, meaning that the setting used for the Default Growth Rate is used for controlling all transitions. If the Very Slow, Slow, Medium or Fast options are selected then it will result in eight, four, two or one layers of cells respectively of the same size being created next to the geometry surfaces in the template.

Creating a New Mesh Pipeline (Assembly Mode)

It is possible to create additional mesh pipelines when working with assemblies, so that different meshing parameters can be assigned to different parts. Initially, all of the SpaceClaim parts will be assigned to the default mesh pipeline. These can be reassigned to alternative mesh pipelines when they are created. A part can only belong to one mesh pipeline. In assembly mode, a Mesh Pipelines manager node is added to the object tree.



User specified mesh controls are located within the particular mesh pipeline.

To create a new mesh pipeline in assembly mode:

- Right-click the Mesh Pipeline node within the object tree and select Add Mesh Pipeline.



During the creation of a new mesh pipeline, it is necessary to assign to it the SpaceClaim parts that are to be meshed using its parameters.

Working with Mesh Controls

Mesh controls allow you to define custom mesh settings for a particular section of the model. The following types of mesh control are provided:

- Face Mesh Control

  This mesh control is applied to one or more faces of the SpaceClaim model.

- Volumetric Mesh Control

  Use this control to define a box or sphere shape within which the mesh settings will apply.

- Part Mesh Control (Assembly mode)

  This control applies custom mesh settings to selected parts.

Multiple controls can be assigned to the same item to control different mesh parameters. If, for example, a face has two mesh controls assigned to it--one that controls the mesh size and another that controls the prism layer thickness, the final mesh will respect both the custom mesh size and the custom prism layer thickness. If two mesh controls set the same mesh parameter, the control created last will take precedence over the first.

Using a Face Mesh Control

A face mesh control allows you to set the surface mesh size and the prism layer parameters on a face. Custom settings made on a face control will take priority over settings in the parent mesh pipeline, and over settings made in any part mesh controls.

To add a face mesh control:

- Right-click the Mesh Controls node belonging to the mesh pipeline node of interest and choose Add Face Mesh Control from the pop-up menu.



The *Face Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning faces to the face mesh control and then setting the required values in the custom mesh size and/or prism layer panels. If the trimmer mesher is selected, the custom boundary growth rate panel will also be available.

Setting Face Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: in the first instance, the face(s) will be assigned values from any part mesh control applied to the part to which the face(s) belong. If this does not apply, values will be taken from any applicable volumetric mesh control, and in the absence of this values will be taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the face(s).
- Disable: No prism layer will be generated on the face(s).

Using a Volumetric Mesh Control

A volumetric control permits mesh refinement within a user-defined box or sphere. Only the mesh size can be modified with a volumetric control; prism layer parameters are not affected.

To add a volumetric mesh control:

- Right-click the Mesh Controls node and select Add Volumetric Mesh Control from the pop-up menu.



The *Volumetric Control Settings* dialog appears.



The *Volumetric Control Settings* dialog provides the following functions:

- Set a custom mesh size
- Create a volume shape
- Edit a volume shape

- Delete a volume shape

Setting a Custom Mesh Size

The mesh size defines the triangle edge length for surface mesh within the volume shape, and the characteristic cell size for volume mesh. The custom mesh size will still affect the volume mesh in cases where the volume shape does not intersect any surfaces of the SpaceClaim geometry.

The mesh size can either be defined relative to the base size of the parent mesh pipeline, or as an absolute dimension with length units.

Creating a Volume Shape

Two types of volumetric shape can be created: a box or a sphere. When a shape has been created, it is added to the list of shapes in the *Shape* box. The custom mesh size will be applied to the combined volume of all shapes listed in the panel; multiple volume shapes can therefore be applied to the same volumetric control.

The name of the volume shape to be created should be entered in the Name entry box.

Box

To create a box, select the Box radio button within the *Shape* box.

X, Y and Z coordinates of the two opposite corners of the box are required by the panel. These coordinates can either be entered manually, or taken from existing points within the SpaceClaim window. To use an existing point, click the Set Point button and use the SpaceClaim point selection operation to select the appropriate point.

After completing the point coordinate entries, click the Create button to add the volume shape to the list.

Sphere

To create a sphere, select the Sphere radio button within the *Shape* box.



The sphere shape is defined by a center point and a radius. The center point can be set either by entering the coordinates directly, or by clicking the Set Point button and using the SpaceClaim point selection procedure to choose an existing point in the SpaceClaim window. The coordinates of the existing point will be used to populate the panel.

When the center coordinates and radius dimension have been correctly set, click the Create button to add the new sphere shape to the list.

Editing a Volume Shape

To edit a volume shape, select it from the list within the *Shapes* panel. Edit the parameters as required and click the Update button.



Deleting a Volume Shape

To delete a volume shape, select it from the list in the *Shape* panel and click the Delete button.

Using a Part Mesh Control (Assembly Mode)

Custom mesh settings (surface mesh size and prism layer parameters) can be applied to one or more parts using a part mesh control. Settings made on a part mesh control will take priority over settings made on the parent mesh pipeline, but will not have priority over settings made on a face mesh control.

To add a part mesh control:

- Right-click the Mesh Controls node and select Add Part Mesh Control from the pop-up menu.



The *Part Mesh Control Settings* dialog appears.



Custom mesh settings are applied by first assigning parts to the part mesh control and then setting the required values in the custom mesh size and/or prism layer panels.

Setting Part Mesh Control Parameters

To set a custom surface mesh size, tick the Enable checkbox within the *Custom Mesh Size* panel and make the required changes. For the prism layer, a drop-down menu is provided with the following options:

- Inherit Values: for this option, prism layers on the surfaces of the part(s) will be generated based on settings taken from the parent mesh pipeline.
- Use Custom Values: allows the prism layer parameters to be set directly on the part(s). These will take priority over settings made on the parent mesh pipeline, but not over settings made on any face controls associated with the same part(s).
- Disable: No prism layer will be generated from the surfaces of these part(s).

Generating the Mesh

It is recommended that you first generate the surface mesh and inspect its quality before generating the volume mesh. Problems in the surface mesh will have a definite impact on the volume mesh, and may result in the volume mesher failing to complete.

Mesh generation will be carried out automatically when you issue the Build and Run command.

Generating the Surface Mesh

To generate the surface mesh:

- Click the Generate Surface Mesh button located in the *STAR-SpaceClaim* ribbon.



Output from the surface mesher will be displayed in the *STAR-SpaceClaim Info App* window.

Generating the Volume Mesh

To generate the volume mesh:

- Click the Generate Volume Mesh button located in the *STAR-SpaceClaim* ribbon.



Information on the progress of the volume mesher will be displayed in the *STAR-SpaceClaim Info App* window.

Clearing Generated Meshes

The *Clear Generated Meshes* action deletes all generated meshes without affecting the settings.

To clear the mesh:

- Select Clear Mesh from the ribbon.



Visualizing the Mesh

Surface and volume meshes will automatically be displayed on first creation. If required, they can be hidden from view by clicking the Show/Hide Mesh button on the *STAR-SpaceClaim* ribbon.



Meshing Behavior at Interfaces

The volume mesher will treat interfaces between two regions differently depending on what type of region (fluid or solid) these are. If the interface lies between two fluid regions, no prism layers will be grown from both sides of the interface. If an interface lies between a solid region and a fluid region, prism layers will be grown from the interface. The default behavior of the prism layer generation can be overridden using a mesh control, so prism layers can be disabled on a solid-fluid interface.

Common Causes for Meshing Failures

There may be instances when the mesh generation fails; in most cases this will occur during the surface meshing phase. STAR-SpaceClaim will try and identify the cause and suggest ways of rectifying this. This information will be printed in the Output tab of the *STAR-SpaceClaim Info App* window.

Some common causes for the mesh generation to fail are listed below.

- The Parts in an Assembly are not Imprinted

The Parts in an Assembly are not Imprinted

The parts in an assembly should be imprinted together so that there is a common contact area between them. A simple example of when imprinting is required is shown below.



In this case the face of Part A (grey box) is splitting the face of Part B (blue box) causing this assembly to fail to mesh.

The output will display an error message.

```
Initializing meshing pipeline...
  All Geometry up to date.
Executing surface meshers...
Preparing surface for surface meshing in regions: Region 1, Region
Error during surface remeshing, deleting resulting mesh
An error occurred while attempting to orient the surface.
This is usually due to self-intersections in the surface. Check for problem(s) at :
  Geometric Range : (-0.0225836 0.00152109 0.0632873) and (0.0501103 0.0703362 0.0647127)
  Centroid of face : (0.0256414 0.0471603 0.064)
  Face type of face : 6
  Boundary = Default.
  Region = Region 1
Command: GenerateSurfaceMesh
   CompletedCommand: GenerateSurfaceMesh
   In: [Machine::main]
   Slave Message:
      Slave Name: Proc[0]
      Slave Rank: 0
   error: Server Error
```

The surface of Part B has to be imprinted with the contact face of Part A as shown below, after which meshing will complete successfully.



Obtaining a Solution

The following section details the execution of the STAR-CCM+ solvers in order to obtain a solution. This includes the operations listed below:

- Setting Solution Controls
- Checking the Model for Errors
- Building and Running the Simulation
- Running the Simulation
- Running more Iterations
- Clearing the Solution
- Visualizing the Solution

Setting Solution Controls

The solution you obtain from the solvers is a compromise between accuracy and processing time. You can manage the speed and accuracy of the solver from the *Solution Controls* dialog. The solution accuracy is also dependant on the mesh density.

To access the *Solution Controls* dialog:

1. Right-click on the STAR-SpaceClaim node within the object tree and select Set Solution Controls.



- Alternatively, from the ribbon, select Solution Controls.



The *Solution Controls* dialog appears.



2. Here, you can:
   - Choose from five pre-defined solver options, from high accuracy, low speed, to low accuracy, high speed.

- Manually enter the solution control parameters in the *Advanced* panel. It is recommended that you initially select the quickest solution control to verify that the simulation is providing reasonable results. Then increase the accuracy to an appropriate level for your system.

3. Click OK to close the dialog.

Setting Advanced Solution Control Parameters

If you want to have more control over the solution, adjust the advanced solution control parameters:



1. In the Number of Iterations field, set the maximum number of iterations for which the solver can run. When the STAR-CCM+ solver reaches the specified number of iterations, it will stop. You can then post-process the results in STAR-CCM+.
2. In the Convergence Level field, set how many orders of magnitude the residuals should drop before the solver stops.
3. In the Solution Accuracy field, select a numerical scheme to be used by the solver. There are two options available:
   - First Order—This option is generally the quicker, but less accurate solver option. The first order scheme introduces a dissipative error that is stabilizing and helps the solver achieve robust convergence. However, it can also lead to the neglect and merging of discontinuities, meaning the solution is not fully captured.
   - Second Order—This option is generally the more accurate, but slower solver option. The second order scheme attempts to minimize both dispersive and dissipative errors, providing a better representation of the simulation. However, in some situations, the lack of numerical dissipation can result in poorer convergence properties.

Note: It is common practice to use the first order scheme initially to ensure the simulation is running as expected and then switch to second order.

4. Adjust the under-relaxation parameters:
   - Momentum
   - Pressure
   - Turbulence (with turbulence modelling defined in Continuum Settings)
   - Temperature (with temperature modelling defined in Continuum Settings)
     - Fluid Temperature Under-Relaxation Factor
     - Solid Temperature Under-Relaxation Factor
5. If you want to use additional STAR-CCM+ physics models with STAR-SpaceClaim, activate the Use Java File checkbox. You can use Java macros with STAR-CCM+ to automate your simulation.

   See Using Java Macros
6. Click OK to close the dialog.

Note: Selecting one of the five predefined options will override any manual entries you have made.

Convergence Level

The numerical method used by the solver requires an iterative process in order to obtain a solution. After each iteration, residuals are produced that indicate how well the governing equations for each solver quantity are being satisfied numerically. In general, the residual magnitude should decrease as the solution converges. When the magnitude of the residuals for all the quantities falls below the convergence level, the solver will stop iterating, and results will be exported for post-processing.

Note: In most cases, a drop of 3-5 orders of magnitude in the residuals is sufficient to obtain an accurate solution.

You can adjust the convergence level parameter to meet accuracy or process time requirements. However, if the residuals are no longer decreasing from one iteration to the next, it is likely that machine accuracy has been reached for that case. For a case that does not exhibit decreasing residuals, it would be necessary to reconsider the setup in case the volume mesh is too coarse, or boundary conditions have been set incorrectly.

Under-Relaxation Parameters

Under-relaxation parameters determine how quickly the solver reaches a solution for a particular equation, on a scale from 0-1, with one being the fastest.

Ideally, all under relaxation parameters would be set to one, however, doing so would likely cause solver instability due to the complex nature of most simulations. Hence the under relaxation parameters should be set to values that compromise speed with stability.

Note: Adjust the under relaxation parameters if changes to other solution controls such as Number of Iterations, Convergence Levels and Solution Accuracy have failed to improve the behavior of the solver.

Using Java Macros

You can use Java macros with STAR-CCM+ to automate your simulation. You can trigger the use of a Java macro when STAR-CCM+ is launched from STAR-SpaceClaim by activating the Use Java File function in the *Solution Controls* dialog. To use a Java macro file in STAR-SpaceClaim:

1. Open the *Solution Controls* dialog.
2. Activate the Use Java File option.

   A text entry field appears.



3. Click the Select Input File button. This action opens a browser dialog.
4. Navigate to your Java macro file and select it. Selected files appear in the text entry field. The path in the entry field is displayed relative to the current CAD geometry file.
5. Click OK to close the dialog.

When you run the simulation, STAR-CCM+ will launch and play the macro before starting the solver. This sequence allows you to set up and use additional STAR-CCM+ physics models that are not available in STAR-SpaceClaim.

Using STAR-CCM+ Continuum Physics with STAR-SpaceClaim

You can use a Java macro to change a setting in STAR-CCM+, such as selecting an alternative turbulence model, before the solver starts running.

An example of how to do this is outlined below:

1. Launch STAR-SpaceClaim and set up the case, defining the continuum settings, boundary conditions, mesh settings and solution control settings.
2. Launch STAR-CCM+ .
3. In STAR-CCM+ , start recording a macro and perform the desired operations, such as changing the physics models.
4. Stop recording the macro. The macro saves at this point.
5. Return to STAR-SpaceClaim, open the *Solution Controls* dialog and activate the Use Java File checkbox.
6. Select the Java macro that you created earlier.
7. Run the case from STAR-SpaceClaim.

Checking the Model for Errors

STAR-SpaceClaim monitors the active case for any errors in real-time and displays a warning triangle ⚠ next to any nodes for which there is a problem.



Three areas are monitored by the live model checker:

- Continua—the continua checker makes sure that all necessary boundaries are present in the case, such as the presence of at least one inlet and one outlet.
- Regions and boundaries—the regions and boundaries checker covers issues specific to boundaries, such as a velocity inlet boundary having a zero value.
- Mesh pipelines—the mesh pipeline checker covers issues that are related to the mesh, such as the presence of multiple parts in a trimmed mesh pipeline.

To see the cause of an error:

1. Right-click on an node displaying the warning triangle, and select Show Errors...



A dialog appears describing the error, as shown below.



2. Click OK to close the dialog.

The Show Errors... right-click menu is available for any node in the STAR-SpaceClaim object tree that contains a warning triangle. Selecting it will show errors for that node and its sub-nodes. To see a complete list of issues for the case, right-click on the STAR-SpaceClaim root node and select Show Errors...

You can also manually check for errors. To perform a manual check:

- Select Check Model for Errors in the *STAR-SpaceClaim* ribbon.



If errors are found in the case set up, these are displayed in the Model Checker tab of the STAR-SpaceClaim Info App window.

Building and Running the Simulation

Once you have defined the model physics, boundary conditions, mesh settings, and solver settings in STAR-SpaceClaim, you can Build and Run the case.

To build and run the case:

- Select Build and Run in the *STAR-SpaceClaim* ribbon.



This action automatically generates the mesh, and runs the solver. When you run the solver, a new Residuals tab is added to the STAR-SpaceClaim Info App window. You can use this residuals plot to monitor solution convergence. When the solver finishes, STAR-CCM+ launches and a scalar scene is automatically created for you.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running

If you have already generated a volume mesh for the model, you can run the case directly. To do this:

- Select Run Solver in the *STAR-SpaceClaim* ribbon.



STAR-CCM+ runs the solvers until the specified number of iterations is reached or until the convergence criteria is reached.

Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, clicking Build and Run clears the solution and starts running from the first iteration. If you want to run for more iterations, change the stopping criteria and click Run Solver, or click Run More Iterations.

Running For More Iterations

You may need to run the STAR-CCM+ solvers for more iterations than you originally set in the Number of Iterations stopping criterion. The Run More Iteration function runs the solver for an additional number of iterations as specified in the stopping criteria. Take, for example, a case where you have already run the simulation for 300 iterations. You specify 500 iterations in the Number of Iterations stopping criterion. When you next click Run for More Iterations, the simulation will run for 500 extra iterations, making 800 iterations in total. The solver runs until one of the stopping criteria is satisfied.

To run for more iterations:

3. Select Run For More Iterations in the *STAR-SpaceClaim* ribbon.



Note: If a solution has already been computed in STAR-CCM+, and the stopping criteria has been reached, this function causes STAR-CCM+ to perform one additional iteration. To run for more iterations, change the stopping criteria.

Clearing the Solution

You can clear all results that were calculated by the STAR-CCM+ solvers using the *Clear Solution* function. All other STAR-SpaceClaim case settings remain intact.

Caution: This action is irreversible.

To clear the solution:

- From the ribbon, select Clear Solution.



The solution that was calculated by the STAR-CCM+ solvers is deleted.

Visualizing the Solution

When STAR-CCM+ launches after the solution is calculated, it automatically displays a scalar scene of the model that shows the `Absolute Pressure` field function. To visualize the results, you can use this scene to display different scalar quantities, or use any of the other visualization tools in STAR-CCM+, such as streamlines, plane sections and vector scenes.

You can import CAD parts into STAR-CCM+ for visualization purposes. These parts will not be taken into account during the analysis. This feature is useful for displaying the original geometry in a post-processing scene in STAR-CCM+.

Exporting Files and Data

The following section details the process for exporting files and data from STAR-SpaceClaim.

- Exporting a Surface Mesh
- Exporting a Volume Mesh
- Exporting an Engineering Report
- Exporting as a STAR-CCM+ Simulation

Exporting the Surface Mesh

You can export the surface mesh generated in STAR-SpaceClaim as a native CD-adapco surface mesh (`.dbs`) file. The `.dbs` file format is the recommended format to use when transferring surface data between CD-adapco programs. The boundary names and boundary types are retained by the file, as well as the interfaces that were automatically created between regions.

To export the surface mesh:

- From the STAR-SpaceClaim ribbon, select Export > Surface Mesh (.dbs).



Exporting the Volume Mesh

You can export the volume mesh that you generated in STAR-SpaceClaim as a native CD-adapco volume mesh (`.ccm`) file format. The `.ccm` file format is the recommended format to use when transferring volume meshes between CD-adapco programs. The region and boundary names, and region and boundary types are retained in this file. The interfaces that were automatically created interfaces between regions are also retained in this file.

To export the volume mesh:

- From the STAR-SpaceClaim ribbon, select Export > Volume Mesh (.ccm).



Exporting an Engineering Report

You can export a summary of the simulation settings and results in HTML format to two individual files:

- The first will be given a user-specified file name, and contains results from the simulation.
- The second will be given a user-specified file name appended with `_caseinfo`, and contains details on the simulation set-up.

Note: You can only export an engineering report after you obtain a solution.

To export an engineering report:

- From the STAR-SpaceClaim ribbon, select Export > Engineering Report (.html,.csv).



Results Report

The following is a summary of the information included in the results report (`<filename>.htm`). This data is provided in the report with reference to each material region in the simulation:

- Case name and time-stamp
- Solution information: CPU time, Elapsed Time and Number of Iterations
- Mass flow data, at each boundary
- Mass flux data, at each boundary
- Area averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)
  - Turbulent Viscosity (Pa-s)
  - Density (kg/m^3)
- Mass flow averaged boundary data, for the following scalars:
  - Pressure (Pa)
  - Total Pressure (Pa)
  - Pressure Coefficient
  - Total Pressure Coefficient
  - Turbulent Kinetic Energy (J/kg)
  - Turbulent Dissipation Rate (m^2/s^3)

    - Turbulent Viscosity (Pa-s)
    - Density (kg/m^3)
- Force data, at each boundary:
    - Net Force [Fx, Fy, Fz]
    - Pressure Force [Fx, Fy, Fz]
    - Shear Force [Fx, Fy, Fz]
- Force coefficient data, at each boundary:
    - Coefficient of Net Force [Cfx, Cfy, Cfz]
    - Coefficient of Pressure Force [Cfx, Cfy, Cfz]
    - Coefficient of Shear Force [Cfx, Cfy, Cfz]

Setup Report

The following is a summary of the information found in the simulation setup report (`<filename>_caseinfo.htm`):

- The parts used in the simulation, and their associated regions
- The surfaces for each part, and their associated boundaries
- The continua used in the simulation, with the following information:
    - Physics models applied to the continuum
    - Materials properties
    - Turbulence properties
    - Segregated flow properties
    - Reference values
    - Initial conditions
- The mesh pipelines used in the simulation, with the following information:
    - Polyhedral mesher properties
    - Prism layer mesher properties (if activated)
    - Surface remesher properties
    - Reference values
    - Volumetric controls
- The regions used in the simulation, and their associated continua and mesh pipeline
- The boundaries for each continua, with the applied conditions and values

- Solver parameters
- Stopping criteria
- Units

Exporting as a STAR-CCM+ Simulation

You can export the current simulation in STAR-SpaceClaim to the native STAR-CCM+ simulation (`.sim`) file format. The `.sim` file format retains all simulation settings defined in STAR-SpaceClaim, as well as the surface and volume meshes.

To export a simulation file:

- From the STAR-SpaceClaim ribbon, select Export > Simulation (.sim).



Resetting STAR-SpaceClaim Settings

STAR-SpaceClaim saves all case setup information to the CAD file. This means that all continuum settings, boundary settings, solution control settings, and mesh settings, are stored within the CAD geometry file. You can remove this information from the CAD geometry file.

Caution: This action is irreversible.

To clear all STAR-SpaceClaim settings:

- From the ribbon, select Clear All.



This action returns the geometry to its original state.

CAD Client Tutorials

This section contains tutorials for each of the CAD Clients. The tutorials are grouped in the following sections:

- STAR-Works
- STAR-CAT5
- STAR-NX
- STAR-Pro/E
- STAR-Creo
- STAR-Inventor
- STAR-SpaceClaim

STAR-Works CAD Client

Several tutorials have been provided for the STAR-Works CAD Client. These provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Works environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starworks\data`.

- Copy directory `data` to your own disk space and open the part file `exhaustManifold.SLDPRT` in STAR-Works.

The geometry is shown below.



The *STAR-Works* ribbon appears at the top of the screen. The STAR-Works object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-Works.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Works. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-Works tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces that are highlighted below. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Mass Flow Inlet from the Type box.
4. Set Boundary Name to `Inlet`.
5. Set Mass Flow Rate to `0.017` kg/s.

   The completed dialog is shown below.



6. Click ✓ to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type box.
4. Set Boundary Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Works tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to `2.5%`.
3. Expand the Prism Layer tab.
4. Set Prism Layer Thickness to `15%`.

The completed *Mesh Pipeline Settings* dialog is shown below.



5. Click to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Click Solution Controls on the STAR-Works ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `200`.
3. Set Convergence Level to `0.0001`.

   The completed property page is shown below.



4. Click to close the dialog.
5. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Works object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Click Surface Mesh from the ribbon.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Works *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the `<T>` key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Works environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starworks\data`.

- Copy directory `data` to your own disk space and open the assembly file `PorousModel.SLDASM` in STAR-Works.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

The geometry is shown below.



The *STAR-Works* ribbon appears at the top of the screen. The STAR-Works object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-Works.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Works. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Works tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click ✔ to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-Works tree and select New.



The *Region Settings* dialog appears.

2. Select the middle part of the assembly as shown below.



3. In the *Region Settings* dialog, set the Region Name to `Porous Region`.
4. Activate the Porous model.
5. Within the *Porous* box, select the Uni-Directional option.
6. Set Alpha α to `25` kg/m^4 and Beta β to `1500` kg/m^3.s.
7. Set Cross-Flow Resistance Factor to `250`.

   The completed *Region Settings* dialog is shown below.


Region Settings
Region
Name
Porous Region
Part Selection
porousmiddle-1@PorousModel_Se
Continuum
Continuum
Porous
Isotropic
Uni-Directional
α
25
kg/m^4
β
1500
kg/m^3.s
250
0
0
1
Use Assembly Coordinate System
Moving Reference Frame
Advanced
Exclude from Analysis

8. Click ✔.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-Works automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type box.
4. Set Boundary Name to `Inlet`.
5. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



6. Click .

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select New.

The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type box.
4. Set Boundary Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click ✓ .

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Works tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Deactivate the Relative to Geometry option in the *Base Size* box.
4. Set Base Size to `0.0025` m.
5. Expand the Prism Layer tab.
6. Activate the Use Absolute Value option for Prism Layer Thickness.
7. Set Prism Layer Thickness to `0.001` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Mesh Pipeline
Name
Mesh Pipeline
Part Selection
porousin-1@PorousModel
porousmiddle-1@PorousMode
porousoutlet-2@PorousModel
Mesh type
Polyhedral
Trimmer
Base Size
0.0025
m
Relative to Geometry
Surface Mesh Size
100
%
Use Absolute Value
Prism Layer
0.001
m
Use Absolute Value
2
1.5
Advanced Settings

8. Click .

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Click Solution Controls on the STAR-Works ribbon.



   The *Solution Controls* dialog appears.

2. Set Number of Iterations to `400`.
3. Set *Solution Accuracy* to First Order.

   The completed property page is shown below.



4. Click .
5. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Works object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Works *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Works environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starworks\data`.

- Copy directory `data` to your own disk space and open the assembly file `piccoloAssem.SLDASM` in STAR-Works.

The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

The geometry is shown below.



The *STAR-Works* ribbon appears at the top of the screen. The STAR-Works object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-Works.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-Works.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-Works tree and select Edit.



   The *Continuum Settings* dialog appears.

2. In the *Continuum Settings* dialog, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Expand the Initialization box by clicking .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.


Continuum Settings
Name
Air
Material
Gas
Liquid
Solid
Acetylene
Air
Aluminum
Aluminum Oxide
Ammonia
Argon
Butane
Carbon Dioxide
Flow Characteristics
Laminar
Turbulent
Equation of State
Constant Density
Ideal Gas



6. Click ✔ .

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-Works tree and select New.



The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material* menu, and select Aluminum Alloy from the list.
4. Expand the Initialization box by clicking .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click .

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-Works tree and select New.



The *Region Settings* dialog appears.

2. Select the wing skin as shown below. Once selected, the part becomes highlighted.



3. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
4. Set Continuum to Al using the drop-down menu.Set Continuum Selection to Al from the selection box.



This action assigns the aluminum alloy that you created earlier as the material for this part.

The completed dialog is shown below.



5. Click .

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-Works tree and select Rename.



2. Enter `Internal Air` as the new name.
3. Click a blank area to accept the name change.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces highlighted below. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type box.
4. Set Boundary Name to `Inlet`.
5. Set Normal Velocity to `90` m/s.
6. Set Temperature to `460` K.
7. Set Turbulence Intensity to `0.135`.
8. Set Turbulent Length Scale to `2.2e-005` m.

The completed dialog is shown below.


Boundary Settings
Type
Velocity Inlet
Mass Flow Inlet
Stagnation Inlet
Flow-Split Outlet
Pressure Outlet
Symmetry Plane
Wall
Custom
Face Selection
Face<1>@internalAir-1
Face<2>@internalAir-1
Face<3>@internalAir-1
Boundary Name
Inlet
Velocity Inlet
Enter Normal Velocity
90
m/s
Specify Velocity Components
Turbulence
Int
0.135
len
2.2e-005
m
Temperature
460
K

9. Click .

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* box.
4. Set Boundary Name to `Outlet`.

   The default settings are suitable for this tutorial.

5. Click .

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select New.

   The *Boundary Settings* dialog appears.

2. Select the four faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries.



3. Leave the selection as Wall in the *Type* box.
4. Set Boundary Name to `Piccolo Wall`.
5. Set Thermal Specification to Temperature.



6. Set Temperature to `460` K.
7. Click .

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select New.

   The *Boundary Settings* dialog appears.

2. Select the two faces highlighted below, one on either side of the internal air region.



3. Select Symmetry Plane from the *Type* box.
4. Set Boundary Name to `Air Symmetry`.
5. Click ✓ .

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.

2. Select the two faces highlighted below, one on either side of the geometry.



3. Select Symmetry Plane from the *Type* box.
4. Set Boundary Name to `Solid Symmetry`.
5. Click ✔ .

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Edit.

   The *Boundary Settings* dialog appears.

2. Set Thermal Specification to Convection.



3. Set Heat Transfer Coefficient HTC to `100` W/m^2.K
4. Set Ambient Temperature to `228` K.

   The completed dialog is shown below.



5. Click .
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Works tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Deactivate the Relative to Geometry option in the *Base Size* box.
4. Set Base Size to `0.007` m.
5. Expand the Prism Layer tab.
6. Set Prism Layer Thickness to `10`%.

The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Mesh Pipeline
Name
Mesh Pipeline
Part Selection
internalAir-1@piccoloAssem
wingSkin-1@piccoloAssem
Mesh type
Polyhedral
Trimmer
Base Size
0.007
m
Relative to Geometry
Surface Mesh Size
100
%
Use Absolute Value
Prism Layer
10
%
Use Absolute Value
2
1.5
Advanced Settings

7. Click ✓ .

Setting the Coincident Faces Detection Tolerance

The geometry that is used in this tutorial has large changes in feature size and requires you to reduce the Coincident Faces Detection Tolerance.

1. Right-click STAR-Works and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.

2. Set Coincident Faces Detection Tolerance to `0.0001`.



3. Click .

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.



The *Volumetric Control* dialog appears.

2. Set Name to `Inlet Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-100.00 mm`
   - Y Coordinate of Point 1 Y_1: `21.00 mm`
   - Z Coordinate of Point 1 Z_1: `-20.00 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-90.00 mm`
   - Y Coordinate of Point 2 Y_2: `10.00 mm`
   - Z Coordinate of Point 2 Z_2: `60.00 mm`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click .

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.

   The *Volumetric Control* dialog appears.
2. Set the Name to `Mid Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-111.00 mm`
   - Y Coordinate of Point 1 Y_1: `34.00 mm`
   - Z Coordinate of Point 1 Z_1: `-20.00 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-64.00 mm`
   - Y Coordinate of Point 2 Y_2: `-2.00 mm`
   - Z Coordinate of Point 2 Z_2: `60.00 mm`
7. Click Create.

The volumetric control region is shown below.



8. Click .
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Works tree and select New Part Mesh Control.

The *Part Mesh Control* dialog appears.

2. Select the solid wing region on the CAD model. The selection is added to the *Part Selection* box.
3. Set Name to `Solid Region`.
4. In the Prism Layer box, set Custom Prism Settings to Disable using the drop-down menu.



The completed dialog is shown below.



5. Click .

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Works tree and select New Face Mesh Control.



The *Face Mesh Control* dialog appears.

2. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



3. Set Name to `Inlet Refinement`.
4. Activate the Custom Surface Mesh Size option.
5. Activate the Use Absolute Value option.
6. Set Custom Surface Mesh Size to `0.00025` m.
7. In the Prism Layer box, set Custom Prism Layer Settings to Use Custom Values using the drop-down menu.
8. Activate the Use Absolute Value option.
9. Set Prism Layer Thickness to `0.0002` m.

   The completed dialog is shown below.



10. Click ✓ .

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Click Solution Controls on the STAR-Works ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `500`.
3. Set Convergence Level to `0.01`.
4. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

   The completed dialog is shown below.



5. Click .
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Works object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Works *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Works environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starworks\data`.

- Copy the `data` directory to your own disk space and open the assembly file `gpuAssem.SLDASM` in STAR-Works.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

The geometry is shown below.



The *STAR-Works* ribbon appears at the top of the screen. The STAR-Works object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-Works.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Works. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Works tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

The remaining default settings are kept for this case.



3. Click ✓ to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-Works tree and select New.



The *Region Settings* dialog appears.

2. Select the disc that is located inside the graphics card housing. It is possible to select this part in the geometry window, through the translucent case. Once selected, the part becomes highlighted as shown below.



3. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
4. Activate the Moving Reference Frame option.
5. Set Rotational Velocity, ω to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.
6. Set Coordinates of Axis Origin to the following (you can set units by typing them in as shown below):
   - X: `-122.1761 mm`
   - Y: `-1.9434 mm`

- Z: -19.0609 mm

7. Set Component of Flow Direction to the following:
   - i : 0
   - j : 1
   - k : 0
8. Make sure that Use Assembly Coordinate System is selected.

   The completed property page is shown below:


Region Settings
Region
Name
Rotating Region
Part Selection
gpuBlowerFan-1@gpuAssem
Continuum
Continuum
Porous
Moving Reference Frame
-2800
rpm
X
-122.1761mm
Y
-1.9434mm
Z
-19.0609mm
Choose Point...
0
1
0
Use Assembly Coordinate System
Advanced
Exclude from Analysis

9. Click ✓ .

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the *Type* box.
4. Set Boundary Name to `Case Inlet`.
5. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



6. Click .

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select New.

The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* box.
4. Set Boundary Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click ✔ .

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* box.

4. Set Boundary Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click ✔ .
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Works tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.
2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Deactivate the Relative to Geometry option in the *Base Size* box.
4. Set Base Size to `0.016` m.
5. Expand the Prism Layer tab.
6. Set Prism Layer Thickness to `20%`.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Mesh Pipeline
Name
Mesh Pipeline
Part Selection
gpuBlowerFan-1@gpuAssem
pcTowerCase-1@gpuAssem
Mesh type
Polyhedral
Trimmer
Base Size
0.016
m
Relative to Geometry
Surface Mesh Size
100
%
Use Absolute Value
Prism Layer
20
%
Use Absolute Value
2
1.5
Advanced Settings

7. Click ✓ .

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.



   The *Volumetric Control* dialog appears.

2. Set Name to `Fan Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-166.00 mm`
   - Y Coordinate of Point 1 Y_1: `4.00 mm`
   - Z Coordinate of Point 1 Z_1: `-61.00 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-79.00 mm`
   - Y Coordinate of Point 2 Y_2: `-34.00 mm`
   - Z Coordinate of Point 2 Z_2: `25.00 mm`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click ✔ .

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.

2. The *Volumetric Control* dialog appears.
3. Set Name to `GPU Box`.
4. Activate the Use Absolute Value option.
5. Set Custom Surface Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383.00 mm`
   - Y Coordinate of Point 1 Y_1: `16.00 mm`
   - Z Coordinate of Point 1 Z_1: `-111.00 mm`
7. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-52.00 mm`
   - Y Coordinate of Point 2 Y_2: `-49.00 mm`
   - Z Coordinate of Point 2 Z_2: `49.00 mm`
8. Click Create.

   The volumetric control region is shown below.



9. Click ✔ .

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.

   The *Volumetric Control* dialog appears.

2. Set the Name to `Case Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383.00 mm`
   - Y Coordinate of Point 1 Y_1: `30.00 mm`
   - Z Coordinate of Point 1 Z_1: `-111.00 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `43.00 mm`
   - Y Coordinate of Point 2 Y_2: `-159.00 mm`
   - Z Coordinate of Point 2 Z_2: `77.00 mm`
7. Click Create.

   The volumetric control region is shown below.



8. Click ✔ .

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Click Solution Controls on the STAR-Works ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `300`.
3. Set Convergence Level to `0.01`.
4. Set Momentum Under-Relaxation Factor to `0.85`.

The completed property dialog is shown below.



5. Click .
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Works object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Works *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Works environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starworks\data`.

- Copy directory `data` to your own disk space and open the part file `car.SLDPRT` in STAR-Works.

The geometry is shown below.



The *STAR-Works* ribbon appears at the top of the screen. The STAR-Works object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-Works.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-Works tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

The remaining default settings are kept for this case.



3. Click .

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Symmetry Plane from the *Type* box.
4. Set Boundary Name to `Symmetry Plane`.

   The completed dialog is shown below.



5. Click .

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Velocity Inlet from the *Type* box.
4. Set Boundary Name to `Inlet`.
5. Change the units of Normal Velocity by clicking and select mph.



6. Set Normal Velocity to `70` mph.
7. Click .

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the back face and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* box.
4. Set Boundary Name to `Outlet`.

   The default settings are acceptable for this case.
5. Click ✔ .

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the top and side walls, and select the faces that are highlighted below.



3. Select Symmetry Plane from the *Type* box.
4. Set Boundary Name to `Boundary Symmetry`.
5. Click .

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the bottom face, and select the face that is highlighted below.



3. Leave the default Wall in the *Type* box.
4. Set Boundary Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click .

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Rename.
2. Enter `Car` as the new name.
3. You may need to click on a blank area of the workspace to accept the name change.

4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Works tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Deactivate the Relative to Geometry option in the *Base Size* box.
3. Set Base Size to `0.03` m.
4. Expand the Prism Layer tab.
5. Activate the Use Absolute Value option for Prism Layer Thickness.
6. Set Prism Layer Thickness to `0.03` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.



7. Click .

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.



   The *Volumetric Control* dialog appears.

2. Set Name to `Car Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50.00` mm
   - Y Coordinate of Point 1 Y_1 : `-650.00` mm
   - Z Coordinate of Point 1 Z_1 : `3432.00` mm
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000.00` mm
   - Y Coordinate of Point 2 Y_2 : `1440.00` mm
   - Z Coordinate of Point 2 Z_2 : `-4567.00` mm
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click .

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Volumetric Control.

   The *Volumetric Control* dialog appears.

2. Set Name to `Wake Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 : `-50.00 mm`
   - Y Coordinate of Point 1 : `-650.00 mm`

- Z Coordinate of Point 1 Z_1 : 3432.00 mm

6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : 2000.00 mm
   - Y Coordinate of Point 2 Y_2 : 3240.00 mm
   - Z Coordinate of Point 2 Z_2 : -14567.00 mm
7. Click Create.

   The volumetric control region is shown below.



8. Click ✔ .

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Works tree and select New Face Mesh Control.



The *Face Mesh Control* dialog appears.

2. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



3. Set Name to `Fluid Domain Face Control`.
4. Activate the Custom Surface Mesh Size option.
5. Activate the Use Absolute Value option.
6. Set Custom Mesh Size to `0.35` m.

The completed dialog is shown below.



7. Click .

Modifying Solution Controls

Modify the solution controls:

1. Click Solution Controls on the STAR-Works ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `400`.
3. Set Convergence Level to `0.0001`.

The completed property page is shown below.



4. Click ✓ .
5. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Works object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Works *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.
3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below.



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-CAT5 CAD Client

Several tutorials have been provided for the STAR-CAT5 CAD Client. These provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-CAT5 environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcat5\data`.

- Copy directory `data` to your own disk space and open the part file `exhaustManifold.CATPart` in STAR-CAT5.
- Position the geometry as shown below.



The *STAR-CAT5* menu appears at the top of the screen. The STAR-CAT5 object tree is located on the left of the screen.

- Save the analysis as `exhaustManifold.CATAnalysis`.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-CAT5. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-CAT5 tree and select Continuum object > Definition.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces that are highlighted below. These selections are reported in the *Supports* box in the *Boundary Settings* dialog.



3. Select Mass Flow Inlet from the Boundary Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Mass Flow Rate to `0.017` kg/s.

   The completed dialog is shown below.



6. Click OK to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Boundary Type drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click OK to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-CAT5 tree and select Mesh Pipeline object > Definition.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to `2.5%`.
3. Set Layer Thickness to `15%`.

The completed *Mesh Pipeline Settings* dialog is shown below.



4. Click OK to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Right-click the Solution folder in the STAR-CAT5 tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Click More to expand the *Advanced* menu.
3. Set Iterations to `200`.
4. Set Convergence to `0.0001`.

The completed property page is shown below.



5. Click OK to close the dialog.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-CAT5 object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Surface Mesh* from the menu.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Volume Mesh* from the menu.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-CAT5* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-CAT5 *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the <T> key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-CAT5 environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcat5\data`.

- Copy directory `data` to your own disk space and open the assembly file `PorousModel.CATProduct` in STAR-CAT5.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

- Position the geometry as shown below.



The *STAR-CAT5* menu appears at the top of the screen. The STAR-CAT5 object tree is located on the left of the screen.

- Save the analysis as `exhaustManifold.CATAnalysis`.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-CAT5. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-CAT5 tree and select Continuum object > Definition.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-CAT5 tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the middle part of the assembly as shown below.



3. In the *Region Settings* dialog, set the Name to `Porous Region`.
4. Activate the Porous model.
5. Within the *Porous Settings* box, select the Uni-Directional option from the drop-down menu.



6. Set Alpha to `25` kg/m^4 and Beta to `1500` kg/m^3.s.
7. Set Resistance to `250`.

   The completed *Region Settings* dialog is shown below.



8. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-CAT5 automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below. The selection is reported in the *Supports* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Boundary Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



6. Click OK.

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click OK.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-CAT5 tree and select Mesh Pipeline object > Definition.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Supports* box.
3. Activate the Use Absolute Value checkbox in the *Base Size* box.
4. Set Base Size to `0.0025` m.
5. Activate the Use Absolute Value option for Layer Thickness.
6. Set Layer Thickness to `0.001` m.

The completed *Mesh Pipeline Settings* dialog is shown below.



7. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Right-click the Solution folder in the STAR-CAT5 tree and select Solution Controls.



The *Solution Controls* dialog appears.

2. Click More to expand the *Solution Controls* menu.
3. Set Iterations to `400`.
4. Set *Solution Accuracy* to First Order using the drop-down menu.

   The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-CAT5 object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Surface Mesh* from the menu.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Volume Mesh* from the menu.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-CAT5* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-CAT5 *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-CAT5 environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcat5\data`.

- Copy directory `data` to your own disk space and open the assembly file `piccoloAssem.CATProduct` in STAR-CAT5.

  The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

- Position the geometry as shown below.



The *STAR-CAT5* menu appears at the top of the screen. The STAR-CAT5 object tree is located on the left of the screen.

- Save the analysis as `piccoloAssem.CATAnalysis`.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-CAT5.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-CAT5 tree and select Continuum object > Definition.



The *Continuum Settings* dialog appears.

2. In the *Continuum Settings* dialog, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Click on the Initialization tab.
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click OK.

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-CAT5 tree and select Add Continuum.



The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material* menu, and select Aluminum Alloy from the list.
4. Click on the Initialization tab.
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click OK.

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-CAT5 tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the wing skin as shown below. Once selected, the part becomes highlighted.



3. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
4. Set Select Continuum to Al using the drop-down menu.



This action assigns the aluminum alloy that you created earlier as the material for this part.

The completed dialog is shown below.



5. Click OK.

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-CAT5 tree and select Region object > Definition.



The *Region Settings* dialog will appear.

2. Enter `Internal Air` as the new Name.
3. Click OK to close the dialog.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces highlighted below. These selections are reported in the *Supports* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Boundary Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `90` m/s.
6. Set Temperature to `460` K.
7. Set Intensity to `0.135`.
8. Set Length Scale to `2.2e-005` m

   The completed dialog is shown below.



9. Click OK.

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Boundary Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are suitable for this tutorial.
5. Click OK.

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the six faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries.



3. Leave the selection as Wall in the *Boundary Type* drop-down menu.
4. Set Name to `Piccolo Wall`.
5. Set Thermal Type to Temperature.



6. Set Temperature to `460` K.
7. Click OK.

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the two faces highlighted below, one on either side of the internal air region.



3. Select Symmetry Plane from the *Boundary Type* drop-down menu.
4. Set Name to `Air Symmetry`.
5. Click OK.

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the two faces highlighted below, one on either side of the geometry.



3. Select Symmetry Plane from the *Boundary Type* drop-down menu.
4. Set Name to `Solid Symmetry`.
5. Click OK.

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Default object > Definition.

   The *Boundary Settings* dialog appears.
2. Set Thermal Type to Convection.



3. Set Heat Transfer Coefficient to `100` W/m^2.K
4. Set Temperature to `228` K.

   The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-CAT5 tree and select Mesh Pipeline object > Definition.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Supports* box.
3. Activate the Use Absolute Value checkbox in the *Base Size* box.
4. Set Base Size to `0.007` m.
5. Set Layer Thickness to `10%`.

The completed *Mesh Pipeline Settings* dialog is shown below.



6. Click OK.

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Inlet Box`.
3. Activate the Use Absolute Value option.
4. Set Surface Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X1: `-100.00`
   - Y1: `21.00`
   - Z1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X2: `-90.00`

- Y2: `10.00`
- Z2: `60.00`

7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.

   The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Mid Box`.
3. Activate the Use Absolute Value option.
4. Set Surface Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X1: `-111.00`
   - Y1: `34.00`
   - Z1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X2: `-64.00`
   - Y2: `-2.00`
   - Z2: `60.00`

7. Click Create.

   The volumetric control region is shown below.



8. Click OK.
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-CAT5 tree and select Add Part Mesh Control.



The *Part Mesh Control Settings* dialog appears.

2. Select the solid wing region on the CAD model. The selection is added to the *Supports* box.
3. Set Name to `Solid Region`.
4. In the Prism Layer box, set Custom Settings to Disable using the drop-down menu.



The completed dialog is shown below.



5. Click OK.

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-CAT5 tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



3. Set Name to `Inlet Refinement`.
4. In the Surface Mesh Size box, activate the Custom Mesh Size option.
5. Activate the Use Absolute Value option.
6. Set Mesh Size to `0.00025` m.
7. In the Prism Layer box, set Custom Settings toUse Custom Values using the drop-down menu.
8. Activate the Use Absolute Value option.
9. Set Layer Thickness to `0.0002` m.

   The completed dialog is shown below.



10. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Right-click the Solution folder in the STAR-CAT5 tree and select Solution Controls.



The *Solution Controls* dialog appears.

2. Click More to expand the *Advanced* menu.
3. Set Iterations to `500`.
4. Set Convergence to `0.01`.
5. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

   The completed dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-CAT5 object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-CAT5* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-CAT5 *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-CAT5 environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcat5\data`.

- Copy the `data` directory to your own disk space and open the assembly file `gpuAssem.CATProduct` in STAR-CAT5.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

- Position the geometry as shown below.



The *STAR-CAT5* menu appears at the top of the screen. The STAR-CAT5 object tree is located on the left of the screen.

- Save the analysis as `gpuAssem.CATAnalysis`.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-CAT5. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-CAT5 tree and select Continuum object > Definition.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case.



3. Click OK to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-CAT5 tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the disc that is located inside the graphics card housing. The part can be selected by expanding the Links Manager.1 node in the Analysis Manager tree and selecting Link.1 > PC Tower and Graphics Card Assembly > gpuFan (Fan). Once selected, the part becomes highlighted as shown below.



3. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
4. Activate the Moving Reference Frame option.
5. Set Omega to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.
6. Set Origin to the following:
   - X: `-122.1761 mm`
   - Y: `-1.9434 mm`
   - Z: `-19.0609 mm`
7. Set Axis to the following:
   - i: `0`
   - j: `1`

- k: 0

8. Make sure that Use Assembly Coordinate System is selected.

   The completed property page is shown below:



9. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Supports* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the *Boundary Type* drop-down menu.
4. Set Name to `Case Inlet`.
5. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



6. Click OK.

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and select the face that is highlighted below.



3. Select Pressure Outlet from the *Boundary Type* drop-down menu.
4. Set Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Boundary Type* drop-down menu.
4. Set Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-CAT5 tree and select Mesh Pipeline object > Definition.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Supports* box.
3. Activate the Use Absolute Value checkbox in the *Base Size* box.
4. Set Base Size to `0.016` m.
5. Set Layer Thickness to `20`%.

The completed *Mesh Pipeline Settings* dialog is shown below.



6. Click OK.

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Fan Box`.
3. Activate the Use Absolute Value option.
4. Set Surface Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X1: `-166`
   - Y1: `4`
   - Z1: `-61`
6. Enter the following coordinates for Point 2:
   - X2: `-79`

- Y2: `-34`
- Z2: `25`

7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.

2. The *Volumetric Control Settings* dialog appears.
3. Set Name to `GPU Box`.
4. Activate the Use Absolute Value option.
5. Set Surface Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X1: `-383`
   - Y1: `16`
   - Z1: `-111`
7. Enter the following coordinates for Point 2:
   - X2: `-52`
   - Y2: `-49`
   - Z2: `49`
8. Click Create.

   The volumetric control region is shown below.



9. Click OK.

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.

The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Case Box`.
3. Activate the Use Absolute Value option.
4. Set Surface Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X1: `-383`
   - Y1: `30`
   - Z1: `-111`
6. Enter the following coordinates for Point 2:
   - X2: `43`
   - Y2: `-159`
   - Z2: `77`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Right-click the Solution folder in the STAR-CAT5 tree and select Solution Controls.



The *Solution Controls* dialog appears.

2. Click More to expand the *Solution Controls* menu.
3. Set Iterations to `300`.
4. Set Convergence to `0.01`.
5. Set Momentum Under-Relaxation Factor to `0.85`.

The completed property dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-CAT5 object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select *STAR-CAT5* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-CAT5 *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-CAT5 environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcat5\data`.

- Copy directory `data` to your own disk space and open the part file `car.CATPart` in STAR-CAT5.
- Position the geometry as shown below.



The *STAR-CAT5* menu appears at the top of the screen. The STAR-CAT5 object tree is located on the left of the screen.

- Save the analysis as `car.CATAnalysis`.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-CAT5 tree and select Continuum object > Definition.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

   The remaining default settings are kept for this case.



3. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Supports* box in the *Boundary Settings* dialog.



3. Select Symmetry Plane from the *Boundary Type* drop-down menu.
4. Set Name to `Symmetry Plane`.

   The completed dialog is shown below.



5. Click OK.

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select Add Boundary.
   The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Velocity Inlet from the *Boundary Type* drop-down menu.
4. Set Name to `Inlet`.
5. Change the units of Normal Velocity by clicking ▾ and select mph.



6. Set Normal Velocity to `70` mph.
7. Click OK.

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the back face and select the face that is highlighted below.



3. Select Pressure Outlet from the *Boundary Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case.
5. Click OK.

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the top and side walls, and select the faces that are highlighted below.



3. Select Symmetry Plane from the *Boundary Type* drop-down menu.
4. Set Name to `Boundary Symmetry`.
5. Click OK.

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the bottom face, and select the face that is highlighted below.



3. Leave the default Wall in the *Boundary Type* drop-down menu.
4. Set Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click OK.

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Default object > Definition.

   The *Boundary Settings* dialog will appear.
2. Change the Name to `Car`.
3. Click OK to close the dialog.
4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-CAT5 tree and select Mesh Pipeline object > Definition.



The *Mesh Pipeline Settings* dialog appears.

2. Activate the Use Absolute Value option in the *Base Size* box.
3. Set Base Size to `0.03` m.
4. Activate the Use Absolute Value option for Layer Thickness.
5. Set Layer Thickness to `0.03` m.

The completed *Mesh Pipeline Settings* dialog is shown below.



6. Click OK.

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Car Box`.
3. Activate the Use Absolute Value option.
4. Set Surface Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X1: `-50.00`
   - Y1: `-650.00`
   - Z1: `3432.00`
6. Enter the following coordinates for Point 2:
   - X2: `2000.00`
   - Y2: `1440.00`

- Z2: `-4567.00`

7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Volumetric Control.

   The *Volumetric Control Settings* dialog appears.
2. Set Name to `Wake Box`.
3. Activate the Use Absolute Value option.
4. Set Surface Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X1: `-50.00`
   - Y1: `-650.00`
   - Z1: `3432.00`
6. Enter the following coordinates for Point 2:
   - X2: `2000.00`
   - Y2: `3240.00`
   - Z2: `-14567.00`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-CAT5 tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



3. Set Name to `Fluid Domain Face Control`.
4. Activate the Custom Mesh Size option in the *Surface Mesh Size* box.
5. Activate the Use Absolute Value option.
6. Set Mesh Size to `0.35` m.

   The completed dialog is shown below.



7. Click OK.

Modifying Solution Controls

Modify the solution controls:

1. Right-click the Solution folder in the STAR-CAT5 tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Click More to expand the *Solution Controls* menu.
3. Set Iterations to `400`.
4. Set Convergence to `0.0001`.

The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-CAT5 object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-CAT5* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select *STAR-CAT5* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-CAT5 *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.

3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-NX CAD Client

Several tutorials have been provided for the STAR-NX CAD Client. These provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-NX environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starnx\data`.

- Copy directory `data` to your own disk space and open the part file `exhaustManifold.prt` in STAR-NX.

The geometry is shown below.



The *STAR-NX* menu appears at the top of the screen. The STAR-NX object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-NX.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-NX. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-NX tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces that are highlighted below. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Mass Flow Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Mass Flow Rate to `0.017` kg/s.

The completed dialog is shown below.



6. Click OK to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click OK to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-NX tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to 2.5%.

3. Expand the Prism Layer tab.
4. Set Layer Thickness to `15%`.

   The completed *Mesh Pipeline Settings* dialog is shown below.



5. Click OK to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Right-click the STAR-NX node in the STAR-NX tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Advanced* box.
3. Set Number of Iterations to `200`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK to close the dialog.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-NX object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Select *STAR-NX* > *Mesh* > *Generate Surface Mesh* from the menu.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-NX* > *Mesh* > *Generate Volume Mesh* from the menu.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-NX* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-NX *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the <T> key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-NX environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starnx\data`.

- Copy directory `data` to your own disk space and open the assembly file `PorousModel.prt` in STAR-NX.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

The geometry is shown below.



The *STAR-NX* menu appears at the top of the screen. The STAR-NX object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-NX.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-NX. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-NX tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-NX tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the middle part of the assembly as shown below.



3. In the *Region Settings* dialog, set the Name to `Porous Region`.
4. Activate the Porous model.
5. Within the *Porous Settings* box, select the Uni-Directional option.
6. Set Alpha Coefficient α to `25` kg/m^4 and Beta Coefficient β to `1500` kg/m^3.s.
7. Set Cross-Flow Resistance Factor to `250`.

   The completed *Region Settings* dialog is shown below.


Region Settings
Name
Porous Region
Part Selection
porousmiddle
Continuum Selection
Continuum
Porous Settings
Porous
Isotropic
Uni-Directional
α
25.0000
kg/m^4
β
1500.0000
kg/m^3.s
R
250.0000
i
0.0000
j
0.0000
k
1.0000
Use Assembly Co-ordinate System
Moving Reference Frame Settings
Moving Reference Frame
Energy Source Settings
Energy Source
Advanced
Exclude From Analysis
OK
Cancel

8. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-NX automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



6. Click OK.

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click OK.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-NX tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.0025` m.
5. Expand the Prism Layer tab.
6. Select the Absolute radio button for Layer Thickness.
7. Set Layer Thickness to `0.001` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
porousmiddle
porousoutlet
porousin
Mesh Type
Polyhedral
Trimmer
Base Size
0.0050
m
Relative to Geometry
Absolute
Surface Mesh Size
100.0000
%
Relative to Base Size
Absolute
Prism Layer
Enable
0.0010
m
Relative to Base Size
Absolute
2
1.5000
Advanced Settings
OK
Cancel

8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Right-click the STAR-NX node in the STAR-NX tree and select Set Solution Controls.



   The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `400`.
4. Set *Solution Accuracy* to First Order.

   The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-NX object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-NX* > *Mesh* > *Generate Surface Mesh* from the menu.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-NX* > *Mesh* > *Generate Volume Mesh* from the menu.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-NX* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-NX *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-NX environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starnx\data`.

- Copy directory `data` to your own disk space and open the assembly file `piccoloAssem.prt` in STAR-NX.

  The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

The geometry is shown below.



The *STAR-NX* menu appears at the top of the screen. The STAR-NX object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-NX.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-NX.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-NX tree and select Edit.



The *Continuum Settings* dialog appears.

2. In the *Continuum Settings* dialog, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Expand the Initialization box by clicking .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.


Continuum Settings
Name
Air
Material
Material Type
Gas
Liquid
Solid
Material Name
Air
Flow Characteristics
Laminar Flow
Turbulent Flow
Equation of State
Ideal Gas
Constant Density
Gravity
Enable
Optional Physics Models
Passive Scalar
Temperature
Material Properties
Reference Properties
Initialization
U
0.0000
V
0.0000
W
0.0000
m/s
0.0000
Pa
400.0000
K
Int
0.0100
len
0.0100
m
OK
Cancel

6. Click OK.

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-NX tree and select Add Continuum.



   The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material Type* menu, and select Aluminum Alloy from the drop-down list.
4. Expand the Initialization box by clicking .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click OK.

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-NX tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the wing skin as shown below. Once selected, the part becomes highlighted.



3. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
4. Set Continuum Selection to Al from the selection box.



This action assigns the aluminum alloy that you created earlier as the material for this part.

The completed dialog is shown below.



5. Click OK.

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-NX tree and select Rename.



2. Enter `Internal Air` as the new name.
3. Click a blank area to accept the name change.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.



   The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces highlighted below. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `90` m/s.
6. Set Temperature to `460` K.
7. Set Turbulence Intensity to `0.135`.
8. Set Turbulent Length Scale to `2.2e-005` m.

   The completed dialog is shown below. You may need to use the cursor to see the full entry in the Turbulent Length Scale field.



9. Click OK.

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are suitable for this tutorial.
5. Click OK.

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the four faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries.



3. Leave the selection as Wall in the *Type* drop-down menu.
4. Set Name to `Piccolo Wall`.
5. Set Thermal Specification to Temperature.



6. Set Temperature to `460` K.
7. Click OK.

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

The *Boundary Settings* dialog appears.

2. Select the two faces highlighted below, one on either side of the internal air region.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Air Symmetry`.
5. Click OK.

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the two faces highlighted below, one on either side of the geometry.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Solid Symmetry`.
5. Click OK.

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Edit.

   The *Boundary Settings* dialog appears.
2. Set Thermal Specification to Convection.



3. Set Heat Transfer Coefficient HTC to 100 W/m^2.K
4. Set Ambient Temperature to 228 K.

   The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-NX tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.007` m.
5. Expand the Prism Layer tab.
6. Set Layer Thickness to `10`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
internalAir
wingSkin
Mesh Type
Polyhedral
Trimmer
Base Size
0.0070
m
Relative to Geometry
Absolute
Surface Mesh Size
100.0000
%
Relative to Base Size
Absolute
Prism Layer
Enable
10.0000
%
Relative to Base Size
Absolute
2
1.5000
Advanced Settings
OK
Cancel

7. Click OK.

Setting the Coincident Faces Detection Tolerance

The geometry that is used in this tutorial has large changes in feature size and requires you to reduce the Coincident Faces Detection Tolerance.

1. Right-click STAR-NX and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.

2. Set Coincident Faces Detection Tolerance to `0.0001`.



3. Click OK.

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Inlet Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-100.00 mm`
   - Y Coordinate of Point 1 Y_1: `21.00 mm`
   - Z Coordinate of Point 1 Z_1: `-20.00 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-90.00 mm`
   - Y Coordinate of Point 2 Y_2: `10.00 mm`
   - Z Coordinate of Point 2 Z_2: `60.00 mm`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.
2. Set the Name to `Mid Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-111.00 mm`
   - Y Coordinate of Point 1 Y_1: `34.00 mm`
   - Z Coordinate of Point 1 Z_1: `-20.00 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-64.00 mm`
   - Y Coordinate of Point 2 Y_2: `-2.00 mm`
   - Z Coordinate of Point 2 Z_2: `60.00 mm`
7. Click Create.

The volumetric control region is shown below.



8. Click OK.
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-NX tree and select Add Part Mesh Control.



The *Part Mesh Control Settings* dialog appears.

2. Select the solid wing region on the CAD model. The selection is added to the *Part Selection* box.
3. Set Name to `Solid Region`.
4. In the Prism Layer box, set Prism layer option to Disable using the drop-down menu.



The completed dialog is shown below.



5. Click OK.

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-NX tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



3. Set Name to `Inlet Refinement`.
4. In the Custom Surface Mesh Size box, activate the Enable option.
5. Set Custom Mesh Size Option to Absolute.
6. Set Mesh Size to `0.00025` m.
7. In the Prism Layer box, set Prism Layer Option to Use Custom Values using the drop-down menu.
8. Set Layer Thickness Option to Absolute.
9. Set Layer Thickness to `0.0002` m.

   The completed dialog is shown below.



10. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Right-click the STAR-NX node in the STAR-NX tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Advanced* box.
3. Set Number of Iterations to `500`.
4. Set Convergence Level to `0.01`.
5. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

The completed dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-NX object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-NX* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-NX* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-NX* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-NX *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-NX environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starnx\data`.

- Copy the `data` directory to your own disk space and open the assembly file `gpuAssem.prt` in STAR-NX.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

The geometry is shown below.



The *STAR-NX* menu appears at the top of the screen. The STAR-NX object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-NX.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-NX. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-NX tree and select Edit.



   The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case.



3. Click OK to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-NX tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the disc that is located inside the graphics card housing. The part can be selected from within the PC tower by right-clicking on the disc and selecting Select from List... > Solid Body in GPUBLOWERFAN from the *Quick Pick* dialog. Once selected, the part becomes highlighted as shown below.



3. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
4. Activate the Moving Reference Frame option.
5. Set Rotational Velocity, ω to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.
6. Set Coordinates of Axis Origin to the following:
   - X : `-122.1761`
   - Y : `-1.9434`

- Z: -19.0609

7. Set Component of Flow Direction to the following:
   - i : 0
   - j : 1
   - k : 0
8. Make sure that Use Assembly Coordinate System is selected.

   The completed property page is shown below:



9. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Case Inlet`.
5. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



6. Click OK.

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-NX tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.016` m.
5. Expand the Prism Layer tab.
6. Set Layer Thickness to `20`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
pcTowerCase
gpuBlowerFan
Mesh Type
Polyhedral
Trimmer
Base Size
0.0160
m
Relative to Geometry
Absolute
Surface Mesh Size
100.0000
%
Relative to Base Size
Absolute
Prism Layer
Enable
20.0000
%
Relative to Base Size
Absolute
2
1.5000
Advanced Settings
OK
Cancel

7. Click OK.

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Fan Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-166 mm`
   - Y Coordinate of Point 1 Y_1: `4 mm`
   - Z Coordinate of Point 1 Z_1: `-61 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-79 mm`
   - Y Coordinate of Point 2 Y_2: `-34 mm`
   - Z Coordinate of Point 2 Z_2: `25 mm`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.
2. The *Volumetric Control Settings* dialog appears.
3. Set Name to `GPU Box`.

4. Select the Absolute option.
5. Set Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383 mm`
   - Y Coordinate of Point 1 Y_1: `16 mm`
   - Z Coordinate of Point 1 Z_1: `-111 mm`
7. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-52 mm`
   - Y Coordinate of Point 2 Y_2: `-49 mm`
   - Z Coordinate of Point 2 Z_2: `49 mm`
8. Click Create.

   The volumetric control region is shown below.



9. Click OK.

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Case Box`.

3. Select the Absolute option.
4. Set Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383 mm`
   - Y Coordinate of Point 1 Y_1: `30 mm`
   - Z Coordinate of Point 1 Z_1: `-111 mm`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `43 mm`
   - Y Coordinate of Point 2 Y_2: `-159 mm`
   - Z Coordinate of Point 2 Z_2: `77 mm`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Right-click the STAR-NX node in the STAR-NX tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `300`.
4. Set Convergence Level to `0.01`.
5. Set Momentum Under-Relaxation Factor to `0.85`.

   The completed property dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-NX object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-NX* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-NX* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select *STAR-NX* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-NX *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-NX environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starnx\data`.

- Copy directory `data` to your own disk space and open the part file `car.prt` in STAR-NX.

The geometry is shown below.



The *STAR-NX* menu appears at the top of the screen. The STAR-NX object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree icon.

If these items are not visible, read the help on Launching STAR-NX.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-NX tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

   The remaining default settings are kept for this case.



3. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Symmetry Plane`.

   The completed dialog is shown below.



5. Click OK.

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Inlet`.
5. Change the units of Normal Velocity by clicking and select mph.



6. Set Normal Velocity to `70` mph.
7. Click OK.

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Rotate the view to display the back face and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case.
5. Click OK.

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the top and side walls, and select the faces that are highlighted below.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Boundary Symmetry`.
5. Click OK.

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the bottom face, and select the face that is highlighted below.



3. Leave the default Wall in the *Type* drop-down menu.
4. Set Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click OK.

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Rename.
2. Enter `Car` as the new name.
3. You may need to click on a blank area of the workspace to accept the name change.
4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-NX tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Select the Absolute option in the *Base Size* box.
3. Set Base Size to `0.03` m.
4. Expand the Prism Layer tab.
5. Select the Absolute radio button for Layer Thickness.
6. Set Layer Thickness to `0.03` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.



7. Click OK.

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Car Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50.00` mm
   - Y Coordinate of Point 1 Y_1 : `-650.00` mm
   - Z Coordinate of Point 1 Z_1 : `3432.00` mm
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000.00` mm
   - Y Coordinate of Point 2 Y_2 : `1440.00` mm
   - Z Coordinate of Point 2 Z_2 : `-4567.00` mm
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.

Using the *See-Thru* view function in NX provides the following view of the volumetric control.



8. Click OK.

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.
2. Set Name to `Wake Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50.00` mm
   - Y Coordinate of Point 1 Y_1 : `-650.00` mm
   - Z Coordinate of Point 1 Z_1 : `3432.00` mm
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000.00` mm
   - Y Coordinate of Point 2 Y_2 : `3240.00` mm
   - Z Coordinate of Point 2 Z_2 : `-14567.00` mm

7. Click Create.

   The volumetric control region is shown below.

   Using the See-Thru view function in NX provides the following view of the volumetric control.



8. Click OK.

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-NX tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



3. Set Name to `Fluid Domain Face Control`.
4. Activate the Enable option in the *Custom Mesh Size* box.
5. Select the Absolute option.
6. Set Mesh Size to `0.35` m.

   The completed dialog is shown below.



7. Click OK.

Modifying Solution Controls

Modify the solution controls:

1. Right-click the STAR-NX node in the STAR-NX tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `400`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-NX object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-NX* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-NX* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select *STAR-NX* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-NX *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.
3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below.



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-Pro/E CAD Client

Several tutorials have been provided for the STAR-Pro/E CAD Client. These provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Pro/E environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starproe\data`.

- Copy directory `data` to your own disk space and open the part file `exhaustmanifold.prt` in STAR-Pro/E.

The geometry is shown below.



The *STAR-Pro/E* menu appears at the top of the screen. The STAR-Pro/E object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Pro/E.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Pro/E. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-Pro/E tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Pro/E face selection mechanism.
3. On the CAD model, select the three faces that are highlighted below. To select multiple faces, hold down the `<Ctrl>` key. These selections are reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog window.
5. Select Mass Flow Inlet from the Type drop-down menu.
6. Set Name to `Inlet`.
7. Set Mass Flow Rate to `0.017` kg/s.

   The completed dialog is shown below.



8. Click OK to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click OK to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Pro/E tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to `2.5`%.
3. Expand the Prism Layer tab.
4. Set Layer Thickness to `15`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.



5. Click OK to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Right-click the STAR-Pro/E node in the STAR-Pro/E tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Advanced* box.
3. Set Number of Iterations to `200`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK to close the dialog.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Pro/E object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Surface Mesh* from the menu.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Volume Mesh* from the menu.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-Pro/E* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-Pro/E *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the `<T>` key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Pro/E environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starproe\data`.

- Copy directory `data` to your own disk space and open the assembly file `porousmodel.asm` in STAR-Pro/E.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

The geometry is shown below.



The *STAR-Pro/E* menu appears at the top of the screen. The STAR-Pro/E object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Pro/E.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Pro/E. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Pro/E tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-Pro/E tree and select Add Region.



The *Region Settings* dialog appears.

2. Click (Select Part) to activate the Pro/E part selection mechanism.
3. Select the middle part of the assembly as shown below.



4. Click OK to close the *Select* dialog.
5. In the *Region Settings* dialog, set the Name to `Porous Region`.
6. Activate the Porous model.

7. Within the *Porous Settings* box, select the Uni-Directional option.
8. Set Alpha α to `25` kg/m^4 and Beta β to `1500` kg/m^3.s.
9. Set Cross-Flow Resistance Factor to `250`.

   The completed *Region Settings* dialog is shown below.


Region Settings
Name
Porous Region
Part Selection
part_id98
Continuum Selection
Continuum
Porous Settings
Porous
Isotropic
Uni-directional
α
25.000000
kg/m^4
β
1500.000000
kg/m^3.s
250.000000
0.000000
0.000000
1.000000
Use Assembly Coordinate System
Moving Reference Frame Settings
Moving Reference Frame
Advanced
Exclude From Analysis
OK
Cancel

10. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-Pro/E automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



   The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Pro/E face selection mechanism.
3. Select the face that is highlighted below. The selection is reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog window
5. Select Velocity Inlet from the Type drop-down menu.
6. Set Name to `Inlet`.
7. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



8. Click OK.

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click OK.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Pro/E tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.0025` m.
5. Expand the Prism Layer tab.
6. Select the Absolute radio button for Layer Thickness.
7. Set Layer Thickness to `0.001` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
part_id93
part_id98
part_id103
Mesh Type
Polyhedral
Trimmer
Base Size
0.002500
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000
%
Relative to Base Size
Absolute
Prism Layer
Enable
0.001000
m
Relative to Base Size
Absolute
2
1.500000
Advanced Settings
OK
Cancel

8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Right-click the STAR-Pro/E node in the STAR-Pro/E tree and select Set Solution Controls.



   The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `400`.
4. Set *Solution Accuracy* to First Order.

   The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Pro/E object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Surface Mesh* from the menu.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Volume Mesh* from the menu.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-Pro/E* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-Pro/E *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Pro/E environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starproe\data`.

- Copy directory `data` to your own disk space and open the assembly file `piccoloassem.asm` in STAR-Pro/E.

  The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

The geometry is shown below.



The *STAR-Pro/E* menu appears at the top of the screen. The STAR-Pro/E object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Pro/E.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-Pro/E.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-Pro/E tree and select Edit.



The *Continuum Settings* dialog appears.

2. In the *Continuum Settings* dialog, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Expand the Initialization box by clicking ▶ .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.


Continuum Settings
Name
Air
Material
Material Type
Gas
Liquid
Solid
Material Name
Air
Flow Characteristics
Laminar Flow
Turbulent Flow
Equation of State
Ideal Gas
Constant Density
Gravity
Enable
Optional Physics Models
Passive Scalar
Temperature
Material Properties
Reference Properties
Initialization
U
0.000000
m/s
V
0.000000
W
0.000000
0.000000
Pa
400.000000
K
Int
0.010000
len
0.010000
m
OK
Cancel

6. Click OK.

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-Pro/E tree and select Add Continuum.



   The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material Type* menu, and select Aluminum Alloy from the drop-down list.
4. Expand the Initialization box by clicking ▶ .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click OK.

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-Pro/E tree and select Add Region.



The *Region Settings* dialog appears.

2. Click (Select Part) to activate the Pro/E part selection mechanism.

3. Select the wing skin as shown below. Once selected, the part becomes highlighted.



4. Click OK to close the *Select* dialog.
5. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
6. Set Continuum Selection to Al from the selection box.



This action assigns the aluminum alloy that you created earlier as the material for this part.

The completed dialog is shown below.



7. Click OK.

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-Pro/E tree and select Rename.



2. Enter `Internal Air` as the new name.
3. Click a blank area to accept the name change.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Pro/E face selection mechanism.
3. On the CAD model, select the three faces highlighted below. To select multiple faces hold down the `<Ctrl>` key. These selections are reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog.
5. Select Velocity Inlet from the Type drop-down menu.
6. Set Name to `Inlet`.
7. Set Normal Velocity to `90` m/s.
8. Set Temperature to `460` K.
9. Set Turbulence Intensity to `0.135`.
10. Set Turbulent Length Scale to `2.2e-005` m.

The completed dialog is shown below.



11. Click OK.

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are suitable for this tutorial.
5. Click OK.

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Use the Pro/E face selection mechanism to select the six faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries. To select multiple faces hold down the `<Ctrl>` key.



3. Leave the selection as Wall in the *Type* drop-down menu.
4. Set Name to `Piccolo Wall`.
5. Set Thermal Specification to Temperature.



6. Set Temperature to `460` K.
7. Click OK.

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.
   The *Boundary Settings* dialog appears.

2. Use the Pro/E face selection mechanism to select the two faces highlighted below, one on either side of the internal air region. To select multiple faces hold down the `<Ctrl>` key.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Air Symmetry`.
5. Click OK.

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Pro/E face selection mechanism to select the two faces highlighted below, one on either side of the geometry. To select multiple faces hold down the `<Ctrl>` key.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Solid Symmetry`.
5. Click OK.

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Edit.

   The *Boundary Settings* dialog appears.
2. Set Thermal Specification to Convection.



3. Set Heat Transfer Coefficient HTC to `100` W/m^2.K
4. Set Ambient Temperature to `228` K.

   The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Pro/E tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.007` m.
5. Expand the Prism Layer tab.
6. Set Prism Layer Thickness to `10`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
part_id20
part_id44
Mesh Type
Polyhedral
Trimmer
Base Size
0.007000000
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Base Size
Absolute
Prism Layer
Enable
10.000000000
%
Relative to Base Size
Absolute
2
1.500000
Advanced Settings
OK
Cancel

7. Click OK.

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Inlet Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-100.00`
   - Y Coordinate of Point 1 Y_1: `21.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-90.00`
   - Y Coordinate of Point 2 Y_2: `10.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.
2. Set the Name to `Mid Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-111.00`
   - Y Coordinate of Point 1 Y_1: `34.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-64.00`
   - Y Coordinate of Point 2 Y_2: `-2.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

The volumetric control region is shown below.



8. Click OK.
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Pro/E tree and select Add Part Mesh Control.



The *Part Mesh Control Settings* dialog appears.

2. Click (Select Part) to activate the Pro/E part selection mechanism.
3. Select the solid wing region on the CAD model. The selection is added to the *Part Selection* box.
4. Click OK to close the *Select* dialog.
5. Set Name to `Solid Region`.
6. In the Prism Layer box, set Prism layer option to Disable using the drop-down menu.



The completed dialog is shown below.



7. Click OK.

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Pro/E tree and select Add Face Mesh Control.



   The *Face Mesh Control Settings* dialog appears.

2. Click (Select Face) to activate the Pro/E face selection mechanism.
3. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



4. Click OK to close the *Surface Sets* dialog.
5. Set Name to `Inlet Refinement`.
6. In the Custom Surface Mesh Size box, activate the Enable option.
7. Set Custom Mesh Size Option to Absolute.
8. Set Mesh Size to `0.00025` m.
9. In the Prism Layer box, set Prism Layer Option to Use Custom Values using the drop-down menu.
10. Set Layer Thickness Option to Absolute.
11. Set Layer Thickness to `0.0002` m.

The completed dialog is shown below.



12. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Right-click the STAR-Pro/E node in the STAR-Pro/E tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Advanced* box.
3. Set Number of Iterations to `500`.
4. Set Convergence Level to `0.01`.
5. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

   The completed dialog is shown below.



6. Click OK.

7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Pro/E object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select *STAR-Pro/E* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-Pro/E *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Pro/E environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starproe\data`.

- Copy the `data` directory to your own disk space and open the assembly file `gpuAssem.asm` in STAR-Pro/E.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

The geometry is shown below.



The *STAR-Pro/E* menu appears at the top of the screen. The STAR-Pro/E object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Pro/E.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Pro/E. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Pro/E tree and select Edit.



   The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case.



3. Click OK to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-Pro/E tree and select Add Region.



The *Region Settings* dialog appears.

2. Click (Select Part) to activate the Pro/E part selection mechanism.
3. Select the disc that is located inside the graphics card housing. You can select this part by moving the mouse over the fan area and right-clicking once. This action highlights the fan.Left-click to select the fan. Once selected, the part becomes highlighted as shown below.



4. Click OK to close the *Select* dialog.
5. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
6. Activate the Moving Reference Frame option.
7. Set Rotational Velocity, ω to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.
8. Set Coordinates of Axis Origin to the following:
   - X: `-122.1761`
   - Y: `-1.9434`
   - Z: `-19.0609`
9. Set Component of Flow Direction to the following:

- i : 0
- j : 1
- k : 0

10. Make sure that Use Assembly Coordinate System is selected.

    The completed property page is shown below:


Region Settings
Name
Rotating Region
Part Selection
part_id20
Continuum Selection
Continuum
Porous Settings
Porous
Moving Reference Frame Settings
Moving Reference Frame
-2800.000000
rpm
X
-122.176100
Y
-1.943400
Z
-19.060900
i
0.000000
j
1.000000
k
0.000000
Use Assembly Coordinate System
Advanced
Exclude From Analysis
OK
Cancel

11. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Click (Select Part) to activate the Pro/E part selection mechanism.
3. On the CAD model, select the face that is highlighted below. The selection is reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog.
5. Select Velocity Inlet from the *Type* drop-down menu.
6. Set Name to `Case Inlet`.
7. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



8. Click OK.

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Pro/E tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.016` m.
5. Expand the Prism Layer tab.
6. Set Prism Layer Thickness to `20`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
part_id20
part_id82
Mesh Type
Polyhedral
Trimmer
Base Size
0.016000000
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Base Size
Absolute
Prism Layer
Enable
20.000000000
%
Relative to Base Size
Absolute
2
1.500000
Advanced Settings
OK
Cancel

7. Click OK.

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Fan Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-166`
   - Y Coordinate of Point 1 Y_1: `4`
   - Z Coordinate of Point 1 Z_1: `-61`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-79`
   - Y Coordinate of Point 2 Y_2: `-34`
   - Z Coordinate of Point 2 Z_2: `25`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.
2. The *Volumetric Control Settings* dialog appears.
3. Set Name to `GPU Box`.

4. Select the Absolute option.
5. Set Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `16`
   - Z Coordinate of Point 1 Z_1: `-111`
7. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-52`
   - Y Coordinate of Point 2 Y_2: `-49`
   - Z Coordinate of Point 2 Z_2: `49`
8. Click Create.

   The volumetric control region is shown below.



9. Click OK.

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Case Box`.

3. Select the Absolute option.
4. Set Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `30`
   - Z Coordinate of Point 1 Z_1: `-111`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `43`
   - Y Coordinate of Point 2 Y_2: `-159`
   - Z Coordinate of Point 2 Z_2: `77`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Right-click the STAR-Pro/E node in the STAR-Pro/E tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `300`.
4. Set Convergence Level to `0.01`.
5. Set Momentum Under-Relaxation Factor to `0.85`.

The completed property dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Pro/E object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select *STAR-Pro/E* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-Pro/E *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Pro/E environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starproe\data`.

- Copy directory `data` to your own disk space and open the part file `car.prt` in STAR-Pro/E.

The geometry is shown below.



The *STAR-Pro/E* menu appears at the top of the screen. The STAR-Pro/E object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Pro/E.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-Pro/E tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

   The remaining default settings are kept for this case.



3. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Pro/E face selection mechanism.
3. On the CAD model, select the face that is highlighted below. The selection is reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog window.
5. Select Symmetry Plane from the *Type* drop-down menu.
6. Set Name to `Symmetry Plane`.

   The completed dialog is shown below.



7. Click OK.

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Inlet`.
5. Change the units of Normal Velocity by clicking and select mph.



6. Set Normal Velocity to `70` mph.
7. Click OK.

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the back face and use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case.
5. Click OK.

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the top and side walls, and use the Pro/E face selection mechanism to select the faces that are highlighted below. To select multiple faces, hold down the `<Ctrl>` key while selecting the faces.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Boundary Symmetry`.
5. Click OK.

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the bottom face, and use the Pro/E face selection mechanism to select the face that is highlighted below.



3. Leave the default Wall in the *Type* drop-down menu.
4. Set Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click OK.

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Rename.
2. Enter `Car` as the new name.
3. You may need to click on a blank area of the workspace to accept the name change.

4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Pro/E tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Select the Absolute option in the *Base Size* box.
3. Set Base Size to `0.03` m.
4. Expand the Prism Layer tab.
5. Select the Absolute radio button for Layer Thickness.
6. Set Layer Thickness to `0.03` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.



7. Click OK.

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Car Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `1440`
   - Z Coordinate of Point 2 Z_2 : `-4567`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.
2. Set Name to `Wake Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `3240`
   - Z Coordinate of Point 2 Z_2 : `-14567`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Pro/E tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Click (Select Face) to activate the Pro/E face selection mechanism.
3. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



4. Click OK to close the *Surface Sets* dialog window.
5. Set Name to `Fluid Domain Face Control`.
6. Activate the Enable option in the *Custom Mesh Size* box.
7. Select the Absolute option.
8. Set Mesh Size to `0.35` m.

   The completed dialog is shown below.



9. Click OK.

Modifying Solution Controls

Modify the solution controls:

1. Right-click the STAR-Pro/E node in the STAR-Pro/E tree and select Set Solution Controls.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.

3. Set Number of Iterations to `400`.
4. Set Convergence Level to `0.0001`.

   The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Pro/E object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Surface Mesh* from the menu.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Select *STAR-Pro/E* > *Mesh* > *Generate Volume Mesh* from the menu.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select *STAR-Pro/E* > *Solution* > *Build and Run* from the menu.



The *Residuals* plot appears in the STAR-Pro/E *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.
3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below.



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-Creo CAD Client

Several tutorials are provided for the STAR-Creo CAD Client. These tutorials provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Creo environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcreo\data`.

- Copy directory `data` to your own disk space and open the part file `exhaustmanifold.prt` in STAR-Creo.

The geometry is shown below.



The *STAR-Creo* ribbon appears at the top of the screen. The STAR-Creo object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Creo.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Creo. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-Creo tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Creo face selection mechanism.
3. On the CAD model, select the three faces that are highlighted below. To select multiple faces, hold down the `<Ctrl>` key. These selections are reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog window.
5. Select Mass Flow Inlet from the Type drop-down menu.
6. Set Name to `Inlet`.
7. Set Mass Flow Rate to `0.017` kg/s.

   The completed dialog is shown below.



8. Click OK to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click OK to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Creo tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to 2.5%.
3. Expand the Prism Layer tab.
4. Set Layer Thickness to 15%.



5. Click OK to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Click Solution Controls on the STAR-Creo ribbon.



The *Solution Controls* dialog appears.

2. Expand the *Advanced* box.
3. Set Number of Iterations to `200`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK to close the dialog.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Creo object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Click Surface Mesh from the ribbon.

The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Creo *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the <T> key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Creo environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcreo\data`.

- Copy directory `data` to your own disk space and open the assembly file `porousmodel.asm` in STAR-Creo.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

The geometry is shown below.



The *STAR-Creo* ribbon appears at the top of the screen. The STAR-Creo object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Creo.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Creo. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Creo tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-Creo tree and select Add Region.



The *Region Settings* dialog appears.

2. Click (Select Part) to activate the Creo part selection mechanism.
3. Select the middle part of the assembly as shown below.



4. Click OK to close the *Select* dialog.
5. In the *Region Settings* dialog, set the Name to `Porous Region`.

6. Activate the Porous model.
7. Within the *Porous Settings* box, select the Uni-Directional option.
8. Set Alpha α to `25` kg/m^4 and Beta β to `1500` kg/m^3.s.
9. Set Cross-Flow Resistance Factor to `250`.

   The completed *Region Settings* dialog is shown below.


Region Settings
Name
Porous Region
Part Selection
part_id78
Continuum Selection
Continuum
Porous Settings
Porous
Isotropic
Uni-directional
α
25.000000
kg/m^4
β
1500.000000
kg/m^3.s
250.000000
0.000000
0.000000
1.000000
Use Assembly Coordinate System
Moving Reference Frame Settings
Moving Reference Frame
Advanced
Exclude From Analysis
OK
Cancel

10. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-Creo automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



   The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Creo face selection mechanism.
3. Select the face that is highlighted below. The selection is reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog window
5. Select Velocity Inlet from the Type drop-down menu.
6. Set Name to `Inlet`.
7. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



8. Click OK.

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click OK.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Creo tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.0025` m.
5. Expand the Prism Layer tab.
6. Select the Absolute radio button for Layer Thickness.
7. Set Layer Thickness to `0.001` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
part_id73
part_id78
part_id83
Mesh Type
Polyhedral
Trimmer
Base Size
0.002500000
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Base Size
Absolute
Prism Layer
Enable
0.001000000
m
Relative to Base Size
Absolute
2
1.500000
Advanced Settings
OK
Cancel

8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Click Solution Controls on the STAR-Creo ribbon.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `400`.
4. Set *Solution Accuracy* to First Order.

The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Creo object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.

The surface mesh is shown below.

Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Creo *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Creo environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcreo\data`.

- Copy directory `data` to your own disk space and open the assembly file `piccoloassem.asm` in STAR-Creo.

  The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

The geometry is shown below.



The *STAR-Creo* ribbon appears at the top of the screen. The STAR-Creo object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Creo.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-Creo.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-Creo tree and select Edit.



   The *Continuum Settings* dialog appears.

2. In the *Continuum Settings* dialog, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Expand the Initialization box by clicking ▶.
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.


Continuum Settings
Name
Air
Material
Material Type
Gas
Liquid
Solid
Material Name
Air
Flow Characteristics
Laminar Flow
Turbulent Flow
Equation of State
Ideal Gas
Constant Density
Gravity
Enable
Optional Physics Models
Passive Scalar
Temperature
Material Properties
Reference Properties
Initialization
U
0.000000
m/s
V
0.000000
W
0.000000
0.000000
Pa
400.000000
K
Int
0.010000
len
0.010000
m
OK
Cancel

6. Click OK.

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-Creo tree and select Add Continuum.



   The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material Type* menu, and select Aluminum Alloy from the drop-down list.
4. Expand the Initialization box by clicking ▶.
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click OK.

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-Creo tree and select Add Region.



The *Region Settings* dialog appears.

2. Click (Select Part) to activate the Creo part selection mechanism.
3. Select the wing skin as shown below. Once selected, the part becomes highlighted.



4. Click OK to close the *Select* dialog.
5. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
6. Set Continuum Selection to Al from the selection box.



This action assigns the aluminum alloy that you created earlier as the material for this part.

The completed dialog is shown below.



7. Click OK.

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-Creo tree and select Rename.



2. Enter `Internal Air` as the new name.
3. Click a blank area to accept the name change.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.



   The *Boundary Settings* dialog appears.
2. Click (Select Face) to activate the Creo face selection mechanism.
3. On the CAD model, select the three faces highlighted below. To select multiple faces hold down the `<Ctrl>` key. These selections are reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog.
5. Select Velocity Inlet from the Type drop-down menu.
6. Set Name to `Inlet`.
7. Set Normal Velocity to `90` m/s.
8. Set Temperature to `460` K.
9. Set Intensity to `0.135`.
10. .Set Length Scale to `2.2e-005` m.

The completed dialog is shown below.



11. Click OK.

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are suitable for this tutorial.
5. Click OK.

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Use the Creo face selection mechanism to select the faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries. To select multiple faces hold down the `<Ctrl>` key.



3. Leave the selection as Wall in the *Type* drop-down menu.
4. Set Name to `Piccolo Wall`.
5. Set Thermal Specification to Temperature.



6. Set Temperature to `460` K.
7. Click OK.

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Creo face selection mechanism to select the two faces highlighted below, one on either side of the internal air region.To select multiple faces hold down the `<Ctrl>` key.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Air Symmetry`.
5. Click OK.

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Creo face selection mechanism to select the two faces highlighted below, one on either side of the geometry. To select multiple faces hold down the `<Ctrl>` key.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Solid Symmetry`.
5. Click OK.

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Edit.

   The *Boundary Settings* dialog appears.

2. Set Thermal Specification to Convection.



3. Set Heat Transfer Coefficient HTC to `100` W/m^2.K
4. Set Temperature to `228` K.

   The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Creo tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.007` m.
5. Expand the Prism Layer tab.
6. Set Prism Layer Thickness to `10`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
part_id209
part_id214
Mesh Type
Polyhedral
Trimmer
Base Size
0.007000000
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Base Size
Absolute
Prism Layer
Enable
10.000000000
%
Relative to Base Size
Absolute
2
1.500000
Advanced Settings
OK
Cancel

7. Click OK.

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Inlet Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-100.00`
   - Y Coordinate of Point 1 Y_1: `21.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-90.00`
   - Y Coordinate of Point 2 Y_2: `10.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.

   The *Volumetric Control* dialog appears.
2. Set the Name to `Mid Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-111.00`
   - Y Coordinate of Point 1 Y_1: `34.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-64.00`
   - Y Coordinate of Point 2 Y_2: `-2.00`
   - Z Coordinate of Point 2 Z_2: `60.00`

7. Click Create.

   The volumetric control region is shown below.



8. Click OK.
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Creo tree and select Add Part Mesh Control.



The *Part Mesh Control Settings* dialog appears.

2. Click (Select Part) to activate the Creo part selection mechanism.
3. Select the solid wing region on the CAD model. The selection is added to the *Part Selection* box.
4. Click OK to close the *Select* dialog.
5. Set Name to `Solid Region`.
6. In the Prism Layer box, set Prism layer option to Disable using the drop-down menu.



The completed dialog is shown below.



7. Click OK.

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Creo tree and select New Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Click (Select Face) to activate the Creo face selection mechanism.
3. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



4. Click OK to close the *Surface Sets* dialog.
5. Set Name to `Inlet Refinement`.
6. In the Custom Surface Mesh Size box, activate the Enable option.
7. Set Custom Mesh Size Option to Absolute.
8. Set Custom Surface Mesh Size to `0.00025` m.
9. In the Prism Layer box, set Custom Prism Layer Option to Use Custom Values using the drop-down menu.
10. Set Layer Thickness Option to Absolute.
11. Set Layer Thickness to `0.0002` m.

    The completed dialog is shown below.



12. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Click Solution Controls on the STAR-Creo ribbon.



The *Solution Controls* dialog appears.

2. Expand the *Advanced* box.
3. Set Number of Iterations to `500`.
4. Set Convergence Level to `0.01`.
5. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

   The completed dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Creo object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.

The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Creo *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Creo environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcreo\data`.

- Copy the `data` directory to your own disk space and open the assembly file `gpuAssem.asm` in STAR-Creo.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

The geometry is shown below.



The *STAR-Creo* ribbon appears at the top of the screen. The STAR-Creo object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Creo.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Creo. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Creo tree and select Edit.



   The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case.



3. Click OK to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-Creo tree and select Add Region.



The *Region Settings* dialog appears.

2. Click (Select Part) to activate the Creo part selection mechanism.
3. Select the disc that is located inside the graphics card housing. You can select this part by moving the mouse over the fan area and right-clicking once. This action highlights the fan. Left-click to select the fan. Once selected, the part becomes highlighted as shown below.



4. Click OK to close the *Select* dialog.
5. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
6. Activate the Moving Reference Frame option.
7. Set Rotational Velocity, ω to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.
8. Set Coordinates of Axis Origin to the following:
   - X: `-122.1761`

- Y : -1.9434
- Z : -19.0609

9. Set Component of Flow Direction to the following:
   - i : 0
   - j : 1
   - k : 0
10. Make sure that Use Assembly Coordinate System is selected.

    The completed property page is shown below:


Region Settings
Name
Rotating Region
Part Selection
part_id1284
Continuum Selection
Continuum
Porous Settings
Porous
Moving Reference Frame Settings
Moving Reference Frame
-2800.000000
rpm
X
-122.176100
Y
-1.943400
Z
-19.060900
i
0.000000
j
1.000000
k
0.000000
Use Assembly Coordinate System
Energy Source Settings
Energy Source
Advanced
Exclude From Analysis
OK
Cancel

11. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



   The *Boundary Settings* dialog appears.

2. Click (Select Part) to activate the Creo part selection mechanism.
3. On the CAD model, select the face that is highlighted below. The selection is reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Select Velocity Inlet from the *Type* drop-down menu.
5. Set Name to `Case Inlet`.
6. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



7. Click OK.

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Creo tree and select Edit.



   The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.016` m.
5. Expand the Prism Layer tab.
6. Set Prism Layer Thickness to `20`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
part_id1330
part_id1335
Mesh Type
Polyhedral
Trimmer
Base Size
0.016000000
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Base Size
Absolute
Prism Layer
Enable
20.000000000
%
Relative to Base Size
Absolute
2
1.500000
Advanced Settings
OK
Cancel

7. Click OK.

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Fan Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-166`
   - Y Coordinate of Point 1 Y_1: `4`
   - Z Coordinate of Point 1 Z_1: `-61`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-79`
   - Y Coordinate of Point 2 Y_2: `-34`
   - Z Coordinate of Point 2 Z_2: `25`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.
2. The *Volumetric Control Settings* dialog appears.
3. Set Name to `GPU Box`.

4. Select the Absolute option.
5. Set Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `16`
   - Z Coordinate of Point 1 Z_1: `-111`
7. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-52`
   - Y Coordinate of Point 2 Y_2: `-49`
   - Z Coordinate of Point 2 Z_2: `49`
8. Click Create.

   The volumetric control region is shown below.



9. Click OK.

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Case Box`.

3. Select the Absolute option.
4. Set Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `30`
   - Z Coordinate of Point 1 Z_1: `-111`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `43`
   - Y Coordinate of Point 2 Y_2: `-159`
   - Z Coordinate of Point 2 Z_2: `77`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Click Solution Controls on the STAR-Creo ribbon.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `300`.
4. Set Convergence Level to `0.01`.
5. Set Momentum Under-Relaxation Factor to `0.85`.

   The completed property dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Creo object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.

The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.

The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Creo *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Creo environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starcreo\data`.

- Copy directory `data` to your own disk space and open the part file `car.prt` in STAR-Creo.

The geometry is shown below.



The *STAR-Creo* ribbon appears at the top of the screen. The STAR-Creo object tree is located on the left of the screen.

If these items are not visible, read the help on Launching STAR-Creo.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-Creo tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

The remaining default settings are kept for this case.



3. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Click (Select Face) to activate the Creo face selection mechanism.
3. On the CAD model, select the face that is highlighted below. The selection is reported in the *Included Surfaces* box in the *Surface Sets* dialog.



4. Click OK to close the *Surface Sets* dialog window.
5. Select Symmetry Plane from the *Type* drop-down menu.
6. Set Name to `Symmetry Plane`.

   The completed dialog is shown below.



7. Click OK.

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Inlet`.
5. Change the units of Normal Velocity by clicking and select mph.



6. Set Normal Velocity to `70` mph.
7. Click OK.

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Rotate the view to display the back face and use the Creo face selection mechanism to select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case.
5. Click OK.

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the top and side walls, and use the Creo face selection mechanism to select the faces that are highlighted below. To select multiple faces, hold down the `<Ctrl>` key while selecting the faces.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Boundary Symmetry`.
5. Click OK.

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the bottom face, and use the Creo face selection mechanism to select the face that is highlighted below.



3. Leave the default Wall in the *Type* drop-down menu.
4. Set Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click OK.

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Rename.
2. Enter `Car` as the new name.
3. You may need to click on a blank area of the workspace to accept the name change.
4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Creo tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Select the Absolute option in the *Base Size* box.
3. Set Base Size to `0.03` m.
4. Expand the Prism Layer tab.
5. Select the Absolute radio button for Layer Thickness.
6. Set Layer Thickness to `0.03` m.

The completed *Mesh Pipeline Settings* dialog is shown below.



7. Click OK.

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Car Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `1440`
   - Z Coordinate of Point 2 Z_2 : `-4567`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.
2. Set Name to `Wake Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `3240`
   - Z Coordinate of Point 2 Z_2 : `-14567`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Creo tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Click (Select Face) to activate the Creo face selection mechanism.

3. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



4. Click OK to close the *Surface Sets* dialog window.
5. Set Name to `Fluid Domain Face Control`.
6. Activate the Enable option in the *Custom Mesh Size* box.
7. Select the Absolute option.
8. Set Custom Mesh Size to `0.35` m.

   The completed dialog is shown below.



9. Click OK.

Modifying Solution Controls

Modify the solution controls:

1. Click Solution Controls on the STAR-Creo ribbon.



The *Solution Controls* dialog appears.

2. Expand the *Solution Controls* box.
3. Set Number of Iterations to `400`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Creo object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.

The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Creo *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.
3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-Inventor CAD Client

Several tutorials have been provided for the STAR-Inventor CAD Client. These provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Inventor environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starinventor\data`.

- Copy directory `data` to your own disk space and open the part file `exhaustManifold.ipt` in STAR-Inventor.

The geometry is shown below.



The *STAR-Inventor* ribbon appears at the top of the screen. The STAR-Inventor object tree is located on the left of the screen. To switch to this tree:

- Click Model and select STAR-Inventor.



If these items are not visible, read the help on Launching STAR-Inventor.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Inventor. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-Inventor tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces that are highlighted below. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Mass Flow Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Mass Flow Rate to `0.017` kg/s.

   The completed dialog is shown below.



6. Click OK to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click OK to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Inventor tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to `2.5`%.



3. Click the Prism Layer tab.
4. Set Prism Layer Thickness to `15`%.



5. Click OK to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Click Solution Controls on the STAR-Inventor ribbon.



The *Solution Controls* dialog appears.

2. Click the Advanced tab.
3. Set Number of Iterations to `200`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK to close the dialog.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Inventor object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Expand the Volume Mesh icon and select Surface Mesh from the ribbon.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Expand the Surface Mesh icon and click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Inventor *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the <T> key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Inventor environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starinventor\data`.

- Copy directory `data` to your own disk space and open the assembly file `PorousModel.iam` in STAR-Inventor.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

The geometry is shown below.



The *STAR-Inventor* ribbon appears at the top of the screen. The STAR-Inventor object tree is located on the left of the screen. To switch to this tree:

- Click Model and select STAR-Inventor.



If these items are not visible, read the help on Launching STAR-Inventor.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Inventor. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Inventor tree and select Edit.



   The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-Inventor tree and select New.



The *Region Settings* dialog appears.

2. Select the middle part of the assembly as shown below.



3. In the *Region Settings* dialog, set the Name to `Porous Region`.
4. Activate the Porous model.
5. Within the *Porous Type* box, select the Uni-Directional option.
6. Set Alpha α to `25` kg/m^4 and Beta β to `1500` kg/m^3.s.
7. Set Cross-Flow Resistance Factor to `250`.

   The completed *Region Settings* dialog is shown below.



8. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-Inventor automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



6. Click OK.

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select New.

   The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click OK.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Inventor tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Select Parts* box.
3. Deactivate the Relative to bodies option in the *Base Size* box.
4. Set Base Size to `0.0025` m.



5. Click the Prism Layer tab.
6. Activate the Use Absolute Value option for Prism Layer Thickness.
7. Set Prism Layer Thickness to `0.001` m.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Click Solution Controls on the STAR-Inventor ribbon.



The *Solution Controls* dialog appears.

2. Click the Advanced tab.
3. Set Number of Iterations to `400`.
4. Set *Solution Accuracy* to First Order.

The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Inventor object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Expand ▾ the Volume Mesh icon and select Surface Mesh from the ribbon.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Expand the Surface Mesh icon and click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Inventor *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Inventor environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starinventor\data`.

- Copy directory `data` to your own disk space and open the assembly file `piccoloAssem.iam` in STAR-Inventor.

The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

The geometry is shown below.



The *STAR-Inventor* ribbon appears at the top of the screen. The STAR-Inventor object tree is located on the left of the screen. To switch to this tree:

- Click Model and select STAR-Inventor.



If these items are not visible, read the help on Launching STAR-Inventor.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-Inventor.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-Inventor tree and select Edit.



   The *Continuum Settings* dialog appears.

2. In the Physics tab, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Click on the Initialization tab.
5. Set Temperature to `400` K.

The remaining default settings are kept for this case. The completed dialog is shown below.





6. Click OK.

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-Inventor tree and select New.

The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material* menu, and select Aluminum Alloy from the list.
4. Click on the Initialization tab.

5. Set Temperature to 400 K.

   The remaining default settings are kept for this case. The completed dialog is shown below.





6. Click OK.

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-Inventor tree and select New.



   The *Region Settings* dialog appears.

2. Select the wing skin as shown below. Once selected, the part becomes highlighted.



3. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
4. Set Continuum to Al using the drop-down menu.



This action assigns the aluminum alloy that you created earlier as the material for this part.

The completed dialog is shown below.



5. Click OK.

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-Inventor tree and select Rename.



2. Enter `Internal Air` as the new name.

3. Click a blank area to accept the name change.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces highlighted below. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `90` m/s.
6. Set Temperature to `460` K.
7. Click on the Turbulence tab.
8. Set Turbulence Intensity to `0.135`.
9. Set Turbulent Length Scale to `2.2e-005` m.

   The completed dialog is shown below.



10. Click OK.

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are suitable for this tutorial.
5. Click OK.

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select New.

   The *Boundary Settings* dialog appears.
2. Select the four faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries.



3. Leave the selection as Wall in the *Type* drop-down menu.
4. Set Name to `Piccolo Wall`.
5. Set Thermal Specification to Temperature.



6. Set Temperature to `460` K.
7. Click OK.

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select New.

   The *Boundary Settings* dialog appears.
2. Select the two faces highlighted below, one on either side of the internal air region.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Air Symmetry`.
5. Click OK.

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.

2. Select the two faces highlighted below, one on either side of the geometry.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Solid Symmetry`.
5. Click OK.

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Edit.

   The *Boundary Settings* dialog appears.
2. Set Thermal Specification to Convection.



3. Set Heat Transfer Coefficient HTC to `100` W/m^2.K
4. Set Ambient Temperature to `228` K.

   The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Inventor tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Select Parts* box.
3. Deactivate the Relative to bodies option in the *Base Size* box.
4. Set Base Size to `0.007` m.



5. Click on the Prism Layer tab.
6. Set Prism Layer Thickness to `10`%.



7. Click OK.

Setting the Coincident Faces Detection Tolerance

The geometry that is used in this tutorial has large changes in feature size and requires you to reduce the Coincident Faces Detection Tolerance.

1. Right-click STAR-Inventor and select Geometry Transfer Settings.



The *Geometry Transfer Settings* dialog appears.

2. Set Coincident Faces Detection Tolerance to `0.0001`.



3. Click OK.

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Inlet Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-100.00`

- Y Coordinate of Point 1 Y_1: `21.00`
- Z Coordinate of Point 1 Z_1: `-20.00`

6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-90.00`
   - Y Coordinate of Point 2 Y_2: `10.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.

   The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Mid Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-111.00`
   - Y Coordinate of Point 1 Y_1: `34.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-64.00`
   - Y Coordinate of Point 2 Y_2: `-2.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Inventor tree and select New Part Mesh Control.



The *Part Mesh Control* dialog appears.

2. Select the solid wing region on the CAD model. The selection is added to the *Part Selection* box.
3. Set Name to `Solid Region`.
4. In the Prism Layer box, select Disable from the drop-down menu.



The completed dialog is shown below.



5. Click OK.

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-Inventor tree and select New Face Mesh Control.



The *Face Mesh Control* dialog appears.

2. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



3. Set Name to `Inlet Refinement`.
4. Activate the Custom Surface Mesh Size option.
5. Activate the Use Absolute Value option.
6. Set Custom Surface Mesh Size to `0.00025` m.
7. In the Prism Layer box, select Use Custom Values from the drop-down menu.
8. Activate the Use Absolute Value option.
9. Set Prism Layer Thickness to `0.0002` m.

   The completed dialog is shown below.



10. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Click Solution Controls on the STAR-Inventor ribbon.



The *Solution Controls* dialog appears.

2. Click on the Advanced tab.
3. Set Number of Iterations to `500`.
4. Set Convergence Level to `0.01`.

5. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

   The completed dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Inventor object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Expand the Volume Mesh icon and select Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Expand the Surface Mesh icon and click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Inventor *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Inventor environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starinventor\data`.

- Copy the `data` directory to your own disk space and open the assembly file `gpuAssem.iam` in STAR-Inventor.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

The geometry is shown below.



The *STAR-Inventor* ribbon appears at the top of the screen. The STAR-Inventor object tree is located on the left of the screen. To switch to this tree:

- Click Model and select STAR-Inventor.



If these items are not visible, read the help on Launching STAR-Inventor.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-Inventor. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-Inventor tree and select Edit.



   The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case.



3. Click OK to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-Inventor tree and select New.



The *Region Settings* dialog appears.

2. Select the disc that is located inside the graphics card housing. The part can be selected from within the PC tower by right-clicking on the disc and selecting Select Other... The disc will be highlighted in white, indicating a preliminary selection, left-click to confirm the selection. Once selected, the part becomes highlighted as shown below.



3. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
4. Activate the Moving Reference Frame option.
5. Set Rotational Velocity, to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.

6. Set Coordinates of Axis Origin to the following:
   - X : -122.1761
   - Y : -1.9434
   - Z : -19.0609
7. Set Component of Flow Direction to the following:
   - i : 0
   - j : 1
   - k : 0
8. Make sure that Use Assembly Coordinate System is selected.

   The completed property page is shown below:



9. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Case Inlet`.
5. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



6. Click OK.

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select New.

   The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select New.

   The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-Inventor tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Select Parts* box.
3. Deactivate the Relative to bodies option in the *Base Size* box.
4. Set Base Size to `0.016` m.



5. Click on the Prism Layer tab.
6. Set Prism Layer Thickness to 20%.



7. Click OK.

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.



   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Fan Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-166`
   - Y Coordinate of Point 1 Y_1: `4`
   - Z Coordinate of Point 1 Z_1: `-61`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-79`
   - Y Coordinate of Point 2 Y_2: `-34`
   - Z Coordinate of Point 2 Z_2: `25`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.

2. The *Volumetric Control Settings* dialog appears.
3. Set Name to `GPU Box`.
4. Activate the Use Absolute Value option.
5. Set Custom Surface Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `16`
   - Z Coordinate of Point 1 Z_1: `-111`
7. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-52`
   - Y Coordinate of Point 2 Y_2: `-49`
   - Z Coordinate of Point 2 Z_2: `49`
8. Click Create.

   The volumetric control region is shown below.



9. Click OK.

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.

The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Case Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `30`
   - Z Coordinate of Point 1 Z_1: `-111`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `43`
   - Y Coordinate of Point 2 Y_2: `-159`
   - Z Coordinate of Point 2 Z_2: `77`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Click Solution Controls on the STAR-Inventor ribbon.



The *Solution Controls* dialog appears.

2. Click on the Advanced tab.
3. Set Number of Iterations to `300`.
4. Set Convergence Level to `0.01`.
5. Set Momentum Under-Relaxation Factor to `0.85`.

   The completed property dialog is shown below.



6. Click OK.
7. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Inventor object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Expand the Volume Mesh icon and select Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Expand the Surface Mesh icon and click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Inventor *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-Inventor environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starinventor\data`.

- Copy directory `data` to your own disk space and open the part file `car.ipt` in STAR-Inventor.

The geometry is shown below.



The *STAR-Inventor* ribbon appears at the top of the screen. The STAR-Inventor object tree is located on the left of the screen. To switch to this tree:

- Click Model and select STAR-Inventor.



If these items are not visible, read the help on Launching STAR-Inventor.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-Inventor tree and select Edit.



The *Continuum* dialog appears.

2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

   The remaining default settings are kept for this case.



3. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select New.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Symmetry Plane`.

   The completed dialog is shown below.



5. Click OK.

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Inlet`.
5. Change the units of Normal Velocity by clicking and select mph.



6. Set Normal Velocity to `70` mph.
7. Click OK.

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the back face and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case.
5. Click OK.

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the top and side walls, and select the faces that are highlighted below.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Boundary Symmetry`.
5. Click OK.

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select New.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the bottom face, and select the face that is highlighted below.



3. Leave the default Wall in the *Type* drop-down menu.
4. Set Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click OK.

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Rename.
2. Enter `Car` as the new name.

3. You may need to click on a blank area of the workspace to accept the name change.
4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-Inventor tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Deactivate the Relative to bodies option in the *Base Size* box.
3. Set Base Size to `0.03` m.



4. Click on the Prism Layer tab.
5. Activate the Use Absolute Value option for Prism Layer Thickness .
6. Set Prism Layer Thickness to 0.03 m.



7. Click OK.

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Car Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `1440`
   - Z Coordinate of Point 2 Z_2 : `-4567`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Volumetric Control.

   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Wake Box`.
3. Activate the Use Absolute Value option.
4. Set Custom Surface Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `3240`

- Z Coordinate of Point 2 Z_2 : -14567

7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-Inventor tree and select New Face Mesh Control.



The *Face Mesh Control* dialog appears.

2. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



3. Set Name to `Fluid Domain Face Control`.
4. Activate the Custom Surface Mesh Size option.
5. Activate the Use Absolute Value option.
6. Set Custom Surface Mesh Size to `0.35` m.

The completed dialog is shown below.



7. Click OK.

Modifying Solution Controls

Modify the solution controls:

1. Click Solution Controls on the STAR-Inventor ribbon.



The *Solution Controls* dialog appears.

2. Click on the Advanced tab.
3. Set Number of Iterations to `400`.
4. Set Convergence Level to `0.0001`.

The completed property page is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-Inventor object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Expand the Volume Mesh icon and select Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Expand the Surface Mesh icon and click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-Inventor *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.
3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-SpaceClaim CAD Client

Several tutorials have been provided for the STAR-SpaceClaim CAD Client. These provide an overview of most of the features available in the client, and illustrate some of the applications for which it can be used. These comprise:

- Cold Flow Through a Motorcycle Exhaust Manifold
- Porous Media Flow
- Aircraft Wing Deicing System
- Rotating Flow Inside a Graphics Card
- External Flow Over a Car

Cold Flow Through a Motorcycle Exhaust Manifold

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-SpaceClaim environment. For this tutorial, the model is an exhaust manifold from a three-cylinder motorcycle engine.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starspaceclaim\data`.

- Copy directory `data` to your own disk space and open the geometry file `exhaustManifold.scdoc` in STAR-SpaceClaim.

The geometry is shown below.



The *STAR-SpaceClaim* ribbon appears at the top of the screen. The STAR-SpaceClaim object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree tab.

If these items are not visible, read the help on Launching STAR-SpaceClaim.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-SpaceClaim. To turn temperature calculations off:

1. Right-click the Continuum node in the STAR-SpaceClaim tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Defining Boundary Conditions

You define the inlet and outlet boundary conditions for the simulation.

Begin by creating the inlet boundary first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces that are highlighted below. To select multiple faces, hold down the `<Ctrl>` key. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Mass Flow Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.

5. Set Mass Flow Rate to `0.017` kg/s.

   The completed dialog is shown below.



6. Click OK to close the dialog.

Create the outlet boundary:

1. Right-click Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.

   The default settings are acceptable for this case. The completed dialog is shown below.



5. Click OK to close the dialog.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-SpaceClaim tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Set Base Size to `2.5%`.
3. Set Layer Thickness to `15%`.

   The completed *Mesh Pipeline Settings* dialog is shown below.



4. Click OK to close the dialog.

Modifying Solution Controls

The maximum number of iterations and the solution accuracy are modified.

1. Click Solution Controls on the STAR-SpaceClaim ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `200`.
3. Set Convergence Level to `0.0001`.

The completed property page is shown below.



4. Click OK to close the dialog.
5. Save your progress.

Reviewing the Simulation Set Up

The final STAR-SpaceClaim object tree is shown below.



Generating the Surface Mesh

To generate the surface mesh:

- Click Surface Mesh from the ribbon.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-SpaceClaim *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing two section planes. One plane bisects the inlets, while the other plane bisects the outlet. Pressure is plotted on these planes to show the pressure drop along the exhaust manifold.

To create the scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window. Next, create the section plane:

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`
7. Set the normal vector to `[0 mm, 0 mm, 1000 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.

The completed dialog is shown below.



9. Click Create.

Create the second section plane using the same dialog:

1. In the *Plane Parameters* box, set origin to `[0 mm, 0 mm, 35.0 mm]` and leave normal as `[0 mm, 0 mm, 1000 mm]`.



2. Click Create and then Close.

Define the scalar field to display the pressure contours:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Pressure.

The pressure plot shows the pressure drop inside the exhaust manifold. In order to get a smooth transition between the different pressure values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Make sure that the geometry is positioned on its side by pressing the `<T>` key while *Scalar Scene 2* is active.
4. Click the simulation button to return to the main STAR-CCM+ tree.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the exhaust manifold.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window. Next, modify the opacity of the geometry model in the scene:

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



4. Set Opacity to `0.3`
5. Click simulation.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the Seed Mode section.
6. Set Seed Parts to Regions > Region > Boundaries > Inlet.
7. Make sure that New Streamline Displayer is selected in the *Display* group box.

The completed dialog is shown below.



8. Click Create to create the streamlines and add them to the geometry scene as part of a streamline displayer.

9. Close the *Create Streamline* dialog.

Change the display method of the streamlines to tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.0020`

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity > Lab Reference Frame > Magnitude.
3. Save the simulation .

The finished scene is shown below.



Porous Media Flow

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-SpaceClaim environment. For this tutorial, the model is an idealized catalyst.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starspaceclaim\data`.

- Copy directory `data` to your own disk space and open the geometry file `PorousModel.scdoc` in STAR-SpaceClaim.

  The assembly is made up of three parts: the inlet region, the porous middle region, and the outlet region.

The geometry is shown below.



The *STAR-SpaceClaim* ribbon appears at the top of the screen. The STAR-SpaceClaim object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree tab.

If these items are not visible, read the help on Launching STAR-SpaceClaim.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-SpaceClaim. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-SpaceClaim tree and select Edit.



The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case. The completed dialog is shown below.



3. Click OK to close the dialog.

Setting the Porous Region

Create the porous region and define its properties:

1. Right-click the Regions manager in the STAR-SpaceClaim tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the middle part of the assembly as shown below.



3. In the *Region Settings* dialog, set the Name to `Porous Region`.
4. Activate the Porous model.

5. Within the *Porous Settings* box, select the Uni-Directional option.
6. Set Alpha Coefficient α to `25` kg/m^4 and Beta Coefficient β to `1500` kg/m^3.s.
7. Set Cross-Flow Resistance Factor to `250`.

   The completed *Region Settings* dialog is shown below.


Region Settings
Name
Porous Region
Part Selection
comp0
Continuum Selection
Continuum
Porous Settings
Porous
Isotropic
Uni-Directional
α
25
kg/m^4
β
1500
kg/m^3.s
250
0.000000000
0.000000000
1.000000000
Use Assembly Coordinate System
Moving Reference Frame Settings
Moving Reference Frame
Advanced
Exclude from Analysis
OK
Cancel

8. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions. STAR-SpaceClaim automatically adds interfaces between the different regions. These interfaces are visible in the final STAR-CCM+ simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type drop-down menu.
4. Set Name to `Inlet`.
5. Set Normal Velocity to `5` m/s.

   The completed dialog is shown below.



6. Click OK.

Create the outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the face that is highlighted below.



3. Select Pressure Outlet from the Type drop-down menu.
4. Set Name to `Outlet`.
5. The default settings are acceptable for this case.
6. Click OK.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-SpaceClaim tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.0025` m.
5. Select the Absolute radio button for Layer Thickness.
6. Set Layer Thickness to `0.001` m.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
comp2
comp1
comp0
Mesh Type
Polyhedral
Trimmer
Base Size
0.0025
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Geometry
Absolute
Prisim Layer
Enable
0.001
m
Relative to Base Size
Absolute
2
1.500000000
Advanced Settings
OK
Cancel

7. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy settings.

1. Click Solution Controls on the STAR-SpaceClaim ribbon.



   The *Solution Controls* dialog appears.

2. Set Number of Iterations to `400`.
3. Set *Solution Accuracy* to First Order.

   The completed property page is shown below.



4. Click OK.
5. Save your progress.

Reviewing the Simulation Set Up

The final STAR-SpaceClaim object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-SpaceClaim *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Vector Scene

Create a vector scene with a section plane bisecting the geometry. Velocity vectors are plotted on this plane.

To create the vector scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Vector.



A new vector scene, *Vector Scene 1*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane...



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 mm, 0 mm, 0 mm]`.
7. Set the normal vector to `[0 mm, 1000 mm, 0 mm]`
8. In the *Display* group box, select Existing Displayer and in the drop-down menu select Vector 1.

The completed dialog is shown below.



9. Click Create followed by Close.

   Velocity is automatically set as the active function when creating a vector displayer.

Change the view of the scene to something more suitable:

- Click the camera icon on the toolbar and select Views > +X +Y +Z > Up +Y.

In order to remove the unwanted outlines from the scene do the following:

1. Click the scene/plot button in the explorer pane.
2. Navigate to Vector Scene 1 > Displayers > Outline 1 > Parts.
3. Right-click on the Parts folder and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Select the Derived Parts > Plane Section node.
6. Click OK.
7. Click the simulation button to return to the simulation object tree.
8. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow inside the catalyst.

Begin by creating a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.

   A new geometry scene is added to the *Graphics* window.

2. Click scene/plot.
3. Navigate to Geometry Scene 1 > Displayers > Geometry 1.
4. Set Opacity to `0.3`

To remove the yellow interface boundaries from the scene:

1. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts and select Edit.



2. In the *Parts* dialog that appears, expand the Regions > Porous Region > Boundaries and Regions > Region > Boundaries nodes.
3. Deactivate the following items in each region:
   - Interface [In-Place 1]

- Interface 2 [In-Place 2]



4. Click OK.

Hide the outlines on the geometry model:

1. Right-click on the Geometry Scene 1 > Displayers > Outline 1 > Parts node in the object tree and select Hide All Parts.



2. Click simulation.

To change the view of the scene to something more suitable:

- Click the camera button on the toolbar and select Views > +X +Y +Z > Up +Y.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline...

   The *Create Streamline* dialog appears.
2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Porous Region and Region.
4. Click OK to close the *Select Objects* dialog.
5. Make sure that Part Seed is selected in the *Seed Mode* section.
6. Set Seed Parts to Regions > Region >Boundaries > Inlet.
7. Click OK.

8. Make sure that New Streamline Displayer is selected in the *Display* group box.

   The completed dialog is shown below.



9. Click Create then Close.

Change the display mode of the streamlines to use tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Width to `0.0012` and Mode to Tubes.

Change the color of the streamlines to represent velocity magnitude:

1. Navigate to the Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field.
2. Set Function to Velocity >Lab Reference Frame> Magnitude.
3. Save the simulation .

The finished scene is shown below.



Aircraft Wing Deicing System

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-SpaceClaim environment. For this tutorial, the model is a modern deicing system that is used on aircraft wings. The hot air from the aircraft's jet engine is used to heat the leading edge of the wing.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starspaceclaim\data`.

- Copy directory `data` to your own disk space and open the geometry file `piccoloAssem.scdoc` in STAR-SpaceClaim.

The assembly is made up of two parts: the solid leading edge wing skin, and the internal air region.

The geometry is shown below.



The *STAR-SpaceClaim* ribbon appears at the top of the screen. The STAR-SpaceClaim object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree tab.

If these items are not visible, read the help on Launching STAR-SpaceClaim.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to modify the physics continuum that is used to model air. For this case, the flow is turbulent and thermal.

Note: Temperature calculations are on by default in STAR-SpaceClaim.

Modify the initial conditions:

1. Right-click the Continua > Continuum node in the STAR-SpaceClaim tree and select Edit.



   The *Continuum Settings* dialog appears.

2. In the *Continuum Settings* dialog, enter `Air` for the Name.
3. Set the Equation of State to Ideal Gas.
4. Expand the Initialization box by clicking .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.


Continuum Settings
Name
Air
Material
Material Type
Gas
Liquid
Solid
Material Name
Air
Flow Characteristics
Laminar Flow
Turbulent Flow
Equation of State
Ideal Gas
Constant Density
Gravity
Enable
Optional Physical Models
Passive Scalar
Temperature
Material Properties
Molecular Weight
28.966400000
kg/kmol
Specific Heat
1003.620000000
J/kg-K
Molecular Viscosity
0.000018551
Pa-s
Conductivity
0.026030500
W/m-K
Reference Properties
Initialization
U
0.000000000
m/s
V
0.000000000
W
0.000000000
0.000000000
Pa
400
K
Int
0.010000000
Icn
0.010000000
m
OK
Cancel

6. Click OK.

Adding Solid Physics Models

The next step in setting up this tutorial is to define the physics continuum of the solid to be used.

1. Right-click the Continua manager in the STAR-SpaceClaim tree and select Add Continuum.



   The *Continuum Settings* dialog appears.

2. Enter `Al` for the Name.
3. Select Solid from the *Material Type* menu, and select Aluminum Alloy from the drop-down list.
4. Expand the Initialization box by clicking .
5. Set Temperature to `400` K.

   The remaining default settings are kept for this case. The completed dialog is shown below.



6. Click OK.

Setting the Solid Region

Create a new solid region for the wing skin:

1. Right-click the Regions manager in the STAR-SpaceClaim tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the outer wing skin.
3. In the *Region Settings* dialog, enter `Leading Edge` for the Name.
4. Set Continuum Selection to Al from the selection box.

   This action assigns the aluminum alloy that you created earlier as the material for this part.

   The completed dialog is shown below.



5. Click OK.

Now that the solid region of the wing is defined it would be suitable to rename the default region to something more appropriate:

1. Right-click on the Regions > Region node in the STAR-SpaceClaim tree and select Rename.



2. Enter `Internal Air` as the new name.
3. Click a blank area to accept the name change.

Defining Boundary Conditions for the Fluid Region

Create the inlet and outlet boundary conditions. Use symmetry boundaries on the side walls to extend the simulation domain. Use a wall boundary with a fixed temperature to model the piccolo walls.

Create the inlet boundary first:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the three faces highlighted below. To select multiple faces hold down the `<Ctrl>` key. These selections are reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the Type drop-down menu.

4. Set Name to `Inlet`.
5. Set Normal Velocity to `90` m/s.
6. Set Temperature to `460` K.
7. Set Turbulence Intensity to `0.135`.
8. Set Turbulent Length Scale to `2.2e-005` m.

   The completed dialog is shown below.



9. Click OK.

Create the outlet boundary:

1. Right-click the Regions > Internal Air > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

   The default settings are suitable for this tutorial.
5. Click OK.

Create the wall boundary that models the heated piccolo walls.

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the four faces highlighted below, make sure to select the faces that are perpendicular to the inlet boundaries.



3. Leave the selection as Wall in the *Type* drop-down menu.
4. Set Name to `Piccolo Wall`.
5. Set Thermal Specification to Temperature.



6. Set Temperature to `460` K.
7. Click OK.

Create the symmetry boundaries:

1. Right-click Regions > Internal Air > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the two faces highlighted below, one on either side of the internal air region. To select multiple faces hold down the `<Ctrl>` key.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Air Symmetry`.
5. Click OK.

Defining Boundary Conditions for the Solid Region

Use the convection boundary condition to model the heat propagation from the solid wing.

The convective boundary condition models ambient conditions on the outside of the domain.Use convection to model the heat propagation in the solid wing skin. Also apply symmetry boundaries to the solid region.

Create the symmetry boundary first:

1. Right-click the Regions > Leading Edge > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Select the two faces highlighted below, one on either side of the geometry. To select multiple faces hold down the `<Ctrl>` key.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Solid Symmetry`.
5. Click OK.

Modify the default boundary so that it uses convection:

1. Right-click the Regions > Leading Edge > Boundaries > Default node and select Edit.

   The *Boundary Settings* dialog appears.
2. Set Thermal Specification to Convection.



3. Set Heat Transfer Coefficient HTC to `100` W/m^2.K
4. Set Ambient Temperature to `228` K.

   The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-SpaceClaim tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.007` m.
5. Set Layer Thickness to `10`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
comp1
comp0
Mesh Type
Polyhedral
Trimmer
Base Size
0.007
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Geometry
Absolute
Prisim Layer
Enable
10
%
Relative to Base Size
Absolute
2
1.500000000
Advanced Settings
OK
Cancel

6. Click OK.

Creating the Volumetric Mesh Controls

Create volumetric mesh control regions to capture the finer details of the flow near the air jets that heat the leading edge of the wing.

Create the volumetric mesh control for the inlet:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Inlet Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0005` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-100.00`
   - Y Coordinate of Point 1 Y_1: `21.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-90.00`
   - Y Coordinate of Point 2 Y_2: `10.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the region near the piccolo tube:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.
2. Set the Name to `Mid Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.0015` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-111.00`
   - Y Coordinate of Point 1 Y_1: `34.00`
   - Z Coordinate of Point 1 Z_1: `-20.00`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-64.00`
   - Y Coordinate of Point 2 Y_2: `-2.00`
   - Z Coordinate of Point 2 Z_2: `60.00`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.
9. Save your progress.

Creating a Part Mesh Control

Create a part mesh control to disable the prism layer generation in the solid region. Solid regions do not require prism layers.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-SpaceClaim tree and select Add Part Mesh Control.



The *Part Mesh Control Settings* dialog appears.

2. Select the solid wing region on the CAD model. The selection is added to the *Part Selection* box.
3. Set Name to `Solid Region`.
4. In the Prism Layer box, set Prism layer option to Disable using the drop-down menu.



The completed dialog is shown below.



5. Click OK.

Creating a Face Mesh Control

Create a face mesh control to reduce the size of the prism layer thickness at the inlet boundary.

1. Right-click on Mesh Pipelines > Mesh Pipeline > Mesh Controls in the STAR-SpaceClaim tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Select the cylindrical faces perpendicular to the inlet boundary, as shown below.



3. Set Name to `Inlet Refinement`.
4. In the Custom Surface Mesh Size box, activate the Enable option.
5. Set Custom Mesh Size Option to Absolute.
6. Set Mesh Size to `0.00025` m.
7. In the Prism Layer box, set Prism Layer Option to Use Custom Values using the drop-down menu.
8. Set Layer Thickness Option to Absolute.
9. Set Layer Thickness to `0.0002` m.

   The completed dialog is shown below.



10. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy. To help accelerate convergence, edit the solid under-relaxation factor:

1. Click Solution Controls on the STAR-SpaceClaim ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `500`.
3. Set Convergence Level to `0.01`.
4. Set Solid Temperature Under-Relaxation Factor to `0.9999`.

The completed dialog is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-SpaceClaim object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the simulation:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-SpaceClaim *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the model. Plot temperature on this section to demonstrate how the hot gas coming from the piccolo tube affects the wing skin and the air inside the wing.

To create the scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Leading Edge and Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Set the origin coordinates to `[0.0 mm, 0.0 mm, 18.4 mm]`.
7. Set the normal vector to `[0.0 mm, 0.0 mm, 1.0 mm]`.
8. In the *Display* group box, select Existing Displayer, and select Scalar 1 from the drop-down menu.



The completed dialog is shown below.



9. Click Create followed by Close.

Change the scalar field to display temperature contours:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Temperature.

Modify the contour style to have a smooth transition between the different temperature values in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.

Adjust the outlines in the scene to show only the plane section:

1. Right-click on Scalar Scene 2 > Displayers > Outline 1 > Parts and select Edit.



The *Parts* dialog appears.

2. Right-click on a blank area and choose Deselect All.
3. Select the Derived Parts > plane section node.
4. Click OK.
5. Position the scene to show the side view of the geometry.
6. Click simulation.
7. Save the simulation .

The finished scene is shown below.



Visualizing the Wing Skin Temperature

Create another scalar scene to display the temperature on the wing skin geometry. To create another scalar scene:

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.

A new scalar scene, *Scalar Scene 3*, is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click Scalar Scene 3 > Displayers > Scalar 1 > Parts and select Edit.



The *Parts* dialog appears.

4. Expand the Regions > Leading Edge > Boundaries folder and select the Leading Edge node.
5. Click OK.

6. Navigate to Scalar Scene 3 > Displayers > Scalar 1 > Scalar Field.



7. Set Function to Temperature.

To have a smooth transition between the different temperature values, edit the contour style:

1. Navigate to Scalar Scene 3 > Displayers > Scalar 1.
2. Set Contour Style to Smooth Filled.

Adjust the outlines to appear on the solid region only:

1. Right-click on Scalar Scene 3 > Displayers > Outline 1 > Parts and select Edit.

   The *Parts* dialog appears.
2. Expand the Regions folder and select the Internal Air node.
3. Click OK.
4. Reorient the geometry to get a better view of the leading edge.
5. Click simulation.

6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of how the hot air flows inside the leading edge of the wing.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene, *Geometry Scene 1*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to the Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Next, create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Internal Air.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Internal Air > Boundaries > Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

The streamlines do not yet appear in the geometry scene. Follow the steps that are outlined below to modify the streamlines:

1. Navigate to Derived Parts > Streamline > Source Seed.



2. Activate the Randomize option.



3. Navigate to Derived Parts > Streamline > 2nd Order Integrator.



4. Set Maximum Propagation to `0.3`.



Adjust the appearance of the streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. Set Mode to Tubes and Width to `0.00025`.

Adjust the scalar field of the streamlines to display temperature:

1. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1> Scalar Field.
2. Set Function to Temperature.

Turn off the outlines in the geometry scene to have a better view of the streamlines.

1. Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reorient the geometry to get a better view of the leading edge.
3. Save the simulation .

The finished scene is shown below.



Rotating Flow inside a Graphics Card

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-SpaceClaim environment. For this tutorial, the model is a fully sealed graphics card that is located inside an empty, idealized PC tower case.

The geometry that is used in this tutorial has been simplified on order to decrease the computational cost. As a consequence the mesh is relatively coarse at approximately 200,000 cells compared to the 3-4 million cells used for an industrial analysis.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starspaceclaim\data`.

- Copy the `data` directory to your own disk space and open the geometry file `gpuAssem.scdoc` in STAR-SpaceClaim.

  The assembly is made up of two parts: the PC tower case (including the GPU housing), and the rotating region that encloses the blower fan of the GPU.

The geometry is shown below.



The *STAR-SpaceClaim* ribbon appears at the top of the screen. The STAR-SpaceClaim object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree tab.

If these items are not visible, read the help on Launching STAR-SpaceClaim.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal. Temperature calculations are on by default in STAR-SpaceClaim. To turn temperature calculations off:

1. Right-click the Continua > Continuum node in the STAR-SpaceClaim tree and select Edit.



   The *Continuum Settings* dialog appears.

2. Clear the Temperature option in the Optional Physics Models box.

   The remaining default settings are kept for this case.



3. Click OK to close the dialog.

Setting the Rotating Region

Create the rotating region that contains the GPU fan:

1. Right-click the Regions manager in the STAR-SpaceClaim tree and select Add Region.



The *Region Settings* dialog appears.

2. Select the disc that is located inside the graphics card housing.This part can be selected from the Structure tab by selecting gpuBlowerFan. Once selected, the part becomes highlighted as shown below.



3. In the *Region Settings* dialog, enter `Rotating Region` for the Name.
4. Activate the Moving Reference Frame option.
5. Set Rotational Velocity, ω to `-2800 rpm`. Ensure that the units are set to `rpm` using the drop-down menu.
6. Set Coordinates of Axis Origin to the following:
   - X: `-122.1761`
   - Y: `-1.9434`
   - Z: `-19.0609`

7. Set Component of Axis Vector to the following:
   - i : 0
   - j : 1
   - k : 0
8. Make sure that Use Assembly Coordinate System is selected.

   The completed property page is shown below:



9. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary conditions for this simulation.

Create the inlet boundary first:

1. Right-click the Regions > Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Velocity Inlet from the *Type* drop-down menu.
4. Set Name to `Case Inlet`.
5. Set Normal Velocity to `0.1 m/s`.

   The complete dialog is shown below.



6. Click OK.

Create the case outlet boundary:

1. Right-click the Regions > Region > Boundaries folder and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. On the CAD model, rotate the display and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Case Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

Create the GPU outlet boundary:

1. Right-click Regions > Region > Boundaries and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `GPU Outlet`.

   The remaining default settings are suitable for this tutorial.
5. Click OK.

6. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipelines > Mesh Pipeline node in the STAR-SpaceClaim tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Make sure that all parts of the geometry are selected in the *Part Selection* box.
3. Select the Absolute option in the *Base Size* box.
4. Set Base Size to `0.016` m.
5. Set Layer Thickness to `20`%.

   The completed *Mesh Pipeline Settings* dialog is shown below.


Mesh Pipeline Settings
Name
Mesh Pipeline
Part Selection
comp1
comp0
Mesh Type
Polyhedral
Trimmer
Base Size
0.016
m
Relative to Geometry
Absolute
Surface Mesh Size
100.000000000
%
Relative to Geometry
Absolute
Prisim Layer
Enable
20
%
Relative to Base Size
Absolute
2
1.500000000
Advanced Settings
OK
Cancel

6. Click OK.

Creating the Volumetric Mesh Controls

Use volumetric mesh control regions to capture the finer details of the geometry model. Create three regions in total: a small region that closely surrounds the rotating region containing the fan, a region that encloses the GPU case, and a region that acts as an intermediary between the base size and the other volumetric control regions.

Create the volumetric mesh control for the rotating region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Fan Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.002 m`.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-166`
   - Y Coordinate of Point 1 Y_1: `4`
   - Z Coordinate of Point 1 Z_1: `-61`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-79`
   - Y Coordinate of Point 2 Y_2: `-34`

- Z Coordinate of Point 2 Z_2: 25

7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the GPU region:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.
2. The *Volumetric Control Settings* dialog appears.
3. Set Name to `GPU Box`.

4. Select the Absolute option.
5. Set Mesh Size to `0.005` m.
6. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `16`
   - Z Coordinate of Point 1 Z_1: `-111`
7. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `-52`
   - Y Coordinate of Point 2 Y_2: `-49`
   - Z Coordinate of Point 2 Z_2: `49`
8. Click Create.

   The volumetric control region is shown below.



9. Click OK.

Create the volumetric mesh control for the PC tower case:

1. Right-click the Mesh Pipelines > Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.

2. Set the Name to `Case Box`.

3. Select the Absolute option.
4. Set Mesh Size to `0.009` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1: `-383`
   - Y Coordinate of Point 1 Y_1: `30`
   - Z Coordinate of Point 1 Z_1: `-111`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2: `43`
   - Y Coordinate of Point 2 Y_2: `-159`
   - Z Coordinate of Point 2 Z_2: `77`
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Modifying Solution Controls

Modify the maximum number of iterations and the solution accuracy:

1. Click Solution Controls on the STAR-SpaceClaim ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `300`.
3. Set Convergence Level to `0.01`.
4. Set Momentum Under-Relaxation Factor to `0.85`.

   The completed property dialog is shown below.



5. Click OK.
6. Save your progress.

Reviewing the Simulation Set Up

The final STAR-SpaceClaim object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the solver:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-SpaceClaim *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene containing a cross-section of the GPU and plot velocity magnitude on this plane.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Right-click on the Derived Parts node and select New Part > Section > Plane.



The *Create Plane Section* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. In the *Plane Parameters* box, set the origin vector to `[0 m, -0.012 m, 0 m]` and the normal vector to `[0 m, 1 m, 0 m]`.

   Note that the units are in meters.
7. In the *Display* group box, select Existing Displayer and in the drop-down menu select Scalar 1.



The completed dialog is shown below.



8. Click Create and then Close.

Display velocity magnitude on the section plane:

1. Click scene/plot.
2. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



3. Set Function to Velocity >Lab Reference Frame > Magnitude.

To have a smooth transition between the different velocity values, set the contour style to smooth:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. To hide all outlines, right-click Scalar Scene 2 > Displayers > Outline 1, and select Toggle Visibility.



4. Click the Scalar Scene 2 tab to activate the scene, and press the <S> key to change to the top view of this geometry.
5. Click simulation.
6. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the air flow inside the PC tower case and graphics card.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Navigate to Geometry Scene 1 > Displayers > Geometry 1 node.



4. Set Opacity to `0.1`

Create the streamlines:

1. Click simulation.

2. Right-click on the *Derived Parts* node and select New Part > Streamline...



The *Create Streamline* dialog appears.

3. In the *Input Parts* box, click Select...
4. In the *Select Objects* dialog, expand the Regions node and select Region and Rotating Region.
5. Click OK to close the *Select Objects* dialog.
6. Make sure that Part Seed is selected in the Seed Mode section.
7. Set Seed Parts to Regions > Region > Boundaries > Case Inlet.
8. Set Part U-Resolution to `10`.
9. Set Part V-Resolution to `10`.
10. Make sure that New Streamline Displayer is selected in the Display section.

The completed dialog is shown below.



11. Click Create followed by Close.

Adjust the appearance of streamlines to display tubes instead of the default lines:

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1and set Mode to Tubes and Width to `0.0020`.
3. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Magnitude.

Hide all the outlines in the scene:

1. Right-click Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.
2. Reposition the geometry to get a better view of the flow inside the case.
3. Save the simulation .

The finished scene is shown below.



External Flow Over a Car

This tutorial takes you through the steps that are required to run a CFD calculation in the STAR-SpaceClaim environment. For this tutorial, the model is a car in a fluid domain.



Opening the CAD Model

A directory containing the CAD geometry is provided with the STAR-CCM+ distribution. This directory is located in `[STAR-CCM+_INSTALL_DIR]\doc\startutorialsdata\starspaceclaim\data`.

- Copy directory `data` to your own disk space and open the geometry file `car.scdoc` in STAR-SpaceClaim.

The geometry is shown below.



The *STAR-SpaceClaim* ribbon appears at the top of the screen. The STAR-SpaceClaim object tree is located on the left of the screen. To switch to this tree:

- Click the STAR Tree tab.

If these items are not visible, read the help on Launching STAR-SpaceClaim.

Selecting Fluid Physics Models

The first step in setting up this tutorial is to define the physics of the problem. For this case, the flow is turbulent and isothermal.

1. Right-click the Continuum node in the STAR-SpaceClaim tree and select Edit.



2. Clear the Temperature option in the *Optional Physics Models* box as shown below.

   The remaining default settings are kept for this case.



3. Click OK.

Defining Boundary Conditions

Create the inlet and outlet boundary condition. Use a moving wall to model the relative movement of the road.

Create the symmetry plane first:

1. Right-click the Region > Boundaries manager and select Add Boundary.



The *Boundary Settings* dialog appears.

2. On the CAD model, select the face that is highlighted below. The selection is reported in the *Face Selection* box in the *Boundary Settings* dialog.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Symmetry Plane`.

   The completed dialog is shown below.



5. Click OK.

Create the inlet boundary:

1. Right-click the Regions > Boundaries manager and select Add Boundary. The *Boundary Settings* dialog appears.
2. Select the face that is highlighted below.



3. Select Velocity Inlet from the *Type* drop-down menu.

4. Set Name to `Inlet`.
5. Change the units of Normal Velocity by clicking and select mph.



6. Set Normal Velocity to `70` mph.
7. Click OK.

Create the outlet boundary:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.
2. Rotate the view to display the back face and select the face that is highlighted below.



3. Select Pressure Outlet from the *Type* drop-down menu.
4. Set Name to `Outlet`.

The default settings are acceptable for this case.

5. Click OK.

Create the fluid domain symmetry boundaries:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Rotate the view to display the top and side walls, and select the faces that are highlighted below. To select multiple faces, hold down the `<Ctrl>` key while selecting the faces.



3. Select Symmetry Plane from the *Type* drop-down menu.
4. Set Name to `Boundary Symmetry`.
5. Click OK.

Create the boundary that will model the moving ground:

1. Right-click the Region > Boundaries manager and select Add Boundary.

   The *Boundary Settings* dialog appears.

2. Rotate the view to display the bottom face, and select the face that is highlighted below.



3. Leave the default Wall in the *Type* drop-down menu.
4. Set Name to `Ground`.
5. Set Wall Type to Translating.



6. Set W Component of Wall Velocity Vector to `-70` mph. Ensure that the units are set to mph.
7. Click OK.

Rename the default boundary:

1. Right-click Region > Boundaries > Default and select Rename.
2. Enter `Car` as the new name.

3. You may need to click on a blank area of the workspace to accept the name change.
4. Save your progress.

Setting the Mesh Pipeline

To obtain more accurate results, modify the mesh pipeline:

1. Right-click the Mesh Pipeline node in the STAR-SpaceClaim tree and select Edit.



The *Mesh Pipeline Settings* dialog appears.

2. Select the Absolute option in the *Base Size* box.
3. Set Base Size to `0.03` m.
4. Select the Absolute radio button for Layer Thickness.
5. Set Layer Thickness to `0.03` m.

The completed *Mesh Pipeline Settings* dialog is shown below.



6. Click OK.

Creating the Volumetric Mesh Controls

Create two volumetric mesh control regions to capture the finer details of the flow: one around the car and one that extends into the wake of the car.

Create the volumetric mesh control for the car region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.



The *Volumetric Control Settings* dialog appears.

2. Set Name to `Car Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.09` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`
   - Z Coordinate of Point 1 Z_1 : `3432`
6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : `2000`
   - Y Coordinate of Point 2 Y_2 : `1440`
   - Z Coordinate of Point 2 Z_2 : `-4567`
7. Click Create.

   The completed dialog is shown below.



The volumetric control region is shown below.



8. Click OK.

Create the volumetric mesh control for the wake region:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Volumetric Mesh Control.

   The *Volumetric Control Settings* dialog appears.

2. Set Name to `Wake Box`.
3. Select the Absolute option.
4. Set Mesh Size to `0.177` m.
5. Enter the following coordinates for Point 1:
   - X Coordinate of Point 1 X_1 : `-50`
   - Y Coordinate of Point 1 Y_1 : `-650`

- Z Coordinate of Point 1 Z_1 : 3432

6. Enter the following coordinates for Point 2:
   - X Coordinate of Point 2 X_2 : 2000
   - Y Coordinate of Point 2 Y_2 : 3240
   - Z Coordinate of Point 2 Z_2 : -14567
7. Click Create.

   The volumetric control region is shown below.



8. Click OK.

Creating the Face Mesh Control

Apply a face mesh control to the walls of the fluid domain to coarsen the mesh on these faces.

Create the face mesh control:

1. Right-click the Mesh Pipeline > Mesh Controls node in the STAR-SpaceClaim tree and select Add Face Mesh Control.



The *Face Mesh Control Settings* dialog appears.

2. Select all six boundary faces that form the fluid domain: `Inlet`, `Outlet`, `Boundary Symmetry`, `Symmetry Plane`, and `Ground`.



3. Set Name to `Fluid Domain Face Control`.
4. Activate the Enable option in the *Custom Mesh Size* box.
5. Select the Absolute option.
6. Set Mesh Size to `0.35` m.

   The completed dialog is shown below.



7. Click OK.

Modifying Solution Controls

Modify the solution controls:

1. Click Solution Controls on the STAR-SpaceClaim ribbon.



The *Solution Controls* dialog appears.

2. Set Number of Iterations to `400`.
3. Set Convergence Level to `0.0001`.

   The completed property page is shown below.



4. Click OK.
5. Save your progress.

Reviewing the Simulation Set Up

The final STAR-SpaceClaim object tree is shown below.



Generating the Surface Mesh

Generate the surface mesh:

- Click Surface Mesh from the ribbon.



The completed surface mesh is shown below.



Generating the Volume Mesh

Generate the volume mesh:

- Click Volume Mesh on the ribbon.



The completed volume mesh is shown below.



Building and Running

The case is now ready to run. To run the case:

- Select Build and Run from the ribbon.



The *Residuals* plot appears in the STAR-SpaceClaim *Output* window. This plot can be used to monitor the convergence of the solution.

STAR-CCM+ launches when the solver is done.

Post-Processing Results

A scalar scene, *Scalar Scene 1*, automatically appears in the *Graphics* window when the solver has finished. This scene displays the absolute total pressure.

Further post-processing is carried out in STAR-CCM+.

The results that are contained in the following sections are representative. You may see some variation in the results that you obtain, depending on what platform and package you are using.

- Save the simulation .

Creating a Scalar Scene

Create a scalar scene to display static pressure.

1. Right-click on the Scenes node in the explorer pane and select New Scene > Scalar.



A new scalar scene, *Scalar Scene 2*, is created and selected in the *Graphics* window.

2. Click scene/plot.

3. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

4. Expand the Regions > Region > Boundaries folder and select the following items:
   - Car
   - Ground
   - Symmetry Plane
5. Click OK.

Display the pressure contours in the scene:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



2. Set Function to Static Pressure.

To have a smooth transition between the different pressure values:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Contour Style to Smooth Filled.
3. Rotate the scene to get a better view of the front of the car.
4. Save the simulation .

The finished scene is shown below.



Modifying the Scalar Scene

Modify Scalar Scene 2 to display Turbulent Kinetic Energy and remove the symmetry plane from the scene. You also mirror the car body:

1. Right-click on Scalar Scene 2 > Displayers > Scalar 1 > Parts and select Edit...



The *Parts* dialog appears.

2. Expand the Regions > Region > Boundaries folder and deactivate Symmetry Plane.
3. Click OK.

4. Navigate to Scalar Scene 2 > Displayers > Scalar 1 > Scalar Field.



5. Set Function to Turbulent Kinetic Energy.

Hide the outlines in the scene:

- Right-click on Scalar Scene 2 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Scalar Scene 2 > Displayers > Scalar 1.



2. Set Transform to Symmetry Plane 1.



3. Rotate the scene to get a better view of the wake of the car.
4. Click simulation.
5. Save the simulation .

The finished scene is shown below.



Creating a Streamlines Scene

Create streamlines to get a better understanding of the flow around the car and in the wake area.

First, create a geometry scene:

1. Right-click on the Scenes node in the object tree and select New Scene > Geometry.



A new geometry scene is created and selected in the *Graphics* window.

2. Click scene/plot.
3. Right-click on Geometry Scene 1 > Displayers > Geometry 1 > Parts node and select Edit.



The *Parts* dialog appears.

4. Right-click on a blank area and choose Deselect All.
5. Expand the Regions > Region > Boundaries node and select Car.
6. Click OK.

Hide the outlines in the scene:

- Right-click on Geometry Scene 1 > Displayers > Outline 1 and select Toggle Visibility.



Mirror the car body:

1. Navigate to Geometry Scene 1 > Displayers > Geometry 1.



2. Set Transform to Symmetry Plane 1.



3. Click simulation.

Create the streamlines:

1. Right-click on the *Derived Parts* node and select New Part > Streamline.



The *Create Streamline* dialog appears.

2. In the *Input Parts* box, click Select...
3. In the *Select Objects* dialog, expand the Regions node and select Region.
4. Click OK to close the *Select Objects* dialog.
5. Set Seed Mode to Line Seed.



6. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-123 mm`
     - Z: `2600 mm`
7. Set Resolution to `10`.

8. Make sure that New Streamline Displayer is selected in the Display section.

   The completed dialog is shown below.



9. Click Create.

Create another set of streamlines using the existing dialog:

1. Enter these coordinates for the line:
   - Point 1:
     - X: `0 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`
   - Point 2:
     - X: `1200 mm`
     - Y: `-473 mm`
     - Z: `2600 mm`

The completed dialog is shown below



2. Click Create followed by Close.

Modify the appearance of the streamlines to display ribbons instead of the default lines.

1. Click scene/plot.
2. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1.
3. In the *Properties* window:
   - Set Width to `0.025`.
   - Set Mode to Ribbons.
   - Set Transform to Symmetry Plane 1.

Modify the scalar field of the streamlines to display velocity magnitude:

4. Navigate to Geometry Scene 1 > Displayers > Streamline Stream 1 > Scalar Field and set Function to Velocity > Lab Reference Frame > Magnitude.
5. Repeat the above steps for Streamline Stream 2.
6. Rotate the scene to get a better view of the wake of the car.
7. Save the simulation .

The finished scene is shown below.



STAR-View+ Contents

Using STAR-View+ 1

What is STAR-View+? **2**

Installing STAR-View+ **3**

Installing the Downloaded Package on Windows 3

Installing the Downloaded Package on Linux 6

Exporting a Scene to a STAR-View+ File **7**

Starting STAR-View+ **10**

Working with STAR-View+ **11**

Opening a File 11

Embedding a STAR-View+ File into Microsoft Office 13

Working With Scenes 22

Closing a File and Exiting STAR-View+ 35

Using STAR-View+

This part of the documentation guides you on the use of a viewer for three-dimensional STAR-CCM+ scenes, STAR-View+. This viewer is only supported on Linux and Windows. The following sections explain:

- What is STAR-View+?
- Installing STAR-View+
- Exporting a scene to a file of the correct format
- Starting STAR-View+
- Working with STAR-View+
- Integrating STAR-View+ in Microsoft Office 2007 and 2010

The other parts of the documentation cover the following:

- An introduction to using STAR-CCM+
- Procedures for using 3D-CAD to build geometry from scratch within the STAR-CCM+ environment
- How to generate or import meshes and work with them
- Modeling physics -- the concepts, theory and implementation details of the models
- Using tools in STAR-CCM+ to facilitate your work
- Detailed information on interacting with CAD/CAE products in a STAR-CCM+ simulation
- Examples of applying the models of STAR-CCM+ in sophisticated cases
- Preparation and procedures for running a STAR-CCM+ simulation
- Analyzing results through visualization and plotting
- A wizard guide that shows how STAR-CCM+ features are combined in customized interfaces for special applications
- A tutorial guide that shows you how to use STAR-CCM+ for various applications in a step-by-step format

What is STAR-View+?

STAR-View+ is a light-weight standalone scene viewer for STAR-CCM+. You can use it to examine a scene as you would in STAR-CCM+, when that scene is exported to a three-dimensional `.sce` graphics file.



As an alternative to running the full STAR-CCM+ client, this file format is strictly for interactive viewing. You cannot make scene modifications as you would in STAR-CCM+, such as adding parts to be displayed, changing the coloring or using a different scalar.

Installing STAR-View+

STAR-View+ is included as part of the default STAR-CCM+ installation on Windows and Linux 32- and 64-bit platforms. It is also available to download as a separate installable package from the CD-adapco website. The link to use is:

http://www.cd-adapco.com/download

As no licensing is required for STAR-View+, it is possible to provide this link to those who wish to view the results of STAR-CCM+ simulations without needing to purchase the software.

Instructions are provided on how to install the downloaded package on Windows and Linux. Note that STAR-View+ requires a Java runtime environment. If necessary, you can provide a custom location for the Java installation using the `STARCCM+_JAVA_HOME` environment variable.

Installing the Downloaded Package on Windows

The steps required to install STAR-View+ on Windows 32- and 64-bit platforms are:

- Download the appropriate package from the CD-adapco website.
- Navigate to the download directory using an appropriate file explorer and run right-click on the `.exe` file to launch the installer.

- A *User Account Control* dialog may appear (in Windows XP this will be a *Security Warning* dialog); this can be safely ignored. Click Yes to proceed to the installation.



- The welcome screen of the installer will appear as shown below.



- Click Next to proceed with the installation. Subsequent pages will be presented to ensure the installation is installed in the correct location, and that appropriate options are provided for launching the tool.

- If you downloaded the 32-bit version, you can integrate STAR-View+ in Microsoft Office 2007 and 2010 by installing the required add-in. Simply tick the Install STAR-View+ Office Add-in checkbox on the *Select Additional Tasks* page.



This will install `Visual Studio Tools for the Office system 3.0 Runtime` and `CD-adapco Office Add-Ins` on your machine.

When the installation is complete, STAR-View+ may be launched as described in the section, Starting STAR-View+.

Installing the Downloaded Package on Linux

The steps required to install STAR-View+ on Linux 32- and 64-bit platforms are:

- Download the required package from the CD-adapco website.
- Unpack the installer using the standard `tar` program supplied with Linux. A typical command would be:

  ```
  tar zxvf [STAR-View+_package_name].tar.gz
  ```

- This will create a directory in the same location as the installer file with the name `starview+X.0X.0XX` where the X's depend on the particular version installed (for example, `starview+8.04.005`). This directory contains all the necessary components to launch STAR-View+.
- To launch STAR-View+, simply enter

  ```
  [DIR]/starview+X.0X.00X/starview+
  ```

  where `[DIR]` is the path to the directory containing the downloaded installer.

Exporting a Scene to a STAR-View+ File

You can export the contents of a scene to a scene file (`.sce`) or a VRML file from STAR-CCM+. You can open the `.sce` file in STAR-View+ and work with the same three-dimensional image that you had in STAR-CCM+. Be sure to use the same version of STAR-View+ as the version of STAR-CCM+ that you used to create the scene file. The VRML file can be viewed using a suitable third-party viewer such as Cortona 3D Viewer. In both cases you can interact with the image to rotate, pan, and zoom.

To save the scene contents to these formats, right-click on the scene node or the scene display, and select Export 3D Visualization File... from the pop-up menu.



This will activate the *Save Scene File* dialog, in which the *Files of type:* filter is used to select the format of the saved file. This is set to the STAR-View+ (*.sce) format by default. To export the scene as a VRML file, select the VRML (*.wrl) format from the drop-down list.



Navigate to the location where you want the file saved, enter the file name in the *File Name:* text box, and click Save.



This action is disabled if more than one scene node is selected.

Starting STAR-View+

Assuming you have followed the instructions for the installation, start STAR-View+ one of the following ways, depending on your platform:

- On Windows, in the *Start* menu, launch STAR-View+ by selecting its menu item in the CD-adapco submenu.
- On Linux, point to the proper directory for starting STAR-CCM+, and enter the following command line:

  ```
  starview+
  ```

  Alternatively, you can include the name (and path if necessary) of the specific `.sce` file to open:

  ```
  starview+ file.sce
  ```

Working with STAR-View+

This section describes the following:

- Opening a file
- Embedding a STAR-View+ file into Microsoft Office (Windows only)
- Working with scenes
- Closing a file and exiting STAR-View+

Opening a File

To open a `.sce` file in STAR-View+:

- Start the viewer.
- Select File > Open.



- In the *Read STAR-View+ file* dialog, navigate to the `.sce` file. Click Open.



- The scene will be available for three-dimensional viewing.



STAR-View+ opens one file per session. To view multiple files, run multiple sessions.

Embedding a STAR-View+ File into Microsoft Office

You can incorporate your `.sce` file into a Microsoft Office 2007 and 2010 document. Only 32-bit versions of STAR-View+ support this feature. Previous versions of Microsoft Office are not officially supported. This section explains how to embed a STAR-View+ file in the following packages:

- Microsoft PowerPoint
- Microsoft Word
- Microsoft Excel

Embedding a STAR-View+ File in Microsoft PowerPoint 2007 and 2010

To embed a STAR-View+ (.sce) file in Microsoft Office PowerPoint 2007 or 2010, do the following:

- Launch Microsoft PowerPoint 2007 or 2010.

  The first time PowerPoint is launched after installing STAR-View+, the *Microsoft Office Customization Installer* dialog will appear asking for permission to install the STAR-View+ add-in. An example is shown below.



- Click Install.

- Open the presentation file or create a new one.
- Click on the CD-adapco tab on the ribbon.



If you do not have the CD-adapco tab available in your copy of the software, ensure that you installed the 32-bit version of STAR-View+ with the appropriate plug-in.

- Click on the Insert a Scene button.

- In the *File Open* dialog, find and select the `.sce` file, and click Open.



An ActiveX Control object will appear in the slide, containing the name of the `.sce` file.

- To access the `.sce` file, select the Slide Show tab and activate the *Start Slide Show* feature, either from the beginning or, if you have the correct slide active, from the current slide.

This will access the `.sce` file's content. You can move around the scene as you would in STAR-View+ itself.



However, unlike the standalone version, the only pop-up menu item available is Reset View.



- Press `<Esc>` to return to the PowerPoint GUI.

Embedding a STAR-View+ File in Microsoft Word 2007 and 2010

To embed a STAR-View+ (.sce) file in Microsoft Office Word 2007 or 2010, do the following:

- Launch Microsoft Word 2007 or 2010.

The first time Word is launched after installing STAR-View+, the *Microsoft Office Customization Installer* dialog will appear asking for permission to install the STAR-View+ add-in. An example is shown below.



- Click Install.
- Open the document file or create a new one.
- Click on the CD-adapco tab on the ribbon.



If you do not have the CD-adapco tab available in your copy of the software, ensure that you installed the 32-bit version of STAR-View+ with the appropriate plug-in.

- Click on the Insert a Scene button.

- In the *File Open* dialog, find and select the `.sce` file, and click Open.



The scene will be imported into the document and can be arranged using the Page Layout tab. You can move around the scene as you would in STAR-View+ itself.



However, unlike the standalone version, the only pop-up menu item available is Reset View.



Embedding a STAR-View+ File in Microsoft Excel 2007 and 2010

To embed a STAR-View+ (.sce) file in Microsoft Office Excel 2007 or 2010, do the following:

- Launch Microsoft Excel 2007 or 2010.

The first time Excel is launched after installing STAR-View+, the *Microsoft Office Customization Installer* dialog will appear asking for permission to install the STAR-View+ add-in. An example is shown below.



- Click Install.
- Open the spreadsheet or create a new one.
- Click on the CD-adapco tab on the ribbon.



If you do not have the CD-adapco tab available in your copy of the software, ensure that you installed the 32-bit version of STAR-View+ with the appropriate plug-in.

- Click on the Insert a Scene button.

- In the *File Open* dialog, find and select the `.sce` file, and click Open.



The scene will be imported into the spreadsheet as an image. You can move around the scene as you would in STAR-View+ itself.



However, unlike the standalone version, the only pop-up menu item available is Reset View.

Working With Scenes

Once you have the scene open you can:

- Move around a scene
- Hide and show scene objects
- Restore a scene

Moving Around a Scene

The techniques for examining a scene in STAR-View+ are very similar to interacting with the scene display in the STAR-CCM+ *Graphics* window.



Specifically, the following mouse actions in STAR-View+ replicate those in a STAR-CCM+ scene display:

- Panning
- Rotating
- Rolling
- Zooming

However, vectors in STAR-View+ will look as they did when you exported the scene; the vector lengths will not update when you zoom in. This is due to the fact that exported vector scenes do not retain all vector data.

Panning

To pan the image in the visualization display, hold down the right mouse button and drag the mouse in the direction where you want to move the object.



The tiny colored square that appears while you click and drag represents the point at which you began the movement. This is helpful for knowing where you started the movement. However, the point from which you start panning has no effect on the movement; the movement is governed exclusively by how you reposition the mouse.

The image appears to change its view angle as you pan. This is because the scene's view is in the `Perspective` projection mode, which depicts three-dimensional space. In the `Parallel` projection mode, the view angle stays the same.



You can achieve the same effect with the middle mouse button if you hold down the `<Shift>` key at the same time.

Rotating

STAR-CCM+ uses a point on a part to rotate (left image below). If you select a point somewhere other than on a part (right image below), STAR-CCM+ uses the depth to that last point; the depth is normal to the view plane.





The scene can be rotated around various axes, depending on the direction you move the mouse.

Rotating around a Vertical Axis

Choose a point on or near the object in the visualization display, and click and drag (holding down the left mouse button) in a horizontal direction from that point.



This action rotates the object around the vertical axis. The small colored square represents the point at which you began the click-and-drag motion. The colored number near it indicates degrees of rotation.

Rotating around a Horizontal Axis

Now try clicking and dragging in a vertical direction.



The axis is horizontal.

Rotating Around an Arbitrary Axis

Click and drag at an angle.



The two directions of the axes are combined.

This technique rotates the object around the axis perpendicular to the direction of your drag. You set the location of the axis when you chose a point to begin clicking and dragging.

Rolling

To roll the scene, hold down the `<Ctrl>` key, and click and drag the mouse toward you (downward).



Using the `<Ctrl>` key rotates the object along an axis perpendicular to your screen. The downward movement of your mouse creates a clockwise rolling motion. The movement stops when your mouse pointer goes outside the edges of the visualization display. To roll the object further, simply release the mouse button and move the mouse back to its earlier position.

Now click and drag in the opposite direction.



This results in a counter-clockwise roll.

If you try clicking and dragging left or right in this mode, there is no effect.

Zooming

There are different ways to zoom with the mouse:

- Choosing a focal point
- Rubber banding
- Using the mouse wheel

Choosing a Focal Point

Choose a point inside the display, then click and drag downward using the middle mouse button.



This brings your view closer to that point.

Moving the mouse straight up has the opposite effect.



Rubber Banding

To zoom to a specific area in the display, click the (Rubberband Zoom) button, and then click two points that define the opposing corners of the zoomed area.



Using the Mouse Wheel

Turning the mouse wheel toward you brings your view closer to the visualization display.



The opposite wheel movement takes your view farther away.

The mouse pointer must be within the visualization display for this operation to work. Its particular position within the display determines the direction of this type of zooming.

Hiding and Showing Scene Objects

To hide an object in a scene, you can right-click on that object and select Hide Object.



Alternatively, to hide all the objects associated with a particular displayer, you can right-click and select Hide Displayer.

To restore hidden objects, you can right-click in the scene and select Show all.

Restoring a Scene

To restore the default settings of the view, right-click on the display and select Reset view.



Closing a File and Exiting STAR-View+

To close a file, you can open another file or exit STAR-View+.

You can exit STAR-View+ by:

- Selecting File > Quit from the menu



- Pressing the `<Alt><F4>` keyboard shortcut

Optimate Contents

Optimate 1

Capabilities of Optimate and Optimate+ 2

Optimate Run Modes 3

The Optimate Workflow 6

Run Mode Panel Reference 12

Optimate Variables 14

Outputs Panel Reference 19

The Assembly Panel 21

Run Panel Reference 27

Customizing the Optimate Workflow Using Macros 28

How Does Optimate Work? 30

Running Optimate Jobs 31

Optimate Environment Variables 32

Running Local Jobs in Batch 33

Running Optimate Studies on a Cluster 35

Modifying Project Settings After Running a Job 44

Sharing Designs 45

Running a CAD Robustness Study 46

The SHERPA Algorithm 47

Saving and Loading Optimate Projects 49

Optimate Licensing 50

Post-Processing Optimate Results 51

Optimate Options 53

Optimate Options Panel 54

File Descriptions 56

Troubleshooting 58

Optimate Tutorials 61

Design Sweeps Tutorial 62
Prerequisites 63
Starting Optimate 64
Selecting the Study Type 65
Creating the Variables 66
Selecting Data to Post-Process 67
Defining the Run Settings 68
Running the Study 69
Post-Processing 70

Design of Experiments Tutorial 71
Prerequisites 73
Part 1: Creating a Response Surface 74
Part 2: Running a Full Factorial Design of Experiments Study 79

Pareto Optimization Tutorial 82
Prerequisites 83
Starting Optimate 84
Selecting the Study Type 85
Creating the Variables 86
Selecting Data to Post-Process 87
Defining the Run Settings 88
Running the Study 89
Post-Processing 90

Use Existing Designs Tutorial 93
Prerequisites 94
Preparing An Existing Designs File 95
Selecting the Study Type 97
Creating the Variables 98
Selecting Data to Post-Process 99
Defining the Run Settings 100
Running the Study 101
Post-Processing 102

2D Flap Optimization with Macro Tutorial **104**
Prerequisites 105
Preparing the Post-Processing Macro 106
Starting Optimate 108
Defining the Run Settings 109
Running the Study 110
Post-Processing 111

Optimate

Optimate is an add-on for STAR-CCM+ that allows you to perform design exploration studies (which includes design sweeps, optimization, design of experiments, and robustness analyses) within the STAR-CCM+ environment.

Optimate is installed alongside STAR-CCM+ and is accessed from the toolbar. Optimate uses technology from Red Cedar Technology® to execute, monitor, and post-process the design exploration studies. The SHERPA algorithm is used for the Optimization and Pareto Optimization modes of Optimate+.

Note: Optimate is designed for use when STAR-CCM+ is the only analysis tool in the automation or optimization process. If you want to use an additional tool, you must use an external optimization package. In this case, Optimate can help create the necessary batch scripting to run STAR-CCM+.

See Also

The SHERPA Algorithm

Capabilities of Optimate and Optimate+

There are two variants of Optimate: Optimate and Optimate+.

Optimate provides a fully automated process to gain insight into a product's behavior. Optimate+ provides a fully automated process to find a set of optimum designs that match engineering objectives. The following table provides a list of supported features in each variant:

Functionality	Optimate	Optimate+
Design Sweeps	Yes	Yes
Design of Experiments	Yes	Yes
Stochastic Analysis	Yes	Yes
Optimization using SHERPA	No	Yes
Pareto Optimization using MO-SHERPA	No	Yes

Optimate Run Modes

You can perform a number of studies with Optimate.

Design Sweeps

Design Sweeps are used for design exploration. You can use Design Sweeps mode to replicate test results, create aerodynamic loads databases, create compressor maps, or with any application where you know the configuration for each job before running the analysis. In Design Sweeps mode, you provide the variables (which are stored in a `.csv` file) before starting the analysis. No optimization takes place in Design Sweeps mode.

Design of Experiments

A Design of Experiments study is used for local exploration around a particular design. In this type of analysis, you supply the variable values and Optimate uses Design of Experiments principles to create a response surface. In this way you can assess the sensitivities of design factors and their interaction with each other.

For each DoE, you can select one of the following sensitivity and screening methods:

- *2 Level Full Factorial Runs* — A two-level full factorial DoE studies the various factors, and all possible combinations of the factors, at a specified number of levels. From this type of study, you can gain insight into the sensitivity of each factor and every combination of factors to the design responses. A two-level full factorial study examines all the factors at two levels: low and high. The number of runs that are required for a 2-level full factorial study is 2n, where *n* is the number of DoE factors.
- *3 Level Full Factorial Runs* — A three-level full factorial DoE study is conceptually similar to a two-level full factorial study, with the main difference being that all factors and combinations of factors are evaluated at three levels: low, mid, and high. The number of runs that are required for a three-level full factorial study is 3n, where *n* is the number of DoE factors.

For each DoE, you can select one of the following response surface modelling methods:

- *Latin Hypercube* — The Latin Hypercube sampling option is an efficient method to generate a response surface for the design space. You can use this generated response surface for higher-level DoE (Design of Experiments) study types, like two- or three-level full

factorial studies, since running these studies on a response surface is computationally inexpensive. The minimum number of evaluations to fit a quadratic surface is:

$$\frac{(n+1)\cdot(n+2)}{2}$$

where `n` is the number of DoE factors. You are recommended to evaluate at least two to four more evaluations than the minimum, up to a maximum of two times the minimum.

- *Using Existing Designs* — Allows you to create a response surface that fits a design set specified by you. The designs are specified through a `.csv` file which you can export through HEEDS Post. To export a set of designs from HEEDS Post, open the design table and select **Save Plot Data**.

Stochastic Analysis

A Stochastic Analysis study is used to assess the impact of minor variations in the input values on a particular design. These variations could come from installation or manufacturing tolerances. In a Stochastic Analysis, Optimate applies a specified probability distribution to the inputs. The resulting distribution of the outputs helps you assess the robustness of the design. You can also specify constraints on design responses to obtain a prediction of failure rate.

Optimization Using SHERPA

In an Optimization study, the optimization algorithm selects the input values to meet the specific goals of the analysis. There are two types of optimization studies: Single-Objective and Multi-Objective.

In Single-Objective Optimization, typical goals include minimizing pressure drop, maximizing outlet uniformity, minimizing component temperatures, and minimizing stress. You define the range of permissible values for each variable and the search algorithm chooses any value within the bounds during the analysis. Single-Objective Optimizations use the SHERPA search algorithm by Red Cedar Technology®. See The SHERPA Algorithm.

In Multi-Objective optimization, you optimize based on two competing objectives. Pareto Optimization is a formulation for investigating multi-objective optimization problems. This type of optimization is suited to cases where there are two objectives which are competitive in nature. In such cases, there is no single optimum design but rather there is a curve along which all designs are optimum in one objective for a given value in the other objective. This behavior is called non-dominated sorting. The resulting curve is referred to as a Pareto Front and defines the

optimum trade-off relationship between two objectives. In Pareto Optimization, you are not required to define the weighting for the competing objectives before running the analysis. The study returns a set of all optimal designs that lie on the Pareto Front. Multi-Objective Optimizations use the MO-SHERPA search algorithm by Red Cedar Technology®.

When performing an optimization study, you must define at least one objective. If you define multiple objectives, a linear weighting is used to combine the objectives. Because the weighting of possibly competing objectives is defined a priori, running an optimization analysis returns a single best design.

The Optimate Workflow

This section contains the recommended workflow to set up a study in Optimate.

The recommended process to follow when performing a study in Optimate is provided below:

1. Identify the goals of your study.
2. Create a case in STAR-CCM+, focusing on what you want to achieve in Optimate.

   This stage involves setting up your case in STAR-CCM+ with a suitable mesh and physics setting.

   Create the reports that you want Optimate to extract data from.

   Finally, make sure that you create the relevant scenes that you want Optimate to use for post-processing.
3. Launch Optimate:

   a) From the STAR-CCM+ user interface, click (**Start Optimate**). The *Load Study?* dialog appears.



   b) Select the option that is relevant to your situation.
4. Select the type of study that you want Optimate to run:

   a) In the *Run Mode* tab, set *Run Mode* to the type of run that you want to perform in Optimate. See Run Modes.

b) For each Run Mode, Optimate presents a set of Run Options in the lower part of the panel. See Run Mode Panel Reference.



5. Create the variables that you want Optimate to modify in the study:

   a) In the *Variables* tab, select the desired variable from the appropriate drop-down list and click **Add**.

      For a list of supported variables, see Optimate Variables.

      The selected variable is added to the *Variables* group-box. All active variables are displayed in this group-box.

   b) From the *Variables* group-box, select the variable that you want to configure.

      The options for the selected variable appear to the right of the *Variables* group box.

   c) Set *Definition Mode* to the desired option. See Variable Panel Reference.

   d) Set the parameters for the selected definition mode. See Variable Panel Reference.

   e) Repeat these steps for all other active variables in Optimate.



6. Select the reports, plots, and scenes that you want Optimate to export:

   a) In the *Outputs* tab, select the report that you want Optimate to extract data from and click **Add**.

   If you want to extract data from all the available reports, click **Add All**.

Note: You must select at least one report.

   The selected report is added to the *Recorded Reports* group-box. All active reports are displayed here.

   b) From the *Recorded Reports* group-box, select the report that you want to define in Optimate.

   c) Specify the output options. See Outputs Panel Reference.

   d) Select the desired plots and click **Add**.

   The selected plots are added to the *Exported Plots/Scenes* group-box.

   e) Select the desired scenes and click **Add**.

   The selected scenes are added to the *Exported Plots/Scenes* group-box.



7. Define the run settings for the Optimate study:
   a) In the *Assembly* tab, select the number of cores to run the study on, as well as how many jobs to run at the same time. See Job-Related Options.
   b) Define the STAR-CCM+ run options. See STAR-CCM+ Run Options.
   c) Define the licensing options. See Licensing Options.
   d) Click **Build Sutdy**.

      All necessary files for this study are created. A pop up window appears while the study is building.



8. Run the Optimate study:

   a) In the *Run* tab, from the ribbon, click (**Run**).
   The current STAR-CCM+ server closes and the jobs are submitted in the background. The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the STAR-CCM+ server list.

   b) Click (**Open HEEDS | Post**) to launch HEEDS Post to post-process the results from the study.



9. Post-process the study results in HEEDS Post.

Run Mode Panel Reference

The run options for the various run modes are provided below.

Design Sweeps Run Options

Use CSV file	Specifies the `.csv` file that contains the variables to be used in the Design Sweep.
	An example format is provided below:
	`Eval #, InletPressure, OutletStDev, depth, TranslateZ`
	`1, 1.1737733459e+02, 4.4361197948e-01, 4.0000000000e-02, 0`
	`2, 1.0233065796e+03, 1.3013350964e-01, 8.8000000000e-02, 2.0000000000e-03`

Design of Experiments Run Options

Sensitivity and Screening		
2 Level Full Factorial	Carries out a two-level full factorial DoE (Design of Experiments) study. See Design of Experiments: 2 Level Full Factorial Runs.	
3 Level Full Factorial	Carries out a three-level full factorial DoE (Design of Experiments) study. See Design of Experiments: 3 Level Full Factorial Runs.	
Response Surface Modeling		
Latin Hypercube	Uses the Latin Hypercube sampling option to generate a response surface for the design space. See Design of Experiments; Latin Hypercube.	
Use Existing Designs	Allows you to create a response surface that fits a design set specified by you. The designs are specified through a `.csv` file which you can export through HEEDS Post. To export a set of designs from HEEDS Post, open the design table and select **Save Plot Data**. See Design of Experiments: Using Existing Designs.	

	Existing Design File	Specifies the `.csv` file that contains the designs. An example format is provided below:
		`Eval #, InletPressure, OutletStDev, depth, TranslateZ`
		`1, 1.1737733459e+02, 4.4361197948e-01, 4.0000000000e-02,`

		`0`
		`2, 1.0233065796e+03,`
		`1.3013350964e-01, 8.8000000000e-02,`
		`2.0000000000e-03`

Optimization Using SHERPA Run Options

Single-objective	Specifies a single objective to optimize using the SHERPA algorithm. See Optimization Using SHERPA.	
Multi-objective	Specifies multiple objectives to optimize using the MO-SHERPA algorithm. See Optimization Using SHERPA.	
	Archive size	Specifies the desired number of non-dominating designs in your Pareto Front or trade-off curve. A larger archive size typically produces better resolved trade-off curves, but only if the number of evaluations is increased as well. Typically, you can calculate the optimal archive size as: `Number of evaluations / 100`.
	Stop on Pareto Convergence?	Specifies whether to stop the solver when the Pareto Front converges. Activate this option to force the optimization solver to stop as soon as the Pareto Front converges. If the Pareto Front does not converge, Optimate carries out the number of evaluations as specified in the *Assembly* panel.

Optimate Variables

In Optimate, variables are the values from the baseline STAR-CCM+ simulation that it can modify to verify the effect of the change.

Design Parameters

In 3D-CAD models, the design parameters that you expose are available in Optimate. Changing a design parameter pushes the modified geometry though the mesh pipeline so that it is remeshed. Make sure that range for the variables allows for the geometry to remesh successfully. If a change to a design parameter causes the geometry to become invalid, the Optimate analysis fails to run.

You can create variables from the following 3D-CAD design parameters:

- Scalar quantities — values such as dimension distances, radii, extrusion lengths, taper angles.
- Translation components — Optimate attempts to translate a coordinate system with the same name as the current design parameter. This feature is useful if you want to rotate a component after the translation. It is important to create the translation variables first, before any rotation variables.

Design parameters that originate from CAD Clients are also supported.

Parts for Swapping

You can create part variables using parts that are imported directly into the **Parts** node in the source simulation file. To use parts as variables, you must supply at least two CAD files in a specified directory. If swapping a part causes a mesh operation to return an invalid result, the Optimate analysis fails to run. If you select a NASTRAN (`.nas`) file, the simulation preferred length unit is selected as the import unit.

Motions

Motion variables are created from defined motions in the simulation file. If you are using reference frames for boundary conditions, variables are not created. You must create a motion and use the automatically created reference frame as the boundary condition for the region. You can create variables from the following motion types:

- Rotation rate
- Translation velocity

Regions

Region variables are created from regions that are stored within the source simulation file. Selecting a region populates the *Regions Physics Values* property with the physics values of that region.

Regions Physics Values

You can create variables from physics values that are stored in the *Regions* property within the source simulation file. Only the physics values of regions whose *Method* property is set to **Constant** are supported.

The following region physics values are supported:

- Volumetric heat source

Mesh Values

You can perform mesh dependency studies by varying the base size in your mesh continuum (regions-based meshing) or your automated mesh operation (parts-based meshing).

Boundaries

Boundary variables are created from boundary data that is stored within the source simulation file. Selecting a boundary as a variable populates the *Boundary Physics Values* property with the physics values of that boundary.

Boundary Physics Values

You can create variables from boundary physics values. Only the physics values of boundaries whose *Method* property is set to **Constant** are supported.

The following boundary physics values are supported:

- Velocity magnitude
- Mass flow rate
- Static temperature
- Total temperature
- Static pressure
- Total pressure
- Mach number
- Flow angles

- Mass fractions
- Wall roughness
- Turbulence scalars
- Thermal scalars

Flow Angles

If you want to vary the inlet flow angles, Optimate can modify a local Cartesian coordinate system labeled `wind`. To use this feature, create a coordinate system in the source simulation file and name it `wind`. Also make sure that you align this coordinate system so that the flow with zero pitch angle, zero roll angle, and zero yaw angle is aligned along the positive X axis. Pitch, yaw , and roll are defined as rotations about the Y, Z, and X axis respectively. All reports that remain relative to the *wind* coordinate system must reference the *wind* coordinate system.

Scalar User Field Functions

You can create variables from user field function that are stored within the source simulation file. Only scalar field functions whose *Definition* property contain numeric values are supported. Examples of allowable definitions are:

- `10`
- `-1.325`

Variables Panel Reference

Adjust the variable parameters according to the type of study you want to perform.

Variables

Optimate can access a number of simulation objects in STAR-CCM+ and make changes to their properties or settings.

Design Parameters	In 3D-CAD models, the design parameters that you expose are available in Optimate. See Design Parameters. Design parameters that originate from CAD Clients are also supported.
Parts for Swapping	You can create part variables using parts that are imported directly into the **Parts** node in the source simulation file. See Parts for Swapping.

Motions	Motion variables are created from defined motions in the simulation file. See Motions.
Regions	Region variables are created from regions that are stored within the source simulation file. See Regions.
Region Physics Values	You can create variables from physics values that are stored in the *Regions* property within the source simulation file. See Region Physics Values.
Mesh Values	You can perform mesh dependency studies by varying the *Base Size* in your mesh continuum (regions-based meshing) or your automated mesh operation (parts-based meshing).
Boundaries	Boundary variables are created from boundary data that is stored within the source simulation file. See Boundaries.
Boundary Physics Values	You can create variables from boundary physics values. See Boundary Physics Values.
Scalar User Field Functions	You can create variables from user field function that are stored within the source simulation file. See Scalar User Field Functions.

Definition Mode

The Definition Mode specifies how the variables are defined. The following modes are available for most run types:

Minimum, Maximum, Resolution	Specifies the minimum, maximum, and resolution for the selected variable.	
	Minimum	Specifies the minimum value that Optimate can set the selected variable.
	Maximum	Specifies the maximum value that Optimate can set the selected variable.
	Resolution	Specifies the number of different values to look at for the selected variable.
	Increment	Displays the increment size.
Minimum, Maximum, Increment	Specifies the minimum, maximum, and increment for the selected variable.	
	Minimum	Specifies the minimum value that Optimate can set the selected variable.
	Maximum	Specifies the maximum value that Optimate can set the selected variable.

	Resolution	Displays the number of different values that Optimate will look at for the selected variable.
	Increment	Specifies the increment size for the selected variable.
List	Specifies the values to use for the selected variable. The values must be specified in a comma-separated list. Optimate runs a case for each value in the list. This method is useful if the values are not evenly spaced.	
CSV File Column	Specifies the values for the selected variable using an external file with comma-separated values. You can import the file in the Run Mode panel, when the run mode is set to Design Sweeps. See Design Sweeps. With the file imported, select the column that contains the list of values to use for the variable.	

Definition Mode for Stochastic Analyses

In Stochastic Analyses, you choose the type of distribution you want to define the variable values. There are two types of distribution:

Uniform	Specifies a uniform distribution.	
	Maximum Variation	Specifies the maximum variation, in either direction, from the mean value.
	Value	Specifies the mean value of the variable.
Gaussian	Specifies a Gaussian distribution.	
	Standard Deviation	Specifies the amount of variation from the mean value.
	Value	Specifies the mean value of the variable.

Outputs Panel Reference

In Optimate, an output is the data that you want to extract during the analysis.

The following requirements must be satisfied for the analysis to complete successfully:

- For all analyses, you must record data from at least one report as either an output, objective, or constraint.
- For optimization cases, you must select at least one objective.
- For Pareto optimization cases, you must select at least two objectives.

Outputs

A list of output types is listed below. Select the desired output from the relevant drop-down list and click **Add**.

Reports	Optimate can record data from reports that are stored within source simulation file. This data is recorded at the end of each analysis. The values that are used for the reports are always in units that are specified in the report properties.
Plots	Optimate can record data from plots that are stored within the source simulation file. Optimate records the selected plots/scenes at the end of each analysis.
Scenes	Optimate can record data from scenes that are stored within the source simulation file. Optimate records the selected plots/scenes at the end of each analysis.

Recording Mode Options for Reports

Select an appropriate recording mode for the specified report and adjust the output settings.

Output	Instructs Optimate to record the value of the report.		
Objective	Instructs Optimate to record the value of the report and to evaluate how the report value compares to previous designs.		
	Goal	Specifies the goal of the comparison.	
		Minimize	Specifies that the goal is to minimize the value in the selected report.

		Maximize	Specifies that the goal is to maximize the value in the selected report.
	Normalization	Specifies the normalization. Only use this property if there are multiple objectives. Set this property so that all objectives are the same order of magnitude.	
	Linear Weight	Determines how the objectives are combined to evaluate the performance of each design. Only use this property if there are multiple objectives.	
	Use Baseline	Sets the normalization values automatically for each variable by using the baseline run.	
Constraint	Instructs Optimate to record the value of the report and to determine whether this design is feasible or not. If a particular set of inputs creates an infeasible design, the Optimate solver avoids this part of the design space in subsequent analyses. For each constraint, you set the *Type* and *Limit*.		
	Type	Specifies the type of constraint to define.	
		Less than	Instructs Optimate to class designs with a report value that is less than the value specified in the *Limit* property as infeasible.
		Greater than	Instructs Optimate to class designs with a report value that is greater than the value specified in the *Limit* property as infeasible.
	Limit	Specifies the value to use with the constraint.	
Import from RSM	Specifies the response surface to be used as the source of the study. A dialog appears when you activate this option. Choose the response surface to associate with the selected output.		

The Assembly Panel

The Assembly panel contains the settings that control how the STAR-CCM+ jobs run.

Job-Related Options

These options control how the Optimate jobs are run, and provide information on the job.

Simultaneous Jobs	Specifies how many jobs are run at the same time. When running in Optimization mode, it is recommended that you keep this value as small as practically possible. The SHERPA algorithm works best when it can learn as many times as possible.	
Cores per STAR-CCM+ Job	Specifies how many cores to run with each STAR-CCM+ job.	
Macro Insertion Point	Specifies the insertion point for custom macros into the `StarDriver.java` macro. This option is only available after you activate the **Insert Macro** option. The custom macros are displayed in the *Macros To Be Inserted* list.	
	Before Meshing	Displays the custom macros that are selected to be inserted immediately before executing the meshing pipeline.
	After Meshing	Displays the custom macros that are selected to be inserted after generating the simulation mesh, but before running the solver.
	After Running	Displays the custom macros that are selected to be inserted immediately after the solver has been completed, and before any results (reports, scenes, and plots) are returned to the Optimate solver.
Macros To Be Inserted	Displays a list of macros that you insert into the `StarDriver.java` macro. This list updates according to what you set option you select for *Macro Insertion Point*.	
Number of STAR-CCM+ Runs	Specifies the total number of cases to run for the project. For Design Sweeps, the number of runs is determined by the number of variables and how many values are used for each variable. For Optimization studies, you set the number of runs based on how many cases there is time to run. For a Pareto optimization study, you must run at least 40 cases. For two-level full factorial or three-level full	

	factorial runs, the number of runs is determined by the algorithm and is populated automatically.	
Check Frequency [sec]	Specifies how often Optimate checks each STAR-CCM+ job to see if it has finished. This value is stored locally, so that next time you launch Optimate, the value from your previous session is used. Modify this value according to the expected run times for your studies.	
Maximum Run Time [sec]	Specifies how long Optimate waits for one job to complete before submitting the next one. This value is stored locally, so that next time you launch Optimate, the value from your previous session is used. Modify this value so that it is slightly longer than the time it takes the job to run. When setting this value, keep in mind queue times and possible solver slowdowns.	
Tokens Required	Displays how many Power Tokens are required to run a project with the current settings. This value is determined automatically by multiplying the number of simultaneous jobs by the number of cores per job. If you are using a cluster, this value does not reflect the correct number of tokens you require unless you set *Cores per STAR-CCM+ Job* to reflect the command you provide in the queue configuration window.	
Optimate Save Options	Specifies what designs to save.	
	Last Best Design	Instructs the Optimate solver to remove all design directories except for the one that is identified as the best design (based on the specified objectives). This option is only available when running optimization studies.
	All Best Designs	Instructs the Optimate solver to remove all design directories except for those that are identified as the best design during the history of the project run. For example, the first run of an optimization project is always identified as the first best design, however, as subsequent jobs run, the Optimate solver identifies improved designs and the first design is no longer classed as the best design. Activating this option retains the directory for the first design because it was identified at one point as the best design. This option is only available when running optimization studies.

	All Designs	Instructs Optimate to retain all design directories.
	None	Instructs Optimate to remove all design directories after the outputs of a completed simulation (report values, plots, and scenes) are extracted.

STAR-CCM+ Run Options

The STAR-CCM+ run options specify how STAR-CCM+ jobs are submitted, executed, and saved.

Mesh Mode	Controls whether the design is remeshed.	
	AUTO	Checks the variables and decides whether a remesh is required. If the variables make a direct change to the geometry or a mesh parameter, a remesh is required.
	FORCE	Forces a remesh. Sometimes you have a field function or macro that affects the geometry or mesh, however Optimate does not detect that a remesh is required. If you know that your design requires a remesh, use this option to force the remesh.
	NONE	Does not remesh the geometry.
Remove Invalid Cells	Instructs STAR-CCM+ to remove invalid cells on each individual job that is submitted. It is recommended that you activate this option if the mesh settings in the source simulation file occasionally generate invalid meshes. Invalid meshes are more likely to occur for projects where widely varying geometries are generated and when you use surface wrapping on complex geometry.	
Convert to 2D	Converts each submitted job into a 2D case. The same restrictions for running 2D cases in STAR-CCM+ apply when running Optimate analyses in 2D. The mesh must have a face that is aligned with the X-Y plane. The physics continua, interfaces, and reports are automatically updated to the appropriate forms for 2D simulations. Any scene that is selected in the output tab is automatically updated so that the displayer parts are the newly created 2D regions.	

Save all Simulation Files	Adds the save command to `StarDriver.java` macro which is executed for each simulation in the analysis. The result is that any `star_0/Design<N>/star/` (where `N` is the design number) directories that are retained by Optimate contain the completed simulation file. Any other files (for example `*.trk`) for that particular simulation are also retained. If the **Autosave Batch Runs** option is activated in the simulation file (found under the **File > Auto Save** menu), then a simulation file is retained even if this option is deactivated. It is recommended that you double-check this option before starting the project.
Clear Solution	Specifies whether to clear the solution in STAR-CCM+ before running each study. Activate this option if you want to clear the current solution before running each job. Deactivate this option if you want to use the baseline solution as a starting point for the next job, which can significantly reduce the number of iterations that are required to reach a converged solution.
Insert Macro	Activates the macro insertion portion of the *Assembly* panel. This option allows you to insert custom macros into the `StarDriver.java` macro that is used to execute each STAR-CCM+ simulation. See File Descriptions.
Use Windows Queue	Instructs Optimate to submit the jobs using a Windows Scheduler. See Submitting Optimate Jobs to a Windows Cluster with Remote Execution.
Use Linux Queue	Instructs Optimate to submit the jobs using a generic queuing system for Linux platforms. See Submitting Optimate Jobs to a Linux Cluster with Remote Execution.

Licensing Options

The Optimate licensing options control how licenses are checked out for the study. During run time, Optimate jobs require an Optimate or Optimate+ license in addition to the necessary licenses for starting the STAR-CCM+ jobs. You can use existing STAR-CCM+ license features in addition to Power Tokens. Power Token licensing offers the flexibility of launching multiple simultaneous designs while only requiring a single `ccmpsuite` feature license.

Use Power Tokens	Instructs Optimate to check out Power Tokens for each STAR-CCM+ job based on the number of cores each job will be run on. This option is the default licensing

	scheme for Optimate. When this option is activated, Optimate checks out a single `ccmpsuite` feature license for the duration of the study. Note that if you are using Power Tokens in addition to other license types, such as serial and HPC, or Power Sessions, this `ccmpsuite` feature license is additional to any other feature licenses that are required by individual designs being run using these other licensing schemes.
Use Serial and HPCs	Allows Optimate to check out STAR-CCM+ session licenses and HPC licenses for each STAR-CCM+ job.
Use Power Sessions	Allows Optimate to check out Power Session licenses for each STAR-CCM+ job.
Use Power on Demand	Allows Optimate to check out a Power Session license from the CD-adapco Power-on-Demand server. Activating this option brings up a dialog where you are prompted for a Power-on-Demand key. See Power on Demand.The license server is set to `1999@flex-cd-adapco.com`. You can change the server port and address by modifying the `Process1.in` file after building the project. See File Descriptions.

Directory Options

The directory options specify the location where the study files are stored. You can also specify the location where you have stored your custom macros and CAD geometry files.

Macro Directory	Specifies the directory containing the custom macros. This option is only available when you are using custom macros, and you have activated the **Insert Macro** STAR-CCM+ run option. You can specify the directory by browsing the local file system or manually if the run directory is not currently mounted. Relative paths are acceptable when manually setting this directory.
CAD Directory	Specifies the directory that contains the CAD parts. This option is only available if you have created parts variables in the *Variables* tab and want to swap parts during the project. By default, the directory is set to the same directory as the source simulation file. If your CAD parts are stored in a different directory, update the path to point to the correct directory. You can specify the directory by browsing the local file system or manually if the run

	directory is not currently mounted. Relative paths are acceptable when manually setting this directory. Relative paths are acceptable when manually setting this directory.
Specify Custom Study Directory	Activating this option allows you to specify a custom directory where the project files are stored.
Custom Study Directory	Specifies the custom directory for the project.

Run Panel Reference

The Run panel contains the options to run Optimate.

Open HEEDS | Post	Launches HEEDS Post for post-processing the Optimate study results. See Post-Processing Optimate Results.
Run	Runs the Optimate study.
Run CAD Robustness	Runs a CAD Robustness study to check that a valid geometry can be generated with the specified variables. See Running a CAD Robustness Study.
Stop	Stops the Optimate study. Stopping an Optimate study does not kill the STAR-CCM+ jobs that are already running.
Pause/ Resume	Pauses the Optimate study. Resumes the study from the paused stated.
Restart	Restarts the Optimate study.
Unlock Study	Unlocks the Optimate study to allow you to modify the study parameters. Each study is locked when you run the Optimate solver. See Modifying Project Settings After Running a Job.
Share Designs	Allows you to share your own geometry designs with the Optimate solver. See Sharing Designs.
Print Message File	Prints the contents of the Optimate solver message file, `Agent_Group_1.mes`, to the Optimate Output window. This file contains information about the Optimate job and is useful for debugging jobs that fail to start.
Command	Displays the command that Optimate submits when you run the solver. The information displayed in this window is read-only.
Output	Displays information about the Optimate project. Information about the project is displayed in this window when you build the project and run the Optimate solver.

Customizing the Optimate Workflow Using Macros

You can customize the Optimate workflow using macros to perform additional actions in the STAR-CCM+ simulation.

When you run the Optimate solver, it runs through a set of default steps which are defined in the `StarDriver.java` macro. You can expand the workflow by adding custom macros to the Optimate study. This modifies the `StarDriver.java` macro to reflect your custom workflow. There are three areas where you can expand the workflow: before meshing, after meshing, and after running.

One example of how you can customize the workflow is to convert a 3-D mesh into a 2-D mesh. You can record a macro in STAR-CCM+, where you manually convert the mesh from 3-D to 2-D, then insert this macro into the `StarDriver.java` macro after the meshing stage. When Optimate reaches this stage, it performs the custom steps and then proceeds with the remaining default workflow.

To customize the workflow:

1. In STAR-CCM+, record the steps that you want to include in the custom workflow.
2. In Optimate, set up your study parameters.
3. In the *Assembly* tab, activate the **Insert Macro** option that is located under *STAR-CCM+ Run Options*.

   The macro insertion segment of the panel activates.
4. Set *Macro Insertion Point* to the option that is relevant to your situation:
   - **Before Meshing** — select this option if you want to perform the actions immediately before executing the meshing pipeline.
   - **After Meshing** — select this option if you want to perform the actions immediately after executing the meshing pipeline.

- **After Running** — select this option if you want to perform the actions immediately after the STAR-CCM+ solver completes, and before any results are returned to the Optimate solver.

5. Click ... next to *Macros To Be Inserted*.
6. In the *Select STAR-CCM+ Macros to be Inserted* dialog navigate to the directory where you stored the custom macros and select them.

   The macros are added to the *Macros To Be Inserted* list. Only macros for the selected insertion point are displayed. To view the macros from the other insertion points, switch to appropriate*Macro Insertion Point*.

The custom macros are added to the relevant part of the `StarDriver.java` macro. When you run the Optimate study, the custom steps are run at the appropriate stage of the workflow.

How Does Optimate Work?

Optimate is a plug-in that works together with STAR-CCM+ to help you understand your design space.

Every Optimate study is based on a STAR-CCM+ simulation, so you must first set up your case in STAR-CCM+ before launching Optimate. In Optimate, you choose the type of study to perform. You then create variables from STAR-CCM+ simulation objects that Optimate can modify throughout the course of the study. Then define the aims and goals for your design exploration and set up the run parameters.

When you run an Optimate study, the STAR-CCM+ server for the baseline simulation shuts down to free up the license. The STAR-CCM+ solver launches in batch mode and starts solving the first job in the study. You control the number of jobs that Optimate runs in a study, as well as how many jobs run at the same time, in the *Assembly* panel. When running a particular job, Optimate modifies the simulation settings in STAR-CCM+ (which can include changes to the geometry design, or changes to the physics conditions) and starts recording the outputs. Optimate checks the running jobs to see if they are complete. The frequency with which Optimate checks the jobs is specified in the *Assembly* tab. When a job completes, Optimate records the outputs.

For each analysis, Optimate writes out certain files. See File Descriptions.

Running Optimate Jobs

You can run Optimate jobs locally, on a remote machine, interactively, or in batch mode.

Optimate Environment Variables

Optimate requires certain environment variables in order to run correctly.

When you run Optimate from the interface, these variables are automatically created for you. If you run Optimate from the command prompt, you must create and set these variables manually.

Optimate Environment Variables for Linux Platforms

On Linux platforms, set the following variables:

- `PYTHONHOME=[STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/Python27`
- `HEEDS_ROOT=[STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/LX64/solver`
- `LD_LIBRARY_PATH=[STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/LX64:[STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/LX64/solver`
- `PATH=$PATH:[STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/LX64`

Optimate Environment Variables for Windows Platforms

On Windows platforms, set the following variables:

- `PYTHONHOME=[STAR-CCM+_INSTALL_DIR]\optimate\[VERSION]\Python27`
- `HEEDS_ROOT=[STAR-CCM+_INSTALL_DIR]\optimate\[VERSION]\Win64\solver`
- `PATH=PATH;[STAR-CCM+_INSTALL_DIR]\optimate\[VERSION]\Win64`

Running Local Jobs in Batch

You can run submit jobs in batch by executing the Optimate solver from the command prompt.

To run a job in batch:

1. Make sure that the environment variables that Optimate requires are set. See Optimate Environment Variables.
2. Open a command prompt.
3. Change to the directory that contains the baseline STAR-CCM+ simulation file and the study files.
4. Enter the following command:
   - For Windows:

     ```
     > [STAR-CCM+_INSTALL_DIR]\optimate\[VERSION]
     \Win64\solver\Optimate.exe i=Agent_Group_1.in
     overwrite skipEvalCheck
     ```

   - For Linux:

     ```
     % [STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/LX64/
     solver/Optimate i=Agent_Group_1.in overwrite
     skipEvalCheck
     ```
5. If the previous Optimate job did not exit correctly, delete the lock file named `star0.lock`.
6. Change the extension of `UserDesigns0_initialze` file from `.in_backup_at_[n]_evals` to `.in`.
7. Execute the command again.

Resuming Jobs

If a job aborts for any reason, you can resume it where it picks up from the last completed design.

You can only resume a job from the command prompt and not via the Optimate interface.

To resume a job:

1. Make sure that the environment variables that Optimate requires are set. See Optimate Environment Variables.

2. Type the following in the command prompt:

```
[STAR-CCM+_INSTALL_DIR]/optimate/Ver2014.06/LX64/
solver/Optimate i=Agent_Group_1.in useHistory
skipEvalCheck
```

The study continues from the last completed design.

Extending Jobs

If a job completes but further designs are needed to achieve the desired results, you can extend the job.

You can only resume a job from the command prompt and not via the Optimate interface.

To extend a job:

1. Locate the `Agent_Group_1.in` file in your study directory.
2. Open `Agent_Group_1.in` in a text editor of your choice.
3. Locate the following line:

```
$Start Agent ID, End Agent ID, Number of Designs
Evaluated
0, 0, 20
```

4. Modify the value that corresponds to `Number of Designs Evaluated` and enter the number of additional designs that you require. In the following example, 10 additional designs are requested:

```
0,0,10
```

5. Modify the end of the file to include an `*EXTENDING` command just before the `*END` command:

```
*EXTENDING
*END
```

6. Save the file and close it.
7. Make sure that the environment variables that Optimate requires are set. See Optimate Environment Variables.
8. In your command prompt, enter the following command to extend the job:

```
[STAR-CCM+_INSTALL_DIR]\optimate
\Ver2014.06\Win64\solver\Optimate.exe
i=Agent_Group_1.in
```

Running Optimate Studies on a Cluster

There are two methods that you can use to run an Optimate study on a cluster. In both methods, the STAR-CCM+ solver runs on the cluster. The difference is where the Optimate solver executable is run.

Running the Optimate Solver on a Cluster

With direct execution, you run the Optimate solver executable on the cluster head node, or any compute node which can submit jobs to other nodes. This method is useful if your local machine restarts or shuts down during the night, or if you do not have access to a password-free SSH tunnel to a cluster. Optimate is installed together with STAR-CCM+ on the cluster, meaning that no additional installation or configuration is required to run an Optimate job on a cluster.

Running the Optimate Solver on the Local Machine using Remote Execution

With remote execution, you run Optimate on the local workstation and send STAR-CCM+ data to and from the remote environment. Use remote execution when you have large cases that require STAR-CCM+ to run on a remote workstation or cluster.

With remote execution, you run Optimate on the local workstation and send STAR-CCM+ data to and from the remote environment. Functionally, this means that you can set up and run Optimate locally, while the STAR-CCM+ jobs are automatically run on a separate workstation or cluster.

To perform a remote execution, make sure that you satisfy the following requirements:

- You must have a password-free SSH tunnel set up between the local machine and the cluster. You must be able to directly access the head node (or a compute node that can submit jobscan on a cluster. For detailed instructions on how to set this up, contact your dedicated support engineer.
- You cannot restart or shut down the local machine during the study.
- All files must exist on the local machine so that they can be transferred over to the remote machine as necessary. Consequently, the script that launches Optimate on the remote workstation must also be in the local directory where you build the project.

Running the Optimate Solver on a Linux Cluster

To run an Optimate study on a cluster with a generic queuing system:

1. On your local machine, launch STAR-CCM+ and set up your baseline simulation.
2. On your local machine, launch Optimate and set up your study. See The Optimate Workflow.

Note: Do not run the Optimate solver from the user interface.

3. Build the Optimate study and transfer the study folder to your run directory on the cluster.
4. Copy your run script into the study folder that you copied over to your run directory on the cluster.

   Make sure that you place the script in the folder that contains the study files, `Agent_Group_1.in`, `Agent_Group_1.mes`, etc.
5. On the cluster, set the variables that Optimate requires. See Optimate Environment Variables
6. From the study directory on the cluster, submit the job:

```
[STAR-CCM+_INSTALL_DIR]/optimate/[VERSION]/LX64/
solver/Optimate i=Agent_Group_1.in overwrite
skipEvalCheck
```

   Make sure that you enter the full path to the Optimate solver executable.

Running the Optimate Solver on a Windows Cluster

To run an Optimate study on a cluster with a Microsoft queuing system:

1. On your local machine, launch STAR-CCM+ and set up your baseline simulation.
2. On your local machine, launch Optimate and set up your study. See The Optimate Workflow.

   Do not run the Optimate solver from the user interface.
3. Build the Optimate study and transfer the study folder to your run directory on the cluster.
4. Copy your run script from the cluster into the study folder that you copied over to your run directory on the cluster.

   Make sure that you place the script in the folder that contains the study files, `Agent_Group_1.in`, `Agent_Group_1.mes`, etc.

5. On the cluster, set the variables that Optimate requires. See Optimate Environment Variables
6. From the study directory, submit the job:

```
[STAR-CCM+_INSTALL_DIR]\optimate\[VERSION]
\Win64\solver\Optimate.exe i=Agent_Group_1.in
overwrite skipEvalCheck
```

Make sure that enter the full path to the Optimate solver executable.

Submitting Optimate Jobs to a Linux Cluster with Remote Execution

To submit an Optimate job to a cluster with a generic queuing system using remote execution:

1. On your local machine, launch STAR-CCM+ and set up your baseline simulation.
2. On your local machine, launch Optimate and set up your study variables and outputs. See The Optimate Workflow.
3. If you do not already have a remote profile, create one. See Creating a Remote Profile.
4. Select **Tools > Configure Linux Queue**.
5. In the *Linux Queue Configuration* dialog, activate the **Remote Execution** option.
6. Select your remote profile.
7. Configure the remaining parameters for the queuing system. See Linux Queue Configuration Panel .
8. Click **Save and Close** to close the *Linux Queue Configuration* dialog.
9. In Optimate, click the *Assembly* tab and activate the **Use Linux Queue** option.
10. If you want to use custom macros in your study, transfer these macros to the remote machine and select the macros from that location.

    Make sure that the path displayed in *Macro Directory* points to the directory that contains the macros on the remote machine.
11. Click **Build Study**.

    The study folder is generated on your local machine.
12. Copy the run script from the cluster to your local study folder.

    Make sure that you place the run script in the same folder as `Agent_Group_1.in`, `Agent_Group_1.mes`, etc.

When setting up a generic queue case, Optimate creates an additional file named `optimate.txt`. This file contains an ID flag that must be passed to STAR-CCM+ to access the Power Token licensing scheme. Make sure that the run script incorporates this ID. See Creating the Run Script for Submitting Jobs to a Linux Cluster.

13. Click the *Run* tab and run the Optimate study.

    The study files are sent to the run folder on the cluster. The STAR-CCM+ simulation files remain on the cluster, however the image files are sent back to your local machine where you can view the results and perform post-processing in real-time. The *Output* window displays information about the run, where you can keep track of the jobs locally.

Creating the Run Script for Submitting Jobs to a Linux Cluster

When setting up a generic queue case, Optimate creates an additional file named `optimate.txt`. This file contains an ID flag that must be passed to STAR-CCM+ to access the Power Token licensing scheme. You must incorporate this ID into the run script that is used to call STAR-CCM+.

An example of a c-shell run script is provided below:

```
#!/bin/tcsh
set UUID=`cat optimate.txt`
starccm+ -doeuuid $UUID -batch StarDriver.java -np 2
test.sim >>& test.log
```

Where:

- `set UUID=`cat optimate.txt`` creates the variable with the licensing information that is required to access the Power Token licensing scheme.
- `-doeuuid $UUID` is the command line option that passes the variable to STAR-CCM+.
- `-batch StarDriver.java` is the batch script to run STAR-CCM+.

By default, the run script that Optimate calls is:

```
./ccmprun.csh
```

It is likely that you will append this script to include machine information that is related to the specific cluster you want to use. Also note that the command that executes STAR-CCM+ is set as a background process. This is necessary for the Optimate solver to function correctly.

Linux Queue Configuration Panel Reference

Set up the commands to run on a generic queuing system. The default values reflect the commands to run on a Linux cluster running OpenPBS. Modify the values to reflect the configuration that you plan on running.

There are two methods available.

Using the Generic Scheduler

The options to run using the generic scheduler are displayed below.

Job Submit	Command that is used to submit each job on the cluster. The default value is set up for a cluster running OpenPBS where each job is submitted on one node running four cores per node. ***Note:*** Modifying the number of nodes per job and cores per node does not affect the number of Power Tokens displayed in the assembly tab. However, it does affect the number of Power Tokens that are required to run.
Job Name Specifier	Identifies the name for each submitted job. The default value is set up for a cluster running OpenPBS.
Command to execute STAR-CCM+	Command that is given for each submitted job. By default each job is started by executing the contents of the shell script `./ccmprun.csh`. You must create this shell script for a specific machine. See Creating the Run Script for Submitting Jobs to a Linux Cluster.

Direct Parallel Execution

The options to run using Direct Parallel Execution are displayed below.

Command to submit job to cluster	Specifies the exact command to start a STAR-CCM+ simulation minus any files associated with the submission. Typically this is the command you use in a command prompt to submit to a job scheduler, minus the simulation file, for example: `starccm+ -np 3 -batch StarDriver.java`

	Alternatively, you can use a scheduling tool and a script to submit a job. In this case, include everything but the script file itself. As an example, using OpenPBS you enter: `qsub -l nodes=1:ppn=8`
Run script	Specifies the name of any file that is required in addition to the command you enter in *Command to submit job to cluster*. For the first example above, enter the name of the simulation file: `<simulation file name>.sim` For the second example, enter the name of the run script. The file specified here is copied from the project directory to the directory that is created by Optimate to launch the individual design variant. The complete command executed by Optimate from this directory is the combined command from both this text box and *Command to submit job to cluster*.
Command executed	Specifies the STAR-CCM+ command or the batch script that is used to launch STAR-CCM+ on the cluster. It is important to provide the shared path of the executable for the full installation of STAR-CCM+.
Remote Execution	Toggles the remote execution command. When activated, select your profile from the **Profile** drop-down menu.
Profile	Specifies the profile to use when performing a remote execution. Profiles that you create in the *Remote Host Profiles Manager* dialog are displayed here. See Remote Host Profiles Manager Panel.

Submitting Optimate Jobs to a Windows Cluster with Remote Execution

To submit an Optimate job to a cluster with a Microsoft queuing system using remote execution:

1. On your local machine, launch STAR-CCM+ and set up your baseline simulation.
2. On your local machine, launch Optimate and set up your study variables and outputs. See The Optimate Workflow.
3. Create a remote profile. See Creating a Remote Profile.
4. Select **Tools** > **Configure Microsoft Queue**.

5. In the *Configure Windows HPC Queue* dialog, activate the **Remote Execution** option.
6. Select the remote profile that you created in the previous steps.
7. Configure the remaining parameters for the queuing system. See Configure Windows HPC Queue Panel.
8. Click **Save and Close** to close the *Configure Windows HPC Queue* dialog.
9. In Optimate, click the *Assembly* tab and activate the **Use Windows Queue** option.
10. If you want to use custom macros in your study, transfer these macros to the run folder on the remote machine and select the macros from that location.

    Make sure that the path displayed in *Macro Directory* points to the directory that contains the macros on the remote machine.
11. Click **Build Study**.

    The study folder is generated on your local machine.
12. Copy the run script from the cluster to your local study folder.

    Make sure that you place the run script in the same folder as `Agent_Group_1.in`, `Agent_Group_1.mes`, etc.
13. Click the *Run* tab and run the Optimate study.

    The study files are sent to the run folder on the cluster. The STAR-CCM+ simulation files remain on the cluster, however the image files are sent back to your local machine where you can view the results and perform post-processing in real-time. The *Output* window displays information about the run, where you can keep track of the jobs locally.

Configure Windows HPC Queue Panel Reference

This panel contains the options to configure jobs on a Microsoft cluster.

Queue Submission Command	System command to put a job in the queue. You can use the default value, `job submit` in almost all cases.
Job Name Specifier	Command line argument to identify the name for each submitted job. You can use the default value, `/jobname:`, in almost all cases.
Job Name Prefix	The prefix of each STAR-CCM+ job that is launched. The full name of the job is this prefix, appended by the evaluation number that corresponds to that job.

STAR-CCM+ Execution Command	The STAR-CCM+ command or the batch script that is used to launch STAR-CCM+ on the cluster. It is important to provide the shared path of the executable for the full installation of STAR-CCM+.
Script	Activate this option if you are providing a batch script to launch STAR-CCM+.
Full Command	Displays a preview of the submission command.
Remote Execution	Toggles the remote execution command. When activated, select your profile from the **Profile** drop-down menu.
Profile	Specifies the profile to use when performing a remote execution. Profiles that you create in the *Remote Host Profiles Manager* dialog are displayed here. See Remote Host Profiles Manager Panel.

Creating a Remote Profile

Before performing a remote execution, you must first create a remote profile. This profile stores all the information about the remote machine.

To create a remote profile:

1. In the Optimate environment, select **Tools** > **Manage Remote Profiles**.

   The *Remote Host Profiles Manager* dialog appears.



2. Click **ADD**.

   A new entry is added to the *Remote Host Profiles* group box.

3. Define the profile parameters. See Remote Host Profiles Manager Panel.

4. Click **Save and Close**.

Once you create a remote profile, it is stored so that you can use it for future studies.

Remote Host Profiles Manager Panel Reference

You can manage the remote host profiles from this panel.

Remote Host Profiles	Contains a list of active remote host profiles.
Remote Host Type	Specifies the operating system of the remote machine.
Profile Name	Specifies the name to store for the remote profile in Optimate.
Remote Host Name	Specifies the name of the remote host. If you require the domain name to connect from your local machine, include it here.
Remote User Name	Specifies the user name that is required to log into the remote machine. A password-free SSH must be set up to this user account on the remote machine.
Shared Drive	Specifies whether your local workstation shares the same drive with your remote machine. Activate this option if you can access the same files from both local and remote machines.
Study Path	Specifies the path to the directory that contains your study.
SSH Command	The command that is used to SSH from the local machine to the remote machine.
SCP Command	The command that is used to SCP (copy files) from the local machine to the remote machine.

Modifying Project Settings After Running a Job

After running a study, the Optimate project settings are locked until you reload the project.

You can modify the study settings by unlocking the study. Unlocking a study removes the results from the current session, but not from the disk.

To unlock a project:

- From the *Run* tab, click (**Unlock Study**).

  The *Optimate Message* dialog appears.



- Click **Unlock**.

The results for this session are removed and you can modify the study settings. If you run the solver in this session, you are prompted to back up the old results before running the modified study.

Sharing Designs

You can share your own design with the study within Optimate.

This option allows you to insert your own design into the study. You can insert a design at any time into the study during the run. You can then study the current set of results and make a best guess at a good design. Optimate then evaluates this design and uses the results for the rest of the run.

To share a design:

- From the *Run* tab, click (**Share Design**).

  The *Share Designs Dialog* appears.



- Set the desired values for each variable and click **Share**.

This design runs after the set of current evaluations complete.

Running a CAD Robustness Study

The CAD robustness study checks whether the CAD geometries that are created from the parameters you specify in the Optimate study are valid.

To run a CAD robustness test:

- In Optimate, select **Tools > Options**.
- In the *Options* dialog, set the parameters for the CAD Robustness Study. See CAD Robustness Study Options.
- From the *Run* tab, click **Run > CAD Robustness**.

  Optimate samples the design space using a Latin Hypercube sampling technique.

After you run a CAD Robustness analysis, you can post-process the analysis:

- Launch HEEDS Post.
- In HEEDS Post, locate the objective called **CAD Robustness Report**.

  Any design with a value of 0 for this report means that the design was unable to regenerate the CAD correctly.

Note: If a geometry manages to generate successfully, this does not guarantee that it can be meshed successfully.

The SHERPA Algorithm

SHERPA is a unique search algorithm that is developed by Red Cedar Technology®.

During a single search, SHERPA (which stands for Simultaneous Hybrid Exploration that is Robust, Progressive, and Adaptive), uses multiple search methods simultaneously rather than sequentially. This approach uses the best attributes of each search method. If a particular search method is deemed to be ineffective, SHERPA reduces its participation.

A comparison of a traditional optimization approach and of the SHERPA approach is displayed below.



While running the SHERPA algorithm, a combination of global and local search methods are used. At any given time, the number of

different methods that are used can range between two and ten. Unlike traditional optimization algorithms that require you to tune parameters manually, the tuning parameters in each method that SHERPA uses are modified automatically during the search according to knowledge gained about the nature of the design space. As SHERPA learns more about the design space, it determines when and to what extent to use each search method.

The general advantages of using SHERPA are:

- You do not need to spend time understanding the design space before choosing a suitable algorithm. SHERPA learns about the design space.
- You do not need to have any expertise in optimization algorithms and applications. SHERPA makes all the decisions about which method to use and how to tune them.
- You can define realistic problems that are based on real engineering or business costs and benefits, without being constrained by the limitations of a particular search method. With SHERPA, your analysis can contain much broader goals and can include a larger number of variables compared to traditional approaches.

Saving and Loading Optimate Projects

The Optimate study settings are stored in an external file, called the Optimate project (`.optm`) file.

You can save or load the Optimate project file.

Saving the Optimate Project

To save an Optimate project:

- In Optimate, select **File** > **Save**.
- In the *Save* dialog, enter the name and location where you want to save the project.
- Click **Save**.

To save a copy of the project:

- In Optimate, select **File** > **Save As**.
- In the *Save* dialog, enter the new name and location.
- Click **Save**.

Loading an Optimate Project

To load an Optimate project:

- Load STAR-CCM+ and launch Optimate.
- In the *Load Study?* dialog, click **Open an Existing Optimate Study**.

Optimate Licensing

Optimate and Optimate+ are add-on products for STAR-CCM+.

By default, all Optimate and Optimate+ sessions use Power Tokens (`DOEtoken` feature licenses). Power Tokens are a unique and flexible approach to STAR-CCM+ licensing. See Power Tokens.

The following licenses are required to run an Optimate study within STAR-CCM+:

- 1 STAR-CCM+ session license (`ccmpsuite` feature license). Used for the duration of your Optimate study.
- 1 Optimate (`doesuite`) or Optimate+ (`doesuiteplus`) session license. Used for the Optimate solver.
- 1 license for each concurrent configuration. This may be a Power Token (`doetoken`) or STAR-CCM+ session license (`ccmpsuite`)
- You can run each configuration in parallel. Each process requires a Power Token (`doetoken`) or STAR-HPC (`hpcdomains`) license.

You can control the order the licenses are checked out. You can also use Power Tokens for Pre- and Post-Processing sessions in STAR-CCM+.

See Also

Controlling the Check Out Order of DOE Parallel Licensing
Using Power Tokens for Pre- and Post-Processing Sessions
Using Power Tokens as a Parallel License Feature for HPC Runs

Post-Processing Optimate Results

Once a series of jobs have completed, you can post-process the results in HEEDS Post. If you want to have a look at a particular design in more detail, it is recommended that you perform the post-processing in STAR-CCM+.

STAR-CCM+

If you choose to save all simulation files during your Optimate analysis, then the individual simulation files are stored in the following directories:

```
star_0/DesignN/star/
```

Where `N` is the design ID for each job. Each simulation file is appended with a three digit *design ID* so that you can copy them into the same directory for archiving, but can still reference the ID for post-processing in HEEDS Post. For more information on how to save all simulation files, see Assembly Panel.

HEEDS Post

HEEDS Post is a tool for post-processing series of jobs. You can Typically you launch HEEDS Post from within the Optimate environment.

To launch HEEDS Post:

- Launch Optimate and load your project file. See Launching Optimate.
- In the *Assembly* panel, rebuild the project.
- In the *Run* panel, click (**Open HEEDS | post**).

  The active study in Optimate is loaded into HEEDS Post.

If you move the study to a new folder, you must reset the association in HEEDS Post.

To re-associate the study folder:

- Right-click the study name in the tree, select **Find Study Folder**.
- Navigate to the study folder and click **Accept**.

Launching HEEDS Post Using a Shortcut

If you use HEEDS Post often, you can create a shortcut to the HEEDS Post executable and launch it directly with the shortcut, without having to launch Optimate first.

To create a shortcut to HEEDS Post:

- Navigate to `[INSTALL_DIR]\optimate\[VERSION]\Win64\`.
- Right-click the `HEEDSPost.exe` executable and select **Send to > Desktop (create shortcut)**.

  A shortcut appears on your desktop.
- Double-click the shortcut to launch HEEDS Post.
- In HEEDS Post, select **File > Open** and select the desired Optimate project file (`*.optm`).

Optimate Options

You can configure the behavior of Optimate through the Options dialog.

From the Options dialog, you can instruct Optimate to use an external solver, specify the study options fo the CAD robustness study, and add STAR-CCM+ command line options.

To access the the dialog:

- In Optimate, click **Tools** > **Options**.

  The *Options* dialog appears.



- Modify the options. See Optimate Options Panel.
- Click **Save and Close**.

Optimate Options Panel

You can control the behavior of Optimate through the Options panel.

Optimate Options

The Optimate options affect the behavior of Optimate.

Always On Top	Controls whether the Optimate dialog always sits on top of all other dialogs and windows, even when you click outside the Optimate panel.
Hide Messages	Controls whether Optimate displays error messages that are related to incorrect inputs and execution errors in a pop-up dialog that you must close interactively. If you do not want Optimate to display these types of messages, deactivate this option.
Use External Optimizer	Allows you to choose an external optimizer. Activating this option clears all current variables.
Use Custom Solver	Allows you to choose a different version of the Optimate solver. In the dialog that appears, navigate to the directory that contains the desired version of the Optimate solver executable and select it. When you next run an Optimate study, this version of the solver will be used.

CAD Robustness Study Options

The CAD Robustness Study Options affect the behavior of the CAD robustness study.

Sampling Intensity	Controls the number of samples that are run for a CAD robustness study. You control the number of samples with the slider. The higher the sampling intensity, the greater the number of samples that are picked throughout the design space.
Concurrent Runs	Controls the number of CAD robustness evaluations that are run in parallel.

STAR-CCM+ Command Line Options

The STAR-CCM+ command line options specify additional command line arguments for launching each STAR-CCM+ evaluation. By default, Optimate uses power tokens before using HPC's for each additional

CPU withing an evaluation. For more information on how to modify licensing behavior, see Power Tokens.

File Descriptions

File descriptions are provided below to help you debug cases that are not executing as expected.

Build Projects

When you build an Optimate project, the following files are found in the simulation directory.

File	Description
`Agent_Group_1.in`	Contains information about the run mode and the number of runs to be completed.
`Assembly.in`	Contains information about the run mode.
`Definitions.in`	Contains information variables are changed at run time. Here, you can find the name of the variable, the minimum value, the maximum value, and the baseline value.
`M_{SIM FILE NAME}.in`	Specifies how Optimate modifies the input file for each run such that STAR-CCM+ runs use the updated variable values.
`Performance0.in`	Contains information about the outputs, objectives, and constraints for each project, and tells Optimate where it can find these values after each STAR-CCM+ job completes.
`{SIM FILE NAME}.in`	Contains the variable names and values that are read by STAR-CCM+ at run time. This file is copied to the design directory at run time. Optimate modifies this file with new values. The STAR-CCM+ batch script reads this file and then updates the fields in STAR-CCM+ before running the analysis.
`{SIM FILE NAME}.out`	Written by each STAR-CCM+ job at completion and contains the values of the outputs that are read by Optimate. If the STAR-CCM+ analysis

File	Description
	does not complete successfully, this file is not written and Optimate marks the analysis as a failure.
`Process1.in`	Contains the command to start each STAR-CCM+ analysis, either via a queuing system or directly. Also contains information about how often to check the project for completion and the named files to be copied into each design directory.
`Representation0.in`	Contains information about the resolution of each variable.
`StarDriver.java`	Batch script that each STAR-CCM+ job executes. It is responsible for reading the input file, adjusting parameters in STAR-CCM+, remeshing if necessary, running the analysis, and exporting the results.
`UserDesigns0_Initialize.in`	Contains the value of each variable to be used for each STAR-CCM+ job. This file is only written for analyses that are run using Design Sweeps mode. See Design Exploration.

Running and Completed Projects

The following files are found in the project directory.

File	Description
`Agent_Group_1.mes`	Output from the Optimate process. Contains information about when each STAR-CCM+ job is started and completed. If there are any errors parsing the other input files, you can find a detailed message here.

`star0.gph`
`star0.hst`
`star0.res`
`star0.rpt`
	Post-processing files that are loaded into HEEDS Post. If you run a project on a computer that does not have HEEDS Post, you must copy these four files and the `star0.hds` file to the machine with HEEDS Post.

`star_0/`	Directory that is created at run time and contains all the design directories.

Troubleshooting

You can find information about failed jobs in the files written by Optimate and STAR-CCM+.

Message File

During an Optimate run, Optimate writes to the `Agent_Group_1.mes` file. This file is located in the same directory as the original simulation file. Errors parsing the project files written by Optimate are stored in this location.

A list of commonly encountered errors with suggested solutions is provided below:

Type of Error	Suggested Solution
Licensing errors	• Check the `CDLMD_LICENSE_FILE` variable • Make sure that the license server is running
Errors executing	• Check the command in the `Process1.in` file

STAR-CCM+ Log File

During an Optimate run, Optimate creates directories for each STAR-CCM+ job. The directory for each run are located in the following directory:

```
star_0/DesignN/star/
```

where `N` is the design ID of each job, starting at one and ending with the number of jobs to be run.

If you activate the **Save all Simulation Files** option in the *Assembly* Panel, these folders are not deleted and persist after the run completes. If you deactivate this option, the directories only persist for the duration of the run.

If any job fails, it is placed in the directory:

`star_0/DesignN-ERROR/star/`

The above directories contain all the input files, outputs files, a log file, and the simulation file for each run. If STAR-CCM+ fails, an error message is added to the `star.log` file.

A list of commonly encountered errors with suggested solutions is provided below:

Type of Error	Suggested Solution
Errors getting appropriate licenses	• Change the number of simultaneous jobs in the `Process1.in` file • Change the number of cores per STAR-CCM+ job in the `Process1.in` file • Check the `CDLMD_LICENSE_FILE` variable • Make sure that the license server is running
Errors modifying the 3D-CAD model	• Change values to avoid invalid geometry
Errors meshing	• Check the mesh settings in the original simulation file
Floating point errors	• Review the solver settings • Review the initial conditions from the original simulation file
Errors importing CAD files	• Check the `directoryCAD` variable in the `StarDriver.java` file
Errors exporting scenes	• Check the `directoryCAD` variable in the `StarDriver.java` file
Errors saving simulation file	• Check the `directoryPOST` variable in `StarDriver.java` file
Incompatible version errors	• Modify the `path` environmental variable such that the correct version of STAR-CCM+ is listed first in the path

HEEDS Post

If HEEDS Post fails to open or does not work as expected, you can take a look at the HEEDS Post log file for further information. The log file is called `HEEDSPost.log` and is created in the home directory.

Optimate Is Running But No Designs Are Launched

This issue is typically caused by an incorrect or inappropriate STAR-CCM+ command line option in the **Tools > Options** command line textbox.

To troubleshoot issues like these:

1. Navigate to the `star_0/Design<N>/star/` folder.
2. Open `Tool_star_output.msg` in a text editor of your choice.

   This file contains the output from the command prompt or shell that was used to launch the STAR-CCM+ process. Typically, when a design fails to launch, it includes a message from STAR-CCM+ that indicates the command line arguments cannot be interpreted. This file also provides a reference to the acceptable command line arguments.
3. In Optimate, go to **Tools > Options** and check the contents of the command line.
4. Make sure that any additional command line arguments conform to the guide that is provided in the `Tool_star_output.msg` file.

Optimate Tutorials

This section contains tutorials for Optimate and Optimate+.

Design Sweeps Tutorial

This tutorial demonstrates how to use Optimate to setup a Design Sweeps study.

In this type of study, you describe all of the inputs for each case for STAR-CCM+ to run. This type of study is applicable when recreating test results or creating loads databases.

In this tutorial, you vary the depth of a blockage in a pipe to see how it affects the flow. The depth of the blockage is a design parameter that Optimate can modify during the study.



Prerequisites

This tutorial assumes that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

This tutorial requires an Optimate+ license in order to use the **Optimization using SHEPRA** run mode.

Starting Optimate

Load the starting simulation file and launch Optimate.

To launch Optimate for the study:

1. In your workspace, create a folder named `designSweeps`.
2. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/optimate/data`
3. Copy `pipeBlockage.sim` to your workspace.
4. Launch STAR-CCM+.
5. Select **File > Load Simulation...**
6. In the *Load Simulation* dialog, click **Browse**.
7. In the *Open* dialog, select `pipeBlockage.sim` from your workspace and click **Open**.
8. In the *Load Simulation* dialog, click **OK**.
9. From the STAR-CCM+ toolbar, click (**Start Optimate**).
   The *Load Study?* dialog appears.



10. Click **Create a new Optimate Study**.
    The Optimate workspace appears.

Selecting the Study Type

To select the type of study to perform:

1. In the Run Mode tab, make sure that *Run Mode* is set to **Design Sweeps**.
2. Save the Optimate study file by selecting **File > Save As** from the menu.
3. In the *Save* dialog, save the file as `designSweeps.optm`.

Creating the Variables

Create the variables for your study. In this tutorial, you want to vary the depth of the blockage in the pipe from 0.08 meters to 0.011 meters.

1. Click the *Variables* tab.
2. Select **depth** from the *Design Parameters* drop-down list and click **Add**.

   The design parameter, **depth**, is added to the *Variables* group box and is identified as a variable that Optimate can modify.
3. Make sure that *Definition Mode* is set to **Minimum, Maximum, N**.

   This option allows you to specify the minimum and maximum size for the chosen variable, as well as the number of increments to use.
4. Set *Minimum* to `0.08`.

   The units that are used are taken from the simulation file. In this case, the units are meters.
5. Set *Maximum* to `0.11`.
6. Set *N* to `7`.

   This property instructs Optimate to execute seven STAR-CCM+ runs for this variable. The values are equally spaced starting at 0.08 meters and ending at 0.11 meters.

Selecting Data to Post-Process

Select what data you want Optimate to save during the analysis. This data is used for post-processing in a later step.

1. Click the *Outputs* tab.
2. Click **Add All** next to the *Reports* drop-down menu.

   All reports from the simulation file are added to the *Recorded Reports* group box. Optimate will save data from these reports during the analysis.
3. From the *Recorded Reports* group box, select **InletPressure**.

   The *Recording Mode* properties for this report are displayed.
4. Set *Recording Mode* to **Objective**.

   The *Goal*, *Normalization*, and *Linear Weight* properties appear.
5. Make sure that *Goal* is set to **Minimize**.

   Even though there will be no optimization for this study, the results are sorted based on the performance of each case, ranging from the case with the smallest inlet pressure to the case with the largest inlet pressure.
6. Select **Residuals** from the *Plots* drop-down list and click **Add**.

   The **Residuals** plot is added to the *Exported Plots/Scene* group box. Optimate exports data from this plot after each STAR-CCM+ run.
7. Select **Velocity** from the *Scenes* drop-down list and click **Add**.

   The **Velocity** scene is added to the *Exported Plots/Scene* group box. Optimate exports data from this scene after each STAR-CCM+ run.

Defining the Run Settings

Specify the run settings for the solver. To reduce the time it takes to complete the study, you can run multiple jobs at the same time.

1. Click the *Assembly* tab.
2. Set *Simultaneous Jobs* to `2`.

   This property instructs Optimate to run two jobs at the same time. Note that your machine must have at least two cores spare. If you have more cores, and sufficient tokens, you can increase this value or increase the number of cores that are used per job through the *Cores per STAR-CCM+ job* property. The total number of tokens required to run your study is displayed under the *Tokens Required* property.
3. Set *Check Frequency* to `10`.

   This property instructs Optimate to check for completed jobs every ten seconds, which for this small analysis is reasonable.
4. Make sure that **Save all Simulation Files** is activated.

   This option instructs STAR-CCM+ to save a separate simulation file for each case.
5. Activate the **Clear Solution** option.

   This option instructs STAR-CCM+ to clear the solution before each new case is run.
6. Make sure that *Optimate Save Option* is set to **All Designs**.

   This option instructs Optimate to save all design iterations.
7. Save the study.
8. Click **Build Study**.

   All necessary files for this study are created. A pop up window appears while the study is building.

Running the Study

The study is now ready to be run. To start the job on your local machine:

1. Click the *Run* tab.
2. Click (**Run**) and select **Run**.

   The current STAR-CCM+ server closes and the jobs are submitted in the background. The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the server list.

Post-Processing

Post-process the results in HEEDS Post.

Create a Model View plot to visualize the different designs:

1. Click (**Open HEEDS | Post**) to launch HEEDS Post.
2. From the HEEDS Post ribbon, in the **General Plots** panel, click (**Model View**).

   The *Create New Plot Wizard* dialog appears.
3. From the *Available files* group box, select **Velocity-scene.png**.
4. Click **Finish**.

   A velocity scene is added to the HEEDS Post environment.
5. Select **Model_5** from the **Plots** folder,
6. In the HEEDS Post ribbon, in the *Series Data* panel, cycle through the designs to see how the changes to the design affect the flow in the pipe.



Design of Experiments Tutorial

This tutorial demonstrates how to use Optimate set up a Design of Experiments study.

In this tutorial, you first create a response surface that is representative of the physical problem and then proceed to carrying out a full factorial Design of Experiments study on the generated response surface. The aim of this study is to gain sensitivity information about the problem.

The geometry in this tutorial consists of an aircraft wing with a double-slotted flap system.



Optimate can modify the X and Y translation of each flap, as well as the rotation of the flap.


Leading Flap
Trailing Flap
Translation
Rotation

Prerequisites

This tutorial assumes that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

Part 1: Creating a Response Surface

In the first part of this tutorial you create the response surface.

Starting Optimate

Load the starting simulation file and launch Optimate.

To launch Optimate for the study:

1. In your workspace, create a folder named `DoE`.
2. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/optimate/data`
3. Copy `2dAirfoilFlapDoE.sim` and `macroToUpdatePlots.java` to your workspace.
4. Launch STAR-CCM+.
5. Select **File > Load Simulation...**
6. In the *Load Simulation* dialog, click **Browse**.
7. In the *Open* dialog, select `2dAirfoilFlapDoE.sim` from your workspace and click **Open**.
8. In the *Load Simulation* dialog, click **OK**.
9. From the STAR-CCM+ toolbar, click (**Start Optimate**).
   The *Load Study?* dialog appears.



10. Click **Create a new Optimate Study**.
    The Optimate workspace appears.

Selecting the Study Type

To select the type of study to perform:

1. In the *Run Mode* tab, set *Run Mode* to **Design of Experiments**.

2. In the *Run Options* panel, make sure that *Response Surface Modeling* is set to **Latin Hypercube**.
3. Save the Optimate study by selecting **File > Save As** from the menu.
4. In the *Save* dialog, save the file as `responseSurface.optm`.

Creating the Variables

Specify the design parameters for your study. In this tutorial, you want to vary the position of the leading and trailing flaps of the wing.

1. Click the *Variables* tab.
2. Select **LeadingFlapRotateAngle** from the *Design Parameters* drop down menu and click **Add**.

   This value corresponds to the rotation angle of the leading flap.

   The **LeadingFlapRotateAngle** variable is added to the *Variables* group-box. All active variables are displayed here.
3. Make sure that *Definition Mode* is set to **Minimum, Maximum, Resolution**.
4. Set *Minimum* to `-4.0`.
5. Set *Maximum* to `4.0`.
6. Leave the value of *Resolution* to its default value.
7. Using the same technique, define the following variables:

Variable	Definition Mode	Minimum	Maximum
TrailingFlapRotateAngle	**Minimum, Maximum, Resolution**	`-15.0`	`25.0`
TranslateTrailingFlapX	**Minimum, Maximum, Resolution**	`-0.02`	`0.02`
TranslateTrailingFlapY	**Minimum, Maximum, Resolution**	`-0.01`	`0.01`

Selecting Data to Post-Process

For this study, you examine the sensitivity to all of the available outputs.

1. Click the *Outputs* tab.
2. Click **Add All** next to the *Reports* drop-down menu.

All available reports are added to the *Recorded Reports* group box.

3. Select a report form the *Recorded Reports* group box to update the entries in the Optimate workspace.
4. Click **Add All** next to the Plots drop-down menu.

   All available plots are added to the *Exported Plots/Scenes* group box. Recording reports is useful for assessing the convergence of each study.
5. Select **Scalar - Mach** from the *Scenes* drop-down menu and click **Add**.

   The selected scene is added to the *Exported Plots/Scenes* group box.
6. Using the same technique, select **Scalar - Pressure**.

Defining the Run Settings

Specify the run settings for the solver. To reduce the time it takes to complete the study, you can run multiple jobs at the same time.

1. Click the *Assembly* tab.
2. Set *Simultaneous Jobs* to `2`.

   This property instructs Optimate to run two jobs at the same time. Note that your machine must have at least two cores spare. If you have more cores, and sufficient tokens, you can increase this value or increase the number of cores that are used per job through the *Cores per STAR-CCM+ job* property. The total number of tokens required to run your study is displayed under the *Tokens Required* property.
3. Set *Check Frequency [sec]* to `10`.

   This property instructs Optimate to check for completed jobs every ten seconds, which for this small analysis is reasonable.
4. Activate the **Clear Solution** option.

   This option instructs STAR-CCM+ to clear the solution before it runs each new case.
5. Set *Number of STAR-CCM+ Runs* to `20`.

   This option instructs STAR-CCM+ to run 20 different cases. The minimum number of evaluations for 4 factors is (4+1)*(4+2)/2, which is 15. To ensure a good fit, however, you can run 20 evaluations.
6. Save the study.
7. Click **Build Study**.

All necessary files for this study are created. A pop up window appears when this process is complete.

Running the Study

The study is now ready to be run. To start the job on your local machine:

1. Click the *Run* tab.
2. Click (**Run**) and select **Run**.

   The current STAR-CCM+ server closes and the jobs are submitted in the background. The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the server list.

Post-Processing

Visualize the 3D response surface in HEEDS Post.

1. Click (**Open HEEDS | Post**) to launch HEEDS Post.
2. From the HEEDS Post ribbon, in the **DOE Plots** panel, click (**3D Response Surface**).

   The *Create New Plot Wizard* dialog appears.
3. Set *X-Axis Parameter* to **TrailingFlapRotateAngle**.
4. Set *Y-Axis Parameter* to **TranslateTrailingFlapX**.
5. Select the coefficient of lift **Cl** as the response.
6. Leave *Fit model* as **Linear**.
7. Click **Finish**.

   The 3D response surface is created.



You can now use the created response surface for further DoE or Stochastic analyses. Close HEEDS Post, saving the results, but leave Optimate running.

Part 2: Running a Full Factorial Design of Experiments Study

Use the response surface that you generated in part 1 to run a full factorial Design of Experiments study.

Modifying the Study

Modify the study to run a 2 level full factorial study. After running a study, Optimate locks the study parameters. You can unlock the study to modify the parameters.

1. In Optimate, in the *Run* tab, click (**Unlock Study**).
2. In the dialog that appears, click **Unlock**.

   The study unlocks and you can now modify the study parameters.
3. Click the *Run Mode* tab and set *Sensitivity and Screening* to **2 Level Full Factorial**.
4. Click the *Variables* tab and make sure that the variables are as defined in Part 1.

   The variables and variable ranges must match those in the response surface.
5. Click the *Outputs* tab.
6. In the *Recorded Reports* group box, select **Cl** and activate the **Import from RSM** option.
7. In the *Open* dialog, navigate to the `2dAirfoilFlapDoE_responseSurface.optm` folder and open `Cl-Quadratic-0.mdl`.
8. Using the same technique, import the RSM data for the remaining reports:

Report	RSM Data
Cd	`Cd-Quadratic-0.mdl`
Mass Cons	`Mass Cons-Quadratic-0.mdl`

   A message appears informing you that all outputs are calculated from the generated response surfaces. As you are not running the CFD analysis again, the plots and scenes are no longer valid.
9. Click **OK** to close the message.

   All plots and scenes are removed from the project automatically.

10. Click the *Assembly* tab.

    Simultaneous Jobs and *Number of STAR-CCM+ Runs* are automatically defined to match the requirements of a two-level full factorial run.
11. Activate the **Insert Macro** option.

    This option activates the macro insertion segment of the panel.
12. Set *Macro Insertion Point* to **After Meshing**.

    This option runs the inserted macros after the simulation mesh is generated.
13. Click (**Custom Editor**) next to *Macro To Be Inserted Point*.
14. In the *Select STAR-CCM+ Macros to be Inserted* dialog, navigate to your workspace folder for this tutorial.
15. Select `macroToUpdatePlots.java` from your workspace and click **Select**.

    This macro converts the 3D mesh into a 2D mesh. The path to this Java macro is automatically added to *Macro Directory*.
16. Save the study as `responseSurface_FullFactorial.optm`.
17. Click **Build Study**.

    All necessary files for this study are created. A pop up window appears while the study is building.

Running the 2 Level Full Factorial Study

The 2 Level Full Factorial study is now ready to be run. To start the job on your local machine:

1. Click the *Run* tab.
2. Click (**Run**) and select **Run**.

   The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the server list.

Post Processing the Full Factorial Study

Post-process the results of the full factorial study.

1. Click (**Open HEEDS | Post**) to launch HEEDS Post.
2. In HEEDS Post, select the **Plots > Effects Cl** node.

   This graph shows the effect of each factor on Cl, as the factors go from their low to high values. The higher the slope of a specific curve, the more effect that factor has on Cl.



Similar graphs are automatically generated for Cd and Mass Conservation as well.

Pareto Optimization Tutorial

This tutorial demonstrates how to use Optimate to set up a Pareto Optimization study.

In this type of study, you define the range for the geometry modifications and competing objectives. The MO-SHERPA algorithm then finds the set of designs that are on the Pareto front.

In this tutorial, Optimate varies the depth of a blockage in a pipe, as well as the position of the blockage along the pipe, with the aim of minimizing the pressure at the inlet.



Prerequisites

This tutorial assumes that you are already familiar with certain techniques in STAR-CCM+.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+

This tutorial requires an Optimate+ license in order to use the **Optimization using SHEPRA** run mode.

Starting Optimate

Load the starting simulation file and launch Optimate.

To launch Optimate for the study:

1. In your workspace, create a folder named `paretoOptimization`.
2. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/optimate/data`
3. Copy `pipeMixingBlockage.sim` to your workspace.
4. Launch STAR-CCM+.
5. Select **File > Load Simulation...**
6. In the *Load Simulation* dialog, click **Browse**.
7. In the *Open* dialog, select `pipeMixingBlockage.sim` from your workspace and click **Open**.
8. In the *Load Simulation* dialog, click **OK**.
9. From the STAR-CCM+ toolbar, click (**Start Optimate**).
   The *Load Study?* dialog appears.



10. Click **Create a new Optimate Study**.
    The Optimate workspace appears.

Selecting the Study Type

To select the type of study to perform:

1. In the *Run Mode* tab, set *Run Mode* to **Optimization using SHERPA**.
2. Set *Run Options* to **Multi-Objective**.
3. Make sure that *Archive* is set to `20`.
4. Make sure that the **Stop on Pareto Convergence?** option is deactivated.
5. Save the Optimate study file by selecting **File > Save As** from the menu.
6. In the *Save* dialog, save the file as `paretoOptimization.optm`.

Creating the Variables

Specify the design parameter for your study. In this tutorial, you want to vary the depth of the blockage in the pipe as well as the position of the blockage along the pipe.

1. Click the *Variables* tab.
2. Select **depth** from the *Design Parameters* drop-down list and click **Add**.

   The design parameter, **depth**, is added to the *Variables* group box and is identified as a variable that Optimate can modify.
3. Make sure that *Definition Mode* is set to **Minimum, Maximum, Resolution**.

   This option allows you to specify the minimum and maximum size for the chosen variable, as well as the number of increments to use.
4. Set *Minimum* to `0.04`.

   The units that are used are taken from the simulation file. In this case, the units are meters.
5. Set *Maximum* to `0.14`.
6. Set *Resolution* to `101`.

   This property instructs Optimate to execute 101 STAR-CCM+ runs for this variable. The values are equally spaced starting at 0.04 meters and ending at 0.14 meters.
7. Using the same technique, define the following variable:

Variable	Definition Mode	Minimum	Maximum	Resolution
TranslateZ	**Minumim, Maximum, Resolution**	`0.0`	`0.2`	`101`

Selecting Data to Post-Process

Select what data you want Optimate to save during the analysis. This data is used for post-processing in a later step.

1. Click the *Outputs* tab.
2. Select **InletPressure** from the *Reports* drop-down list and click **Add**.

   The **InletPressure** report is added to the *Recorded Reports* group box.
3. From the *Recorded Reports* group box, select **InletPressure**.

   The *Recording Mode* properties for this report are displayed.
4. Set *Recording Mode* to **Objective**.

   The *Goal*, *Normalization*, and *Linear Weight* properties appear.
5. Make sure that *Goal* is set to **Minimize** and that the **Use Baseline** option is activated.

   These options instruct Optimate to aim towards minimizing the pressure at the inlet. The **Use Baseline** option automatically sets the normalization values for each variable by using the baseline simulation.
6. Using the same technique, define the following report output:

Report	Recording Mode	Goal	Use Baseline Option
OutletStDev	**Objective**	**Minimize**	Activated

7. Select **Residuals** from the *Plots* drop-down list and click **Add**.

   The **Residuals** plot is added to the *Exported Plots/Scene* group box. Optimate will export data from this plot after each STAR-CCM+ run.
8. Select **Velocity** from the *Scenes* drop-down list and click **Add**.

   The **Velocity** scene is added to the *Exported Plots/Scene* group box. Optimate will export data from this scene after each STAR-CCM+ run.

Defining the Run Settings

Specify the run settings for the solver. For Pareto optimization studies, you require a minimum of 40 jobs. This values is automatically set in the *Number of STAR-CCM+ Runs* parameter. To reduce the time it takes to complete the study, you can run multiple jobs at the same time.

1. Click the *Assembly* tab.
2. Set *Simultaneous Jobs* to `2`.

   This property instructs Optimate to run two jobs at the same time. Note that your machine must have at least two cores spare. If you have more cores, and sufficient tokens, you can increase this value or increase the number of cores that are used per job through the *Cores per STAR-CCM+ job* property. The total number of tokens required to run your study is displayed under the *Tokens Required* property.
3. Set *Check Frequency* to `10`.

   This property instructs Optimate to check for completed jobs every ten seconds, which for this small analysis is reasonable.
4. Activate the **Clear Solution** option.

   This option instructs STAR-CCM+ to clear the solution before it runs each new case.
5. Save the study.
6. Click **Build Study**.

   All necessary files for this study are created. A pop up window appears when this process is complete.

Running the Study

The study is now ready to be run. To start the job on your local machine:

1. Click the *Run* tab.
2. Click (**Run**) and select **Run**.

   The current STAR-CCM+ server closes and the jobs are submitted in the background. The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the server list.

Post-Processing

For this analysis, the Pareto front is automatically plotted as **InletPressure vs. OutletStDev**.

1. Click (**Open HEEDS | Post**) to launch HEEDS Post.
2. In HEEDS Post, click the **Plots > InletPressure vs. OutletStDev** node to display the Pareto front.



Create a 2D relation plot to show the effect of the inputs on the outputs.

3. From the HEEDS Post ribbon, click the arrow next to the **Create Plot** icon and select **2D Relation**.

   The *Create New Plot Wizard* dialog appears.
4. Set *X-Axis Parameter* to **depth**.
5. Set *Y-Axis Parameter* to **InletPressure** and **OutletStDev**.



6. Click **Finish**.
   The 2D relation plot is created.



7. Repeat the same process, using **TranslateZ** as the *X-Axis Parameter*.



8. Click **File** > **Save** to save the results in HEEDS Post.

Use Existing Designs Tutorial

This tutorial demonstrates how to use Optimate to setup a Design of Experiments tutorial using existing designs.

In this tutorial, you first create a response surface representative of the physical problem, and then conduct a full factorial DoE on the generated response surface to gain sensitivity information about the problem.

Prerequisites

This tutorial assumes that you are already familiar with certain techniques in STAR-CCM+ and Optimate.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+
Running a Pareto optimization study in Optimate	Pareto Optimization Tutorial

This tutorial makes use of the result files from the Pareto Optimization Tutorial. Make sure that you have completed that tutorial first before attempting the current tutorial.

Preparing An Existing Designs File

1. In your workspace, create a folder named `existingDesigns`.
2. Navigate to the `paretoOptimization` folder that you created in the Pareto Optimization tutorial.
3. Copy over the `paretoOptimization.optm` and `pipeMixingBlockage.sim` files into your `existingDesigns` folder.
4. Launch STAR-CCM+ and load `pipeMixingBlockage.sim`.
5. From the STAR-CCM+ toolbar, click (**Start Optimate**).

   The *Load Study?* dialog appears.



6. Click **Open an Existing Optimate Study**.
7. In the *Select Optimate Study File* dialog, select `paretoOptimization.optm` and click **Open**.

   The Optimate workspace appears and the settings that are contained in the project file are loaded.
8. In the *Assembly* tab, rebuild the study.
9. In the *Run* tab, click (**Open HEEDS | Post**) to launch HEEDS Post.
10. In HEEDS Post, select the **Plots** > **Design Table** node.

    The design table is displayed.
11. From the ribbon, in the **Tools** panel, click **Export Data** > **Save Plot Data...**.



12. In the dialog that appears, save the design data as `designTableData.csv`.

    Save this data in your workspace.
13. Close HEEDS Post without saving the changes.

Selecting the Study Type

Modify the run mode so that the project uses the data from the exported design table.

1. Click the *Run Mode* tab.
2. Set *Run Mode* to **Design of Experiments**.
3. In the dialog that appears, click **Yes** to confirm the changes.

   All variables are removed.
4. Set *Response Surface Modeling* to **Use Existing Designs**.
5. In the *Open* dialog, navigate to your workspace and import `designTableData.csv`.

   The full path to the design file is displayed in *Existing Design File*.
6. Save the study as `designOfExperiments_existingDesigns.optm`.

Creating the Variables

Create the same variables as contained in the existing designs table:

1. Click the *Variables* tab.
2. Select **depth** from the *Design Parameters* drop-down list and click **Add**.

   The design parameter, **depth**, is added to the *Variables* group box and is identified as a variable that Optimate can modify.
3. Select **depth** from the *Variables* group box and set *Definition Mode* to **CSV File Column**.

   The *CSV File Column* property appears.
4. Set *CSV File Column* to **depth**.
5. Using the same technique, create the **TranslateZ** variable:

Variable	Definition Mode	CSV File Column
TranslateZ	**CSV File Column**	**TranslateZ**

Selecting Data to Post-Process

Set the outputs to match those in the existing designs table:

1. Click the *Outputs* tab.
2. Select **InletPressure** from the *Reports* drop-down list and click **Add**.
   The **InletPressure** report is added to the *Recorded Reports* group box.
3. Select **InletPressure** from the *Recorded Reports*.
   Definition Mode is set to **Output**.
4. Set *CSV File Column* to **InletPressure**.
5. Using the same technique, define the **OutletStDev** output:

Report	Recording Mode	CSV File Column
OutletStDev	**Output**	**OutletStDev**

Defining the Run Settings

Rebuild the study:

1. Click the *Assembly* tab.
2. Save the study.
3. Click **Build Study**.

   All necessary files for this study are created. A pop up window appears while the study is building.

Running the Study

The study is now ready to be run. To start the job on your local machine:

1. Click the *Run* tab.
2. Click (**Run**) and select **Run**.

   The current STAR-CCM+ server closes and the jobs are submitted in the background. The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the server list.

Post-Processing

Visualize the 3D response surface in HEEDS Post.

1. Click (**Open HEEDS | Post**) to launch HEEDS Post.
2. From the HEEDS Post ribbon, in the **DOE Plots** panel, click (**3D Response Surface**).
   The *Create New Plot Wizard* dialog appears.
3. Set *X-Axis Parameter* to **depth**.
4. Set *Y-Axis Parameter* to **TranslateZ**.
5. Select the inlet pressure **InletPressure** as the response.
6. Leave *Fit model* as **Linear**.
7. Click **Finish**.
   The 3D response surface is created.



To view the discrete design data overlaid on the surface:

8. From the HEEDS Post ribbon, in the *Series Data* panel, click (**Add New Series**).
9. In the *Add Series to Plot* dialog, set *Type* to **Design data**.
10. Set *Design Set* to **All Designs**.
11. Click **OK**.

The design data is overlaid on the 3D response surface.



2D Flap Optimization with Macro Tutorial

This tutorial demonstrates how to use Optimate to set up an optimization on a 2D simulation.

As part of the tutorial, you will record a macro and insert it into the project. You start with the provided Optimate+ project and add the functionality to insert a macro for post-processing purposes.

Prerequisites

This tutorial assumes that you are already familiar with certain techniques in STAR-CCM+ and Optimate.

Before attempting this tutorial, make sure that you are familiar with the following techniques:

Techniques	Associated Tutorial
The STAR-CCM+ workflow	Introduction to STAR-CCM+
Using visualization tools, scenes, and plots	Introduction to STAR-CCM+
Recording and editing Java macros	Simple Java Macros: Post-Processing Objects
The Optimate+ workflow	Design Sweeps Pareto Optimization

This tutorial requires an Optimate+ license in order to use the **Optimization using SHEPRA** run mode.

Preparing the Post-Processing Macro

Prepare the post-processing macro that you want to use in Optimate.

1. In your workspace, create a folder named `optimizationWithMacro`.
2. Navigate to `[INSTALL_DIR]/doc/startutorialsdata/optimate/data`
3. Copy `2dAirfoilFlapOptimization.sim` and `macroOptimization.optm` to your workspace.
4. Launch STAR-CCM+.
5. Select **File > Load Simulation...**
6. In the *Load Simulation* dialog, click **Browse**.
7. In the *Open* dialog, select `2dAirfoilFlapOptimization.sim` and click **Open**.
8. In the *Load Simulation* dialog, click **OK**.

Create a macro that will be executed by STAR-CCM+ after each design is converted to a 2D mesh. The next step is to convert the 3D region into 2D.

9. Select **Mesh > Convert to 2D**
10. In the *Convert Regions To 2D* dialog, activate the **Delete 3D Regions After Conversion** option and click **OK**.

    The 3D region is converted to 2D.

Note: The **Cp plot** and **Y+ plot** plots don't have any parts assigned to them. This is because the surfaces of interest did not exist in the 3D model, but do exist in the current 2D model.

11. Save the simulation as `2dFlapOptimization_2d.sim`.

Record a macro adding all the lifting surfaces of the airfoil to **Cp** and **Y+** plots.

12. Click (**Start Recording...**) to start recording the Java macro.
13. In the *Save* dialog, save the macro as `macroToUpdatePlots.java` and click **Save**.
14. Multi-select the **Cp plot** and **Y+ plot** nodes.
15. Right-click on one of the selections and choose **Edit...**.
16. In the *Multiple Objects* dialog, click (**Custom Editor**) next to the **Parts** node.
17. In the *Multiple Objects - Parts* dialog, expand the **Regions > Airfoil 2D > Boundaries** and select the following nodes:

- **01 Airfoil Trailing Edge**
- **01 FF Trailing Edge**
- **01 RF Trailing Edge**
- **Default**

18. Click **OK**.
19. Click **Close**.
20. Click (**Stop Recording**) to stop recording the macro.
21. Save the simulation.
22. Close the simulation.

Starting Optimate

1. Launch STAR-CCM+ and load `2dFlapOptimization.sim` from your workspace.
2. From the STAR-CCM+ toolbar, click (**Start Optimate**).
   The *Load Study?* dialog appears.



3. Click **Open an Existing Optimate Study**.
4. In the *Select Optimate Study File* dialog, select `macroOptimization.optm` and click **Open**.
   The Optimate workspace appears and the settings that are contained in the project file are loaded.
5. Examine the *Run Mode*, *Variables*, and *Outputs* panels to familiarize yourself with the case setup.

Defining the Run Settings

Insert the Java macro into the workflow so that it converts the 3D mesh into a 2D mesh and sets up the Cp and Y+ plots.

1. Click the *Assembly* tab.
2. Activate the **Insert Macro** option under *STAR-CCM+ Run Options*.

   The *Macro Insertion Point* tool activates.
3. Set *Macro Insertion Point* to **After Meshing**.
4. Click (**Custom Editor**) next to *Macro To Be Inserted*.
5. In the *Select STAR-CCM+ Macros to be Inserted* dialog, navigate to the folder where you saved your macro from the previous section.
6. Select `macroToUpdatePlots.java` and click **Select**.

   The path to this Java macro is automatically added to *Macro Directory*.
7. Save the study as `macroOptimization_run.optm`.
8. Click **Build Study**.

   All necessary files for this study are created. A pop up window appears while the study is building.

Running the Study

The study is now ready to be run. To start the job on your local machine:

1. Click the *Run* tab.
2. Click (**Run**) and select **Run**.

   The current STAR-CCM+ server closes and the jobs are submitted in the background. The output from Optimate is displayed in the *Output* window. The STAR-CCM+ servers for each job are visible in the server list.

Post-Processing

Check the Coefficient of Pressure and Y+ plots in the results directory.

To access the solution data:

1. Navigate to your workspace folder for this tutorial.
2. Expand the `2dAirfoilFlapOptimization_run.optm\images` folder.
3. Check that the **Cp** and **Y+** plots are generated for each simulation.